Online Fair Division: A Survey

Martin Aleksandrov, Toby Walsh
Technical University Berlin and UNSW Sydney
Germany and Australia
martin.aleksandrov@tu-berlin.de, tw@cse.unsw.edu.au

Abstract

We survey a burgeoning and promising new research area that considers the online nature of many practical fair division problems. We identify wide variety of such online fair division problems, as well as discuss new mechanisms and normative properties that apply to this online setting. The online nature of such fair division problems provides both opportunities and challenges such as the possibility to develop new online mechanisms as well as the difficulty of dealing with an uncertain future.

Introduction

Fair division (Brams and Taylor 1996) is an important problem facing society today as increasing economical, environmental, and other pressures require us to try to do more with limited resources. Much previous work in fair division assumes the problem is offline and fixed. That is, we suppose that the agents being allocated resources, and the resources being allocated to these agents are all known and fixed. But practical reality is often quite different (Walsh 2014a; 2015). Fair division problems are often online, with either the agents, or the resources to be allocated, or both not being fixed and potentially changing over time. This presents a number of technical challenges and opportunities.

Consider allocating deceased organs to patients, donated food to charities, electric vehicles to charging stations, viewing slots to a telescope, 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, when allocating charging slots to electric cars, we may not know when or where cars will arrive for charging in the future but must commit to providing charging slots to cars now. As a fourth example, when allocating landing slots at an airport for the next hour, there might be considerable uncertainty about the demands in six hours time. As a fifth example, when scheduling viewing slots on an expensive telescope, we might need to commit to the current period, before future viewing conditions and demands are fully known.

The *online* nature of such fair division problems changes *the mechanisms* available to allocate items. For example, consider the well-known and widely used *sequential allocation* mechanism studied in (Bouveret and Lang 2011; Kalinowski et al. 2013; Kalinowski, Narodytska, and Walsh 2013; Aziz, Walsh, and Xia 2015; Aziz et al. 2016; Walsh 2016). This mechanism has agents taking turns to pick their most preferred remaining item. This is used in many offline settings due to its simplicity and nice normative properties. For instance, it can be used to return allocations where an agent's envy for the allocation of another agent is limited to at most one item. However, this mechanism cannot be used in an online setting as an agent's most preferred item may not be currently (or even ever) available. We, therefore, need to develop new *online* mechanisms that account for the fact that some items might never become available.

The *online* nature of such fair division problems also changes *the normative properties* that we might look to demand. For example, suppose we are interested in *strategy-proof* mechanisms that encourage agents to report their sincere preferences. In deciding if agents have any incentive to misreport preferences in an online setting, we may take into account the online nature of the problem, and thus consider fixed past decisions and unknown future. This leads to an *online* and weaker form of strategy-proofness. It is weaker because agents have committed to their past decisions and uncertainty about the future reduces their strategic options. On the plus side, this means that it may be easier to achieve strategy-proofness in the online setting than in the offline setting. On the minus side, this also means that it could be harder in the online setting to achieve other properties. For instance, items that arrive in the future may require us to change the allocation of items made in the present to ensure Pareto efficiency.

Dimensions of online fair division

We can characterize models of problems concerning online fair division along a number of orthogonal dimensions.

Resource types

We can classify online fair division problems according to whether resources are divisible or indivisible. For instance, (Walsh 2011) considers an online cake cutting problem in

which the resource ("the cake") is divisible, whilst (Aleksandrov et al. 2015; Aleksandrov and Walsh 2017c; Mattei, Saffidine, and Walsh 2017a; Aleksandrov and Walsh 2017a; 2017b) consider fair division problems where items arrive in an online fashion but are indivisible.

Further, we can characterize online fair division problems along the number and nature of those resources. Do we have one or multiple resources, and are those resources homogeneous or heterogeneous? For instance, (Walsh 2011) considers a single, heterogeneous and divisible cake, whilst (Kash, Procaccia, and Shah 2014) considers multiple, homogeneous and divisible resources (e.g. disk space, memory and CPUs in a server farm).

And, online fair division problems can be characterized on the number of resources allocated to each agent. Do we allocate a single resource to each agent (e.g. a single organ to each patient in (Mattei, Saffidine, and Walsh 2017b; 2018))? Or do we allocate multiple items to each agent (e.g. multiple food products in (Aleksandrov et al. 2015), multiple charging slots in (Gerding et al. 2019))?

Online features

Another dimension on which to decompose online fair division problems is whether agents, resources or both are online. For instance, (Walsh 2011) considers the resource ("the cake") to be fixed, and the agents being allocated this resource to arrive and depart over time. Think of people turning up over time for a birthday party and cutting off slices of the cake for people as they arrive. On the other hand, (Aleksandrov et al. 2015) considers the agents to be fixed, but the items to arrive over time and to be allocated immediately on arrival. Think of foods products being donated to a food bank over the course of the day but being allocated immediately to charities on donation. And, in the third case, both agents and items arrive over time. Think of an online organ matching problem (e.g. (Mattei, Saffidine, and Walsh 2018)) in which both the patients and the organs can arrive at particular points in time and, in addition, patients can also depart.

Mechanism features

Another dimension to consider is whether mechanisms are informed or uninformed about the future. In the informed setting (e.g. (He et al. 2019)), the mechanism might have information about items yet to arrive. In the uninformed setting, such information is not known. In this case, we may make an adversarial risk-averse assumption that agents act supposing the worst possible future (Walsh 2011). We can also distinguish between centralized and distributed mechanisms. Finally, we can distinguish between mechanisms that make decisions one by one (e.g. (Aleksandrov et al. 2015)) or in batches (e.g. (Benade et al. 2018)).

Formal background

We suppose there are n agents being allocated m items. Allocations may be whole (e.g. in the case of indivisible items), fractional (e.g. in the case of divisible items), or randomized (e.g. in the case of indivisible items, this represents a probability distribution over whole allocations). Often, as in many

other areas, we will assume additive utilities. Additive utilities offer a compromise between simplicity and expressivity. However, some work in this area has considered more general utilities (e.g. (Aleksandrov and Walsh 2019a) studies the more general class of monotone utilities). We consider a number of classical normative properties such as Pareto efficiency and envy-freeness. An allocation is *Pareto efficient* iff there is no other allocation where all agents have as much utility, and at least one agent has strictly more. An allocation is *envy-free* iff no agent has strictly greater utility for the items allocated to another agent than the utility for their own items. Envy-freeness is a desirable but often unachievable fairness property (consider two agents and one indivisible item that they both like). Therefore, we also consider a relaxation that can be always achieved: *envy-free up to one item* (EF1). An allocation is EF1 iff no agent has envy for another agent's bundle, supposing we can remove one item from this bundle. EF1 can be achieved with a simple round-robin mechanism that allocates the most desired item left to the next agent in a round-robin fashion. We say that a mechanism is Pareto efficient / envy-free / EF1 iff it only returns allocations that are Pareto efficient / envy-free / EF1. See (Aleksandrov and Walsh 2019b) for more details.

Technical approaches

There are a number of technical responses to the challenges and opportunities introduced by the online nature of such fair division problems.

Online mechanisms

One approach to deal with the challenges of online fair division is to propose new mechanisms that exploit the online nature of the fair division problem. Indeed, many offline mechanisms cannot be used in an online setting because they make assumptions, such as all the items are available at one time, that are violated. We, therefore, often need to "flip" mechanisms around to work in an online fashion.

Many mechanisms for the offline fair division of indivisible items take an agent-centered view where one or more agents get to pick an item or items from a set of items. For example, in offline fair division, the probabilistic serial mechanism (Bogomolnaia and Moulin 2001) has every agent "eating" their most preferred item at uniform speed, generating a randomized allocation of items to agents. Such offline mechanisms need to be adapted to work in an online setting as an agent's most preferred item might not be currently, or even ever available.

An online setting where items arrive over time also naturally invites an item-centric view. For example, (Aleksandrov et al. 2015) propose the LIKE mechanism where an item on arrival is allocated uniformly at random between agents that declare strictly positive utility for the item. This online mechanism has nice normative properties. For example, it is strategy-proof and envy-free in expectation. The LIKE mechanism can be seen as the online analog of the (offline) probabilistic serial mechanism with agents "eating" each next item which they like.

Many other online mechanisms for the fair division of indivisible items can be described as choosing uniformly be-

tween some subset of agents that are feasible for an arriving item. Thus, an agent from this set receives the item with *conditional probability* that is uniform with respect to the other agents that are feasible for the item. Such mechanisms return a probability distribution over allocations, and an actual allocation with some positive probability that is obtained as a product of the conditional randomizations. Some such mechanisms are studied in (Aleksandrov and Walsh 2019b):

- ONLINE SERIAL DICTATOR: there is a strict priority order σ of the agents prior to round one, and the unique feasible agent for each next item is the first agent in σ that bids positively for it.

- ONLINE RANDOM PRIORITY: this draws uniformly at random a strict priority order σ of the agents prior to round one, and runs ONLINE SERIAL DICTATOR with it.

- PARETO LIKE: an agent is feasible for an item if allocating the item to the agent is Pareto efficient ex post.

- LIKE: an agent is feasible for an item if they bid positively for it.

- BALANCED LIKE: an agent is feasible for an item if they bid positively for it, and they are amongst the agents bidding positively with the fewest items in the current allocation.

- MAXIMUM LIKE: an agent is feasible for an item if they are amongst the agents making the largest bid for it.

It should be clear that the online setting introduces a new design space in which to define mechanisms for fair division. It is likely that there are many interesting parts of this space yet to explore.

Online properties

Another technique to deal with the challenges of online fair division is to consider relaxed normative properties that are applicable to the online setting. For example, (Aleksandrov and Walsh 2019b) relax the definition of strategy-proofness to suppose past decisions are fixed while future decisions could still be strategic. A mechanism is **strategy-proof** if, an agent cannot improve their (expected) utility by bidding insincerely, supposing the entire information (i.e. the items, their arriving order and the mechanism) about the allocation process is available to each agent. That is, for each problem and agent a_i, $\overline{u}(u_{i1}, \ldots, u_{im}) \geq \overline{u}(v_{i1}, \ldots, v_{im})$ for any (possibly strategic) bids $v_{i1}, \ldots, v_{im}$ of agent a_i for the m items. On the other hand, we can relax this definition to a weaker online version of strategy-proofness. We say that a mechanism is **online strategy-proof** if, for each problem and round $j \in [1, m]$, $\overline{u}(u_{i1}, \ldots, u_{i(j-1)}, u_{ij}) \geq \overline{u}(u_{i1}, \ldots, u_{i(j-1)}, v_{ij})$ for any (possibly strategic) bid v_{ij} of agent a_i for the final jth item.

Online strategy-proofness is less onerous to achieve than strategy proofness. It supposes all earlier decisions are now fixed and the agent can only be strategic about the current decision. It is not hard to see that a strategy-proof mechanism is further online strategy-proof but the opposite may not hold. Indeed, whilst the only online mechanisms that are strategy-proof are rather unresponsive mechanisms such as

the random mechanism which just allocates items randomly irrespective of the bids of agents, the class of mechanisms that are online strategy-proof is far larger and can be responsive to the bids of agents. In fact, Theorem 2 in (Aleksandrov and Walsh 2019b) characterizes the class of online mechanisms that are online strategy-proof as "step" mechanisms that allocate items uniformly between agents that bid over some threshold of utility for an arriving item.

Stronger normative properties can also be useful in the online setting. For instance, (Aleksandrov and Walsh 2019b) introduce **shared envy-freeness** which requires that each pair of agents are envy-free of each other over only the items that both agents in the pair like in common. If you get an item that I don't like, it does not cause me any envy. Shared envy-freeness implies envy-freeness but not the other way around. Shared envy-freeness is one of the normative properties that distinguishes between the LIKE and BALANCED LIKE mechanisms. Table 1, which is taken from (Aleksandrov and Walsh 2019b), summarizes many of the differences in normative properties of the different online mechanisms discussed in the previous section.

Impossibility results

Another approach to understand the challenges of online fair division is to identify normative properties which are impossible to achieve in the online setting which can be achieved in the offline setting. For example, in offline fair division, Pareto efficiency and envy-freeness ex ante 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, these two normative properties are also possible to achieve simultaneously with 0/1 utilities, For instance, the LIKE mechanism, allocating each next item uniformly at random to any agent who declares strictly positive utility, is Pareto efficient and envy-free ex ante. But, with general (i.e. non-0/1) utilities, there is no online mechanism that can be both Pareto efficient and envy-free ex ante (proved first in (Aleksandrov and Walsh 2019b)).

Theorem 1 *With general utilities, no online mechanism is envy-free ex ante and Pareto efficient ex ante.*

Proof: Let us consider the below simple fair division problem with two agents and two arriving items.

agent	item 1	item 2
1	1	2
2	2	1

To ensure envy-freeness ex ante, a mechanism must give to each agent a probability of 1/2 for each item. This is not Pareto efficient ex ante. Giving each agent just their most valued item Pareto dominates. ◇

In Figure 1, we summarize a range of possibility and impossibility results that characterize whether there exist mechanisms that have combinations of properties such as strategy proofness (SP), envy-freeness ex ante (EFA), Pareto efficiency ex ante (PEA) and Pareto efficiency ex post (PEP).

Mechanism	SP	OSP	EFA	SEFA	EFP	SEFP	BEFP	PEA	PEP
	general cardinal utilities								
ONLINE RANDOM PRIORITY	✓	✓	✓	✓	×	×	×	×	✓
ONLINE SERIAL DICTATOR	✓	✓	×	×	×	×	×	✓	✓
MAXIMUM LIKE	×	×	×	×	×	×	×	✓	✓
PARETO LIKE	×	×	×	×	×	×	×	×	✓
LIKE	✓	✓	✓	✓	×	×	×	×	×
BALANCED LIKE	×	✓	×	×	×	×	×	×	×
	identical cardinal utilities								
LIKE	✓	✓	✓	✓	×	×	×	✓	✓
BALANCED LIKE	×	✓	✓	✓	×	×	×	✓	✓
	binary cardinal utilities								
LIKE	✓	✓	✓	✓	×	×	×	✓	✓
BALANCED LIKE	×	✓	✓	×	×	×	✓	✓	✓

Table 1: Normative properties of different online mechanisms. SP is strategy proof. OSP is online strategy proof. EFA is envy-free ex ante. SEFA is shared envy-free ex ante. EFP is envy-free ex post. SEFP is shared envy-free ex post. BEFP is bounded envy-free ex post. PEA is Pareto efficient ex ante. PEP is Pareto efficient ex post.

Intractability results

Another angle to understanding online fair division is to study the complexity of relevant computational problems such as computing outcomes or strategic actions. For example, consider again the LIKE mechanism and the related BALANCED LIKE mechanism which allocates the next item uniformly at random to any agent who declares strictly positive utility for the item from amongst the agents currently allocated the fewest items (Aleksandrov et al. 2015). The BALANCED LIKE mechanism is designed to balance the number of items allocated to each agent.

Now, the LIKE mechanism is strategy-proof. If an agent has strictly positive utility for an item, they should bid for it. On the other hand, the BALANCED LIKE mechanism is not strategy-proof, even when restricted to 0/1 utilities. Agents may decide not to bid now, even when they have strictly positive utility for an arriving item, as it may be more profitable to wait for an item that arrives later which is less competitive. Furthermore, they may thus receive multiple items in future.

In response, Aleksandrov and Walsh ((Aleksandrov and Walsh 2017a)) consider the computational complexity for an agent to compute strategic bids to improve their (expected) outcome. Such computational problems are intractable in general. So the fact that the BALANCED LIKE mechanism is not strategy-proof is tempered by the intractability of computing strategic actions. In more general terms, there is a trade-off between computability of strategic actions and achievability of fairer allocations.

Asymptotic guarantees

Another response to the challenge of online fair division is to try to achieve properties asymptotically. Even in the offline setting, desirable fairness property such as envy-freeness cannot be guaranteed. Consider, for instance, two agents and one indivisible item that both agents like. Envy-freeness is even harder to guarantee in the online setting. Suppose two agents both like the next indivisible item. If we want to ensure envy-freeness, we cannot allocate this item fairly without ultimately knowing what items will arrive next, that might compensate the agent not allocated the current item.

However, we can look for 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 envy that vanishes in expectation. Unfortunately, it is not possible to do better than such a "blind" mechanism. Indeed, the random mechanism is asymptotically optimal up to logarithmic factors.

Reallocating items

A final technique to deal with the online nature of fair division is to consider allocations that can be adjusted once subsequent items are revealed. By re-allocating items allocated in the past when new items arrive, we can perhaps restore desirable normative properties. For example, with two agents, additive valuations and m online items, (He et al. 2019) prove that any uninformed algorithm requires $\Theta(m)$ items to be re-allocated to ensure an EF1 allocation. Unsurprisingly, an informed algorithm requires no re-allocations to ensure EF1. For the case of three or more agents, (He et al. 2019) prove that even informed algorithms require $\Omega(m)$ items to be reallocated to ensure an EF1 allocation, and design an uninformed algorithm that makes do with $O(m^{3/2})$. These results leave open the question of whether there is a separation between the number of items that informed and uninformed algorithms must re-allocate in order to achieve EF1 (like the strict separation proven in the case of two agents) in the general case of three or more agents.

Conclusions

Online fair division is a promising and active research area that looks to take account of the online nature of many resource allocation problems. Whilst recent work has identi-
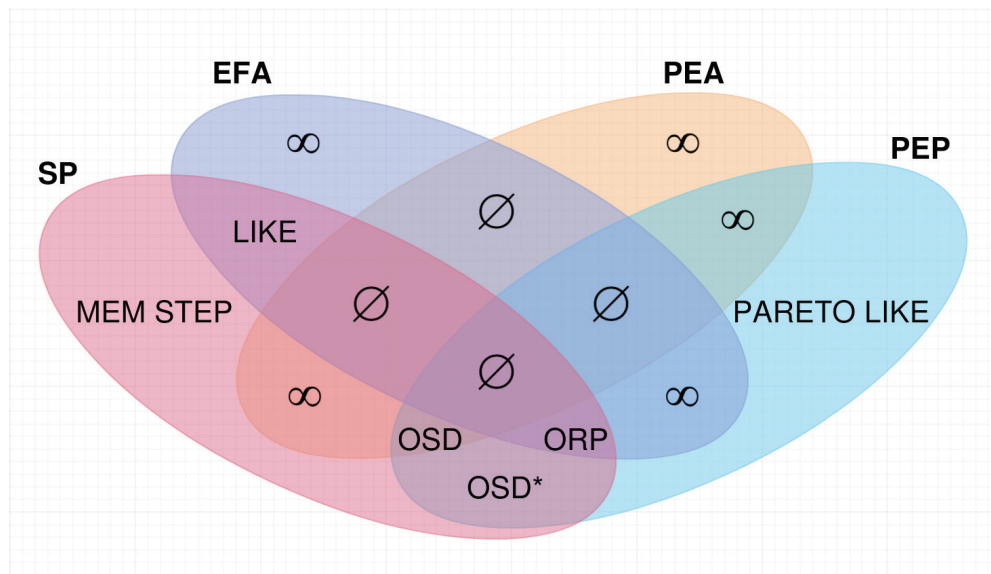
Figure 1: Summary of results about online mechanisms satisfying (combinations of) properties such as strategy proofness (SP), envy-freeness ex ante (EFA), Pareto efficiency ex ante (PEA) and Pareto efficiency ex post (PEP). Any part of the Venn diagram marked "∅" is impossible (e.g. no mechanism is envy-free ex ante and Pareto efficient ex post). Any part of the Venn diagram marked "∞" represents a combination of properties which an infinite number of different mechanisms can satisfy (e.g. any probabilistic combination of the LIKE and ORP mechanisms returns a probability distribution of allocations that is envy-free ex ante). OSD* represents any probability distribution of online serial dictator (OSD) mechanisms. MEM STEP is any "memoryless step" mechanism where the current item is allocated irrespective of the past (i.e. memoryless), and whether an agent is feasible is simply a step function of their bid. This figure is taken from (Aleksandrov and Walsh 2019b).

fied many important features of such problems, there are many research directions still to be considered. For example, how do we extend recent results such as (Aziz et al. 2015b; 2015a; Aziz, Schlotter, and Walsh 2016; Aziz et al. 2017; Gerding et al. 2019; Aziz et al. 2019; Aleksandrov and Walsh 2018) about the (offline) fair division of goods and bads to the online setting? As a second example, the preferences of agents are often highly correlated. How can we exploit this fact in an online setting? As a third example, we can apply online mechanisms like those discussed here to an offline setting by imposing an "artificial" arrival order to the agents or items. Can this help achieve desirable normative properties? And, as a fourth example, can we extend these ideas to the online versions of other problems in social choice such as capacitated facility location (Aziz et al. 2020) or peer assessment (Walsh 2014b).

Acknowledgments

References

Aleksandrov, M., and Walsh, T. 2017a. Expected outcomes and manipulations in online fair division. In *KI 2017: Advances in Artificial Intelligence - 40th Annual German Conference on AI, Dortmund, Germany, September 25-29, 2017, Proceedings*, 29–43.

Aleksandrov, M., and Walsh, T. 2017b. Most competitive mechanisms in online fair division. In *KI 2017: Advances in Artificial Intelligence - 40th Annual German Conference on AI, Dortmund, Germany, September 25-29, 2017, Proceedings*, 44–57.

Aleksandrov, M., and Walsh, T. 2017c. Pure Nash equilibria in online fair division. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 42–48.

Aleksandrov, M., and Walsh, T. 2018. Group envy freeness and group Pareto efficiency in fair division with indivisible items. In *KI 2018: Advances in Artificial Intelligence - 41st German Conference on AI, Berlin, Germany, September 24-28, 2018, Proceedings*, 57–72.

Aleksandrov, M., and Walsh, T. 2019a. Monotone and online fair division. In Benzmüller, C., and Stuckenschmidt, H., eds., *42nd German Conference on AI (KI 2019)*, volume 11793 of *Lecture Notes in Computer Science*, 60–75. Springer.

Aleksandrov, M., and Walsh, T. 2019b. Strategy-proofness, envy-freeness and pareto efficiency in online fair division with additive utilities. In Nayak, A. C., and Sharma, A., eds., *16th Pacific Rim International Conference on Artificial Intelligence (PRICAI 2019)*, volume 11670 of *Lecture Notes in Computer Science*, 527–541. Springer.

Aleksandrov, M.; Aziz, H.; Gaspers, S.; and Walsh, T. 2015.

Online fair division: analysing a food bank problem. In Yang, Q., and Wooldridge, M., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, (IJCAI 2015)*, 2540–2546.

Aziz, H.; Gaspers, S.; Mackenzie, S.; Mattei, N.; Narodytska, N.; and Walsh, T. 2015a. Manipulating the probabilistic serial rule. In Weiss, G.; Yolum, P.; Bordini, R.; and Elkind, E., eds., *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015*, 1451–1459.

Aziz, H.; Gaspers, S.; Mackenzie, S.; and Walsh, T. 2015b. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71 – 92.

Aziz, H.; Kalinowski, T.; Walsh, T.; and Xia, L. 2016. Welfare of sequential allocation mechanisms for indivisible goods. In Kaminka, G.; Fox, M.; Bouquet, P.; Hüllermeier, E.; Dignum, V.; Dignum, F.; and van Harmelen, F., eds., *ECAI 2016 - 22nd European Conference on Artificial Intelligence*, Frontiers in Artificial Intelligence and Applications, 787–794. IOS Press.

Aziz, H.; Rauchecker, G.; Schryen, G.; and Walsh, T. 2017. Algorithms for max-min share fair allocation of indivisible chores. In Singh, S., and Markovitch, S., eds., *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 335–341. AAAI Press.

Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2019. Fair allocation of indivisible goods and chores. In Kraus, S., ed., *Proceedings of 28th International Joint Conference on Artificial Intelligence*, 53–59.

Aziz, H.; Chan, H.; Lee, B.; Li, B.; and Walsh, T. 2020. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In Conitzer, V., and Sha, F., eds., *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*. AAAI Press.

Aziz, H.; Schlotter, I.; and Walsh, T. 2016. Control of fair division. In Kambhampati, S., ed., *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*, 67–73. IJCAI/AAAI Press.

Aziz, H.; Walsh, T.; and Xia, L. 2015. Possible and necessary allocations via sequential mechanisms. In Yang, Q., and Wooldridge, M., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, (IJCAI 2015)*.

Benade, G.; Kazachkov, A. M.; Procaccia, A. D.; and Psomas, C. 2018. How to make envy vanish over time. In Tardos, É.; Elkind, E.; and Vohra, R., eds., *Proceedings of the 2018 ACM Conference on Economics and Computation*, 593–610.

Bogomolnaia, A., and Moulin, H. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100(2):295 – 328.

Bouveret, S., and Lang, J. 2011. A general elicitation-free protocol for allocating indivisible goods. In Walsh, T., ed., *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, 73–78. IJCAI/AAAI.

Brams, S., and Taylor, A. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge: Cambridge University Press.

Gerding, E. H.; Perez-Diaz, A.; Aziz, H.; Gaspers, S.; Marcu, A.; Mattei, N.; and Walsh, T. 2019. Fair online allocation of perishable goods and its application to electric vehicle charging. In Kraus, S., ed., *Proceedings of 28th International Joint Conference on Artificial Intelligence*, 5569–5575.

He, J.; Procaccia, A. D.; Psomas, A.; and Zeng, D. 2019. Achieving a fairer future by changing the past. In Kraus, S., ed., *Proceedings of 28th International Joint Conference on Artificial Intelligence*, 343–349.

Kalinowski, T.; Narodytska, N.; Walsh, T.; and Xia, L. 2013. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*. AAAI Press.

Kalinowski, T.; Narodytska, N.; and Walsh, T. 2013. A social welfare optimal sequential allocation procedure. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-2013)*. International Joint Conference on Artificial Intelligence.

Kash, I. A.; Procaccia, A. D.; and Shah, N. 2014. No agent left behind: Dynamic fair division of multiple resources. *J. Artif. Intell. Res.* 51:579–603.

Mattei, N.; Saffidine, A.; and Walsh, T. 2017a. Mechanisms for online organ matching. In *Proceedings of the 26th International Joint Conference on AI*. International Joint Conference on Artificial Intelligence.

Mattei, N.; Saffidine, A.; and Walsh, T. 2017b. Mechanisms for online organ matching. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 345–351.

Mattei, N.; Saffidine, A.; and Walsh, T. 2018. Fairness in deceased organ matching. In *Proceedings of 1st AAAI/ACM Conference on AI, Ethics, and Society*.

Walsh, T. 2011. Online cake cutting. In *Algorithmic Decision Theory, Second International Conference, ADT 2011. Proceedings*, Lecture Notes in Computer Science. Springer.

Walsh, T. 2014a. Allocation in practice. In *Proceedings of 37th German Conference on Artificial Intelligence (KI-2014)*, Lecture Notes in Artificial Intelligence, 13–24. Springer.

Walsh, T. 2014b. The PeerRank Method for peer assessment. In Schaub, T., ed., *Proc. of the 21st European Conference on Artificial Intelligence (ECAI-2012)*, Frontiers in Artificial Intelligence and Applications. IOS Press.

Walsh, T. 2015. Challenges in resource and cost allocation. In *Proceedings of the 29th AAAI Conference on AI*, 25–30. Association for Advancement of Artificial Intelligence.

Walsh, T. 2016. Strategic behaviour when allocating indivisible goods. In Schuurmans, D., and Wellman, M., eds., *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 4177–4183. AAAI Press.