Abstract:
We investigate computational and mechanism design aspects of optimal scarce resource allocation, where the primary rationing mechanism is through waiting times. Specifically we consider the problem of allocating medical treatments to a population of patients. Each patient has demand for exactly one unit of treatment, and can choose to be treated in one of k hospitals, $H_1, …, H_k$. Different hospitals have different costs, which are fully paid by a third party ---the "payer"--- and do not accrue to the patients. The payer has a fixed budget B and can only cover a limited number of treatments in the more expensive hospitals. Access to over-demanded hospitals is rationed through waiting times: each hospital H_i will have waiting time w_i. In equilibrium, each patient will choose his most preferred hospital given his intrinsic preferences and the waiting times. The payer thus computes the waiting times and the number of treatments authorized for each hospital, so that in equilibrium the budget constraint is satisfied and the social welfare is maximized.

We show that even if the patients' preferences are known to the payer, the task of optimizing social welfare in equilibrium subject to the budget constraint is NP-hard. We also show that, with constant number of hospitals, if the budget constraint can be relaxed from B to $(1+\epsilon)B$ for an arbitrarily small constant ϵ, then the original optimum under budget B can be approximated very efficiently.

Next, we study the endogenous emergence of waiting time from the dynamics between hospitals and patients, and show that there is no need for the payer to explicitly enforce the optimal equilibrium waiting times. When the patients arrive uniformly along time and when they have generic types, all that the payer needs to do is to enforce the total amount of money he would like to pay to each hospital. The waiting times will simply change according to the demand, and the dynamics will always converge to the desired waiting times in finite time.

We then go beyond equilibrium solutions and investigate the optimization problem over a much larger class of mechanisms containing the equilibrium ones as special cases. In the setting with two hospitals, we show that under a natural assumption on the patients' preference profiles, optimal welfare is in fact attained by the randomized assignment mechanism, which allocates patients to hospitals at random subject to the budget constraint, but avoids waiting times.

Finally, we discuss potential policy implications of our results, as well as follow-up directions and open problems.