Complexity-Theoretic Barriers in Economics

Tim Roughgarden[1]

Abstract

We survey several unexpected connections between computational complexity and fundamental economic questions that appear unrelated to computation.

1 Introduction

Computational complexity theory is a branch of theoretical computer science that studies the amount of resources (time, space, communication, etc.) required to solve different computational problems (see e.g. [14].) Computational complexity has already had plenty to say about economic equilibria. Some of the most celebrated results in algorithmic game theory give strong evidence that there is no general efficient algorithm for computing various types of equilibria, including mixed Nash equilibria in finite games and market equilibria in markets with divisible goods and concave utility functions (see e.g. [7, 9, 6]).

Recent work has established several unexpected connections between computational complexity and fundamental economic questions that seem to have little to do with computation (as opposed to the explicitly computational questions mentioned above). These new connections suggest the possibility of a rich theory with many applications. We briefly describe a few of these connections, with the hope of inspiring further work along similar lines.

2 Barriers for Simple Auctions

The goal of auction design is to develop auction formats that enjoy good properties when the participants act in a self-interested manner (i.e., reach an "equilibrium"). One fundamental and high-stakes example is the design of multi-item auctions for selling wireless spectrum licenses, as used by the Federal Communications Commission (see e.g. [8]). Many practitioners in multi-item auction design abide by the following rule of thumb: simple auctions, such as selling each item separately in parallel, can perform well if and only if there are not strong synergies across items. An example of such a synergy would be between licenses for spectrum in northern and southern California—a company looking to roll out a new statewide data plan only wants one of the licenses if it can get both. (Or for a more prosaic example, think of a left shoe and a right shoe.)

This 25-year-old rule of thumb has only been formalized over the past few years. Proving that simple auctions can work well without item synergies boils down to proving "price of anarchy (POA)" bounds, which state that all of the equilibria of a simple auction are almost as good as an optimal outcome (e.g., in terms of the social welfare achieved). The theory of "smooth games" developed in [16], and its subsequent extensions (e.g. [20]), turns out to

[1]Department of Computer Science, Stanford University. Email: `tim@cs.stanford.edu`.

be well suited to this goal. Many good POA bounds for different auction formats are now known (see the survey paper [18]).

Proving the converse, that simple auctions *cannot* work well when there are item synergies, seems more difficult. While a positive result only requires analyzing one simple auction format, a negative result needs to rule out *every* possible simple auction format. Fortunately, computational complexity theory has a rich toolbox for proving such impossibility results. The main result in [15] uses lower bounds in communication complexity (see e.g. [12, 17]) to formally prove the impossibility of good simple auctions when there are strong item synergies.

One of many intriguing open research directions is to prove impossibility results that are stronger than those implied by the communication complexity connection established in [15]. The POA lower bounds proved in [15] rely on two properties of equilibria: guaranteed existence and efficient verifiability. Equilibria possess additional properties—can these be exploited to prove stronger impossibility results? Specifically, every equilibrium implicitly solves the best-response problem for every player. These single-agent problems are a second source of potential impossibility, orthogonal to the difficulty posed by approximating the underlying multi-agent optimization problem.

3 Barriers for Competitive Equilibria

One of the most basic notions in economics is that of a competitive equilibrium, which comprises prices for items that equalize supply and demand. In a market with divisible goods (milk, wheat, etc.), a competitive equilibrium exists under reasonable assumptions. Many markets with indivisible goods (houses, wireless spectrum licenses, airport landing slots, etc.), however, can fail to possess a competitive equilibrium. Recent work [19] uses computational complexity to explain the rampant non-existence of such equilibria.

To be more concrete, fix a set of m (non-identical and indivisible) items and let $\mathcal{V}$ denote a set of valuation functions (i.e., real-valued set functions with ground set $\{1, 2, \ldots, m\}$). For example, $\mathcal{V}$ could be the set of additive valuations, unit-demand valuations, valuations satisfying the gross substitutes condition [11], or general monotone valuations. One of the main results in [19] proves a necessary condition for the guaranteed existence of a competitive equilibrium, which compares the computational tractability of two different problems. The *welfare-maximization problem* for $\mathcal{V}$ is: given valuations $v_1, \ldots, v_n \in \mathcal{V}$ for n agents (described either succinctly or through oracle access), compute the allocation $(S_1, \ldots, S_n)$ of items to agents that maximizes the social welfare $\sum_{i=1}^{n} v_i(S_i)$. The *utility-maximization problem* for $\mathcal{V}$ is: given a single valuation $v \in \mathcal{V}$ and item prices $p_1, \ldots, p_m$, compute a utility-maximizing bundle (i.e., a bundle in $\operatorname{argmax}_S v(S) - \sum_{j \in S} p_j$). The welfare-maximization problem (with n agents) is generally only harder than the utility-maximization problem (with 1 agent). The necessary condition in [19] for the guaranteed existence of a competitive equilibrium in every market with valuations in $\mathcal{V}$ is that the welfare-maximization problem is no harder than the utility-maximization problem. By "no harder," we mean that given a subroutine that solves the utility-maximization problem, it is possible to solve also the welfare-maximization problem (with the number of invocations of the subroutine and the additional work performed bounded above by some polynomial function of n and m).

One of many open questions in this direction is to develop analogous impossibility results for settings without prices or money, such as markets for matching (marriage, college admissions, etc.) or kidney exchange. One possible plan for proving such impossibility results is to derive a contradiction from statements of the following types: (i) Pareto optimality, which states that no one can be made better off without making someone else worse off, is computationally difficult to verify; (ii) a competitive equilibrium is computationally easy to verify (assuming that the single-agent utility-maximization problem is tractable); (iii) "first and second welfare theorems," which provide a close correspondence between competitive equilibria and Pareto-optimal outcomes.

4 Barriers for Tractable Characterizations

Border's theorem is a fundamental result about probability distributions that has important applications in auction design [2]. One way to phrase Border's theorem is that it gives an explicit description of the facets of a certain polytope, which is the projection of the space of all single-item auctions onto their "interim allocation rules" with respect to a prior distribution over bidders' valuations for an item. Border's theorem is useful in auction design both for understanding the structure of revenue-optimal auctions (as one varies the prior), and in efficiently computing such an auction (given a description of the prior as input). Generalizations of Border's theorem beyond the single-item auction setting have been hard to come by (though see [1, 3, 5]), unless one resorts to approximation (as in [3, 4]). Recent work [10] uses computational complexity theory to explain this state of affairs: under standard complexity assumptions (that the "polynomial hierarchy" does not collapse), Border's theorem cannot be extended significantly beyond single-item auctions. The techniques introduced in [10] do not appear specific to auction design, and show the promise of ruling out tractable characterizations of many other mathematical objects. For example, in recent years there has been tremendous progress in ruling out "extended linear programming formulations" for various computational problems, including graph matching (see e.g. [13]). Can the techniques introduced in [10] advance this research agenda further, under suitable complexity assumptions?

References

[1] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason D. Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce, EC '12*, page 17, 2012.

[2] K. C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59(4):1175–1187, 1991.

[3] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC*, pages 459–478, 2012.

[4] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Proceedings of the 53rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 130–139, 2012.

[5] Y.-K. Che, J. Kim, and K. Mierendorff. Generalized reduced form auctions: A network flow approach. *Econometrica*, 81:2487–2520, 2013.

[6] X. Chen, D. Dai, Y. Du, and S.-H. Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 273–282, 2009.

[7] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):14, 2009.

[8] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.

[9] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[10] P. Gopalan, N. Nisan, and T. Roughgarden. Public projects, Boolean functions, and the borders of Border's theorem. In *Proceedings of the 16th Annual ACM Conference on Economics and Computation (EC)*, 2015.

[11] Alexander S. Kelso and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

[12] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge, 1996.

[13] T. Rothvoß. The matching polytope has exponential extension complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 263–272, 2014.

[14] T. Roughgarden. Computing equilibria: A computational complexity perspective. *Economic Theory*, 42(1):193–236, 2010.

[15] T. Roughgarden. Barriers to near-optimal equilibria. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 71–80, 2014.

[16] T. Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):32, 2015.

[17] T. Roughgarden. *Communication Complexity (for Algorithm Designers)*. Now Publishers, 2016.

[18] T. Roughgarden, V. Syrgkanis, and É. Tardos. The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101, 2017.

[19] T. Roughgarden and I. Talgam-Cohen. Why prices need algorithms. In *Proceedings of the 16th Annual ACM Conference on Economics and Computation (EC)*, 2015.

[20] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 211–220, 2013.