# Fast Algorithms for Electronic Markets

Avrim Blum Anupam Gupta
Computer Science Department
Carnegie Mellon University
{avrim,anupamg}@cs.cmu.edu

The rise of electronic markets involving large numbers of buyers, large numbers of goods, complex preference interactions, and unknown and changing market conditions necessitates an algorthmic, computational perspective on classic economic questions. Buyers and sellers are computational agents and should be treated as such, and fast algorithms are needed for performing economic actions such as pricing and resource allocation in these complex markets. Indeed, this has led to the development of the entire research field of electronic commerce and algorithmic mechanism design. This work aims to bring this field in several exciting new directions. In particular, can one *automate* tasks such as pricing resources under uncertainty over future demand, and can one *explain* common real-world practices in a formal algorithmic model? Below we present two specific directions this work aims to explore.

**1.** Nearly all algorithmic work in the area of combinatorial auctions (the selling or allocating of a large number of items to buyers who have complex valuations over subsets of these items) has focused on two specific cases. These are the cases of *limited supply* where the seller has just one copy of each item (like in a traditional auction), and that of *unlimited supply* where the seller has an unbounded number of copies of each item (modeling the sale of digital goods such as software). The former can be viewed as a setting where the marginal cost per item to the seller jumps from 0 to infinity after the first copy, and is filled with mostly impossibility results.[1] The latter can be viewed as a case where the marginal cost per item to the seller is fixed (e.g., a grocery store that can purchase apples, bananas, etc from its supplier for 25 cents per apple and 5 cents per banana), and there are many positive statements (including many results by MSR researchers) in that setting. However, there is a large area between these extremes that is important and ubiquitous, and has yet remained unexplored: this is when the marginal costs of items increase according to some curve that lies between a flat line and an immediate jump to infinity as more and more items are sold. For instance, such a model would be appropriate when considering resources such as network bandwidth or computing power: as these resources are depleted, the network provider may have to lease additional, expensive bandwith from other networks. In proposed work, we aim to address the question: if our marginal costs increase at some reasonable rate, can one achieve the same (or similar) strong guarantees as achievable in the usual setting of unlimited supply? Furthermore, can we do so via a natural item pricing scheme? In particular, we would like to achieve a strong guarantee for buyers arriving *online*, with *arbitrary* valuation functions.

[1] At a high level, the difficulty is that if there is the potential for late arrival of a buyer who has high value but only on receiving the entire collection of $n$ items, then one cannot sell small pieces of the collection to early buyers who want just one or two items.

We are working on this question with graduate student Ankit Sharma, and currently have a number of exciting preliminary results on this problem.

A related question concerns the opposite direction of settings with *economies of scale*. Building a million widgets is often substantially cheaper than a million times the cost of building a single widget. And if you don't know the future demand, is it worth getting into the market for selling widgets at all given that you are likely to be initially selling at a loss? This problem in the case of just a single good (widgets) and a single irrevocable binary decision (build widget factory or not) can be viewed as a version of the classic rent-or-buy problem in online algorithms. But what happens under the more complex structure of combinatorial pricing? Even formulating this problem in a precise and convincing way appears to be a challenge. This is a fairly open-ended direction where we would very much like to work together with MSR researchers and believe such collaborations could be extremely fruitful.

**2.** Another exciting and important area of investigation is that of Bayesian Mechanism Design—loosely speaking, these are mechanism design problems where we know stochastic information about the buyers, and we would like to use them to design better mechanisms. As an example, imagine we are selling a set of goods, and as usual we don't know what the precise valuations of the buyers for the goods are; however, we posit that these valuations are random variables, drawn from a probability distribution which we do know. (E.g., in online markets, we ofrten know a lot of information about the demographic of the buyers, and hence can use historical data to get this probability distribution.) Given this information, can we do much better than in the worst-case? For simple cases (the so-called "single-parameter" setting), we know how to obtain the best revenue possible, via a beautiful but somewhat mysterious mechanism of Myerson. However, when buyers want to buy several objects, and their valuations for these objects are correlated,[2] then these techniques do not work any more. Even when each user has independent valuations for items, these problems get difficult to solve if there are auxiliary constraints on which items can be simultaneously sold (or even which ones can simultaneuously be shown to buyers—such constraints often arise from budgets and capacities). Our goal in this area is to develop broad-spectrum techniques and flexible algorithms in this general area. The hope is to use ideas from linear programming, which not only demystifies current results, but also gives a powerful specification and computational language in which to encode the desiderata of the auctions of the future.

**Collaboration with MSR:** There are multiple potential collaborators at Microsoft Research in this effort, including Kamal Jain at MSR Redmond, and Kunal Talwar, Moshe Babaioff, Alex Slivkins, and Andrew Goldberg at MSR SVC. The PIs, and their past students, have had fruitful collaborations with researchers at MSR SVC on related topics in the past.

**Budget Information:** We request a budget of $80k/year for two years. This will primarily be used to support one graduate student, Ankit Sharma, who is jointly supervised by the two PIs, and also to support some minor travel and sundry expenses.

[2]This is likely to happen: imagine a buyer's valuations for a flight ticket and a hotel room.