



Distributionally robust optimization with multiple time scales: valuation of a thermal power plant

Wim van Ackooij¹ · Debora Daniela Escobar² · Martin Glanzer² · Georg Ch. Pflug^{2,3}

Received: 9 April 2019 / Accepted: 15 October 2019 / Published online: 20 November 2019
© The Author(s) 2019

Abstract

The valuation of a real option is preferably done with the inclusion of uncertainties in the model, since the value depends on future costs and revenues, which are not perfectly known today. The usual value of the option is defined as the maximal expected (discounted) profit one may achieve under optimal management of the operation. However, also this approach has its limitations, since quite often the models for costs and revenues are subject to model error. Under a prudent valuation, the possible model error should be incorporated into the calculation. In this paper, we consider the valuation of a power plant under ambiguity of probability models for costs and revenues. The valuation is done by stochastic dynamic programming and on top of it, we use a dynamic ambiguity model for obtaining the prudent minimax valuation. For the valuation of the power plant under model ambiguity we introduce a distance based on the Wasserstein distance. Another highlight of this paper is the multiscale approach, since decision stages are defined on a weekly basis, while the random costs and revenues appear on a much finer scale. The idea of bridging stochastic processes is used to link the weekly decision scale with the finer simulation scale. The applicability of the introduced concepts is broad and not limited to the motivating valuation problem.

Keywords Model ambiguity · Distributionally robust decision making · Multistage stochastic optimization · Multiscale stochastic optimization · Dynamic programming · Wasserstein distance · Nested distance

✉ Martin Glanzer
martin.glanzer@univie.ac.at

¹ OSIRIS, EDF R&D, 7 Boulevard Gaspard Monge, 91120 Palaiseau, France

² Department of Statistics and Operations Research, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Wien, Austria

³ International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria

1 Introduction

Since the deregulation of the energy market, the question of how to determine the value of a power plant can be asked. The traditional approach of valuing it within a given portfolio of other assets in a coordinated way against one's customer load is one possibility. A second approach is to adopt the ideas of real option pricing in finance. In the first case one ends up with models resembling unit commitment (e.g., van Ackooij et al. 2018) but at a long time scale. Although the actual operation of the power plant can be presented in great detail, it will be harder to incorporate other features in the model. This will typically be the case for uncertainty, where one ends up with multi-stage mixed-integer programs which are not easily solved. One can also argue that it is unreasonable to model the system as fully coordinated. In contrast, when modelling the power plant as a real option, thus operating it in the face of a set of market signals, the setting becomes that of perfect competition. Uncertainty is also naturally modelled, but it comes at the expense of modelling the plant as an independent production unit and thus with less realism in that sense.

However, the price of the real option may well serve as a financial reference base between two parties. For example between the power plant owner and a trading entity actually operating on the market. Taking the option pricing perspective, it must be emphasized that energy markets are by far not as "granular" as the equity markets. For instance, on the electricity market one cannot buy a contract of delivery for a given hour 6 months from now. The classical pricing-hedging duality argument is thus not feasible. Moreover, when operating a power plant, generation will be bound locally by a given power output level. This can be either the result of ramping conditions or minimum up/down times. It is therefore reasonable to try to model the power plant with sufficient realism for the above discrepancies to be minimal. This is the stance that we have taken in the current work. It will lead us to consider a multiscale stochastic program in the sense of Glanzer and Pflug (2019), i.e., a multistage stochastic optimization problem where each stage itself is subdivided into a given set of time instants.

To account for uncertainty, we start out with a set of typical stochastic models for underlying prices, which are based on multi-factor models (e.g., Clelow and Strickland 2000) driven by Brownian motions. Clearly, such (commonly used) idealized modelling assumptions are rather unrealistic. It is thus the aim and the core part of the present paper to relax such strong assumptions by computing distributionally robust solutions to the studied operational problem and to investigate how the resulting valuation deviates when considering model ambiguity. Distributionally robust optimization is a field which has recently gained a lot of popularity in the literature (see (Pflug and Pichler 2014, pp. 232–233) for a review of different approaches). In particular, ambiguity sets based on distance concepts between probability measures (such as the Wasserstein distance) are well-supported by theory and frequently applied (e.g., Pflug and Wozabal 2007; Esfahani and Kuhn 2018; Glanzer et al. 2019; Duan et al. 2018; Gao and Kleywegt 2016). However, to the best of our knowledge, the effects of distributional robustness in (especially multistage) real-world applications, have not been investigated yet.

In order to solve the formulated problem numerically, the given uncertainty model will be discretized on a scenario lattice. The multiscale structure could then simply

mean that uncertainty is lost within a given stage (cf., e.g., Moriggia et al. 2018). More advanced approaches do consider some uncertainty [e.g., the so-called multi-horizon approach originally suggested in Kaut et al. (2014) and subsequently studied and applied in Seljom and Tomasgard (2017), Skar et al. (2016), Werner et al. (2013), Zhonghua et al. (2015), Maggioni et al. (2019)], but the resulting paths do not necessarily connect with subsequent elements in the scenario tree/lattice. Hence, the multi-horizon approach is not appropriate for the present problem, as the key requirement of two time scales which may be assumed to run completely independent from each other, is not given. Indeed, we deal with two different granularities associated with one and the same stochastic process reflecting the evolution of the underlying market prices. A framework for such situations, where sub-stage paths in the lattice are carefully connected, has recently been proposed in Glanzer and Pflug (2019). We test the multiscale stochastic programming approach suggested in Glanzer and Pflug (2019) in the context of the present real-world application.

Although the resulting ideas will be illustrated through the power plant real option framework, their potential usage is readily seen to be beyond this specific application. In terms of contributions we can state:

- For the application of real option pricing, we investigate more reasonable exercise patterns. In order to keep computational burden low, this naturally leads to multiscale stochastic programs. We also consider model ambiguity to mitigate the fairly ideal models for market prices. From a high-level perspective, we thus extend the literature on real-world applications of dealing with two fundamental problems in stochastic programming, namely the problem of time scales with multiple granularities as well as the problem of model ambiguity.
- With respect to multistage model ambiguity, we propose a new concept based on the Wasserstein distance. It is tailored with a computational intention, namely in such a way that (on a discrete scenario tree/lattice) the applicability of a classical backward dynamic programming recursion can be maintained. In particular, the suggested framework leads to solutions that are robust w.r.t. model misspecification in a ball around each conditional transition probability distribution. The size of these balls may be controlled uniformly by a single input parameter. We also link the concept to the nested distance in such a way that it inherits a favourable stability property of the latter.
- In the context of Wasserstein ambiguity sets, we propose a state-dependent metric as a basis for the Wasserstein distance. Thereby we account for more realistic worst-case scenarios. We discuss that the well-appreciated statistical motivation for using Wasserstein balls is not invalidated by doing so.

The paper is organized as follows. Section 2 describes the valuation model and the uncertainty model. As typical for real-world energy applications, a sound mathematical framework reflecting all peculiarities of the problem requires carefulness in all details. The underlying uncertainty factors are modelled by a continuous time stochastic process. However, in the light of the nature of the decision problem, we will eventually apply a stochastic dynamic programming algorithm which operates backwards in (discrete) time. To prepare for the computational solution, we therefore discuss all discretization steps required by the multiscale stochastic programming framework

that we adopt. Section 3 is dedicated to model ambiguity. We introduce and discuss a new concept which is tailor-made for incorporating model ambiguity into dynamic stochastic optimization models on discrete structures. All numerical experiments and aspects of the computational solution algorithm are given in Sect. 4. Section 5 concludes. Some technical details and examples are deferred to the “Appendix”.

2 The model

Our valuation problem belongs to the class of discrete time sequential decision problems with finite horizon T , decisions u_t , state variables z_t , and a Markovian driving process ξ_t :

$$\begin{aligned} \max_{\{u_t\}_{t=0}^{T-1}} \mathbb{E} \left[\sum_{t=0}^{T-1} \beta_t h_t(z_t, u_t, \xi_t) \right] \\ \xi_{t+1} \quad \forall t = 0, \dots, T-1 \\ u_t \in \mathcal{U}_t(z_t) \text{ a.s. } \forall t = 0, \dots, T-1, \quad u_t = u_t(z_t, \xi_t), \\ z_t \in \mathcal{Z}_t \text{ a.s. } \forall t = 0, \dots, T-1 \end{aligned} \quad (1)$$

Here T is the number of decision stages and $g_t(z_t, u_t, \xi_t)$ is the state transition function. The driving stochastic process ξ_t is assumed to belong to $L_1(\Omega_t, \mathcal{F}_t; \mathbb{R}^m)$ and the feasible decision variables at stage t are defined by the set $\mathcal{U}_t(z_t) \subseteq \mathbb{R}^m$. The set of all reachable state variables is denoted by $\mathcal{Z}_t \subseteq \mathbb{R}^{d_1}$. The stage-wise profit function $h_t : \mathbb{R}^{d_1} \times \mathbb{R}^m \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$ is continuous and satisfies the following growth condition:

$$|h_t(z, u, x)| \leq K \cdot (1 + \|z\| + \|u\| + \|x\|),$$

for all $(z, u, x) \in \mathbb{R}^{d_1+m+d_2}$ and some constant K . We choose the discount factor $\beta_t = \beta^t$ for some constant $\beta \in (0, 1]$ throughout the paper. Any decision u_t to be made at time t may only depend on the current state z_t and the most recent observation of exogenous information ξ_{t-1} . This is the non-anticipativity condition. The initial conditions for the random process ξ and the state vector z are that ξ_0 and z_0 are assumed to be constant.

In our application, the decisions u_t represent the weekly electricity production plan for a thermal power plant. The latter is characterized by many technical constraints, such as minimum up/down times or ramping constraints. Fine grain constraints can be incorporated into the model by increasing the dimension of the state vector and accounting for the number of hours the plant has been offline/online. Such state-representations of constraints on generation assets have received attention in the literature (see, e.g., Martínéz et al. 2008; Frangioni and Gentile 2006; Frangioni et al. 2008 and the references therein). Finer granularity of the time dimension and/or the state variable would result in a significant increase of time steps T (and reduction of the time step size Δt) as well as an increase in the dimension of z_t . For this reason, we introduce here the idea of a multiscale model: While the production decisions

u_t are made on a weekly scale, the production costs and revenues are calculated on a finer time scale. To make the dynamic optimization algorithm tractable, we make the assumption that the decisions, i.e. the production profiles, must be chosen from a pre-specified set with finite cardinality. The profiles are set up such that they reflect realistic operating conditions and key choices, such as generating at minimal stable generation (MSG) at off peak hours.

Just prior to presenting the specific instantiation of (1), let us emphasize once more that the idea of subdividing a “stage” to mitigate (the curse of) dimensionality goes largely beyond the presented application. Typical other energy problems with similar mechanisms are cascaded reservoir management problems [e.g., see the extensive discussion in van Ackooij et al. (2014) as well as Escudero et al. (1996, 1999), Zéphyr et al. (2015), Cervellera et al. (2006), Aasgård et al. (2014), Séguin et al. (2017), Fleten et al. (2011)].

2.1 Instantiation of the problem: the valuation model

In our instantiation of problem (1) the time horizon is spanned by T weeks. Each week $t = 1, \dots, T$ is subdivided into S equally sized blocks of hours. With respect to our earlier introduced notation, we now present the following specific versions:

- the price process $\xi_{t,s} = (\xi_{t,s}^e, \xi_{t,s}^f, \xi_{t,s}^c) \in \mathbb{R}^3$ represents the electricity price in GBP per megawatt hour (£/MWh), the fuel price in USD per tonne (\$/tonne(fuel)) and the CO₂ allowances price in EUR per tonne (€/tonne(carbon)), for each block s within week t , for $s = 0, \dots, S$. The information up to stage t is the information up to $\xi_{t,0}$. The values within the weeks are $\xi_{t,s}$ for $s = 1, \dots, S$ with the convention that $\xi_{t,S} = \xi_{t+1,0}$ coincides with the initial prices of the next stage. With this convention we ensure continuity of prices in between weeks. In this way, the information up to stage t is to be understood as the information up to the value $\xi_{t,0}$.
- the control $u_t = \{u_{t,s}\}_{s=0}^{S-1} \in \mathcal{U} \subseteq \mathbb{R}_+^S$ represents the production profile vector for week t , where $u_{t,s}$ is given in megawatt (MW) and denotes the production at block s . Before the beginning of intermediate values of week t , we determine u_t . Then, u_t is $\xi_{t,0}$ -measurable.
- the state vector z_t is two-dimensional, i.e., $z_t = (x_t, y_t)$ with
 - $x_t \in \mathbb{R}_+$ representing the amount of CO₂ allowances (measured in tonnes of carbon), that are left for week t .
 - $y_t \in \mathbb{Z}_+$ representing the number of hours the power plant was offline before the beginning of week t .

The objective function $h_t : \mathbb{R} \times \mathbb{Z}_+ \times \mathbb{R}_+^S \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is given by

$$h_t([x_t, y_t], u_t, \xi_{t,0}) = \mathbb{E} \left[\sum_{s=0}^{S-1} f_s(x_t, y_t, u_t, \xi_{t,s}) \middle| \xi_{t,0} \right]. \quad (2)$$

Table 1 Description of the constants

Notation	Unit	Description
H_1	(£/\$)·(tonne(fuel))/MWh	$H_2 \cdot H_7$
H_2	£/\$	Exchange rate
H_3	£/€	Exchange rate
H_4	Tonne(carbon)/MWh	$H_7 \cdot J$
H_5	MWh/GJ	
H_6	Tonne(fuel)/GJ	$H_5 \cdot H_7$
H_7	Tonne(fuel)/MWh	Heat rate
J	tonne(carbon)/tonne(fuel)	Amount of CO ₂ emitted due to fuel burnt

The profit at each block s within week t is defined as follows:

$$f_s(x_t, y_t, u_t, \xi_{t,s}) = \begin{cases} (\xi_{t,0}^e - H_1 \xi_{t,0}^f) u_{t,0} \Delta s - f^{\text{CO}_2}(x_t, \bar{u}_t, \xi_{t,0}^c) \\ \quad - f^{\text{start}}(0, y_t, u_t, \xi_{t,0}^e, \xi_{t,0}^f) - f^{\text{tr}}(u_t) & \text{if } s = 0 \\ (\xi_{t,s}^e - H_1 \xi_{t,s}^f) u_{t,s} \Delta s - f^{\text{start}}(s, y_t, u_t, \xi_{t,s}^e, \xi_{t,s}^f) & \text{if } s > 0, \end{cases}$$

where $\bar{u}_t := \sum_{s=0}^{S-1} u_{t,s}$. Costs incurred are based on the following component functions:

- $f^{\text{CO}_2} : \mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ gives the cost of buying more CO₂ allowances at the beginning of week t (before the values within week t are known);
- $f^{\text{start}} : \mathbb{Z}_+^2 \times \mathbb{R}_+^S \times \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ gives the start-up cost if the power plant has been offline prior to (one of its arguments) block s ;
- $f^{\text{tr}} : \mathbb{R}_+^S \rightarrow \mathbb{R}_+$ represents fuel transportation costs linked to a selected production profile at the beginning of week t .

Table 1 summarizes constants used above or in the sequel.

The way in which each state variable is updated will be described now. First we will focus on the variables regarding the CO₂ allowances. Although in the past, a given set of allowances was allocated for free, in principle, they are now obtained from a non modelled auction process. Within our model, the variable I_t will represent the number of additional CO₂ allowances received from the regulator at the beginning of week t (measured in tonnes of carbon). Note that this variable will typically be equal to zero but sometimes it will take a relatively high value. The latter happens exactly at the rare events when new allowances are obtained.

Now, $H_4 \bar{u}_t \Delta s$ is the amount of generated CO₂ during stage t . Hence, together with x_t being the remaining stock level and I_t the “inflows”, the amount of allowances one needs to buy at stage t is $\alpha_t = [x_t + I_t - H_4 \bar{u}_t \Delta s]_-$.¹ In the case where α_t is positive, we follow a procurement strategy based on a low/middle/high price range partition resulting from some pre-market analysis. Prices in the interval $[\underline{b}, \bar{b}]$, with $0 \leq \underline{b} < \bar{b}$, are considered middle range. This is formalized as follows:

¹ We use here the definition $[f(x)]_- := -\min\{f(x), 0\}$ for the negative part of a function.

$$A(x_t, \bar{u}_t, \xi_{t,0}^c) = \alpha_t \cdot (1 + C_\alpha \cdot \min \{ \max \{ \bar{b} - \xi_{t,0}^c, 0 \} / (\bar{b} - \underline{b}), 1 \}), \quad (3)$$

where C_α is a constant that determines the size of the extra amount to be bought.

Recall, our implicit assumption is that new allowances are always bought before the prices within week t are known. The cost of buying more certificates for week t is then given by

$$f^{\text{CO}_2}(x_t, \bar{u}_t, \xi_{t,0}^c) = A(x_t, \bar{u}_t, \xi_{t,0}^c) \cdot H_3 \xi_{t,0}^c,$$

The amount x_{t+1} of remaining allowances after the previous purchase, is updated as follows:

$$g^{(1)}(x_t, u_t, \xi_{t,0}^c) = A(x_t, \bar{u}_t, \xi_{t,0}^c) + x_t + I_t - H_4 \bar{u}_t \Delta s.$$

The second state variable accounts for start-ups and related costs. The latter depend on the amount of time the power plant was offline. In our model this time frame will be partitioned into C different intervals of hours denoted by $(c_j, c_{j+1}]$, (c_C, ∞) , for $j = 1, \dots, C - 1$, $c_1 = 0$, over which the start-up costs are assumed to be constant. The associated costs are in terms of power, fuel burnt and extra costs. Depending on y_t and the chosen profile u_t , one can readily figure out in which interval each start-up of u_t falls.

The induced start-up costs at block s within week t are given by:

$$f^{\text{start}}(s, y_t, u_t, \xi_{t,s}^e, \xi_{t,s}^f) = \bar{W}_s(y_t, u_t) \xi_{t,s}^e + \bar{B}_s(y_t, u_t) H_2 H_6 \xi_{t,s}^f + \bar{E}_s(y_t, u_t), \quad (4)$$

where

- $\bar{W}_s(y_t, u_t)$ is the amount of works power (MWh) for a start-up at s ;
- $\bar{B}_s(y_t, u_t)$ is the amount of solid fuel burnt (GJ) during a start-up at s ;
- $\bar{E}_s(y_t, u_t)$ denotes engineering and imbalance costs (£) during a start-up at s .

The updated state y_{t+1} is given by:

$$g^{(2)}(u_t) = (S - \max \{ s : u_{t,s} \neq 0 \}) \cdot \Delta s,$$

where $\max \{ \emptyset \} = 0$.

As a further cost factor, we account for the fuel transportation costs associated to each profile:

$$f^{\text{tr}}(u_t) = C^{\text{tr}} H_7 \bar{u}_t \Delta s, \quad (5)$$

where the unit transportation cost (in £ per tonne of fuel) is given by the constant factor C^{tr} .

2.2 Underlying price processes

To model the underlying uncertainties, i.e., the stochastic price-evolution of electricity, fuel and CO₂ allowances, we postulate a version of a classical two-factor model. The latter are commonly used for the modeling of commodity markets (cf. Clelow and Strickland 2000; Ewald et al. 2018; Ribeiro and Hodges 2004; Farkas et al. 2017). More specifically, in our model the electricity price behaviour is governed by a long term and a short term factor, whereas fuel and CO₂ allowances prices evolve according to a one factor model. In summary, we get a three-dimensional geometric Brownian motion model driven by four correlated one-dimensional Brownian components $B^{e,sh}$, $B^{e,lo}$, B^f , B^c . In particular, the dynamics of the underlying stochastic process F are described by the SDE

$$\begin{pmatrix} dF_{t,t}^e/F_{t,t}^e \\ dF_{t,t}^f/F_{t,t}^f \\ dF_{t,t}^c/F_{t,t}^c \end{pmatrix} = \begin{pmatrix} \sigma_t^{e,sh} & 0 & 0 & \sigma_t^{e,lo} \\ 0 & \sigma_t^f & 0 & 0 \\ 0 & 0 & \sigma_t^c & 0 \end{pmatrix} \begin{pmatrix} dB_t^{e,sh} \\ dB_t^f \\ dB_t^c \\ dB_t^{e,lo} \end{pmatrix}, \quad (6)$$

where the superscripts *sh* and *lo* refer to short-term and long-term, respectively. Volatility is allowed to be time-dependent but deterministic. The double-index notation $F_{t,t}$ expresses the fact that we model the spot price as a special case of the forward price. In particular, the forward price $F_{0,t}$ (as observed in the market at time 0) will enter the solution of (6) at time t . In this way, we account for the well-known seasonality (peak-hours and off-peak-hours) inherent in electricity prices. Figure 1 visualizes this typical effect. To avoid notational clutter, we henceforth write ξ_t with one index for the spot price as a short hand for $F_{t,t}$. Note that we are dealing with a continuous time




Fig. 1 Typical structure of electricity forward prices. Each curve represents one week of EFA block power prices (in £/MWh, observed every 4 h), from Saturday, 3 am, until Friday, 11 pm. Five weeks of forward price data from the beginning of February to the first week of March, 2017

stochastic model here. The index notation should not be confused with the discrete-time multiscale indexes used in the valuation model; the context will always make clear what is meant.

Regarding the dependence structure between the underlying assets, we allow for a time dependent correlation matrix

$$\rho_t = \begin{pmatrix} 1 & \varrho_t^{e^{sh},f} & \varrho_t^{e^{sh},c} & \varrho_t^{e^{sh},e^{lo}} \\ \varrho_t^{e^{sh},f} & 1 & \varrho_t^{f,c} & \varrho_t^{e^{lo},f} \\ \varrho_t^{e^{sh},c} & \varrho_t^{f,c} & 1 & \varrho_t^{e^{lo},c} \\ \varrho_t^{e^{sh},e^{lo}} & \varrho_t^{e^{lo},f} & \varrho_t^{e^{lo},c} & 1 \end{pmatrix}.$$

Using the (lower triangular) matrix L_t resulting from a Cholesky decomposition of ρ_t , we may replace the Brownian factors $[dB_t^{e,sh}, dB_t^f, dB_t^c, dB_t^{e,lo}]^\top$ by the matrix-vector product $L_t \times [dW_t^{(1)}, dW_t^{(2)}, dW_t^{(3)}, dW_t^{(4)}]^\top$, such that the underlying prices are driven by independent Wiener processes W_1^s, W_2, W_3, W_1^l . Multiplying the volatility matrix in (6) with L_t , we can write the model in the form

$$\begin{pmatrix} d\xi_t^e / \xi_t^e \\ d\xi_t^f / \xi_t^f \\ d\xi_t^c / \xi_t^c \end{pmatrix} = \begin{pmatrix} a_{11}(t) & a_{12}(t) & a_{13}(t) & a_{14}(t) \\ a_{21}(t) & a_{22}(t) & 0 & 0 \\ a_{31}(t) & a_{32}(t) & a_{33}(t) & 0 \end{pmatrix} \begin{pmatrix} dW_t^{(1)} \\ dW_t^{(2)} \\ dW_t^{(3)} \\ dW_t^{(4)} \end{pmatrix}. \quad (7)$$

The non-zero components of the above coefficient matrix involve nasty terms with combinations of the various correlations. The precise parameters can be found in the “Appendix”.

The solution of SDEs of such a form as in (7) is well known to be of the geometric Brownian motion type (e.g., see (Oksendal 2000, p. 62)). In particular, the random vector $\xi_t = [\xi_t^e, \xi_t^f, \xi_t^c]$ follows a three-dimensional log-normal distribution. The corresponding parameters can again be found in the “Appendix”.

2.2.1 Discretization and the associated bridge process

For our numerical solution framework, which is discussed in detail in Sect. 4, the process ξ will first be discretized in all decision stages. Then, an approximate solution of the problem will be computed by stochastic dynamic programming with a backward recursion. In each decision stage, the algorithm relies on the expected profit/loss associated with any decision to be made for the upcoming observation blocks of the following week. To compute such values, we will exploit the structure of the valuation model, the uncertainty model and the backwards recursion. In particular, we are able to compute the expected profits by an analytical formula.

Let us start with the discretization step. To account for the two different time scales explained in Sect. 2.1 above, namely the weekly decision scale and the much finer intra-week observation scale, we use the notation $\xi_{t,s}$, where $t = 0, \dots, T$ runs in weeks and $s = 0, \dots, S$ in hour-blocks of equal size Δs (such that $t + S \cdot \Delta s = t + 1$).

Considering the fact that intra-week data of fuel and CO₂ allowances prices typically show – if even available – a rather stable evolution with low fluctuations, we assume those prices to be constant from Monday to Sunday. On the contrary, the electricity price dimension is truly stochastic even on a fine time-scale. Looking at the expected profit function (2) at some block s during a week t , it turns out that (on the basis of our assumptions) the problem boils down to the expected value of the electricity price $\xi_{t,s}^e$ given both the initial value $\xi_{t,0}^e$ as well as the final value $\xi_{t,S}^e$ of week t . This is due to the fact that the function $h_t(\cdot)$ is linear in $\xi_{t,s}^e$. Mathematically speaking, we are left with the computation of the conditional expected value at time $t + s \cdot \Delta s$ of the stochastic bridge process linking the values $\xi_{t,0}^e$ and $\xi_{t,S}^e$, for all $s = 1, \dots, S$. All other parts can be computed in a straightforward way.

The one-dimensional process $\xi_{t,s}^e$ follows a univariate lognormal distribution. Thus, its transition density δ is available in analytical terms and the transition density of the associated bridge process can be computed explicitly. Let an *initial* value η_1 of the process at the beginning of some week t and a *final* value η_2 at the end of that week be given (i.e., $\xi_{t,0}^e = \eta_1$ and $\xi_{t,S}^e = \eta_2$). Then, the bridge process transition density, at time $s \in [0, S]$, is given by

$$\begin{aligned} \delta(x, t + s \cdot \Delta s | \eta_1, t, \eta_2, t + 1) &= \frac{\delta(\eta_2, t + 1 | x, t + s \cdot \Delta s) \cdot \delta(x, t + s \cdot \Delta s | \eta_1, t)}{\delta(\eta_2, t + 1 | \eta_1, t)} \\ &= \frac{1}{\sqrt{2\pi\hat{\sigma}_{s|t}^2}x} \exp\left(-\frac{(\log(x) - \log(\eta_1) - \hat{\mu}_{s|t})^2}{2\hat{\sigma}_{s|t}^2}\right), \end{aligned}$$

where

$$\begin{aligned} \hat{\mu}_{s|t} &= \frac{\int_t^{t+s \cdot \Delta s} \sigma^2(u) du}{\int_t^{t+1} \sigma^2(u) du} \log\left(\left(\frac{F_{0,t}^e}{F_{0,t+1}^e}\right) \cdot \frac{\eta_2}{\eta_1}\right), \\ \hat{\sigma}_{s|t} &= \frac{\left(\int_t^{t+s \cdot \Delta s} \sigma^2(u) du\right) \left(\int_t^{t+1} \sigma^2(u) du\right)}{\int_t^{t+1} \sigma^2(u) du}. \end{aligned}$$

In particular, we get for the conditional expectation

$$\mathbb{E}[\xi_{t,s}^e | \xi_{t,0}^e = \eta_1, \xi_{t,S}^e = \eta_2] = \eta_1 \cdot \left(\frac{F_{0,t+s \cdot \Delta s}^e}{F_{0,t}^e}\right) \cdot \exp(\hat{\mu}_{s|t}). \quad (8)$$

Let us emphasize that the above analytical tractability is not due to our restriction of the intra-week stochasticity to one dimension (see Glazer and Pflug 2019 for a treatment of the more general multi-dimensional case). This restriction is purely motivated by data.




Fig. 2 Electricity forward price data (solid line) versus 5 simulated trajectories (dashed lines). Simulation based on the bridge process dynamics

Figure 2 illustrates a set of sample paths from the bridge process, which starts and ends in the forward prices corresponding to two consecutive weeks. The intermediate forward prices are shown for comparison of the seasonal behaviour.

3 Ambiguity for dynamic stochastic optimization models

It is an application of classical stochastic dynamic programming theory to solve (1) backwards in time on the basis of the following recursion scheme:

$$\begin{aligned} V_t(z_t, \xi_t) &= \max_{u_t \in \mathcal{U}_t(z_t)} h_t(z_t, u_t, \xi_t) + \beta \mathbb{E}[V_{t+1}(z_{t+1}, \xi_{t+1}) | \xi_t] \\ \text{s.t. } z_{t+1} &= g_t(z_t, u_t, \xi_{t+1}), \end{aligned} \quad (9)$$

where $V_T(z_T, \xi_T) \equiv 0$, z_0 and ξ_0 are given.

Let ξ be a Markovian process defined on a finite state space $\Xi_0 \times \dots \times \Xi_T$, where on each Ξ_t there is a distance d_t . Let the cardinality of Ξ_t be N_t with $N_0 = 1$ (typically nondecreasing in t). Then the transition matrices P_t , $t = 0, \dots, T-1$ are of the form $N_t \times N_{t+1}$, where the i -th row of the matrix P_t is denoted by $p_t(i)$, for all $i = 1, \dots, N_t$. Notice that each row $p_t(i)$ describes a probability measure on the metric space (Ξ_{t+1}, d_{t+1}) .

Let $\xi_t^i \in \Xi_t$ be given. Then, the conditional probability to transition to $\xi_{t+1}^j \in \Xi_{t+1}$ is given by the j th element of the row vector $p_t(i)$, denoted as $p_t(i, j)$ for $j = 1, \dots, N_{t+1}$ and $i = 1, \dots, N_t$. In this discrete case, the objective of the recursion in (9) can be written as

$$V_t(z_t, \xi_t^i) = \max_{u_t \in \mathcal{U}_t(z_t)} h_t(z_t, u_t, \xi_t^i) + \beta \sum_{j=1}^{N_{t+1}} p_t(i, j) \cdot V_{t+1}(g_t(z_t, u_t, \xi_{t+1}^j), \xi_t^j). \quad (10)$$

3.1 A new concept: uniform Wasserstein distances

In order to consider model ambiguity, we look for alternative transition matrices Q_t , which are close to a given matrix P_t . Let us first recall the general definition of the Wasserstein distance for discrete models.

Definition 3.1 Let $P = \sum_{i=1}^n P_i \delta_{\xi^i}$ and $Q = \sum_{j=1}^{\tilde{n}} Q_j \delta_{\tilde{\xi}^j}$ be two discrete measures sitting on the points $\{\xi^1, \dots, \xi^n\} \subset \Xi$ and $\{\tilde{\xi}^1, \dots, \tilde{\xi}^{\tilde{n}}\} \subset \tilde{\Xi}$, respectively. Then, the Wasserstein distance between P and Q is defined as

$$\mathfrak{W}(P, Q) := \min_{\pi_{ij}} \sum_{i,j} \pi_{ij} \cdot D_{ij},$$

where $\pi = \sum_{i,j} \pi_{i,j} \cdot \delta_{\xi^i, \tilde{\xi}^j}$ is a probability measure on $\Xi \times \tilde{\Xi}$ with marginals P and Q , and where D_{ij} is a distance between the resp. atoms.

We will now measure the closeness between (discrete) multistage models \mathbb{P} and \mathbb{Q} by a uniform Wasserstein distance concept. The rows of alternative matrices Q_t are denoted by $q_t(i)$, $i = 1, \dots, N_t$. The measure $q_t(i)$ is sitting on at most N_{t+1} points in Ξ_{t+1} and is such that $\text{supp}(q_t(i)) \subseteq \text{supp}(p_t(i))$. Then, we define the distance

$$\mathfrak{W}^\infty(\mathbb{P}, \mathbb{Q}) = \max_{0 \leq t \leq T-1} \max_{1 \leq i \leq N_t} \mathfrak{W}(p_t(i), q_t(i)), \quad (11)$$

which can be interpreted as a uniform version of scenariowise Wasserstein distances. An ε ball around \mathbb{P} is characterized by the fact that all members \mathbb{Q} satisfy

$$\mathfrak{W}(p_t(i), q_t(i)) \leq \varepsilon,$$

for all $t = 0, \dots, T-1$ and all $i = 1, \dots, N_t$.

When introducing ambiguity into the model, we would like to solve problem (1) wherein the objective function is replaced with:

$$\max_{\{u_t\}_{t=0}^{T-1}} \min_{\mathbb{Q} : \mathfrak{W}^\infty(\mathbb{P}, \mathbb{Q}) \leq \varepsilon} \mathbb{E}_{\mathbb{Q}} \left[\sum_{t=0}^{T-1} \beta_t h_t(z_t, u_t, \xi_t) \right],$$

where $\mathbb{E}_{\mathbb{Q}}$ denotes the expectation with respect to the measure \mathbb{Q} .

Our choice of the multistage distance makes it possible to keep the decomposed structure of the backward recursion, which reads:

$$\begin{aligned} V_t(z_t, \xi_t^i) &= \max_{u_t} \min_{\mathfrak{W}(p_t(i), q_t(i)) \leq \varepsilon} h_t(z_t, u_t, \xi_t^i) + \beta \sum_{j=1}^{N_{t+1}} q_t(i, j) V_{t+1}(z_{t+1}, \xi_{t+1}^j) \\ \text{s.t. } z_{t+1} &= g_t(z_t, u_t, \xi_{t+1}). \end{aligned} \quad (12)$$

Hence, the ambiguity approach just extends the max model to a maximin model. The ambiguous model can also be seen as a risk adverse model in contrast to the basic risk neutral model. If the distance is not of the decomposable form, then the backward recursion does not decompose scenariowise and one has to find all optimal decisions in one very big stagewise but not scenariowise decomposed algorithm. However, decomposability is the key feature of successful methods for dynamic decision problems. Hence, our concept is strongly motivated by its favourable computational properties. However, as we will discuss now, under a mild regularity condition (in the sense that it will always hold for discrete models, which are the basis of the whole computational framework) it can still be shown that optimal solutions are close if the underlying models are close w.r.t. the uniform Wasserstein distance.

The general distance concept for stochastic processes (including their discrete representation in the form of scenario trees) is the nested distance introduced in Pflug (2010), Pflug and Pichler (2012) as a multistage generalization of the classical Wasserstein distance.² In our case we have a Markov process which can be seen as a lattice process. Notice that a lattice can be interpreted as a compressed form of a tree. It can always be “unfolded” to a tree representing the same filtration structure, by splitting each node according to the number of incoming arcs. Thus, all results applying for trees do hold for lattices as well. The uniform Wasserstein distance introduced above is given by the maximum Wasserstein distance over all conditional transitions. The subsequent stability result holds.

Proposition 3.1 *Let \mathbb{P} and $\tilde{\mathbb{P}}$ be two discrete Markovian probability models defined on the filtered space $(\Omega, \sigma(\xi))$. Assume the following Lipschitz condition regarding \mathbb{P} to hold for all $t = 0, \dots, T-1$ and all values ξ_t^i, ξ_t^j , where $i, j = 1, \dots, N_t$:*

$$\mathfrak{W}\left(P_t(\cdot|\xi_t^i), P_t(\cdot|\xi_t^j)\right) \leq K_t \cdot \|\xi_t^i - \xi_t^j\|,$$

for $K_t \in \mathbb{R}$. Consider the generic multistage stochastic optimization problem

$$v(\mathbb{Q}) := \inf \mathbb{E}^{\mathbb{Q}}[c(x, \xi)],$$

² The general definition of the nested distance can be found in the “Appendix”.

where the (nonanticipative) decisions x lie in some convex set and where the function $c(\cdot, \cdot)$ is convex in x and 1-Lipschitz w.r.t. ξ . Then the relation

$$\left| v(\mathbb{P}) - v(\tilde{\mathbb{P}}) \right| \leq \mathfrak{d}\mathbb{I}(\mathbb{P}, \tilde{\mathbb{P}}) \leq K \cdot \mathfrak{W}^\infty(\mathbb{P}, \tilde{\mathbb{P}})$$

holds, where $\mathfrak{d}\mathbb{I}(\cdot, \cdot)$ denotes the nested distance, and where the constant K is given by

$$K := \sum_{t=0}^{T-1} \prod_{j=t+1}^T (1 + K_j).$$

Proof The first inequality is a well-known result from (Pflug and Pichler 2012, Th. 11). The statement then follows readily from (Pflug and Pichler 2014, Lem. 4.27), by using $\mathfrak{W}^\infty(\mathbb{P}, \tilde{\mathbb{P}})$ as a uniform bound for $\mathfrak{W}(P_t(\cdot|\xi_t), \tilde{P}_t(\cdot|\xi_t))$, over all t . \square

Remark Notice that for discrete Markov chain models the assumption in Proposition 3.1 always holds, as one can simply choose the ergodic coefficient

$$K_t = \max_{\xi_t^i \neq \xi_t^j} \frac{\mathfrak{W}\left(P_t(\cdot|\xi_t^i), P_t(\cdot|\xi_t^j)\right)}{\|\xi_t^i - \xi_t^j\|}.$$

Remark In the above construction, all models contained in the ambiguity set share exactly the same tree structure and node values. Thus, one might conjecture at a first glance that it would be possible to bound the nested distance by a simple sum of the stagewise maximum of conditional Wasserstein distances, weighted by the number of subtrees at the respective stage. A simple example in the “Appendix” shows that such a construction does not work in general.

3.2 State-dependent distances

In practice, the worst-case model for an upcoming period may often depend on the current state. In the model considered in the present paper (cf. Sect. 2.1), we decide only at the beginning of each stage about the procurement of additional CO₂ allowances. In particular, we restrict ourselves not to buy any if the current stock is sufficient for whatever we may do during the subsequent week; regardless of their market price. If we neglect this consideration when searching for the optimal distributionally robust production profile, the worst case may reflect a variation in the CO₂ allowances price dimension which in fact will not have an impact on our optimal decision. Thus, we modify the distance on the underlying three-dimensional space by projecting to the electricity price and the fuel price dimension only, given that our stock of CO₂ allowances is sufficient. Otherwise, we keep the usual L_1 norm. More formally, we define

$$D\left([\xi_t^e, \xi_t^f, \xi_t^c], [\tilde{\xi}_t^e, \tilde{\xi}_t^f, \tilde{\xi}_t^c]\right) := \begin{cases} w^e |\xi_t^e - \tilde{\xi}_t^e| + w^f |\xi_t^f - \tilde{\xi}_t^f| & \text{if } \alpha_t = 0 \\ w^e |\xi_t^e - \tilde{\xi}_t^e| + w^f |\xi_t^f - \tilde{\xi}_t^f| + w^c |\xi_t^c - \tilde{\xi}_t^c| & \text{if } \alpha_t > 0, \end{cases} \quad (13)$$

with α_t defined in Sect. 2.1 and positive weights w^e, w^f, w^c . Notice that D is not a distance, as it does not separate points. However, this fact does not entail any restrictions for our considerations.

When basing the uncertainty model on historical observations, there is a strong statistical argument for using balls w.r.t. the Wasserstein distance as ambiguity sets (cf. Esfahani and Kuhn 2018). In particular, large deviations results are available (see Bolley et al. 2007; Fournier and Guillin 2015 for the case of the Wasserstein distance and Glanzer et al. (2019) for the case of the nested distance) which provide probabilistic confidence bounds for the true model being contained in the ambiguity set around the (smoothed) empirical measure. Observe that such results are not invalidated by the state-dependency that we introduce: it is evident that a given confidence bound is directly inherited if one neglects some dimension. Notice however that a general state-dependent weighting of the dimensions would require a more careful treatment.

4 A case study

In the following, we will test the framework elaborated in Sects. 2.1 and 3 for a specific power plant. For the present application, each week t is subdivided in $S = 42$ blocks, where each block has 4 h. We solve the problem for a quarter ahead, thus we take the horizon to be $T = 13$ weeks.

In this case, the control variables are given vectors of dimension $S = 42$. The set of production profiles we use for this case study consists of 10 different production schedules. This set is denoted by $\mathcal{U} = \{u^{(i)}\}_{i=1}^{10}$ (see Fig. 3) and it will remain constant for every stage.

The discrete evolution of prices is given as a lattice process (ξ) defined on $\Xi_0 \times \dots \times \Xi_T$, where each space Ξ_t has N_t elements correspondent to the number of nodes, i.e., $\Xi_t = \{\xi_t^1, \dots, \xi_t^{N_t}\}$ and each node $\xi_t^i = (\xi_t^{e,i}, \xi_t^{f,i}, \xi_t^{c,i}) \in \mathbb{R}^3$ for all $i = 1, \dots, N_t$ and all $t = 0, \dots, T$. As explained in Sect. 3, P_t will denote the probability transition matrices from stage t to stage $t + 1$ of dimensions $N_t \times N_{t+1}$. The description of the lattice construction will be explained in Sect. 4.1.1.

The profit function h_t at stage t is defined by the expected profit during the upcoming week (see (2)). Profits at every block s in the week are quantified by the functions f_s , for $s = 0, \dots, S - 1$. Costs of buying additional allowances and transportation costs are quantified only at the beginning of each week, while start-up costs need to be assigned at each block s . We proceed to the description of the state variables. The costs of buying new allowances depend on the strategy A defined in (3). For its computation we consider $[\underline{b}, \bar{b}] = [4.4, 9.6]$, where the latter values were obtained by applying a simple quantile rule to the available data set. Moreover, we set $C_\alpha = 2$.




Fig. 3 All different profiles in \mathcal{U} . Production is given in MW




Fig. 4 If we choose profile $u^{(10)}$, the strategy A to follow is illustrated for different prices $\xi_{t,0}^c$ and all possible left allowances x_t in the partition. The horizontal lines indicate the lower and upper bounds for the prices \underline{b} and \bar{b} , respectively

For the partition of the amount of available CO₂ allowances x_t , we consider different possible values from 0 tonne(carbon) to 10^5 tonne(carbon), and we also take into account the allowances needed for each profile. All in all, the partition of state x_t has 16 different elements.

For every state in the partition, an example of the procurement strategy is shown in Fig. 4 for different prices. Note that we illustrate this example when we choose full

Table 2 Classification of offline hours and associated start-up costs

Offline hours	Solid fuel burnt (\bar{B}) in GJ	Works power (\bar{W}) in MWh	Engineering costs (EC)	Imbalance costs (IC)
(0, 20]	412	14	10,000	1600
(20, 36]	454	20.5	10,800	1600
(36, 48]	596	30	11,800	1600
(48, 65]	596	40.5	12,800	1600
(65, ∞)	629	60.9	13,800	1600

production, i.e., $u^{(10)}$. The strategy when choosing a different profile is similar, the only change is that we do not need to buy as many allowances as with $u^{(10)}$.

The second state variable y_t describes the hours the power plant was off since the last time it was on. The costs associated with restarting the production depend on y_t and the chosen profile u_t . The cost function is a step function given in Table 2.

Once a profile is chosen for any week, the first time the profile is different than zeros is where we consider initial start up costs. Then, for the initial start up costs we consider the hours the power plant was offline before week t starts (i.e., y_t) in addition to the hours the chosen profile is off before it starts producing. For the rest of the blocks the costs will only depend on the profile. As for the notation of the elements in f^{start} , see (4), \bar{E}_s is the sum of the last two columns of Table 2. We illustrate the initial start up costs for a profile that is on in the beginning of the week for all possible values of y_t (see Fig. 5).

Given that we have a finite set of profiles, we can calculate the value of y_t for each profile. The partition of y_t will have these values and the limit numbers of the classes in Table 2.

Regarding the transportation cost function in (5), we assume $C^{tr} = 40$ (£/tonne(fuel)).

Finally, the values of the constants in Table 1 are $H_2 = 0.78$ (£/\$), $H_3 = 0.9$ (£/€), $H_5 = 0.0975$ (MWh/GJ), $H_7 = 0.45$ (tonne(fuel)/MWh) and $J = 2.31$ (tonne(carbon)/tonne(fuel)).

4.1 The solution algorithm

We numerically solve the power plant valuation problem by a stochastic dynamic programming algorithm. A lattice structure is used as a discrete representation of the uncertainty model in all decision stages.

4.1.1 On the lattice construction

The state-of-the art approach for the construction of scenario lattices is based on optimal quantization techniques (cf. Bally and Pagès 2003; Löhndorf and Wozabal 2018). For a given number of discretization points, such methods select the optimal locations as well as the associated probabilities in such a way that the Wasserstein distance




Fig. 5 Initial start-up costs (£) for all possible initial states $y_t \in \{0, 1, \dots, 168\}$, when the profile is on in the beginning of the week. The costs are calculated for a specific node with electricity price 40 €/MWh and fuel price 90 \$/tonne(fuel)

(or some other distance concept for probability measures) with respect to a (continuous) target distribution is minimized. For the present study, we have implemented a stochastic approximation algorithm for the quantization task, following (Pflug and Pichler 2014, Algorithm 4.5). Referring to the latter algorithm, we first applied the iteration step (ii) in order to find the atoms of all marginal distributions, separately for each stage. We then formed a lattice out of these sets of points by fixing the structure of allowed transitions and then applied step (iv) to determine all conditional transition probabilities. Eventually, this also determines the absolute probabilities of each node in the lattice. The Wasserstein distance of order two has been used as a target measure for the minimization. We use a ternary lattice, i.e., each node has (at most) three successors with a positive transition probability.

As for the notation, recall that we distinguish between stages (weeks) and intra-week blocks. Decisions are taken only at each stage but the profit will be calculated taking into account the random evolution of the prices during the entire week. The discretized process in each stage t and node i is denoted as $\xi_t^i = (\xi_t^{e,i}, \xi_t^{f,i}, \xi_t^{c,i})$, and the values of the process within week t starting in node i , are denoted by $\xi_{t,s}^i = (\xi_{t,s}^{e,i}, \xi_{t,s}^{f,i}, \xi_{t,s}^{c,i})$, for $s = 0, \dots, S$. At $s = 0$, $\xi_{t,0}^i = \xi_t^i$ takes the value of node i ; and at $s = S$, $\xi_{t,S}^i = \xi_{t+1}^j$ takes the value of node j in the next stage with probability $p_t(i, j)$, for $j = 1, \dots, N_{t+1}$. Note that if the lattice structure does not make a link between two nodes in consecutive stages, then the probability of such a transition will be zero.

4.2 Computing the expected profit between two decisions

The valuation of the power plant is obtained by solving (10) backwards in time from $t = T$ to $t = 0$. In this section, we specify how to solve (10) and its robust version (12) at any stage t . We specifically concentrate on the calculation of the expected profit within each week given the current node and a successor node.

As discussed in Sect. 2.2.1, electricity prices within weeks will be modeled by the bridge process that we described. Fuel and CO₂ allowances prices are assumed to remain constant between stages.

We start now with the computation of the expected profit, as defined in (2). In classical stochastic dynamic programming problems, h_t exclusively depends on the values observed at time t . In contrast, in Sect. 2.1 we instantiated (1) in such a way that the function h_t is defined as an expected value of the the random profits within week t . Hence, given the values of node i at stage t (at block $s = 0$), as well as initial states (x_t, y_t) ; the weekly profit h_t will be calculated as follows

$$\begin{aligned} h_t([x_t, y_t], u_t, \xi_{t,0}^i) &= \mathbb{E} \left[\sum_{s=0}^{S-1} f_s(x_t, y_t, u_t, \xi_{t,s}^i) \middle| \xi_t^i = \xi_{t,0}^i \right] \\ &= \sum_{j=1}^{N_{t+1}} \left(\sum_{s=0}^{S-1} \mathbb{E} \left[f_s(x_t, y_t, u_t, \xi_{t,s}^i) \middle| \xi_t^i = \xi_{t,0}^i, \xi_{t+1}^j = \xi_{t+1,0}^j \right] \right) \cdot p_t(i, j). \end{aligned}$$

We define

$$h_t^j(x_t, y_t, u_t, \xi_{t,0}^i) = \sum_{s=0}^{S-1} \mathbb{E} \left[f_s(x_t, y_t, u_t, \xi_{t,s}^i) \middle| \xi_t^i = \xi_{t,0}^i, \xi_{t+1}^j = \xi_{t+1,0}^j \right],$$

for all $j = 1, \dots, N_{t+1}$. Then, $h_t([x_t, y_t], u_t, \xi_{t,0}^i) = \sum_{j=1}^{N_{t+1}} h_t^j(x_t, y_t, u_t, \xi_{t,0}^i) \cdot p_t(i, j)$. We compute now h_t^j as follows

$$\begin{aligned} h_t^j(x_t, y_t, u_t, \xi_{t,0}^i) &= \sum_{s=0}^{S-1} (u_{t,s} \Delta s - \bar{W}_s(y_t, u_t)) \cdot \mathbb{E}[\xi_{t,s}^{e,i} | \xi_t^e = \xi_{t,0}^{e,i}, \xi_{t+1}^e = \xi_{t+1,0}^{e,j}] \\ &\quad - \xi_{t,0}^{f,i} \sum_{s=0}^{S-1} (H_1 u_{t,s} \Delta s + H_2 H_6 \bar{B}_s(y_t, u_t)) \\ &\quad - A(x_t, \bar{u}_t, \xi_{t,0}^{c,i}) H_3 \xi_{t,0}^{c,i} - \sum_{s=0}^{S-1} \bar{E}_s(y_t, u_t) - H_7 C^{tr} \bar{u}_t \Delta s. \end{aligned}$$

At stage T we set the terminal condition $V_T = 0$. Given an initial state (x_0, y_0) , going backwards in time from $t = T - 1$ to $t = 0$, we obtain the power plant value at $t = 0$. The latter is calculated with respect to the baseline multistage model \mathbb{P} and will

be denoted as $v_0(\mathbb{P})$. The policy associated with $v_0(\mathbb{P})$ is denoted with $u_{\mathbb{P}}^*$. It can be represented as a probabilistic tree of profiles.

If we incorporate ambiguity in the lattice process using the uniform Wasserstein distance, for all $0 \leq t \leq T - 1$ we solve:

$$V_t([x_t, y_t], \xi_t^i) = \max_{u_t} \min_{\mathfrak{W}(p_t(i), q_t(i)) \leq \varepsilon} \left\{ \sum_{j=1}^{N_{t+1}} \left[h_t^j(x_t, y_t, u_t, \xi_{t,0}^i) + \beta V_{t+1}([g^{(1)}(x_t, u_t, \xi_{t,1}^c), g^{(2)}(u_t)], \xi_{t+1}^j) \right] q_t(i, j) \right\}.$$

The optimal value is reached at $t = 0$ and it will be denoted as $v_0(\mathcal{Q}^\varepsilon)$, where \mathcal{Q}^ε denotes the ambiguity set defined as

$$\mathcal{Q}^\varepsilon = \{\mathbb{Q} : \mathfrak{W}^\infty(\mathbb{P}, \mathbb{Q}) \leq \varepsilon\}.$$

A worst-case model $\mathbb{Q}^{\varepsilon*}$ is any multistage probability model contained in \mathcal{Q}^ε such that $v_0(\mathcal{Q}^\varepsilon) = v_0(\mathbb{Q}^{\varepsilon*})$. More concrete, the optimal value is reached at a saddle point $(u_{\mathbb{Q}^{\varepsilon*}}^*, \mathbb{Q}^{\varepsilon*})$ where $u_{\mathbb{Q}^{\varepsilon*}}^*$ is the policy associated with the worst-case model.

At each node i , the objective function of the minimization problem is linear in $q_t(i)$ under linear constraints. Define

$$c_j = h_t^j(x_t, y_t, u_t, \xi_{t,0}^i) + \beta V_{t+1}([g^{(1)}(x_t, u_t, \xi_{t,1}^c), g^{(2)}(u_t)], \xi_{t+1}^j),$$

for $j = 1, \dots, N_{t+1}$. Then, the minimization problem can be written as

$$\begin{aligned} \min_{q_t(i,j), \pi_{k,l}} \quad & \sum_{j=1}^{N_{t+1}} c_j \cdot q_t(i, j) \\ \text{s.t.} \quad & \sum_{k=1}^{N_{t+1}} \pi_{k,l} = q_t(i, l) \quad \forall l = 1, \dots, N_{t+1} \\ & \sum_{l=1}^{N_{t+1}} \pi_{k,l} = p_t(i, k) \quad \forall k = 1, \dots, N_{t+1} \\ & \sum_{k,l} D_{kl} \pi_{k,l} \leq \varepsilon \\ & \sum_{k,l} \pi_{k,l} = 1 \\ & \pi_{k,l} \geq 0, \forall k, l = 1, \dots, N_{t+1}, \end{aligned}$$

where $D_{kl} = D\left([\xi_{t+1}^{e,k}, \xi_{t+1}^{f,k}, \xi_{t+1}^{c,k}], [\xi_{t+1}^{e,l}, \xi_{t+1}^{f,l}, \xi_{t+1}^{c,l}]\right)$ is the distance between nodes k and l at stage $t + 1$, as defined in (13) with specific weights $w^e = 1$, $w^f = H_1$ and $w^c = H_3 \cdot H_4$.

4.3 Impact of model ambiguity

4.3.1 The value of the power plant

We describe the optimal valuation of the power plant when we compute the iterative system of backward equations in (10) and (12). We assume that the initial state x_0 provides enough allowances to execute any of the profiles and the power plant was not offline before we start, i.e., $y_0 = 0$. Moreover, the terminal condition for both problems is set to be $V_T = 0$. With the baseline model we obtain an expected profit of approximately $v_0(\mathbb{P}) = 2.3 \cdot 10^6$ (£). The optimal decision at $t = 0$ is to turn off the power plant by choosing $u^{(1)}$. The valuation of the power plant including ambiguity is obtained for different radii $\varepsilon \in [0, 2]$. The different optimal values $v_0(\mathcal{Q}^\varepsilon)$ are

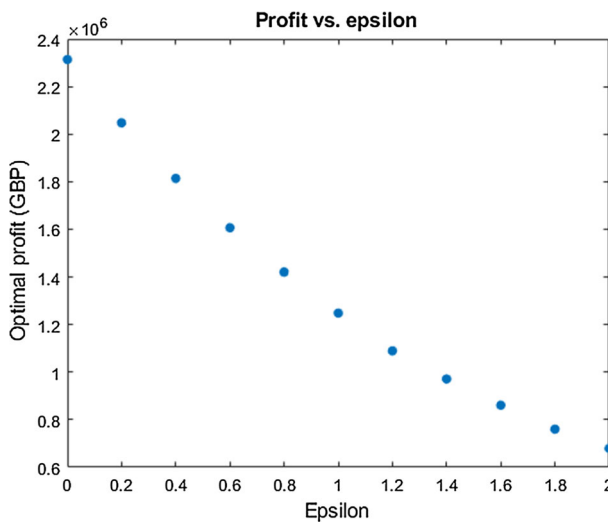


Fig. 6 The impact of the ambiguity radius ε on the optimal profit over 13 weeks

Table 3 Percentage of change in electricity prices for every stage and different ambiguity radii

ε	Stage-wise change in allowances prices								
	1	2	3	6	7	8	11	12	13
0.4	0.0072	0.0034	0.0094	0.0179	0.0121	0.0109	0.0387	0.0692	0.0442
1	0.0179	0.0183	0.0311	0.0437	0.0410	0.0179	0.0806	0.1633	0.0932
1.4	0.0274	0.0066	0.0316	0.0413	0.0532	0.0263	0.1287	0.2511	0.1121
2	0.0415	0.0173	0.0493	0.0678	0.0892	0.0359	0.1879	0.2991	0.1277

illustrated in Fig. 6 with respect to the ambiguity level. We observe the valuation of the power plant decreases when the ambiguity radius is higher. For $\varepsilon = 2$, the valuation of the power plant decreases to $v_0(Q^\varepsilon) = 6.8 \cdot 10^5$.

In order to get an insight into the change of prices in the ambiguity model, we report the changes of electricity prices for the worst case models $Q^{\varepsilon*}$ in Table 3. Let B_0 be a vector containing the stagewise expectations of electricity prices with respect to the baseline model and let B_ε be the vector containing the expectations with respect to each worst-case model $Q^{\varepsilon*}$. The percentage of change at each stage t in the prices are denoted with the parameter θ_t , such that $B_\varepsilon(t) = (1 - \theta_t)B_0(t)$. We observe that the largest change in prices is given in stage 12, where electricity prices decrease up to 30% for $\varepsilon = 2$.

4.3.2 Forward in time

With the iterative solution of the backwards equations we eventually obtain an initial optimal profile at $t = 0$, namely $u_0^* = u^{(1)}$. With this initial decision we go forward in time and create a probabilistic tree $u_{\mathbb{P}}^*$ of the optimal decisions together with their profits. Starting with the given states and the optimal profile at $t = 0$, the updated states at stage $t = 1$ are completely determined by the knowledge of x_0 , y_0 and u_0^* . The choice of the optimal profile in $t = 1$, for each node $i = 1, \dots, N_1$, will be made by looking at the nearest location of the updated states in the grid and taking the correspondent profile chosen in the backwards algorithm. We proceed in this way until we obtain all the optimal profiles at stage $T - 1$. Eventually, we obtain a probabilistic tree with 3^T possible paths. Following the same procedure, we calculate




Fig. 7 Stagewise distribution of the optimal profiles $u_{\mathbb{P}}^*$




Fig. 8 Stage distribution of the optimal profile with respect to the ambiguity radius. For each ε we plot the distribution of $u_{Q^{\varepsilon*}}^*$

the probabilistic tree of optimal profiles $u_{Q^{\varepsilon*}}^*$ for each worst-case model $Q^{\varepsilon*}$, and the corresponding profits under the worst-case models for different radii ε .

Starting with $u^{(1)}$ is optimal for all models at $t = 0$. For the subsequent stages the choices of optimal profiles change. Figure 7 shows the stagewise distribution of the optimal profiles chosen with the baseline model \mathbb{P} . Figure 8 shows the changes of profile choices when we incorporate ambiguity in the model.

With no ambiguity there is a probability greater than 0 to choose full production in stages 6, 7, 8, 10 and 12. When we start increasing the radius of ambiguity these chances drop to 0. The larger the ambiguity radius is the more we choose to be offline or not to produce in the weekends choosing profiles like $u^{(2)}$, $u^{(4)}$, $u^{(6)}$. A different option, but with less probability is not to produce in peak hours, by choosing $u^{(3)}$ or $u^{(5)}$.

Given the optimal profiles for the baseline model \mathbb{P} and the alternative models $Q^{\varepsilon*}$ we can calculate the profits we make along the decision tree. To be precise, we denote by $i_\tau \in N_\tau$ any node index at stage $\tau = 0, \dots, T$. Since $N_0 = 1$, a possible path to follow forward in time up to stage t , will go through any sequence of nodes




Fig. 9 On the left the profit tree by following the optimal profiles and on the right the distribution of the final profits obtained following every path

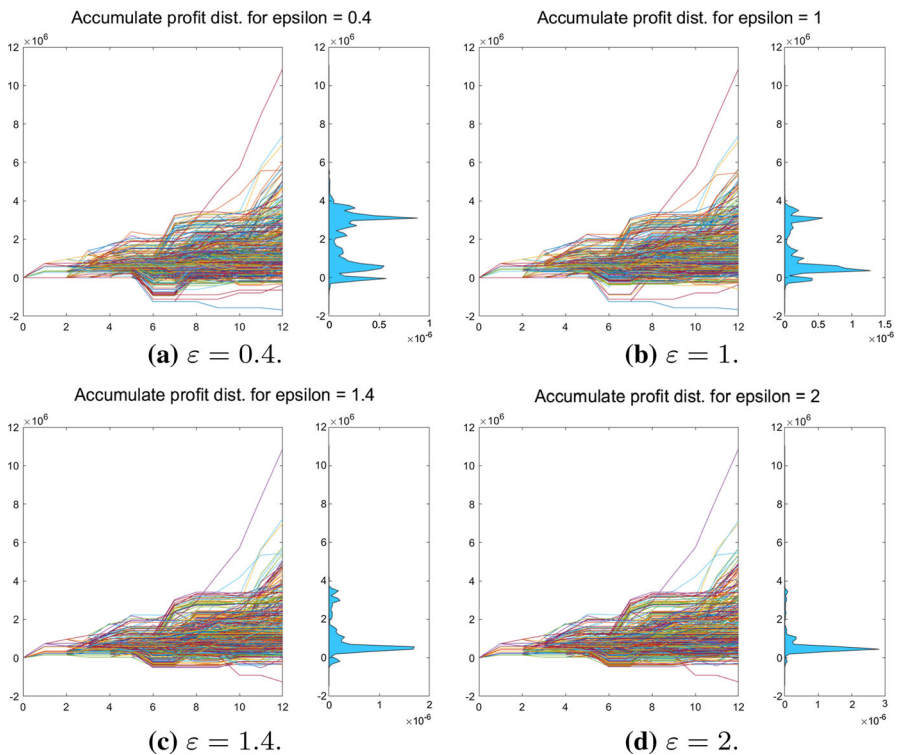


Fig. 10 Stagewise accumulated profit distribution for each solution $u_{\mathbb{Q}^{\varepsilon}}^*$

$(1, i_1, \dots, i_{t-1}, i_t)$. The profit from i_τ to $i_{\tau+1}$ is the profit made in stage τ and is written as $h_\tau^{i_\tau, i_{\tau+1}}$. This profit is obtained with probability $p_\tau(i_\tau, i_{\tau+1})$. Therefore, the accumulated profit until stage $t - 1$ is $(h_0^{0, i_1} + \dots + h_{t-1}^{i_{t-1}, i_t})$ with probability $p_0(1, i_1) \cdots p_{t-1}(i_{t-1}, i_t)$, when we end in node $i_t \in N_t$. Figure 9 shows the accumulated profits following the tree of optimal profiles $u_{\mathbb{P}}^*$ as well as the distribution of the final profits.

If we include ambiguity, then the optimal profits change as well as their probabilities. Figure 10 shows the profit trees together with the final distribution with respect to the correspondent alternative model. We observe that for larger ε the alternative models put more weight at lower profits.

5 Conclusion

In this paper, we have shown how a realistic valuation of a power plant can be done by solving a multistage Markovian decision problem. The value is defined as the (discounted) expected net profit, that one can get from the operation of the plant, if an optimal production plan is implemented. In this valuation process, all relevant purchasing costs and selling prices are included in the model. The number of feasible production plans is finite and thus a discrete multistage optimization problem has to be solved. We use the classical backward algorithm for the Markovian control problem and a forward algorithm for determining an estimate of the achievable profit and its distribution. The novelty of the paper is twofold. First, we adopt a multiscale approach, where decisions are made on a coarser scale than costs are calculated. This allows us to keep the computational effort tractable. Second, we do not only consider the baseline model for the random factors, but rather a set of models (the ambiguity set) which are close to the baseline model. This allows to incorporate the fact that probability distributions for future costs and revenues are not known precisely. The more models, and especially the more unfavourable models are included in the ambiguity set, the smaller is the robust value of the plant. We demonstrate how the final value under model ambiguity depends on the degree of uncertainty about the correct price and cost model. Our distance model for the ambiguity set depends on the state of the system, taking into account that what is close for two price vectors depends also on the fact whether these prices are relevant for the state at hand. We also noticed that the optimal production strategy not only depends on the degree of ambiguity, but also gets more diversified for larger ambiguity, in contrast to some bang-bang solutions in unambiguous models.

Acknowledgements Open access funding provided by University of Vienna. The authors would like to acknowledge financial support from the Gaspard-Monge program for Optimization and Operations Research (FMJH/PGMO) project with number 2017-0094H and title “Incorporating model error in the management of electricity portfolios”.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

Parameters for the model

The parameters corresponding to Eq. (7) in Sect. 2.2 are given by:

$$\begin{aligned}
 a_{11}(t) &= \sigma_t^{e,sh} + \sigma_t^{e,lo} \varrho_t^{esh,e^{lo}}, & a_{12}(t) &= \sigma_t^{e,lo} \frac{\varrho_t^{e^{lo},f} - \varrho_t^{esh,f} \varrho_t^{esh,e^{lo}}}{\sqrt{1 - (\varrho_t^{esh,f})^2}}, \\
 a_{13}(t) &= \sigma_t^{e,lo} \frac{\varrho_t^{e^{lo},c} - \varrho_t^{esh,c} \varrho_t^{esh,e^{lo}} - \frac{(\varrho_t^{esh,f} \varrho_t^{esh,e^{lo}} - \varrho_t^{e^{lo},f})(\varrho_t^{esh,f} \varrho_t^{esh,c} - \varrho_t^{f,c})}{1 - (\varrho_t^{esh,f})^2}}{\sqrt{1 - (\varrho_t^{esh,c})^2 - \frac{(\varrho_t^{f,c} - \varrho_t^{esh,f} \varrho_t^{esh,c})^2}{1 - (\varrho_t^{esh,f})^2}}}, \\
 a_{14}(t) &= \sigma_t^{e,lo} \sqrt{\left\{ 1 - (\varrho_t^{esh,e^{lo}})^2 - \frac{(\varrho_t^{e^{lo},f} - \varrho_t^{esh,f} \varrho_t^{esh,e^{lo}})^2}{1 - (\varrho_t^{esh,f})^2} \right.} \\
 &\quad \left. - \frac{(\varrho_t^{esh,f} \varrho_t^{esh,c} \varrho_t^{e^{lo},f} + \varrho_t^{e^{lo},c} - (\varrho_t^{esh,f})^2 \varrho_t^{e^{lo},c} - \varrho_t^{esh,c} \varrho_t^{esh,e^{lo}} - \varrho_t^{e^{lo},f} \varrho_t^{f,c} + \varrho_t^{esh,f} \varrho_t^{esh,e^{lo}} \varrho_t^{f,c})^2}{((\varrho_t^{esh,f})^2 - 1)((\varrho_t^{esh,f})^2 + (\varrho_t^{esh,c})^2 + (\varrho_t^{f,c})^2 - 2\varrho_t^{esh,f} \varrho_t^{esh,c} \varrho_t^{f,c} - 1)} \right\}}, \\
 a_{21}(t) &= \sigma_t^f \varrho_t^{esh,f}, & a_{22}(t) &= \sigma_t^f \sqrt{1 - (\varrho_t^{esh,f})^2}, & a_{31}(t) &= \sigma_t^c \varrho_t^{esh,c}, \\
 a_{32}(t) &= \sigma_t^c \frac{\varrho_t^{f,c} - \varrho_t^{esh,f} \varrho_t^{esh,c}}{\sqrt{1 - (\varrho_t^{esh,f})^2}}, & a_{3,3} &= \sigma_t^c \sqrt{1 - (\varrho_t^{esh,c})^2 - \frac{(\varrho_t^{f,c} - \varrho_t^{esh,f} \varrho_t^{esh,c})^2}{(\varrho_t^{esh,f})^2 - 1}}.
 \end{aligned}$$

The parameters of the multivariate lognormal distribution of the GBM process (7), i.e., the expectation vector $\mu(t)$ and the (components of the) variance-covariance matrix $\Sigma(t)$ of the associated multivariate normal distribution are given by:

$$\begin{aligned}
 \mu(t) &= \begin{pmatrix} \log(F_{0,t}^e) - \frac{1}{2} \int_0^t a_{11}^2(s) + a_{12}^2(s) + a_{13}^2(s) + a_{14}^2(s) \, ds \\ \log(F_{0,t}^f) - \frac{1}{2} \int_0^t a_{21}^2(s) + a_{22}^2(s) \, ds \\ \log(F_{0,t}^c) - \frac{1}{2} \int_0^t a_{31}^2(s) + a_{32}^2(s) + a_{33}^2(s) \, ds \end{pmatrix}, \\
 \Sigma^{11}(t) &= \int_0^t a_{11}^2(s) + a_{12}^2(s) + a_{13}^2(s) + a_{14}^2(s) \, ds, \\
 \Sigma^{12}(t) &= \Sigma^{21}(t) = \int_0^t a_{11}(s)a_{21}(s) + a_{12}(s)a_{22}(s) \, ds, \\
 \Sigma^{13}(t) &= \Sigma^{31}(t) = \int_0^t a_{11}(s)a_{31}(s) + a_{12}(s)a_{32}(s) + a_{13}(s)a_{33}(s) \, ds, \\
 \Sigma^{23}(t) &= \Sigma^{32}(t) = \int_0^t a_{21}(s)a_{31}(s) + a_{22}(s)a_{32}(s) \, ds, \\
 \Sigma^{22}(t) &= \int_0^t a_{21}^2(s) + a_{22}^2(s) \, ds, & \Sigma^{33}(t) &= \int_0^t a_{31}^2(s) + a_{32}^2(s) + a_{33}^2(s) \, ds.
 \end{aligned}$$




Fig. 11 Two versions of a scenario tree with different transition probabilities. The nested distance is 10. The Wasserstein distance between the first stage subtrees is 4, that between each of the two pairs of second stage subtrees is 2

A remark on the nested distance

Definition 5.1 The nested distance $\text{dl}(\mathbb{P}, \tilde{\mathbb{P}})$ between two \mathbb{R}^m -valued filtered stochastic processes $(\mathbb{P}, \mathcal{F})$ and $(\tilde{\mathbb{P}}, \tilde{\mathcal{F}})$ is defined as the optimal value of the following mass transportation problem, which optimizes over the set of all joint distributions that respect the given conditional marginals:

$$\begin{aligned} & \inf_{\pi} \iint \|\omega - \tilde{\omega}\| \pi(d\omega, d\tilde{\omega}) \\ & \text{s.t. } \pi\left(A \times \mathbb{R}^m \mid \mathcal{F}_t \otimes \tilde{\mathcal{F}}_t\right) = \mathbb{P}\left[A \mid \mathcal{F}_t\right] \quad A \in \mathcal{F}_T; \forall 0 \leq t \leq T \\ & \quad \pi\left(\mathbb{R}^m \times B \mid \mathcal{F}_t \otimes \tilde{\mathcal{F}}_t\right) = \tilde{\mathbb{P}}\left[B \mid \tilde{\mathcal{F}}_t\right] \quad B \in \tilde{\mathcal{F}}_T; \forall 0 \leq t \leq T. \end{aligned}$$

Figure 11 illustrates that the nested distance (between two trees resulting from a variation of the transition probabilities) cannot be bounded by only considering the Wasserstein distance between the subtrees with matching node values. In the given example, the maximum Wasserstein distance in the second stage is 2, the first case Wasserstein distance is 4. Thus, $4 + 2 \cdot 2 = 8$ would still be smaller than the nested distance between the two trees, which is 10.

References

- Aasgård EK, Andersen GS, Fleten S, Haugstvedt D (2014) Evaluating a stochastic-programming-based bidding model for a multireservoir system. *IEEE Trans Power Syst* 29(4):1748–1757
- Bally V, Pagès G (2003) A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli* 9(6):1003–1049
- Bolley F, Guillin A, Villani C (2007) Quantitative concentration inequalities for empirical measures on non-compact spaces. *Probab Theory Relat Fields* 137(3–4):541–593
- Cervellera C, Chen V, Wen A (2006) Optimization of a large-scale water reservoir network by stochastic dynamic programming with efficient state space discretization. *Eur J Oper Res* 171(3):1139–1151
- Clelow L, Strickland C (2000) *Energy derivatives: pricing and risk management*. Lacima Group, London
- Duan C, Fang W, Jiang L, Yao L, Liu J (2018) Distributionally robust chance-constrained approximate AC-OPF with Wasserstein metric. *IEEE Trans Power Syst* 33(5):4924–4936
- Escudero LF, de la Fuente JL, García C, Prieto FJ (1996) Hydropower generation management under uncertainty via scenario analysis and parallel computation. *IEEE Trans Power Syst* 11(2):683–689

- Escudero LF, Quintana FJ, Salmeron J (1999) CORO, a modeling and an algorithmic framework for oil supply, transformation and distribution optimization under uncertainty. *Eur J Oper Res* 114(3):638–656
- Esfahani PM, Kuhn D (2018) Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Math Program* 171(1–2):115–166
- Ewald C-O, Zhang A, Zong Z (2018) On the calibration of the Schwartz two-factor model to WTI crude oil options and the extended Kalman filter. *Ann Oper Res* 282(1–2):119–130
- Farkas W, Gourier E, Huitema R, Necula C (2017) A two-factor cointegrated commodity price model with an application to spread option pricing. *J Bank Financ* 77(C):249–268
- Fleten S-E, Haugstvedt D, Steinsbø JA, Belsnes M, Fleischmann F (2011) Bidding hydropower generation: Integrating short- and long-term scheduling. MPRA Paper 44450, University Library of Munich, Germany
- Fournier N, Guillin A (2015) On the rate of convergence in Wasserstein distance of the empirical measure. *Probab Theory Relat Fields* 162(3–4):707
- Frangioni A, Gentile C (2006) Solving non-linear single-unit commitment problems with ramping constraints. *Oper Res* 54(4):767–775
- Frangioni A, Gentile C, Lacalandra F (2008) Solving unit commitment problems with general ramp constraints. *Int J Electr Power Energy Syst* 30:316–326
- Gao R, Kleywegt AJ (2016) Distributionally robust stochastic optimization with Wasserstein distance. *arXiv preprint arXiv:1604.02199*
- Glanzer M, Pflug GCh (2019) Multiscale stochastic optimization: modeling aspects and scenario generation. *Comput Optim Appl*. <https://doi.org/10.1007/s10589-019-00135-4>
- Glanzer M, Pflug GCh, Pichler A (2019) Incorporating statistical model error into the calculation of acceptability prices of contingent claims. *Math Program* 174(1):499–524
- Kaut M, Midthun K, Werner A, Tomasgard A, Hellemo L, Fodstad M (2014) Multi-horizon stochastic programming. *Comput Manag Sci* 11(1):179–193
- Löhndorf N, Wozabal D (2018) Gas storage valuation in incomplete markets. *Optimization*. http://www.optimization-online.org/DB_HTML/2017/02/5863.html
- Maggioni F, Allevi E, Tomasgard A (2019) Bounds in multi-horizon stochastic programs. *Ann Oper Res*. <https://doi.org/10.1007/s10479-019-03244-9>
- Martínez MG, Diniz AL, Sagastizábal C (2008) A comparative study of two forward dynamic programming techniques for solving local thermal unit commitment problems. In: 16th PSCC Conference, Glasgow Scotland, p 1–8
- Moriggia V, Kopa M, Vitali S (2018) Pension fund management with hedging derivatives, stochastic dominance and nodal contamination. *Omega*
- Oksendal B (2000) Stochastic differential equations, 5th edn. Springer, Berlin
- Pflug GCh, Pichler A (2012) A distance for multistage stochastic optimization models. *SIAM J Optim* 22(1):1–23
- Pflug GCh, Pichler A (2014) Multistage stochastic optimization, 1st edn. Operations research and financial engineering. Springer, Berlin
- Pflug GCh, Wozabal D (2007) Ambiguity in portfolio selection. *Quant Financ* 7(4):435–442
- Pflug GCh (2010) Version-independence and nested distributions in multistage stochastic optimization. *SIAM J Optim* 20(3):1406–1420
- Ribeiro DR, Hodges SD (2004) A two-factor model for commodity prices and futures valuation. In: EFMA 2004 Basel Meetings Paper
- Séguin S, Fleten S-E, Sôté P, Pichler A, Audet C (2017) Stochastic short-term hydropower planning with inflow scenario trees. *Eur J Oper Res* 259(1):1156–1168
- Seljom P, Tomasgard A (2017) The impact of policy actions and future energy prices on the cost-optimal development of the energy system in Norway and Sweden. *Energy Policy* 106(C):85–102
- Skar C, Doorman G, Pérez-Valdés GA, Tomasgard A (2016) A multi-horizon stochastic programming model for the European power system. Censes working paper 2/2016, NTNU Trondheim. ISBN: 978-82-93198-13-0
- van Ackooij W, Lopez I Danti, Frangioni A, Lacalandra F, Tahanan M (2018) Large-scale unit commitment under uncertainty: an updated literature survey. *Ann Oper Res* 271(1):11–85
- van Ackooij W, Henrion R, Möller A, Zorghi R (2014) Joint chance constrained programming for hydro reservoir management. *Optim Eng* 15:509–531
- Werner AS, Pichler A, Midthun KT, Hellemo L, Tomasgard A (2013) Risk measures in multi-horizon scenario trees. Springer, New York, pp 177–201

- Zéphyr L, Lang P, Lamond BF (2015) Controlled approximation of the value function in stochastic dynamic programming for multi-reservoir systems. *Comput Manag Sci* 12(4):539–557
- Zhonghua S, Egging R, Huppmann D, Tomasgard A (2015) A Multi-Stage Multi-Horizon Stochastic Equilibrium Model of Multi-Fuel Energy Markets. Censes working paper 2/2016, NTNU Trondheim. ISBN: 978-82-93198-15-4

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.