2020

Specifying and Validating Probabilistic Inputs for Prescriptive Models of Decision Making over Time

Sarah McAllister Ryan
Iowa State University, smryan@iastate.edu

Specifying and Validating Probabilistic Inputs for Prescriptive Models of Decision Making over Time

Abstract
Optimization models for making decisions over time in uncertain environments rely on probabilistic inputs, such as scenario trees for stochastic mathematical programs. The quality of model outputs, i.e., the solutions obtained, depends on the quality of these inputs. However, solution quality is rarely assessed in a rigorous way. The connection between validation of model inputs and quality of the resulting solution is not immediate. This chapter discusses some efforts to formulate realistic probabilistic inputs and subsequently validate them in terms of the quality of solutions they produce. These include formulating probabilistic models based on statistical descriptions understandable to decision makers; conducting statistical tests to assess the validity of stochastic process models and their discretization; and conducting re-enactments to assess the quality of the formulation in terms of solution performance against observational data. Studies of long-term capacity expansion in service industries, including electric power, and short-term scheduling of thermal electricity generating units provide motivation and illustrations. The chapter concludes with directions for future research.

Disciplines
Operational Research | Systems Engineering

Comments
This is a post-peer-review, pre-copyedit version of a book chapter published as Ryan, Sarah McAllister. "Specifying and Validating Probabilistic Inputs for Prescriptive Models of Decision Making over Time." In: Smith A. (ed.) *Women in Industrial and Systems Engineering*. *Women in Engineering and Science*. Springer, Cham (2020): 49-69. The final authenticated version is available online at DOI: 10.1007/978-3-030-11866-2_3. Posted with permission.

Noname manuscript No.
(will be inserted by the editor)

Specifying and validating probabilistic inputs for prescriptive models of decision making over time

Sarah McAllister Ryan

April 2018

Abstract Optimization models for making decisions over time in uncertain environments rely on probabilistic inputs, such as scenario trees for stochastic mathematical programs. The quality of model outputs; i.e., the solutions obtained, depends on the quality of these inputs. However, solution quality is rarely assessed in a rigorous way. The connection between validation of model inputs and quality of the resulting solution is not immediate. This chapter discusses some efforts to formulate realistic probabilistic inputs and subsequently validate them in terms of the quality of solutions they produce. These include formulating probabilistic models based on statistical descriptions understandable to decision makers; conducting statistical tests to assess the validity of stochastic process models and their discretization; and conducting re-enactments to assess the quality of the formulation in terms of solution performance against observational data. Studies of long-term capacity expansion in service industries, including electric power, and short-term scheduling of thermal electricity generating units provide motivation and illustrations. The chapter concludes with directions for future research.

1 Introduction

Each day, a flowergirl must decide on a quantity of fresh flowers to purchase at the wholesale cost before finding out how many she is able to sell at the retail price that day. Unlike her brother, the newsboy, she is able to hold onto some

This chapter is dedicated to the memory of my mother, Janice Crawford McAllister (1929-2017).

S. Ryan
Iowa State University
Tel.: +1-515-294-4347
Fax: +1-515-294-3524
E-mail: smryan@iastate.edu

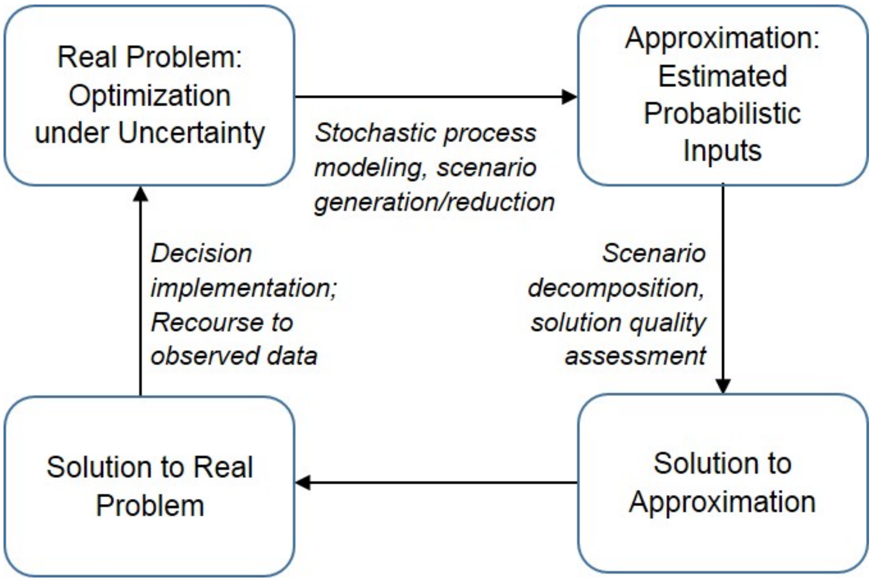

Fig. 1 Stochastic modeling and optimization process.

fraction of unsold inventory – flowers that remain fresh enough to sell on the following day. Her daily problem is to choose a purchase quantity to maximize profit over a sequence of days. If she buys too many she wastes money on flowers that cannot be sold. If she buys too few then she either incurs an opportunity cost of uncollected revenue (Casimir, 1990) or is forced to pay the retail price to make up the difference (Pflug and Pichler, 2014). Having operated this business for a while, she has data on past demand that she can combine with knowledge about future events such as holidays to estimate a joint probability distribution for demand in the days ahead. The ability to carry inventory means that, unlike the newsboy, she cannot simply compute and purchase a critical fractile of the demand distribution for a single day.[1]

A rolling strategy for solving problems such as the flowergirl's is depicted in Figure 1. At each time, t, given a decision problem with uncertain parameters, we fit a stochastic process model using the data available at that time. Next, we approximate the space of possible realizations, keeping solution procedures in mind, and solve the approximated problem. The solution procedure itself may introduce further levels of approximation. Finally, we implement the decisions that must be taken at time t, roll forward to time $t+1$, and repeat the process.

Errors can enter the process at any step – in fact, they are deliberately introduced in the form of approximations employed for computational tractability. For the decision maker to accept the sequence of decisions suggested, she must be persuaded that the whole process of modeling, approximation, and

[1] Generally I prefer the gender-neutral term "newsvendor," but in this book chapter I wish to emphasize that the professional problem solved by the woman may be more complicated than the one her male counterpart faces!

solution is sound. Given that, in an uncertain world, more or less sound decisions can be followed by either good or bad consequences, the best decision justification we can offer is that we continually did the best we could with the information we had available. How can we support this claim with data?

The following sections describe research on various aspects of this rolling process for decision making under uncertainty conducted over the past three decades. Recurring themes that pervade this work include the use of data-driven methods to specify instances, an emphasis on optimizing the initial one of a sequence of decisions, and a reliance on approximation and solution methods that decompose the stochastic optimization problem according to the possible future realizations of the uncertain parameters. A section is devoted to each step of the rolling process, culminating in some recent efforts to comprehensively assess the whole cycle. Forecasters test the quality of outputs by *backcasting*; i.e., testing how accurately their methods would have predicted values in a historical dataset, while financial traders test trading strategies by *backtesting* them on historical datasets. Similarly, I argue that the process depicted in Figure 1 should be assessed by *re-enactment* over a sequence of historical instances. This process is distinguished from simulation by the use of actual obervations rather than randomly generated data. The chapter concludes with research needs.

1.1 Recurring Applications

Much of the work described in this chapter has focused on two types of resource management problems, one with a long term orientation and the other having a short term perspective. Planning for the long term inherently involves uncertainty because of the difficulty of forecasting demands and costs in advance. Uncertainty is also present in short-term scheduling of assets when they rely on variable inputs while demand is governed by both physical processes and consumer decisions. Both applications concern the provision of services where the impracticality of either storing inventory or delaying delivery mandates that demand is satisfied when it is experienced.

Capacity expansion: Planning capacity additions to meet a growing demand for service is challenging because the rate of future demand growth is difficult to predict, while the facilities that can be employed to meet the demand typically are not continuously expandable. Capacity comes in chunks because of physical constraints and/or scale economies. Different types of facilities may exist with various combinations of investment levels, operational costs, and lead times for building or installing them. Demand may vary continously over time and be spatially distributed with physical limits on transportation from the facilities to the demand locations. The time value of money affects the relationship between immediate and future costs. Ryan et al (2011) summarized challenges of resource planning in the electric power industry.

Unit commitment: In electric power systems, thermal generating units are subject to operating constraints and cost characteristics that limit their flexibility to change production levels. Unit commitment is the problem of scheduling which units will run at each future time point in view of these considerations, which take the form of minimum up- and downtime constraints; limits on how quickly units can start up or shut down and how fast production can ramp up or down when they are running; and fixed costs for producing at any positive level, which result in minimum economically feasible production levels in addition to upper limits imposed by their capacities. Different types of units vary in the severity of these constraints as well as in their variable production costs, which mostly depend on fuel cost. The deepening penetration of variable renewable generation, such as wind and solar energy, has increased the uncertainty in the net demand for electricity that thermal units are required to supply. At the same time, wholesale market rules impose strict limits on the amount of time available for optimizing the unit commitment schedule. The challenge is to build a credible model of uncertainty and produce high-quality solution within the time limits.

2 Stochastic Process Modeling

Ryan (1988) used an indirect approach to addressing uncertainty in optimization problems, such as production planning or capacity expansion, formulated over indefinite time horizons. This work began with a deterministic formulation of a dynamic optimization problem over an infinite time horizon and used the concepts of a *forecast horizon* and a *solution horizon*. A forecast horizon is a time horizon sufficiently long that any information pertaining to events after it ends has no impact on the decision that is optimal to implement immediately. Under the assumption of a unique optimal unitial decision, a forecast horizon was defined in Ryan and Bean (1989) as a time horizon length such that, if the problem is solved over a finite horizon at least that long, the (unique) optimal decision will coincide with the optimal initial decision for the infinite horizon problem. Relaxing the uniqueness assumption, Ryan et al (1992) defined a (general) solution horizon as a time horizon long enough that, if the problem is solved over a finite horizon at least that long, *any* optimal initial decision identified is guaranteed to be optimal for the infinite horizon problem. The underlying assumption in this work is that events can be forecast with reasonable accuracy over some initial time period. The question was, how long must that initial time period be for a decision maker to confidently implement an initial decision? If future costs are discounted sufficiently then future uncertainties diminish in importance.

Ryan et al (1992) continued in this vein by developing a method for breaking ties among alternative optimal solutions so that a solution horizon could be identified. This paper included a numerical study of capacity expansion assuming a finite set of facility types characterized by their capacities and fixed costs, as well as a known function for the cumulative demand for capacity over time, eventually bounded by an exponential function with a rate less than the interest rate used to continuously discount future costs. Because facility costs exhibited economies of scale, eventually a “turnpike” facility, offering the most cost-efficient capacity growth per unit time, would be adopted. However, finding the initial sequence of optimal installations before this turnpike policy took effect required optimization. The numerical results indicated that solution horizons could be substantially shorter than forecast horizons, so that an optimal initial decision could be identified easily even under almost complete uncertainty about demand growth in the long run.

Experience with forecasting demand growth for use in expansion planning in the utilities division of a large chemical manufacturer prompted me to consider how to model uncertainty in a realistic way that managers would appreciate. This division was responsible for producing steam used for process heat, to run mechanical equipment, and to cogenerate electric power. While electricity generation was supplemented by outside purchases, steam production was entirely internal to the plant. Because having insufficient steam pressure could result in product quality degradation or even force a partial plant shutdown, planning sufficient boiler capacity was critical. The existing planning process was to annually generate a five year forecast for demand growth. To protect against forecast errors and allow for the lead time required to procure and install equipment, a fixed margin was added to the forecast. The need to expand capacity was signaled if the augmented demand growth was projected to equal or exceed existing capacity within this five-year planning horizon.

Two major features appeared to be important to include in a model for uncertain demand growth for use in such an environment. One was to explicitly represent forecast revisions that occurred in response to demand data observed each year. The other was to replace the fixed margin with a probabilistic envelope around the forecast that would reflect the increase in uncertainty associated with forecasts of more remote time periods. The utility managers were comfortable with statistical confidence intervals and had retained historical records of monthly steam usage. Some questions that arose were how often to update the forecast and what would be a suitable confidence level to use in constructing prediction intervals. These two model parameters are related because less frequent forecast revisions could decrease the accuracy of the prediction limits, while requiring a higher confidence level would magnify the effect of inaccurate point forecasts. Ryan (1998) describes an empirical study in which a time series model was fit to the historical data and the process of repeatedly generating forecasts with the latest available data was *re-enacted*, augmented by either a fixed or a probabilistic margin. Upon each forecast revision, the optimization problem to determine the optimal timing of additional boiler installation was solved using the augmented demand and implementation of any capacity increment to occur before the next roll forward was recorded. The resulting re-enacted capacity expansion policies were compared according to their combinations of total discounted cost and measures of insufficient capacity. Based on an efficient frontier constructed with these two performance measures, the value of more frequent forecasts became apparent while cost-risk profiles of the different capacity margins could be assessed. By assigning a penalty value to capacity shortages, McAllister and Ryan (2000) used first-order stochastic dominance in an expanded simulation study to select the best combination of forecast frequency and capacity margin.

Although the use of fixed or probabilistic capacity margins was partially motivated by the lead times required to expand capacity, those lead times had not been modeled explicitly in optimization models, going all the way back to classical work by Manne (1961) and Freidenfelds (1981). The model formulated in Ryan (2003) considered them as fixed constants while representing demand according to a time series model with parameters that could be estimated from data. The choice of an integrated moving average model allowed different aspects of the uncertainty in demand; namely, its autocorrelation which results in nonstationary expected growth and its random variation about the expectation, to be isolated. The presence of lead times suggested that expansions should be based on the capacity *position*, similar to the inventory position commonly used in inventory models, that includes not only existing capacity but any capacity in the process of being added. The optimality of an expansion timing policy based on the proximity of uncertainty-inflated demand to this capacity position could then be proved. An approximation for the optimal expansion size was found by adapting a continuous-time optimal expansion policy to the discrete time setting compatible with the demand growth model. Simulation studies revealed the effects of autocorrelation and randomness in the demand growth on the threshold of excess capacity position that would optimally trigger an expansion. The main conclusion was that failing to account for autocorrelation (or nonstationarity) in the demand growth model could lead to overestimating the randomness and expanding capacity too early, resulting in higher than necessary discounted costs.

To explore the effect of expansion lead times analytically, Ryan (2004) modeled demand growth as following a geometric Brownian motion (GBM) process, in line with earlier work by Manne (1961); Bean et al (1992), and showed how to optimize expansions to minimize their expected discounted cost subject to a service level constraint. Meanwhile, real options models had been rapidly increasing in popularity as a way to assess the value of the flexibility provided by some investment alternatives to respond to uncertain future events. Motivated by the success of the Black-Scholes formula for the value of a European call option on a stock along with analogies between financial options and operational flexibility, many authors formulated models for investment or even operational decision problems that relied on an explicit or implicit assumption that the value of some "underlying asset" would follow a GBM process. In the engineering economic analysis literature, some examples of such GBM-following variables cited by Marathe and Ryan (2005) included both sales volume and price of a product; internal production, outsourcing and delivery costs of a product; prices of commodities derived from natural resources; present values of cash flows pertaining to equipment operation; and various other physical asset values. Such assumptions did not frequently appear to be based on any analysis of data. Relying heavily on Ross (1999), we proposed simple statistical tests of historical time series data that would verify the fundamental assumptions of the GBM model and allow for reliable estimates of the model's parameters. Applying them to historical data concerning electricity consumption, airline passenger enplanement, revenue from cell phone subscriptions and number of Internet hosts, we found that the data were consistent with the GBM assumptions in the first two instances but not the last two. It would be interesting to re-examine updated data concerning demand for capacity in the two then-nascent industries for which we found the GBM model to not fit well. Meanwhile, Marathe and Ryan (2009) employed formulas for pricing exotic options to evaluate the potential for shortage during the lead time required to add capacity, assuming GBM demand growth.

The capacity expansion studies described above employed stochastic process models directly in continuous-time dynamic programming problem formulations. The formulations were simple enough that at least some aspects of the form of an optimal policy could be derived analytically and only modest computation was necessary to find optimal solutions. Sample paths of the stochastic process models were generated only for the purpose of simulating or re-enacting the process of estimating parameters, constructing forecasts and probability limits, and computing the corresponding decision sequences. In situations where operational costs vary widely according to the investment decisions chosen, stochastic programming models are more suitable. Efforts to discretize stochastic process models in order to generate scenarios are described more in section 3. While the emphasis shifts to finding a relatively small set of scenario paths that well represent the whole space of possible future realizations, it is important to not neglect the identification of an appropriate stochastic process. For example, Jin et al (2011) applied the statistical tests suggested by Marathe and Ryan (2005) to validate the use of GBM models for both demand for electricity and the price of natural gas before applying a moment-matching procedure to generate scenarios for a two stage model of electricity generating capacity expansion.

This section concludes with a recent effort to build a stochastic process model that can be used to generate probabilistic scenarios for short-term planning. Feng and Ryan (2016) combined various methods including a functional regression method based on epi-splines (Royset and Wets, 2014) to develop a model of demand for electricity based on a day-ahead weather forecast while capturing typical temporal patterns and accounting for seasonal and geographic information. While the accuracy of the forecast can be assessed according to the usual measures of mean squared error and mean absolute percentage error, the shape of the distribution of forecast errors plays an important role in generating probabilistic scenarios based on the model. Our approach resulted in both tighter and less skewed error distributions than commonly used benchmark models.

3 Discretization

Once a stochastic process model for the evolution of uncertain parameters is identified, the next step is to find a tractable representation of it for use in optimzation. With a few exceptions, such as the infinite horizon generic capacity expansion models described in Section 2, a continuous-time and -state stochastic process model cannot be used directly in optimization. The process of formulating a discrete set of future realizations, with associated probabilities, has been investigated under the label of scenario generation. One popular approach is to randomly generate a large collection of sample paths and then apply so-called scenario reduction procedures to identify a representative subset. While some methods for stochastic optimization embed the scenario generation or sampling process within the optimization procedure, I focus attention on methods where the representation of uncertainty is completed before the solution procedure commences.

Study of a medium-term energy planning problem sparked my interest in this issue. Quelhas et al (2007) had formulated a multiperiod generalized network flow model for bulk energy flows in the US and Quelhas and McCalley (2007) had validated it against actual utilization of different primary energy sources to meet the demand for electric energy over a year. While coal prices were quite stable, the volatility in the prices of natural gas and crude oil made the assumption of deterministic fuel costs seem unrealistic. In the first few years of the twenty-first century, natural gas generation had grown to account for a significant share of electricity generation in the US because of the relative flexibility and lower emissions of gas-fired generating units compared to coal-fired ones. However, before innovations in shale gas extraction took hold, the price of natural gas was generally increasing with considerable volatility from year to year. Our goal in Wang and Ryan (2010) was to add a representation of the fuel cost uncertainty to the network flow model and investigate the impact of this uncertainty on resource utilization decisions. Because the model of Quelhas et al (2007) was formulated on a discrete-time basis, a natural approach was to form the deterministic equivalent of a two-stage stochastic program where flows for one period composed the first-stage decisions and flows for later periods composed the recourse decisions that could be delayed until after the fuel prices for those periods were realized. We adopted a receding horizon approach to simulate the process of monthly decision making with updates on the fuel price forecasts. We used just three possible values of natural gas price in each month, corresponding to the point forecast and its lower and upper confidence limits according to forecasts published by the US Department of Energy. Even so, the assumption of independence between periods resulted in a large number of scenario time series. When they were combined with the thousands of nodes and arcs in the original deterministic formulation, the extensive form of the deterministic equivalent became prohibitively large.

Several approaches exist for managing the computational difficulties associated with solving with the large-scale deterministic equivalent, and most realistic applications require a combination of them. One is to apply decomposition methods, such as those based on scenario decomposition as described in Section 4 or Benders decomposition as applied by Wang and Ryan (2010). Another is to limit the number of scenarios used, as described in Section 3.1. Both of these approaches assume the deterministic formulation is fixed and approximate either the joint probability distribution of the uncertain parameters or the optimal solution of the problem based on a given discrete distribution (i.e., a scenario tree). A third approach, discussed in Section 3.2, is to consider more carefully the relative value of detail in the deterministic formulation as opposed to the representation of uncertainty.

3.1 Scenario reduction

Wang (2010) investigated the existing scenario reduction approaches based on probability metrics (Dupačová et al, 2003; Heitsch and Römisch, 2003) and deemed them unsatisfactory because they operate entirely within the probability space of the stochastic process realizations without considering the optimization context. In fact, they are motivated by results concerning the stability of the optimal first-stage decisions with respect to the discrete approximations they provide of the continuous probability distributions for the uncertain parameters. However, there are two levels of approximation present in these, by now, "classical" methods of scenario reduction. First, proximity of the solution found using the reduced scenario set to the true optimal solution is expressed in terms of an upper bound on the distance in cost, not the distance itself. Second, the optimization problem to find a reduced scenario set that minimizes this upper bound is only approximately solved using fast heuristics such as fast forward selection (Heitsch and Römisch, 2007) – otherwise, the scenario reduction procedure could be less tractable even than the optimization problem it is intended to simplify. To inject some information about the optimization context into the reduction procedure, Wang developed a heuristic approach that employed the forward selection heuristic within clusters of scenarios identified on the basis of their similarity in terms of their optimal first-stage decisions.

Feng and Ryan (2013) elaborated this idea and applied it to the electricity generation expansion planning model of Jin et al (2011). The moment-matching procedure was simplified to take advantage of the stationarity property of the GBM processes. However, even with only two or three branches in the scenario tree each period, the number of scenario paths was too large to allow for solving the extensive form of the deterministic equivalent over a realistic time horizon. As Wang had proposed, we solved the deterministic "wait-and-see" subproblem for each scenario, characterized the optimal decision in terms of a few summary descriptors, and then clustered the scenarios based similarity of these descriptors for the resulting optimal decisions. By applying fast forward selection to choose one scenario from within each cluster, we obtained a reduced set of scenarios that performed similarly to a set of the same size found by applying fast forward selection to the whole set of scenarios. However, the time required for our reduction procedure was much lower and, unlike forward selection, remained approximately constant regardless of the desired number of scenarios in the reduced set.

For stochastic unit commitment, limiting the number of scenarios is critical for obtaining high quality commitment schedules in the limited time allowed by market rules on the day before the target day. Feng et al (2015) combined segmentation of similar days with epi-spline functional regression to develop stochastic process models for the hourly load, incorporating the uncertainty associated with weather forecasts. Rather than generating randomly sampled paths, we carefully constructed probabilistic scenarios by approximation using conditional expectations. Probabilistic scenarios for wind energy generation were obtained from a commercial vendor based on numerical weather prediction models. The net load scenarios, representing possible time series of load less the wind generation amounts, were formed by crossing the two sets of scenarios. Thus, although the sets of scenarios had been carefully constructed to be small, we still ended up with large sets of net load scenarios. To reduce their number, Feng and Ryan (2016) further developed the approach of Feng and Ryan (2013). In this variant, scenarios were clustered based on the major components of the objective function; namely, the production cost and the positive or negative imbalance between energy produced and the net load in each hour. Compared with the unit commitment schedules found by using fast forward selection, those produced by optimization with our reduced scenario set provided more reliable electricity delivery and were more similar to the schedules produced by using the whole set of scenarios.

3.2 Comparative granularity

Quelhas and McCalley (2007) validated their deterministic model by comparing its optimal network flows with the actual amounts of fuel transported and utilized for electricity generation as well as the electricity transmitted among regions in case studies of two separate past years. They attributed differences between the optimal and actual network flows to the spatial and temporal aggregation necessitated by limitations in the available data and the absence of market interactions in the model, as well as the lack of representation of uncertainty and future expectations by decision makers. Wang and Ryan (2010) attempted to represent uncertainty in fuel costs, as well as changing expectations concerning them, by re-enacting the solution of a stochastic program where the scenarios represented both the forecasts and the associated levels of uncertainty, with forecast updates included in the receding horizon procedure. When this was done, the multiperiod flows comprising the sequence of first-stage decisions that would be implemented in the receding horizon procedure were quite similar to the actual decisions that had been made. As we concluded in the paper, "When model validation is unsatisfactory, analysts frequently strive to include more temporal or spatial detail. Our results suggest that incorporating stochastic variability may be another practical way to improve model fidelity, especially when historical forecasts are available but disaggregated temporal and spatial data are not."

Similar issues arise in infrastructure planning, specifically electicity generation and transmission expansion planning. Practitioners advocate a procedure called scenario planning, where they define a *scenario* as a description of possible future conditions under which the infrastructure would be operated, usually at a single future time point. Electricity system resource planners sometimes use the word "future" instead, where a future could describe global system characteristics and policy choices such as degree of penetration of renewable energy; and the presence or absence of carbon emission regulations, large-scale energy storage, and demand response mechanisms. A detailed deterministic operational model is used to optimize investments in infrastructure for each future, with the goal of identifying investment decisions that are common across all futures. Muñoz et al (2014) provide a clear description and critique of this approach, as compared with stochastic programming, in a transmission expansion planning case study for the western US. The weaknesses of scenario planning include the lack of any assessment of the relative likelihood of the futures considered and the possibility that a decision that is optimal for each scenario individually is not optimal when they are considered simultaneously.

However, the intuitive appeal of this approach has led to its widespread adoption and the related assumption that operational models must be sufficiently detailed to accurately assess the value of infrastructure investments. Including a high level of operational detail produces a large scale multiperiod optimization model, with both high-dimensional decision variables, some of which are discrete to represent nonconvexities, and many constraints to capture the details of system operation under temporal variation. As a result, planners are reluctant to consider many different futures or scenarios because simulating operation with each one is so expensive computationally. In such a context, a stochastic program with multiple probabilistic scenarios to be considered simultaneously appears impractical. Jin et al (2014) formulated a stochastic program for thermal generation expansion planning with probabilistic scenarios representing availability of wind power in a typical year. To control the size of the extensive form, we compared the results of different simplifications. One was to decrease the stochastic granularity by reducing the number of wind energy scenarios considered and the other was to decrease the temporal granularity by dropping the nonconvex unit commitment constraints while retaining the continuous ramping restrictions. In case studies comparing the results of both approximations with the full model, we found that the more granular stochastic representation combined with coarse-grained operational constraints resulted in more accurate solutions and more efficient computation than than the coarse-grained stochastic representation combined with highly detailed operational constraints. Accuracy of the solution was judged according to similarity with the solution obtained by solving the full model with high detail in both the stochastic and temporal representations.

4 Solution Methods

Both capacity expansion and unit commitment are naturally formulated as stochastic mixed integer programs (SMIPs) because of the discrete character of the primary decisions. In capacity expansion, increments of capacity typically are not available in continuous sizes because of economies of scale and other design considerations for durable equipment or the construction of major facilities. The decision variables that describe operations may also be discrete because of minimum run-time or production level constraints, discontinuities in marginal cost, or nonlinearities that are approximated as piecewise linear. In unit commitment, binary decision variables are used to express the fundamental on/off decisions as well as nonlinear or nonconvex operational features. In realistically scaled instances, the deterministic subproblem for a single scenario may be challenging to solve in a reasonable amount of time. In both application contexts, considerable research has been devoted to devising reformulations and decomposition methods to solve the deterministic instances efficiently.

Including multiple probabilistic scenarios for parameter values exacerbates the computational challenge and motivates the development of approximate solution methods. Various decomposition methods have been explored including Benders (stage-wise) decomposition and Dantzig-Wolfe decomposition (column generation), as well as Lagrangian relaxation of "complicating constraints." We have focused on scenario decomposition, which can be viewed as relaxation of the nonanticipativity that is expressed either implicitly or explicitly in the formulation of a SMIP. Nonanticipativity is expressed implicitly by formulating the problem in terms of decision stages, where the decision variables in a given stage can depend on realizations of uncertain parameters observed in that stage or earlier, but not on values to be revealed in future stages. In a scenario formulation, all decision variables are scenario-dependent, but explicit nonanticipativity constraints are introduced to force agreement in a given stage for all decision variables corresponding to scenarios that agree up to that stage. When the nonanticipativity constraints are relaxed, the problem decomposes into separate deterministic scenario subproblems that can be solved efficiently using all the solution technology developed for deterministic instances in that application. For example, software for solving unit commitment combines mixed integer programming solvers with specialized constraint management and acceleration techniques such as warm starting.

Scenario decomposition algorithms for solving SMIPs typically produce approximate solutions because exact methods based on branch-and-bound (Carøe and Schultz, 1999) converge too slowly to be practical or because guarantees of convergence to optimality that exist in the continuous case (Rockafellar and Wets, 1991) fail to hold for nonconvex problems. Focusing without loss of generality on cost-minimization problems, lower bounds on the optimal objective function value are essential, either to employ in branch-and-bound algorithms or to assess the quality of a terminal solution. In the scenario decomposition method known as progressive hedging, Gade et al (2016) derived a lower bounding approach using the information available in any iteration of the algorithm and demonstrated its practical use in two-stage stochastic server location as well as stochastic unit commitment. Cheung et al (2015) employed these lower bounds, in stochastic unit commitment instances of the scale typically solved daily by US independent system operators, to demonstrate that parallel progressive hedging could obtain high-quality solutions in a practical length of time. For two-stage SMIPs, Guo et al (2015) exploited the correspondence between this progressive hedging lower bound and one based on Lagrangian relaxation of the nonanticipativity constraint to speed up convergence of the exact branch-and-bound algorithm of Carøe and Schultz (1999). Guo and Ryan (2017) extended the progressive hedging lower bound to certain time-consistent formulations of risk-minimizing multi-stage stochastic programs.

5 Comprehensive Assessment

Following the sequence of activities discussed in the previous three sections, we have

1. formulated a stochastic process model for uncertain parameters in our optimization model, informed by observational data and allowing parameter estimates to be updated as additional data are collected;
2. carefully discretized the models to produce a modest number of probabilistic scenarios, considering tradeoffs the amount of detail included in operational considerations and the granularity of the stochastic discretization; and
3. developed a method to assess the quality of an approximate solution to the resulting stochastic mixed integer program.

Steps 2 and 3 have emphasized the role of scenario subproblems. Scenario reduction methods developed for use in Step 2 employed them to characterize and cluster scenarios in terms of the optimal decisions for the associated deterministic subproblems. The lower bound in Step 3 was developed for solution procedures based on scenario decomposition. This section describes approaches to assess the quality of scenario sets and the solutions obtained by optimizing against them. As in the previous work, we employ scenario decomposition and emphasize the influence of different scenarios on the decisions to be implemented at once. In settings where instances of the same problem are solved repeatedly with continually updated parameter values, we argue that *re-enactment* is an appropriate data-driven approach for assessment and develop computationally efficient shortcuts for it. Here we use the term *scenario generation method* (SGM) to denote “any combination of stochastic process modeling, approximation, sampling and reduction techniques that results in a set of probabilistic scenarios based on the information available at the time [when] the [stochastic program] is to be solved” (Sarı Ay and Ryan, 2018).

5.1 Direct Assessment of Scenario Generation Methods

Before describing methods for assessing scenario generation methods, let us consider some related concepts that have been rigorously defined and tested in the closely related, but not identical, context of probabilistic forecasting. As defined by Gneiting and Katzfuss (2014), "a probabilistic forecast takes the form of a predictive probability distribution over future quantities or events of interest." A probabilistic forecast is called *calibrated*, or equivalently, *reliable* if the probabilities associated with predicted values correspond closely to their observed frequencies. The goal for a probabilistic forecaster is to produce predictive distributions that are as concentrated, i.e., *sharp* as possible, subject to reliability. The combination of reliability and sharpness is called *skill* (Pinson and Girard, 2012). Precise definitions and metrics for these and other desirable characteristics of probabilistic forecasts of scalar quantities have been developed. Various "scoring functions," which measure the distance between a probabilistic forecast and the observed value, are used to compare the predictive performance of competing forecasting methods. Although the observed value could be viewed as a random variable with a degenerate distribution, the probability metrics used for scenario reduction are not mentioned in the probabilistic forecast assessment literature. Moreover, as Gneiting and Katzfuss (2014) note, corresponding metrics and scoring functions for assessing probabilistic forecasts of multidimensional quantities (e.g., scenarios for stochastic programs) are lacking. Many of those that exist were developed in the context of weather forecasting where, typically, equally likely sample paths, called ensemble forecasts, are generated by running multiple replications of numerical weather prediction – simulation – models under different conditions or assumptions. Pinson and Girard (2012) applied some statistical metrics for reliability and skill to evaluate equally likely scenario time series for wind energy production over the short term.

It is important here to note a distinction between the so-called "probability metrics" used in scenario reduction in the stochastic programming literature and the "statistical metrics" used in probabilistic forecast verification. For stability of the optimal solution to a stochastic program, the discretized or reduced scenario set should minimize the distance to the "true" distribution in terms of the Wasserstein distance. Given two cumulative distribution functions (CDFs), F and G for a real-valued random variable, the simplest variant of the Wasserstein distance is (Pflug, 2001):

$$d_W(F,G) = \int_{-\infty}^{\infty} |F(u) - G(u)|du, \tag{1}$$

that is, the total absolute deviation between the CDFs. This distance measure is often called the mass transportation or earth mover's distance because, for discrete distributions, it can be computed by solving a linear transportation problem to move the probability mass from one distribution to the other with minimal work (defined as mass times distance). On the other hand, in the nonparametric goodness-of-fit testing literature, the distance between empirical distributions is often measured using the energy distance (Székely and Rizzo, 2013):

$$d_E(F,G) = \int_{-\infty}^{\infty} (F(u) - G(u))^2 du = 2E|X-Y| - E|X-X'| - E|Y-Y'|, \quad (2)$$

where X and X' are independent random variables distributed as F and Y and Y' are independent random variables distributed as G. The name comes from a relation to Newtonian potential energy within a gravitational space. The energy score used to evaluate probabilistic forecasts is based on the energy distance between the forecast CDF and the observation (Gneiting and Raftery, 2007). When probabilistic forecasts and the corresponding observations are available for a collection of historical instances, the skill of the forecasting method can be evaluated in terms of the average energy score over the instances. Both the Wasserstein distance and the energy distance can be computed easily for joint distributions of several discrete variables, such as time series, by solving the corresponding mass transportation problem or evaluating the corresponding expectations as probability-weighted sums.

Another distance-based approach for assessing the reliability of ensemble forecasts of multidimensional quantities, which can be seen as multiple equally likely scenarios, is based on minimum spanning trees (Wilks, 2004). Given a collection of historical instances, the idea is to quantitatively assess the degree to which the observation is indistinguishable from an ensemble member. For each instance $d = 1, \ldots, D$, a complete graph is constructed with nodes for each ensemble member, $s = 1, \ldots, S$, as well as the observation where edge lengths are computed according to a suitable distance measure, usually Euclidean distance. Next, a minimum spanning tree (MST) is constructed to connect all the ensemble members and its total edge length is recorded, say as ℓ_0^d. Then, for each ensemble member $s = 1, \ldots, S$, the observation is substituted for member s and the length of the resulting MST over those S nodes, not including member s, is recorded as ℓ_s^d. The $S + 1$ MST lengths for instance d are sorted in increasing order and the rank of ℓ_0 is recorded as r_d. Finally, a histogram with bins for the possible values $1, \ldots, S + 1$ of the ranks $\{r_d, d = 1, \ldots, D\}$ is constructed and evaluated for uniformity. A flat histogram indicates that the observation is equally likely to fall in the middle of the ensemble or its outer reaches. Overpopulation of the lower-valued bins occurs if the ensemble is either underdispersed or biased because the observation tends to be more distant from the ensemble members than they are from each other. A disproportionate number of higher rank values indicates that the ensemble is overdispersed so that the observation falls in the middle too often. Uniformity of the rank distribution can be quantified using a goodness-of-fit statistic but the graphical histogram is appealing because its shape helps diagnose the nature of the errors in ensemble forecasts (or sets of equally likely scenarios).

When a scenario generation method employs approximation rather than generating sample paths of the stochastic process model, or when scenario reduction methods are used, the resulting scenarios generally are not equally likely. To assess the reliability of unequally likely scenarios, Sarı et al (2016) developed a rank histogram based on the Wasserstein distance. The mass transportation distance (MTD) rank histogram (Sarı and Ryan, 2016) is constructed similarly to the MST rank histogram with the following three differences. First, ℓ_0^d is computed as the minimum cost of transporting the probability mass from the scenarios to the observation. Second, when the observation is substituted for scenario s, it is assigned the probability of that scenario and ℓ_s^d is computed as the minimum cost of transporting all the probability mass, including that mass having been re-assigned to the observation, to scenario s. Finally, MTDs are sorted in decreasing order to find r^d as the rank of ℓ_0^d. In simulation studies, we demonstrated that the MTD rank histogram has a similar shape to the MST histogram under the same conditions of bias, overdispersion or underdispersion. The MTD values can be computed directly (even more efficiently than greedy-algorithm-based MST lengths) as the sum of probability-weighted distances. We applied the MTD rank histogram, as well as energy scores and event-based scores, to assess two different methods for generating wind power scenario time series on the day ahead and found that it could distinguish among scenario sets based on their autocorrelation levels as well as their bias and dispersion.

5.2 Assessing Solutions by Re-enactment

While reliability of scenario sets may be seen as a necessary condition for obtaining good solutions to stochastic programs, it may not be sufficient. In fact, there seem to be few studies that have "closed the loop" and examined how well the solution to a stochastic programming performs in the target context. The stochastic process modeling step can be assessed by comparing sample paths generated by the model to observed realizations, but studies of this type are rarely reported. Scenario reduction procedures, operating entirely in the realm of probability models, aim to approximate a continuous or highly granular discrete model with a coarse-grained discrete one. We return to the idea of re-enactment as a data-driven approach for assessing the quality of solutions obtained by the whole process of formulating a stochastic program, generating scenarios and obtaining approximate solutions.

The term re-enactment has been used recently, to describe a procedure to assess prediction intervals for wind energy generation, as "a walk forward through date-times in the past, computing prediction intervals using only data available prior to that date-time. In doing so, we compute prediction intervals using only relevant historical information, and are able to assess prediction interval quality using actual observations not used in the computation of those prediction intervals" (Nitsche et al, 2017). Staid et al (2017) used a similar procedure to evaluate scenarios for wind power time series in terms of energy scores, MST rank histograms, and other metrics. In the context of stochastic unit commitment, Sarı and Ryan (2017) extended this idea to re-enact the process of not only generating scenarios but also solving the extensive forms of the stochastic programs. For each historical day d, we generated scenarios by competing methods, including some variants, using the data available through day $d-1$, then solved the stochastic program to obtain an optimal commitment schedule, and finally simulated dispatching the committed units to meet the observed net load on day d. We found that the variant of the scenario generation method that would be selected according to energy score, MTD rank histogram and some event-based scores produced the lowest average cost over the set of historical days.

Encouraged by these empirical results but cognizant of the computational burden of repeatedly solving stochastic programs to conduct this type of re-enactment, Sarı Ay and Ryan (2018) proposed solution assessment methods for two-stage stochastic programs (SPs) based on MTD rank histograms of the costs of solutions to scenario subproblems. As described in that paper, "for each [historical] instance, a single-scenario version of the SP is solved to find a candidate first-stage solution. Then, for each scenario as well as the observation, the second-stage solution is optimized assuming the candidate solution has been implemented, and the total cost for the scenario is computed. Reliability assessment is then applied to these costs. Variants of this approach differ according to whether the expected value (EV) scenario, perfect information (PI, i.e., the observation), or a randomly selected (RS) scenario is used to find the candidate solution." The use of a RS scenario is consistent with the notion that members of a reliable scenario set are statistically indistinguishable from the corresponding observations. We simulated this process using synthetic data for stochastic server location as well as stochastic unit commitment instances and then applied it to a case study of stochastic unit commitment with uncertain wind energy production. "Simulation studies demonstrate that reliability of SGMs can be assessed accurately by the EV-based method. The stochastic unit commitment case study indicated that the PI- and RS-based methods can be used to distinguish between higher and lower quality SGMs, as have been identified by re-enactment" (Sarı Ay and Ryan, 2018).

6 Conclusion and Further Research

My current interest in re-enactment as a data-driven strategy for evaluating the entire modeling and solution process depicted in Fig. 1 arose while conducting a project on stochastic unit commitment for the Advanced Research Projects Agency-Energy (ARPA-E) of the US Department of Energy. Because of the funding source, the project emphasized engagement with end users to enable transfer of the technology developed. Our team, which included personnel from two universities, a software developer and a national laboratory in partnership with an independent system operator, readily identified two major barriers to adoption of stochastic optimization by electricity system operators. One was mistrust in the scenario generation process and the other was doubt that high quality solutions could be found within realistic time limits. Some of the research described in Section 3.1 was aimed at overcoming the former barrier while the work outlined in Section 4 addressed the latter. The pair of papers by Feng et al (2015) and Cheung et al (2015) summarize this project's major accomplishments. However, the real test of our project came when we were asked to demonstrate the cost savings that the system operator might enjoy by replacing their current deterministic optimization with our proposed stochastic programming procedure. To estimate them, our team conducted a detailed and careful re-enactment of daily unit commitment over a year's time. For the stochastic programming model, this process included stochastic process modeling and scenario generation using the data available up to the target day, followed by dispatch of the committed units to satsify the observed net load on the target day. The results demonstrating savings of a few percent have been presented at conferences but, unfortunately, not documented in a published paper.

While writing this chapter I was surprised to recall that one of the first papers I independently conceived and wrote (Ryan, 1998) involved a similar process of re-enactment with actual data to test different ways of modeling uncertainty in a capacity expansion problem. About a decade later, after spending some time on queuing models of manufacturing systems, I began working on stochastic mathematical programming and, again, used re-enactment to explore the impact of uncertainty on an optimization model intended to simulate actual decision making (Wang and Ryan, 2010).

For this form of validation of the stochastic modeling and optimization process to be widely accepted and used, it must developed rigorously. For this development, we need an underlying probability model for the observed data that first inform the stochastic modeling process and later are used to evaluate solutions. The detailed re-enactment procedure is predicated on the idea that a higher quality scenario generation method should result in lower re-enacted costs. Because the cumulative cost over the re-enactment period is a random variable that depends on the observed data collected, comparisons of the costs incurred by different modeling and solution approaches can only be claimed in probabilistic terms. Proofs that the faster approach for scenario and solution assessment proposed in Sarı Ay and Ryan (2018) is itself reliable await completion.

In the long run I envision open-source software tools that could streamline the conduct of re-enactment studies. Our R package to compute the MTD rank histogram (Sarı and Ryan, 2016) is a tiny step in this direction. The PySP package in Pyomo (Hart et al, 2017) has helped to structure the way I think about a stochastic program, as a deterministic model accompanied by a scenario tree. This structure facilitates the repeated re-formulation of problem instances that differ only according to the scenarios included, which is also necessary for re-enactment. The easy parallelization and inclusion in Pyomo of an extension that computes the progressive hedging lower bound both facilitate scenario decomposition as a fast and effective strategy for repeatedly solving re-enacted optimization models. Transparent validation methods, made easier by software tools, could expand the use of stochastic optimization and result in better decisions in a world of uncertainty.

Biography

I am the youngest of three daughters born to an accountant who retrained to become a systems analyst and a chemistry major who never worked outside the home after she married. In the 1970s my parents noted the rising divorce rate in the US, along with the high grades my sisters and I achieved in math and science, and concluded that engineering could be a profession that would support us in case our future spouses did not. I followed their urging and my sisters' examples and entered an engineering program. I chose The University of Virginia for its strength in a diversity of fields in case I wanted to switch to a different field. I selected systems engineering in part because of its requirements to take courses in psychology and economics. Then I discovered the field of operations research, where I could combine mathematical analysis with real world considerations. I was fortunate to have both Chip White, who was then department chair, and Carl Harris as mentors who encouraged me to pursue graduate study. I completed my graduate degrees in industrial and operations engineering at The University of Michigan, and have enjoyed faculty positions in industrial engineering at the University of Pittsburgh, the University of Nebraska-Lincoln, and Iowa State University. The research described here has been supported by the National Science Foundation, including a CAREER award, the US Department of Energy, and industry consortia among other sources. I am now a Fellow of the Institute of Industrial & Systems Engineers, the Joseph Walkup Professor of Industrial and Manufacturing Systems Engineering at Iowa State University, and the 2017 recipient of its College of Engineering's D. R. Boylan Eminent Faculty Award for Research. I am deeply grateful to my husband, Steve Ryan, for embracing the role of primary caregiver to our children for many years.

Some say we choose the professional specialization that helps us most personally. I do view life as a series of decisions made under uncertainty. My research and experience have taught me to distinguish what I can control from what I cannot, content myself with having made the best choice I knew with the information available at the time, and eschew regret as a criterion for evaluating past decisions.

References

Bean JC, Higle J, Smith RL (1992) Capacity expansion under stochastic demands. Operations Research 40:S210–S216

Carøe CC, Schultz R (1999) Dual decomposition in stochastic integer programming. Operations Research Letters 24:37–45

Casimir RJ (1990) The newsboy and the flower-girl. OMEGA International Journal of Management Science 18(4):395–398

Cheung K, Gade D, Ryan S, Silva-Monroy C, Watson JP, Wets R, Woodruff D (2015) Toward scalable stochastic unit commitment - part 2: Assessing solver performance. Energy Systems 6(3):417–438, DOI 10.1007/s12667-015-0148-6

Dupačová J, Gröwe-Kuska N, Römisch W (2003) Scenario reduction in stochastic programming. Mathematical Programming 95(3):493–511, DOI 10.1007/s10107-002-0331-0

Feng Y, Ryan SM (2013) Scenario construction and reduction applied to stochastic power generation expansion planning. Computers & Operations Research 40(1):9–23

Feng Y, Ryan SM (2016) Day-ahead hourly electricity load modeling by functional regression. Applied Energy 170:455–465, DOI 10.1016/j.apenergy.2016.02.118

Feng Y, Rios I, Ryan SM, Spurkel K, Watson JP, Wets RJB, Woodruff DL (2015) Toward scalable stochastic unit commitment - part 1: Load scenario generation. Energy Systems DOI 10.1007/s12667-015-0146-8

Freidenfelds J (1981) Capacity Expansion: Analysis of Simple Models with Applications. North-Holland, New York

Gade D, Hackebeil G, Ryan SM, Watson JP, Wets RJB, Woodruff DL (2016) Obtaining lower bounds from the progressive hedging algorithm for stochastic mixed-integer programs. Mathematical Programming, Series B 157(1):47–67, DOI 10.1007/s10107-016-1000-z

Gneiting T, Katzfuss M (2014) Probabilistic forecasting. Annual Review of Statistics and Its Application 1:125–151, DOI 10.1146/annurev-statistics-062713-085831

Gneiting T, Raftery AE (2007) Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association 102(477):359–378, DOI 10.1198/016214506000001437

Guo G, Hackebeil G, Ryan SM, Watson JP, Woodruff DL (2015) Integration of progressive hedging and dual decomposition in stochastic integer programs. Operations Research Letters 43(3):311–316, DOI 10.1016/j.orl.2015.03.008

Guo GC, Ryan SM (2017) Progressive hedging lower bounds for time consistent risk-averse multistage stochastic mixed-integer programs, URL https://works.bepress.com/sarah_m_ryan/93/

Hart WE, Laird CD, Watson JP, Woodruff DL, Hackebeil GA, Nicholson BL, Siirola JD (2017) Pyomo – Optimization Modeling in Python, 2nd edn. Springer

Heitsch H, Römisch W (2003) Scenario reduction algorithms in stochastic programming. Computational Optimization and Applications 24:187–206

Heitsch H, Römisch W (2007) A note on scenario reduction for two-stage stochastic programs. Operations Research Letters 35(6):731–738, DOI 10.1016/j.orl.2006.12.008

Jin S, Ryan S, Watson JP, Woodruff D (2011) Modeling and solving a large-scale generation expansion planning problem under uncertainty. Energy Systems 2:209–242, DOI 10.1007/s12667-011-0042-9, URL http://dx.doi.org/10.1007/s12667-011-0042-9

Jin S, Botterud A, Ryan SM (2014) Temporal vs. stochastic granularity in thermal generation capacity planning with wind power. Power Systems, IEEE Transactions on 29(5):2033–2041, DOI 10.1109/TPWRS.2014.2299760

Manne AS (1961) Capacity expansion and probabilistic growth. Econometrica 29(4):632–649

Marathe R, Ryan SM (2005) On the validity of the geometric brownian motion assumption. The Engineering Economist 50:159–192, DOI 10.1080/00137910590949904

Marathe RR, Ryan SM (2009) Capacity expansion under a service level constraint for uncertain demand with lead times. Naval Research Logistics 56(3):250–263

McAllister C, Ryan SM (2000) Relative risk characteristics of rolling horizon hedging heuristics for capacity expansion. The Engineering Economist 45(2):115–128

Muñoz FD, Hobbs BF, Ho JL, Kasina S (2014) An engineering-economic approach to transmission planning under market and regulatory uncertainties: Wecc case study. Power Systems, IEEE Transactions on 29(1):307–317

Nitsche S, Silva-Monroy C, Staid A, Watson JP, Winner S, Woodruff D (2017) Improving wind power prediction intervals using vendor-supplied probabilistic forecast information. In: IEEE Power & Energy Society General Meeting

Pflug GC (2001) Scenario tree generation for multiperiod financial optimization by optimal discretization. Mathematical Programming Series B 89:251–271, DOI 10.1007/s101070000202

Pflug GC, Pichler A (2014) Multistage Stochastic Optimization. Springer

Pinson P, Girard R (2012) Evaluating the quality of scenarios of short-term wind power generation. Applied Energy 96:12–20, DOI 10.1016/j.apenergy.2011.11.004

Quelhas A, McCalley JD (2007) A multiperiod generalized network flow model of the u.s. integrated energy system: Part iisimulation results. Power Systems, IEEE Transactions on 22(2):837–844

Quelhas A, Gil E, McCalley JD, Ryan SM (2007) A multiperiod generalized network flow model of the u.s. integrated energy system: Part i – model description. IEEE Transactions on Power Systems 22(2):829–836

Rockafellar RT, Wets RJB (1991) Scenarios and policy aggregation in optimization under uncertainty. Mathematics of Operations Research 16(1):119–147

Ross S (1999) An Introduction to Mathematical Finance. Cambridge University Press, Cambridge, UK

Royset JO, Wets RJB (2014) From data to assessments and decisions: Epi-spline technology. INFORMS Tutorials in Operations Research pp 27–53, DOI 10.1287/educ.2014.0126

Ryan SM (1988) Degeneracy in discrete infinite horizon optimization. Ph.D. dissertation, The University of Michigan

Ryan SM (1998) Forecast frequency in rolling horizon hedging heuristics for capacity expansion. European Journal of Operational Research 109(3):550–558

Ryan SM (2003) Capacity expansion with lead times and correlated random demand. Naval Research Logistics 50(2):167–183, DOI 10.1002/nav.10055

Ryan SM (2004) Capacity expansion for random exponential demand growth with lead times. Management Science 50(6):740–748, DOI 10.1287/mnsc.1030.0187

Ryan SM, Bean JC (1989) Degeneracy in infinite horizon optimization. Mathematical Programming 43:305–316

Ryan SM, Bean JC, Smith RL (1992) A tie-breaking rule for discrete infinite horizon optimization. Operations Research 40(Supplement 1):S117–S126

Ryan SM, McCalley JD, Woodruff D (2011) Long term resource planning for electric power systems under uncertainty. In: Eto JH, Thomas RJ (eds) Computational Needs for the Next Generation Electric Grid, U. S. Department of Energy, pp 6–1–41, URL http://energy.gov/sites/prod/files/FINAL_CompNeeds_Proceedings2011.pdf

Sarı D, Ryan SM (2016) URL https://cran.r-project.org/package=MTDrh

Sarı D, Ryan SM (2017) Statistical reliability of wind power scenarios and stochastic unit commitment cost. Energy Systems DOI 10.1007/s12667-017-0255-7

Sarı D, Lee Y, Ryan S, Woodruff D (2016) Statistical metrics for assessing the quality of wind power scenarios for stochastic unit commitment. Wind Energy 19:873893, DOI 10.1002/we.1872

Sarı Ay D, Ryan SM (2018) Observational data-based quality assessment of scenario generation for stochastic programs, URL https://works.bepress.com/sarah_m_ryan/94/

Staid A, Watson JP, Wets RJB, Woodruff DL (2017) Generating short-term probabilistic wind power scenarios via nonparametric forecast error density estimators. Wind Energy 20(12):1911–1925, DOI 10.1002/we.2129

Székely GJ, Rizzo ML (2013) Energy statistics: A class of statistics based on distances. Journal of Statistical Planning and Inference 143(8):1249–1272, DOI 10.1016/j.jspi.2013.03.018

Wang Y (2010) Scenario reduction heuristics for a rolling stochastic programming simulation of bulk energy flows with uncertain fuel costs. Ph.D. dissertation, Iowa State University, URL http://search.proquest.com/docview/848503163

Wang Y, Ryan SM (2010) Effects of uncertain fuel costs on optimal energy flows in the u.s. Energy Systems 1:209–243

Wilks DS (2004) The minimum spanning tree histogram as a verification tool for multidimensional ensemble forecasts. Monthly Weather Review 132:1329–1340