Algorithmically Driven Shared Ownership Economies

Vincent Conitzer and Rupert Freeman

Duke University, Durham, NC 27708, USA

Abstract. Resource sharing is a natural, and increasingly common, way to efficiently utilize resources. In many situations today, however, resources are owned by a single agent who has sole control over the usage of the resource. In this article, we examine the feasibility of an alternative model where ownership of resources is shared, and usage schedules are determined algorithmically by a fixed set of rules. We explore several design parameters for shared ownership algorithms, surveying existing work in each area and proposing directions for future research.

1 Introduction

While the broader concept of the sharing economy has had significant impact, truly joint ownership of scarce resources is fraught with difficulties. For it to have a chance of succeeding, there generally need to be clear rules. Shareholders need to have clear voting rights. Some resources may be owned by the state, and therefore in a sense are collectively owned; here, too, ideally there are clear rules for how public decisions concerning these resources are made. Going beyond ownership in the narrow sense, joint custody arrangements often involve carefully stated rules. In all these cases, disputes may still arise as ambiguities in the rules are discovered, and settling these disputes can be very costly.

The costs of drawing up precise rules and adjudicating disputes make it simply not worthwhile to pursue joint ownership arrangements in many contexts. While in principle, jointly owning a car with one's neighbor may seem like a good idea in a city where one rarely needs a car, the risk of conflict because of disagreement about the rules is likely to make one think twice. The more typical outcome will be for one of the neighbors to own the car and the other to perhaps sometimes borrow it (or rent it, or borrow it with a vague expectation of providing something in return at some point), for both to own a car, or for neither to own a car and both to rely on taxis, Uber/Lyft, or other modes of transportation.

However, as devices become increasingly connected and run by algorithms—i.e., as the Internet of Things takes off—shared ownership may become a more realistic alternative. The reason is that the rules of shared ownership can now be formally encoded in an algorithmic manner. When someone attempts to use the jointly owned car, the algorithm can simply check the rules to see whether that person is currently allowed to drive the car, and if the answer is "no" refuse him permission to start the car. Such rules can take an endless variety of forms, ranging from first-come first-served, to sophisticated systems based on virtual currencies for keeping track of how much of her allotment an agent has already spent. The best precise rules presumably depend on the application and the people involved; for example, the rules should be *understandable* to the people involved.

There is at least one environment where this vision is already realistic, and that is in the context of allocating shared computing resources. This allocation is naturally controlled by algorithms, making it an ideal test case for the broader vision we have laid out above. As computation becomes increasingly ubiquitous in our environment, we expect these issues to become increasingly pervasive.

2 Allocating over time

One of the most common motivations for shared ownership is that the relevant agents do not need the resource in question at every point in time, and the times at which one agent needs the resource differ from those of another. For example: today I really would like to have the car to go to a birthday party, but I do not expect to have any need for it tomorrow. A similar situation can occur for tasks: I have an important exam tomorrow so I would very much appreciate it if you could do the dishes today, and after the exam I will be able to vacuum the house. Clearly we already make such arrangements informally. But in the allocation of computing resources, we already see the use of formal protocols.

In many settings, computing resources are allocated on-the-fly in a truly online manner, such as an operating system allocating resources among different applications as demand arises. (Here, "online" is used not in the sense of being connected to a computer network, but rather in the sense of making these decisions dynamically as new information arrives, rather than solving a large scheduling problem up front from whose solution we do not expect to deviate later.) Algorithmic schedulers receive demands from different applications, and allocate resources accordingly [Lipton and Tomkins, 1994, Kolen et al., 2007, Davis and Burns, 2011]. The demands may in turn be dependent on higher-level, human user behavior, which is difficult for the applications to accurately predict, prohibiting advance planning. In other shared ownership situations, however, we may want some flexibility to plan ahead. For instance, people with shared ownership of a car may want to schedule an hour every week to go grocery shopping, or book the car in advance to get to a particular appointment. For joint ownership of a vacation home, we may want the capacity to book the home months, or even years, in advance. In fact, this may also happen in the allocation of computing resources. Consider, for example, a researcher who is expecting a trove of data to come in from a spacecraft one week from now as it passes by a planet's moon; she may wish to reserve time on a supercomputer in advance for analyzing the data.

This consideration alone provides a rich design space. At one extreme, we can require users to express their preferences for a unit of time far in advance.

This has the obvious benefit of allowing us to optimize the schedule for whatever objective we want (since at any point in time we have preference data far into the future), at the cost of requiring users to accurately predict their demand for the item in the future. At the other extreme, we can ask that users express their demand for only the unit of time immediately following the current time period. For example, users would declare at 1pm whether or not they want the car from 2-3pm, with the car only then allocated to someone.

Both extremes are obviously impractical. We want to both allow agents to book usage in advance, and accommodate demand that appears only at the last minute. But it is not clear how to optimally design such a protocol. Currently, it is common for similar schemes (such as existing car-sharing services) to allow users to book usage for any unreserved time slot in a first-come first-served style. However, we would not want a routine, easily rescheduled shopping trip that was booked in advance to take priority over a medical emergency, say, even though the latter is impossible to predict in advance. Further, we do not want a user to *fake* a medical emergency to obtain use of the car, when in fact she simply forgot to book the car for a scheduled meeting with friends. Such strategic considerations cannot be ignored in the context of systems used by multiple self-interested agents – indeed, they feature prominently in the literature on scheduling [Tian, 2013, Im and Kulkarni, 2016, Chawla et al., 2017, Babaioff et al., 2017]. These agents are focused on their own objectives, and generally cannot be expected to honestly reveal the importance or urgency of their request for resources, if doing so is not in their own interest. One possibility for addressing this is to implement some sort of pricing scheme, allowing usage to be prioritized along socially desirable axes such as importance, advance booking, demand, and any other factors considered important by the relevant stakeholders.

3 Money or No Money?

Money is a natural tool for allocating scarce resources. Users with greater need for the item are willing to pay more for its use, suggesting that the item is more likely to end up with someone who needs it more. For example, a simple scheme would run an auction that closes some time (say, a day) before the designated use period for the item. The agent that wins the auction then has the right to use the item in the specified slot. In the case that some emergency arises between the close of the auction and the designated usage period, some high 'override' price could be fixed in advance, that would allow users to capture use of the item even if they did not originally win the auction, also allowing the original winner to be compensated for the unexpected loss of the resource.

Unfortunately, it is not always clear that the greatest willingness to pay corresponds to the greatest need. A wealthy person's willingness to pay can exceed a poor person's ability to pay, even if the poor person has much greater need for the item than the wealthy person. Additionally, in many situations, there may be other disadvantages to the use of money. Distributing payments amongst a group of friends can be awkward, and if the amounts are small, often not worthwhile. Moral or legal issues can also prevent the use of money, when the shared resource is deemed to be such that every person should be granted equal access, regardless of wealth. Kidney exchange (see, among many others, [Roth et al., 2004]) is an example setting where non-monetary mechanism design has flourished due largely to the illegality of buying and selling organs, with the use of money being seen as "repugnant" [Roth, 2007, Danovitch and Delmonico, 2008].

For these reasons, we may be interested in the design of protocols that do not involve the use of money. This is generally a harder problem, since money allows agents to substantively back up claims that they want or need use of the item. Without money, we must introduce other incentives for the agents not to overstate their preferences, the most natural approach being that agents 'pay' for resources by sacrificing their right to the resource at some later point in time, or the right to some other resource (now or later). One way to achieve this is to introduce some *artificial* currency that can be equally distributed to agents for the sake of fairness, and treat it similarly to if it were real money [Gorokh et al., 2016, 2017]. Such an approach is commonly used for course allocation in business schools [Sönmez and Ünver, 2010], although poor performance of the artificial currency version of the real money mechanisms has forced some schools to move to a more sophisticated mechanism (still using artificial currency) [Budish, 2011, Budish and Kessler, 2014]. Of course, if the artificial currency is transferable among agents, there is the risk that they will exchange it for real money, thereby bringing it back into the picture. Some approaches do not exactly introduce artificial currency, but still allow an agent to obtain extra resources now at the cost of fewer future resources. By keeping track of how much utility an agent should, in expectation, receive in the future, it is sometimes possible to achieve strong guarantees in a way that resembles the use of artificial currency [Athey and Bagwell, 2001, Guo et al., 2009, Abdulkadiroğlu and Bagwell, 2013]. Even within a single round only, it is sometimes possible to design mechanisms for trading resources (without money) that obtain some formal approximation guarantees [Guo and Conitzer, 2010, Han et al., 2011, Maya and Nisan, 2012, Cole et al., 2013a,b, Cheung, 2016, Amanatidis et al., 2017].

Another line of work in online resource allocation without money restricts users to reporting only binary (like/dislike) valuations, with some (randomized) tie-breaking rule in place to decide how to allocate the resource [Aleksandrov et al., 2015, Aleksandrov and Walsh, 2017, Freeman et al., 2018]. These schemes have the benefit of being easy to understand for users, but come at the expense of being unable to distinguish users whose need is great from those whose need is only small. There also exists literature on *dynamic fair division* [Walsh, 2011, Kash et al., 2014, Friedman et al., 2017], but this work predominantly focuses on agents arriving and departing over time rather than dynamic preferences. Allowing agents to buy in and out of participation over time is of course a desirable feature, but exploiting dynamic preferences is crucial to the profitability of shared ownership.

4 Public vs Private Goods

Sometimes, jointly owned items are useful to only a single agent at a time (*private* goods), while other items can be freely used by many agents simultaneously (*public* goods). For instance, the allocation of CPU time to computer applications is a private good, given that the CPU can only service one application at a time, whereas the storage of data in a shared cache is a public good, since many applications can access the data once stored [Kunjir et al., 2017]. Outside of existing computer systems, we could consider a neighborhood jointly owning a maintenance robot that performs tasks such as clearing trash, repairing pot holes, and trimming trees, with the robot's tasks being collectively decided by the neighborhood.

While public goods are a particularly natural candidate for joint ownership, allowing for multiple agents to utilize a resource simultaneously and thus preventing duplicate cost and effort, they also present additional difficulties. For instance, a basic and problematic form of strategizing when goods are public is the *free-rider* problem [Samuelson, 1954, Green and Laffont, 1979], where an agent pretends to dislike some popular course of action in the hope that the action will proceed anyway, but they will not have to pay for it (in the sense of real or artificial currency, or future consumption). For instance, if I know that all of my neighbors are desperate for a particular pot hole to be repaired, then my best action may be to report that I would rather the robot did some other task, like removing a tree. Since all the other neighbors wanted the pot hole repaired, the robot will still perform that task (depending on the exact form of the protocol), and I get to enjoy the benefits from that, while also profiting from the fact that the protocol believes me to be unhappy with the outcome and will attempt to compensate me later.

Building on the classic work of Groves and Ledyard [1977], who propose quadratic pricing in public good economies, Lalley and Weyl [2018a] propose *Quadratic Voting* for binary collective decisions, where voters purchase votes at a price that is quadratic in the number of votes bought. Under a set of reasonable assumptions, voters buy a number of votes proportional to how much they care about the issue at hand [Posner and Weyl, 2014, Weyl, 2017, Lalley and Weyl, 2018b]. Potentially, agents could 'buy' their votes using not only real money but also artificial currency or even 'effort' (say, buying votes for the maintenance robot to pick up trash requires you to spend some time clearing trash yourself). In other situations, the free-rider problem can be avoided if it is possible to restrict access to otherwise public goods in some way, say by restricting access to a shared facility to members only, or denying access to cached data after a certain number of requests. Goods for which this is possible are known as *club goods* [Buchanan, 1965]. For such goods, it can be possible to prevent free-riding by disallowing access to users that report only a low utility for use of the good [Moulin, 1994, Pu et al., 2016].

Another difference between public and private goods lies in how we may wish to define what a fair solution looks like. For private goods, fairness definitions can utilize the fact that each agent receives a well-defined allocation, be it in terms of what items the agent has control of, or for how much time. We may wish to guarantee each agent a proportional share of her utility for all resources combined [Steinhaus, 1948], or a share that she (almost) prefers to that of every other agent [Foley, 1967, Lipton et al., 2004]. For public goods, no agent is ever in sole possession of an item, so asking whether an agent envies another is less natural, but fairness notions based on proportionality are still applicable [Conitzer et al., 2017], as are those based on stability with respect to coalitional deviations (the *core*) [Foley, 1970, Fain et al., 2016, 2018]. The Maximum Nash Welfare solution [Nash, 1950], which maximizes the *product* of agents' utilities, (approximately) satisfies several of these fairness properties [Caragiannis et al., 2016, Conitzer et al., 2017], making it an appealing objective for collective decision making. In online settings, greedy algorithms for maximizing the Nash welfare provide both good approximations to the optimal Nash welfare, and satisfy desirable properties as algorithms in their own right [Freeman et al., 2017].

5 Communication outside the protocol

Often, we would expect that joint owners have the ability to communicate with each other outside of the formal allocation protocol. They may therefore be able to modify the allocation chosen by the algorithm. An agent allocated usage of the item may be able to pass use to a different agent (potentially for something in return), or to share the item with other agents (say, by carpooling, or storing data in a cache that will be useful to other applications).

While in some cases, such outside communication is desirable, in others it has the capacity to subvert the shared protocol. Secondary resale markets could easily devolve into simply having an exchange rate between artificial and real currency, effectively re-introducing money to the system. Private communication could also lead to collusion; in bidding rings, for instance, groups of buyers agree to share information prior to an auction, and then behave in a pre-determined way accordingly, sharing any profits [McAfee and McMillan, 1992, Leyton-Brown et al., 2000, 2002, Marshall and Marx, 2007]. For sufficiently general settings, it has been shown that the space of collusion-resistant mechanisms is extremely narrow [Schummer, 2000, Goldberg and Hartline, 2005], but collusion-resistant mechanisms have been designed for certain special cases such as spectrum auctions [Wu et al., 2009], or when some amount of agent type verification is allowed [Penna and Ventre, 2014]. Resistance to such group-strategic behavior will clearly play an important role in the design of shared ownership protocols moving forward.

6 Who chooses the protocol, and what should be the objective?

One clear theme emerges throughout in the study of collective decision making: that there is unlikely to be a single 'best' rule for all purposes. Thus, it is worth thinking about how a shared ownership protocol will be chosen among competing alternatives. As we have already seen, a wealthy person who wants to be able to use a shared resource at little notice may have quite different preferences over candidate protocols than a poor person with a fixed, predictable schedule.

When the resource in question is truly a publicly owned resource, such as a facility at a public park, then we may expect that the government could choose a protocol that optimizes in certain socially-desirable directions. However, when the resource is shared among a smaller group of people, then the protocol must be agreed upon by them, for example by voting over candidate protocols. This could be an onerous task, since an algorithmic protocol must be precisely defined to handle all eventualities. For informal shared ownership arrangements that we observe in practice currently, rare and minor exceptions can be handled informally with relative ease (e.g., if someone falls ill, they can have priority use of the car provided it does not happen often and they "make up for it" at some point in the future). To design protocols that take all such eventualities into account is difficult and the resulting protocols are unlikely to be generally applicable to many situations. Additionally, the space of protocols would become extremely large, possibly unmanageably so.

This perspective permits a more positive view of side deals resulting from communication outside the mechanism. Rather than handling all possible scenarios within the protocol, we may wish to allow, and even facilitate, users making informal trades that are not visible to the protocol. For instance, having established usage rights for a particular time slot, we could allow an agent to do as she wishes with those rights. Rare and exceptional cases would then be naturally handled in a way that allows for a greater degree of flexibility than would be convenient to explicitly encode into a protocol for shared ownership.

One could imagine a time where goods are sold with ready-to-use algorithms for shared usage, one of which can be activated upon purchase. While convenient, this possibility leads us to further questions. Moving from a predominantly individual ownership model to joint ownership will not benefit everyone; in particular, some manufacturers may see a decreased demand for their goods if not everyone needs a separate copy. Is it in their interest to help facilitate joint ownership by building (good) sharing protocols into their products? The transformation of society towards a true shared-ownership economy will require not only improvements in algorithmic economic design, but also buy-in from other parties, from corporations to government to individuals.

7 Conclusion

Algorithmic approaches for allocating shared resources among multiple self-interested parties are already being used today in the allocation of computing resources. But this may be but a harbinger of a much broader societal development. More and more resources are starting to be algorithmically controlled and connected to computer networks. Many of these resources either need to be assigned to a user at any given point in time, or decisions need to be made about their use, where multiple parties have an interest in these decisions. This makes them candidates to have their usage governed by the types of algorithms discussed in this chapter, possibly ushering in an age of a true shared-ownership economy. Many open questions remain, and by collaborating on them, computer scientists, economists and others may help to realize this vision in a desirable fashion.

Acknowledgments

Conitzer is thankful for support from NSF under awards IIS-1527434 and CCF-1337215. Freeman is thankful for support from a Facebook PhD Fellowship.

Bibliography

Atila Abdulkadiroğlu and Kyle Bagwell. Trust, reciprocity, and favors in cooperative relationships. *American Economic Journal: Microeconomics*, 5(2): 213–259, 2013.

Martin Aleksandrov and Toby Walsh. Pure Nash equilibria in online fair division. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 42–48, Melbourne, Australia, 2017.

Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. Online fair division: Analysing a food bank problem. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2540–2546, Buenos Aires, Argentina, 2015.

Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the Eighteenth ACM Conference on Economics and Computation (EC)*, pages 545–562, Cambridge, MA, USA, 2017.

Susan Athey and Kyle Bagwell. Optimal collusion with private information. *RAND Journal of Economics*, 32(3):428–465, 2001.

Moshe Babaioff, Yishay Mansour, Noam Nisan, Gali Noti, Carlo Curino, Nar Ganapathy, Ishai Menache, Omer Reingold, Moshe Tennenholtz, and Erez Timnat. ERA: A framework for economic resource allocation for the cloud. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 635–642, 2017.

James M Buchanan. An economic theory of clubs. *Economica*, 32(125):1–14, 1965.

Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

Eric Budish and Judd Kessler. Changing the course allocation mechanism at Wharton. *Chicago Booth Research Paper,(15-08)*, 4, 2014.

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the Seventeenth ACM Conference on Economics and Computation (EC)*, pages 305–322, Maastricht, The Netherlands, 2016.

Shuchi Chawla, Nikhil Devanur, Janardhan Kulkarni, and Rad Niazadeh. Truth and regret in online scheduling. In *Proceedings of the Eighteenth ACM Conference on Economics and Computation (EC)*, pages 423–440, Cambridge, MA, USA, 2017.

Yun Kuen Cheung. Better strategyproof mechanisms without payments or prior - an analytic approach. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 194–200, New York, NY, USA, 2016.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division: allocating divisible items without payments. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC)*, pages 251–268, Saint Paul, Minnesota, USA, 2013a.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Positive results for mechanism design without money. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1165–1166, Saint Paul, Minnesota, USA, 2013b.

Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, pages 629–646, Cambridge, MA, USA, 2017.

Gabriel M Danovitch and Francis L Delmonico. The prohibition of kidney sales and organ markets should remain. *Current Opinion in Organ Transplantation*, 13(4):386–394, 2008.

Robert I Davis and Alan Burns. A survey of hard real-time scheduling for multiprocessor systems. *ACM Computing Surveys (CSUR)*, 43(4):35, 2011.

Brandon Fain, Ashish Goel, and Kamesh Munagala. The core of the participatory budgeting problem. In *Proceedings of the Twelth International Conference on Web and Internet Economics (WINE)*, pages 384–399, 2016.

Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, pages 575–592, Ithaca, NY, USA, 2018.

Duncan K. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.

Duncan K Foley. Lindahl's solution and the core of an economy with public goods. *Econometrica: Journal of the Econometric Society*, pages 66–72, 1970.

Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair and efficient social choice in dynamic settings. In *Proceedings of the Twenty Sixth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4580–4587, Melbourne, Australia, 2017.

Rupert Freeman, Seyed Majid Zahedi, Vincent Conitzer, and Benjamin C. Lee. Dynamic proportional sharing: A game-theoretic approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2(1):3:1–3:36, April 2018.

Eric Friedman, Christos-Alexandros Psomas, and Shai Vardi. Controlled dynamic fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, pages 461–478, Cambridge, MA, USA, 2017.

Andrew V Goldberg and Jason D Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 620–629, Vancouver, Canada, 2005.

Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. Near-efficient allocation using artificial currency in repeated settings. 2016.

Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. From monetary to non-monetary mechanism design via artificial currencies. 2017.

Jerry R Green and Jean-Jacques Laffont. Incentives in public decision making. 1979.

Theodore Groves and John Ledyard. Optimal allocation of public goods: A solution to the "free rider" problem. *Econometrica: Journal of the Econometric Society*, pages 783–809, 1977.

Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 881–888, Toronto, Canada, 2010.

Mingyu Guo, Vincent Conitzer, and Daniel M. Reeves. Competitive repeated allocation without payments. In *Proceedings of the Fifth Workshop on Internet and Network Economics (WINE)*, pages 244–255, Rome, Italy, 2009.

Li Han, Chunzhi Su, Linpeng Tang, and Hongyang Zhang. On strategy-proof allocation without payments or priors. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE)*, pages 182–193, Singapore, 2011.

Sungjin Im and Janardhan Kulkarni. Fair online scheduling for selfish jobs on heterogeneous machines. In *Proceedings of the Twenty-Eighth ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 185–194, Moneterey, California, USA, 2016.

Ian Kash, Ariel D Procaccia, and Nisarg Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51:579–603, 2014.

Antoon WJ Kolen, Jan Karel Lenstra, Christos H Papadimitriou, and Frits CR Spieksma. Interval scheduling: A survey. *Naval Research Logistics (NRL)*, 54 (5):530–543, 2007.

Mayuresh Kunjir, Brandon Fain, Kamesh Munagala, and Shivnath Babu. ROBUS: Fair cache allocation for data-parallel workloads. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD)*, pages 219–234, Chicago, IL, USA, 2017.

Steven P Lalley and E Glen Weyl. Nash equilibria for quadratic voting. *arXiv preprint arXiv:1409.0264*, 2018a.

Steven P. Lalley and E. Glen Weyl. Quadratic voting: How mechanism design can radicalize democracy. *American Economic Association Papers and Proceedings*, 1(1), 2018b.

Kevin Leyton-Brown, Yoav Shoham, and Moshe Tennenholtz. Bidding clubs: institutionalized collusion in auctions. In *Proceedings of the Second ACM Conference on Electronic Commerce (EC)*, pages 253–259, Minneapolis, MN, USA, 2000.

Kevin Leyton-Brown, Yoav Shoham, and Moshe Tennenholtz. Bidding clubs in first-price auctions. In *Proceedings of the Eighteenth AAAI Conference on Artificial Intelligence (AAAI)*, pages 373–378, Edmonton, Canada, 2002.

Richard J Lipton and Andrew Tomkins. Online interval scheduling. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, volume 94, pages 302–311, Arlington, VA, USA, 1994.

Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the Fifth*

ACM Conference on Electronic Commerce (EC), pages 125–131, New York, NY, USA, 2004. ACM.

Robert C Marshall and Leslie M Marx. Bidder collusion. *Journal of Economic Theory*, 133(1):374–402, 2007.

Avishay Maya and Noam Nisan. Incentive compatible two player cake cutting. In *Proceedings of the Eighth Workshop on Internet and Network Economics (WINE)*, pages 170–183, Liverpool, UK, 2012.

R Preston McAfee and John McMillan. Bidding rings. *The American Economic Review*, pages 579–599, 1992.

Hervé Moulin. Serial cost-sharing of excludable public goods. *The Review of Economic Studies*, 61(2):305–325, 1994.

John F Nash. The bargaining problem. *Econometrica: Journal of the Econometric Society*, pages 155–162, 1950.

Paolo Penna and Carmine Ventre. Optimal collusion-resistant mechanisms with verification. *Games and Economic Behavior*, 86:491–509, 2014.

Eric A Posner and E Glen Weyl. Quadratic voting as efficient corporate governance. *The University of Chicago Law Review*, 81(1):251–272, 2014.

Qifan Pu, Haoyuan Li, Matei Zaharia, Ali Ghodsi, and Ion Stoica. FairRide: Near-optimal, fair cache sharing. In *Proceedings of the Thirteenth USENIX Symposium on Networked Systems Design and Implementation*, pages 393–406, Santa Clara, CA, USA, 2016.

Alvin E Roth. Repugnance as a constraint on markets. *The Journal of Economic Perspectives*, 21(3):37–58, 2007.

Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.

Paul A Samuelson. The pure theory of public expenditure. *The Review of Economics and Statistics*, 36(4):387–389, 1954.

James Schummer. Manipulation through bribes. *Journal of Economic Theory*, 91(2):180–198, 2000.

Tayfun Sönmez and M Utku Ünver. Course bidding at business schools. *International Economic Review*, 51(1):99–123, 2010.

Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

Yuan Tian. Strategy-proof and efficient offline interval scheduling and cake cutting. In *Proceedings of the Ninth International Conference on Web and Internet Economics (WINE)*, pages 436–437, Cambridge, MA, USA, 2013.

Toby Walsh. Online cake cutting. In *Proceedings of the Third International Conference on Algorithmic Decision Theory (ADT)*, pages 292–305, Piscataway, NJ, USA, 2011.

E. Glen Weyl. The robustness of quadratic voting. *Public choice*, 172(1-2): 75–107, 2017.

Yongle Wu, Beibei Wang, KJ Ray Liu, and T Charles Clancy. A scalable collusion-resistant multi-winner cognitive spectrum auction game. *IEEE Transactions on Communications*, 57(12), 2009.