Inventory control with partial observations and inspections

This project studies optimal inventory control problems whose states are not directly observed. What is observed instead are surrogate measures called signals, which are used to characterize the probability distribution of the state variables. Our project draws upon useful concepts developed for nonlinear filtering and stochastic control to group the problems into two classes. In the first class, we find a finite-dimensional sufficient statistic whose evolution describes the system behavior completely for the purpose of obtaining an optimal solution. Such cases occur for inventory models where the information about the inventory level is delayed, and only the conditional distribution of the inventory level, given the inventory level in an earlier period, is known. In these cases, we show that a reference inventory position can be constructed as a sufficient statistic and that a base-stock policy in terms of this position is optimal. In other words, there exists an order-up-to level such that it is optimal to order the difference when the reference position is below this level and not to order any when the reference position is above the level. In the second class, no finite-dimensional statistic exists. An example of the second class concerns lost sales, where the inventory level is fully observed only when it is zero. Otherwise all we know is that the inventory level is positive. Thus, the signal at any time reveals whether the inventory is zero or positive. When positive, only the conditional distribution of the inventory level given that it is positive is known. Thus, one needs to work with the conditional probability, which in general is an infinite-dimensional entity. In addition, the conditional probability evolves in a highly nonlinear fashion. We develop a weighting scheme–similar to the Zakai equation in the filtering literature–that converts the conditional probability into an unnormalized probability which evolves linearly. The resulting linear system facilitates considerably our study of the inventory problem and the associated dynamic programming equations. It enables us to obtain the existence of an optimal feedback inventory ordering policy and its characterization. It helps in developing practical and approximate optimal policies as finite-dimensional functions, and procedures to compute them. We also study inventory problems with important features, such as censored demand (only sales but not the full demand are observed), random spoilage, etc.

We develop new approaches for the analysis of inventory control problems under partial observations in particular and of applications in broader management and engineering contexts with incomplete information in general. Our multi-disciplinary project contributes to optimal decision making and risk management in a variety of systems with partial observations, such as those found in manufacturing, supply chains, machine maintenance, quality control and finance. The partial observations in these systems are often associated with the inventory level, customer demand, machine-tool ware, production yield, and the parameters of the underlying distributions. We should note that while a substantial literature exists in each of these domains, much of it assumes that the system states are completely observed. Thus, the policies and procedures obtained in this project have the potential to improve the inventory control methodologies used in the industry, which would result in cost savings and better customer service. The new understanding and mathematical methods that stem from our project should be valuable contributions to the operations research, management science and finance literatures, and should facilitate the use of such methods for many interesting problems in these areas.