

**Measuring the unfairness
feeling in allocation problems**

L.N. Hoang, F. Soumis,
G. Zaccour

G-2014-70

September 2014

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2014.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*.

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2014.

Measuring the unfairness feeling in allocation problems

Lê Nguyễn Hoàng^a

François Soumis^a

Georges Zaccour^b

^a GERAD & Polytechnique Montréal, Montréal (Québec)
Canada, H3C 3A7

^b GERAD & HEC Montréal, Montréal (Québec) Canada,
H3T 2A7

le.nguyen.hoang@gerad.ca
francois.soumis@gerad.ca
georges.zaccour@gerad.ca

September 2014

**Les Cahiers du GERAD
G–2014–70**

Copyright © 2014 GERAD

Abstract: In this paper, we introduce a new measure of social fairness based on unfairness feelings of the players involved in an allocation problem, e.g., cake-cutting problem or shift scheduling. We only require that each player be described by a von Neumann — Morgenstern utility function. Next, we propose a social normalization of each player's utility function, based on how each player sees the other players' shares through her own utility function. Further, we extend this normalization idea to a setting where the players are represented by a weighted oriented graph, where the weights assess the relatedness of (or similarities between) the agents. Among other results, we establish some links between our measures of fairness and those classically used in the cake-cutting-problem literature.

Key Words: Fairness, social normalization, utility, allocation problem, social network.

Résumé: Dans cet article, nous introduisons une nouvelle mesure d'équité dans un problème d'allocations, e.g., un problème de partage de gâteaux ou de construction d'horaires. Cette mesure se fonde sur les sentiments d'injustice des joueurs impliqués. Pour ce faire, nous proposons une normalisation sociale des fonctions d'utilités des joueurs, qui repose sur la manière selon laquelle chaque joueur juge les parts des autres à travers sa propre fonction d'utilité. Par la suite, nous étendons cette idée de normalisation au cas où les joueurs sont liés dans un graphe pondéré et orienté, dont les poids décrivent les interactions (ou les similarités) entre les joueurs. Entre autres, nous établissons des liens entre nos mesures d'équité et celles utilisées dans la littérature classique du problème de partage de gâteaux.

Mots clés: Équité, normalisation sociale, utilité, problème d'allocation, réseau social.

1 Introduction

The problem of the fair division of a cake, a metaphor used to designate a common resource, has been the topic of a large body of literature in the last six decades or so; see, e.g., Steinhaus (1948), Brams and Taylor (1995), Brams and Taylor (1996), Robertson and Webb (1998). Typically in this literature, two assumptions have been made on the individual utility functions of the stakeholders in the cake, namely: (i) they are additive; and (ii) the utility of the whole cake is (normalized to) one for all players. These assumptions have been instrumental in designing cake-cutting algorithms and deriving some properties. Further, solving such problems requires us to specify from the outset what is meant by a *fair* division. Here, the literature has proposed a series of definitions of fairness, e.g., exact, proportional, envy-free and equitable fairness, each having its pros and cons.¹

In this paper, we introduce a new measure of fairness without requiring that the utility functions be additive or that the whole-cake value be normalized to one for all players. The original motivation for developing this measure was a research contract to provide a methodology for shift-scheduling problems where a manager wishes to implement the (technically feasible) schedule that minimizes a certain unfairness criterion. The starting assumption is that the manager can obtain the employees' preferences regarding a small² set of acceptable schedules.³ In such a context, the additivity assumption naturally does not hold anymore, that is, the utility of the sum of two shifts is clearly not equal to the sum of their utilities, and the whole-cake-normalization assumption is meaningless. The idea that the utility function is not necessarily additive, but rather super- or sub-additive, is by no means akin to shift scheduling but is a standard assumption in economics. The implications of abandoning the additivity assumption are important. In particular, Mirchandani (2013) showed that most existing fair-division procedures are incompatible with nonadditive utility functions.

In this paper, we only require that each player have a von Neumann — Morgenstern (vNM) utility function. To be able to compare players' payoffs and adequately assess the fairness of any division of the cake, we propose a normalization of the individual utility functions. As we will see, this normalization is centered on the idea that each player compares her allocation to other players' allocations through the lens of her own utility function. For this reason, we call it a social normalization. Next, using socially normalized utility functions, we introduce the concept of the *unfairness feeling*. A division will then be called socially fair if all players have no unfairness feeling. We will relate our social fairness to classical cake-cutting-literature definitions. In problems with a large number of agents, or when the agents are heterogenous in some way, it may become intuitively appealing to suppose that each stakeholder in the cake is only sensitive to how "similarly" or "closely" players are treated. To handle such a case, we extend our definition of fairness to a setting where the players are located on a weighted oriented graph. This formulation captures the idea of a social network where individuals are only interested in the fate of those in their circles.

The rest of the paper is organized as follows: Section 2 provides some background and preliminaries. Sections 3 and 4 introduce our normalization of utilities and concepts of fairness, respectively. Section 5 deals with local fairness, and Section 6 briefly concludes.

2 Background and preliminaries

We start by recalling some of the most commonly used definitions of fairness in the cake-cutting-problem literature, to which we will link our fairness criterion:

Exact Fairness: A division is exact if all players' allocations are identical, i.e., exchanging shares will not affect any player's outcome.

¹Of course, many other measures of fairness and equity exist and are based on different premises. For an interesting discussion in the context of resource allocation, the interested reader may refer to Bertsimas et al. (2011).

²Pragmatically, it does not make sense to ask the employees to rate all feasible schedules.

³By acceptable, we mean a technically feasible schedule that does not involve a too high additional cost with respect to the least-cost one. Put differently, management is willing to forgo some revenues in order to please the employees.

Proportional Fairness (PF): A division is proportionally fair if every player prefers her allocation to an allocation from an exact division. Another interpretation is that players prefer their allocation to the average of what they would get if allocations were given away uniformly randomly. See, e.g., Procaccia (2009), Mossel and Tamuz (2010), Bei et al. (2012).

Envy-Freeness (EF): A division is envy-free if every player prefers her allocation to any other player's allocation. See the early contributions in, e.g., Foley (1967), Varian (1974), Varian (1976), Arnsperger (1994), and the more recent ones focusing on cake-cutting procedures in, e.g., Stromquist (2007), Cohler et al. (2011), Chen et al. (2013).

Equitable Fairness: A division is equitable if all players have the same utility for their respective shares.

To illustrate some of the drawbacks in these definitions, we consider a few anecdotal examples. When the cake is made of indivisible pieces, e.g., where a car and a summer cottage to be fairly shared following a divorce, an exact division is obviously not implementable. Even when it is, exactness can be very restrictive and lead to some counter-intuitive results. For instance, if a cake that contains chocolate on a half and nuts on the other half is to be shared between a person allergic to nuts and another who hates chocolate, then imposing an exact division would be peculiar and obviously not Pareto-optimal (assuming that this feature is of interest). Proportional fairness and envy-free fairness are much less demanding than exactness, but may also be infeasible when the goods are indivisible. Equitable fairness may be questionable on some grounds. To see this, consider a sugar cake with three cherries on top, to be divided among four individuals, three of whom have no interest at all in the cake but love cherries, while the fourth person only wants the sugar cake. One division is to give to the first three players one cherry, and the cake to the fourth person. This solution is intuitively fair, and it is fair according to the definitions of PF and EF, but it is not equitable. This highlights that equitable fairness may fail to achieve a fair solution according to common sense. In these examples, the focus was on a “physical” division of the cake rather than on dividing the corresponding total value of the cake. For this, we must define a utility function for each player that has certain properties. This is where the assumptions mentioned in the introduction come into play, namely, the additivity and normalization to one of the whole cake. We now introduce the notation and formally discuss these issues.

In a cake-cutting problem, a finite set of players $N = \{1, \dots, n\}$ and a cake **CAKE** are given. A division, or allocation, of the cake is a vector $x = (x_1, \dots, x_n)$, where $x_i \subset \text{CAKE}$ is the share of player i , and $\cup_{i \in N} x_i = \text{CAKE}$. Each player i has a utility function u_i that associates a real number to any x_i . Player i prefers share x_i to x'_i , if, and only if, $u_i(x_i) \geq u_i(x'_i)$. The utility function is additive; that is, for any disjoint subsets x_k and x_l , we have

$$u_i(x_k \cup x_l) = u_i(x_k) + u_i(x_l). \quad (1)$$

In particular, this implies that the utility of an empty allocation is equal to zero, i.e., $u_i(\emptyset) = 0$, and by the normalization of the whole-cake value, we have $u_i(\text{CAKE}) = 1$, $\forall i \in N$. With this notation, we can rephrase the above definitions of fairness as follows:

Exact Division: A division is exact if for any player $i \in N$, $u_i(x_j) = 1/n$, $\forall j \in N$.

Proportional Fairness: A division is proportionally fair (PF) if any player gets at least $1/n$, i.e., $u_i(x_i) \geq 1/n$, $\forall i \in N$.

Envy-Freeness: A division is envy-free (EF) if any player i prefers her allocation to any other player's allocation, i.e., $u_i(x_i) \geq u_i(x_j)$, $\forall j \in N$.

Equitable: A division is equitable if all players obtain the same utility, i.e., $u_i(x_i) = u_j(x_j)$, $\forall i, j \in N$.

Suppose now that a given set of feasible divisions of the cake are proposed to the agents and they are asked to rate them according to their utility functions. For any allocation $x = (x_1, \dots, x_n)$, the resulting evaluation can be represented by the following utility matrix:

$$U = \begin{pmatrix} u_1(x_1) & \dots & \dots & \dots & u_1(x_n) \\ \vdots & \ddots & & & \vdots \\ \vdots & & u_i(x_i) & & \vdots \\ \vdots & & & \ddots & \vdots \\ u_n(x_1) & \dots & \dots & \dots & u_n(x_n) \end{pmatrix} = \begin{pmatrix} U_{11} & \dots & \dots & \dots & U_{1n} \\ \vdots & \ddots & & & \vdots \\ \vdots & & U_{ii} & & \vdots \\ \vdots & & & \ddots & \vdots \\ U_{n1} & \dots & \dots & \dots & U_{nn} \end{pmatrix}, \quad (2)$$

where $(u_i(x_j))_{1 \leq i, j \leq n}$ (or U_{ij}) gives the utility of player i for player j 's share. An interesting feature of the utility matrix is that it shows at a glance what would happen if two players decided to trade shares. Further, the normalization condition $u_i(\text{CAKE}) = 1$, for all $i \in N$ implies that U is a row stochastic matrix, i.e., all entries are nonnegative and entries in any row add up to 1.

We note that this utility matrix contains most (if not all) of the relevant information needed to judge the fairness of an allocation, and it will play a central role in the rest of the paper. To illustrate, consider the following sugar-cake example where the total resource to be allocated can be written as $\text{CAKE} = \text{cherry}_1 \cup \text{cherry}_2 \cup \text{cherry}_3 \cup \text{sugar_cake}$, with $u_i(\text{CAKE}) = 1$ for all $i \in N$. As the first three players have no utility for the `sugar_cake` and the fourth player, no utility for a cherry, it seems reasonable to recommend the allocation $x_i = \text{cherry}_i$ for $1 \leq i \leq 3$ and $x_4 = \text{sugar_cake}$. The corresponding utility matrix is then the following:

$$U = \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Now, we can easily check for the satisfaction of the different criteria (definitions) of fairness. In an exact division, all entries must be equal to $1/n$. For a division to be PF, the diagonal, which represents the players' utilities for their own allocation, must only contain numbers greater than or equal to $1/n$. Meanwhile, envy-free fairness holds when each diagonal term is the largest in its row. Finally, a division is equitable when all diagonal terms are the same. Inspecting the above utility matrix, we conclude that the considered allocation is PF and EF, but not exact or equitable.

2.1 Generalization to non-additive utility functions

From now on, we relax the additivity assumption of the utility function in (1) to only require that each player has a von Neumann — Morgenstern (vNM) utility function.⁴ Further, we set aside the condition $u_i(\text{CAKE}) = 1$, for all $i \in N$ and consequently, the utility matrix U can contain any numbers. We recall the following well-known theorem stating that a vNM utility function is defined up to a positive affine transformation:

Theorem 1 (Von Neumann and Morgenstern (1944)) *If u_i is the vNM utility function of player i , then, for any $d_i > 0$ and any $b_i \in \mathbb{R}$, $d_i u_i + b_i$ is also the utility function of player i .*

The above theorem implies that a utility matrix is also defined up to an affine positive transformation of its rows. This means that, for any positive diagonal matrix D and any matrix B with identical columns, U

⁴It is well known that the existence of a vNM utility function is equivalent to the players' preferences being complete, transitive, continuous and independent. A preference of player i is an order $\succeq$ on $\Delta(X)$. It is

Complete: if, for all $\tilde{x}_i, \tilde{y}_i$, we have $\tilde{x}_i \succ \tilde{y}_i$, $\tilde{y}_i \succ \tilde{x}_i$ or $\tilde{x}_i = \tilde{y}_i$ (both lotteries are equally preferred).

Transitive: if, whenever $\tilde{x}_i \succeq \tilde{y}_i$ and $\tilde{y}_i \succeq \tilde{z}_i$, then $\tilde{x}_i \succeq \tilde{z}_i$.

Continuous: if, whenever $\tilde{x}_i \succeq \tilde{y}_i \succeq \tilde{z}_i$, there exists $p \in [0, 1]$ such that $p\tilde{x}_i + (1-p)\tilde{z}_i = \tilde{y}_i$.

Independent: if, $\tilde{x}_i \succ \tilde{y}_i$, then for any $p \in [0, 1]$ and $\tilde{z}_i$, we have $p\tilde{x}_i + (1-p)\tilde{z}_i \succ p\tilde{y}_i + (1-p)\tilde{z}_i$.

and $DU + B$ are equivalent, that is,

$$U \text{ is equivalent to } \begin{pmatrix} d_1 & 0 & \cdots & \cdots & 0 \\ 0 & \ddots & \ddots & 0 & \vdots \\ \vdots & \ddots & d_i & \ddots & \vdots \\ \vdots & 0 & \ddots & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & d_n \end{pmatrix} U + \begin{pmatrix} b_1 & \cdots & b_1 & \cdots & b_1 \\ \vdots & & \vdots & & \vdots \\ b_i & \cdots & b_i & \cdots & b_i \\ \vdots & & \vdots & & \vdots \\ b_n & \cdots & b_n & \cdots & b_n \end{pmatrix}, \quad (3)$$

where all entries d_i are strictly positive. Without the additivity assumption, the definitions of exact and PF fairness are changed as follows:

Exact Division: A division $x = (x_1, \dots, x_n)$ is exact if for all $i \in N$, $u_i(x_j) = u_i(x_l)$, for $j, l = 1, \dots, n$.

Proportional Fairness: A division is PF if any player i gets at least the average of the others, i.e.,

$$u_i(x_i) \geq u_i(\tilde{x}) = \frac{1}{n} \sum_{j \in N} u_i(x_j), \quad (4)$$

where $\tilde{x}$ is the uniform distribution on the allocations $x_1, \dots, x_n$.

The definitions of envy-freeness and equitable fairness remain as before. In terms of the utility matrix, we note that exact division, PF and EF are characterized by the following inequalities:

$$\begin{aligned} \text{Exact division:} & \quad U_{ii} - U_{ij} \geq 0 \quad \text{and} \quad U_{ij} - U_{ii} \geq 0 \\ \text{Proportional fairness:} & \quad U_{ii} - U_i^T \mathbf{1}_n / n \geq 0, \\ \text{Envy-free fairness:} & \quad U_{ii} - U_{ij} \geq 0, \end{aligned} \quad (5)$$

for all players $i, j \in N$. To be consistent, we expect any concept of fairness to yield the same conclusion for U and $DU + B$. This is shown in the following proposition.

Proposition 1 *A utility matrix U is exact (respectively, PF and EF) if and only if $DU + B$ is.*

Proof. Each inequality defining exactness, PF and EF is of the form $U_i C \geq 0$, where C is a column. For instance, for the j^{th} constraint of EF, by fixing $C_i = 1$ and $C_j = -1$ and $C_k = 0$ if $k \notin \{i, j\}$, the constraint $U_i C \geq 0$ corresponds to $U_{ii} - U_{ij} \geq 0$. Crucially, for all definitions above, entries of C add up to 0, i.e., $\mathbf{1}_n^T C = 0$. Now, consider positive affine transformations of U 's i -th row written as $U_i \mapsto d_i U_i + b_i \mathbf{1}_n^T$. The fairness inequality becomes $(d_i U_i + b_i \mathbf{1}_n^T) C = d_i (U_i C) + b_i (\mathbf{1}_n^T C) = d_i (U_i C)$, since $\mathbf{1}_n^T C = 0$. As $d_i > 0$, we have $(d_i U_i + b_i \mathbf{1}_n^T) C \geq 0$ if and only if $U_i C \geq 0$. In other words, inequalities that define fairness criteria hold for U if and only if they hold for any positive affine transformation of U . $\square$

Let us provide a visual standpoint of the definitions of fairness. Exact division corresponds to symmetric allocations. Since each column represents the utility of an allocation as regarded by different players, exact division holds true when all columns of U are the same. Equivalently, this means that all the elements in any row are equal, that is, $u_i(x_j) = u_i$, $j = 1, \dots, n$. The positive affine transformation defined by $\tilde{u}_i(x_j) = u_i(x_j) - u_i$ yields $\tilde{u}_i(x_j) = 0$ for all $j \in N$. By doing so for all columns, we see that a utility matrix U is exact if, and only if, its columns are affine positive transformations of the zero matrix. For instance, the utility matrix

$$U = \begin{pmatrix} 1 & 1 & 1 \\ -2 & -2 & -2 \\ 4 & 4 & 4 \end{pmatrix} \quad (6)$$

represents an exact division. Indeed, by choosing $D = I_3$, where I_3 is the identity matrix, and $(b_1, b_2, b_3) = (-1, 2, -4)$, we obtain $DU + B = 0$.

Proportional fairness means that the diagonal element U_{ii} of a utility matrix U is greater or equal to the average of the other row entries, that is, $U_{ii} \geq \frac{1}{n-1} \sum_{j \neq i} U_{ij}$. Finally, U is envy-free if, and only if, each diagonal element is the greatest of its row. To illustrate, consider the utility matrix

$$U = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 1 & 2 \\ 0 & 3 & 3 \end{pmatrix}. \quad (7)$$

Since the columns of U do not match, the division is not exact. It is not EF either, as the diagonal elements are not all the greatest of their rows. Indeed, the second player prefers the third player's allocation, as $u_2(x_2) = U_{22} < U_{23} = u_2(x_3)$. However, the allocation is proportionally fair because each diagonal element is greater than or equal to the average of its row.

To wrap up, we have rephrased the most commonly used fairness definitions for a cake-cutting problem in the case where the utility functions are not additive and the whole cake value is not normalized to one. Further, we characterized these definitions in terms of the utility matrix.

3 Social normalization

To compare players' assessments for different allocations, we need to normalize their utility functions. Let $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$. Denote by $\mu_i(u_i, x_{-i})$ the average valuation that player i has for other players' shares, and by $\sigma_i(u_i, x_{-i})$ the corresponding standard deviation, that is,

$$\mu_i(u_i, x_{-i}) = \frac{1}{n-1} \sum_{j \neq i} u_i(x_j) \quad \text{and} \quad \sigma_i^2(u_i, x_{-i}) = \frac{1}{n-1} \sum_{j \neq i} (u_i(x_j) - \mu_i(u_i, x_{-i}))^2.$$

In terms of the utility matrix U , $\mu_i(u_i, x_{-i})$ is the average of all the entries in row i except the diagonal element, and $\sigma_i(u_i, x_{-i})$ is the standard deviation. These terms have a social flavor, as they correspond to player i 's judgment of others' allocations. They lead us to define the social normalization of utility functions.

Proposition 2 *Suppose that $\sigma_i(u_i, x_{-i}) > 0$ for all $i \in N$. The utility function*

$$\bar{u}_i(\cdot) = \frac{u_i(\cdot) - \mu_i(u_i, x_{-i})}{\sigma_i(u_i, x_{-i})}, \quad (8)$$

is socially normalized, that is,

$$\mu_i(\bar{u}_i, x_{-i}) = 0 \quad \text{and} \quad \sigma_i(\bar{u}_i, x_{-i}) = 1.$$

Proof. Straightforward computations lead to the results. Indeed,

$$\begin{aligned} \mu_i(\bar{u}_i, x_{-i}) &= \frac{1}{n-1} \sum_{j \neq i} \bar{u}_i(x_j) = \frac{1}{n-1} \sum_{j \neq i} \frac{u_i(x_j) - \mu_i(u_i, x_{-i})}{\sigma_i(u_i, x_{-i})} \\ &= \frac{1}{\sigma_i(u_i, x_{-i})} \left(\left(\frac{1}{n-1} \sum_{j \neq i} u_i(x_j) \right) - \left(\frac{1}{n-1} \sum_{j \neq i} \mu_i(u_i, x_{-i}) \right) \right) \\ &= \frac{1}{\sigma_i(u_i, x_{-i})} (\mu_i(u_i, x_{-i}) - \mu_i(u_i, x_{-i})) = 0, \end{aligned}$$

and, similarly,

$$\begin{aligned} \sigma_i(\bar{u}_i, x_{-i}) &= \frac{1}{n-1} \sum_{j \neq i} (\bar{u}_i(x_j) - \mu_i(\bar{u}_i, x_{-i}))^2 = \frac{1}{n-1} \sum_{j \neq i} \left(\frac{u_i(x_j) - \mu_i(u_i, x_{-i})}{\sigma_i(u_i, x_{-i})} - 0 \right)^2 \\ &= \frac{1}{\sigma_i(u_i, x_{-i})^2} \left(\frac{1}{n-1} \sum_{j \neq i} (u_i(x_j) - \mu_i(u_i, x_{-i}))^2 \right) = \frac{\sigma_i(u_i, x_{-i})^2}{\sigma_i(u_i, x_{-i})^2} = 1. \end{aligned}$$

Therefore, $\bar{u}_i$ is socially normalized. □

Player i 's socially normalized utility function is independent of other players' assessments of the division x and has no measurement unit. It gives how many standard deviations above (or below) average a player considers her allocation to be compared to what the others have. This is illustrated in Figure 1 for the case where the number of players is large enough that each player's utility for the other players' allocations can be described by a probability density function.




Figure 1: Socially normalized utility is a comparison of one's utility with a benchmark given by utilities of trades.

The following proposition shows that normalizing a utility function or any admissible affine transformation of that function yields the same result:

Proposition 3 Suppose that $\sigma_i(u_i, x_{-i}) > 0$ for all $i \in N$. For any $d_i > 0$ and any $b_i \in \mathbb{R}$, the social normalizations of u_i and $d_i u_i + b_i$ coincide.

Proof. We first note that $\mu_i(d_i u_i + b_i, x) = d_i \mu_i(u_i, x) + b_i$ and $\sigma_i(d_i u_i + b_i, x) = d_i \sigma_i(u_i, x)$. Then, applying (8) to $d_i u_i + b_i$ gives

$$\frac{(d_i u_i + b_i)(\cdot) - \mu_i(d_i u_i + b_i, x)}{\sigma_i(d_i u_i + b_i, x)} = \frac{d_i u_i(\cdot) + b_i(x_i) - (d_i \mu_i(u_i, x) + b_i)}{d_i \sigma_i(u_i, x)} = \frac{u_i(\cdot) - \mu_i(u_i, x)}{\sigma_i(u_i, x)}, \quad (9)$$

which shows that the social normalizations of $d_i u_i + b_i$ and u_i coincide. $\square$

From now on, we denote by $\bar{U}$ the socially normalized utility matrix, with $\bar{U}_{ij} = \bar{u}_j(x_i)$ for all $i, j \in N$. Next, we show that checking proportional fairness is a very straightforward test when working with socially normalized utility matrices.

Proposition 4 Suppose that $\sigma_i(u_i, x_{-i}) > 0$ for all $i \in N$. A utility matrix U is PF if and only if $\bar{U}_{ii} \geq 0$ for all $i \in N$.

Proof. By construction, for any socially normalized utility matrix, the average of all off-diagonal elements in a row is 0. Recalling that PF requires that each diagonal element be greater than or equal to the average of all off-diagonal elements, we see that PF is equivalent to the criterion $\bar{U}_{ii} \geq 0$ in the Proposition. $\square$

The above propositions are stated under the condition that $\sigma_i(u_i, x_{-i}) > 0$ for all $i \in N$; otherwise, the socially normalized utility function (8) does not exist. Further, observe that in any two-player cake-cutting problem, our social normalization is useless as the standard deviations are computed for a scalar, and are

therefore equal to zero. An illustrative case where $\sigma_i(u_i, x_{-i}) > 0$ is not satisfied for all $i \in N$ is in our sugar-cake example for $x_i = \text{cherry}$ for $1 \leq i \leq 3$ and $x_4 = \text{sugar_cake}$. Indeed, player 4 always gets a cherry when trading with others. Consequently, $\sigma_4(u_4, x) = 0$ and this means that this player has no socially normalized utility function. We will deal later on with this “degenerate” case of $\sigma_i(u_i, x_{-i}) = 0$.

We will henceforth omit the arguments of μ_i and σ_i when no ambiguity may arise.

Remark 1 *In any 3-player cake-cutting problem, all off-diagonal entries of $\bar{U}$ are $\pm 1/2$. To illustrate, consider the following utility matrix:*

$$U = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 3 & 2 \\ 0 & 3 & 3 \end{pmatrix}. \quad (10)$$

It is easy to check that $\mu_1 = 1/2$, $\sigma_1 = 1$, $\mu_2 = 1$, $\sigma_2 = 2$, $\mu_3 = 3/2$ and $\sigma_3 = 3$, and consequently, the socially normalized utility matrix is given by

$$\bar{U} = \begin{pmatrix} 3/2 & 1/2 & -1/2 \\ -1/2 & 1 & 1/2 \\ -1/2 & 1/2 & 1/2 \end{pmatrix}. \quad (11)$$

4 Social fairness

In this section, we introduce our new measure of fairness. To do so, let us start by stating, very informally, three intuitive principles regarding the complaints that could be formulated by the agents against a proposed allocation: (i) if no player complains about a division of the cake, then this division is considered fair; (ii) no player will complain if she gets more than what *she thinks* she should get, based on what others have been given; and (iii) a complaint put forward by any player has to be assessed in relative terms, i.e., with respect to other players' complaints or to a certain standard. Based on these simple ideas, we introduce a measure of the *unfairness feeling* of a player.

Definition 1 *Assume that $\sigma_i(u_i, x_{-i}) > 0$ for all $i \in N$. Player i 's unfairness feeling is defined by*

$$\mathcal{F}_i = (h(\bar{u}_i(x_i)) - \bar{h})^+, \quad (12)$$

where $h(\cdot)$ is a nonnegative and convex decreasing function, and $\bar{h}$ is the complaint-potential reference given by

$$\bar{h} = \min \left(h(0), \frac{1}{n} \sum_{i \in N} h(\bar{u}_i(x_i)) \right). \quad (13)$$

Definition 2 *Assume that $\sigma_i(u_i, x_{-i}) > 0$ for all $i \in N$. A division $x = (x_1, \dots, x_n)$ feels fair for player i if $\mathcal{F}_i = 0$. When all players feel that the division is fair, we say that the division is socially fair (SF).*

A few comments are in order regarding the ingredients in the above definitions. Our measure of the unfairness feeling is based, as it should be, on the comparable players' socially normalized utilities. However, to take into account the fact that differences in *positive* socially normalized utilities are less of a problem than differences in *negative* ones, we rescale the socially normalized utilities when measuring complaints by the mapping h , to which we refer as the *complaint potential*. The assumptions made on $h(\cdot)$ ensure that negative variations of negative socially normalized utilities will result in a much greater complaint-potential increase than negative variations of positive socially normalized utilities. Now, as pointed out before, a complaint by a player has to be judged with respect to a norm or a reference point. This is the role played by $\bar{h}$, which is defined as the minimum between the average of complaint potentials and the complaint potential evaluated at $\bar{u}_i(x_i) = 0$. The idea behind requiring $\bar{h}$ to be at most $h(0)$ is that the players will necessarily complain if their socially normalized utility is less than 0, that is, if the allocation is not PF. Given this, we say that player i feels unfairness if her complaint potential $h(\bar{u}_i(x_i))$ is greater than the complaint potential reference $\bar{h}$. We measure this feeling of unfairness as the difference between these two values if it is positive. In other words, player i 's feeling of unfairness is $(h(\bar{u}_i(x_i)) - \bar{h})^+$, where $t^+ = \max(0, t)$ is the positive part of t .

The terminology “social fairness” is justified by social comparisons at different levels. First, the computation of unfairness feelings requires the social normalization of utility functions, which, for each player, we recall, is dependent on others’ allocations. Second, this unfairness feeling is computed by comparing one’s complaint potential $h(\bar{u}_i(x_i))$ to the average $\bar{h}$ of others’ complaint potentials. Finally, social fairness requires that the division feel fair to all players.

An important difference between social fairness and the other mentioned before is that social fairness assumes that the players are not only sensitive to the average (as for PF) or the best of the others’ allocations (as for EF), but they are also concerned by the degree of dispersion in the shares through $\bar{u}_i(x_i)$.

Remark 2 *In the rest of the paper, we adopt the following functional for $h(\cdot)$:*

$$\text{Complaint potential function : } h(t) = \exp(-t),$$

which clearly satisfies the above properties and has the merit of being simple.

Now, suppose that the agents are shown a number of proposals for dividing the cake and are asked to evaluate them using the utility matrices. If one of them happens to be SF, then we are done. If not, that is, if there is at least one $\mathcal{F}_i > 0$ for each proposed allocation $x = (x_1, \dots, x_n)$, then it becomes relevant to compare these non-SF allocations. We suggest doing so based on the degree of social fairness.

Definition 3 *The degree of social fairness of an allocation $x = (x_1, \dots, x_n)$ is measured by the average of the unfairness feelings, that is,*

$$\mathcal{G}_{SF}(u, x) = \frac{1}{n} \sum_{i \in N} \left(h(\bar{u}_i(x_i)) - \bar{h} \right)^+. \quad (14)$$

The lower-bound value of $\mathcal{G}_{SF}(u, x)$ is 0, a case that occurs if and only if $\mathcal{F}_i = 0$ for all $i \in N$, that is, the allocation is SF. To illustrate, let us reconsider the three-player example described by the utility matrix (10) and socially normalized utility matrix (11). First, we have

$$\frac{1}{3} \sum_{i \in N} h(\bar{u}_i(x_i)) = \left(e^{-3/2} + e^{-1} + e^{-1/2} \right) / 3 = 0.39918,$$

and

$$\begin{aligned} \bar{h} &= \min \left(h(0), \frac{1}{n} \sum_{i \in N} h(\bar{u}_i(x_i)) \right) \\ &= \min(1, 0.39918) = 0.39918. \end{aligned}$$

Second, we compute $\mathcal{F}_i, i \in N$ to get

$$\mathcal{F}_1 = (e^{-3/2} - 0.39918)^+ = 0, \quad \mathcal{F}_2 = (e^{-1} - 0.39918)^+ = 0, \quad \mathcal{F}_3 = (e^{-1/2} - 0.39918)^+ = 0.20735.$$

This shows that players 1 and 2 do not have any unfairness feeling, whereas player 3 feels the division to be unfair. As $\mathcal{F}_3 > 0$, it becomes relevant to compute the degree of social fairness of this allocation, which is given by

$$\mathcal{G}_{SF}(u, x) = \frac{1}{3} (0 + 0 + 0.20735) = 0.069117.$$

How to interpret the above value? As our measure of the degree of social fairness is not upper bounded, there is no answer in absolute terms to this question. However, in relative terms, an answer can be given by comparing the above value to the corresponding values of other feasible allocations, and then selecting the one that minimizes $\mathcal{G}_{SF}(u, \cdot)$.

4.1 Links to other fairness definitions

In the rest of this section, we establish some relationships between our fairness criterion and those commonly used in the cake-cutting-problem literature. The following theorem provides an upper bound to the degree of social fairness for allocations that are proportionally fair or envy-free. Its proof uses Lemma 1, which is given in Appendix 7.

Theorem 2 Assume $\sigma_i > 0$ for all players i . If a division x is PF or EF, then $\mathcal{G}_{SF}(u, x) \leq 1/4$.

Proof. Assume the division is proportionally fair. Then, for any player i , we have $\bar{u}_i(x_i) \geq 0$. Hence, $0 \leq h_i \leq 1$, where $h_i = \exp(-\bar{u}_i(x_i))$. In particular, the average of these values is at most one. Hence, $\bar{h}$ is the average of h_i for $i \in N$. Then,

$$\mathcal{G} = \mathcal{G}_{SF}(u, x) = \frac{1}{n} \sum_{i \in N} (h_i - \bar{h})^+ = \frac{1}{n} \sum_{h_i \geq \bar{h}} (h_i - \bar{h}). \quad (15)$$

Let us now introduce

$$\mathcal{H} = \frac{1}{n} \sum_{\bar{h} > h_i} (\bar{h} - h_i). \quad (16)$$

Then, we have $\mathcal{G} - \mathcal{H} = \frac{1}{n} \sum_{i \in N} h_i - \bar{h} = 0$. Thus, $\mathcal{G} = \mathcal{H}$. Moreover, using Lemma 1, we have

$$\mathcal{G} = \frac{1}{2} (\mathcal{G} + \mathcal{H}) = \frac{1}{2n} \sum_{i \in N} |h_i - \bar{h}| \leq 1/4. \quad (17)$$

This proves the Proposition for PF. The fact that any envy-free division is PF completes the proof. $\square$

As $\mathcal{G}_{SF}(u, x)$ is nonnegative by construction, the above theorem ensures that the degree of social fairness of any PF or EF lies in the bounded interval $[0, 1/4]$. Back to our above example where we had $\mathcal{G}_{SF}(u, x) = 0.069117$, we now know that this division x does much better than the worst PF (or EF) division that would have resulted in a $\mathcal{G}_{SF}(u, x)$ of $1/4$. The following example gives an envy-free division that scores badly in terms of our degree of social fairness, that is, its $\mathcal{G}_{SF}(u, x)$ is close to $1/4$:

Example 1 Consider the utility matrix given by

$$U = \begin{pmatrix} M & 1 & \cdots & \cdots & 1 & 0 \\ 0 & \ddots & \ddots & & & 1 \\ 1 & \ddots & M & \ddots & & \vdots \\ \vdots & \ddots & \ddots & 1 & \ddots & \vdots \\ \vdots & & \ddots & \ddots & \ddots & 1 \\ 1 & \cdots & \cdots & 1 & 0 & 1 \end{pmatrix}. \quad (18)$$

This matrix has been constructed as follows:

- *Diagonal terms:* For $i \leq n/2$, set $U_{ii} = M > 1$, and for $i > n/2$, set $U_{ii} = 1$.
- *Off-diagonal terms:*
 - For $1 \leq i < n$, set $U_{i+1,i} = 0$.
 - Set $U_{1n} = 0$.
 - All other entries are $U_{ij} = 1$.

As the diagonal elements are the greatest of their rows, the division represented by U is EF. Excluding the diagonal, each row has $n - 2$ cells with 1 as an entry, and one cell with 0 as an entry. Consequently, for any player i , we have

$$\mu_i = \frac{n-2}{n-1} = 1 - \frac{1}{n-1} \quad \text{and} \quad \sigma_i^2 = \frac{1}{n-1} \left(\frac{(n-2)^2}{(n-1)^2} + (n-2) \frac{1}{(n-1)^2} \right) = \frac{n-2}{(n-1)^2}. \quad (19)$$

The socially normalized utilities are given by

$$\bar{u}_i(x_i) = \begin{cases} \left(M - \frac{n-2}{n-1} \right) \frac{n-1}{\sqrt{n-2}}, & \text{if } i \leq n/2, \\ \frac{1}{\sqrt{n-2}}, & \text{if } i > n/2. \end{cases} \quad (20)$$

For n large enough, we have $h(\bar{u}_i(x_i)) = e^{-(M - \frac{n-2}{n-1}) \frac{n-1}{\sqrt{n-2}}} \leq e^{-M}$ if $i \leq n/2$, and $h(\bar{u}_i(x_i)) = e^{-1/\sqrt{n-2}}$, for $i > n/2$. When M and n are very large, then half of the complaint potentials are nearly 0, and the other half are nearly 1. Hence, the degree of SF is nearly $1/4$.

In the next proposition, we have a reciprocal result for PF.

Proposition 5 Assume $\sigma_i > 0$ for all players i ; then, social fairness implies proportional fairness.

Proof. By definition, we have $\bar{h} \leq h(0)$. If $\mathcal{G}_{SF}(u, x) = 0$, then $h(\bar{u}_i(x_i)) = \bar{h} \leq h(0)$ for all $i \in N$. Yet, h is decreasing. Thus, $\bar{u}_i(x_i) = \bar{U}_{ii} \geq 0$ for all players i . Proposition 4 enables the conclusion. $\square$

Figure 2 recapitulates the relationships between the fairness criteria.




Figure 2: Exactness, EF, SF and PF and their interrelations. The dotted arrow “bound” from fairness criterion A to SF means that A implies a bound of $1/4$ on the degree of SF.

4.2 Non-existence of socially normalized utilities

The results above are stated under the assumption that $\sigma_i(u_i, x_{-i}) > 0$ for all $i \in N$. Now, we deal with the case where $\sigma_i(u_i, x_{-i}) = 0$ for at least one player, which occurs when $u_i(x_j) = \mu_i(u_i, x_{-i})$ for $j \neq i$. Denote by $\tilde{N} \subseteq N$ the subset of players who have a socially normalized utility function, and by $N \setminus \tilde{N}$ the subset of players for whom $\sigma_i(u_i, x_{-i}) = 0$. For any player $i \in N \setminus \tilde{N}$, three possibilities can occur when she compares her allocation to the average of the other players’ allocations, namely:

1. $u_i(x_i) > \mu_i(u_i, x_{-i})$, i.e., player i prefers her allocation to the average of the other players’ allocations. Denote by N^+ the subset of players for whom this inequality holds. For these players, we set their socially normalized utility $\bar{u}_i(x_i)$ equal to $+\infty$. Consequently, the lower bound of complaint potential is given by $h(\bar{u}_i(x_i)) = e^{-\infty} = 0$.
2. $u_i(x_i) = \mu_i(u_i, x_{-i})$, i.e., player i is indifferent between her allocation and the average of the other players’ allocations. Denote by N^0 the subset of players for whom this equality holds. For these players, we set their socially normalized utility equal to 0.

3. $u_i(x_i) < \mu_i(u_i, x_{-i})$, i.e., player i prefers the average of the other players' allocations to her allocation. Let N^- be the set of players for whom this equality holds. For these players, we set their socially normalized utility equal to $-\infty$.

Consequently, we have the following partition of the set of players:

$$N = \tilde{N} \cup N^+ \cup N^0 \cup N^-. \quad (21)$$

Obviously, players in N^- will complain about their shares, and any reasonable fairness criterion must declare a division unfair whenever N^- is non-empty. To illustrate the construction of a socially normalized utility matrix when $\sigma_i(u_i, x_{-i})$ is equal to zero for at least one player, we reconsider the example of the sugar cake with three cherries. Recall that the utility matrix is given by

$$U = \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Clearly, for player 4, we have $\sigma_4(u_4, x_{-4}) = 0$, and $u_4(x_4) > \mu_4(u_4, x_{-4})$. As this player is in N^+ , we set $\bar{U}_{44} = +\infty$, and compute the rest using

$$\mu_i(u_i, x_{-i}) = \frac{1}{n-1} \sum_{j \neq i} u_i(x_j) \quad \text{and} \quad \sigma_i^2(u_i, x_{-i}) = \frac{1}{n-1} \sum_{j \neq i} (u_i(x_j) - \mu_i(u_i, x_{-i}))^2,$$

to obtain the following socially normalized utility matrix:

$$\bar{U} = \begin{pmatrix} \sqrt{2}/2 & \sqrt{2}/2 & \sqrt{2}/2 & -\sqrt{2} \\ \sqrt{2}/2 & \sqrt{2}/2 & \sqrt{2}/2 & -\sqrt{2} \\ \sqrt{2}/2 & \sqrt{2}/2 & \sqrt{2}/2 & -\sqrt{2} \\ 0 & 0 & 0 & +\infty \end{pmatrix}. \quad (22)$$

Using the above modifications for players in the set $N \setminus \tilde{N}$, we introduce the following generalization of the degree of social fairness.

Definition 4 *The degree of social fairness is defined by*

$$\mathcal{G}_{SF}(u, x) = \begin{cases} +\infty & \text{if } N^- \neq \emptyset, \\ 0 & \text{if } N^0 = N, \\ \frac{1}{|\tilde{N} \cup N^+|} \sum_{i \in \tilde{N} \cup N^+} (h(\bar{u}_i(x_i)) - \bar{h})^+ & \text{otherwise,} \end{cases} \quad (23)$$

where

$$\bar{h} = \frac{1}{|\tilde{N} \cup N^+|} \sum_{i \in \tilde{N} \cup N^+} h(\bar{u}_i(x_i))$$

In the sugar-cake example, we have $\bar{h} = (3e^{-\sqrt{2}/2} + 0)/4 \approx 0.3$, and consequently, we obtain the following SF degree for the allocation $x = (\text{cherry}_1, \text{cherry}_2, \text{cherry}_3, \text{sugar_cake})$:

$$\mathcal{G}_{SF}(u, x) = \frac{1}{4} \left(3(e^{-\sqrt{2}/2} - \bar{h}) + 0 \right) = \frac{3}{16} e^{-\sqrt{2}/2} \approx 0.092. \quad (24)$$

The following proposition is a straightforward characterization of socially fair divisions of a cake:

Proposition 6 *A division is socially fair if and only if one of the two following conditions is satisfied:*

1. $N = N^+ \cup N^0$ (or, equivalently, $\tilde{N} = N^- = \emptyset$).
2. $N^+ = N^- = \emptyset$ and $\bar{u}_i(x_i) = \bar{u}_j(x_j) \geq 0$ for all $i, j \in \tilde{N}$.

Proof. First, note that if $N = N^+ \cup N^0$, then $\exp(-\bar{u}_i(x_i)) = 0$ for all $i \in \tilde{N} \cup N^+$ since $\tilde{N} = \emptyset$. Thus, their variance is zero, which proves that the division is SF. Second, let (u, x) be an SF division such that $\tilde{N}$ is non-empty. Obviously, we need to have $N^- = \emptyset$. Now, let $i \in \tilde{N}$, then $\exp(-\bar{u}_i(x_i)) > 0$. For the variance to be zero, we thus cannot have $\exp(-\bar{u}_j(x_j)) = 0$, which means that N^+ is empty. What is more, the values $\exp(-\bar{u}_i(x_i))$ must all have the same values for $i \in \tilde{N}$, which corresponds to the equalities $\bar{u}_i(x_i) = \bar{u}_j(x_j)$ for all $i, j \in \tilde{N}$. Finally, we must have $h(\bar{u}_i(x_i)) = \bar{h}$, which requires $\bar{u}_i(x_i) \geq 0$. $\square$

We end this section by making the following observations regarding exact divisions: (i) their socially normalized utility matrices are zero matrices, i.e., $\bar{U} = 0$; and (ii) they are socially fair (this is a consequence of the second item in the above proposition).

5 Local fairness

In this section, we assume that each player is more sensitive to the shares given to “similar” or “close” players, than by the shares obtained by the other players. To illustrate, a full professor is probably more concerned by how much of the total research budget the head of the department is allocating to another full professor than by the share reserved to a starting assistant professor. To handle such a case, we represent the interactions (or connectedness) between the players by a weighted directed graph, and redefine fairness in local terms.

Consider a weighted directed graph where w_{ij} is the weight on arc (i, j) , that is, the link between players i and j . The larger the value of w_{ij} , the more player i cares about player j 's share, and $w_{ij} = 0$ means that what player j is getting is of no concern to player i . We suppose that for each player i there exists at least a player j with $w_{ij} > 0$, such that $\sum_{j \in N} w_{ij} \neq 0$. Note that w_{ij} is not necessarily equal to w_{ji} . We set $w_{ii} = 0$ for any player i , and consequently, the matrix of weights W contains nonnegative numbers with a zero diagonal. We will refer to W as the social network matrix. Clearly, if $w_{ij} = 1$ for all $i \neq j$ and $w_{ii} = 0$ for all i , then we recover the case studied in the previous section.

As before, we will normalize the utilities, with the difference here being that we take into account the links, or their absence, between the players. Denote by $\mu_i(u_i, x, w_i)$ the weighted average utility for player i given by

$$\mu_i(u_i, x, w_i) = \left(\sum_{j \in N} w_{ij} u_i(x_j) \right) / \left(\sum_{j \in N} w_{ij} \right). \quad (25)$$

Observe that the above quantity only involves terms of the i -th rows of the utility matrix U and of the social network matrix W . Unlike previously, there is no need to write x_{-i} because $w_{ii} = 0$. The weighted standard deviation for player i is defined by

$$\sigma_i^2(u_i, x, w_i) = \left(\sum_{j \in N} w_{ij} (u_i(x_j) - \mu_i)^2 \right) / \left(\sum_{j \in N} w_{ij} \right). \quad (26)$$

Proposition 7 *Suppose that $\sigma_i(u_i, x, w_i) > 0$ for all $i \in N$. The utility function*

$$\bar{u}_i(\cdot) = \frac{u_i(\cdot) - \mu_i(u_i, x, w_i)}{\sigma_i(u_i, x, w_i)}, \quad (27)$$

is locally socially normalized, that is,

$$\mu_i(\bar{u}_i, x, w_i) = 0 \quad \text{and} \quad \sigma_i(\bar{u}_i, x, w_i) = 1.$$

Proof. Straightforward computation gives

$$\begin{aligned}\mu_i(\bar{u}_i, x, w_i) &= \left(\sum_{j \in N} w_{ij} \bar{u}_i(x_j) \right) / \left(\sum_{j \in N} w_{ij} \right) = \left(\sum_{j \in N} w_{ij} \frac{u_i(x_j) - \mu_i(u_i, x_{-i})}{\sigma_i(u_i, x_{-i})} \right) / \left(\sum_{j \in N} w_{ij} \right) \\ &= \frac{1}{\sigma_i(u_i, x_{-i})} \left(\left(\sum_{j \in N} w_{ij} u_i(x_j) \right) / \left(\sum_{j \in N} w_{ij} \right) - \mu_i(u_i, x_{-i}) \right) = 0,\end{aligned}$$

and, similarly,

$$\begin{aligned}\sigma_i(\bar{u}_i, x_{-i}) &= \left(\sum_{j \in N} w_{ij} (\bar{u}_i(x_j) - \mu_i(\bar{u}_i, x_{-i}))^2 \right) / \left(\sum_{j \in N} w_{ij} \right) \\ &= \left(\sum_{j \in N} w_{ij} \left(\frac{u_i(x_j) - \mu_i(u_i, x_{-i})}{\sigma_i(u_i, x_{-i})} - 0 \right)^2 \right) / \left(\sum_{j \in N} w_{ij} \right) \\ &= \frac{1}{\sigma_i(u_i, x_{-i})^2} \left(\sum_{j \in N} w_{ij} (u_i(x_j) - \mu_i(u_i, x_{-i}))^2 \right) / \left(\sum_{j \in N} w_{ij} \right) = \frac{\sigma_i(u_i, x_{-i})^2}{\sigma_i(u_i, x_{-i})^2} = 1.\end{aligned}$$

Therefore, $\bar{u}_i$ is locally socially normalized. $\square$

In the rest of the section, we assume that $\sigma_i(u_i, x, w_i) > 0$ for all $i \in N$ (the case $\sigma_i(u_i, x, w_i) = 0$ can be dealt with following the same approach as in the previous section). Next, we straightforwardly extend the measures of player i 's complaint-potential reference $\bar{h}_i$ and unfairness feeling $\mathcal{F}_i$ as follows:

$$\begin{aligned}\bar{h}_i &= \min \left(1, \left(\sum_{j \in N} w_{ij} h(\bar{u}_j(x_j)) \right) / \left(\sum_{j \in N} w_{ij} \right) \right), \\ \mathcal{F}_i &= \left(h(\bar{u}_i(x_i)) - \bar{h}_i \right)^+.\end{aligned}$$

Note that now, player i 's complaint-potential reference is specific to her, which is consistent with the idea that each player is only concerned by the outcomes for people to whom she is connected. A definition of the degree of local social fairness follows.

Definition 5 *The degree of local social fairness (LSF) is given by*

$$\mathcal{G}_{LSF}(u, x, w) = \frac{1}{n} \sum_{i \in N} (h(\bar{u}_i(x_i)) - \bar{h}_i)^+. \quad (28)$$

A division is locally socially fair (LSF) if $\mathcal{G}_{LSF}(u, x, w) = 0$.

Now, we can adapt the definitions of fairness used in the cake-cutting-problem literature to our “local” or social network setting.

Definition 6 *A division is*

1. *locally exact* if $u_i(x_i) = u_i(x_j)$ whenever j is a neighbor of i , that is, $w_{ij} > 0$;
2. *locally proportionally fair (LPF)* if $u_i(x_i) \geq \mu_i$ for all i who have at least one neighbor;
3. *locally envy-free (LEF)* if $u_i(x_i) \geq u_i(x_j)$ whenever j is a neighbor of i .

The following theorem characterizes the relationships between these definitions.

Theorem 3 *We have the following implications:*

1. $EF \Rightarrow LEF \Rightarrow LPF \Rightarrow \mathcal{G}_{LSF} \leq 1/4$.
2. $LSF \Rightarrow LPF$.
3. $exact \Rightarrow locally\ exact \Rightarrow LEF, LPF\ and\ LSF$.

Proof. The implications $EF \Rightarrow LEF \Rightarrow LPF$ are immediate. The bound on LSF is then derived similarly to what we did in Theorem 2. Similarly, the second implication is based on the same grounds as Proposition 5. Now, the third implications ($exact \Rightarrow locally\ exact \Rightarrow LEF$ and LPF) are trivial. But then, all local social normalizations are undefined, which, following Section 4.2, implies LSF . $\square$

We recapitulate all connections between the fairness concepts in Figure 3.




Figure 3: Relations between the different fairness definitions.

Remark 3 *An important observation is that proportional fairness does not imply local proportional fairness. To illustrate this counterintuitive result, consider the following utility and social-network (weights) matrices:*

$$U = \begin{pmatrix} 0 & 1 & -5 & -5 \\ 1 & 0 & -5 & -5 \\ -5 & -5 & 0 & 1 \\ -5 & -5 & 1 & 0 \end{pmatrix} \quad \text{and} \quad W = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (29)$$

It is easy to check that the division is PF but not LPF. This result has an interesting interpretation in behavioral terms. It shows that a division, while being globally fair in the usual PF sense, fails to satisfy some players who only care about how other, specific players are treated.

Let us finish by illustrating LSF in the sugar-cake example. Recall that in this example, players 1, 2 and 3 had the same preferences, while player 4's preferences were the opposite of the three other players'. A formal way to see this is to compute correlations between any two players' rows in the utility matrix. Given this, it makes sense to assume that players 1, 2 and 3 are socially connected, but not 4. In other words, we have the following social network matrix

$$W = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (30)$$

The fourth player has no player to compare herself too and thus cannot feel that the division is unfair. Meanwhile, the division is locally exact for players 1, 2 and 3; hence, they feel that the division is fair too. Overall, the division is thus LSF.

6 Conclusion

In this paper, we proposed an alternative approach to existing ways of measuring fairness. The first step was to introduce a new way of normalizing the players' utility functions. Our social normalization for each player consists of using other players' allocations as a benchmark to which the player will compare her own allocation. Importantly, she does so through her own lens, that is, her utility function. The socially normalized utility function of a player is the affine positive transformation of her utility function, such that the average utility for others' allocations equals 0, and the standard deviation is 1. Thereby, the player's socially normalized utility counts the number of standard deviations above the average that she thinks her allocation is, compared to the distribution of others' allocations. We argue that such a normalization gives a relevant meaning to utility values, hence making two different utility functions comparable. Then, roughly, a division will be socially fair if all players have equal socially normalized utilities.

Now, in many large-scale optimization problems, a principal agent is in charge of solving the allocation problem. Yet, there may be no feasible fair solution. For instance, in shift scheduling, one employee will have to close the shop on Saturday evening, even if no employee wants to. Besides, even if there are fair solutions, the NP-hardness of many allocation problems indicates that computing these fair solutions may not be feasible in a reasonable amount of time. Moreover, concerns other than fairness may enter in play, e.g., social efficiency. These aspects then require a comparison of unfair allocations. We have proposed to do so by measuring players' unfairness feelings. These unfairness feelings are computed in two steps. First, for a given player, we compute her complaint potential as a decreasing and convex function of her socially normalized utility. We then argue that a player will actually feel unfairness only if her complaint potential exceeds that of others. If it does, then her unfairness feeling is defined as the difference between her complaint potential and the average of the others'. The degree of social fairness of an allocation is then given by the average unfairness feeling.

While this approach to measuring fairness is reasonable in many cases, it may not be implementable when the number of players is very large. To illustrate, Air Canada's crew-scheduling problem involves 11,000 pilots. It is then more realistic to consider that in such cases, players do not actually compare their allocations to all others, including for information-access reasons. Most likely, players instead feed their complaint potential by focusing on the allocations of "socially close" or "similar" players. By introducing a social network and by adapting concepts to this setting where interactions between players are limited, we have introduced local fairness concepts that appear to be much more suitable for and relevant to describing fairness in large-scale contexts. In particular, the correlation of players' preferences can be used as an indicator of how much they will be using one another's allocation to derive their complaint potentials. In an upcoming research project, we intend to apply these concepts to a concrete instance of shift scheduling, where a manager aims to establish a economically competitive, socially efficient and socially fair shift-allocation for her employees.

7 Appendix

Lemma 1 (for Theorem 2) *Let $h_1, \dots, h_n \in [0, 1]$, and denote $\bar{h}$ their average. Then,*

$$\sum_{i \in N} |h_i - \bar{h}| \leq n/2. \quad (31)$$

Proof. Let N_+ (respectively, $N_=$ and N_-) the subsets of N such that $h_i > \bar{h}$ (respectively, $h_i = \bar{h}$ and $h_i < \bar{h}$). Denote $f(h_1, \dots, h_n) = \sum_{i \in N} |h_i - \bar{h}|$. If $h_i \geq \bar{h}$, then the right partial derivative of f with respect to h_i equals

$$\frac{\partial f}{\partial h_i^+} = 1 - \frac{|N_+| - |N_+| - |N_-|}{n} > 0. \quad (32)$$

Similarly, if $h_i \leq \bar{h}$, then the left partial derivative with respect to h_i is going to be negative. This means that the first-order condition cannot hold for $h_i \in]0, 1[$. Thus, the maximum of f is necessarily reached for $(h_1, \dots, h_n) \in \{0, 1\}^n$. But then, $\bar{h} = |N_+|/n$. Therefore,

$$f(h_1, \dots, h_n) = |N_+| \left(1 - \frac{|N_+|}{n}\right) + |N_-| \frac{|N_+|}{n} = \frac{2|N_+|(n - |N_+|)}{n}, \quad (33)$$

whose maximum is reached for $|N_+| = n/2$, in which case the right term equals $n/2$. This gives us the upper bound of the Lemma. $\square$

References

- C. Arnsperger. Envy-freeness and distributive justice. *Journal of Economic Surveys*, 8(2):155–186, 1994.
- X. Bei, N. Chen, X. Hua, B. Tao, and E. Yang. Optimal proportional cake cutting with connected pieces. In *AAAI*, 2012.
- D. Bertsimas, V.F. Farias, and N. Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.
- J.B. Brams and A.D. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102:9–12, 1995.
- J.B. Brams and A.D. Taylor. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996.
- Y. Chen, J.K. Lai, D.C. Parkes, and A.D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.
- Y.J. Cohler, J.K. Lai, D.C. Parkes, and A.D. Procaccia. Optimal envy-free cake cutting. In *AAAI*, 2011.
- D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- R.S. Mirchandani. Superadditivity and subadditivity in fair division. *Journal of Mathematics Research*, 5(3):78, 2013.
- E. Mossel and O. Tamuz. Truthful fair division. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory*, pages 288–299, 2010.
- A.D. Procaccia. Thou shalt covet thy neighbor’s cake. In *IJCAI*, pages 239–244, 2009.
- J. Robertson and W. Webb. *Cake-Cutting Algorithms: Be Fair If You Can*. AK Peters Ltd, 1998.
- H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- W. Stromquist. Envy-free cake divisions cannot be found by finite protocols. In *Fair Division*, 2007.
- H. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.
- H.R. Varian. Two problems in the theory of fairness. *Journal of Public Economics*, 5(3):249–260, 1976.
- J. Von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1944.