



Minimising inequality in multiagent resource allocation

Structural analysis of a distributed approach

Sebastian Schneckeburger¹ · Britta Dorn¹ · Ulle Endriss²

Accepted: 16 February 2022 / Published online: 23 March 2022
© The Author(s) 2022

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 who agree 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 there are significant computational hurdles to be overcome in the worst case. In particular, finding an optimal allocation is computationally intractable and under the distributed approach a large number of structurally complex deals, possibly involving many agents and items, may be required before convergence to a socially optimal allocation. This remains true even in severely restricted resource allocation scenarios where all agents have the same utility function. From a methodological point of view, while much work in multiagent resource allocation relies on combinatorial arguments, here we instead use insights from basic calculus.

Keywords Multiagent resource allocation · Fair division · Inequality indices

Mathematics Subject Classification (2010) 91B32

An early version of this paper appears in the proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems [33].

✉ Sebastian Schneckeburger
sebastian.schneckeburger@posteo.de

Britta Dorn
britta.dorn@uni-tuebingen.de

Ulle Endriss
ulle.endriss@uva.nl

¹ University of Tübingen, Sand 13, 72076 Tübingen, Germany

² ILLC, University of Amsterdam, 1090 GE Amsterdam, The Netherlands

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 [11]. 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, such as the monograph by Moulin [26].

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 [32]. 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. Examples include *egalitarian social welfare*, measuring quality as utility of the worst-off agent, and its refinement the *leximin-ordering*, where you try to maximise the utility of each agent subject to the constraint that no agent already worse off suffers a loss of utility in the process [7, 16]; *Nash social welfare*, measuring quality as the product of the individual utilities [28, 30]; and the *absence of envy*, in the sense of no agent preferring another agent’s bundle over her own [12, 22]. However, fairness criteria based on measuring inequality, which are widely used in the social sciences [2, 20, 34], to date have received almost no attention in AI and Computer Science (exceptions include the works of Lesca and Perny [25], Endriss [15], and Gemici et al. [19]).

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* [2], and analyse how to achieve allocations of resources to agents that are optimal relative to this criterion. Our main contributions concern the challenge of ensuring convergence to 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, starting from a given initial allocation [12, 14, 16, 31]. In addition, we analyse the *price of minimising inequality*, i.e., the loss in economic efficiency incurred as a result of attempting to minimise inequality, and the *computational complexity* of computing an optimal allocation that minimises inequality. These results, both of which are independent of the specific approach chosen for performing multiagent resource allocation, highlight the fact that minimising inequality is a challenging task. Our results regarding the distributed approach to minimising inequality 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 deals of arbitrarily high structural complexity (i.e., deals involving a large number of agents and/or items) and we may require an exponential number of deals to be implemented in sequence. 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, and some of the methods we use and develop here might also be applicable to other problems in the area of multiagent resource allocation.

This paper extends our original work on the Atkinson index in multiagent resource allocation [33] in a number of ways, most notably by showing that our results illustrating the computational challenges associated with minimising inequality in a distributed manner hold up also in highly restricted resource allocation scenarios where all agents have the same utility function.

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. Towards the end of Section 2, we prove two baseline results regarding the *price of minimising inequality* and the computational complexity of minimising inequality. Our main contributions are presented in Section 3, 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. We further show that these limitations persist even for very restricted scenarios. Section 4 concludes with a brief outlook on future directions of research in this domain.

2 The model and basic results

In this section, we first introduce the basic model of multiagent resource allocation we are going to work with. In this model, which is widely used in the multiagent systems literature—see, e.g., the surveys by Bouveret et al. [6] and Chevaleyre et al. [11]—a number of indivisible goods need to be allocated to a group of agents who each have their own preferences over which bundles of goods they would like to obtain. We then review relevant definitions regarding inequality measurement from the literature on welfare economics—pioneered by authors such as Atkinson [2] and Sen [34], and reviewed in depth, for instance, by Moulin [26]. Following Endriss [15], we then adapt these notions to the setting of indivisible goods typically considered in the multiagent systems literature.

Finally, we analyse basic features of our model in view of our objective of minimising inequality by stating basic results regarding, first, the impact of this objective on economic efficiency and, second, the computational complexity of implementing this objective. These basic results underscore the fact that minimising inequality is a demanding requirement when choosing an allocation of indivisible goods, be it by means of a centralised mechanism or the kind of distributed mechanism we are going to focus on in later parts of this paper.

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 $i \neq j$ and $A(1) \cup \dots \cup A(n) = \mathcal{G}$. 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* $u(A) = (u_1(A), \dots, u_n(A))$. The collection of the utility functions of all agents is denoted by \mathcal{U} . We refer to triples $(\mathcal{N}, \mathcal{G}, \mathcal{U})$ as *scenarios*.

Some of our results apply only to scenarios with specific types of utility functions. We call a utility function u *normalised* if $u(\emptyset) = 0$. Furthermore, u is called *monotone* if $B \subseteq B'$ implies $u(B) \leq u(B')$; it is called *additive* if $u(B) = \sum_{x \in B} u(\{x\})$ for all bundles B ; it is called *modular* if $u(B) + u(B') = u(B \cup B') + u(B \cap B')$ for all bundles B and B' , and it is called *submodular* if $u(B) + u(B') \geq u(B \cup B') + u(B \cap B')$ for all bundles B and B' . Observe that every modular utility function is also submodular. Furthermore, by a well-known fact, u is modular if and only if $u(B) = u(\emptyset) + \sum_{x \in B} (u(\{x\}) - u(\emptyset))$ for all bundles B . Thus, u is additive if and only if it is both modular and normalised. We call a scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ *normalised* (or *monotone*, or *additive*, or *modular*, or *submodular*) if all utility functions in \mathcal{U} have that property. Finally, we call a scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ *symmetric* if all agents have the same utility function, i.e., if $u_i = u_j$ for all $i, j \in \mathcal{N}$. When dealing with utility functions, we sometimes work with their derivatives on an *open* interval. We use the notation $]a, b[$ or $]a, b]$ to denote the open interval of real numbers between a and b or the corresponding half open interval, respectively.

A number of criteria for assessing the quality of a given allocation of resources, be it in terms of economic efficiency or fairness, have been developed in the literature on welfare economics [26], some of which are now routinely used in the multiagent systems literature as well [11]. The *utilitarian social welfare* of allocation A is defined as $sw_{util}(A) = \sum_{i \in \mathcal{N}} u_i(A)$. Maximising utilitarian social welfare, or equivalently maximising the *mean value* $\mu(A) = \frac{1}{n} sw_{util}(A)$, amounts to maximising total (or average) utility. Thus, this is a measure of economic efficiency. The *Nash social welfare* of A is defined as $sw_{nash}(A) = \prod_{i \in \mathcal{N}} u_i(A)$. Like utilitarian social welfare, Nash social welfare increases when we increase the utility of individual agents. But it also encodes a notion of fairness by also rewarding certain equality-increasing redistributions of utility. For example, in a scenario with three agents, switching from a state with utility vector $(1, 2, 6)$ to a state with utility vector $(1, 4, 4)$ increases Nash social welfare from 12 to 16 but does not affect utilitarian social welfare, which is equal to 9 in both cases.

2.2 Inequality indices

vectors. This is an intuitively appealing idea that goes straight to the core of the concept of fairness. But it is not easy to give a clear definition of ‘inequality’ and a sizeable literature in the social sciences, stretching back more than a century, has been dedicated to this challenge. A common approach is to measure inequality using 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). High values close or equal to 1 stand for high inequality amongst the agents, while values close to 0 are reserved for allocations that, intuitively speaking, are ‘almost equitable’.¹ Well-known examples of inequality indices include the *Gini index* [20], the *Robin Hood index* [24] (also called the *maximum relative mean deviation*), and in particular the family of *Atkinson indices* [2].

¹An alternative approach to formalising this idea of an allocation being close to equitable has been studied very recently by Gourvès et al. [21] and Freeman et al. [17], amongst others: an allocation A is *equitable up to one good* (EQ1) if, for any two agents $i, j \in \mathcal{N}$, we can find an item $a \in \mathcal{G}$ such that $u_i(A(i)) \geq u_j(A(j) \setminus \{a\})$. That is, i is at least as happy as j would be if we were to remove a from j ’s bundle. Thus, while ‘near equitability’ as characterised by an inequality index is defined in terms of the utility levels of the agents involved, EQ1 instead defines ‘near equitability’ in terms of the goods inducing those utility levels. For this reason, at the technical level, there is no immediate link between the two approaches.

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 $u(A)$ of an allocation A , we first compute the so-called *equally distributed equivalent level of income* $\mu_{sw}(A)$, defined in such a way that the vector $(\mu_{sw}(A), \dots, \mu_{sw}(A))$ has the same social welfare as $u(A)$. The Atkinson index based on the function sw is then defined as $I_{sw}(A) := 1 - \frac{\mu_{sw}(A)}{\mu(A)}$ [2]. Under certain technical assumptions, which are satisfied by all standard notions of social welfare, we can ensure that $0 \leq \mu_{sw}(A) \leq \mu(A)$ [34], i.e., $I_{sw}(A) \in [0, 1]$ as required. Observe that $I_{sw}(A) = 0$ whenever A exhibits perfect equality, as then $\mu_{sw}(A) = \mu(A)$.

So why is this a reasonable approach to measuring inequality? The first thing to be said here is that, of course, the adequacy of I_{sw} depends, in part, on the adequacy of sw as a means of measuring social welfare. If sw provides a notion of social welfare that we are willing to accept as relevant, then, for any given allocation A , we can think of $(\mu_{sw}(A), \dots, \mu_{sw}(A))$ as the utility vector of an imaginary allocation A' that is as desirable (i.e., that delivers the same social welfare) as A while also being perfectly equitable. But A' achieves this level of social welfare ‘using’ less total utility than A does. For instance, to recast an example originally given by Atkinson [2], if it is the case that $\frac{\mu_{sw}(A)}{\mu(A)} = 0.7$, then this can be interpreted as indicating that, if utility were transferable and divided equally amongst all agents, then we would require only 70% of the total utility generated by A to obtain the same social welfare as we do for A .

In this paper, we are going to focus on the most important representative of this family of inequality indices, the Atkinson index based on 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 entire family and for this specific instance of an index, from here on we only use it in this latter sense. In what follows, we write \mathcal{I} instead of I_{nash} .

We stress that Nash social welfare and the Atkinson index do not measure the same thing. For instance, two allocations with utility vectors $(2, 2)$ and $(1, 4)$ yield the same Nash social welfare but exhibit very different levels of inequality, as measured by the Atkinson index: 0 vs. 0.2. Under the social welfare perspective, the loss in equality when moving from $(2, 2)$ to $(1, 4)$ is made up for by the gain in total utility, while the Atkinson index does not permit any such compensation between different social criteria.

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 following inequality for the arithmetic and the geometric mean, which holds for any nonnegative real numbers x_1, \dots, x_n :

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

Equality holds if and only if $x_1 = x_2 = \dots = x_n$. A proof of this fact can be found in Cauchy’s *Analyse Algébrique* [10, p. 457]. □

We focus on the Atkinson index because of its importance in the literature in the social sciences [1, 2, 26, 34]. 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 [9, 26, 27, 34]. Furthermore, the Atkinson index fulfils the common basic axioms for inequality indices which include the transfer principle, symmetry, and scale invariance [1, 2, 13].

The *transfer principle* states that transfers from an agent with a high utility to one with low utility shall not increase the inequality (provided that transfer is not so significant as to invert the relative ranking of these two agents). An inequality index is *symmetric* if it does not depend on the ordering of the agents, i.e., if it is invariant under permutations of the utility vector. Finally, *scale invariance* requires that the level of inequality measured should not change if we multiply the utility of every agent with the same positive constant. Thus, inequality should not depend on the ‘currency’ we use to measure individual utility.

Before we turn to our discussion of the challenges involved in finding an allocation of goods that minimises inequality in view of the Atkinson index, let us go over an extended example that, not only, illustrates the process of computing the Atkinson index for a given allocation but that also allows us to contrast this solution concept with the solution concepts based on the maximisation of either utilitarian or Nash social welfare.

Example 1 Consider the scenario $(\mathcal{N}, \mathcal{G}, \mathcal{U})$, with agents $\mathcal{N} = \{1, 2, 3\}$ and goods $\mathcal{G} = \{a, b, c\}$. The collection \mathcal{U} of utility functions is defined in terms of the values the agents assign to each of the eight possible bundles:

\mathcal{U}	\emptyset	$\{a\}$	$\{b\}$	$\{c\}$	$\{a, b\}$	$\{a, c\}$	$\{b, c\}$	$\{a, b, c\}$
$\mathbf{1}$	0	2	1	1	3	3	2	10
$\mathbf{2}$	0	0	2	3	2	3	5	10
$\mathbf{3}$	0	3	0	2	3	5	2	10

Observe that these utility function are ‘almost additive’. When going over the calculations that follow, the reader is invited to focus on the bundles with one or three goods, respectively; the three bundles with two goods each do not play a significant role. In the above scenario ...

- ... the allocation A with $A(1) = \{a, b, c\}$, $A(2) = \emptyset$, and $A(3) = \emptyset$ maximises sw_{util} .
- ... the allocation A' with $A'(1) = \{b\}$, $A'(2) = \{c\}$, and $A'(3) = \{a\}$ maximises sw_{nash} .
- ... the allocation A'' with $A''(1) = \{a\}$, $A''(2) = \{b\}$, and $A''(3) = \{c\}$ minimises \mathcal{I} .

The corresponding utility vectors and values of sw_{util} , sw_{nash} , and \mathcal{I} are shown in the following table:

	u	sw_{util}	sw_{nash}	\mathcal{I}
A	: (0, 0, 10)	10	0	$1 - \frac{\sqrt[3]{0}}{\frac{1}{3} \cdot 10} = 1$
A'	: (1, 3, 3)	7	9	$1 - \frac{\sqrt[3]{9}}{\frac{1}{3} \cdot 7} \approx 0.11$
A''	: (2, 2, 2)	6	8	$1 - \frac{\sqrt[3]{8}}{\frac{1}{3} \cdot 6} = 0$

This example shows that the solution concept of utilitarian social welfare can be highly unfair. The Nash social welfare is good mixture of efficiency and fairness, but there are contexts where inequality measurements provide an intuitively fairer solution.

2.3 The price of minimising inequality

As we have seen, despite the fact that the Atkinson index is based on the notion of Nash social welfare, these two approaches to assessing allocations do not always rank allocations in the same way. While the Atkinson index focuses on the avoidance of inequality—and thus on fairness—alone, Nash social welfare combines fairness and efficiency concerns. Unsurprisingly, fairness and efficiency demands will sometimes be in conflict and require a certain tradeoff. Caragiannis et al. [8] have introduced the *price of fairness* as a means of quantifying this tradeoff. For a given notion of fairness and a given scenario, it is defined as the ratio between (i) the utilitarian social welfare of the most efficient allocation amongst all allocations and (ii) the utilitarian social welfare of the most efficient allocation amongst all fair allocations. The price of fairness of a class of scenarios (e.g., defined in terms of certain properties of the utility functions) is defined as the maximum price of fairness across all scenarios belonging to this class. A small price of fairness is desirable, and the best value the price of fairness can take is 1, which happens when the best allocation according to the considered fairness concept is also maximally efficient.

Amongst other things, Caragiannis et al. [8] consider the fairness concept of *equitability*. An allocation is equitable if the utilities the agents enjoy from the goods received are equal. For indivisible goods (and additive utility functions that assign utility 1 to the full bundle of all goods), the price of equitability is finite only for the case of two agents, whereas it is infinite for $n \geq 3$ agents [8]. This result immediately extends to our setting and the fairness concept of minimising inequality according to the Atkinson index:

Proposition 2 *The price of minimising inequality is infinite for scenarios with $n \geq 3$ agents, even when restricted to scenarios with additive utility functions that assign utility 1 to the full bundle of all goods.*

Proof The claim is an immediate consequence of the quoted result by Caragiannis et al. [8, Theorem 15] regarding the price of equitability, given that for scenarios for which an equitable allocation exists that allocation also minimises inequality according to the Atkinson index. In their proof, Caragiannis et al. construct a family of scenarios for which the only equitable allocation is as far removed as possible from the most efficient one: For n agents and n goods, it is the allocation in which every agent receives a good that has utility ε , where ε is an arbitrarily small positive real number. The social welfare of this allocation is then $n\varepsilon$, whereas the most efficient allocation has a social welfare of $n - (n + 1)\varepsilon$ by their construction. This leads to a price of equitability of $\Omega(\frac{1}{\varepsilon})$, which implies the claim. \square

We interpret this result as a first piece of evidence that finding an allocation that minimises inequality, our stated objective for this paper, is not an easy task.

We note that Bertsimas et al. [5], independently from Caragiannis et al. [8], also introduced the notion of a price of fairness, albeit defined in a slightly different manner. They consider the relative reduction in utilitarian social welfare under a fair allocation compared to the most efficient allocation, expressed as a real number in the interval $[0, 1]$, where 0 corresponds to the most desirable price of fairness. If we use their definition, the price of minimising inequality for three or more agents is 1.

2.4 The computational complexity of minimising inequality

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.² Next, 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:	Is there an allocation A such that $\mathcal{I}(A) = 0$?

Unfortunately, it turns out that this problem is strongly NP-hard:³

Proposition 3 *The decision problem PERFECT INDEX OPTIMISATION is strongly NP-hard, even for additive scenarios with symmetric utility functions.*

Proof Before we present the proof, observe that to specify an instance of the problem with additive utility functions, we need to specify one number for every pair in $\mathcal{N} \times \mathcal{G}$. That is, an additive utility function can be specified by fixing the utility for each individual good (rather than for each bundle of goods).

Now, note that, by Lemma 1, we have $\mathcal{I}(A) = 0$ if and only if all agents enjoy the same level of utility. We use a reduction from the strongly NP-hard 3-PARTITION problem [18], which is defined as follows:

Partition	
Instance:	A finite set X , with $ X = 3q$ ($q \in \mathbb{N}$), a bound $T \in \mathbb{Z}_{>0}$ and a size s with $s(x) \in \mathbb{Z}_{>0}$ for each $x \in X$ such that $T/4 < s(x) < T/2$ and $\sum_{x \in X} s(x) = q \cdot T$.
Question:	Is there a partition $X_1 \cup \dots \cup X_q$ of X into sets of size 3 such that $\sum_{x \in X_i} s(x) = T$ for all $i \in \{1, \dots, q\}$?

Given an instance $\langle X, T, s \rangle$ of the 3-PARTITION problem, we construct an instance $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ of the PIO problem, where the set of agents is $\mathcal{N} = \{1, \dots, q\}$, the set of goods corresponds to the elements of X to be partitioned, i.e., $\mathcal{G} = X$, and the collection \mathcal{U} of utility functions is defined as follows: the utility of each good is given by the size function s , i.e., $u_i(x) = s(x)$ for all $x \in X, i \in \mathcal{N}$, and we further set $u_i(B) = \sum_{x \in B} s(x)$ for all bundles $B \subseteq \mathcal{G} = X$ and $i \in \mathcal{N}$. Hence, the constructed scenario is additive and symmetric.

Now $\langle X, T, s \rangle$ is a yes-instance if and only if $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ is: First, assume that $\langle X, T, s \rangle$ is a yes-instance with partition $X_1 \cup \dots \cup X_q$ of X , then assign to each agent $i \in \mathcal{N}$ the goods

²Prior work by Endriss [15] has also considered the computational complexity of the task of *reducing* (but not necessarily minimising) inequality, albeit not for the specific notion of inequality encoded by the Atkinson index.

³While we are not aware of a proof of this result in the literature, Proposition 3 is not surprising. Indeed, remarks by Freeman et al. [17, footnote 7] amount to essentially the same result, hinting at the same kind of reduction to obtain that result.

corresponding to the elements in the set X_i . Since $\sum_{x \in X_i} s(x) = T$ for all $i \in \{1, \dots, q\}$, each agent receives the same utility, hence the allocation has Atkinson index 0. Conversely, assume there exists an allocation A of goods with $\mathcal{I}(A) = 0$. By Lemma 1, this means that each agent enjoys the same utility $\mu(A)$ from the goods assigned to her by A . For $i \in \mathcal{N}$, let B_i be the bundle assigned to agent i by A . Then, for all $i \in \mathcal{N}$, we set $X_i = B_i$, and we have $\mu(A) = u_i(A) = \sum_{x \in B_i} u_i(x) = \sum_{x \in X_i} s(x) = T$. Hence there exists a 3-partition of X as required. \square

In our earlier work [33], we instead used a reduction from the PARTITION problem to show NP-hardness. The advantage of the reduction from 3-PARTITION given here is that it shows that PIO is NP-hard in the strong sense, meaning that there is no pseudo-polynomial algorithm to solve it and that there can be no fully polynomial-time approximation scheme (FPTAS) for this problem, unless $P = NP$ [36].

3 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. [16], relying on ideas originally introduced by Sandholm [31]. 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 reach 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 [31], egalitarian social welfare [16], Nash social welfare [28], and envy-freeness [12].

After defining the notion of a deal formally (in Section 3.1), we 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 3.2). However, we are then going to see that a very mild relaxation of this notion of locality is sufficient to obtain a convergence result (see Section 3.3). 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 3.4), where the notion ‘complexity’ refers to the number of agents and goods involved in a single deal. Second, we show that we must allow for the possibility of exponentially long sequences of deals before convergence is realised (see Section 3.5). We conclude by demonstrating that these difficulties cannot be significantly ameliorated even for highly restricted scenarios that are symmetric and in which the utility functions are monotone and submodular (see Section 3.6).

3.1 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 2 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 Fig. 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 the allocations A° and A^* of Fig. 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.

3.2 No convergence by local deals

We are looking for a criterion to select admissible deals such that (i) the agents involved in any given deal are able to determine locally whether that deal is admissible and (ii) any sequence of admissible deals eventually leads to an optimal allocation. Regarding the latter requirement, we are specifically interested in sequences of deals for which inequality decreases monotonically along the way, so as to obtain a mechanism with an ‘anytime’ character, meaning that we can guarantee that the situation will continue to improve as long as the agents keep on agreeing on deals. But how should we define ‘locality’ in this context? Endriss et al. [16] call a criterion for determining the admissibility of a deal $\delta = (A, A')$ *local* if and only if the question of whether δ is admissible 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 suitable deal selection criterion that is local in this sense. As we are going to show next, if we restrict ourselves to deals that satisfy some criterion that is local, then the agents involved in a deal will not always be able to determine




Fig. 1 Illustration of Example 2: Starting from allocation A° which assigns goods a and b to agent 1, good c to agent 2, and good d to agent 3, respectively, the agents 1 and 2 involved in the deal δ exchange goods b and c , which results in allocation A^* . Agent 3 is not involved in this deal

whether a given deal would decrease inequality. Hence, no such local deal criterion can possibly be used to adequately guide our search for an optimal allocation.

Proposition 4 *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 $\mathcal{N} = \{1, 2, 3\}$ 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:

$$\left. \begin{array}{l} \mathcal{U}_1 \ a \ b \ c \ d \\ \mathbf{1} : 2 \ 1 \ 3 \ 4 \\ \mathbf{2} : 2 \ 5 \ 2 \ 1 \\ \mathbf{3} : 1 \ 2 \ 1 \ 6 \end{array} \right\} \text{and} \left. \begin{array}{l} \mathcal{U}_2 \ a \ b \ c \ d \\ \mathbf{1} : 2 \ 1 \ 3 \ 4 \\ \mathbf{2} : 2 \ 5 \ 2 \ 1 \\ \mathbf{3} : 3 \ 2 \ 3 \ 2 \end{array} \right\}.$$

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 Fig. 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. □

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, as demonstrated by Endriss et al. [16] and Ramezani and Endriss [28]. 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 [12]. The solution proposed by Chevaleyre et al. [12], which concerns a model like ours but permitting monetary side payments, is to allow agents not involved in a deal to receive (but never make) payments. As all of the agents not involved are required to get *the same* amount, this means that the agents involved (who make and compute these payments) must have the information of how many agents there are overall, i.e., they must know n . In the sequel, we are going to take a similar route and also break the locality requirement in a very subtle way by providing some additional information of a global nature to the agents involved in a local deal.

3.3 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.⁴ On the other hand, the underlying reason for the impossibility stated in Proposition 4 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. 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:

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

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:

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

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 5 *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 the 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. □

Similar convergence results have been proved for a number of other criteria for social optimality [12, 16, 28, 31]. In some cases, notably for utilitarian social welfare and envy-freeness [12, 31], 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 [16, 28], just as for our result here,

⁴This 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.

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 5 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.

3.4 Necessity of complex deals

Theorem 5 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.

Next, we show that, unfortunately, for every deal that is not ‘independently decomposable’ (to be defined shortly) there exists a scenario such that this deal is necessary for reaching an allocation with minimal inequality. In this context, following Endriss et al. [16], we call a deal $\delta = (A, A'')$ *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 $\mathcal{N}^{\delta_1} \cap \mathcal{N}^{\delta_2} = \emptyset$. To prove necessity of all independently decomposable deals, we make use of the fact that the Atkinson index can assume any arbitrary value in the interval $[0, 1]$, which can be shown by proving surjectivity of a correspondingly defined function. This is done in the following technical lemma.

Lemma 6 *For every $n \in \mathbb{N}_{>1}$ and $1 \leq d \leq n$, the function*

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

$$x \mapsto 1 - \frac{\sqrt[n]{(1-x)^d}}{1 - \frac{x \cdot d}{n}}$$

is strictly monotonically increasing and thus bijective.

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

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

Theorem 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 in such a way that $\mathcal{I}(A') = 0$, and that $\mathcal{I}(A)$ is strictly smaller than the inequality of any allocation different from A and A' . This will imply that, starting from allocation A , the deal $\delta = (A, A')$ is the only Atkinson deal reducing \mathcal{I} , and hence necessary.

Since A and A' are different, there is at least one agent j with $A(j) \neq A'(j)$. We fix this j and let $0 \leq x \leq 1$. We now define the utility functions for any given bundle $B \in 2^G$ as

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

This implies $\mathcal{I}(A') = 0$ and

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

We now compare $\mathcal{I}(A)$ to the inequality of all other possible allocations and show that the value of the variable x appearing in $\mathcal{I}(A)$ can be set such that $\mathcal{I}(A)$ is strictly greater than 0, but strictly smaller than the inequality of any other allocation. Apart from A and A' , there are two different types of allocations with respect to the designated agent j and the allocation A . Allocations of the first type coincide with allocation A for agent j , allocations of the second type do not.

Consider an allocation of the first type, which we will denote by \tilde{A} . As x is not fixed yet, we can interpret each $\mathcal{I}(\tilde{A})$ as a function $\mathcal{I}(\tilde{A}) : [0, 1] \rightarrow [0, 1]$ with

$$\mathcal{I}(\tilde{A})(x) = 1 - \frac{\sqrt[n]{(1-x) \cdot \prod_{i \in \mathcal{N} \setminus \{j\}} t_i}}{\frac{1}{n} \left(1 - x + \sum_{i \in \mathcal{N} \setminus \{j\}} t_i \right)},$$

where $t_i \geq 1$ holds for all $i \neq j$, and $t_i > 1$ holds for at least one $i \neq j$. We have $0 < \mathcal{I}(\tilde{A})(0)$, which can be shown with the inequality of the arithmetic and geometric mean as in the proof of Lemma 1. Furthermore, it is easy to check that $\frac{d}{dx} \mathcal{I}(\tilde{A})(x) > 0$ for all $x \in [0, 1]$, hence $\mathcal{I}(\tilde{A})$ is a strictly monotonically increasing function on $[0, 1]$ with $\mathcal{I}(\tilde{A})(0) > 0$, which means that the family of functions $\{\mathcal{I}(\tilde{A})\}_{\tilde{A}}$ is bounded from below, i.e., there is a real number $\varepsilon_1 > 0$ such that $0 < \varepsilon_1 < \mathcal{I}(\tilde{A})(x)$ for all $\mathcal{I}(\tilde{A})$ and all $x \in [0, 1]$.

Next, we show that for any allocation A^* of the second type 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^* coincides 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_2 with $0 < \varepsilon_2 < \min_{A^* \neq A, A'} \mathcal{I}(A^*)$ and then set x such that $\mathcal{I}(A) = \varepsilon := \min\{\varepsilon_1, \varepsilon_2\}$, which is possible due to Lemma 6. Hence, we have $0 = \mathcal{I}(A') < \mathcal{I}(A) \leq \varepsilon_2 < \min_{A^* \neq A, A'} \mathcal{I}(A^*)$ as well as $\mathcal{I}(A) \leq \varepsilon_1 < \mathcal{I}(\tilde{A})$ for any allocation \tilde{A} of the first type. Thus, in this scenario, starting from allocation A , $\delta = (A, A')$ is the only deal reducing \mathcal{I} , and thus the only Atkinson deal. \square

Theorem 7 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. Our construction used in the proof of Theorem 7 is similar to the construction used to derive necessity results for utilitarian and egalitarian social welfare [16] as well as Nash social welfare [28]. In those other settings, not only are all non-independently decomposable deals necessary, but these are the *only* such deals. Surprisingly, in the present setting the situation is worse and even deals that *are* independently decomposable are necessary, as the following example demonstrates:

Example 3 Consider the additive scenario $(\mathcal{N}, \mathcal{G}, \mathcal{U})$, with $\mathcal{N} = \{1, 2, 3, 4\}$ 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:

$$\begin{array}{c|cccc} \mathcal{U} & a & b & c & d \\ \hline \mathbf{1} & 4 & 10 & 4 & 4 \\ \mathbf{2} & 10 & 3 & 3 & 3 \\ \mathbf{3} & 2 & 2 & 2 & 10 \\ \mathbf{4} & 1 & 1 & 10 & 1 \end{array}$$

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^{i1}, A') and (A, A^{i2}, A') with $A^{i1} = (\{a\}, \{b\}, \{d\}, \{c\})$ and $A^{i2} = (\{b\}, \{a\}, \{c\}, \{d\})$.

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^{i1} , A^{i2} , and A' have an inequality different from the one of A . The values of \mathcal{I} for these four allocations are as follows:

$$\begin{aligned} \mathcal{I}(A) &: 1 - \frac{\sqrt[4]{4 \cdot 3 \cdot 2 \cdot 1}}{\frac{1}{4} \cdot (4 + 3 + 2 + 1)} \approx 0.115 \\ \mathcal{I}(A^{i1}) &: 1 - \frac{\sqrt[4]{10 \cdot 10 \cdot 2 \cdot 1}}{\frac{1}{4} \cdot (10 + 10 + 2 + 1)} \approx 0.346 \\ \mathcal{I}(A^{i2}) &: 1 - \frac{\sqrt[4]{4 \cdot 3 \cdot 10 \cdot 10}}{\frac{1}{4} \cdot (4 + 3 + 10 + 10)} \approx 0.128 \\ \mathcal{I}(A') &: 1 - \frac{\sqrt[4]{10 \cdot 10 \cdot 10 \cdot 10}}{\frac{1}{4} \cdot (10 + 10 + 10 + 10)} = 0 \end{aligned}$$

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

In fact, we are able to strengthen Theorem 7 and show that *every* deal is necessary. We again begin by establishing a technical lemma.

Lemma 8 For $a, b \in \mathbb{R}_{>0}$ and two numbers $d, n \in \mathbb{N}$ with $1 \leq d < n$ we have that

$$\lim_{x \rightarrow \infty} \frac{\sqrt[n]{a \cdot x^d}}{\frac{1}{n}(b + d \cdot x)} = 0.$$

Proof As $\sqrt[n]{a \cdot x^d} \xrightarrow{x \rightarrow \infty} +\infty$ and $\frac{1}{n}(b + d \cdot x) \xrightarrow{x \rightarrow \infty} +\infty$ we can apply L'Hôpital's rule:

$$\begin{aligned} \lim_{x \rightarrow \infty} \frac{\sqrt[n]{a \cdot x^d}}{\frac{1}{n}(b + d \cdot x)} &= \lim_{x \rightarrow \infty} \frac{\frac{d}{dx} \left(\sqrt[n]{a \cdot x^d} \right)}{\frac{d}{dx} \left(\frac{1}{n}(b + d \cdot x) \right)} \\ &= \lim_{x \rightarrow \infty} \frac{\frac{d}{n} \cdot \sqrt[n]{a} \cdot x^{\frac{d}{n}-1}}{\frac{d}{n}} = \lim_{x \rightarrow \infty} \sqrt[n]{a} \cdot x^{(\frac{d}{n}-1)} \stackrel{d < n}{=} 0. \end{aligned}$$

This proves the claim. □

Now we can prove a stronger version of Theorem 7, by building both on that theorem and the lemma above:

Theorem 9 *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.*

Proof If δ is not independently decomposable, the statement is covered by Theorem 7. So we suppose δ is independently decomposable. Then it is always possible to find a sequence

$$(A = A^1, A^2, \dots, A^d, A^{d+1} = A')$$

of allocations such that each pair $(A^\ell, A^{\ell+1})$ consisting of two consecutive allocations of the sequence is a deal that is not independently decomposable, and furthermore for any two deals δ_a and δ_b , each consisting of two consecutive allocations of the sequence, we have $\mathcal{N}^{\delta_a} \cap \mathcal{N}^{\delta_b} = \emptyset$. Such a sequence can be constructed iteratively in the following manner. We start with the sequence (A, A') . If a pair $(A^\ell, A^{\ell+1})$ of consecutive allocations in the current sequence is independently decomposable, there is an allocation $A^{\ell'}$ $\notin \{A^\ell, A^{\ell+1}\}$ such that $\mathcal{N}^{(A^\ell, A^{\ell'})} \cap \mathcal{N}^{(A^{\ell'}, A^{\ell+1})} = \emptyset$ and $\mathcal{N}^{(A^\ell, A^{\ell'})} \cup \mathcal{N}^{(A^{\ell'}, A^{\ell+1})} = \mathcal{N}^{(A^\ell, A^{\ell+1})}$, then we insert this allocation $A^{\ell'}$ into the sequence between A^ℓ and $A^{\ell+1}$ and relabel the allocations. We repeat this until the sequence fulfills the conditions.

For any pair $(A^\ell, A^{\ell+1})$, we choose an agent $j \in \mathcal{N}^{(A^\ell, A^{\ell+1})}$ and denote the set of all these agents by \mathcal{D} (by construction no agent can be chosen twice and $|\mathcal{D}| = d$), and let x, y be real numbers with $0 \leq x \leq 1 < y$ (the exact values of x and y will be defined later). We now define the utility functions in a bundlewise way as

$$u_i(B) = \begin{cases} y & \text{if } A'(i) = B, \\ 1 & \text{if } (i \notin \mathcal{D}) \text{ and } A(i) = B, \\ 1 - x & \text{if } (i \in \mathcal{D}) \text{ and } A(i) = B, \\ i + 1 & \text{otherwise.} \end{cases}$$

Analogously to the construction in the proof of Theorem 7, we have $\mathcal{I}(A') = 0$ and

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

In this proof, we will distinguish three types of allocations (apart from A and A') with respect to the designated agents $j \in \mathcal{D}$ and the allocations A and A' : allocations of the first type coincide with A for some of the agents in \mathcal{D} and furthermore do not coincide with A' for any agent $i \in \mathcal{N}$; allocations of the second type do not coincide with A for any of the

agents in \mathcal{D} and furthermore do not coincide with A' for any agent $i \in \mathcal{N} \setminus \mathcal{D}$; the third type of allocations contains all allocations which coincide with A' for at least one agent.

Concerning the variables x and y , this means the following: variable x is relevant for type 1, variable y is not; neither x or y are relevant for type 2; finally, even though the variable x might occur, only variable y will be relevant for type 3.

Consider an allocation of the first type, which we will denote by \tilde{A} . We have $\tilde{A}(j) \neq A'(j)$ for all $j \in \mathcal{D}$ and $\tilde{A} \neq A$, but $\tilde{A}(j) = A(j)$ for at least one $j \in \mathcal{D}$. Analogously to the proof of Theorem 7, it can be shown that there is a real number $\varepsilon_1 > 0$ such that $0 < \varepsilon_1 < \mathcal{I}(\tilde{A})$ for all possible values of x . (Please note that the values of the utilities in this case are restricted to 1, $1 - x$, and $i + 1$ for some $i \in \mathcal{N}$, and that the latter type of value occurs at least once.)

Next, we show that for any allocation A^* of the second type, we have $\mathcal{I}(A^*) > 0$. Again, there is at least one pair of agents k, ℓ with $u_k(A^*(k)) \neq u_l(A^*(\ell))$, which implies that $\mathcal{I}(A^*) \neq 0$. As there are only finitely many allocations, we have

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

We now choose some ε_2 with $0 < \varepsilon_2 < \min_{A^* \neq A, A'} \mathcal{I}(A^*)$ and then set x such that $\mathcal{I}(A) = \varepsilon := \min\{\varepsilon_1, \varepsilon_2\}$, which is possible due to Lemma 6. Hence, we have $0 = \mathcal{I}(A') < \mathcal{I}(A) \leq \varepsilon_2 < \min_{A^* \neq A, A'} \mathcal{I}(A^*)$ as well as $\mathcal{I}(A) \leq \varepsilon_1 < \mathcal{I}(\tilde{A})$ for any allocation \tilde{A} of the first type.

Now we focus on the remaining type 3 of allocations which we denote by A° . These are the allocations in which some, but not all agents receive the same bundle as in allocation A' . For any of these allocations A° , we have

$$\mathcal{I}(A^\circ) = 1 - \frac{\sqrt[n]{y^{d'} \cdot \prod_{i=1}^{n-d'} t_i}}{\frac{1}{n}(y \cdot d' + \sum_{i=1}^{n-d'} t_i)}$$

for some $1 < d' < d$ and $0 < t_i$. We remark that $t_i = 1 - x$ is possible for some i , but at this stage of the construction x is a constant. Since this term converges to 1 as y goes to infinity by Lemma 8, we can set y such that $\mathcal{I}(A^\circ) > \varepsilon = \mathcal{I}(A)$. Therefore, in this scenario, if the allocation A is given, the deal (A, A') is the only deal reducing \mathcal{I} . □

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. This mirrors a known result by Ramezani and Endriss [28] for the case of Nash social welfare.

Theorem 10 *For any $n \in \mathbb{N}$ there are additive scenarios $(\mathcal{N}, \mathcal{G}, \mathcal{U})$ with $|\mathcal{N}| = n = |\mathcal{G}|$ and an allocation A such that a deal involving all agents and goods is necessary for reaching an allocation that minimises inequality, as defined by the Atkinson index, by means of Atkinson deals only.*

Proof Consider the additive scenario $(\mathcal{N}, \mathcal{G}, \mathcal{U})$, with $|\mathcal{N}| = n = |\mathcal{G}|$ for some $n \in \mathbb{N}$ with $n > 1$. We denote $\mathcal{N} = \{1, \dots, n\}$ and $\mathcal{G} = \{g_1, \dots, g_n\}$. The collection \mathcal{U} of additive utility functions is defined in terms of the values the agents assign to each of the items via

$$u_i(g_k) = \begin{cases} \frac{1}{i+1} & , \text{if } k \neq (i+1) \pmod n, \\ y & , \text{otherwise,} \end{cases}$$

for some $y \in \mathbb{R}$. Now consider the two scenarios A and A' with $A(i) = \{g_i\}$ and $A'(i) = \{g_{(i+1) \bmod n}\}$. Obviously $0 < \mathcal{I}(A) < 1$ and $\mathcal{I}(A') = 0$. The deal (A, A') is illustrated in Fig. 2. For all allocations A^* with $A^*(i) = \emptyset$ for some agent $i \in \mathcal{N}$, we have $\mathcal{I}(A^*) = 1$. For all allocations A^* in which no agent receives the same item as in A' , but all agents receive exactly one item, we have $\mathcal{I}(A^*) = \mathcal{I}(A)$. For all other allocations A° , we have

$$\mathcal{I}(A^\circ) = 1 - \frac{\sqrt[n]{y^{d'} \cdot \prod_{i=1}^{n-d'} t_i}}{\frac{1}{n}(y \cdot d' + \sum_{i=1}^{n-d'} t_i)}$$

for some $1 < d' < n$ and $0 < t_i < 1$. Analogously to the proof of Theorem 9, this term converges to 1 as y goes to infinity, so again we can set y such that $\mathcal{I}(A^\circ) > \mathcal{I}(A)$ for all allocations A° of this particular type. Therefore, in this scenario, if the allocation A is given, the deal (A, A') (which involves all agents and all goods) is the only deal reducing \mathcal{I} . \square

The proof of Theorem 10 can easily be generalised to scenarios where the number of goods exceeds the number of agents.

3.5 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 11, before we proceed to the general case of n agents in Lemma 14.

Lemma 11 *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 [28] 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




Fig. 2 An illustration of the deal used in the proof of Theorem 10. Each agent transfers the one good she owns to another agent in a cyclic manner. Hence, all agents and all goods are involved in this deal

the $(j^i_A)^{th}$ bundle in his ordering. It is easy to see that any two allocations A and A' have different Nash social welfare, since

$$sw_{nash}(A) = 2^{j^1_A} \cdot 3^{j^2_A} = 2^{j^1_{A'}} \cdot 3^{j^2_{A'}} = sw_{nash}(A')$$

would imply directly $j^1_A = j^1_{A'}$ and $j^2_A = j^2_{A'}$, due to the unique prime factorisation of every integer.

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

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

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 $j^1_A = j^1_{A'}$ and $j^2_A = j^2_{A'}$ which implies $A = A'$. □

The proof of Lemma 11 cannot easily be generalised to more than two agents, as the argumentation with the modulo calculation does not hold any longer if we use more than two prime numbers. For instance, for three agents, $(2^i + 3^j + 5^k) \pmod 5$ can be congruent to 0 for $i = j = k$, e.g., $(2^1 + 3^1 + 5^1) \pmod 5 = 0$. To nevertheless obtain a generalisation, we require the following two technical lemmas, the first of which can be proven by proceeding analogously to the reasoning in the proof of Lemma 11.

Lemma 12 *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 if and only if $j_1 = j_2$ and $k_1 = k_2$.

Lemma 13 *Real functions $g_{a,b}^{(k)}: \mathbb{R}_{>0} \rightarrow \mathbb{R}$ (here $a, b > 0$ and $k \in \mathbb{N}$) given by*

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

can be interpolated exactly by using just two points $(x_1, c_1), (x_2, c_2)$ of the graph of the function.

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 (i.e., the equation has a unique solution $(a, b) \in \mathbb{R}_{>0}^2$). \square

This means that, if two functions of the above type agree on their values for two 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.

Now we are ready to generalise Lemma 11 to arbitrary numbers of agents.

Lemma 14 *For any numbers $n, m \in \mathbb{N}$, 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 have a different value of \mathcal{I} .*

Proof We consider the scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ and construct utility functions that fulfill 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 $p_{ij} = u_i(B_j)$. 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 12, that $p = q$.⁵

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

⁵This means, sloppily speaking, that similarly as in Lemma 11, 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—furthermore the considered bundles might intersect. We will use the name ‘potential partial allocations’. Note that allocations are a special case of potential partial allocations, therefore properties which hold with respect to potential partial allocations in particular hold with respect to allocations. 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.

$P^{(\ell-1)}$ (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 ‘potential partial allocations’ would exhibit a different value of inequality (again, we cannot really speak of \mathcal{I} yet). More formally, this means in the case of $k = 1$ (when the first entry of each row is computed) that

$$1 - \frac{\ell^{-1}\sqrt{\prod_p}}{\frac{1}{\ell-1}(\sum_p)} = 1 - \frac{\ell^{-1}\sqrt{\prod_q}}{\frac{1}{\ell-1}(\sum_q)}$$

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

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

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

We now give a precise definition of our recursion. Let a denote the number of entries of P already fixed.

Recursion start $a_0 = 2 \cdot 2^m$: We already constructed the first two rows of the matrix as a feasible set. So our recursion starts with $a_0 = 2 \cdot 2^m$. As the first two rows contain $2 \cdot 2^m = (3 - 1) \cdot 2^m + 1 - 1$ elements, the correct values for ℓ and k are $\ell_0 = 3$ and $k_0 = 1$.

For the recursion step $a \mapsto a + 1$, we have two different cases, depending whether or not $k = 1$.⁶ In the case of $k = 1$, we set the first entry of a new row; in the case of $k > 1$, the first $k - 1$ entries of the specific row already have been set properly.

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 11, just taking powers of primes is not helpful. We therefore define real functions that feature the property used in Lemma 13.

Recursion step $a \mapsto a + 1$ **if** $k = 1$: 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

$$f_p : [0, \infty[\rightarrow [0, 1]$$

$$x \mapsto 1 - \frac{\sqrt{\prod_p \cdot x}}{\frac{1}{\ell}(\sum_p + x)}.$$

Obviously, $f_p(0) = 1$ for any $p \in P^{(\ell-1)}$, but restricted to the interval $]0, \infty[$, any pair of distinct functions of this family cannot intersect more than once. To see this, we observe the connection to Lemma 13. For $p, q \in P^{(\ell-1)}$ and $x \in]0, \infty[$, the equation

$$1 - \frac{\sqrt{\prod_p \cdot x}}{\frac{1}{\ell}(\sum_p + x)} = 1 - \frac{\sqrt{\prod_q \cdot x}}{\frac{1}{\ell}(\sum_q + x)}$$

⁶Please note that this is not a two-variable recursion; we do not independently iterate over ℓ and k .

Table 1 The partially filled matrix P in the recursion step in the proof of Lemma 14

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

is equivalent to

$$\frac{\prod_p \cdot x}{(\sum_p + x)^\ell} = \frac{\prod_q \cdot x}{(\sum_q + x)^\ell}.$$

We are thus in the situation described in Lemma 13. Therefore, if $p, q \in P^{(\ell-1)}$ are given, and $f_p(x_1) = f_q(x_1)$ and $f_p(x_2) = f_q(x_2)$ for two distinct values $x_1, x_2 \in]0, \infty[$, we have $p = q$.

Now, considering all functions of the family $(f_p)_{p \in P^{(\ell-1)}}$, let π be the largest x -value for which two of these functions intersect. By choosing a value greater than π for $p_{\ell,1}$, 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 $p, q \in P^{(\ell-1)}$.

Recursion step $a \mapsto a + 1$ if $k > 1$: 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 , but also greater than some other lower bound implicitly given by the set

$$W = \left\{ 1 - \frac{\sqrt[\ell]{\prod_q \cdot p_{\ell,i}}}{\frac{1}{\ell}(\sum_q + p_{\ell,i})} \right\}_{q \in P^{(\ell-1)}, 1 \leq i < k}.$$

As $\max W < 1$, choosing x large enough will result in $1 > f_p(x) > w$ for all $p \in P^{(\ell-1)}$ and $w \in W$. This can be done since

- (i) $\lim_{x \rightarrow \infty} f_p(x) = 1$, and
- (ii) $f_p(0) = 1$ if and only if $x = 0$ holds for all $f_p \in (f_p)_{p \in P^{(\ell-1)}}$.

So, choosing x large enough to obtain $1 > f_p(x) > w$ for all $p \in P^{(\ell-1)}$ and $w \in W$ is possible. We refer to Fig. 3 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

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




Fig. 3 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 14 have a similar shape, in particular we use that $\lim_{x \rightarrow \infty} f_p(x) = 1$ for all f_p

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

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

Now let us check $P = (p_{i,j})_{i=1, \dots, n}^{j=1, \dots, 2^m}$. By construction, the function

$$\mathcal{F}: P^{(n)} \rightarrow [0, 1], p \mapsto 1 - \frac{\sqrt[n]{\prod_p}}{\frac{1}{n}(\sum_p)}$$

is injective. We define $\mathcal{U}_{\mathcal{N}, \mathcal{G}} := \{u(A) : A \text{ is an allocation in } \langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle\}$. 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 the 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 15 *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 14, 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 3.5. \square

3.6 Symmetric scenarios with monotone and submodular utilities

While we have seen that, in principle, an allocation minimising inequality can be achieved by a sequence of uncoordinated Atkinson deals, we have also seen that, in practice, it will usually be difficult to take advantage of this theoretical possibility, given that structurally highly complex deals as well as very large numbers of deals might be required to attain an optimal outcome. To some extent, these negative results can be traced back to the high generality of our framework: agents may have arbitrary utility functions. Next, we explore whether better results can be obtained when we impose certain restrictions on utility functions. Specifically, we investigate restrictions that are directly inspired by the kinds of scenarios for which inequality indices were developed originally, namely the assessment of different policies for income distribution [20, 34]. Thus, for the restricted scenario studied here, we assume that all agents have the same utility function u (symmetry), which is normalised (i.e., $u(\emptyset) = 0$), that agents always prefer to obtain additional goods (monotonicity), and agents who already own a lot of goods derive less marginal utility from a new item than agents who own only very few goods (submodularity). Indeed, all of these assumptions would be reasonable in the context of allocating money rather than arbitrary indivisible goods.

It is useful to speak of *equivalent* allocations in this context. We say two allocations A and A' are equivalent if there is a permutation $\pi: \mathcal{N} \rightarrow \mathcal{N}$ such that for any $i \in \mathcal{N}$ we have $A(\pi(i)) = A'(i)$. Due to the symmetry of the restricted framework we consider, no social welfare criterion or inequality index can distinguish between equivalent allocations.

We first remark that PIO remains an NP-hard problem in the considered restricted scenario: in Proposition 3, we have shown hardness of PIO even for additive utility functions, and these are a special case of submodular utility functions. As before, we want to reach an optimal allocation starting from an arbitrary one by means of Atkinson deals. How long may this take, i.e., how many Atkinson deals might be necessary to reach an allocation with minimal inequality?

Next, we will see in Theorem 16 that also in this restricted framework, deals of high structural complexity might be necessary. Furthermore, we show that the worst case number of necessary deals from Theorem 15 is lower in this framework—this is due to the symmetry of the utility functions. However, as Theorem 18 shows for the special case of two agents, this number still is exponential in the number of goods. If we abstain from the requirement of normalisation (i.e., if we allow $u(\emptyset) \neq 0$), then we can expand the latter result to Theorem 17 which covers the case of an arbitrary number of agents and modular utilities (while we still assume the utility functions to be symmetric and monotone).

We start our investigation of the restricted framework with a result on the necessity of deals. As it turns out, again highly complex deals involving all agents may be necessary.

Theorem 16 *For any $n \in \mathbb{N}$ there are symmetric scenarios $(\mathcal{N}, \mathcal{G}, \mathcal{U})$ with $|\mathcal{N}| = n$ and $|\mathcal{G}| = 2 \cdot n$ and utility functions that are normalised, monotone, and submodular as well as an allocation A such that a deal including all n agents and n of the $2n$ goods is necessary for reaching an allocation that minimises inequality, as defined by the Atkinson index, by means of Atkinson deals only.*

sketch The proof follows a similar analysis as the proof of Theorem 10. We omit the details and only describe the main necessary technical adaptations. Consider the scenario $(\mathcal{N}, \mathcal{G}, \mathcal{U})$, with $|\mathcal{N}| = n$ and $|\mathcal{G}| = 2 \cdot n$ for some $n \in \mathbb{N}$ with $n > 1$. Let $\mathcal{N} = \{1, \dots, n\}$ and $\mathcal{G} =$

$\{g_1, \dots, g_{2n}\}$. As we want \mathcal{U} to be symmetric, we only have to define one utility function u . We do this in a bundlewise way. For each $B \in \mathcal{G}$, let

$$u(B) = \begin{cases} 0 & \text{if } B = \emptyset \\ 1 & \text{if } |B| = 1 \\ 2 - \varepsilon & \text{if } |B| = 2 \text{ and } \exists i \in \{1, \dots, n\} : B = \{g_i, g_{i+n}\} \\ 2 - \varepsilon_i & \text{if } |B| = 2 \text{ and } B = \{g_i, g_{n+1+(i \bmod n)}\} \\ 2 - \varepsilon_{i,j} & \text{if } |B| = 2 \text{ and } B = \{g_i, g_{i+k}\} \\ & \text{for some } k \neq n \\ 3 - \delta & \text{if } |B| > 3 \end{cases} .$$

With similar techniques as used in the proofs of Theorems 7, 9 and 10, we can derive values for $\varepsilon, \varepsilon_i, \varepsilon_{i,j}$ and δ such that the above function u is monotone and submodular: One first has to fix values for all $\varepsilon_{i,j}$, then for all ε_i , then for ε , and at last for δ , with $0 < \varepsilon \ll \varepsilon_i \ll \varepsilon_{i,j} \ll \delta \ll 1$. Given the allocation A where $A(i) = \{g_i, g_{n+1+(i \bmod n)}\}$ we define $A'(i) = \{g_i, g_{n+i}\}$. The only deals reducing inequality as measured by the Atkinson index are of the form $\delta = (A, \tilde{A}')$ where \tilde{A}' is a allocation equivalent to A' . In any of those deals all agents and at least half of the goods are included.

For the construction, monotonicity and continuity of the induced functions are needed as well. We omit the technical details. Please see Fig. 4 for an illustration of the highly complex deal δ . □

As an aside, we note that Theorem 16 is also true for symmetric scenarios with utility functions that are normalised, monotone, and supermodular (instead of submodular), i.e., when $u(B) + u(B') \leq u(B \cup B') + u(B \cap B')$ for all bundles B, B' . In the construction used in the proof of Theorem 16, replacing the minus signs in the definition of the function values of $u(B)$ by plus signs leads to the desired result.

For the remainder of this section, we are going to restrict the class of scenarios we consider even further and assume that the single utility function modelling the preferences of all agents not only is submodular but even modular. Interestingly, for symmetric scenarios with modular utility functions the Nash social welfare criterion and the Atkinson index always agree on the manner in which they rank alternative allocations.

Lemma 17 *For symmetric scenarios with modular utility functions, the Nash social welfare sw_{nash} and the Atkinson index \mathcal{I} are equivalent in the sense that for any two allocations A_1, A_2 we have*

$$sw_{nash}(A_1) \leq sw_{nash}(A_2) \iff \mathcal{I}(A_1) \geq \mathcal{I}(A_2).$$




Fig. 4 An illustration of the highly complex deal δ defined in the proof of Theorem 16. Each agent initially holds two goods, one of which is transferred to another agent in a cyclic manner in the course of the deal

Proof Recall how \mathcal{I} is defined in terms of sw_{nash} and μ . We are done if we can show that μ is a constant function, returning the same mean value $\mu(A)$ for every allocation A . But this is clearly the case for symmetric scenarios with a single modular utility function u :

$$\begin{aligned} \mu(A) &= \frac{1}{n} \sum_{i \in \mathcal{N}} u(A(i)) \\ &= \frac{1}{n} \sum_{i \in \mathcal{N}} \left(u(\emptyset) + \sum_{x \in A(i)} (u(\{x\}) - u(\emptyset)) \right) \\ &= u(\emptyset) + \frac{1}{n} \sum_{x \in \mathcal{G}} (u(\{x\}) - u(\emptyset)) \end{aligned}$$

Observe that this sum, indeed, does not depend on A . □

As we have seen, in the general setting, there are n^m possible allocations ($n = |\mathcal{N}|$, $m = |\mathcal{G}|$), each of which may be part of a sequence of deals in the worst case. In the restricted scenario, due to symmetry, we cannot distinguish all n^m allocations, but just all possible *partitions* of the set \mathcal{G} into n bundles. Moreover, since for the case of a normalised utility function an agent derives zero utility from the empty set (which leads to $\mathcal{I}(A) = 1$ if an agent receives no item), we cannot distinguish between allocations where one or more agents receive the empty set.

Our next result shows that, nevertheless, all of these seemingly strong restrictions still can give rise to exponentially long sequences of deals (see Theorem 15). We prove this for the special case of two agents.

Theorem 18 *For symmetric scenarios with additive (i.e., normalised and modular) utility functions and two agents, a sequence of Atkinson deals leading to an allocation that minimises inequality, as defined by the Atkinson index, can consist of up to $2^{m-1} - 1$ deals, but not more.*

Proof Let $\mathcal{N} = \{1, 2\}$ and $\mathcal{G} = \{g_1, \dots, g_m\}$. There are 2^m subsets of an m -element set, which means that there are 2^m allocations. Due to the symmetry requirement, only half of them are distinguishable. Hence, there are 2^{m-1} equivalent classes of allocations.

As in the proof of Lemma 14, we show that there exist scenarios for which each of these allocations (modulo equivalence of the allocations) has a unique value of \mathcal{I} , so that this upper bound of deals is attained. We construct the utility function as follows. Any bundle $B \subseteq \mathcal{G}$ can be identified with a binary vector $\mathbf{x} = (x_1, \dots, x_m) \in \{0, 1\}^m$ of length m by setting $x_i = 1$ if and only if $g_i \in B$. Motivated by this, we identify any bundle B with a number $x_B \in \{0, \dots, 2^m - 1\}$ via $x_B = \sum_{g_i \in B} 2^{i-1}$. So, for example, the empty set is identified with the number 0, and the bundle containing the second and the third item corresponds to the number $6 = 2 + 4 = 2^{2-1} + 2^{3-1}$. The (additive) utility function u is then defined as $u(B) = x_B$.

Since we only have two agents, any allocation A is fully described by the bundle that the first agents holds. For $A = B_1 \cup B_2$ we have $B_2 = \mathcal{G} \setminus B_1$ and $x_{B_2} = 2^m - 1 - x_{B_1}$. Hence, for some allocation $A = B_1 \cup B_2$, the Nash social welfare $sw_{nash}(A)$ only depends on the value of $x := x_{B_1}$, and we obtain

$$sw_{nash}(A) = \prod_{i \in \mathcal{N}} u_i(A) = x \cdot (2^m - 1 - x)$$

In order to show that the function sw_{nash} is injective on the set of allocations where $x \in \{0, \dots, 2^{m-1} - 1\}$ (which is a representative system for the equivalent classes), we consider the corresponding real (and continuous) function $N_m: \mathbb{R} \rightarrow \mathbb{R}$ with $N_m(x) = x \cdot (2^m - 1 - x)$, which is a second-degree parabola with maximum turning point at $x^* = 2^{m-1} - \frac{1}{2}$. Therefore, when restricted to the interval $[0, 2^{m-1} - 1]$, N_m is monotonically increasing and hence injective. Figure 5 provides an intuition.

Since sw_{nash} and \mathcal{I} are equivalent as described in Lemma 17, this completes the proof. □

On top of our restriction to symmetric scenarios with a modular utility function, Theorem 18 imposes the additional restriction of that single utility function being normalised and it only applies to the case of two agents. For our final result, we drop the latter two restrictions and show that the maximum number of deals in a path is still exponential. This result crucially depends on a lemma, which we prove first, that shows there are scenarios where any two *partial allocations* differ in the resulting Nash social welfare, unless the two partial allocations are equivalent.

In this context, a function $Z: \mathcal{N} \rightarrow 2^{\mathcal{G}}$ is called a *partial allocation* if $Z(i) \cap Z(j) = \emptyset$ for any $i \neq j$. In contrast to an allocation, we do not require all goods to be allocated (i.e., we allow $Z(1) \cup \dots \cup Z(n) \subsetneq \mathcal{G}$). Every partial allocation Z induces a *utility vector* $u(Z) = (u_1(Z), \dots, u_n(Z))$, and the Nash social welfare of a partial allocation Z is $sw_{nash}(Z) = \prod_{i \in \mathcal{N}} u_i(Z)$. Our equivalence relation for allocations also directly transfers to partial allocations (i.e., we say two partial allocations Z and Z' are equivalent if there is a permutation $\pi: \mathcal{N} \rightarrow \mathcal{N}$ such that for any $i \in \mathcal{N}$, we have $Z(\pi(i)) = Z'(i)$).

Lemma 19 *For any numbers $n, m \in \mathbb{N}$, there exists a symmetric scenario $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ with $|\mathcal{N}| = n$, $|\mathcal{G}| = m$ and modular utility functions such that any two distinct partial allocations are either equivalent or differ in the resulting Nash product.*




Fig. 5 Sketch of the function $N_3: \mathbb{R} \rightarrow \mathbb{R}$ with $N_3(x) = x \cdot (2^3 - 1 - x)$ as a member of the family of functions N_m used in the proof of Theorem 18, providing an intuition of why the constructed functions in N_m are injective on the considered intervals. The depicted function N_3 is injective on the interval $[0, 2^{3-1} - 1] = [0, 3]$. The values of the corresponding function $sw_{nash}: \{0, \dots, 2^3 - 1\} \rightarrow [0, 1]$ are marked with circles

The structure of the following proof is similar to the structure of the proof of Lemma 14. For given agents \mathcal{N} and goods \mathcal{G} we construct a suitable collection \mathcal{U} of utility functions in a recursive way. In the proof of Lemma 14, we identified \mathcal{U} with an $n \times 2^m$ matrix P and filled in the entries of P with the help of a family of functions induced by the already fixed entries of P . Lemma 19 treats scenarios with a single modular utility function (which is identical for all agents. In order to fully determine this utility function it is sufficient to specify the value of $u(\emptyset)$ and for any $g \in B$ the value of $u(\{g\})$. We again make use of a family of functions constructed recursively, here the recursion is on the elements of \mathcal{G} . The values will be induced by those utility vectors which are already fully specified at the current stage of the recursion.

In the following, for the sake of readability, we will use a simplified notation. For a partial allocation Z , we denote by z the induced utility vector, i.e., $u(Z) = z = (z_1, \dots, z_n)$. We introduce the equivalence relation \sim on the set of utility vectors via $z \sim z'$ if and only if Z and Z' are equivalent. (Note that this implies that z' is a permutation of z .) In order to prove the lemma, we have to show that there is a scenario with

$$\prod_{i=1}^n z_i = \prod_{i=1}^n z'_i \implies z \sim z'. \tag{1}$$

for any pair Z, Z' of partial allocations (with induced utility vectors $z = (z_1, \dots, z_n)$ and $z' = (z'_1, \dots, z'_n)$, respectively) in this scenario.

of Lemma 19 Consider the scenario $(\mathcal{N}, \mathcal{G}, \mathcal{U})$, with $\mathcal{N} = \{1, \dots, n\}$ and $|\mathcal{G}| = m$. We assume a fixed ordering of the goods $\mathcal{G} = \{g_1, \dots, g_m\}$. We construct \mathcal{U} by finding feasible values for the marginal utilities of the goods.

Let Z^ℓ be the space of all partial allocations that allocate the goods $\{g_1, \dots, g_\ell\}$ to the agents. Hence, by $Z \in Z^\ell$, we mean that $\bigcup_{i \in \mathcal{N}} Z(i) \subseteq \{g_1, \dots, g_\ell\}$.

For some ℓ and already fixed values $u(\emptyset), u(\{g_1\}), \dots, u(\{g_\ell\})$, we say Z^ℓ is feasible if and only if Statement (1) is true for any pair $Z, Z' \in Z^\ell$.

Let Z_0 be the empty allocation ($Z_0(i) = \emptyset$ for all $i \in \mathcal{N}$) and Z_1 such that $Z_1(1) = \{g_1\}$ and $Z_1(i) = \emptyset$ for all $i \neq 1$. The partition of Z^1 into equivalence classes contains only two equivalence classes of partial allocations. The first is $[Z_0]$ and only contains the empty allocation, the second is $[Z_1]$ and contains the n partial allocations that allocate the good g_1 to one of the agents and the empty set to the remaining agents.

We start the recursion with $\ell = 1$: If we set $u(\emptyset) = 1$ and $u(\{g_1\}) = 2$, then Z^1 is feasible. We hence have the following two cases.

$$\begin{aligned} Z \in [Z_0]: \quad Z(i) &\stackrel{\forall i}{=} \emptyset, \quad z = z_0 = (1, 1, \dots, 1), \quad sw_{nash}(Z) = \prod_{i \in \mathcal{N}} z_i = 1, \\ Z \in [Z_1]: \quad Z(j) &\stackrel{\exists j}{=} \{g_1\}, \quad z \sim z_1 = (2, 1, \dots, 1), \quad sw_{nash}(Z) = \prod_{i \in \mathcal{N}} z_i = 2. \\ Z(i) &\stackrel{\forall i \neq j}{=} \emptyset \end{aligned}$$

To gain an intuition for the recursion step, we set $u(\{g_2\}) = \mathbf{x} + u(\emptyset) = \mathbf{x} + 1$ for a real number \mathbf{x} yet to be defined and show the current state of Z^2 . The partition of Z^2 into equivalence classes contains five equivalence classes: $[Z_0]$ and $[Z_1]$ are the same as before. Furthermore, we have $[Z_2], [Z_3], [Z_4]$, where

$$\begin{aligned} Z_2: \quad Z_2(1) &= \{g_2\} & Z_2(i) &\stackrel{\forall i \neq 1}{=} \emptyset, \\ Z_3: \quad Z_3(1) &= \{g_2\} \quad Z_3(2) = \{g_1\} & Z_3(i) &\stackrel{\forall i \neq 1, 2}{=} \emptyset, \\ Z_4: \quad Z_4(1) &= \{g_1, g_2\} & Z_4(i) &\stackrel{\forall i \neq 1}{=} \emptyset. \end{aligned}$$

Hence, we obtain

$$\begin{aligned}
 Z \in [Z_0]: \quad z = z_0 = (1, 1, \dots, 1) & \quad sw_{nash}(Z) = \prod_{i \in \mathcal{N}} z_i = 1, \\
 Z \in [Z_1]: \quad z \sim z_1 = (2, 1, \dots, 1) & \quad sw_{nash}(Z) = \prod_{i \in \mathcal{N}} z_i = 2, \\
 Z \in [Z_2]: \quad z \sim z_2 = (1 + \mathbf{x}, 1, \dots, 1) & \quad sw_{nash}(Z) = \prod_{i \in \mathcal{N}} z_i = 1 + \mathbf{x}, \\
 Z \in [Z_3]: \quad z \sim z_3 = (1 + \mathbf{x}, 2, \dots, 1) & \quad sw_{nash}(Z) = \prod_{i \in \mathcal{N}} z_i = (1 + \mathbf{x}) \cdot 2, \\
 Z \in [Z_4]: \quad z \sim z_4 = (2 + \mathbf{x}, 1, \dots, 1) & \quad sw_{nash}(Z) = \prod_{i \in \mathcal{N}} z_i = 2 + \mathbf{x}.
 \end{aligned}$$

Clearly, we cannot set $\mathbf{x} = 0$ or $\mathbf{x} = 1$ if we want to fulfill (1). But if, for instance, we set $u(\{g_2\}) = 2$, then Z^2 is feasible.

The idea is hence to fix new values of \mathbf{x} that are ‘big enough’ (again similar to the approach we used in Lemma 14).

We now describe the recursion step. Assume we have $v(g_1), \dots, v(g_\ell)$ such that Z^ℓ is feasible. We have to define a value for $u(\{g_{\ell+1}\})$ such that $Z^{\ell+1}$ is feasible as well. For any $Z \in Z^\ell$ and $i \in \{1, \dots, n\}$ we define a function F_z^i as

$$\begin{aligned}
 F_z^i: \mathbb{R} & \rightarrow \mathbb{R}, \\
 x & \mapsto \left(\prod_{j \neq i} z_j \right) \cdot (z_i + x).
 \end{aligned}$$

Thus, $F_z^i(x) = \prod_{i=1}^n z_i + \left(\prod_{j \neq i} z_j \right) \cdot x$, so any member of the family $\{F_z^i\}$ is of the form $ax + b$ with $a, b \geq 1$ and therefore a strictly increasing linear function. Because of this, for any $Z, Z' \in Z^\ell$ and $1 \leq j_z, j_{z'} \leq n$, the functions $F_z^{j_z}$ and $F_{z'}^{j_{z'}}$ only intersect twice if they are identical, which means that their coefficients are the same, hence only if

$$\prod_{i=1}^n z_i = \prod_{i=1}^n z'_i \quad \text{and} \quad \prod_{i \neq j_z} z_i = \prod_{i \neq j_{z'}} z'_i \tag{2}$$

holds. So by assumption, in this case we already have $z \sim z'$.

Let x_1 be the biggest value such that two elements of $\{F_z^i\}$ which do not fulfill (2) intersect in x_1 . We calculate $x_0 \in \mathbb{R}_{\geq 1}$ with

$$x_0 = \max_{z=u(Z), Z \in Z^{\leq \ell}} \prod_{i=1}^n z_i \tag{3}$$

and set $v(g_{\ell+1}) = \max \{x_0, x_1 + 1\}$.⁷ To show that $Z^{\ell+1}$ is feasible, we have to show that for any pair $Z, Z' \in Z^{\ell+1}$, Statement (1) is true. For any pair $Z, Z' \in Z^\ell \subsetneq Z^{\ell+1}$, this is true by assumption.

For the following, note that by construction, for any Z and i with $g_\ell \in Z(i)$, we have $sw_{nash}(Z) = \prod_{j \neq i} z_j \cdot (z_i + v(g_\ell))$. As $v(g_{\ell+1}) \geq x_1$, Statement (1) is also true for any pair $Z, Z' \in Z^{\ell+1} \setminus Z^\ell$. Finally, as $v(g_{\ell+1}) \geq x_0$ for any $Z \in Z^\ell$ and $Z' \in Z^{\ell+1} \setminus Z^\ell$, we have

$$sw_{nash}(Z) = \prod_{i=1}^n z_i \leq x_0 < x_0 + 1 \leq \prod_{i=1}^n z'_i = sw_{nash}(Z').$$

This completes the proof. □

⁷For $\ell = 1$ this would be $v(g_2) = \max \{2, 0 + 1\} = 2$.

We are now ready to state our result on the maximal path length to an optimal allocation for restricted scenarios with arbitrary numbers of agents. Note that the number of equivalence classes of allocations is equal to the number of possible partitions of \mathcal{G} of size at most $|\mathcal{N}|$. Under the assumption of $|\mathcal{G}| = m > n = |\mathcal{N}|$ we denote this number by \mathcal{B}_m^n . We have

$$\mathcal{B}_m^n = \sum_{k=1}^n S_k^m,$$

where S_k^m is the number of ways to partition a set of cardinality m into exactly k nonempty subsets. Note the connection to the Stirling numbers and the Bell numbers [4, 23, 29]. In particular \mathcal{B}_m^n grows exponentially (at least with respect to m).

Theorem 20 *In a symmetric scenario $(\mathcal{N}, \mathcal{G}, \mathcal{U})$ with modular utility functions and $|\mathcal{G}| = m > n = |\mathcal{N}|$, a sequence of Atkinson deals leading to an allocation that minimises inequality, as defined by the Atkinson index, can consist of up to $\mathcal{B}_m^n - 1$ deals, but not more.*

Proof The set of all allocations corresponds to the set $Z^m \setminus Z^{m-1}$ described in the proof of Lemma 19. Hence, the Nash product of any two allocations differs, unless the two allocations are equivalent. As the scenario fulfils the conditions of Lemma 17, the Nash product and the Atkinson index rank allocations consistently, so it also is the case that any two allocations differ in terms of their Atkinson index unless they are equivalent. The claim then follows from the fact that \mathcal{B}_m^n is the number of equivalence classes of allocations. \square

4 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 5) 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 3), and the fact that the essence of what it means to reduce inequality cannot be captured locally (Proposition 4). While most other social criteria studied in the context of multiagent resource allocation also require us to solve computationally intractable optimisation problems [11], 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 [12].

While Theorem 5 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 9). 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 15). These negative results do not change significantly even in very restricted scenarios, motivated by models of income distribution, where all agents have the same utility function. 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, and research in AI and Computer Science more generally. The very few exceptions include the works of Lesca and Perny [25], Endriss [15], and Gemici et al. [19]. This is so despite the fact that inequality indices are widely studied and used in practice across much of the social sciences.

Our work suggests a number of very concrete avenues for future research. First, we may ask whether a similar analysis is possible for other inequality indices. For the Gini index [20]—which is the most widely used index in practice, even if it is generally considered inferior to the Atkinson index from a normative point of view [34]—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 [35], 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, one may ask how obstructive our negative results are in practice. To address this question, one might generate a large set of scenarios (using synthetic preferences or preferences extracted from a specific real-world problem) and simulate what happens when, at every stage in the process, the agents choose one of the Atkinson deals currently available to them. This choice could be random (possibly giving higher weight to structurally simpler deals) or reflect some suitable behavioural assumptions about agents. 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 9). Similarly, one could analyse the average structural complexity of deals, such as the average maximal number of agents involved in any one deal in a sequence leading to an optimal allocation, or one could count the average number of deals contracted in such a system (to assess the practical relevance of Theorem 15). All of these are important questions of a more empirical nature than the fundamental analysis we have carried out in this paper. Addressing these questions would require significant original research to arrive at reasonable models about agent preferences and agent behaviour in specific application domains of interest, but we believe that such an investment would have the potential to be very fruitful and encourage the research community to take up this challenge.

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

Funding Open Access funding enabled and organized by Projekt DEAL. This work was partly supported by COST Action IC1205 on Computational Social Choice.

Financial/Non-financial interests The authors have no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Allison, P.D.: Measures of inequality. *Am. Sociol. Rev.* **43**(6), 865–880 (1978)
2. Atkinson, A.B.: On the measurement of inequality. *J. Econ. Theory* **2**, 244–263 (1970)
3. Atkinson, A.B.: *Inequality* Harvard University Press (2015)
4. Bell, E.T.: Exponential numbers. *Am. Math. Mon.* **41**(7), 411–419 (1934)
5. Bertsimas, D., Farias, V.F., Trichakis, N.: The price of fairness. *Oper. Res.* **59**(1), 17–31 (2011)
6. Bouveret, S., Chevaleyre, Y., Maudet, N.: 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)
7. Bouveret, S., Lemaître, M.: Computing leximin-optimal solutions in constraint networks. *Artif. Intell.* **173**(2), 343–364 (2009)
8. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: The efficiency of fair division. *Theory Comput. Syst.* **50**(4), 589–610 (2012)
9. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. In: *Proceedings of the 17th ACM Conference on Economics and Computation (EC-2016)*, pp. 305–322 ACM (2016)
10. Cauchy, A.L.: *Cours d'analyse. OEuvres Complètes Bd 3* (1821)
11. Chevaleyre, Y., Dunne, P.E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J.A., Sousa, P.: Issues in multiagent resource allocation. *Informatica* **30**, 3–31 (2006)
12. Chevaleyre, Y., Endriss, U., Maudet, N.: Distributed fair allocation of indivisible goods. *Artif. Intell.* **242**, 1–22 (2017)
13. Dalton, H.: The measurement of the inequality of incomes. *Econ. J.* **30**(119), 348–361 (1920)
14. Dunne, P.E., Wooldridge, M., Laurence, M.: The complexity of contract negotiation. *Artif. Intell.* **164**(1–2), 23–46 (2005)
15. Endriss, U.: Reduction of economic inequality in combinatorial domains. In: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2013)*, pp. 175–182 IFAAMAS (2013)
16. Endriss, U., Maudet, N., Sadri, F., Toni, F.: Negotiating socially optimal allocations of resources. *J. Artif. Intell. Res. (JAIR)* **25**, 315–348 (2006)
17. Freeman, R., Sikdar, S., Vaish, R., Xia, L.: Equitable allocations of indivisible goods. In: S. Kraus, editor *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-2019)*, pp. 280–286. IJCAI (2019)
18. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York (1979)
19. Gemici, K., Koutsoupias, E., Monnot, B., Papadimitriou, C.H., Piliouras, G.: Wealth inequality and the price of anarchy. In: *Proceedings of the 36th International Symposium on Theoretical Aspects of Computer Science (STACS-2019)* (2019)
20. Gini, C.: *Variabilità e Mutabilità*. C. Cuppini, Bologna (1912)
21. Gourvès, L., Monnot, J., Tlilane, L.: Near fairness in matroids. In: T. Schaub, G. Friedrich, and B. O'Sullivan, editors, *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI-2014)*, vol. 263 of *Frontiers in Artificial Intelligence and Applications*, pp. 393–398. IOS Press (2014)
22. Huhns, M.N., Malhotra, A.K.: Negotiating for goods and services. *IEEE Internet Comput.* **3**(4), 97 (1999)
23. Knuth, D.E.: *The Art of Computer Programming, vol. 1: Fundamental Algorithms, 3rd Edition*. Addison-wesley Professional, 3 edition 7 (1997)
24. Kondor, Y.: An old-new measure of income inequality. *Econometrica* **39**(6), 1041–42 (1971)
25. Lesca, J., Perny, P.: LP Solvable models for multiagent fair allocation problems. In: *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-2010)*, pp. 393–398 (2010)
26. Moulin, H.: *Axioms of cooperative decision making* Cambridge University Press (1988)
27. Nash Jr, J.F.: The bargaining problem. *Econometrica* **18**(2), 155–162 (1950)
28. Ramezani, S., Endriss, U.: Nash social welfare in multiagent resource allocation In *Agent-Mediated Electronic commerce: Designing Trading Strategies and Mechanisms for Electronic Markets*, vol 59 of *Lecture Notes in Business Information Processing*, pp. 117–131 Springer-Verlag (2010)
29. Rennie, B.C., Dobson, A.J.: On Stirling numbers of the second kind. *J. Comb. Theory* **7**(2), 116–121 (1969)
30. Rosenschein, J.S., Zlotkin, G.: *Rules of encounter: designing conventions for automated negotiation among computers* MIT press (1994)

31. Sandholm, T.W.: Contract types for satisficing task allocation: I Theoretical results Inproceedings of the 1998 AAAI Spring Symposium on Satisficing Models (1998)
32. Sandholm, T.W.: Distributed rational decision making. In: G. Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pp. 201–258. MIT Press (1999)
33. Schneckenburger, S., Dorn, B., Endriss, U.: The Atkinson inequality index in multiagent resource allocation. In: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017) IFAAMAS (2017)*
34. Sen, A.: *On economic inequality* Oxford University Press (1973)
35. Theil, H.: The information approach to demand analysis. *Econometrica* **33**(1), 67–87 (1965)
36. Vazirani, V.V.: *Approximation algorithms* Springer-Verlag (2001)

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.