

The Atkinson Inequality Index in Multiagent Resource Allocation

Sebastian Schneckenburger
University of Tübingen
Germany
sebastian.schneckenburger
@uni-tuebingen.de

Britta Dorn
University of Tübingen
Germany
britta.dorn
@uni-tuebingen.de

Ulle Endriss
University of Amsterdam
The Netherlands
ulle.endriss@uva.nl

ABSTRACT

We analyse the problem of finding an allocation of resources in a multiagent system that is as fair as possible in terms of minimising inequality between the utility levels enjoyed by the individual agents. We use the well-known Atkinson index to measure inequality and we focus on the distributed approach to multiagent resource allocation, where new allocations emerge as the result of a sequence of local deals between groups of agents agreeing on an exchange of some of the items in their possession. Our results show that it is possible to design systems that provide theoretical guarantees for optimal outcomes that minimise inequality, but also that in practice there are significant computational hurdles to be overcome: finding an optimal allocation is computationally intractable—independently of the approach chosen—and large numbers of potentially highly complex deals may be required under the distributed approach. From a methodological point of view, while much work in multiagent resource allocation relies on combinatorial arguments, here we use insights from basic calculus.

Keywords

Multiagent Resource Allocation; Fair Division; Inequality Indices

1. INTRODUCTION

“What thoughtful rich people call the problem of poverty, thoughtful poor people call with equal justice a problem of riches.”

—Anthony B. Atkinson (1944 – 2017), *Inequality* [3]

Allocating resources to agents is one of the central tasks arising in most multiagent systems [9]. This is true not only for systems of economic agents who need to share the value they have generated together, but also for distributed systems of problem-solving agents who need to share the computational resources available to them. What makes a ‘good’ allocation heavily depends on the application at hand, but there is broad consensus in the multiagent systems research community that, rather than coming up with new *ad hoc* criteria for optimality for every new application, it is fruitful to base

the design of a multiagent system on well-understood formal criteria originally proposed in the literature on social choice theory and welfare economics [20].

For instance, if an *efficient* allocation is sought, both the notion of *utilitarian social welfare*, measuring quality in terms of the sum of the individual utilities, and the weaker notion of *Pareto optimality* have been found to be useful [26]. If *fairness* is a relevant design objective, there is a much wider range of concepts to choose from, several of which have been analysed in the literature on multiagent systems in some detail: e.g., *egalitarian social welfare*, measuring quality as utility of the worst-off agent, and its refinement the *leximin-ordering* [14, 6], *Nash social welfare*, measuring quality as the product of the individual utilities [24, 23], and the *absence of envy* [17, 10]. However, fairness criteria based on measuring inequality, which are widely used in the social sciences [16, 2, 27], to date have received almost no attention in the multiagent systems literature [19, 13].

To help close this gap, in this paper, we focus on one of the most important representatives of this family of criteria, the *Atkinson inequality index*, and analyse how to achieve allocations of resources that are optimal relative to this criterion. Our main contributions concern the challenge of ensuring convergence of an optimal allocation under the *distributed approach*, where the goal is to obtain a good allocation by means of a sequence of local exchanges of items between (typically small) groups of agents [25, 14, 12, 10]. In addition, we analyse the computational complexity of computing an optimal allocation that minimises inequality, which is relevant independently of the specific approach chosen for performing multiagent resource allocation. Our results show that, in principle, an appropriately designed system can be made to guarantee outcomes with minimal inequality amongst the agents, although in practice significant computational hurdles may have to be overcome. Specifically, we may require arbitrarily complex deals and we may require an exponential number of deals. From a methodological point of view, while much work in multiagent resource allocation relies on combinatorial arguments, here we specifically rely on insights from basic calculus.

The remainder of this paper is organised as follows. In Section 2 we introduce the model of multiagent resource allocation with indivisible goods we shall be working with and then recall the relevant definitions from the theory of inequality measurement. Section 3 contains a simple complexity result that clarifies the computational challenges involved in minimising inequality amongst agents in a multiagent resource allocation problem. Our main contributions

are presented in Section 4, where we set up a resource allocation framework that allows agents to compute an optimal allocation minimising inequality in a distributed manner, by means of implementing a number of local deals. Our technical results concern the guaranteed convergence to an optimal outcome as well as the aforementioned limitations of the framework. Section 5 concludes with a brief outlook on future directions of research in this domain.

2. PRELIMINARIES

In this section, we first introduce the basic setting of multi-agent resource allocation widely used in the literature [9, 5], where a number of indivisible goods need to be distributed amongst a number of agents who each have their own preferences over which bundles of goods to obtain. We then review the relevant definitions regarding inequality measurement from the literature on welfare economics [2, 27, 20], adapting them to the setting of indivisible goods [13].

2.1 Multiagent Resource Allocation

Let $\mathcal{N} = \{1, \dots, n\}$ be a finite set of *agents*, i.e., $n = |\mathcal{N}|$, and let \mathcal{G} be a finite set of *goods*, with $m = |\mathcal{G}|$. We refer to the elements of the power set $2^{\mathcal{G}}$ as *bundles*. An *allocation* is a function $A : \mathcal{N} \rightarrow 2^{\mathcal{G}}$, mapping agents to the bundles they obtain, with $A(i) \cap A(j) = \emptyset$ for any $i \neq j$.

Every agent $i \in \mathcal{N}$ is equipped with a *utility function* $u_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}_{\geq 0}$, mapping any bundle she might receive to the (nonnegative) utility she attaches to that bundle. We use $u_i(A)$ as a shorthand for $u_i(A(i))$, the utility enjoyed by agent i under allocation A . Every allocation A induces a *utility vector* $\mathbf{u}(A) = (u_1(A), \dots, u_n(A))$. The collection of all n utility functions is denoted by \mathcal{U} . A *scenario* is a triple $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$. A collection \mathcal{U} (and also the corresponding scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$) is called *additive* if $u_i(B) = \sum_{x \in B} u_i(\{x\})$ for all agents $i \in \mathcal{N}$ and bundles $B \subseteq \mathcal{G}$.

The *utilitarian social welfare* of allocation A is defined as $sw_{util}(A) = \sum_{i \in \mathcal{N}} u_i(A)$, and the closely related *mean value* of A as $\mu(A) = \frac{1}{n} sw_{util}(A)$. They reflect the economic efficiency of A . The *Nash social welfare* $sw_{nash}(A) = \prod_{i \in \mathcal{N}} u_i(A)$ of allocation A can be used to measure fairness.

2.2 Inequality Indices

One way of comparing allocations consists in considering the *inequality* induced by the corresponding utility vectors. This can be measured by a so-called *inequality index*, which is a function mapping allocations (or, equivalently, utility vectors) to the interval $[0, 1]$, where 0 stands for perfect equality (meaning that all agents receive the same utility, which must not be 0). High values close or equal to 1 stand for high inequality amongst the agents.

Famous examples of inequality indices are the *Gini index* [16], the *Robin Hood index* [18] (also called the *maximum relative mean deviation*), and the family of *Atkinson indices* [2].

Every Atkinson index relies on a notion of social welfare: For a given function sw mapping utility vectors to their social welfare and for a utility vector $\mathbf{u}(A)$ of an allocation A , first compute the so-called *equally distributed equivalent level of income* $\mu_{sw}(A)$, such that the vector $(\mu_{sw}(A), \dots, \mu_{sw}(A))$ has the same social welfare as $\mathbf{u}(A)$. The Atkinson index based on the function sw is then defined as $I_{sw}(A) := 1 - \frac{\mu_{sw}(A)}{\mu(A)}$ [2].

We will focus on the most important representative of this family, the Atkinson index based on the Nash social welfare:

$$I_{nash}(A) = 1 - \frac{\sqrt[n]{sw_{nash}(A)}}{\mu(A)} = 1 - \frac{\sqrt[n]{\prod_{i \in \mathcal{N}} u_i(A)}}{\frac{1}{n} \sum_{i \in \mathcal{N}} u_i(A)},$$

with $I_{nash}(A) = 0$ if all individual utilities are 0.

While in the literature the term ‘Atkinson index’ is used both for the family and for this concrete one, here we only use it in this latter sense. From now on, we will use the notation \mathcal{I} instead of I_{nash} . It is easy to see that \mathcal{I} returns 0 if all the agents receive the same utility. Furthermore, we can show that it never returns 0 in any other case:

LEMMA 1. *If $\mathcal{I}(A) = 0$ for an allocation A , then all agents receive the same utility, i.e.,*

$$\mathcal{I}(A) = 0 \implies \forall i \in \mathcal{N} : u_i(A) = \mu(A).$$

PROOF. The assertion follows from the inequality for the arithmetic and the geometric mean, i.e.,

$$\frac{1}{n} \left(\sum_{k=1}^n x_k \right) \geq \sqrt[n]{\prod_{k=1}^n x_k}$$

for any nonnegative real numbers x_1, \dots, x_n , with equality if and only if $x_1 = x_2 = \dots = x_n$. A proof can be found in Cauchy’s *Analyse Algébrique* [8, pp. 457]. \square

We focus on the Atkinson index, because of its importance in the literature in the social sciences [2, 27, 1, 20]. While some other indices, notably the Gini index, are more widely used, the Atkinson index is often considered to be preferable on normative grounds, due to its principled formulation in terms of a notion of social welfare—in our case, Nash social welfare, which itself enjoys sound axiomatic foundations, going back all the way to the seminal work of Nash [21, 27, 20, 7]. Furthermore, the Atkinson index fulfils the common basic axioms for inequality indices which include the transfer principle, symmetry, and scale invariance [11, 2, 1]. The *transfer principle* states that transfers from an agent with a high utility to one with low utility shall not increase the inequality (if their order is maintained). An inequality index I is called *symmetric* if $I(\mathbf{u}(A)) = I(p_{\mathbf{u}(A)})$ holds for any permutation $p_{\mathbf{u}(A)}$ of a utility vector $\mathbf{u}(A)$ (meaning that the entries of $\mathbf{u}(A)$ are permuted). Finally, *scale invariance* means that multiplication of all utilities by a (positive) constant factor has no effect on the measured inequality.

3. COMPUTATIONAL COMPLEXITY

It is clearly desirable to find allocations that minimise the inequality amongst the agents. One might in particular ask whether, for a given scenario, there exists an allocation that is perfectly equal. In this section, we consider the computational complexity of this problem when inequality is measured in terms of the Atkinson index. The PERFECT INDEX OPTIMISATION problem is defined as follows:

PERFECT INDEX OPTIMISATION (PIO)	
Instance:	$\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$
Question:	\exists allocation $A : \mathcal{I}(A) = 0?$

Unfortunately, it turns out that this problem is NP-hard:

PROPOSITION 2. *The decision problem PIO is NP-hard, even for additive scenarios with just two agents.*

PROOF. First, note that, by Lemma 1, we have $\mathcal{I}(A) = 0$ if and only if the two agents enjoy the same level of utility. We use a reduction from the NP-hard PARTITION problem [15], which is defined as follows:

PARTITION	
Instance:	A finite set X , and a size $s(x) \in \mathbb{Z}_{\geq 0}$ for each $x \in X$.
Question:	Is there a subset $X' \subseteq X$ such that $\sum_{x \in X'} s(x) = \sum_{x \in (X \setminus X')} s(x)$?

Given an instance $\langle X, \{s(x) \mid x \in X\} \rangle$ of the PARTITION problem, we construct an instance $\langle \{1, 2\}, \mathcal{G}, \mathcal{U} \rangle$ of PIO, where $\mathcal{G} = X$ and $u_1(B) = u_2(B) = \sum_{b \in B} s(b)$ for all $B \subseteq \mathcal{G} = X$. Then it is easily checked that $\langle X, \{s(x) \mid x \in X\} \rangle$ is a YES-instance if and only if $\langle \{1, 2\}, \mathcal{G}, \mathcal{U} \rangle$ is. \square

The PARTITION problem has previously been used to analyse the complexity of problems arising in multiagent resource allocation, such as the problem of finding an allocation that maximises egalitarian social welfare when utilities are additive [4, 22] and the problem of determining whether a given allocation permits an inequality reduction by means of a transfer between two agents with additive utilities [13]. We remark that NP-hardness for a less restricted scenario (unbounded number of agents, symmetric utilities) of PIO can also be shown by reduction from the EXACT COVER BY 3-SETS problem [15].

4. THE DISTRIBUTED APPROACH

As we have seen, the problem of deciding whether there exists an allocation with perfect equality is computationally intractable already for very restricted instances. Thus, computing such an allocation will be just as hard. Nevertheless, we are interested in minimising inequality amongst the agents. To this end, we will now explore adapting the so-called *distributed approach* formulated by Endriss et al. [14], relying on ideas originally introduced by Sandholm [25]. Under this approach, starting from some initial allocation, the agents can decide to arrange exchanges of some of the goods between some of them by means of so-called *deals*. The key idea is that the agents are supposed to only use local information: only some (preferably small number of) agents may be involved in a deal and they only have access to information on the goods they own and on the goods they exchange, not on the overall allocation. The goal is to devise a protocol for the agents to follow that, despite this limitation to local deals, permits them to negotiate an allocation with good global properties. This approach has been successfully applied to compute, in a distributed manner, allocations that are optimal in view of, amongst others, utilitarian social welfare [25], egalitarian social welfare [14], Nash social welfare [23], and envy-freeness [10].

After defining the notion of a deal formally (in Section 4.1), we will first prove that achieving convergence to an allocation with minimal inequality is impossible for deals that are *local* in the narrow sense in which this term has been defined in the literature before (see Section 4.2). However, we will then see that a very mild relaxation of this notion of locality is sufficient to obtain a convergence result (see




Figure 1: Scenario of Example 1 with a deal.

Section 4.3). This positive result is then tempered by two further results. First, we show that we must admit arbitrarily complex (yet semi-local) deals (see Section 4.4), and we must allow for the possibility of exponentially long sequences of deals before convergence is realised (see Section 4.5).

4.1 Deals and Sequences of Deals

A *deal* $\delta = (A, A')$ is a pair of two (distinct) allocations A and A' . The set of agents *involved* in the deal δ is denoted by \mathcal{N}^δ , i.e., $\mathcal{N}^\delta := \{i \in \mathcal{N} \mid A(i) \neq A'(i)\}$.

Example 1. Consider the set of goods $\mathcal{G} = \{a, b, c, d\}$, the set of agents $\mathcal{N} = \{1, 2, 3\}$, and the two allocations $A^\circ = (\{a, b\}, \{c\}, \{d\})$ and $A^* = (\{a, c\}, \{b\}, \{d\})$. The deal $\delta = (A^\circ, A^*)$ with involved agents $\mathcal{N}^\delta = \{1, 2\}$, in which agent 1 gives item b to agent 2 and receives item c in return, is visualised in Figure 1.

Note that a single deal may include any number of agents and goods (even if we think of a *typical* deal as involving just a few of each). We would like the agents to agree on a sequence of deals that—somehow—converges to an allocation that minimises inequality. Let us first exclude two approaches that are definitely not useful. First, we could give the agents complete freedom what deals to negotiate. This protocol cannot ensure convergence, as we cannot exclude the possibility of loops (e.g., they may indefinitely alternate between A° and A^* of Figure 1). Second, from any given allocation we could only permit a single deal, namely the deal that takes us straight to the optimal allocation. This also is not useful, as it would not leverage any of the potential power of the distributed approach and simply reduce it to a fully centralised optimisation problem.

4.2 No Convergence by Local Deals

We are looking for a criterion to select admissible deals such that (i) any sequence of admissible deals eventually leads to an optimal allocation and (ii) the agents involved in any given deal are able to determine locally whether that deal is admissible. But how should we define ‘locality’ in this context? Endriss et al. [14] call a criterion for determining the admissibility of a deal $\delta = (A, A')$ *local* if and only if the question of whether δ is local can be answered by looking only at the set $\{(i, u_i(A), u_i(A')) \mid i \in \mathcal{N}^\delta\}$. In other words, admissibility should only depend on the utility levels of the agents involved before and after the deal.

Unfortunately, it is impossible to define a deal selection criterion that is local in this sense and that could be used to guide our search for an optimal allocation by only ever admitting deals that reduce inequality:

PROPOSITION 3. *It is impossible to always decide whether a given deal $\delta = (A, A')$ would decrease inequality as defined by the Atkinson index by only inspecting the utility levels of the agents involved in δ in allocations A and A' .*

PROOF. We construct an example where a given deal would decrease inequality in one scenario but increase it in another, while the local information on the utility levels of the agents involved in that deal is the same in both scenarios. Consider the two scenarios $\langle \mathcal{N}, \mathcal{G}, \mathcal{U}_1 \rangle$ and $\langle \mathcal{N}, \mathcal{G}, \mathcal{U}_2 \rangle$, with $\mathcal{N} = \{\mathbf{1}, \mathbf{2}, \mathbf{3}\}$ and $\mathcal{G} = \{a, b, c, d\}$. The additive collections of utility functions \mathcal{U}_1 and \mathcal{U}_2 are defined in terms of the values the agents assign to each of the items:

$$\left| \begin{array}{c|cccc} \mathcal{U}_1 & a & b & c & d \\ \hline \mathbf{1} & 2 & 1 & 3 & 4 \\ \mathbf{2} & 2 & 5 & 2 & 1 \\ \mathbf{3} & 1 & 2 & 1 & 6 \end{array} \right| \quad \left| \begin{array}{c|cccc} \mathcal{U}_2 & a & b & c & d \\ \hline \mathbf{1} & 2 & 1 & 3 & 4 \\ \mathbf{2} & 2 & 5 & 2 & 1 \\ \mathbf{3} & 3 & 2 & 3 & 2 \end{array} \right|$$

Now consider the deal $\delta = (A^\circ, A^*)$ between allocations $A^\circ = (\{a, b\}, \{c\}, \{d\})$ and $A^* = (\{a, c\}, \{b\}, \{d\})$, which is the same deal we had already considered in Figure 1. Let us compute the Atkinson index for each of the two allocations in each of the two scenarios:

	Scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U}_1 \rangle$	Scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U}_2 \rangle$
$\mathcal{I}(A^\circ)$:	$1 - \frac{\sqrt[3]{3 \cdot 2 \cdot 6}}{\frac{1}{3} \cdot (3+2+6)} \approx 0.099$	$1 - \frac{\sqrt[3]{3 \cdot 2 \cdot 2}}{\frac{1}{3} \cdot (3+2+2)} \approx 0.019$
$\mathcal{I}(A^*)$:	$1 - \frac{\sqrt[3]{5 \cdot 5 \cdot 6}}{\frac{1}{3} \cdot (5+5+6)} \approx 0.004$	$1 - \frac{\sqrt[3]{5 \cdot 5 \cdot 2}}{\frac{1}{3} \cdot (5+5+2)} \approx 0.079$

Thus, in the first scenario, δ decreases inequality, while in the second scenario, δ increases inequality. Nevertheless, the two agents involved in δ cannot distinguish between the two scenarios. Hence, there can be no local criterion for the admissibility of deals that would allow us to always select deals that decrease inequality. \square

For comparison, when optimality is defined in terms of utilitarian social welfare, egalitarian social welfare, or Nash social welfare, local criteria for selecting deals that ensure a social improvement do exist [14, 23]. When the goal is to compute an envy-free allocation, there exists no suitable local criterion, but this hurdle can be overcome by slightly relaxing the requirements [10]. We shall follow a similar route.

4.3 Convergence by Semi-Local Deals

Recall that the computation of the Atkinson index involves both the geometric mean and the arithmetic mean of the utilities of all agents. On the one hand, the local information on the utility levels of the involved agents is sufficient to determine both whether (i) the geometric mean increases or decreases, and whether (ii) the arithmetic mean increases or decreases.¹ On the other hand, the underlying reason for the impossibility stated in Proposition 3 is that, nevertheless, this local information is not sufficient to determine which of these two effects is stronger, and thus whether inequality will increase or decrease.

We now define a *semi-local* criterion for the admissibility of deals that relaxes the constraints on the information available a little and thereby allows us to overcome this problem.

¹This is precisely the reason why it is possible to design local criteria for agents wishing to compute allocations with maximal Nash and utilitarian social welfare, respectively.

The central idea is to allow the agents to also access $\mu(A)$, the (arithmetic) mean of the utilities of *all* agents (not just the involved agents) before the deal. Given $\mu(A)$ and the usual local information, we can compute $\mu(A')$ for another allocation A' reached by the deal $\delta = (A, A')$ as follows:

$$\mu(A') = \mu(A) + \frac{1}{n} \cdot \sum_{i \in \mathcal{N}^\delta} (u_i(A') - u_i(A))$$

We still do not have full access to the geometric mean of all utilities, but only to the extent to which it changes during the deal. As will become clear shortly, this is not a problem.

Let us call a deal $\delta = (A, A')$ an *Atkinson deal* if and only if it satisfies the following condition:

$$\frac{\sqrt[n]{\prod_{i \in \mathcal{N}^\delta} u_i(A)}}{\mu(A)} > \frac{\sqrt[n]{\prod_{i \in \mathcal{N}^\delta} u_i(A')}}{\mu(A) + \frac{1}{n} \cdot \sum_{i \in \mathcal{N}^\delta} (u_i(A') - u_i(A))}$$

Observe that we can determine whether a given deal is an Atkinson deal using semi-local information only: we require the utility levels in A and A' for the involved agents as well as the mean value of the entire society in A . The good news is that this is sufficient to allow us to compute an optimal allocation in a distributed manner:

THEOREM 4. *For every scenario and initial allocation, every sequence of Atkinson deals will eventually result in an allocation that minimises inequality, as defined by the Atkinson index.*

PROOF. First, observe that a deal decreases inequality if and only if it is an Atkinson deal (this is immediate from the definitions of Atkinson index and Atkinson deals).

As there are only a finite number of allocations, any sequence without cycles has to terminate eventually. As every deal in the sequence strictly decreases inequality, there cannot be any cycles, which proves termination. Finally, it is impossible for the terminal allocation A to not have minimal inequality, as then there would have to exist another allocation A' with lower inequality, which would make the deal $\delta = (A, A')$ an Atkinson deal, i.e., A could not have been terminal in the first place. \square

Similar convergence results have been proved for a number of other criteria for social optimality [25, 14, 23, 10]. In some cases, notably for utilitarian social welfare and envy-freeness [25, 10], the admissibility criterion for deals has an attractive interpretation as a *rationality criterion* for *selfish agents*. For example, in the case of utilitarian social welfare, we obtain convergence by means of deals for which myopic agents with quasi-linear utilities can negotiate prices that benefit all agents involved in the deal. In other cases, notably for egalitarian social welfare and Nash social welfare [14, 23], just as for our result here, convergence theorems should be interpreted as showing that *cooperative agents* can collectively compute an optimal outcome without requiring global coordination to guide their search. Specifically, Theorem 4 shows that agents can freely contract deals with their neighbours, safe in the knowledge that every single deal will improve the global situation and no deal will cut them off from a route to an optimal allocation.

4.4 Necessity of Complex Deals

Theorem 4 shows that we will always reach an allocation with minimal inequality, provided we keep on contracting new Atkinson deals as long as any such deals exist. But our result does not say anything about how complex these deals are. Ideally, we would prefer deals that involve the exchange of only a small number of goods between a small number of agents. So we may ask whether a given deal, particularly a deal of high structural complexity, might ever become *necessary* for reaching an allocation with minimal inequality. Unfortunately, this is indeed the case:

THEOREM 5. *For every deal δ , there exist utility functions and a starting allocation, such that δ is necessary for reaching an allocation that minimises inequality, as defined by the Atkinson index, by means of Atkinson deals only.*

Theorem 5 is bad news since it shows that, if we want to reach an optimal allocation by using Atkinson deals, it might be unavoidable to use very complex deals—even involving all agents and all items. The proof is omitted for lack of space. It makes use of Lemma 7 below which covers the special case where δ is not *independently decomposable* (to be defined shortly). For the remaining cases, it uses the fact that any other deal can be decomposed into a sequence of deals each of which is not independently decomposable.

In this context, a deal $\delta = (A, A'')$ is called *independently decomposable* if it concerns two separate sets of transactions between two disjoint sets of agents, i.e., if there exists a third allocation A' such that, for the deals $\delta_1 = (A, A')$ and $\delta_2 = (A', A'')$, it is the case that $\mathcal{N}^{\delta_1} \cap \mathcal{N}^{\delta_2} = \emptyset$ [14].

To prove necessity of all independently decomposable deals, we require the following technical lemma.

LEMMA 6. *For every $n \in \mathbb{N}_{>1}$, the function*

$$T : [0, 1] \rightarrow [0, 1]$$

$$x \mapsto 1 - \frac{\sqrt[n]{1-x}}{1-\frac{x}{n}}$$

is strictly monotonically increasing and thus bijective.

PROOF. T is well-defined and differentiable. Furthermore $T(0) = 0$, $T(1) = 1$, and $\frac{d}{dx}T(x) = \frac{(n-1)x(1-x)^{\frac{1}{n}-1}}{(n-x)^2} > 0$ holds for all $x \in]0, 1[$, which implies the claim. \square

We can now show that every deal that is not independently decomposable is necessary in the above sense:

LEMMA 7. *For every deal $\delta = (A, A')$ that is not independently decomposable, there exist utility functions $(u_i)_{i \in \mathcal{N}}$ and a starting allocation, such that δ is necessary for reaching an allocation that minimises inequality, as defined by the Atkinson index, by means of Atkinson deals only.*

PROOF. For the given deal $\delta = (A, A')$, we construct a utility function for every agent. As A and A' are different, there is at least one agent j with $A(j) \neq A'(j)$. We fix this j and let $0 < x < 1$. We now define the utility functions for any given bundle $B \in 2^{\mathcal{G}}$ as

$$u_i(B) = \begin{cases} 1 & \text{if } A'(i) = B, \\ 1 & \text{if } (i \neq j) \text{ and } A(i) = B, \\ 1-x & \text{if } (i = j) \text{ and } A(i) = B, \\ \frac{1}{i+1} & \text{otherwise.} \end{cases}$$

It is now easy to see that $\mathcal{I}(A') = 0$ and

$$\mathcal{I}(A) = 1 - \frac{\sqrt[n]{1-x}}{1-\frac{x}{n}}.$$

Next, we show that for any different allocation A^* we have $\mathcal{I}(A^*) > 0$. As the deal δ is not independently decomposable, there is at least one pair of agents k, ℓ with $u_k(A^*) \neq u_\ell(A^*)$: otherwise, we would have $u_i(A^*) = 1$ for all $i \in \mathcal{N}$, meaning that A^* agrees with either A or A' for every agent, i.e., δ would be independently decomposable into the deals (A, A^*) and (A^*, A') , contradicting our assumptions. Thus, by Lemma 1, we must have that $\mathcal{I}(A^*) > 0$. As there are only finitely many possible allocations, we get

$$\min_{A^* \neq A, A'} \mathcal{I}(A^*) > 0.$$

We now choose some ε with $0 < \varepsilon < \min_{A^* \neq A, A'} \mathcal{I}(A^*)$ and then set x such that $\mathcal{I}(A) = \varepsilon$, which is possible due to Lemma 6. Hence, we have $0 = \mathcal{I}(A') < \mathcal{I}(A) < \mathcal{I}(A^*)$. Thus, in this scenario, from allocation A , $\delta = (A, A')$ is the only deal reducing \mathcal{I} , and thus the only Atkinson deal. \square

Our construction used in the proof of Lemma 7 is similar to the construction used to derive necessity results for utilitarian and egalitarian social welfare [14] as well as Nash social welfare [23]. In those other settings, not only are all non-independently decomposable deals necessary, but these are the *only* such deals. So Theorem 5 is a bad surprise, as in the present setting the situation is worse and even deals that *are* independently decomposable are necessary. The following example also illustrates this fact:

Example 2. Consider the (additive) scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$, with $\mathcal{N} = \{\mathbf{1}, \mathbf{2}, \mathbf{3}, \mathbf{4}\}$ and $\mathcal{G} = \{a, b, c, d\}$. The collection \mathcal{U} of additive utility functions is defined in terms of the values the agents assign to each of the items:

\mathcal{U}	a	b	c	d
$\mathbf{1}$	4	10	4	4
$\mathbf{2}$	10	3	3	3
$\mathbf{3}$	2	2	2	10
$\mathbf{4}$	1	1	10	1

Now consider the deal $\delta = (A, A')$ between allocations $A = (\{a\}, \{b\}, \{c\}, \{d\})$ and $A' = (\{b\}, \{a\}, \{d\}, \{c\})$. This deal is decomposable; there are two possible decomposition sequences, (A, A^{i_1}, A') and (A, A^{i_2}, A') with $A^{i_1} = (\{a\}, \{b\}, \{d\}, \{c\})$ and $A^{i_2} = (\{b\}, \{a\}, \{c\}, \{d\})$.

As \mathcal{U} is additive, only allocations which assign exactly one item to each agent are not completely unfair. So there are only 4! allocations with inequality not equal to 1, but from these, only A^{i_1} , A^{i_2} , and A' have a different inequality from A . The values of \mathcal{I} for these four allocations are as follows:

$$\mathcal{I}(A) : = 1 - \frac{\sqrt[4]{4 \cdot 3 \cdot 2 \cdot 1}}{\frac{1}{4} \cdot (4+3+2+1)} \approx 0.115$$

$$\mathcal{I}(A^{i_1}) = 1 - \frac{\sqrt[4]{10 \cdot 10 \cdot 2 \cdot 1}}{\frac{1}{4} \cdot (10+10+2+1)} \approx 0.346$$

$$\mathcal{I}(A^{i_2}) = 1 - \frac{\sqrt[4]{4 \cdot 3 \cdot 10 \cdot 10}}{\frac{1}{4} \cdot (4+3+10+10)} \approx 0.128$$

$$\mathcal{I}(A') = 1 - \frac{\sqrt[4]{10 \cdot 10 \cdot 10 \cdot 10}}{\frac{1}{4} \cdot (10+10+10+10)} = 0$$

So in this example, given the allocation A , the (independently decomposable) deal $\delta = (A, A')$ is necessary.

Finally, we are able to show that deals involving all agents and all goods can become necessary even for *additive* scenarios—at least when the number of goods equals or exceeds the number of agents. Again, we omit all details for lack of space.

4.5 Path Length to Convergence

In this section, we are interested in the number of deals needed to reach an optimal allocation. It is clear that, given a starting allocation A , it is always possible to reach an optimal allocation A_{opt} with at most one Atkinson deal: just use the deal as $\delta = (A, A_{\text{opt}})$ —unless A already is optimal and no deal is needed. It thus is more interesting to ask how long a sequence of Atkinson deals from an initial to an optimal allocation can be in the worst case. It is easy to establish an upper bound: First, observe that there are n^m possible allocations (recall that $n = |\mathcal{N}|$ and $m = |\mathcal{G}|$). Second, observe that, since every Atkinson deal strictly reduces inequality, we cannot visit any allocation twice. Hence, there can be at most $n^m - 1$ deals in total. We will show that there are scenarios for which this theoretical maximum can in fact be reached. To do so, we will construct a scenario where no two allocations produce the same inequality. We start showing this for the case of two agents in Lemma 8, before we proceed to the general case of this assertion in Lemma 11.

LEMMA 8. *For two agents and m goods, $m \in \mathbb{N}$, it is possible to define utility functions such that any two distinct allocations have a different value of \mathcal{I} .*

PROOF. The proof of this lemma is inspired by the proof of Lemma 1 in the work of Ramezani and Endriss [23] for the Nash social welfare. We assign to agents 1 and 2 the prime numbers 2 and 3, respectively. Now suppose each agent has an ordering on all possible 2^m bundles, and $u_1(B) = 2^j$ if B is the j^{th} bundle in the first agent's ordering. Analogously, let $u_2(C) = 3^j$ if C is the j^{th} bundle in the second agent's ordering. In an allocation A , agent i receives the $(j_A^i)^{\text{th}}$ bundle in his ordering. It is easy to see that any two allocations A and A' have different Nash social welfare, since

$$sw_{\text{nash}}(A) = 2^{j_A^1} \cdot 3^{j_A^2} = 2^{j_{A'}^1} \cdot 3^{j_{A'}^2} = sw_{\text{nash}}(A')$$

would imply directly $j_A^1 = j_{A'}^1$ and $j_A^2 = j_{A'}^2$ due to the unique prime factorisation of every integer.

Now we will show that also $\mathcal{I}(A) = \mathcal{I}(A')$ implies $A = A'$:

$$\begin{aligned} & \mathcal{I}(A) = \mathcal{I}(A') \\ \implies & 1 - \frac{\sqrt{sw_{\text{nash}}(A)}}{\mu(A)} = 1 - \frac{\sqrt{sw_{\text{nash}}(A')}}{\mu(A')} \\ \implies & \frac{\sqrt{2^{j_A^1} \cdot 3^{j_A^2}}}{\frac{1}{2} \cdot (2^{j_A^1} + 3^{j_A^2})} = \frac{\sqrt{2^{j_{A'}^1} \cdot 3^{j_{A'}^2}}}{\frac{1}{2} \cdot (2^{j_{A'}^1} + 3^{j_{A'}^2})} \\ \implies & 2^{j_A^1} \cdot 3^{j_A^2} \cdot (2^{j_{A'}^1} + 3^{j_{A'}^2})^2 = 2^{j_{A'}^1} \cdot 3^{j_{A'}^2} \cdot (2^{j_A^1} + 3^{j_A^2})^2. \end{aligned}$$

As $(2^j + 3^{j'}) \equiv 0 \pmod{2}$ can never hold for any $(j, j') \in \{1, \dots, 2^m\}^2$ and also $(2^j + 3^{j'}) \equiv 0 \pmod{3}$ can never hold for any $(j, j') \in \{1, \dots, 2^m\}^2$, the unique prime factorisation of each side of the last equation leads again directly to $j_A^1 = j_{A'}^1$ and $j_A^2 = j_{A'}^2$, which implies $A = A'$. \square

We will require the following technical lemma:

LEMMA 9. *Let $n \in \mathbb{N}_{\geq 2}$ and $j_1, j_2, k_1, k_2 > 0$. Then*

$$1 - \frac{\sqrt[n]{2^{j_1} \cdot 3^{k_1}}}{\frac{1}{n}(2^{j_1} + 3^{k_1})} = 1 - \frac{\sqrt[n]{2^{j_2} \cdot 3^{k_2}}}{\frac{1}{n}(2^{j_2} + 3^{k_2})}$$

holds iff $j_1 = j_2$ and $k_1 = k_2$.

PROOF. This lemma can be proved by proceeding analogously to the reasoning in the proof of Lemma 8. \square

The proof of Lemma 8 cannot easily be generalised for more than 2 agents, as the argumentation with the modulo calculation does not hold any longer: $(2^i + 3^j + 5^k) \pmod{5}$ can be equal to 0 for $i = j = k$, e.g., $(2^1 + 3^1 + 5^1) \pmod{5} = 0$. Nevertheless, the closely related Lemma 9 will help us to prove the generalised statement for any number of agents. But first, we prove one further technical lemma:

LEMMA 10. *Real functions $g_{a,b}: \mathbb{R} \rightarrow \mathbb{R}$ given by*

$$g_{a,b}(x) = \frac{a \cdot x}{(b+x)^k}$$

for $a, b > 0$ and $k \in \mathbb{N}$ can be interpolated exactly by using just two points $(x_1, c_1), (x_2, c_2)$ of the graph of the function if we restrict the function to values $x \gg b/k$.

PROOF. Given the two equations $\frac{a \cdot x_1}{(b+x_1)^k} = c_1$ and $\frac{a \cdot x_2}{(b+x_2)^k} = c_2$, eliminating a leads to $\frac{c_1 x_2}{c_2 x_1} = \left(\frac{b+x_2}{b+x_1}\right)^k$. This equation can be solved via

$$\underbrace{\sqrt[k]{\frac{c_1 x_2}{c_2 x_1}}}_{\tau} = \frac{b+x_2}{b+x_1} \implies \frac{\tau x_2 - x_1}{1-\tau} = b,$$

so the interpolation is unique (meaning the equation has a unique solution (a, b)). \square

This means that, if two functions of the above type agree on their values for two (large enough) values of x , then they already have to be identical. We will need this property for constructing utility functions that imply different values of \mathcal{I} for each possible allocation in the corresponding scenario.

LEMMA 11. *For any natural numbers n and m , there exists a scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ with $|\mathcal{N}| = n$ and $|\mathcal{G}| = m$ such that any two distinct allocations differ in inequality, as defined by the Atkinson index.*

PROOF. We consider the scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ and construct utility functions that fulfil the claim. As the elements of \mathcal{U} are the functions $u_i: 2^{\mathcal{G}} \rightarrow \mathbb{R}_{\geq 0}$, it is possible to store all the information of \mathcal{U} in the $n \times 2^m$ matrix $P = (p_{i,j})_{i=1, \dots, n}^{j=1, \dots, 2^m}$ with $p_{i,j} = u_i(B_j)$. Herby we suppose some arbitrary, but given ordering (B_1, \dots, B_{2^m}) of the elements of $2^{\mathcal{G}}$. For given $n, m \in \mathbb{N}$, we fill this matrix recursively to obtain the desired result. We start the recursion with the first two rows and the following entries:

$$P = \begin{pmatrix} 2^1 & 2^2 & \dots & 2^{2^m} \\ 3^1 & 3^2 & \dots & 3^{2^m} \\ * & * & \dots & * \\ \vdots & \vdots & \dots & \vdots \\ * & * & \dots & * \end{pmatrix}.$$

The symbol $*$ means that we have not yet fixed a value for the corresponding entry. We notice an interesting property

of this collection of 2^{m+1} real numbers. Let $p = (p(1), p(2))$ and $q = (q(1), q(2))$ be elements of

$$\{2^1, 2^2, \dots, 2^{2^m}\} \times \{3^1, 3^2, \dots, 3^{2^m}\}.$$

Then, with the shorthand notation $\prod_p = \prod_{i=1}^2 p(i)$ and $\sum_p = \sum_{i=1}^2 p(i)$, we see that

$$1 - \frac{\sqrt[2]{\prod_p}}{\frac{1}{2} \sum_p} = 1 - \frac{\sqrt[2]{\prod_q}}{\frac{1}{2} \sum_q}$$

implies, by Lemma 9, that $p = q$. (This means, sloppily speaking, that similarly as in Lemma 8, if utilities are given by the already existing entries of P , any two different ‘allocations’ exhibit a different level of inequality—they only have the same level if they are the same. At this stage, we cannot speak of real allocations and \mathcal{I} yet, since we are only considering the utilities of bundles for a subset consisting of two agents. We therefore also cannot compute \mathcal{I} , but a similar value by taking into account only those agents that are already involved. We make this more formal below.)

We generalise this property to bigger collections of entries of the matrix P . Let $1 \leq \ell \leq n$ and $1 \leq k \leq 2^m$. Suppose we have already fixed values for the entries of the first $\ell - 1$ rows and for the first $k - 1$ entries of the ℓ th row. For every $1 \leq i \leq \ell - 1$, we define

$$P_i := \{p_{i,j} : 1 \leq j \leq 2^m\} \subset \mathbb{R}$$

as the set of entries in the i th row of P , corresponding to the utilities that agent i assigns to the possible bundles, and the Cartesian product $P^{(\ell-1)} := P_1 \times \dots \times P_{\ell-1}$. For the elements $p = (p(1), p(2), \dots, p(\ell - 1))$ of $P^{(\ell-1)}$ (consisting of one entry from each of the already filled rows of P), we use the shorthand notation $\prod_p = \prod_{i=1}^{\ell-1} p(i)$ and $\sum_p = \sum_{i=1}^{\ell-1} p(i)$.

We call a collection of the first $((\ell - 1) \cdot 2^m + k - 1)$ entries from P *feasible*, if—sloppily speaking—for every choice of one entry from each already filled row, i.e., for each set of utilities for the possible bundles, any two ‘allocations’ would exhibit a different value of inequality (again, we cannot really speak of \mathcal{I} yet, as we only consider a *partial* allocation as long as the matrix is not entirely filled, but this helps for the intuition). More formally, this means in the case of $k = 1$ (when the first entry of each row is computed) that

$$1 - \frac{\sqrt[\ell-1]{\prod_p}}{\frac{1}{\ell-1} \sum_p} = 1 - \frac{\sqrt[\ell-1]{\prod_q}}{\frac{1}{\ell-1} \sum_q}$$

implies $p = q$ and $i = j$ for any $p, q \in P^{(\ell-1)}$ and $1 \leq i, j \leq 2^m$. If $k \geq 2$, we call the collection of the $(\ell - 1) \cdot 2^m + k - 1$ already fixed entries of P *feasible* if

$$1 - \frac{\sqrt[\ell]{\prod_p \cdot p_{\ell,i}}}{\frac{1}{\ell} (\sum_p + p_{\ell,i})} = 1 - \frac{\sqrt[\ell]{\prod_q \cdot p_{\ell,j}}}{\frac{1}{\ell} (\sum_q + p_{\ell,j})}$$

implies $p = q$ and $i = j$ for any $p, q \in P^{(\ell-1)}$ and $1 \leq i, j < k$.

The recursion step now is to fix the value for $p_{\ell,k}$ such that the new collection of the $(\ell - 1) \cdot 2^m + k - 1 + 1$ then fixed entries of P is also feasible. Table 1 illustrates the situation of the recursion step. The entry to be fixed is marked by an \mathbf{x} . As we have seen in the remark after Lemma 9, just taking powers of primes is not helpful. We therefore define real functions that feature the property used in Lemma 10.

Table 1: The partially filled matrix P in the recursion step in the proof of Lemma 11.

$$\begin{pmatrix} p_{1,1} & \dots & \dots & \dots & \dots & \dots & p_{1,2^m} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ p_{\ell-1,1} & \dots & p_{\ell-1,k-1} & p_{\ell-1,k} & p_{\ell-1,k+1} & \dots & p_{\ell-1,2^m} \\ p_{\ell,1} & \dots & p_{\ell,k-1} & \mathbf{x} & * & \dots & * \\ * & \dots & * & * & * & \dots & * \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ * & \dots & * & * & * & \dots & * \end{pmatrix}$$

We will start with the recursion step for $k = 1$. We now have to fix the value for $p_{\ell,1}$. We define the family of functions $(f_p)_{p \in P^{(\ell-1)}}$ with

$$\begin{aligned} f_p : [0, \infty) &\rightarrow [0, 1] \\ x &\mapsto 1 - \frac{\sqrt[\ell]{\prod_p \cdot x}}{\frac{1}{\ell} (\sum_p + x)}. \end{aligned}$$

Any pair of distinct functions of this family cannot intersect more than once if we restrict them to a suitable interval of the form $[\pi_1, \infty[$ for some $\pi_1 \in \mathbb{R}$ (which will be determined later on). To see this, we observe the connection to Lemma 10. The equation

$$1 - \frac{\sqrt[\ell]{\prod_p \cdot x_1}}{\frac{1}{\ell} (\sum_p + x_1)} = 1 - \frac{\sqrt[\ell]{\prod_p \cdot x_2}}{\frac{1}{\ell} (\sum_p + x_2)}$$

is equivalent to

$$\frac{\prod_p \cdot x_1}{(\sum_p + x_1)^\ell} = \frac{\prod_p \cdot x_2}{(\sum_p + x_2)^\ell}.$$

Let $p, q \in P^{(\ell-1)}$ be given. If x_1 and x_2 are large enough (for example $0 < \sum_p / \ell, \sum_q / \ell \ll x_1 < x_2$), then $f_p(x_1) = f_q(x_1)$ and $f_p(x_2) = f_q(x_2)$ imply $p = q$. This is true due to Lemma 10.

Let $\pi_1 = \max_{p \in P^{(\ell-1)}} \sum_p / \ell$. Then if we restrict the family $(f_p)_{p \in P^{(\ell-1)}}$ to values greater than π_1 , any pair of those functions can intersect not more than once by the above analysis. Let π_2 be the largest x -value such that two of these functions intersect. If we choose for $p_{\ell,1}$ a value greater than π_2 , we obtain that

$$1 - \frac{\sqrt[\ell]{\prod_p \cdot p_{\ell,1}}}{\frac{1}{\ell} (\sum_p + p_{\ell,1})} = 1 - \frac{\sqrt[\ell]{\prod_q \cdot p_{\ell,1}}}{\frac{1}{\ell} (\sum_q + p_{\ell,1})}$$

implies $p = q$ for any $p, q \in P^{(\ell-1)}$.

The recursion step for $k > 1$ is almost the same. We basically just have to replace $p_{\ell,1}$ by $p_{\ell,k}$. Furthermore, we have to choose for $p_{\ell,k}$ a value not only greater than the corresponding π_2 , but also greater than some other lower bound implicitly given by the set $Z =$

$$\left\{ 1 - \frac{\sqrt[\ell]{\prod_q \cdot p_{\ell,i}}}{\frac{1}{\ell} (\sum_q + p_{\ell,i})} \right\}_{q \in P^{(\ell-1)}, 1 \leq i < k}. \text{ As } \max Z < 1, \text{ choosing } x$$

large enough will result in $1 > f_p(x) > z$ for all $p \in P^{(\ell-1)}$ and $z \in Z$. This can be done since (i) $\lim_{x \rightarrow \infty} f_p(x) = 1$, and (ii) $f_p(0) = 1$ if and only if $x = 0$ hold for all $f_p \in (f_p)_{p \in P^{(\ell-1)}}$. So, choosing x large enough to obtain $1 > f_p(x) > z$ for all $p \in P^{(\ell-1)}$ and $z \in Z$ is possible.

We also refer to Figure 2 for an intuition: All functions f_p have the shape of the function shown in the plot. In particular, all functions of this type are differentiable with

$$\frac{d}{dx} f_p(x) = -\sqrt[\ell]{\prod_p} \cdot \left(\frac{\sqrt[\ell]{x} \left(\frac{(\sum_p + x)}{\ell \cdot x} - 1 \right)}{\left(\frac{1}{\ell+1} \right)^2 (\sum_p + x)^2} \right),$$

so the sign of $\frac{d}{dx} f_p(x)$ is determined by the term $\left(\frac{(\sum_p + x)}{(\ell+1)x} - 1 \right)$. It is easy to check that

$$\frac{d}{dx} f_p(x) \begin{cases} < 0 & \text{if } x \in]0, \sum_p / \ell[, \\ = 0 & \text{if } x = \sum_p / \ell \text{ and} \\ > 0 & \text{if } x \in]\sum_p / \ell, \infty[. \end{cases}$$




Figure 2: A sketch of the function $f_p \in (f_p)_{p \in P^{(\ell-1)}}$ with $\ell = 3$ and $p = (2, 3)$. All functions used in the proof of Lemma 11 have a similar shape, in particular we use that $\lim_{x \rightarrow \infty} f_p(x) = 1$ for all f_p .

Now let us check $P = (p_{i,j})_{i=1, \dots, n}^{j=1, \dots, 2^m}$. By construction, the function $\mathcal{F}: P^{(n)} \rightarrow [0, 1], p \mapsto 1 - \frac{\sqrt[\ell]{\prod_p}}{\frac{1}{n}(\sum_p)}$ is injective. We define $\mathcal{U}_{\mathcal{N}, \mathcal{G}} := \{\mathbf{u}(A) : A \text{ is an allocation in } \langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle\}$. Then $\mathcal{U}_{\mathcal{N}, \mathcal{G}} \subseteq P^{(n)}$ and $\mathcal{F}|_{\mathcal{U}_{\mathcal{N}, \mathcal{G}}} = \mathcal{I}|_{\mathcal{U}_{\mathcal{N}, \mathcal{G}}}$, completing the proof. \square

The uniqueness property just established now is key to proving the result announced earlier (recall once more that n is the number of agents and m is the number of goods):

THEOREM 12. *A sequence of Atkinson deals leading to an allocation that minimises inequality, as defined by the Atkinson index, can consist of up to $n^m - 1$ deals, but not more.*

PROOF. There are n^m possible allocations (each of the m items may be given to any of the n agents). By Lemma 11, there exist scenarios for which each of these allocations has a unique value of \mathcal{I} . Then, by ordering all allocations in descending order by their value of \mathcal{I} and by defining the corresponding deals between these allocations, we obtain a sequence of $n^m - 1$ deals. Each of these deals decreases inequality and therefore is an Atkinson deal. The argument for why there can never be more than $n^m - 1$ Atkinson deals in a row has been given at the beginning of Section 4.5. \square

5. CONCLUSION

We have shown that the Atkinson index, one of the most important social fairness criteria in the literature, can be

optimised in a distributed manner (Theorem 4) and thus is suitable for implementation as an objective in a multiagent system. We have been able to do so despite two inherent difficulties: the fact that the problem of finding an optimal allocation (with perfect equality) is NP-hard (Proposition 2), and the fact that the essence of what it means to reduce inequality cannot be captured locally (Proposition 3). While most other social criteria studied in the context of multiagent resource allocation also require us to solve computationally intractable optimisation problems [9], the only other such criterion that also shares the second difficulty and that nevertheless has been analysed successfully using the distributed approach is envy-freeness [10].

While Theorem 4 is encouraging, our additional results show that implementing this solution still comes with significant practical challenges. First, agents must be able to agree on arbitrarily complex exchanges of resources, without any limits on either the number of agents or the number of resources involved (Theorem 5). Second, the number of exchanges implemented before an optimal allocation is reached can get very high and in the most extreme case we might end up visiting every logically possible allocation along the way (Theorem 12). For these negative results in particular, we have made use of analytical techniques from the basic calculus toolbox, which is unusual in the field of multiagent resource allocation and which we hope might be useful to others working on related problems.

We also hope that our work will inspire other researchers in multiagent systems, first, to use the formal notion of social inequality in the design of practical multiagent systems and, second, to further advance our state of knowledge regarding the algorithmic challenge of minimising inequality in a multiagent system. Both aspects are currently underrepresented in multiagent systems research (the very few exceptions include the works of Lesca and Perny [19] and Endriss [13]), even though inequality indices are widely studied and used in practice across much of the social sciences.

Our work also suggests a number of very concrete avenues for future research. First, is a similar analysis possible for other inequality indices? For the Gini index [16], we conjecture that it would be difficult to achieve optimisation in a distributed manner without making major concessions regarding the definition of the ‘locality’ of a deal. For the Theil index [28], another popular inequality index, our own preliminary results show that distributed optimisation likely will be possible, but in a less elegant manner than for the Atkinson index. Second, how obstructive are our negative results in practice? To address this question, we might generate a scenario (using synthetic preferences or preferences extracted from a real-world problem) and simulate what happens when agents randomly choose between one of the Atkinson deals currently available to them (possibly giving more weight to structurally simpler deals). One could investigate how often such a system gets stuck in a state where all available deals exceed some given structural complexity threshold (to assess the practical relevance of Theorem 5). One could also count the average number of deals contracted in such a system (to assess the practical relevance of Theorem 12).

6. ACKNOWLEDGMENTS

We thank our anonymous reviewers for their careful reading and valuable feedback. This work was partly supported by COST Action IC1205 on Computational Social Choice.

REFERENCES

- [1] P. D. Allison. Measures of inequality. *American Sociological Review*, 43(6):865–880, 1978.
- [2] A. B. Atkinson. On the measurement of inequality. *Journal of Economic Theory*, 2:244–263, 1970.
- [3] A. B. Atkinson. *Inequality*. Harvard University Press, 2015.
- [4] S. Bouveret. *Allocation et Partage Équitables de Ressources Indivisibles: Modélisation, Complexité et Algorithmique*. PhD thesis, ISAE Toulouse, 2007.
- [5] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2016.
- [6] S. Bouveret and M. Lemaître. Computing leximin-optimal solutions in constraint networks. *Artificial Intelligence*, 173(2):343–364, 2009.
- [7] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC-2016)*, pages 305–322. ACM, 2016.
- [8] A.-L. Cauchy. Cours d’analyse. *OEuvres Complètes*, Bd, 3, 1821.
- [9] 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.
- [10] Y. Chevaleyre, U. Endriss, and N. Maudet. Distributed fair allocation of indivisible goods. *Artificial Intelligence*, 242:1–22, 2017.
- [11] H. Dalton. The measurement of the inequality of incomes. *The Economic Journal*, 30(119):348–361, 1920.
- [12] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 2005.
- [13] U. Endriss. Reduction of economic inequality in combinatorial domains. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2013)*, pages 175–182. IFAAMAS, 2013.
- [14] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research (JAIR)*, 25:315–348, 2006.
- [15] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.
- [16] C. Gini. *Variabilità e Mutabilità*. C. Cuppini, Bologna, 1912.
- [17] M. N. Huhns and A. K. Malhotra. Negotiating for goods and services. *IEEE Internet Computing*, 3(4):97, 1999.
- [18] Y. Kondor. An old-new measure of income inequality. *Econometrica*, 39(6):1041–42, 1971.
- [19] J. Lesca and P. Perny. LP solvable models for multiagent fair allocation problems. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-2010)*, pages 393–398, 2010.
- [20] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- [21] J. F. Nash Jr. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- [22] N.-T. Nguyen, T. T. Nguyen, M. Roos, and J. Rothe. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous Agents and Multiagent Systems*, 28(2):256–289, 2014.
- [23] S. Ramezani and U. Endriss. Nash social welfare in multiagent resource allocation. In *Agent-Mediated Electronic Commerce: Designing Trading Strategies and Mechanisms for Electronic Markets*, volume 59 of *Lecture Notes in Business Information Processing*, pages 117–131. Springer-Verlag, 2010.
- [24] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.
- [25] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proceedings of the 1998 AAAI Spring Symposium on Satisficing Models*, 1998.
- [26] T. W. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 201–258. MIT Press, 1999.
- [27] A. Sen. *On Economic Inequality*. Oxford University Press, 1973.
- [28] H. Theil. The information approach to demand analysis. *Econometrica*, 33(1):67–87, 1965.