

Data-driven decisions: how big should your data really be?

Omar Besbes

Columbia University, Graduate School of Business

Abstract:

We consider two fundamental questions in data-driven decision making: 1) how should a decision-maker construct a mapping from historical data to decisions? 2) how much data is needed to operate “effectively”? We discuss various central applications in pricing and capacity decisions, together with different associated data structures. We present recent results that enable to quantify (robustly) achievable performance across data sizes, small and big. These results yield fundamental practical insights on the economics of data sizes: in many applications, a little data can go a long way in optimizing decisions.