

# Developments in Multi-Agent Fair Allocation

**Haris Aziz**

UNSW Sydney and Data61  
Sydney, Australia  
haziz@cse.unsw.edu.au

## Abstract

Fairness is becoming an increasingly important concern when designing markets, allocation procedures, and computer systems. I survey some recent developments in the field of multi-agent fair allocation.

## 1 Introduction

Many important decisions in our lives are increasingly being made by computers. These decisions include whether we get a certain job or loan. It becomes imperative then that computer systems provide guarantees that the decision has been made transparently and fairly. For many settings such as two-sided matching, fairness considerations based on justified envy or diversity guarantees are enshrined in the letter of the law. Even for smaller day to day decisions, the appetite for fairness guarantees is naturally present among people.<sup>1</sup> At a bigger scale, the Internet leads to peer to peer interactions where fairness comes up as an important concern (see, e.g. Moulin (2019) and Kleinberg, Rabani, and Tardos (2001)).

The issue of what is fair and how to achieve fairness has been studied in several disciplines including economics (Moulin 2003; Thomson 2011), political science (Brams 2008; Balinski and Young 1982) and philosophy (Rescher 2002; Rawls 1971).<sup>2</sup>

I mention some recent developments in fair allocation in multi-agent systems. In particular, I discuss allocation of indivisible items (Section 2), allocation of divisible items (Section 3), two-sided matching (Section 4), and fairness in other social choice settings (Section 5). Finally, I wrap up the discussion in Section 6.

<sup>1</sup>See, for example, the interest in fair allocation portal [www.spliddit.org/](http://www.spliddit.org/).

<sup>2</sup>Despite the richness of fairness concepts and principles in these literatures, I will only focus on a handful of fairness concepts in this paper. Other fairness ideas include indices to measure inequality (see, e.g. Atkinson (1970)) or market-based approaches in which agents have equal budgets (see, e.g. Varian (1974)).

## 2 Allocation of Indivisible Items

When indivisible items are allocated among agents, guaranteeing the existence of fair allocations is impossible for many desirable notions of fairness. For example, if there are two agents and one indivisible valued good, allocating the good to either of the two agents will make the other agent envious. In other words, an envy-free allocation may not exist. Another challenge is computational. Suppose agents have positive additive utilities over the items. Then the problem of checking whether there exists an envy-free allocation is NP-complete even for 1-0 utilities (Aziz et al. 2015). Similarly the problems of computing allocations that maximize the egalitarian welfare (utility of the worst off agent) or Nash welfare (geometric mean of the utilities of agents) is NP-hard (see, e.g. Nguyen et al. (2014)).

Given these challenges, a couple of influential fairness concepts were proposed by Budish (2011). The first one is called *maxmin share (MMS)* fairness. The *maxmin fair share* of an agent is the best she can guarantee for herself if she is allowed to partition the items into bundles for the agents but then receives her least preferred bundle. An allocation satisfies MMS fairness if each agent's allocation is at least as preferred to her as her maxmin fair share. Although an MMS fair allocation is not guaranteed to exist even for additive utilities, the concept gives rise to approximation versions in which each agent aspires to get a certain percentage of her maxmin fair share (Procaccia and Wang 2014). This has led to several results including new algorithms that are faster, simpler, or provide better approximation guarantees, or work for more general families of utilities (see e.g. Ghodsi et al. (2018), Aziz et al. (2017b), Amanatidis et al. (2015), Barman and Krishnamurthy (2017) and Huang and Lu (2019)).

The second solution concept that has been proposed is a relaxation of envy-freeness called *envy-free up to one item (EF1)*. An allocation satisfies EF1 if it is envy-free or any agent's envy for another agent can be removed if some item is ignored. Under additive utilities, EF1 can be achieved by a simple algorithm called the round-robin sequential allocation algorithm. Agents take turn in a round-robin manner and pick their most preferred unallocated item. The interest in EF1 was especially piqued when Caragiannis et al. (2016)

proved that for positive additive utilities, a rule based on maximizing Nash social welfare finds an allocation that is both EF1 and Pareto optimal. The result has led to very interesting followup work. For example, Barman, Krishnamurthy, and Vaish (2018) presented a pseudo-polynomial time algorithm for the same setting to compute an Pareto optimal and EF1 allocation. The returned allocation also provides a 1.45 approximation of the maximum Nash welfare.<sup>3</sup>

A majority of the results on indivisible items assume that the items are goods yielding positive utilities (see e.g. Bouveret, Chevaleyre, and Maudet; Lang and Rothe (2016; 2016)). In recent years, there has been a push to obtain similar existence and algorithmic results for negative utilities (see e.g. Aziz (2016)) or more general utility functions (Plaut and Roughgarden 2018; Oh, Procaccia, and Suksompong 2019). For example, Aziz et al. (2019a) designed an algorithm that returns in polynomial time an EF1 allocation for doubly monotonic preferences.

There are ongoing efforts to develop models and algorithms that take into account real-life features including distributional constraints, dynamic and online settings, distributed settings, and asymmetric agents. An interesting new setting that has been considered is one in which the agents are partitioned into groups and the items are allocated to groups (Segal-Halevi and Suksompong 2019). Researchers have also been exploring stronger or alternative fairness and efficiency concepts. For example, there is a renewed focus on concepts that concern envy between groups of agents (Aziz and Rey 2019; Conitzer et al. 2019; Benabbou et al. 2019). Another direction concerns scenarios where items are viewed as nodes of a graph and only those allocations are considered in which each agent gets a connected coalition (Bouveret et al. 2017).

### 3 Allocation of Divisible Items

When divisible items are allocated among agents, it is easier to guarantee the existence of fair allocations. It can also be relatively easier to compute a fair allocation. Again, consider the example of two agents and one item. If the item is divisible, we can simply allocate half of the item to each agent, guaranteeing many reasonable fairness criteria including envy-freeness.

In contrast to the case of indivisible items, an envy-free and Pareto optimal allocation always exists for the case of divisible items. For positive additive utilities, such an allocation can be computed by maximizing Nash social welfare in polynomial time. When the utilities are additive but negative, there are still important gaps in our understanding of the computational complexity of envy-free and Pareto optimal allocations (see e.g. Bogomolnaia et al. (2017)).

For divisible items, one important class of utility functions is called *Leontief preferences* in which each agent requires fixed proportions of the items (such as 1 CPU and 5 GB RAM to perform one task). Ghodsi et al. (2011) pro-

<sup>3</sup>There is a stream of papers primarily focussed on approximating Nash welfare. For example, Cole and Gkatzelis (2015) presented a polynomial-time 2.889-approximation algorithm for maximum Nash welfare.

posed a solution for the problem of allocating as much of each item as possible while equalizing the agents' use of the fraction of their own dominant (bottleneck) resource. The method satisfies envy-freeness, proportionality, Pareto optimality, and strategyproofness. These solutions have been extended in important ways such as dealing with indivisibilities and dynamic settings (Parkes, Procaccia, and Shah 2013).

Another model of divisible items is called cake-cutting (Robertson and Webb 1998; Procaccia 2016; 2013). A cake is a unit interval representing a heterogeneous item. An agent may have different utilities for different subintervals, even if they are of equal size. The challenge is to compute fair allocations using queried information from agents. Cake-cutting has a very long history dating back to the inception of the mathematical theory of fair division. Two of the most well-studied concepts are envy-freeness and proportionality (each agent gets at least  $1/n$  of her value of the whole cake). While there are well-known and simple algorithms to compute a proportional allocation, finding envy-free allocations is more challenging. In the last few years, new algorithms have been proposed that compute an envy-free allocation in a number of queries that is bounded in the number of agents. The algorithms work for either positive utilities (Aziz and Mackenzie 2016) or negative utilities (Dehghani et al. 2018). There is no known bounded algorithm for the case of mixed utilities. It has also been noted by the authors that if the requirement to allocate the whole cake is relaxed, then there are simpler and faster algorithms to compute an allocation that is both envy-free and proportional. In recent years, several new variants of the cake-cutting problems have been studied (see, e.g. Segal-Halevi, Hassidim, and Aumann (2015)).

Recently, researchers have approached problems where indivisible items are allocated among agents, but a minimal number of items are shared or treated as divisible to guarantee existence of fair allocations (Sandomirskiy and Segal-Halevi 2019). The approach can be viewed as a natural extension of the Adjusted Winner rule (see e.g. Brams and Togman (1996)) that finds an envy-free and Pareto optimal allocation among two agents in which only at most one item is shared.

Money can also be viewed as a divisible item that can be moved around to achieve fairness. A well-studied discrete allocation problem with money is the fair room-rent division problem (see e.g. Alkan, Demange, and Gale; Abdulkadiroğlu, Sönmez, and Ünver; Gal et al. (1991; 2004; 2017)).

### 4 Two-Sided Matching

In two sided matching markets (see e.g. Manlove (2013)), the goal is to match entities from two different sets. A typical example of the problem is one where students/doctors are matched to school/hospitals. Students have preferences over schools. Schools have priorities over students. An outcome in these settings is a matching of students to schools.

In two-sided matching, an important criterion of desirable outcomes is a fairness property that requires that no student should want to replace another student in a hospital when the

hospital has a higher priority for the former student over the latter student. The concept is also based on the idea of envy. The only difference is that an agent can only envy another agent if the former has higher priority.

The property of fairness along with that of non-wastefulness (no student prefers to take some empty slot in a school) are the central properties that underpin much of the work on two-sided matching. For many general two-sided matching markets, the two properties can be simultaneously satisfied by an allocation returned by the well-known Deferred Acceptance algorithm (Roth 2008) or its generalizations. For several important settings such as the allocation of school seats to students (Abdulkadiroğlu and Sönmez 2003), one version of the algorithm additionally satisfies strategyproofness.

An important research agenda extends this kind of result to other two-sided matching problems with distributional constraints. These constraints included upper quotas on regions of hospitals (see e.g. Kamada and Kojima; Biro et al.; Goto et al. (2015; 2010; 2016)) as well as target lower quotas for diverse representation (Aziz et al. 2019b; Ehlers et al. 2014; Kurata et al. 2015). These new settings give rise to fundamental questions of tradeoffs between merit, diversity, and segregation. Yet another direction is taking a time-sharing approach to two-sided matching (see e.g. Aziz and Klaus; Kesten and Unver (2019; 2015)).

## 5 Fairness in other Social Choice Settings

We have mostly focussed on allocations problems where agents get ownerships of their allocation. We can also consider allocation of public goods such as the funding of public projects. Public goods can be selected via voting on the different options. In classical voting settings in which a single alternative is to be selected, it is very difficult to guarantee reasonable representation of sizable minorities. Consider the case where 51% of the agents most prefer alternative  $a$  and 49% of the agents prefer alternative  $b$ . Then alternative  $a$  that is supported by a majority is selected even though the remaining agents may find it unpalatable.

There are other voting settings where fairness concerns become more meaningful. One such setting is a simple extension of single-winner voting to selecting multiple winners (see e.g. Aziz et al.; Faliszewski et al.; Aziz (2017a; 2017; 2019b)). Another setting is probabilistic voting or portioning in which each alternative is given a certain probability weight or a fraction of a budget (see e.g. Airiau et al.; Aziz, Bogomolnaia, and Moulin; Aziz; Brandt (2019; 2019; 2019a; 2017)). In both settings, meaningful fairness axioms have been proposed and algorithms have been designed to satisfy the axioms. A general principle while formalizing the axioms is that of proportional representation: if a large enough group cohesively prefers some outcome(s), then that group should get suitable representation (see e.g. Aziz and Lee (2019)). Recently, fairness concerns are also being examined in participatory budgeting problems in which projects have costs and there is a constraint on the maximum budget that can be used (Fain,

Goel, and Munagala 2016; Aziz, Lee, and Talmon 2018; Fluschnik et al. 2017).

Another collective decision-making setting where fairness becomes more meaningful is the case where there are several issues with each issue having its own candidates outcomes (Conitzer, Freeman, and Shah 2017). Agents have preferences over outcomes over each of the issues. The setting is more general than allocation of indivisible items in following way. Each issue corresponds to the decision of allocation of a particular item. The possible outcomes of an issue correspond to giving the item to one of the agents. It turns out that a solution based on maximizing Nash social welfare satisfies both Pareto optimality and a natural relaxation of proportionality.

## 6 Discussion

I discussed some settings where fairness plays an important role. Fairness concerns can come up in many other multi-agent settings as well such as auctions and coalition formation. Much of the work in multi-agent fair allocation is centered around formal axioms to capture fairness.<sup>4</sup> It will be interesting to use data-centric approaches to understand how humans conceptualize fairness or decide on tradeoffs between societal outcomes.

Another area of computer science where fairness is taking center stage is machine learning. Machine learning algorithms learn from data to make actionable suggestions such as giving a loan. Whether these algorithms are biased is a prominent concern for society (see, e.g. Kleinberg et al. (2018)). A natural approach to incorporate fairness in these algorithms follows the Aristotelian “equal treatment of equals” principle. It is an interesting challenge to explore ways in which ideas from multi-agent fair allocation can contribute to fairness in machine learning. In general, the interplay of fairness with other desirable aspects such as accountability, transparency, and trust is broad research question (see, e.g. Shin and Park (2019)).

As computers start to make important decisions in our lives, it is timely to revisit the lessons learnt. New computer applications will give rise to novel and challenging problems so it is critical to continue updating our toolkit of axioms and methods for designing and enabling fair multi-agent systems.

## Acknowledgements

The survey is based on a presentation given by the author at the Dagstuhl-Seminar on Application-Oriented Computational Social Choice (Seminar Number 19381). Aziz is supported by a UNSW Scientia Fellowship, and Defence Science and Technology (DST) under the project “Auctioning for distributed multi vehicle planning” (DST 9190). The author thanks Alex Lam and the reviewers for useful comments.

<sup>4</sup>In addition to standard axioms such as envy-freeness, a formal theory of fair allocation covers many different axiom for desirable rules or algorithms (see, e.g. Thomson (2016)).

## References

- Abdulkadiroğlu, A., and Sönmez, T. 2003. School choice: A mechanism design approach. *American Economic Review* 93(3):729–747.
- Abdulkadiroğlu, A.; Sönmez, T.; and Ünver, M. U. 2004. Room assignment-rent division: A market approach. *Social Choice and Welfare* 22(3):515–538.
- Airiau, S.; Aziz, H.; Caragiannis, I.; Kruger, J.; Lang, J.; and Peters, D. 2019. Portioning using ordinal preferences: Fairness and efficiency. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Alkan, A.; Demange, G.; and Gale, D. 1991. Fair allocation of indivisible goods and criteria of justice. *Econometrica* 59:1023–1039.
- Amanatidis, G.; Markakis, E.; Nikzad, A.; and Saberi, A. 2015. Approximation algorithms for computing maximin share allocations. In *Proceedings of the 35th International Colloquium on Automata, Languages, and Programming (ICALP)*, 39–51.
- Atkinson, A. 1970. On the measurement of inequality. *J. of Economic Theory* 2(3):244–263.
- Aziz, H., and Klaus, B. 2019. Random matching under priorities: Stability and no envy concepts. *Social Choice and Welfare* 53(2):213–259.
- Aziz, H., and Lee, B. 2019. The expanding approvals rule: Improving proportional representation and monotonicity. *Social Choice and Welfare*.
- Aziz, H., and Mackenzie, S. 2016. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*, 416–427.
- Aziz, H., and Rey, S. 2019. Almost group envy-free allocation of indivisible goods and chores. *CoRR* abs/1907.09279.
- Aziz, H.; Gaspers, S.; Mackenzie, S.; and Walsh, T. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017a. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48(2):461–485.
- Aziz, H.; Rauchecker, G.; Schryen, G.; and Walsh, T. 2017b. Algorithms for max-min share fair allocation of indivisible chores. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 335–341.
- Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2019a. Fair allocation of indivisible goods and chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Aziz, H.; Gaspers, S.; Sun, Z.; and Walsh, T. 2019b. From matching with diversity constraints to matching with regional quotas. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Aziz, H.; Bogomolnaia, A.; and Moulin, H. 2019. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, 753–781.
- Aziz, H.; Lee, B. E.; and Talmon, N. 2018. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, 23–31.
- Aziz, H. 2016. Computational social choice: Some current and new directions. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 4054–4057.
- Aziz, H. 2019a. A probabilistic approach to voting, allocation, matching, and coalition formation. In Laslier, J.-F.; Moulin, H.; Sanver, R.; and Zwicker, W. S., eds., *The Future of Economic Design*. Springer-Verlag. Forthcoming.
- Aziz, H. 2019b. A rule for committee selection with soft diversity constraints. *Group Decision and Negotiation* 1–8.
- Balinski, M., and Young, H. P. 1982. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press. (2nd Edition [with identical pagination], Brookings Institution Press, 2001).
- Barman, S., and Krishnamurthy, S. K. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the 18th ACM Conference on Economics and Computation (ACM-EC)*.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, 557–574.
- Benabbou, N.; Chakraborty, M.; Elkind, E.; and Zick, Y. 2019. Fairness towards groups of agents in the allocation of indivisible items. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Biro, P.; Fleiner, T.; Irving, R. W.; and Manlove, D. F. 2010. The college admissions problem with lower and common quotas. *Theoretical Computer Science* 411(34–36):3136–3153.
- Bogomolnaia, A.; Moulin, H.; Sandomirskiy, F.; and Yanovskaya, E. 2017. Competitive division of a mixed manna. *Econometrica* 85(6):1847–1871.
- Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair division of a graph. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 135–141.
- Bouveret, S.; Chevaleyre, Y.; and Maudet, N. 2016. Fair allocation of indivisible goods. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 12.
- Brams, S. J., and Togman, J. M. 1996. Camp david: Was the agreement fair? *Conflict Management and Peace Science* 15(1):99–112.
- Brams, S. J. 2008. *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton University Press.

- Brandt, F. 2017. Rolling the dice: Recent results in probabilistic social choice. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 1, 3–26.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (ACM-EC)*, 305–322.
- Cole, R., and Gkatzelis, V. 2015. Approximating the Nash social welfare with indivisible items. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, 371–380. ACM Press.
- Conitzer, V.; Freeman, R.; Shah, N.; and Vaughan, J. W. 2019. Group fairness for the allocation of indivisible goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17*, 629–646.
- Dehghani, S.; Farhadi, A.; Hajiaghayi, M. T.; and Yami, H. 2018. Envy-free chore division for an arbitrary number of agents. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, 2564–2583.
- Ehlers, L.; Hafalir, I. E.; Yenmez, M. B.; and Yildirim, M. A. 2014. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory* 153:648–683.
- Fain, B.; Goel, A.; and Munagala, K. 2016. The core of the participatory budgeting problem. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*, 384–399.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. In Endriss, U., ed., *Trends in Computational Social Choice*. chapter 2.
- Fluschnik, T.; Skowron, P.; Triphaus, M.; and Wilker, K. 2017. Fair knapsack. *CoRR* abs/1711.04520.
- Gal, Y.; Mash, M.; Procaccia, A. D.; and Zick, Y. 2017. Which is the fairest (rent division) of them all? *J. ACM* 64(6):39:1–39:22.
- Ghods, A.; Zaharia, M.; Hindman, B.; Konwinski, A.; Shenker, S.; and Stoica, I. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation*.
- Ghods, M.; HajiAghayi, M.; Seddighin, M.; Seddighin, S.; and Yami, H. 2018. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*. ACM Press.
- Goto, M.; Iwasaki, A.; Kawasaki, Y.; Kurata, R.; Yasuda, Y.; and Yokoo, M. 2016. Strategyproof matching with regional minimum and maximum quotas. *Artificial intelligence* 235:40–57.
- Huang, X., and Lu, P. 2019. An algorithmic framework for approximating maximin share allocation of chores. *CoRR* abs/1907.04505.
- Kamada, Y., and Kojima, F. 2015. Efficient matching under distributional constraints: Theory and applications. *American Economic Review* 105(1):67–99.
- Kesten, O., and Unver, U. 2015. A theory of school choice lotteries. *Theoretical Economics* 543–595.
- Kleinberg, J.; Ludwig, J.; Mullainathan, S.; and Rambachan, A. 2018. Algorithmic fairness. *American Economic Association Papers and Proceedings* 108:22–27.
- Kleinberg, J. M.; Rabani, Y.; and Tardos, É. 2001. Fairness in routing and load balancing. *J. Comput. Syst. Sci.* 63(1):2–20.
- Kurata, R.; Hamada, N.; Iwasaki, A.; and Yokoo, M. 2015. Controlled school choice with soft bounds and overlapping types. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 951–957.
- Lang, J., and Rothe, J. 2016. Fair division of indivisible goods. In Rothe, J., ed., *Economics and Computation*. 493–550.
- Manlove, D. F. 2013. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company.
- Moulin, H. 2003. *Fair Division and Collective Welfare*. The MIT Press.
- Moulin, H. 2019. Fair division in the age of internet. *Annual Review of Economics* 11:407–441.
- Nguyen, N.-T.; Nguyen, T.-T.; Roos, M.; and Rothe, J. 2014. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous Agents and Multi-Agent Systems* 28(2):256–289.
- Oh, H.; Procaccia, A. D.; and Suksompong, W. 2019. Fairly allocating many goods with few queries. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2141–2148.
- Parkes, D. C.; Procaccia, A. D.; and Shah, N. 2013. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. In *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM-EC)*, 808–825.
- Plaut, B., and Roughgarden, T. 2018. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- Procaccia, A. D., and Wang, J. 2014. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*, 675–692. ACM Press.
- Procaccia, A. D. 2013. Cake cutting: Not just child’s play. *Communications of the ACM* 56(7):78–87.
- Procaccia, A. D. 2016. Cake cutting algorithms. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia,

- cia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 13.
- Rawls, J. 1971. *A Theory of Justice*. Harvard University Press.
- Rescher, N. 2002. *Fairness: Theory and Practice of Distributive Justice*. Transaction Publishers.
- Robertson, J. M., and Webb, W. A. 1998. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters.
- Roth, A. E. 2008. Deferred acceptance algorithms: history, theory, practice, and open questions. *International Journal of Game Theory* 36:537—569.
- Sandomirskiy, F., and Segal-Halevi, E. 2019. Fair division with minimal sharing. *CoRR* abs/1908.01669.
- Segal-Halevi, E., and Suksompong, W. 2019. Democratic fair allocation of indivisible goods. *Artificial Intelligence* 277:103–167.
- Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2015. Envy-free cake-cutting in two dimensions. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 1021–1028.
- Shin, D., and Park, Y. J. 2019. Role of fairness, accountability, and transparency in algorithmic affordance. *Computers in Human Behavior* 98:277–284.
- Thomson, W. 2011. Chapter twenty-one - fair allocation rules. In Arrow, K. J.; Sen, A.; and Suzumura, K., eds., *Handbook of Social Choice and Welfare*, volume 2 of *Handbook of Social Choice and Welfare*. Elsevier. 393 – 506.
- Thomson, W. 2016. Introduction to the theory of fair allocation. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 11.
- Varian, H. R. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9:63–91.