

Strategic Behaviour when Allocating Indivisible Goods

Toby Walsh

UNSW and Data61 (formerly NICTA)
Sydney, Australia
toby.walsh@nicta.com.au

Abstract

We survey some recent research regarding strategic behaviour in resource allocation problems, focusing on the fair division of indivisible goods. We consider a number of computational questions like how a single strategic agent misreports their preferences to ensure a particular outcome, and how agents compute a Nash equilibrium when they all act strategically. We also identify a number of future directions like dealing with non-additive utilities, and partial or probabilistic information about the preferences of other agents.

Introduction

Resource allocation is a perennial problem facing society. Economical, environmental and political pressures are forcing us to try to do more with fewer resources and to do so more fairly. One important and challenging case is the fair division of indivisible goods. This captures a wide range of problems including allocating workers to shifts, houses to people, classes to students, landing slots to airlines, players to teams, leads to salespeople, and time slots on expensive scientific instruments to scientists.

Unfortunately, many mechanisms for allocating indivisible goods are manipulable in theory, and are manipulated in practice (see, for example, (Budish and Cantillon 2012)). Such manipulation can cause significant welfare loss as well as unfairness in the outcome. Recently, attention has turned to how precisely agents might *compute* such strategic behaviour. What if it is just too computationally difficult to compute a manipulation (Bartholdi, Tovey, and Trick 1989; Bartholdi and Orlin 1991)? If manipulation is hard to compute, perhaps agents might simply behave sincerely? Manipulation has been shown to be computationally hard to compute in many voting situations, e.g. (Davies et al. 2011; Narodytska, Walsh, and Xia 2011; Davies, Narodytska, and Walsh 2012; Narodytska, Walsh, and Xia 2012; Davies et al. 2014)). Also, it may not be just the agents being allocated the goods that can behave strategically. As with voting (Bartholdi, Tovey, and Trick 1992), can the chair also manipulate the outcome? For example, what if they can remove or add items? Or adjust the allocation mechanism? If so, how do they compute the best strategic behaviour?

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

With voting, reasoning about uncertainty is closely related to manipulation (Konczak and Lang 2005; Walsh 2007; Pini et al. 2008; Walsh 2008). The same is true for allocation. For instance, if an outcome is possible given uncertainty in an agent's preferences over the items, then this agent can manipulate the allocation to ensure the given outcome. Questions about possible or necessary outcomes thus relate to questions about manipulation. How do we efficiently compute what items an agent can possibly be allocated? What items does an agent necessarily receive? In this paper, I survey recent work in this area and identify some interesting open challenges that remain. We focus on resource allocation where money is not exchanged and fairness properties are important. If money is involved and efficiency is of concern, there is a large and relevant literature on auctions.

An example

Suppose you are named as one of two captains and are picking a football team. You and the other captain take turns to pick players. This is a fair division problem over indivisible goods. Now suppose your first choice is Bob, the top goal scorer all year. But you know the other captain has fallen out with Bob, and so he'll choose Carol first as she is the best goal keeper by far. You win the toss so get to make the first pick. You strategically pick Carol, even though she isn't your first choice. You know that the other captain will leave Bob sitting on the bench so you can still pick him in a later round. The net result is that you get a stronger team, both the best goal scorer and the best goal keeper. If you had been forced to pick players sincerely, you would have only got the best goal scorer. Of course, the other captain may also behave strategically, not picking players you will leave on the bench till the end of the allocation. But how do you (and the other captain) *compute* your optimal strategy? By the end of reading this paper, you will know a simple, linear time method to do so.

Formal background

We suppose that there are n agents being allocated m items. For simplicity, we suppose m is an integer multiple of n and add dummy items of no utility to achieve this. An allocation is an assignment of items to agents. A special case is house allocation, when $n = m$ and each agent gets exactly one

item. As in much of the literature, we often suppose agents have additive utilities over the items. One important challenge is to relax this assumption and consider non-additive utilities. For instance, in many settings, we have complementarities (e.g. a left shoe is only valuable to us if we are also allocated a right shoe) and substitutabilities (e.g. we either wish to be allocated the car or the bicycle but not both). Such fair division involves several subchallenges including efficiently eliciting such complex preferences, and adapting mechanisms to take account of this additional complexity.

Challenge 1. *Develop allocation mechanisms for the fair division of indivisible goods that efficiently and effectively take account of complementarities and substitutabilities.*

When allocating goods, efficiency is one of several desirable properties. An allocation *Pareto improves* another iff each agent has at least the same utility in the first, and there is at least one agent where the utility is greater. An allocation is *Pareto efficient* iff there is no allocation which Pareto improves it. Fairness notions like envy freeness are also important. An allocation is *envy free* if no agent strictly has greater utility for another agent's allocation. An allocation is *proportional* if each of the n agents assigns at least $\frac{1}{n}$ th of their total utility to their allocation. It is also important to understand the role that strategic behaviour can play. A mechanism is *manipulable* if an agent can misreport their preferences and improve their utility. A mechanism that is not manipulable is *strategy-proof*. Finally, a mechanism is *nonbossy* if an agent cannot change the allocation without changing their own allocation.

All the mechanisms we consider are ordinal. That is, they only require agents to declare ordering over items. Ordinal mechanisms have a number of advantages over cardinal mechanisms that require agents to declare their actual utilities. These include simplicity, low communication complexity, and a more limited action space for behaving strategically. Many of the mechanisms we consider are also non-deterministic, returning a probability distribution over allocations. Non-determinism allows us to break ties fairly. If two agents both want an item, we can toss a coin to decide which one has it. We therefore consider both *ex post* and *ex ante* properties. An *ex post* outcome is any allocation an agent receives with non-zero probability. By comparison, an *ex ante* outcome is what an agent receives in expectation. For example, a mechanism is *proportional ex ante* if every agent assigns at least $\frac{1}{n}$ th of their total utility to their allocation in expectation.

We also consider other orderings over outcomes that work with ordinal preferences and lotteries besides the ordering that comes from comparing (expected) utilities. For instance, the *SD* (stochastic dominance) ordering prefers an allocation p to an agent over q if the probability for the agent to get the top i items in p is at least as large as in q for all $i \in [1, m]$. If an allocation to an agent is *SD-preferred* over another then it has greater or equal expected utility for all utilities consistent with the agent's ordinal preferences. Notions like *SD-efficiency*, *SD-envy freeness* and *SD-strategy proofness* can be defined from the *SD*-preference ordering.

Serial dictatorship

We first consider serial dictatorship mechanisms like random serial dictator (RSD). This is also known as random priority. RSD randomly orders the agents, and agents then take turns to pick their allocation of items all at once. RSD has been widely used in school choice, university housing and many other settings. Serial dictatorships can also use other criteria to order the agents (for example, seniority or exam results). Serial dictatorships are some of the few strategy proof mechanisms available. Indeed, any strategy proof, nonbossy and neutral¹ mechanism is necessarily a serial dictatorship (Svensson 1999).

Unfortunately the strategy proofness of serial dictatorship comes at a price in welfare, fairness and computational cost. The price in welfare is that, whilst the allocation returned by a serial dictatorship is *ex post* efficient, it may not be *ex ante* efficient. There is also a price in terms of fairness. Serial dictatorships will often return allocations in which one agent (e.g. an agent at the end of the priority order) envies another agent (e.g. an agent at the start of the priority order). There is also a computational price. For instance, whilst it is easy to execute the RSD mechanism, it has recently been shown that it is intractable to compute the actual probability distribution over outcomes.

Theorem 1. (Aziz et al. 2014; Sabán and Sethuraman 2013) *Computing the probability that an agent gets an item with RSD is #P-complete.*

On the other hand, computing the RSD probabilities is efficient if we have a bounded number of items, or agent types (Aziz and Mestre 2014).

This analysis of serial dictatorship mechanisms suggests a second challenge.

Challenge 2. *Identify or propose new allocation mechanisms which are computationally difficult to manipulate with good welfare and fairness properties.*

Sequential allocation

Another popular class of mechanisms are those based on *sequential allocation* (Brams and Taylor 1996). In a sequential allocation mechanism, agents simply take turns to pick items. This leaves open the particular order used to take turns (the so called "policy"). For example, with the balanced alternation policy 123321, agent 1 picks first, then agent 2, then agent 3 before we repeat in reverse. There are real world settings like course allocation at the Harvard Business School where the policy is chosen at random from a space of balanced alternating policies as a means of ensuring (procedural) fairness. Whilst sequential allocation is not strategy proof (see the example in the introduction to this paper), it has several other desirable properties. For example, it is easy to execute, requires limited preference information,

¹A mechanism is *neutral* if permuting the names merely permutes the outcome (that is, the outcome does not depend on the names of the indivisible items).

and returns Pareto efficient allocations supposing agents act sincerely.

We first consider the case where one agent or a coalitional of agents try to manipulate the outcome by acting strategically to ensure the best possible outcome, whilst the other agents act sincerely. This is equivalent to the best response problem in Nash dynamics. Bouveret and Lang (2014) prove that a single agent can compute an optimal manipulation in polynomial time, as can a coalition if we allow transfers and side payments. However, if we prevent transfers and/or side payments, then a simple reduction from the PARTITION problem proves that computing a successful manipulation is NP-hard even with just 2 manipulators (Bouveret and Lang 2014).

We next consider the case when all agents act strategically. The sequential allocation procedure naturally lends itself to a game theoretic analysis where we look for a Nash equilibrium where no agent can improve their allocation by deviating unilaterally from their (perhaps insincere) picking strategy. We view the sequential allocation mechanism as a finite repeated sequential game in which all agents have complete information about the preference ordering of the other agents. We could use backward induction to find the subgame perfect Nash equilibrium but this would take exponential time in general. When agents have the same preference ordering, Proposition 6 in (Bouveret and Lang 2011) proves that the subgame perfect Nash equilibrium is sincere picking. On the other hand, when preference orderings are different (as in the example in the introduction to this paper), there exist equilibria where behaviour is not sincere.

With two agents, additive utilities, and the strictly alternating policy, Kohler and Chandraesekaran (1971) prove that the subgame perfect Nash equilibrium can be computed in linear time by simply reversing the policy and preference orderings. This subgame perfect Nash equilibrium is unique provided no agent has the same utility for any pair of items. Surprisingly, this method to compute the subgame perfect Nash equilibrium extends to any policy and not just policies in which agents strictly alternate to pick. The basic intuition is that, by acting strategically, the agent who doesn't pick last can ensure *optimally* that the agent who does pick last gets their least favourite item (which they will certainly leave till last to pick). Similarly, if we delete this item, then the agent who doesn't pick second from last can act strategically to ensure *optimally* that the agent who does pick second from last gets their least favourite remaining item (which they will also certainly leave till then to pick), and so on for earlier picks.

Theorem 2. (Kalinowski et al. 2013) *With two agents and additive utilities, we can compute the unique subgame perfect Nash equilibrium in linear time by reversing the preferences, swapping the agents, and running the policy in reverse.*

As promised at the start of this paper, you now know how to compute your optimal strategic behaviour when captain selecting a team. This leaves open the following interesting challenge.

Challenge 3. *Identify how to compute the subgame perfect Nash equilibria for sequential allocation mechanisms with 2 agents when utilities are not additive.*

With three agents, there may no longer be an unique subgame perfect Nash equilibrium. Indeed, there exists a class of allocation problems with additive Borda utilities in which the number of subgame perfect Nash equilibria grows exponentially. Little more is known about the subgame perfect Nash equilibria with 3 or more agents.

Challenge 4. *Characterise the subgame perfect Nash equilibria for sequential allocation mechanisms with 3 or more agents.*

With an unbounded number of agents, the complexity of the problem of computing a subgame perfect Nash equilibrium remained open for a long time. Brams and Straffin (1979) remarked that

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

Recently, we proved that the problem is indeed computationally intractable.

Theorem 3. (Kalinowski et al. 2013) *With an unbounded number of agents and additive Borda utilities, computing a subgame perfect Nash equilibrium is PSPACE-hard.*

As noted before, strategic behaviour is only worthwhile when agents have different utilities. We say that an item is multi-valued if two agents assign it different utilities. There is a polynomial algorithm to computing a subgame perfect Nash equilibrium when the number of multi-valued items is held constant.

Theorem 4. (Kalinowski et al. 2013) *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*

An interesting open problem remains which is the computational complexity of computing the subgame perfect Nash equilibria with a bounded number of agents and no bound on the number of multi-valued items.

Challenge 5. *Determine the computational complexity of computing the subgame perfect Nash equilibria for sequential allocation mechanisms with a bound on the number of agents but not on the number of multi-valued items.*

Note that these results have so far assumed the manipulating agents have perfect information about the preferences of the other agents. This may be unrealistic in practice. Of course, any of our complexity results provide lower bounds on the complexity in the presence of incomplete information. However, more generally, we might want to consider strategic behaviour when agents have only probabilistic information.

Challenge 6. *Study the Markov perfect Nash equilibria of sequential allocation mechanisms.*

Probabilistic serial mechanism

We turn now to a more recent but increasingly popular non-deterministic mechanism, the probabilistic serial mechanism (Bogomolnaia and Moulin 2001). With this mechanism, agents simultaneously “eat” their most preferred items at a uniform speed, moving onto their next most preferred item whenever an item is fully eaten. This gives a randomized or fractional assignment which can easily be realized as a probability distribution over discrete allocations. Unfortunately, the probabilistic serial mechanism is not strategy proof (Bogomolnaia and Moulin 2001). However, it has good welfare and efficiency properties. It is, for instance, SD-efficient and SD-envyfree.

What do we know about computing strategic behaviour with the probabilistic serial mechanism? First, we have shown that computing an expected utility better response is NP-hard (Aziz et al. 2015b). We also identified two tractable cases: computing a best response for lexicographical utilities, and a best response with just two agents. It follows therefore that computing a manipulation that optimizes the expected utility is also NP-hard in general. Second, Nash deviations under the probabilistic serial mechanism can cycle. Despite the possibilities of cycles, we proved that a pure Nash equilibrium is guaranteed to exist. Unfortunately computing this is intractable in general.

Theorem 5. (Aziz et al. 2015a) *Verifying if a profile of preferences is a pure Nash equilibrium for the probabilistic serial mechanism is coNP-complete, whilst computing a pure Nash equilibrium is NP-hard.*

In this case computational complexity is a good thing. It may act as a shield against agents manipulating the mechanism. On the other hand, manipulation is easy with two agents. A pure Nash equilibrium can be computed in linear time (Aziz et al. 2015a). This linear time algorithm exploits a beautiful mapping between the probabilistic serial mechanism and the sequential allocation mechanism where two agents simply alternate turns to pick items. If we divide items into two halves, and replicate the agents’ preferences, then the allocation returned by the probabilistic serial mechanism is ‘essentially’ the same as the allocation obtained by applying sequential allocation to this half-item problem. An interesting open challenge is to generalize this mapping to more than two agents. With 2 agents, it was easy to see that we only needed “half” items. It is much less obvious what fractional division of items is required for the m agent problem.

Challenge 7. *Generalize the mapping between the probabilistic serial mechanism and the sequential allocation mechanism for 3 or more agents.*

This mapping may also help us solve another challenge which is to understand strategic behaviour where we stop viewing this as a one off game. It is perhaps more realistic to view this as a *repeated* game, with a new subgame each time an agent starts eating a new item. We might also want to consider games in which agents only have partial or probabilistic information about the preferences of other agents.

Challenge 8. *Understand strategic behaviour in the repeated version of the probabilistic serial mechanism game, as well as in the game with incomplete or probabilistic information.*

Scoring rule mechanisms

Nguyen, Baumeister and Rothe (2015) have recently studied the strategy proofness of another type of allocation mechanism. In scoring rule mechanisms (Baumeister et al. 2014), agents specify ordinal preferences over individual items. Ordinal rankings are converted into utilities for the individual items using simple scoring vectors like Borda or lexicographical scores. These are lifted to sets of items using standard extension principles from social choice theory like the Kelly or Gärdenfors extensions. Such mechanisms often have a number of good properties like monotonicity and consistency. Nguyen, Baumeister and Rothe (2015) characterized the strategy-proofness of the scoring rule allocation mechanism that maximizes utilitarian welfare. They prove that such mechanisms are strategy proof in only rather limited circumstances.

Theorem 6. (Nguyen, Baumeister, and Rothe 2015) *The utilitarian scoring rule mechanism is strategy proof (according to the Kelly or Gärdenfors set extension) if and only if the scoring rule vector has at most two different values or the largest value occurs more than $\frac{m}{2}$ times.*

As they note, strategy proofness occurs when the scoring rule mechanism is rather oblivious to the preferences of the agents. In situations where the mechanism is more responsive to the preferences of the agents, manipulation is likely to be possible. This suggests the following research direction.

Challenge 9. *Study strategic behaviour in scoring rule mechanisms.*

Item partition mechanisms

Inspired by our work on online fair division (Walsh 2011a; 2014; 2015; Aleksandrov et al. 2015), we have recently identified a novel class of strategy proof mechanisms called item partition mechanisms. These mechanisms are in some sense half way between sequential allocation and serial dictatorship mechanisms. As in both sequential allocation and serial dictatorship mechanisms, agents in an item partition mechanism take turns to pick items. Like serial dictatorship mechanisms, item partition mechanisms are strategy proof but, unlike serial dictatorship mechanisms (and like sequential allocation mechanisms), item partition mechanisms do not require agents to pick all their items at once. To be more precise, in an *item partition* mechanism, items are partitioned into sets of n or fewer items, and a sequential allocation mechanism is then used to allocate the items within each set. For instance, in the *random item partition* (RIP) mechanism, we partition items at random into sets of n items, and then use RSD to allocate each set of items. Of course, other criteria can also be used to partition items and to order agents within round of sequential allocation.

Theorem 7. *Item partition mechanisms are strategy proof. The RIP mechanism is both strategy proof and proportional ex ante. With 2 agents, the RIP mechanism is also envy free ex ante.*

The strategyproofness of the RIP mechanism somewhat contradicts the claim in (Budish and Cantillon 2012) that

“dictatorships are the only strategyproof mechanisms within the class of random priority mechanisms (i.e., mechanisms where agents take turns choosing objects in some random order)” (page 2204).

The RIP mechanism is a type of priority mechanism that is strategy proof but not dictatorial. Interestingly, Budish and Cantillon (2012) argue the Harvard Business School mechanism mentioned earlier is strategy proof on a domain restriction, “block correlated” preferences. This translates to preferences in which the items partition into blocks, and irrespective of the priority order, items from a single block are chosen in each round. It remains an interesting open problem to discover other item partition mechanisms with good properties.

Challenge 10. *Identify other item partition mechanisms with good fairness and efficiency properties.*

Control of fair division

It is not only the agents being allocated items who can behave strategically. The chair allocating items can also try to influence the outcome. In fact, they may be uniquely placed to do so as they may be the only agent who naturally has complete information. For instance, in allocating courses at Harvard Business School, the agents all submit their ordinal preferences to the chair. The chair is then supposed to decide on a policy at random from the space of all possible balanced alternation policies. However, based on the submitted ordinal preferences, the chair could also decide on a policy that achieves a particular outcome. This could be benevolent (e.g. to improve welfare) or malevolent (e.g. to favor a particular agent).

Recently, we have begun to consider the computational questions of choosing policies to achieve particular outcomes. For instance, if any policy is permitted, it is trivially easy to ensure an agent gets a particular item or subset of items. We just give this agent every pick until they are happy. However, if we are limited to balanced alternation policies, it is not so easy to find a policy that ensures a given outcome. Every time we give one agent a pick, we must give every other agent also a pick. Indeed it is now intractable to decide if one of the (possible super-exponential number of) policies achieves a particular outcome.

Theorem 8. *(Aziz, Walsh, and Xia 2015) Finding a balanced alternation policy which ensures an agent gets a particular item or subset of items is NP-hard.*

Such problems are closely related to questions about what is possible or necessary when the policy is chosen at random. For instance, the problem of deciding if an agent gets an item with non-zero probability when a policy is chosen at random is equivalent to the problem of whether the chair can choose

a policy to ensure an agent gets this item. We may, of course, be interested in properties of the outcome rather than specific outcomes. This suggests another open problem.

Challenge 11. *Study the problem of choosing a policy for sequential allocation to ensure good fairness or welfare properties.*

Choosing the policy is not the only way that the chair can influence the outcome. We may also be interested in other control actions that the chair can perform similar to those considered in voting (Bartholdi, Tovey, and Trick 1992; Faliszewski, Hemaspaandra, and Hemaspaandra 2009). For example, the chair might add one or more items to improve the efficiency of the outcome. As a second example, the chair might consider deleting one or more items to improve the fairness of the outcome.

Challenge 12. *Consider other actions that the chair can perform to control the outcome of a resource allocation mechanism (e.g. adding/deleting/replacing items, adding/deleting/replacing agents, or partitioning items/agents).*

Conclusions

There are many interesting computational questions surrounding strategic behaviour in resource allocation problems. We have surveyed some recent work in this area. For instance, how does an agent (mis)report their preferences to ensure a particular outcome? How do we compute the (subgame perfect) Nash equilibrium of a particular mechanism? And how does the chair strategically choose a policy to ensure an agent gets a particular subset of items? However, many interesting questions remain to be answered. In particular, how do we deal with more complex preferences (e.g. conditional, combinatorial, or non-additive), partial or probabilistic information about the preferences of other agents, mechanisms that involve repeated games, and other types of control. Finally, many results discussed so far have been worst case. It would be interesting to consider the complexity in practice, as well as issues like phase transition behaviour (Gent and Walsh 1995; Walsh 2009; 2010; 2011b)

Acknowledgements

Data61 (formerly NICTA) is supported by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. The author also acknowledges support from the Federal Ministry for Education and Research through the Alexander von Humboldt Foundation, as well as the Asian Office of Aerospace Research and Development.

References

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., and Mestre, J. 2014. Parametrized algorithms for random serial dictatorship. *Mathematical Social Sciences* 72:1–6.
- Aziz, H.; Brandt, F.; Brill, M.; and Mestre, J. 2014. Computational aspects of random serial dictatorship. *SIGecom Exchanges* 13(2):26–30.
- Aziz, H.; Gaspers, S.; Mackenzie, S.; Mattei, N.; Narodytska, N.; and Walsh, T. 2015a. Equilibria under the probabilistic serial rule. In Yang, Q., and Wooldridge, M., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, (IJCAI 2015)*, 1105–1112.
- Aziz, H.; Gaspers, S.; Mackenzie, S.; Mattei, N.; Narodytska, N.; and Walsh, T. 2015b. 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. ACM.
- 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)*.
- Bartholdi, J., and Orlin, J. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8(4):341–354.
- Bartholdi, J.; Tovey, C.; and Trick, M. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6(3):227–241.
- Bartholdi, J.; Tovey, C.; and Trick, M. 1992. How hard is it to control an election. *Mathematical and Computer Modeling* 16(8-9):27–40.
- Baumeister, D.; Bouveret, S.; Lang, J.; Nguyen, T.; Nguyen, N.; and Rothe, J. 2014. Scoring rules for the allocation of indivisible goods. In Schaub, T.; Friedrich, G.; and O’Sullivan, B., eds., *Proceedings of 21st European Conference on Artificial Intelligence (ECAI 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, 75–80. IOS Press.
- 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.
- Bouveret, S., and Lang, J. 2014. Manipulating picking sequences. In Schaub, T.; Friedrich, G.; and O’Sullivan, B., eds., *Proceedings of 21st European Conference on Artificial Intelligence (ECAI 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, 141–146. IOS Press.
- Brams, S., and Straffin, P. 1979. Prisoners’ dilemma and professional sports drafts. *The American Mathematical Monthly* 86(2):pp. 80–88.
- Brams, S., and Taylor, A. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge: Cambridge University Press.
- Budish, E., and Cantillon, E. 2012. The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. *American Economic Review* 102(5):2237–71.
- Davies, J.; Katsirelos, G.; Narodytska, N.; and Walsh, T. 2011. Complexity of and algorithms for Borda manipulation. In Burgard, W., and Roth, D., eds., *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*. AAAI Press.
- Davies, J.; Katsirelos, G.; Narodytska, N.; Walsh, T.; and Xia, L. 2014. Complexity of and algorithms for the manipulation of Borda, Nanson’s and Baldwin’s voting rules. *Artificial Intelligence* 217(0):20 – 42.
- Davies, J.; Narodytska, N.; and Walsh, T. 2012. Eliminating the weakest link: Making manipulation intractable? In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*. AAAI Press.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2009. Multimode control attacks on elections. In Boutilier, C., ed., *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-2009)*, 128–133.
- Gent, I., and Walsh, T. 1995. Phase transitions from real computational problems. In *Proceedings of the 8th International Symposium on Artificial Intelligence*, 356–364.
- 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.
- Kohler, D., and Chandrasekaran, R. 1971. A Class of Sequential Games. *Operations Research* 19:270–277.
- Konczak, K., and Lang, J. 2005. Voting procedures with incomplete preferences. In *Proceedings of the IJCAI-2005 workshop on Advances in Preference Handling*.
- Narodytska, N.; Walsh, T.; and Xia, L. 2011. Manipulation of Nanson’s and Baldwin’s rules. In Burgard, W., and Roth, D., eds., *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*. AAAI Press.
- Narodytska, N.; Walsh, T.; and Xia, L. 2012. Combining voting rules together. In de Raedt, L., ed., *Proc. of the 20th European Conference on Artificial Intelligence (ECAI-2012)*, *Frontiers in Artificial Intelligence and Applications*. IOS Press.
- Nguyen, N.; Baumeister, D.; and Rothe, J. 2015. Strategy-proofness of scoring allocation correspondences for indivisible goods. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*. IJCAI.
- Pini, M.; Rossi, F.; Venable, K.; and Walsh, T. 2008. Dealing with incomplete agents’ preferences and an uncertain agenda in group decision making via sequential majority voting. In Brewka, G., and Lang, J., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference (KR 2008)*, 571–578. AAAI Press.
- Sabán, D., and Sethuraman, J. 2013. The complexity of computing the random priority allocation matrix. In Chen, Y., and Immorlica, N., eds., *9th International Conference on*

Web and Internet Economics, WINE 2013, volume 8289 of *Lecture Notes in Computer Science*. Springer.

Svensson, L.-G. 1999. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare* 16(4):557–567.

Walsh, T. 2007. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd National Conference on AI*. Association for Advancement of Artificial Intelligence.

Walsh, T. 2008. Complexity of terminating preference elicitation. In Padgham, L.; Parkes, D. C.; Müller, J. P.; and Parsons, S., eds., *7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, 967–974. IFAAMAS.

Walsh, T. 2009. Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In *Proceedings of 21st IJCAI*, 324–329. International Joint Conference on Artificial Intelligence.

Walsh, T. 2010. An empirical study of the manipulability of single transferable voting. In Coelho, H.; Studer, R.; and Wooldridge, M., eds., *Proc. of the 19th European Conference on Artificial Intelligence (ECAI-2010)*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, 257–262. IOS Press.

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

Walsh, T. 2011b. Where are the hard manipulation problems? *Journal of Artificial Intelligence Research* 42:1–39.

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

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.