Queueing to learn

Chiara Margaria*

April 15, 2016

I study a dynamic resource allocation problem in a queueing setting. A continuum of forward-looking agents compete for a unit flow of resource, and decide whether and when engage in costly queueing to be served. Valuations fluctuate over time, independently across agents; each agent faces an experimentation problem inasmuch as payoffs are informative about the prevailing valuation.

I address the problem of designing a *service discipline* to maximize long-run average payoffs. A service discipline specifies how to allocate the resource across agents as a function of the time since the last check-in. The analysis concentrates on steady state equilibria in which agents have no incentive to abandon the queue before being served. The problem of the designer reduces to the choice of the distribution of the time by which an agent who joins the queue is served in the steady state.

First, I show that any service discipline can be identified by two moments that are a sufficient statistic for it in terms of agents' payoffs. Second, I characterize the set of the feasible summary statics. Third, I solve for the optimal service discipline in a relaxed problem and show that it satisfies the incentive constraints of the original problem. A key observation is that any service time distribution that disincentivizes reneging is a NBUE (new better than used in expectation) distribution: the residual expected waiting time faced by an agent in the queue must be weakly lower than his expected waiting time at check-in.

Combining these results, I show that the service discipline that maximizes welfare is a mixture of first-come first-served and a processor sharing discipline. The mixture depends on the parameter values, and highlights the trade-off between minimizing waiting cost and targeting the "right agents". On the one hand, if the waiting cost is high, the first-come first-served discipline is optimal because it induces, in equilibrium, a shorter queue. On the other hand, for lower waiting costs, the processor sharing discipline performs better because it allows targeting "returning" agents who have precise information about their valuation.

This trade-off is likely to arise in the allocation of scarce resources such as researchers time and lab equipment across different projects that proceed through bursts and setbacks.

*Department of Economics, Yale University, 28 Hillhouse Ave., New Haven, CT 06520, USA, chiara.margaria@yale.edu