Principled Autonomous Decision Making for Markets (Doctoral Consortium)

Enrique Areyan Viqueira
Brown University
115 Waterman St, Providence,
RI, United States, 02906
eareyan@brown.edu

Optimal Allocation of Resources

A quintessential problem at the intersection of *economics* and *computer science* is that of *resource allocation*. Broadly construed, this problem involves a set of *resources* that must be distributed to a set of *agents* without violating a set of constraints. Abstractly, a resource could be thought of as any input that an agent needs to fulfill its goal as it acts on a given environment. A constraint could be thought of as any limit imposed by the environment or other agents, each of which tries to fulfill their independent, and often conflicting, goals.

From an *economics* point of view, this problem is of paramount importance. On the one hand, an economy that is unable to properly allocate resources is inefficient, and thus fails to serve its constituents. In very extreme cases, it could result in catastrophes and destruction of entire societies[1]. On the other hand, an economy that is capable of properly allocating its resources is more likely to be able to sustain its participants and propel growth and stability.

From a *computational* and *mathematical* point of view, the resource allocation problem poses interesting and important challenges, as it is often the case that these types of problems are intractable due to their combinatorial nature. In fact, in almost all modern applications these problems are NP-hard, which means that there is no known algorithm that will take a reasonable amount of time to solve every conceivable instance. To make matters worse, even if a proposed allocation (an assignment of resources to agents) is given, there might not be a fast way of checking if this is, in fact, an optimal (or near optimal) allocation.

A recent and important example of combinatorial resource allocation problems is that of spectrum auctions. In a spectrum auction, a government or central planner uses an auction mechanism to allocate licenses for the right to use scarce electromagnetic spectra to transmit signals. These signals are usually used by radio or television stations and are highly valuable. To briefly illustrate the complexity of this type of resource allocation problem, consider the recent 2014 700 MHz Canadian auction [2], where 98 licenses were up for auction. An agent trying to compete for these resources would have to reason about 2^{98} or roughly 316 billion billion billion different possible combinations of available licenses. Any computational device used by the auctioneer or market participants cannot rely on faulty or slow methods, as there is much at stake; some rounds of the Canadian auction reached a total of over $7 billion in trades.

Another example of a problem in combinatorial resource allocation is ad-exchanges. In particular, we have been studying how to construct an autonomous bidding agent for the Trading Agent Competition Ad Exchange game (TAC AdX) [5]. This competition models online ad exchanges where agents face the challenge of bidding for display-ad impressions needed to fulfill advertisement contracts, after which they earn the amount the advertiser budgeted. The agents in this game have combinatorial valuations since they have to procure enough impressions (ad display opportunities) from *different* market segments but are *indifferent* between market segments of similar characteristics. This is a typical resource allocation problem where resources are indivisible (ad display opportunities are discrete) and agents face two challenges: scarcity (there are a finite amount of ad display opportunities) and adversarial competition (other agents compete for the same resources).

In our attempt to build a bidding agent for this competition, we wondered if it would be possible to compute a Walrasian equilibrium [6] for this type of environment and bid based on it. The combinatorial structure of this problem makes it impossible to guarantee the existence of such outcomes, even under relatively mild conditions. To make matters worse, even if such equilibrium exists, it would be computationally inefficient to compute for all but very small instances. Nonetheless, we remained interested in solving the resource allocation problem posed by this game and our investigation resulted in our paper *On Approximate Welfare- and Revenue-Maximizing Equilibria for Size-Interchangeable Bidders* accepted as an extended abstract at AAMAS17 (Sixteenth International Conference on Autonomous Agents and Multi-agent Systems) and currently under review for EC17 (Eighteenth ACM conference on Economics and Computation), where we propose and analyze a relaxation of Walrasian equilibrium, which we call *Restricted-Envy-Free Pricing* (REFP).

Our relaxation generalizes the concept of an *Envy-Free Pricing* (EFP) [1, 3, 4] where market participants must be envy-free (i.e., utility-maximizing) at the set prices but the market does not need to clear, by requiring that only *winners* (participants that are part of the allocation) are envy-free.

[1] Consider the current situation in Venezuela: `https://swarajyamag.com/world/stop-in-the-name-of-welfare`

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

A REFP has important theoretical guarantees: it always exists, unlike both a Walrasian equilibrium and an EFP. Moreover, for the case of agents in the TAC game and for the case of single-minded bidders (a fundamental benchmarked used to investigate combinatorial markets) we propose efficient algorithms to obtain REFP outcomes and investigate their properties both theoretically and experimentally. Experimentally, we found that for different distributions of input markets, our solution concept retains most of the guarantees of EFP outcomes *despite* the fact that only winners are guaranteed to be envy-free. Moreover, our algorithms outperform other algorithms in the literature on several important metrics such as welfare and revenue. Some of the tested distributions were based on real-world web-usage data, showing the potential real-world applicability of our methods.

Our AAMAS paper is thus an example that illustrates the core of my research agenda: achieving economic efficiency (for example, a *Walrasian Equilibrium*) within the realm of computational efficiency (correspondingly, *Restricted-Envy-Free Pricing*) as it pertains to resource allocation problems and equilibrium computation. From an applied point of view, the goal of this research is to support automated decision makers in increasingly complex environments.

The Future (and Present) of Decision Making

Real-world markets are among the most important and complex systems known to man. Their complexity entails that decision making cannot rely on faulty reasoning. Moreover, modern markets operate on *real-time*, which means that decisions must be made in fractions of a second. Our economies today rely to some extent on automated decision makers, but the evidence suggests that the economies of the future will rely even more on automated agents. My research will help us create economically-*principled* and computationally-*efficient* agents for complex market environments.

Our initial results show that this is a promising research direction. In the case of the TAC AdX game, our first prototype of an automated agent bidding based only on a REFP[2] was capable of accumulating a positive score against other non-principled bidding agents (see [7] for a similar approach to principled agents). In ongoing experiments,we are testing the hypothesis that bidding based on an approximation of a *Walrasian Equilibrium* can produce a robust agent against other agents that are highly optimized for the specific rules of this game, shedding light on the functioning of the game and thus, on the functioning of real ad-exchange markets.

TAC AdX is just one example of the potential applicability of our research. Other examples include the already mentioned spectrum auctions. In this case, one can also think of constructing an autonomous agent that bids based on a REFP and in fact, we plan to build such an agent using a value-generator based on real-world data from the Canadian spectrum auction. These examples illustrate how my research will pave the way for the construction of *principled* bidding agents capable of making *fast* and *near-optimal* decisions.

[2]Code is available in `http://github.com/eareyan/envy-free-prices`

Acknowledgments

I would like to express my gratitude to my advisor, Amy Greenwald, for her support in writing this manuscript and discussing the ideas contained in it.

REFERENCES

[1] S. Brânzei, A. Filos-Ratsikas, P. B. Miltersen, and Y. Zeng. Envy-free pricing in multi-unit markets. *arXiv preprint arXiv:1602.08719*, 2016.

[2] M. Cave and R. Nicholls. The use of spectrum auctions to attain multiple objectives: Policy implications. *Telecommunications Policy*, 2017.

[3] M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 532–549, New York, NY, USA, 2012. ACM.

[4] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1164–1173. Society for Industrial and Applied Mathematics, 2005.

[5] M. Schain and Y. Mansour. Ad exchange–proposal for a new trading agent competition game. In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, pages 133–145. Springer, 2013.

[6] L. Walras. *Elements of Pure Economics: Or the Theory of Social Wealth*. Elements of Pure Economics, Or the Theory of Social Wealth. Routledge, 2003.

[7] M. P. Wellman, D. M. Reeves, K. M. Lochner, and Y. Vorobeychik. Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, 21:19–36, 2004.