Mechanism Design For Fair Division

Vasilis Gkatzelis, (Stanford University)

We revisit the classic problem of fair division from a mechanism design perspective, using proportional fairness as a benchmark. In particular, we aim to allocate a collection of divisible items to a set of agents while incentivizing the agents to be truthful in reporting their valuations. For the very large class of homogeneous valuations, we design a truthful mechanism that provides every agent with at least a $1/e \approx 0.368$ fraction of her proportionally fair valuation. To complement this result, we show that no truthful mechanism can guarantee more than a 0.5 fraction, even for the restricted class of additive linear valuations. To guarantee truthfulness, our mechanisms discard a carefully chosen fraction of the allocated resources; we conclude by uncovering interesting connections between our mechanism and celebrated solutions from the mechanism design literature that use money instead.

(Joint work with Richard Cole and Gagan Goel.)