

# Fair Division: The Computer Scientist’s Perspective

Toby Walsh

UNSW Sydney and Data61  
tw@cse.unsw.edu.au

## Abstract

I survey recent progress on a classic and challenging problem in social choice: the fair division of indivisible items. I discuss how a computational perspective has provided interesting insights into and understanding of how to divide items fairly and efficiently. This has involved bringing to bear tools such as those used in knowledge representation, computational complexity, approximation methods, game theory, online analysis and communication complexity.

## 1 Introduction

Fair division is an important problem facing society today as increasing economical, environmental, and other pressures require us to try to do more with limited resources. Consider allocating donated organs to patients, pupils to schools, doctors to hospitals, donated food to charities, or water quotas to farmers. How do we do this both fairly and efficiently?

Fair division is a problem with a long history. The “divide and choose” mechanism for fair division, for example, is mentioned in the Book of Genesis when Abraham and Lot use it to divide the land of Canaan. However, the theoretical study of fair division dates back only to the end of the Second World War when it attracted the attention of mathematicians such as Hugo Steinhaus, Bronisław Knaster and Stefan Banach [Steinhaus, 1948]. It later attracted the attention of economists like Steven Brams and Vincent Crawford (e.g. [Brams and Taylor, 1996]).

Much more recently, the fair division problem has attracted the attention of computer scientists, in particular those working within artificial intelligence. We argue here that this has led to a range of computational questions to be asked, and computational tools to be applied, that have advanced both our understanding of and ability to find fair and efficient allocations. I focus this survey on the fair division of indivisible goods, but I could have just as easily chosen the fair division of divisible goods. An emerging trend, which I discuss briefly in the conclusion, is to the fair division of indivisible bads (aka “chores”) and of mixed manna (goods and bads).

## 2 Formal model

I briefly give some notation and concepts that will be used in the rest of this survey. We have  $n$  agents and  $m$  indivisible items,  $a_1$  to  $a_m$ . For each agent  $i \in [1, n]$  and  $j \in [1, m]$ , the agent  $i$  has utility  $u_i(a_j) \geq 0$  for the item  $a_j$ . We will mostly limit ourselves to additive utilities where the utility to agent  $i$  for a set of items  $B$  is  $u_i(B) = \sum_{a \in B} u_i(a)$ . One special case of additive utilities is Borda scores where the  $j$ th most preferred item for agent  $i$  has utility  $m - j + 1$ . Another special case of additive utilities is lexicographical scores where the  $j$ th most preferred item for agent  $i$  has utility  $2^{m-j}$ . We will also mostly consider fair division problems involving just goods where items have non-negative utility. This contrasts to bads where all agents have negative utility for the items. An allocation is a partition of the  $m$  items into  $n$  distinguished sets,  $A_1$  to  $A_n$  where  $A_i$  represents the items allocated to agent  $i$  and  $A_i \cap A_j = \emptyset$  for  $i \neq j$ . We consider a number of welfare notions for an allocation  $A_1$  to  $A_n$ : the utilitarian welfare which is  $\sum_{i=1}^n u_i(A_i)$ , the egalitarian welfare which is  $\min_{i \in [1, n]} u_i(A_i)$ , and the Nash welfare which is  $\prod_{i=1}^n u_i(A_i)$ .

A fundamental fairness notion in fair division is envy-freeness. The “divide and choose” mechanism mentioned earlier will return allocations that are envy-free assuming risk averse agents. An allocation  $A_1$  to  $A_n$  is envy-free (EF) iff for any pair of agents  $i, j \in [1, n]$  we have  $u_i(A_i) \geq u_i(A_j)$ . Note that an envy-free allocation may not exist. Consider two agents that both have positive non-zero utility for a single good. A relaxation of envy-freeness for indivisible goods that can always be achieved is envy-freeness up to one good. An allocation  $A_1$  to  $A_n$  is envy-free up to one good (EF1) iff for any pair of exists  $i, j \in [1, n]$  there exists a good  $o_k$  with  $u_i(A_i) \geq u_i(A_j \setminus \{o_k\})$ .

Another fundamental notion in fair division is Pareto efficiency. We don’t want an allocation which all agents think can be improved upon or at least not made worse. We say that an allocation  $A_1$  to  $A_n$  is Pareto dominated iff there is another allocation  $B_1$  to  $B_n$  with  $u_i(B_i) \geq u_i(A_i)$  for all agents  $i \in [1, n]$  and for one agent  $j \in [1, n]$  we have  $u_j(B_j) > u_j(A_j)$ . An allocation is Pareto efficient iff there is no other allocation that Pareto dominates it.

We consider mechanisms that allocate items to agents. For instance, with the sequential allocation mechanism, agents take turns to pick an item according to a picking order. This

can implement mechanisms like round robin (aka “draft”) where agents pick one item each in a fixed order, and this fixed order is then repeated until no more items are left. Another interesting mechanism is the probabilistic serial mechanism [Bogomolnaia and Moulin, 2001]. Agents simultaneously “eat” their most preferred remaining item at some uniform rate. In general, this results in a randomized allocation, which can be implemented as a probability distribution over discrete allocations. An important property of mechanisms is whether they are strategy proof. Do agents have any incentive to behave strategically? We say that a mechanism is strategy proof iff it is in the best interests of an agent to act sincerely (e.g. in the case of the sequential allocation mechanism, to pick their most preferred item at every turn).

### 3 Knowledge representation questions

A starting point for much research in AI/CS are questions of representation and reasoning. How, for instance, do represent a problem *compactly* so that we can reason about it *efficiently*? In fair division, such questions abound. For example, how do we compactly represent an agent’s preferences over the power set of different possible subsets of items that might be allocated so that we can determine tractably an efficient and fair outcome? This power set of items creates an immediate problem. We need to represent (and elicit) an agent’s preferences over an exponential number of possible subsets of items. Often then a domain restriction is supposed like dichotomous preferences or additive utilities to ensure preference can be expressed compactly. However, this alone may not be enough to ensure reasoning is efficient.

For example, Bouveret and Lang [2008] consider simple dichotomous preferences in which a subset of items is acceptable or it isn’t and this is specified by (models of) a propositional logical formula. Even in this simple and compact setting, deciding if an efficient and envy-free allocation exists is intractable in general. Disappointingly, this problem remains intractable under strong domain restrictions such as identical dichotomous preferences with just 2 agents.

**Theorem 1** (Proposition 6 and 7 in [Bouveret and Lang, 2008]). *Deciding the existence of an efficient and envy-free allocation with general dichotomous preferences specified by models of a propositional formula is  $\Sigma_2^P$ -complete in general. The problem remains NP-hard with just 2 agents and identical, monotonic and dichotomous preferences.*

Another common domain restriction is to additive utilities (e.g. [Beviá, 1998; Brams *et al.*, 2003; Chevaleyre *et al.*, 2008; de Keijzer *et al.*, 2009]). Additive utilities offer a compromise between simplicity and expressivity. Preferences are now compact to represent and elicit. However, we can no longer represent certain settings such as those involving diminishing returns (e.g. a second bicycle is of less utility to me than the first bicycle) or complementarities (e.g. the cricket bat is only of value to me with the cricket ball). Again, reasoning about even such a compact representation may still be challenging. For instance, finding an envy-free allocation with additive utilities is NP-hard even under relatively specialized conditions such as additive 0/1 utilities [Aziz *et al.*, 2014; Aziz *et al.*, 2015c].

### 4 Computational questions

One response to the sort of intractability results discussed in the last section is to consider simple polynomial time mechanisms for allocating items. Such mechanisms might return a “good” allocation, even if there is no guarantee of exact optimality. Such mechanisms can also address issues like trust (when they are decentralized) and manipulation (when they are strategy proof). Consider again the simple sequential allocation mechanism. Variants of this mechanism are used in many real world settings such as assigning students to courses at Harvard Business School [Budish, 2011]. In the sequential allocation mechanism, agents simply take turns to pick their most preferred item. Whilst simple, sequential allocation has some nice theoretical properties. For example, if the mechanism is operated in a round robin fashion then the allocation returned is guaranteed to be envy-free up to one good [Budish, 2011].

Such mechanisms for fair division invite a range of computational questions around their application. Consider again the sequential allocation mechanism. This is not strategy proof in general. It may not always be in your best interest sometimes to pick sincerely an item that is your most preferred remaining item. What then is the computational complexity of computing optimal play (e.g. the subgame perfect Nash equilibrium of the repeated sequential allocation game)? This problem remained open for a long time. Brams and Straffin [1979] first identified this computational question over forty years ago:

*“no algorithm is known which will produce optimal [strategic] play more efficiently than by checking many branches of the game tree*

It took more than three decades before computer scientists proved that this computational problem was intractable in general even under restrictions on preferences of the agents.

**Theorem 2** (Theorem 3 in [Kalinowski *et al.*, 2013b]). *With an unbounded number of agents and additive Borda utilities, the subgame perfect Nash equilibrium for the sequential allocation game is unique and is PSPACE-hard to compute.*

With just two agents, optimal play can be computed in linear time (Theorem 1 in [Kalinowski *et al.*, 2013b]). Unfortunately, this optimal play may give rather poor outcomes. The agent not making the final pick can play strategically earlier in the game to ensure that their opponent taking the last (and thus only remaining) item has their least preferred item. Their opponent won’t pick this item given any other choice, but strategic play will ensure that this is their final choice. The agent not taking the penultimate pick can play strategically earlier in the game to ensure that their opponent choosing an item in this penultimate pick has their least preferred item, discounting the item that will be picked in the final pick. And so on for earlier picks. As a result, optimal strategic play results in agents getting some of their least preferred items.

With three or a larger but bounded number of agents, the computational complexity of computing optimal play remains stubbornly open (Challenge 5 in [Walsh, 2016]). It was recently shown that an agent’s best response can be computed in polynomial time for any bounded number of agents using

dynamic programming [Xiao and Ling, 2020]. This, however, leaves open the fixed point.

## 5 Calculation questions

Alongside such computational questions, there are many natural questions about calculation. For example, how do we calculate the welfare that agent can expect to have from applying a given mechanism? What is the probability that an allocation returned by a given mechanism is envy free? What is the least utility an agent can get? How do you calculate the best picking order?

Consider again the sequential allocation mechanism, and a setting where the preferences of agents are fully independent (i.e. all possible preference orders are equiprobable and independent of each other). Bouveret and Lang [2011] consider how to calculate the expected social welfare of allocations returned by applying the sequential allocation mechanism in this setting. They conjecture that calculating the expected welfare is NP-hard for a range of additive utilities such as Borda and lexicographic scores given the super-exponential number of possible preference profiles that need to be considered. Surprisingly, this conjecture turned out to be false. Calculating the expected egalitarian, utilitarian or Nash welfare takes only low-polynomial time for any type of additive utilities (Lemmas 1 and 2 in [Kalinowski *et al.*, 2013a]).

A related question is calculating the picking order that maximizes the expected social welfare. For small numbers of agents and items, Bouveret and Lang [2011] calculate the optimal picking orders to maximize the egalitarian or utilitarian welfare supposing Borda utilities. They conjecture that calculating the optimal picking order is NP-hard in general. For two agents and Borda utilities, it was subsequently shown that the optimal picking order simply alternates between the two agents.

**Theorem 3** (Theorem 1 in [Kalinowski *et al.*, 2013a]). *The expected utilitarian social welfare of the allocation returned by the sequential allocation mechanism is maximized by the strict alternating picking order for two agents supposing Borda utilities and the full independence assumption.*

This result holds even if the two agents act strategically [Kalinowski *et al.*, 2013a]. For three (or more) agents, the complexity of computing the optimal picking order remains open. For  $k$ -approval utilities (where each agent has  $k$  items with utility 1 and the rest with zero) it is not hard to see that the strict alternating picking order may not be optimal even with just two agents. Kalinowski *et al.* [2013a] conjecture that the strict alternating picking order is optimal for all convex scoring functions (which includes both Borda and lexicographical scores).

## 6 Parameterized complexity questions

Another tool in the computer scientist’s toolkit in the face of intractability is parameterized complexity. Can we identify parameters under which the computational problem becomes polynomial? Consider again computing the subgame perfect Nash equilibrium for the sequential allocation game. Recall that this is PSPACE-hard in general. Strategic behaviour is

only worthwhile when agents have different utilities for an item. For example, if you value an item that I don’t, you may strategically delay choosing it since it might still be available in a later round. We say that an item is multi-valued if two agents assign it different utilities, otherwise we say that the item is single valued. Kalinowski *et al.* [2013b] prove that, for any number of agents, we can compute a subgame perfect Nash equilibrium in  $O(k!m^{k+1})$  time where  $k$  is the number of multi-valued items, and  $m$  is the number of single valued items. Thus, when the number of multi-value items  $k$  is bounded, the problem is polynomial.

As a second and more recent example, Brederick *et al.* [2019] provide fixed-parameter tractability results for a broad set of problems concerned with the envy-free and efficient fair division of indivisible items. For instance, they show that computing an envy-free and efficient allocation of indivisible goods is fixed-parameter tractable for the combined parameter of the number of agents together with the maximum value of the utility of an item supposing additive utilities (Theorem 1 in [Brederick *et al.*, 2019]). In principle, this implies polynomial time envy-free and efficient allocations supposing we have few agents and low maximum utility values. These results exploit a famous theorem of Lenstra concerning (the fixed-parameter tractability of) Integer Linear Programs (ILPs) with bounded dimension. It remains to be shown if the constructed ILPs are actually easy to solve in practice. As the authors remark, “algorithms designed for specific problems usually outperform this approach [ILPs of bounded dimension] both in theory and practice”.

## 7 Game theoretic questions

Not surprisingly, game theory provides an important perspective for fair division. We have already mentioned some computational questions around computing strategic actions. Recall, for instance, that the subgame perfect Nash equilibrium for the sequential allocation game is PSPACE-hard to compute in general, but linear time when limited to just two agents. More generally, our understanding of fair division has advanced by asking questions from the interface between computer science and game theory.

Consider, for instance, the probabilistic serial mechanism [Bogomolnaia and Moulin, 2001]. This is envy free and Pareto efficient ex ante but is not strategy proof. The probabilistic serial mechanism is not strategy proof as an agent may not “eat” their most preferred remaining item if other agents do not value this item highly, in deference to eating earlier a less preferred item that might otherwise be eaten by other agents. Although the probabilistic serial mechanism is manipulable, finding an optimal manipulation is in general computationally intractable even if an agent has complete knowledge about the preferences of other agents. Computational complexity may in this case be a welcome barrier against agents behaving strategically.

**Theorem 4** (Theorems 2 and 3 in [Aziz *et al.*, 2015b]). *Computing a best response to the probabilistic serial mechanism game in order to maximize an agent’s expected utility is NP-hard in general. However, it takes linear time for just two agents.*

More generally, a pure Nash equilibrium is guaranteed to exist. Not surprisingly, given the complexity of best response, computing such an equilibrium is NP-hard. Indeed, even verifying whether the preferences reported by agents constitute a pure Nash equilibrium is coNP-complete (Theorem 3 in [Aziz *et al.*, 2015a]). On the other hand, with just two agents, a pure Nash equilibrium can be computed in linear time (Theorem 5 in [Aziz *et al.*, 2015a]). Pleasingly, this results in the same outcome as sincere reporting. Thus, whilst strategic behaviours are easy to compute in the case of just two agents, they do not change the final outcome.

## 8 Online questions

Traditional models of fair division usually suppose all information about the problem is present at the same time. However, computer scientists are used to dealing with online problems where the problem may be revealed over time. Indeed, many real world fair division problems have such a form with either the agents, or the resources to be allocated, or both not being fixed and potentially changing over time (e.g. [Walsh, 2014; Aleksandrov *et al.*, 2015; Walsh, 2015; Kash *et al.*, 2014; Mattei *et al.*, 2017; Mattei *et al.*, 2018; Gerding *et al.*, 2019; He *et al.*, 2019; Aleksandrov and Walsh, 2020]). Consider allocating deceased organs to patients, donated food to charities, water rights to farmers, etc. We often cannot wait till all resources are available, preferences known or agents present before starting to allocate the resources. For example, when a kidney is donated, it must be allocated to a patient within a few hours. As a second example, when food items arrive at a food bank, they must be allocated to charities promptly. As a third example, we might have to start allocating water rights to farmers before we know how much rain will come. Such online problems introduce a range of technical and other questions.

First, how do we adapt offline mechanisms to an online setting? Consider again the sequential allocation mechanism. Agents take turns to pick their most preferred item. But what if an agent is not present? Or if it is, their most preferred item is not? Second, how do adjust normative properties to take account of the online nature of such fair division problems? For instance, Aleksandrov and Walsh [2019] relax the definition of strategy-proofness to suppose past decisions are fixed while future decisions could still be strategic. This online form of strategy-proofness is less onerous to achieve. On the other hand, the online nature of such fair division problems can make other properties harder to achieve. For example, in offline fair division, Pareto efficiency and envy-freeness are possible to achieve simultaneously, e.g. the allocations returned by the probabilistic serial mechanism are Pareto efficient and envy-free ex ante [Bogomolnaia and Moulin, 2001]. In online fair division, however, such a combination is impossible to achieve in general.

**Theorem 5** (Theorem 9 in [Aleksandrov and Walsh, 2019]). *No online mechanism for fair division is Pareto efficient and envy-free ex ante.*

The reason that efficiency and fairness are not simultaneously possible is that the future is uncertain. An online mechanism must commit now to an allocation for the current item

which ensures an efficient but possibly not fair future, or a fair but possibly not efficient future.

## 9 Asymptotic questions

Another tool in the computer scientist’s toolbox is asymptotic analysis. Even in the offline setting, desirable fairness property such as envy-freeness cannot be guaranteed. Consider two agents and one indivisible item that both agents value. However, we can identify online mechanisms that limit how envy grows over time. For example, [Benade *et al.*, 2018] consider whether mechanisms can achieve *vanishing envy*. That is, after  $m$  items have arrived in an online fashion, if the maximum amount of envy an agent has for any other agent is  $envy_m$ , can we design mechanisms so that the ratio of  $\frac{envy_m}{m}$  goes to zero as  $m$  goes to infinity? This is indeed possible and, in fact, easy to achieve. [Benade *et al.*, 2018] show that randomly allocating items gives such an envy ratio that vanishes in expectation against an adaptive adversary. Unfortunately, it is not possible to do better than such a “blind” mechanism. Indeed, this random mechanism is asymptotically optimal up to logarithmic factors.

Asymptotic analysis can also identify the welfare cost on insisting on a fair allocation as we let the number of items grow. For instance, Caragiannis *et al.* [2009] define the price of envy-freeness as the ratio between the optimal utilitarian welfare and the maximum utilitarian welfare of any envy-free allocation and prove that this grows as  $O(n)$  for  $n$  items normalized to lie within  $[0, 1]$ . Note that this result is for problems in which an envy-free allocation exists.

**Theorem 6** (Theorem 5 in [Caragiannis *et al.*, 2009]). *For  $n$  agents, the price of envy-freeness is at least  $\frac{3n+7}{9} - O(\frac{1}{n})$ .*

## 10 Approximation questions

When faced with an intractable problem, computer scientists will often look for approximate solutions that can be computed efficiently. This has proved popular in fair division, starting with the work of Lipton *et al.* [2004] and an attractive relaxation of envy-freeness: envy-freeness up to one good (EF1). As noted earlier, the sequential allocation mechanism run in a round-robin fashion returns an allocation that is EF1. Unfortunately, whilst the allocation it returns is EF1, it may not be Pareto efficient. The mechanism is limited to allocations in which agents get a similar number of items. It may, however, be mutually beneficial to swap one item allocated to one agent for multiple items allocated to a second agent.

Caragiannis *et al.* [2016] have proposed instead computing an allocation that maximizes the Nash welfare. Theorem 3.2 in [Caragiannis *et al.*, 2016] proves that this is both EF1 and Pareto efficient. They also prove it provides a good approximation to another popular (yet sometimes unachievable) fairness property, the maximin share guarantee. The maximin share is the least amount of utility an agent receives if they divide items into bundles and they get the least valuable bundle. The Nash welfare mechanism is implemented and in daily use in the spliddit.org website. This leads Caragiannis *et al.* [2016] to conclude:

“... the MNW [maximal Nash welfare] solution exhibits an elusive combination of fairness and efficiency properties, and can be easily computed in practice. It provides the most practicable approach to date — arguably, the ultimate solution — for the division of indivisible goods under additive valuations...”

This claim is perhaps a little strong since it is NP-hard to compute an allocation with optimal Nash welfare [Nguyen *et al.*, 2014]. As to whether it is the “ultimate” solution, Barman, Murthy and Vaish [2018] recently gave a pseudo-polynomial-time algorithm for computing an allocation that is both EF1 and Pareto efficient. In addition this mechanism provides an 1.45-approximation of the optimal Nash welfare.

An weaker approximation than EF1 is envy-freeness up to any good (EFX) [Caragiannis *et al.*, 2016]. This lies strictly between EF and EF1. An allocation that is EF is EFX, one that is EFX is EF1, but not the reverse. Intriguingly, it is not known whether an EFX allocation always exists. There are, however, some special cases where we know that EFX allocations always exist. For instance, Amanatidis *et al.*, [2020] prove that an allocation that maximizes the Nash welfare is always EFX as long as there are at most two possible values for the goods. With three or more distinct values this implication no longer holds. As a consequence, an EFX allocation always exists with two valued goods.

Other fairness properties can be approximated in a similar way. For instance, another fairness property is proportionality (does an allocation exist in which each of the  $n$  agents get at least  $\frac{1}{n}$ th of their total utility?). With additive utilities, envy-freeness implies proportionality (but not vice versa). Although it is a weaker property, it is not weak enough to ensure that proportional allocations always exist. However, PROP1 allocations that are proportional up to one good always exist [Conitzer *et al.*, 2017]. Indeed, Barman and Krishnamurthy [2019] have recently given a polynomial time mechanism for computing an allocation that is both PROP1 and Pareto efficient.

## 11 Communication complexity questions

Another tool in the computer scientist’s toolkit is communication complexity. Plaut and Roughgarden [2019] initiate the use of the tools of communication complexity to study fair division problems with indivisible items. Their primary question is to determine whether players need to exchange an exponential amount of information to compute a fair allocation, or whether the problem can be solved using only polynomial communication between the agents. Interestingly these results show a significant difference between two and more than two agents. For more than two agents, exponential communication is needed to achieve an envy-free or proportional allocation even when randomization is allowed and valuations are submodular. This exponential continues to hold even when we demand allocations which are within a constant factor of envy-free or proportional. With two agents, there are more promising results. For example, with submodular valuations, there exists a deterministic protocol to achieve a proportional

allocation with communication that is polynomial in the number of items.

## 12 Empirical questions

The final tool in the computer scientist’s toolkit we will consider is empirical analysis. We can code up these mechanisms for fair division and run them on real and synthetic data. There are many empirical questions that we can then address. For instance, how hard are fair and efficient allocations to find in practice? Even if it is not guaranteed, does a mechanism return envy-free allocations with high probability? Do we observe hard instances around some phase boundary as in other combinatorial domains like satisfiability?

Such empirical analysis has been driven in part by access to real world data from websites such as spliddit.org [Goldman and Procaccia, 2014] and PrefLib.org [Mattei and Walsh, 2013]. Computers also make it easy to generate synthetic data on which to test fair division mechanisms. For instance, Dickerson *et al.* [2014] demonstrate that there is a sharp phase transition from non-existence to existence of envy-free allocations, and that finding such allocations is hardest at that transition. They also prove asymptotically that, even when the number of items is larger than the number of agents by a linear fraction, envy-free allocations are unlikely to exist under modest assumptions on the distribution of instances. On the other hand, when the number of items is larger by a logarithmic factor, envy-free allocations exist with high probability again under modest assumptions.

## 13 Conclusions

The computer scientist’s toolbox of questions and methods has enriched our understanding and ability to fairly and efficiently divide resources between multiple agents. It has let us identify the computational complexity of finding fair and efficient allocations, and to seek alternatives such as approximations to envy-freeness like EF1 and PROP1 that can be achieved tractably when exact methods are intractable in general. It has also led to the identification of parameters that give fixed parameter tractability. Other tools in the computer scientist’s toolbox like asymptotic analysis and communication complexity have also proved useful in understanding fair division mechanisms. But, as can be expected, applying these tools and answering these computational questions has led inevitably to new and important questions.

One area of recent attention is in the fair division of bads (sometimes called “chores”), and in the fair division of goods and bads (sometimes called “mixed manna”). See, for instance, [Bogomolnaia *et al.*, 2017; Aziz *et al.*, 2017; Aziz *et al.*, 2018; Aziz *et al.*, 2019; Bogomolnaia *et al.*, 2019]. The fair division of bads is not symmetric to the fair division of goods. In particular, giving no goods to an agent is pessimal, whilst giving no bads to an agent is optimal. For such reasons, results for the fair division of goods do not always map onto analogous results for the fair division of bads.

As well as interesting questions about the fair division of bads and of mixed manna, there are many other important open questions puzzling the research community. I shall end by mentioning just a few:

*Do EFX allocations always exist, even with 3 agents and additive utilities? How do we compute them efficiently?*

*Are there (computational) phase transitions in other properties besides envy-freeness? If so, how do their locations compare?*

*How do we compute optimal play efficiently for the sequential allocation mechanism and 3 agents?*

*Are there other appealing approximations of fairness that can be achieved tractably (especially in combination with Pareto efficiency)?*

*Can we apply other ideas from CSAI to help us understand and solve fair division problems (e.g. automated reasoning methods like SAT for theorem discovery, or machine learning methods for the design of fair division mechanisms)?*

*Does a computational lens help us understand and solve other problems in social choice?*

## References

- [Aleksandrov and Walsh, 2019] Martin Aleksandrov and Toby Walsh. Strategy-proofness, envy-freeness and pareto efficiency in online fair division with additive utilities. In Abhaya C. Nayak and Alok Sharma, editors, *16th Pacific Rim Int. Conf. on Artificial Intelligence (PRICAI 2019)*, volume 11670 of *LNCS*, 527–541. Springer, 2019.
- [Aleksandrov and Walsh, 2020] M. Aleksandrov and T. Walsh. Online fair division: A survey. In Vincent Conitzer and Fei Sha, editors, *Proc. of 34th AAAI Conf. on Artificial Intelligence*. AAAI Press, 2020.
- [Aleksandrov et al., 2015] M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. Online fair division: analysing a food bank problem. In Q. Yang and M. Wooldridge, editors, *Proc. of 24th Int. Joint Conf. on Artificial Intelligence, (IJCAI 2015)*, 2540–2546, 2015.
- [Amanatidis et al., 2020] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris. Maximum nash welfare and other stories about EFX. *CoRR*, abs/2001.09838, 2020.
- [Aziz et al., 2014] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. In A. Bazzan, M. Huhns, A. Lomuscio, and P. Scerri, editors, *Int. conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14*, 1305–1312, 2014.
- [Aziz et al., 2015a] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Equilibria under the probabilistic serial rule. In Q. Yang and M. Wooldridge, editors, *Proc. of 24th Int. Joint Conf. on Artificial Intelligence, (IJCAI 2015)*, 1105–1112, 2015.
- [Aziz et al., 2015b] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Manipulating the probabilistic serial rule. In G. Weiss, P. Yolum, R.H. Bordini, and E. Elkind, editors, *Proc. of the 2015 Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, 1451–1459. ACM, 2015.
- [Aziz et al., 2015c] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71 – 92, 2015.
- [Aziz et al., 2017] H. Aziz, G. Rauchecker, G. Schryen, and T. Walsh. Algorithms for max-min share fair allocation of indivisible chores. In S. Singh and S. Markovitch, editors, *Proc. of 31st AAAI Conf. on Artificial Intelligence*, 335–341. AAAI Press, 2017.
- [Aziz et al., 2018] Haris Aziz, Ioannis Caragiannis, and Ayumi Igarashi. Fair allocation of combinations of indivisible goods and chores. In *AI Workshop on agents and incentives in artificial intelligence, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*.
- [Aziz et al., 2019] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. In Sarit Kraus, editor, *Proc. of 28th Int. Joint Conf. on Artificial Intelligence*, 53–59, 2019.
- [Barman and Krishnamurthy, 2019] Siddharth Barman and Sanath Kumar Krishnamurthy. On the proximity of markets with integral equilibria. In *Proc. of 33rd AAAI Conf. on Artificial Intelligence, AAAI 2019*, 1748–1755. AAAI Press, 2019.
- [Barman et al., 2018] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Greedy algorithms for maximizing nash social welfare. In *Proc. of 17th Int. Conf. on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, 7–13, 2018.
- [Benade et al., 2018] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. How to make envy vanish over time. In Éva Tardos, Edith Elkind, and Rakesh Vohra, editors, *Proc. of 2018 ACM Conf. on Economics and Computation*, 593–610, 2018.
- [Beviá, 1998] Carmen Beviá. Fair allocation in a general model with indivisible goods. *Review of Economic Design*, 3(3):195–213, 1998.
- [Bogomolnaia and Moulin, 2001] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295 – 328, 2001.
- [Bogomolnaia et al., 2017] Anna Bogomolnaia, Hervé Moulin, Fedor Sandomirskiy, and Elena Yanovskaya. Competitive division of a mixed manna. *Econometrica*, 85(6):1847–1871, 2017.
- [Bogomolnaia et al., 2019] Anna Bogomolnaia, Hervé Moulin, Fedor Sandomirskiy, and Elena Yanovskaia. Dividing bads under additive utilities. *Social Choice and Welfare*, 52(3):395–417, Mar 2019.
- [Bouveret and Lang, 2008] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research (JAIR)*, 32:525–564, 2008.
- [Bouveret and Lang, 2011] S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In Toby Walsh, editor, *Proc. of 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011)*, 73–78. IJCAI/AAAI, 2011.
- [Brams and Straffin, 1979] S.J. Brams and P.D. Straffin. Prisoners’ dilemma and professional sports drafts. *The American Mathematical Monthly*, 86(2):pp. 80–88, 1979.

- [Brams and Taylor, 1996] S.J. Brams and A.D. Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, Cambridge, 1996.
- [Brams et al., 2003] Steven J. Brams, Paul H. Edelman, and Peter C. Fishburn. Fair division of indivisible items. *Theory and Decision*, 55(2):147–180, Sep 2003.
- [Bredereck et al., 2019] Robert Bredereck, Andrzej Kaczmarczyk, Dusan Knop, and Rolf Niedermeier. High-multiplicity fair allocation: Lenstra empowered by N-fold integer programming. In Anna Karlin, Nicole Immorlica, and Ramesh Johari, editors, *Proc. of 2019 ACM Conf. on Economics and Computation (EC 2019)*, 505–523. ACM, 2019.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis et al., 2009] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. The efficiency of fair division. In *5th Int. Workshop on Internet and Network Economics (WINE 2009)*, volume 5929 of *LNCS*, 475–482. Springer, 2009.
- [Caragiannis et al., 2016] I. Caragiannis, D. Kurokawa, H. Moulin, A. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. In *Proc. of 2016 ACM Conf. on Economics and Computation, EC '16*, 305–322, New York, NY, USA, 2016. ACM.
- [Chevalyere et al., 2008] Yann Chevalyere, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Multiagent resource allocation in k-additive domains: preference representation and complexity. *Annals of Operations Research*, 163(1):49–62, Oct 2008.
- [Conitzer et al., 2017] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin, editors, *Proc. of 2017 ACM Conf. on Economics and Computation, EC '17*, 629–646, 2017.
- [de Keijzer et al., 2009] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proc. of Algorithmic Decision Theory, 1st Int. Conf., ADT*, 98–110, 2009.
- [Dickerson et al., 2014] J.P. Dickerson, J.R. Goldman, J. Karp, A.D. Procaccia, and T. Sandholm. The computational rise and fall of fairness. In C.E. Brodley and P. Stone, editors, *Proc. of 28th AAAI Conf. on Artificial Intelligence*, 1405–1411, 2014.
- [Gerding et al., 2019] Enrico H. Gerding, Alvaro Perez-Diaz, Haris Aziz, Serge Gaspers, Antonia Marcu, Nicholas Mattei, and Toby Walsh. Fair online allocation of perishable goods and its application to electric vehicle charging. In Sarit Kraus, editor, *Proc. of 28th Int. Joint Conf. on Artificial Intelligence*, 5569–5575, 2019.
- [Goldman and Procaccia, 2014] Jonathan R. Goldman and Ariel D. Procaccia. Spliddit: unleashing fair division algorithms. *SIGecom Exchanges*, 13(2):41–46, 2014.
- [He et al., 2019] Jiafan He, Ariel D. Procaccia, Alexandros Psomas, and David Zeng. Achieving a fairer future by changing the past. In Sarit Kraus, editor, *Proc. of 28th Int. Joint Conf. on Artificial Intelligence*, 343–349, 2019.
- [Kalinowski et al., 2013a] T. Kalinowski, N. Narodytska, and T. Walsh. A social welfare optimal sequential allocation procedure. In *Proc. of the 23rd Int. Joint Conf. on Artificial Intelligence (IJCAI-2013)*. 2013.
- [Kalinowski et al., 2013b] T. Kalinowski, N. Narodytska, T. Walsh, and L. Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proc. of 27th AAAI Conf. on Artificial Intelligence (AAAI 2013)*. AAAI Press, 2013.
- [Kash et al., 2014] Ian A. Kash, Ariel D. Procaccia, and Nisarg Shah. No agent left behind: Dynamic fair division of multiple resources. *J. Artif. Intell. Res.*, 51:579–603, 2014.
- [Lipton et al., 2004] R.J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In J.S. Breese, J. Feigenbaum, and M.I. Seltzer, editors, *Proc. 5th ACM Conf. on Electronic Commerce (EC-2004)*, 125–131. ACM, 2004.
- [Mattei and Walsh, 2013] N. Mattei and T. Walsh. PrefLib: A library of preference data. In *Algorithmic Decision Theory, 3rd Int. Conf., (ADT 2013)*, LNAI. Springer, 2013.
- [Mattei et al., 2017] N. Mattei, A. Saffidine, and T. Walsh. Mechanisms for online organ matching. In *Proc. of 26th Int. Joint Conf. on Artificial Intelligence, IJCAI 2017*, 345–351, 2017.
- [Mattei et al., 2018] N. Mattei, A. Saffidine, and T. Walsh. Fairness in deceased organ matching. In *Proc. of 1st AAAI/ACM Conf. on AI, Ethics, and Society*, 2018.
- [Nguyen et al., 2014] Nhan-Tam Nguyen, Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous Agents and Multi-Agent Systems*, 28(2):256–289, 2014.
- [Plaut and Roughgarden, 2019] B. Plaut and T. Roughgarden. Communication complexity of discrete fair division. In *Proc. of 13th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '19*, 2014–2033, USA, 2019.
- [Steinhaus, 1948] H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- [Walsh, 2014] T. Walsh. Allocation in practice. In *Proc. of 37th German Conf. on Artificial Intelligence (KI-2014)*, LNAI, 13–24. Springer, 2014.
- [Walsh, 2015] T. Walsh. Challenges in resource and cost allocation. In *Proc. of 29th AAAI Conf. on AI*, 25–30. Association for Advancement of Artificial Intelligence, 2015.
- [Walsh, 2016] T. Walsh. Strategic behaviour when allocating indivisible goods. In D. Schuurmans and M. Wellman, editors, *Proc. of 30th AAAI Conf. on Artificial Intelligence*, 4177–4183. AAAI Press, 2016.
- [Xiao and Ling, 2020] M. Xiao and J. Ling. Algorithms for manipulating sequential allocation. In *Proc. of 34th AAAI Conf. on Artificial Intelligence*. AAAI Press, 2020.