Optimization under Uncertainty: Combining Statistics and Stochastic Optimization

Marco Laumanns[1], Yingdong Lu[2], Mark S. Squillante[2]
[1]IBM Research, Saeumerstrasse 4, Rueschlikon 8803, Switzerland
[2]IBM Research, 1101 Kitchawan Road, New York, NY 10598, USA

Abstract
The proliferation of data, advances in statistical methods, and growth of computing power create a wealth of opportunities for optimization and optimal control under uncertainty to address large-scale problems involving stochastic systems. We consider a general approach, building on data and on statistics and moving from stochastic models of uncertainty to stochastic optimization/control, that provides a mathematical foundation for the optimal design and control of complex stochastic systems. A couple of examples serve to illustrate different instances of our overall approach, each involving general classes of resource allocation problems with broad applications. In one case, we devise a simulation-based framework that yields optimal resource allocations in an efficient and effective manner, whereas in another case we derive an optimal control policy that includes easily and efficiently implementable algorithms for governing dynamic resource allocations over time. Computational experiments demonstrate the significant benefits of our approach of combining statistics and stochastic optimization/control over alternative methods previously published in the research literature.

Key Words: Predictive analytics; statistical analysis; learning; stochastic optimization; stochastic control; decision making and reasoning under uncertainty.

1. Introduction

Uncertainties abound throughout our world and unexpected things happen all the time. Nevertheless, we have to deal with uncertainty and risks in making the best decisions and defining the best strategies we can in every aspect of our lives, despite our limited ability to understand, learn, reason, optimize and control around these uncertainties.

Critical decisions and strategies are often based on intuition, and even guessing, notwithstanding the deep complexities of uncertainty, risks and associated trade-offs. The significant lack of understanding and the underestimation of these complexities, however, result in poor decisions and strategies that are far from optimal, far from the best management of risks, and far from being robust.

Even when analytics are used in decision making processes, the approach is more often based on averages, point predictions, worst-case scenarios, rules of thumb, or even ignoring uncertainties altogether. Such approaches are woefully inadequate to address the deep complexities of uncertainty, risks and associated trade-offs. They fail to take advantage of everything available and possible within the field of mathematics.

The proliferation of various sources and types of data, the significant advances in statistics and predictive analytics, and the considerable growth in computing power now create completely new opportunities to improve decision making and reasoning through mathematical optimization under uncertainty. At the same time, these tremendous opportunities bring even greater and deeply complex challenges.

In this paper, we consider aspects of mathematical optimization under uncertainty from the perspective of combining statistics and stochastic optimization. The ultimate goal is to provide better decisions and strategies that fundamentally transform business, society and many aspects of our world by improving the management of uncertainty, risks and volatility and reducing the negative impact of such phenomena.

The remainder of this paper is organized as follows. Section 2 provides additional motivation for opportunities to combine statistics and stochastic optimization, together with a high-level description of a general framework for mathematical optimization under uncertainty. Sections 3 and 4 then respectively present two examples of our general framework for optimization and optimal control under uncertainty. We close with some concluding remarks.

2. Opportunities for Combining Statistics and Stochastic Optimization

From a business perspective, many industries today exhibit very high degrees of uncertainty, risk and volatility arising from multiple diverse sources, as considered and discussed in the recent Harvard Business Review report [DFL, 2014] on the industries plagued by the most uncertainty. The report identifies numerous industries with very high degrees of uncertainty across two dimensions: so-called demand uncertainty (arising from the unknowns associated with customer demand for products/services) and so-called technology uncertainty (arising from the unknowns associated with technology and the ability to deliver desirable solutions). This creates a wealth of opportunities for the optimization and control of decision making processes under uncertainty and risks across a broad spectrum of industries and time horizons.

From a technology perspective, some advanced analytics techniques are reaching a level of maturation whereas others are only starting to form the next disruptive technological wave, as considered and discussed in the recent Gartner report [L, 2015] on the so-called hype cycle for advanced analytics and data science. The report identifies predictive analytics and machine learning technologies as having past the peak of their expectation and reaching a plateau of productivity within a few years, while identifying prescriptive analytics and optimization technologies to be in their innovation phase with benefits expected to be reached within ten years. This creates a wealth of opportunities for combining statistics and stochastic optimization/control to significantly improve decision making processes under uncertainty and risks.

The time is now ripe for a new wave of mathematical optimization under uncertainty technologies to create solutions that optimize and control decision making processes under uncertainty and risks at scale across industries, building on the proliferation of data, the advances in statistics and data analytics, and the growth in computational power. To realize this goal, our approach includes a general framework for combining statistics and stochastic optimization that comprises two primary elements which are intimately related. The first concerns mathematical (probabilistic) models of how the system of interest and its associated decision making processes behave over the time horizon of interest, which are built upon the results of statistical analysis of various data sets to understand and characterize extrinsic sources of uncertainty (e.g., errors in data and noise in measurements) and inherent sources of uncertainty (e.g., demand, weather, economic conditions and success/failure risks). The second element concerns the mathematical (stochastic) formulation of the corresponding optimization or optimal control problem and the mathematical methods for solving this problem to determine a set of decisions that gives rise to the best possible rewards over the time horizon within the context of the probabilistic system models and subject to various constraints.

The next two sections provide specific examples of our general mathematical optimization under uncertainty framework that combines statistics and stochastic optimization.

3. Simulation-Optimization of Stochastic Networks

One example, from the work of Dieker et al. [DGS, 2016], considers capacity planning of different types of resources that are connected and interact through a network structure given uncertainty around the demand for each path of the network of resources. In the context of cloud computing applications, the resources represent different computing servers, the stochastic network represents the topology of the cloud computing infrastructure together with the uncertainty of both the customer demand at each station of the network and the overall demand, and the performance metrics are related to the time for completing customer requirements; whereas in the context of business process management, the resources represent activities performed within an organization, the stochastic network represents the series of such activities to be performed to achieve a common business goal together with the uncertainty of both the activity demand at each station and the overall demand, and the performance metrics are related to the time for completing business process operations; refer to [DGS, 2016].

The first key aspect of this example concerns the mathematical model of the system of interest. Since analytical solutions of stochastic (queueing) networks are mathematically intractable in general, simulation methods are often used to estimate system performance metrics as a function of the set of resource capacities. More specifically, for a given topology of the stations composing the stochastic network, let $Z_i^{\boldsymbol{\beta}}$ denote the steady-state queue length at the i-th station (which can be alternatively replaced by steady-state sojourn times), where $\boldsymbol{\beta}$ represent the vector of service rates for all stations. The dependence of each queue length on $\boldsymbol{\beta}$ is made explicit since we are interested in comparing a functional of the steady-state vector $\boldsymbol{Z}^{\boldsymbol{\beta}}$ as we change the service-rate vector $\boldsymbol{\beta}$. Further let $\boldsymbol{\gamma}$ represent the vector of effective arrival rates for all stations. (Additional parameters of the stochastic network model, such as the routing matrix and the external interarrival and service distributions, need not be specified to present our approach and thus we do not introduce them here.)

The second key aspect of this optimization under uncertainty example concerns the mathematical optimization of a simulation model of stochastic networks to determine the capacity of every station that gives rise to the best possible rewards over the time horizon subject to various constraints. We consider a long-run time horizon modeled as a stochastic network of multiple stations serving customers of a single class under uncertain demand for each path through the network, where simulation is used to determine steady-state performance metrics such as $\boldsymbol{Z}^{\boldsymbol{\beta}}$. We formalize our approach in a setting where the goal is to minimize the sum of the weighted expected steady-state queue lengths subject to a budgetary cost constraint (or, alternatively, the dual formulation where the aim is to optimize the expectation of financial metrics subject to a bound on the sum of the weighted expected steady-state queue lengths). Let c_i denote the cost for each unit of resource capacity at station i, with cost vector $\boldsymbol{c}$, and let C denote the total budget for allocating resources in the network.

The following formulation of a corresponding stochastic optimization problem then can be expressed as

$$(OPT) \qquad \begin{array}{ll} \underset{\boldsymbol{\beta}\in(0,\infty)^L}{\text{minimize}} & \sum_{i=1}^{L} w_i \mathbb{E}Z_i^{\beta} \\ subject\ to & \langle \boldsymbol{c}, \boldsymbol{\beta} \rangle \le C \\ & \beta_i > \gamma_i, \qquad i = 1, \dots, L. \end{array}$$

In other words, we seek to minimize the expected steady-state queue lengths weighted by a vector w, subject to the constraints that the total costs cannot exceed the budget C and the queueing system is stable. We shall assume $\langle c, \beta \rangle < C$ throughout, so that the above optimization problem is feasible, where one can (correctly) expect the solution to satisfy $\langle c, \beta \rangle = C$.

Defining $\tau_i(\boldsymbol{\beta}) := (\beta_i - \gamma_i)\mathbb{E}Z_i^{\beta}$, the objective function of (OPT) then takes the form $t\big(\boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{\tau}(\boldsymbol{\beta})\big)$ for some function $\tau(\boldsymbol{\beta})$ where, for $\boldsymbol{\beta} - \boldsymbol{\gamma}, \boldsymbol{w}, \boldsymbol{\tau} > 0$, we have

$$t(\boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{\tau}) = \sum_{k=1}^{L} w_k \frac{\tau_k}{\beta_k - \gamma_k}. \qquad (3.1)$$

For any station in a single-class product-form network, τ is known to be equal to γ and 1 for expected queue length and sojourn time, respectively. Furthermore, τ_k correspondingly equals $\frac{\gamma(c_A^2 + c_S^2)}{2}$ and $\frac{(c_A^2 + c_S^2)}{2}$ in a single-class Brownian product-form network of *GI/GI/1* queues, where c_A^2 and c_S^2 denote the second-order variation terms for the arrival and service process, respectively. However, in general stochastic networks, $\boldsymbol{\tau}(\boldsymbol{\beta})$ is mathematically intractable. The explicit incorporation of $\beta_i - \gamma_i$ in the denominator is motivated by the above product-form results and as a nearly universal phenomenon in stochastic networks under a wide range of queueing dynamics.

By applying standard Lagrangian methods, the minimum of $t(\boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{\tau})$ over the feasible region of (OPT) is readily proven to be $\boldsymbol{\beta}^*(\boldsymbol{w}, \boldsymbol{\tau})$ where, for $\ell = 1, \dots, L$,

$$\beta_\ell^*(\boldsymbol{w}, \boldsymbol{\tau}) = \gamma_\ell + (C - \langle \boldsymbol{c}, \boldsymbol{\gamma} \rangle) \frac{\sqrt{w_\ell \tau_\ell / c_\ell}}{\sum_{i=1}^{L} \sqrt{w_i \tau_i c_i}}.$$

This then becomes an essential ingredient in our analysis. More specifically, we determine the capacity allocation $\boldsymbol{\beta}^*$ through the system of nonlinear equations

$$\beta_\ell^* = \gamma_\ell + (\mathrm{C} - \langle \mathbf{c}, \boldsymbol{\gamma} \rangle) \frac{\sqrt{w_\ell \tau_\ell(\boldsymbol{\beta}^*) / c_\ell}}{\sum_{i=1}^{L} \sqrt{w_i \tau_i(\boldsymbol{\beta}^*) c_i}},$$

for $\ell = 1, \dots, L$, which is guaranteed to have a unique solution for a certain precisely defined class of stochastic networks.

To numerically find a vector $\boldsymbol{\beta}^*$ that satisfies the above equation, assuming existence, one can use the fixed-point iteration scheme with iterates $\{\boldsymbol{\beta}^{(k)}: k \geq 0\}$ given by

$$\beta_\ell^{(k+1)} = \gamma_\ell + (\mathrm{C} - \langle \mathbf{c}, \boldsymbol{\gamma} \rangle) \frac{\sqrt{w_\ell \tau_\ell(\boldsymbol{\beta}^{(k)})/c_\ell}}{\sum_{i=1}^{L} \sqrt{w_i \tau_i(\boldsymbol{\beta}^{(k)}) c_i}},$$

which implies

$$\frac{\beta_i^{(k+1)} - \gamma_i}{\beta_j^{(k+1)} - \gamma_j} = \sqrt{\frac{\beta_i^{(k)} - \gamma_i}{\beta_j^{(k)} - \gamma_j} \times \frac{w_i \mathbb{E} Z_i^{\boldsymbol{\beta}^{(k)}}/c_i}{w_j \mathbb{E} Z_j^{\boldsymbol{\beta}^{(k)}}/c_j}}. \qquad (3.2)$$

Equation (3.2) establishes an important connection with a resource capacity iteration algorithm based on observed queue-length information. Since we must allocate capacity of at least γ_i to station i, the ratio $\frac{\beta_i - \gamma_i}{\beta_j - \gamma_j}$ reflects the "additional" resource capacities allocated to stations i and j, respectively. This ratio is expressed by (3.2) in terms of the ratio of mean queue lengths, so that more capacity is allocated in the next iteration to stations with disproportionally long queue lengths in the current iteration. The right-hand side of (3.2) can be interpreted as the geometric mean of two fractions, which arises from the assumed functional form (3.1). The effect of building in the asymptote into our algorithm is that the iterates avoid the boundary of the feasible region.

The above iterative algorithm based on (3.1) and (3.2) is shown to converge to a unique limit point [DGS, 2016], which is close to optimal in several network settings as indicated by computational experiments. These limit points can be used directly for basic searches of the entire design space. More generally, our approach leverages such limit points as nearly optimal starting points as input to a more standard stochastic gradient descent approach for very quick convergence to true optimal solutions. Computational experiments further demonstrate that our solution framework provides significant reductions in computation over a purely standard simulation-optimization approach based on stochastic gradient descent, including orders of magnitude computational reductions in several problem instances. Refer to [DGS, 2016] for additional details.

4. Stochastic Optimal Control of Dynamic Resource Allocation

Another example, from the work of Gao et al. [GLSSB, 2013], considers stochastic optimal control of the dynamic allocation of different types of resources in order to best serve uncertain demand so that expected net-benefit is maximized over a time horizon based on the rewards and costs associated with the different resource types. In the context of workforce management applications, the different resource types can represent different human capital sourcing options (e.g., internal, business partner, contractor) that are allocated to satisfy the time-varying and uncertain demand, the rewards can represent revenue and quality of service, and the costs can represent salary and other compensation together with resource allocation adjustment penalties; whereas in the context of energy-aware computing environments, the different resource types can represent computing servers operating at different levels that are allocated to satisfy time-varying and uncertain demand, the rewards can be related to the performance properties of the server allocation levels, and the costs can be related to the energy consumption properties of the server allocations together with resource allocation adjustment penalties; refer to [GLSSB, 2013].

The first key aspect of this example concerns the mathematical model of the system of interest. Here we use a stochastic control framework to model the system dynamics of the general resource allocation problem. To simplify the exposition, consider the stochastic model of satisfying demand over time with two types of resources, namely a primary resource allocation option that has the highest net-benefit and a secondary resource allocation option that has the lowest net-benefit. A dynamic control policy defines at every time $t \in \mathbb{R}_+$ the level of primary resource allocation $P(t)$ and the level of secondary resource allocation $S(t)$ that are used in combination to satisfy the uncertain demand $D(t)$, where $S(t) = [D(t) - P(t)]^+$.

The demand process is defined by the linear diffusion model

$$dD(t) = bdt + \sigma dW(t),$$

where $b \in \mathbb{R}$ is the demand growth/decline rate, $\sigma > 0$ is the demand volatility/variability, and $W(t)$ is a one-dimensional standard Brownian motion, whose sample paths are nondifferentiable [KS, 1991]. The demand process is served by the combination of primary and secondary resource allocations $P(t) + S(t)$.

Let $R_p(t)$ and $C_p(t)$ respectively denote the reward and cost function associated with the primary resource allocation $P(t)$ at time t. The rewards are modeled as linear functions of the primary resource allocation and the demand, whereas the costs are modeled as linear functions of the primary resource allocation. Hence, we have $R_p(t) = r_p \times [P(t) \wedge D(t)]$ and $C_p(t) = c_p \times P(t)$ with reward and cost rates $r_p > 0$ and $c_p > 0$ capturing all per-unit rewards and costs for serving demand with the primary resource allocation and $r_p > c_p$. Observe that the risks associated with the primary resource allocation at time t concern lost reward opportunities whenever $P(t) < D(t)$ on one hand and concern incurred costs whenever $P(t) > D(t)$ on the other hand. Define $N_p(t) \coloneqq R_p(t) - C_p(t)$.

The corresponding reward and cost functions associated with the secondary resource allocation $S(t)$ at time t are given by $R_s(t) = r_s \times [D(t) - P(t)]^+$ and $C_s(t) = c_s \times [D(t) - P(t)]^+$, respectively, with $r_s > 0$ and $c_s > 0$ analogous to r_p and c_p. Observe that the secondary resource allocation at time t is riskless in the sense that rewards and costs are both linear in the resources actually used. Define $N_s(t) \coloneqq R_s(t) - C_s(t)$.

Furthermore, any adjustments to the primary resource allocations have associated costs, where we write $\mathcal{I}_p$ and $\mathcal{D}_p$ to denote the per-unit costs of increasing and decreasing the primary resource allocation $P(t)$, respectively. In other words, $\mathcal{I}_p$ represents the per-unit cost whenever the primary resource allocation is being increased while $\mathcal{D}_p$ represents the per-unit cost whenever the primary resource allocation is being decreased.

The second key aspect of this example of optimal control under uncertainty concerns the stochastic control problems associated with the above system model. This stochastic control problem allows the dynamic control policy to adjust its primary and secondary allocation positions based on the demand realization observed up to the current time, which we call the risk-hedging position of the dynamic control policy. More formally, the decision process $P(t)$ is adapted to the filtration $\mathcal{F}_t$ generated by $\{D(s) : s \leq t\}$. The objective of the optimal dynamic control policy is to maximize the expected discounted net-benefit over an infinite horizon, where net-benefit at time t consists of the difference between the rewards and costs from the primary resource allocation and the secondary resource allocation minus the additional costs for adjustments to the decision process $P(t)$.

Let $\dot{P}(t)$ denote the derivative of $P(t)$ with respect to time. We impose a lower-bound θ_l and an upper-bound θ_u on this decision variable $\dot{P}(t)$ at every time t to capture the inability of the control policy to make unbounded adjustments in the primary resource allocation at any instant in time. Then the formulation of a corresponding stochastic optimal control problem can be expressed as

$$\begin{aligned} &\mathit{maximize} \quad \mathbb{E}\int_0^\infty e^{-\alpha t}[N_p(t)+N_s(t)]\,dt - \mathbb{E}\int_0^\infty e^{-\alpha t}\left[\mathcal{J}_p \cdot I_{\{\dot{P}(t)>0\}}\right]dP(t) \\ &\qquad - \mathbb{E}\int_0^\infty e^{-\alpha t}\left[\mathcal{D}_p \cdot I_{\{\dot{P}(t)<0\}}\right]d\big(-P(t)\big) \qquad (4.1) \\ &\mathit{subject\ to} \quad -\infty < \theta_l \le \dot{P}(t) \le \theta_u < \infty \qquad (4.2) \\ &\qquad dD(t) = b\,dt + \sigma dW(t), \qquad (4.3) \end{aligned}$$

where α is the discount factor and $I_{\{A\}}$ denotes the indicator function returning 1 if A is true and 0 otherwise. Note that the second expectation in (4.1) causes a decrease with rate $\mathcal{J}_p$ in the value of the objective function whenever the control policy increases $P(t)$, and the third expectation in (4.1) causes a decrease with rate $\mathcal{D}_p$ in the value of the objective function whenever the control policy decreases $P(t)$.

Suppose the per-unit cost for decreasing the primary resource allocation is strictly less than the corresponding discounted overage cost and suppose the per-unit cost for increasing the primary resource allocation is strictly less than the corresponding discounted shortage cost. Then, as rigorously established in [GLSSB, 2013], the solution to the stochastic optimal control problem (4.1), (4.2), (4.3) has the following simple form for governing the dynamic adjustments to $P(t)$ over time. Namely, there are two threshold values L and U with $L < U$ such that the optimal dynamic control policy seeks to maintain $X(t) = P(t) - D(t)$ within a risk-hedging interval $[L, U]$ at all times t, taking no action as long as $X(t) \in [L, U]$. Whenever $X(t)$ falls below L, the optimal dynamic control policy pushes toward the risk-hedging interval at the fastest possible rate θ_u, thus increasing the primary resource allocation $P(t)$. Similarly, whenever $X(t)$ exceeds U, the optimal dynamic control policy pushes toward the risk-hedging interval at the fastest possible rate θ_l, thus decreasing the primary resource allocation $P(t)$. Here, the optimal threshold values L and U are uniquely determined by two nonlinear equations that depend upon the parameters of the formulation, including b, σ, r_p, c_p, r_s, c_s, $\mathcal{J}_p$, $\mathcal{D}_p$.

In the special case when the dynamic control policy incurs no costs for making adjustments in the primary resource allocation $P(t)$, the risk-hedging interval $[L, U]$ collapses to a single point δ, for which explicit expressions can be derived, and then the optimal dynamic control policy seeks to maintain $X(t) = P(t) - D(t)$ at the position δ at all time t. Whenever $X(t)$ falls below δ, the optimal dynamic control policy increases the primary resource allocation $P(t)$ toward the critical point at the fastest possible rate θ_u. Similarly, whenever $X(t)$ exceeds δ, the optimal dynamic control policy decreases the primary resource allocation $P(t)$ toward the critical point at the fastest possible rate θ_l. Refer to [GLSSB, 2013] for additional technical details on the above and related results.

Computational experiments demonstrate that our optimal online dynamic control algorithm provides significant benefits in comparison with other dynamic resource allocation schemes proposed in the research literature. This includes outperforming, by as much as a factor of two and more, an optimal offline algorithm that consists of making optimal provisioning decisions in a clairvoyant anticipatory manner based on known average demand within each slot of a discrete-time model where the slot length is chosen to match the timescale at which the system can adjust its resource capacity and so that demand activity within a slot is sufficiently nonnegligible in a statistical sense. Refer to [GLSSB, 2013] for additional details.

4. Conclusions

We have considered in this paper a general approach that combines statistics and stochastic optimization/control to provide a mathematical foundation for the optimal design and control of complex stochastic systems. The two key elements of this framework consist of the stochastic models of the system of interest and the solution methods for the corresponding optimization or optimal control under uncertainty problem. Statistical analysis and learning from various sources of data play crucial roles in helping to understand and characterize the uncertainty, risks and associated trade-offs within the context of the stochastic system models and the stochastic optimization/control solution methods. This was illustrated through two examples of our overall approach, involving general classes of resource allocation problems with broad applications. In the example of simulation-optimization of stochastic networks, statistical analysis and learning techniques are applied to data in order to determine the properties and distributional parameters for the exogenous arrival processes and the service processes at each station, for the (probabilistic) routing matrix of customer flows through the network, and for the financial rewards, costs and penalties. Analogously, in the example of stochastic optimal control of dynamic resource allocation, statistical analysis and learning techniques are applied to data in order to determine the properties and distributional parameters for the exogenous demand process, for the upper and lower bounds on the resource allocation control variable, and for the financial rewards, costs and penalties.

The application of statistical techniques to the relevant available data make it possible to develop high fidelity stochastic models and high fidelity stochastic optimization/control problem formulations. This in turn renders high quality solutions for decision making and reasoning under uncertainty across a wide variety of application domains. By effectively combining statistics and stochastic optimization/control, our general approach provides significant benefits over existing methods (such as those ignoring structural properties and those based on point predictions), which was clearly demonstrated and quantified by the representative computational experiments discussed herein. This includes orders of magnitude computational reductions in comparison with the most well-developed simulation-optimization approach and as much as a factor of two and more improvement in solution quality over a previously published offline-optimal dynamic resource allocation policy.

Acknowledgements

We thank Bob Sutor and Eleni Pratsini for sponsoring the Center for Optimization under Uncertainty Research (COUR, pronounced “kor”) in IBM Research and for all of their strong support. We thank our colleagues and fellow members of the COUR Global

Leadership Team, namely Chid Apte, Bruno Da Costa Flach, Soumyadip Ghosh, Michael Masin, Carlos Raoni De Alencar Mendes, Giovanni Russo, and Stefan Woerner. We also thank our many colleagues and fellow members of COUR throughout IBM.

References

[DGS, 2016] A.B. Dieker, S. Ghosh, M.S. Squillante. Optimal Resource Capacity Management for Stochastic Networks. Operations Research, to appear 2016.

[DFL, 2014] J. Dyer, N. Furr, C. Lefrandt. The Industries Plagued by the Most Uncertainty. Harvard Business Review, September 2014.

[GLSSB, 2013] X. Gao, Y. Lu, M. Sharma, M.S. Squillante, J.W. Bosman. Stochastic Optimal Control for a General Class of Dynamic Resource Allocation Problems. Proceedings of the IFIP W.G. 7.3 Performance Conference, 2013.

[KSh, 1991] I. Karatzas, S. E. Shreve. Brownian Motion and Stochastic Calculus, Second edition. New York: Springer-Verlag. 1991.

[L, 2015] A. Linden. Hype Cycle for Advanced Analytics and Data Science. Gartner Report, July 2015.