Mechanism Design with Uncertain Inputs

Uriel Feige

We consider a task of scheduling with a common deadline on a single machine. Every player reports to a scheduler the length of his job and the scheduler needs to finish as many jobs as possible by the deadline. For this simple problem, there is a truthful mechanism that achieves maximum welfare in dominant strategies. The new aspect of our work is that in our setting players are uncertain about their own job lengths, and hence are incapable of providing truthful reports (in the strict sense of the word).

For a probabilistic model for uncertainty we show that even with relatively little uncertainty, no mechanism can guarantee a constant fraction of the maximum welfare. To remedy this situation, we introduce a new measure of economic efficiency, based on a notion of a {\em fair share} of a player, and design mechanisms that are $\Omega(1)$-fair. In addition to its intrinsic appeal, our notion of fairness implies good approximation of maximum welfare in several cases of interest.

Joint work with Moshe Tennenholtz