



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Meeting Corporate Renewable Power Targets

Alessio Trivella, Danial Mohseni-Taheri, Selvaprabu Nadarajah

To cite this article:

Alessio Trivella, Danial Mohseni-Taheri, Selvaprabu Nadarajah (2022) Meeting Corporate Renewable Power Targets.
Management Science

Published online in Articles in Advance 25 Mar 2022

. <https://doi.org/10.1287/mnsc.2022.4354>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.




For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Meeting Corporate Renewable Power Targets

Alessio Trivella,^a Danial Mohseni-Taheri,^b Selvaprabu Nadarajah^{b,*}

^aIndustrial Engineering and Business Information Systems, University of Twente, 7500 AE Enschede, The Netherlands; ^bCollege of Business Administration, University of Illinois at Chicago, Chicago, Illinois 60607

*Corresponding author

Contact: a.trivella@utwente.nl,  <https://orcid.org/0000-0002-2614-5051> (AT); smohse3@uic.edu,  <https://orcid.org/0000-0002-1945-2221> (DM-T); selvan@uic.edu,  <https://orcid.org/0000-0001-6218-5862> (SN)

Received: December 2, 2018

Revised: April 17, 2020; July 25, 2021; January 10, 2022

Accepted: January 15, 2022

Published Online in Articles in Advance: March 25, 2022

<https://doi.org/10.1287/mnsc.2022.4354>

Copyright: © 2022 INFORMS

Abstract. Several corporations have committed to procuring a percentage of their electricity demand from renewable sources by a future date. Long-term financial contracts with renewable generators based on a fixed strike price, known as virtual power purchase agreements (VPPAs), are popular to meet such a target. We formulate rolling power purchases using a portfolio of VPPAs as a Markov decision process, accounting for uncertainty in generator availability and in the prices of electricity, renewable energy certificates, and VPPAs. Obtaining an optimal procurement policy is intractable. We consider forecast-based reoptimization heuristics consistent with practice that limit the sourcing of different VPPA types and the timing of new agreements. We extend these heuristics and introduce an information-relaxation based reoptimization heuristic, both of which allow for full sourcing and timing flexibilities. The latter heuristic also accounts for future uncertainties when making a decision. We assess the value of decision flexibility in rolling power purchases to meet a renewable target by numerically comparing the aforementioned policies and variants thereof on realistic instances involving a novel strike price stochastic process calibrated to data. Policies with full timing flexibility and no sourcing flexibility reduce procurement costs significantly compared with one with neither type of flexibility. Introducing sourcing flexibility in the former policies results in further significant cost reduction, thus providing support for using VPPA portfolios that are both dynamic and heterogeneous. Computing near-optimal portfolios of this nature entails using our information-relaxation based reoptimization heuristic because portfolios constructed via forecast-based reoptimization exhibit higher suboptimality.

History: Accepted by Ilia Tsetlin, behavioral economics and decision analysis.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2022.4354>.

Keywords: power purchase agreements • corporate renewable energy • climate targets • Markov decision processes • information relaxation and duality • reoptimization

1. Introduction

Corporations are playing an increasing leadership role in promoting sustainability and social responsibility around the globe. More than half of the Fortune 500 companies have publicly announced commitments to meet sustainability and climate goals, which include greenhouse gas emissions reduction, energy efficiency improvements, and renewable power procurement (CDP et al. 2017). We focus on companies that have committed to targets that entail procuring a specified percentage of annual electricity demand from renewable power sources by a future date and sustain this level of renewable procurement thereafter. For example, Johnson & Johnson and Intel have committed to targets of 30% and 100% by 2030, respectively. More than 300 companies have pledged such targets as part of the RE100 initiative.¹ Technology companies are also setting analogous targets specific to their data centers in the United States (Shankleman

2020) and in Europe.² To meet these ambitious procurement goals, electricity purchases need to be coupled with renewable energy certificates (RECs), where each REC allows its owner to validate the use of one megawatt hour (MWh) of renewable energy.

In 2019, 80% of corporate renewable power procurement was through virtual (a.k.a. synthetic) power purchase agreements (VPPAs), which are long-term financial agreements on generation capacity (Baker McKenzie 2018; BNEF 2018a, 2020). Under this agreement, the producer sells the electricity generated to the grid and transfers the associated RECs to the corporation; the firm buys electricity at a short-term (i.e., uncertain) price³; and payments for differences between a preagreed strike price and the short-term price are made to ensure a price hedge (RMI 2016, WBCSD 2018).⁴ Despite this price hedge, the amount of electricity that is generated from the contracted

capacity is uncertain because of operations (e.g., maintenance) beyond the control of the buyer and the intermittency of renewable resources such as wind and solar. In other words, typical VPPAs ensure that the buyer pays a fixed price per MWh on an uncertain quantity of electricity. This feature is a blessing, as it avoids VPPAs being classified as a derivative and adhering to financial regulation (Davies et al. 2018, p. 5) and a curse because it does not eliminate the company's exposure to electricity price and supply uncertainties.

VPPAs are quite different from traditional long-term contracts that ensure certainty in price and supply, and they are signed in a rapidly evolving renewable energy market. In response, there has been significant recent experimentation in corporate renewable procurement policies. Traditionally, long-term power contract lengths were from 20 to 25 years. Recent VPPAs include contracts with tenors ranging from 5 to 15 years, and companies have begun actively managing VPPA portfolios by rolling power purchases over time. There is thus significant flexibility in constructing and adapting such portfolios. To the best of our knowledge, the extant academic literature and practitioner reports do not formally investigate or assess the value of this flexibility to a firm trying to meet a renewable procurement target. Our goal is to take a meaningful step toward reducing this knowledge gap by studying procurement policies involving two dominant procurement options (BNEF 2018b): (i) enter into VPPAs and (ii) buy short-term electricity and purchase, as needed, unbundled RECs.⁵

We formulate a Markov decision process (MDP) that formalizes the use of a dynamic VPPA procurement portfolio to meet a target. This MDP minimizes expected procurement costs over a planning horizon that is divided into a reach period where the target does not have to be fulfilled (but contracts can be signed) and a sustain period where the target must be satisfied. At each stage, the company decides whether to enter into new VPPAs. The set of available VPPAs depends on the contracts offered by generators, which is unpredictable over time. Moreover, when entering a VPPA, the associated purchase quantity needs to satisfy minimum and maximum requirements. These realistic features translate to the procurement decisions from an optimal MDP policy (i) accounting for the *stochastic evolution of prices and supply* and (ii) capturing the *timing flexibility* to sign VPPAs and the *sourcing flexibility* to allocate capacity across VPPAs with different tenors. Nevertheless, computing this policy is intractable owing to our MDP having a high-dimensional state space and a nonconvex action set.

Practice-based procurement policies that use forecasts and take a rolling power purchase approach can

be viewed as special cases of our MDP. In its traditional form, such a policy signs VPPAs of a single fixed tenor (e.g., 20 years), maintains only one VPPA at any time, and determines the capacity of the new VPPA for the next block of the planning horizon (e.g., years 21–40) once the incumbent VPPA is close to expiry. The VPPA capacity is computed by solving easy-to-implement-deterministic models that use forecasts of stochastic quantities in the MDP and limit the forecasting effort to the VPPA tenor. This strategy has low timing flexibility and no sourcing flexibility. We consider two extensions: one that increases timing flexibility by allowing the signing of new VPPAs each year and another that increases timing flexibility in this manner and sourcing flexibility by allowing heterogeneous VPPAs to be signed each year.

The aforementioned policies use forecasts and thus do not directly capture the impact of the future evolution of uncertainty on the current procurement decision. This is conceptually undesirable because the MDP optimal policy does account for such impact of future uncertainty. Moreover, understanding if this discrepancy results in poor empirical policy performance requires a lower bound on the optimal policy value. We develop a novel procurement heuristic based on information relaxations that provides both a lower bound and decisions that account for future uncertainty. It requires more setup than the forecast-based reoptimization approach but continues to rely on solving deterministic models. To elaborate, our reoptimization heuristic solves deterministic (hindsight) optimization models along sample paths in Monte Carlo simulation, with costs corrected by a dual penalty term based on the information relaxation and duality approach (Andersen and Broadie 2004, Haugh and Kogan 2004, Brown et al. 2010). It extracts a nonanticipative decision from a distribution of anticipative decisions across sample paths using a function that we refer to as a decision measure. Examples of decision measures include the mean, median, and mode of a distribution. We specify when a decision measure leads to a feasible procurement policy and leverage the theory on information relaxations to show that this policy is optimal when using an ideal dual penalty.

We conduct numerical experiments on realistic instances with VPPA contract lengths ranging from 5 to 25 years, a planning horizon of 40 years, and electricity demand corresponding to that of two large data centers. The strike prices of VPPAs are specified by a new latent variable model that we develop and calibrate to annualized strike price data based on recent deals. We calibrate known stochastic processes for the evolution of electricity/REC prices as well as supply using market data and the practitioner literature. We uncover the following insights:

- **Cost of renewable target and value of VPPAs:** The cost of a target depends on the procurement policy. When using only short-term electricity purchases and unbundled RECs (i.e., no VPPAs), the procurement costs increase by 21% when the target is 100% (1% represents roughly 5 million USD) compared with having no target. In contrast, when using our best VPPA policy with full timing and sourcing flexibilities, meeting a 100% target leads to a procurement cost increase of 11%. VPPA portfolios are thus valuable in substantially reducing the procurement costs when committing to a target. In addition, procurement costs increase at a slower rate with the target level when using VPPAs than without these agreements. Therefore, the benefit of VPPAs is higher for corporations with more aggressive targets.

- **Value of timing and sourcing flexibilities in VPPA portfolios:** The timing flexibility to potentially sign new VPPAs of the same tenor every year reduces procurement costs by 3.5% on average in comparison with a policy that waits until the current contract expires to sign a new VPPA. The procurement costs decrease by an additional 4.0% on average when there is sourcing flexibility to sign new VPPAs of varying tenors each year. There is thus significant value in dynamically updating a heterogeneous VPPA portfolio.

- **Value of advanced reoptimization:** Procurement costs under the policy obtained using forecast-based reoptimization can be reduced by an average of 3.6% using instead the policy from our information-relaxation based reoptimization heuristic, which accounts for uncertainty when computing decisions. Moreover, the latter policies are near optimal, where an assessment of suboptimality is possible because of the availability of information-relaxation lower bounds. In other words, although dynamic VPPA portfolios can reduce procurement costs significantly, achieving the full extent of this reduction requires an advanced reoptimization technique, such as the one we propose.

- **Impact of strike prices:** Changes in technology and renewable energy policy can result in more volatile VPPA strike prices in the future. Dynamic portfolios of VPPAs result in more significant cost reductions in this case because they provide flexibility for the corporate buyer to asymmetrically benefit from low strike prices when they are available while insulating the buyer from getting locked into long-term agreements with high strike prices by instead signing shorter term ones. In addition, the price differences between short- and long-tenor VPPAs (i.e., tenor-dependent premiums) have a significant impact on the composition of VPPA portfolios. This highlights the value of data on tenor-dependent strike prices to corporate buyers as it would allow them to estimate such premiums. The continued efforts by governments and research agencies to improve

the availability of data on renewable power deals (BerkeleyLab 2020a, b) is thus useful for corporate procurement analytics.

1.1. Novelty and Related Work

We build on the extant literature that studies commodity procurement using spot purchases and forward contracts (Li and Kouvelis 1999, Kleindorfer and Wu 2003, Boyabatli et al. 2011, Secomandi and Kekre 2014) and procurement in supply chains via short-term and long-term contracts, including dual- and multisourcing options (Martínez-de Albéniz and Simchi-Levi 2005, Tomlin and Wang 2005, Veeraraghavan and Scheller-Wolf 2008, Allon and Van Mieghem 2010). Our study of renewable power procurement adds to this line of work. Specifically, our focus on constructing procurement portfolios to meet a target, the strike price stochastic model, the dynamic policies, and related insights are new to this literature. Moreover, the long-term contracts that we consider, that is, VPPAs, have unique structure. For instance, VPPAs deliver energy at each period over the tenor of the contract and their payoff depends on both price and supply uncertainties, which together differs from the long-term contracts considered in the aforementioned papers.

Our work indeed contributes to the growing literature on renewable energy. Several studies in this area study important market level issues related to supply intermittency (Wu and Kapuscinski 2013, Hu et al. 2015, Aflaki and Netessine 2017, Zhou et al. 2019), power supply equilibria (Al-Gwaiz et al. 2016, Sunar and Birge 2019), support schemes and their impact on renewable energy investments (İşlegen and Reichelstein 2011, Drake et al. 2016, Singh and Scheller-Wolf 2017), and market-based or equilibrium-based pricing of feed-in tariffs and VPPAs (Wu and Babich 2012, Alizamir et al. 2016, Ritzenhofen et al. 2016, Bruck et al. 2018). We instead investigate a problem faced by a corporation, that is, a firm-level decision problem as opposed to a market level issue. We also do not focus on pricing VPPAs nor do we use an equilibrium model for this purpose. Rather, we take a corporate buyer's perspective and develop a novel latent variable model to specify the evolution of VPPA strike prices in our MDP and calibrate it to annualized strike price data based on recent deals.

A more closely related research subarea focuses on individual players in the renewable power market. In particular, this stream of research studies the valuation and operations of renewable generators and operators of storage and transmission assets (Denholm and Sioshansi 2009, Kim and Powell 2011, Jiang and Powell 2015, Pandžić et al. 2015, Zhou et al. 2019), as well as the management of consumer incentive

programs such as demand response (Chao and Chen 2005, Webb et al. 2017, and references therein). To the best of our knowledge, a study of the electricity procurement problem faced by a corporation with a target is new to the renewable energy literature. Moreover, our comparison of both traditional and contemporary procurement policies from the unifying lens of an MDP enriches this literature.

The procurement policies we consider add to existing rolling-horizon planning approaches, also known as certainty equivalent control or reoptimization methods, which have been successfully used for decision making under uncertainty in engineering and business applications (Chand et al. 2002, Bertsekas 2005). In the context of energy, reoptimization models are popular for determining the next day unit commitment and real-time economic dispatch of power generators (Weber et al. 2009, Meibom et al. 2011, Milligan et al. 2012). They have also been used in real option settings, most notably for managing energy storage (Lai et al. 2010, Wu et al. 2012, Nadarajah and Secomandi 2018). Our information-relaxation reoptimization heuristic introduces a new reoptimization scheme to this literature. Moreover, our extensive numerical study expands the set of applications for which reoptimization has been considered and shows the value of our proposed reoptimization scheme compared with existing forecast-based approaches.

In addition to the reoptimization literature, we contribute to the active research on the information relaxation and duality approach (Brown and Smith 2011, 2014; Brown and Haugh 2017; Haugh and Lacedelli 2018; Nadarajah and Secomandi 2018; Ye et al. 2018), which does not directly provide control policies. Therefore, Desai et al. (2012) design an auxiliary procedure to obtain decisions in this framework. Specifically, they estimate a value function approximation by regressing on value function estimates computed by solving dual optimization problems in Monte Carlo simulation. This approximation is then used along with the MDP Bellman operator to compute decisions. It is not easy to extend the approach of Desai et al. (2012) to our setting: Estimating a value function approximation and computing decisions using the Bellman operator are both challenging because of the large controllable part of the state space in our MDP arising from tracking the generation capacity contracted via VPPAs. Our new reoptimization heuristic thus adds a direct way to obtain nonanticipative controls when using the information relaxation and duality approach.

More broadly, our information relaxation reoptimization heuristic adds to approximate dynamic programming (Bertsekas 2005), an area of stochastic optimization dealing with the solution of high-dimensional MDPs. Several methods tackle MDPs

where either the endogenous state or the exogenous state is high-dimensional. Well-known examples include least squares Monte Carlo (Longstaff and Schwartz 2001, Tsitsiklis and Van Roy 2001), approximate linear programming (De Farias and Van Roy 2003, Lin et al. 2019), and stochastic dual dynamic programming (Pereira and Pinto 1991, Shapiro 2011). However, methods to approximately solve MDPs with high-dimensional endogenous and exogenous state components are limited (Salas and Powell 2017, Nadarajah and Secomandi 2018) and approaches that handle nonconvex action sets are even more scarce. Combining reoptimization with information relaxations has potential value for solving other problems beyond our specific procurement application. For example, recent and independent work by Min et al. (2021) leverages information relaxations in conjunction with Thompson sampling to obtain control policies for multiarmed bandit problems with encouraging performance.

1.2. Paper Structure

In Section 2, we formulate an MDP to meet a renewable target using rolling power purchases. We present procurement heuristics that approximate this model in Sections 3 and 4. The dynamics of the VPPA strike prices and other uncertainties are discussed in Section 5. We conduct an extensive numerical study and discuss insights in Section 6. We conclude in Section 7. All proofs and additional material related to our models and numerical study can be found in an online appendix.

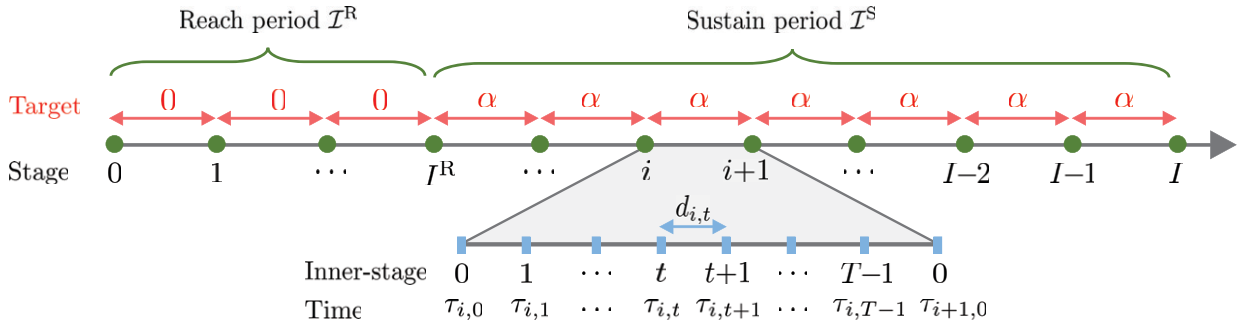
2. MDP View of Rolling Power Purchases

Two different procurement perspectives for meeting a target are to (i) use short-term electricity and REC purchases and (ii) use a single large long-term VPPA with shortfalls procured via the short-term market. The former *short-term* policy exposes the buyer fully to price volatility. The latter long-term VPPA policy faces the risk of locking in uncertain supply at terms that may become unfavorable in a rapidly evolving power market. To balance these risks, we consider *rolling power purchases* to meet a target, which is a flexible strategy that maintains a portfolio of VPPA types with short and long tenors and signs new VPPAs, as needed, on a regular basis. We formulate in this section an MDP model that optimizes this strategy and provides a unifying lens to understand the procurement policies we consider in Sections 3 and 4.

2.1. Planning Horizon

The planning horizon in our MDP is composed of two nested intervals as illustrated in Figure 1. The coarser intervals are delimited by *stages*, which are points in

Figure 1. (Color online) Planning Horizon, Stages, and Inner Stages



time when VPPA contracting decisions are made and we assume any RECs needed to meet the target are purchased from the short-term market at this time. The finer intervals are delimited by inner stages, which are points in time at which the corporation (i) uses short-term electricity purchases to meet electricity demand and (ii) exchanges cash flows with the generator to hedge the price for electricity generated from capacity contracted via VPPAs. The time between inner stages depends on the nature of the short-term electricity purchase, which could be for example from a utility or a retail energy supplier and vary monthly, quarterly or annually. In summary, the enforcement of the target takes place at the decision stages, whereas VPPA cash flow settlements take place at inner stages. To ease exposition, we assume equidistant annual stages and equidistant inner stages.

Formally, the horizon is comprised of I annual intervals, and hence $I + 1$ stages indexed by i belonging to set $\mathcal{I} := \{0, \dots, I\}$. Year i denotes the annual interval that starts at stage i and ends at stage $i + 1$. There are T intervals between stages i and $i + 1$. We index the inner-stages defining these intervals by t with support in set $\mathcal{T} := \{0, \dots, T-1\}$, where we exclude inner-stage T since it coincides with inner-stage 0 of year $i + 1$. The time associated with stage i and inner-stage t is denoted $\tau_{i,t}$. Therefore, $\tau_{i,0}$ is now. A target is enforced annually from year I^R onward, that is, a percentage $\alpha \in [0, 1]$ of the annual demand in each year $i \in \mathcal{I}^S := \{I^R, \dots, I-1\}$ must be satisfied by renewable sources. The target α does not have to be fulfilled in the remaining part of the planning horizon, that is, $\mathcal{I}^R := \{0, \dots, I^R-1\}$, but VPPAs can be signed. The demand that needs to be satisfied in the interval $(\tau_{i,t}, \tau_{i,t+1}]$ in year i is denoted by $d_{i,t}$ (MWh).

2.2. Actions

At a given stage i , the set of potentially available VPPAs are indexed by m with ground set \mathcal{M} . We assume these contracts are differentiated by their duration so that m can be interpreted as the length of

a contract and $M = \max\{m \in \mathcal{M}\}$ is the length of the longest contract.⁶ A VPPA of length m signed at stage i delivers electricity from stages $i + 1$ to $i + m$ at a fixed strike price of $K_{i,m}$ USD/MWh. Thus, the last time VPPAs can be signed is stage $I - 2$. These agreements deliver electricity from stage $I - 1$ to I , whereas no VPPA can be signed at stages $I - 1$ and I . The company can choose to enter into a new VPPA of type m , if it is available, by determining a capacity level in MW that is within minimum and maximum allowable limits (often imposed by the generator) represented by z_m^{\min} and z_m^{\max} , respectively.

We model contract availability at stage i using a binary vector $a_i := (a_{i,m} \in \{0, 1\}, m \in \mathcal{M})$, where $a_{i,m}$ equals one, if contract m is available, and is zero, otherwise. The continuous-valued procurement decision vector is $z_i := (z_{i,m}, m \in \mathcal{M})$, where $z_{i,m}$ is the contracted capacity in MW⁷ of the VPPA of length m years signed at year i . Given a contract availability vector a_i , the vector z_i belongs to set

$$\mathcal{Z}_i(a_i) := \{z_i \in \mathbb{R}_+^{|\mathcal{M}|} \mid z_{i,m} = 0, \text{ if } a_{i,m} = 0, \text{ and } z_{i,m} \in \{0\} \cup [z_m^{\min}, z_m^{\max}], \text{ otherwise, } \forall m \in \mathcal{M}\}, \quad (1)$$

which includes $|\mathcal{M}|$ continuous dimensions (e.g., $|\mathcal{M}| = 5$ in our numerical study), and is nonconvex when minimum purchase quantities are strictly positive (i.e., $z_m^{\min} > 0$).

2.3. State Space

The information required to make procurement decisions at a stage i is described in the MDP state, which contains two components. The first component is a vector $x_i := (x_{i,l}, l = 0, \dots, M-1)$ representing the on-hand power capacity from VPPA contracts, where $x_{i,l}$ is the total capacity in MW available at year $i + l$ by the on-hand VPPAs. The second component $w_{i,t}$ contains the stochastic factors at stage i and inner-stage t needed to determine the Markovian evolution of the

electricity price $P_{i,t}$ (USD/MWh), REC price R_i (USD/MWh), strike price $K_{i,m}$ (USD/MWh), contract availability a_i , and renewable capacity availability $\eta_{i,t}$ (hours), which varies between zero and the number of hours between consecutive inner-stages, for example, 730 for monthly inner stages. To ease notation, we use the shorthand w_i to denote $w_{i,t=0}$ in the following. The complete stage i MDP state is represented by the pair $(x_i, w_i) \in \mathcal{X}_i \times \mathcal{W}_i$, which can be high dimensional in realistic settings involving long-tenor contracts and stochastic models for the evolution of the uncertainties that require vector information. In our numerical experiments, for example, the state space has 51 continuous dimensions: (i) the endogenous part of the state x_i has 25 dimensions because our longest VPPA has a tenor of 25 years (i.e., $M = 25$), and (ii) the exogenous component w_i has 26 dimensions because the strike price dynamics depend on a latent variable, 12 electricity prices, and 12 REC prices (totaling 25 factors), and an additional factor drives uncertain energy supply.

2.4. Transition and Reward Functions

Executing procurement decisions $z_i \in \mathcal{Z}_i(a_i)$ at stage i and state $(x_i, w_i) \in \mathcal{X}_i \times \mathcal{W}_i$ results in an update of the VPPA capacity vector x_i to

$$x_{i+1,l} = f_i(x_i, z_i)_l = \begin{cases} x_{i,l+1} + \sum_{m \in \mathcal{M}: m > l} z_{i,m}, & \text{if } l \in \{0, \dots, M-2\}; \\ z_{i,M}, & \text{if } l = M-1, \end{cases}$$

where $f_i(x_i, z_i)$ is a vector transition function and $f_i(x_i, z_i)_l$ represents its l th element. We assume that procurement decisions do not affect the evolution of the vector of stochastic factors w_i , which evolves according to a Markovian stochastic process. In particular, w_i at stage i and this process determine the distribution of w_j at a future stage j ($> i$) and thus the distributions of supply and the prices of electricity, RECs, and VPPAs. We describe in Section 5 a specific choice for this stochastic process.

For each stage $i < I-1$, the expected procurement cost accrued when entering into VPPAs is

$$c_i(x_i, w_i, z_i) = \mathbb{E}_i \left[\sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \gamma_a^l \gamma^t (K_{i,m} - P_{i+l,t}) \eta_{i+l,t} z_{i,m} + \sum_{t \in \mathcal{T}} \gamma^t P_{i,t} d_{i,t} + \gamma_a R_{i+1} \max\{\alpha d_i - \eta_i x_{i,0}, 0\} \mathbf{1}_{\{i \in \mathcal{I}^S\}} \right], \quad (2)$$

where $\mathbb{E}_i[\cdot] \equiv \mathbb{E}[\cdot | w_i]$, $L_{i,m} := \min\{m, I-i\}$ equals the number of stages of energy delivery within the planning horizon, $\gamma \in [0, 1)$ denotes the discount factor over the interval defined by inner stages, $\gamma_a = \gamma^T \in [0, 1)$ is the annualized discount factor, and $d_i := \sum_{t \in \mathcal{T}} d_{i,t}$ and $\eta_i := \sum_{t \in \mathcal{T}} \eta_{i,t}$ represent the aggregate

(annual) electricity demand and supply, respectively. The first term in (2) models the sum of the settlement cash flows between the generator and the company at each interval over the tenor of the VPPAs signed from set \mathcal{M} . Here $\eta_{i+l,t} z_{i,m}$ is the energy generated in the interval $(\tau_{i,t}, \tau_{i,t+1}]$ of stage $i+l$ by the $z_{i,m}$ MW VPPA signed at stage i with tenor m . The second term is the cost of procuring the known demands $d_{i,t}$, $t \in \mathcal{T}$, from the short-term market at prices $P_{i,t}$, $t \in \mathcal{T}$. The third term accounts for the REC purchasing costs at a price R_{i+1} USD/MWh to cover any shortfall between the target αd_i in year $i \in \mathcal{I}^S$ and the renewable energy supply $\eta_i x_{i,0}$ in this year. At stage $I-1$ only short-term procurement of electricity and RECs is possible, which means

$$c_{I-1}(x_{I-1}, w_{I-1}) = \mathbb{E}_{I-1} \left[\sum_{t \in \mathcal{T}} \gamma^t P_{I-1,t} d_{I-1,t} + \gamma_a R_I \max\{\alpha d_{I-1} - \eta_{I-1} x_{I-1,0}, 0\} \right], \quad (3)$$

and a zero cash flow is associated with the terminal stage I .

2.5. Policies

A stage i dynamic procurement policy π_i is a collection of stage-dependent decision rules $\{Z_j^{\pi_i}, j \in \mathcal{I}_i\}$, each mapping states to actions, where $\mathcal{I}_i := \{i, \dots, I-2\}$. A decision rule $Z_j^{\pi_i}$ in stage i is feasible if it associates with each state $(x_i, w_i) \in \mathcal{X}_i \times \mathcal{W}_i$ an action $z_i(x_i, w_i)$ that belongs to set $\mathcal{Z}_i(a_i)$. We denote by Π_i the set of all feasible stage i policies. Given an initial state (x_i, w_i) in stage i , an optimal procurement policy in Π_i minimizes the expected electricity costs

$$V_i(x_i, w_i) := \min_{\pi_i \in \Pi_i} \mathbb{E}_i \left[\sum_{j \in \mathcal{I}_i} \gamma_a^{j-i} c_j(x_j^{\pi_i}, w_j, Z_j^{\pi_i}(x_j^{\pi_i}, w_j)) + \gamma_a^{I-1-i} c_{I-1}(x_{I-1}^{\pi_i}, w_{I-1}) \right], \quad (4)$$

where $V_i(x_i, w_i)$ is the MDP value function at stage i and state (x_i, w_i) and $x_j^{\pi_i}$ is the endogenous state reached in stage j when following the policy π_i starting from (x_i, w_i) . The decisions associated with an optimal policy are always finite because the set $\mathcal{Z}_i(a_i)$ is bounded. The objective function involves an expectation of a nonlinear function of the state because of the max term in the cost function $c_i(x_i, w_i, z_i)$ and the dependence of the decision rule on the state. Therefore, the objective depends on the properties of the Markovian process driving the evolution of stochastic factors.

An optimal policy of MDP (4) maintains a dynamically evolving portfolio of VPPAs to meet a target. The MDP policy decision at a stage i state accounts for

(i) the future stochasticity of prices and supply, (ii) the full flexibility to time VPPA purchases, and (iii) sourcing flexibility, that is, the access to a variety of VPPA types. Although these properties are appealing, computing the optimal MDP policy is challenging because of the well-known curses of dimensionality associated with MDPs (Bertsekas 2011, Powell 2011). This is indeed the case for MDP (4) in our numerical study as previously discussed. A common strategy to overcome this intractability relies on approximating the value function of the MDP, where convexity of the value function plays an important role in approximate dynamic programming methods (Brown and Smith 2014, Salas and Powell 2017, Nadarajah and Secomandi 2018). The value function of MDP (4) is, however, not convex in general as we show via an example in Online Appendix B. Ensuring convexity requires additional conditions, such as no procurement minimums (i.e., $z_m^{\min} = 0$), that may not hold in practice, as shown in Proposition 1.

Proposition 1. Suppose (i) $z_m^{\min} = 0$ and $z_m^{\max} < \infty$ for all $m \in \mathcal{M}$, and (ii) price expectations are finite, that is, $\mathbb{E}_i[[P_{j,t}]] < \infty$ and $\mathbb{E}_i[[R_{j+1}]] < \infty$ for each $(i, j, t) \in \{0, \dots, I-1\} \times \{i, \dots, I-1\} \times \mathcal{T}$. Then the value function $V_i(\cdot, w_i)$ is convex in the endogenous state x_i for each $(i, w_i) \in \mathcal{I} \times \mathcal{W}_i$.

Thus, in addition to dimensionality issues, another source of intractability in MDP (4) stems from the nonconvex structure of the action space $\mathcal{Z}_i(a_i)$ and of the value function as a consequence. Because of the difficulties of solving this MDP, we pursue simpler strategies. In Section 3, we consider policies based on forecasts that do not account for future uncertainty. We then introduce in Section 4 policies based on information relaxation and duality (Andersen and Broadie 2004, Haugh and Kogan 2004, Brown et al. 2010) that additionally account for stochasticity when obtaining decisions at a given stage and state.

3. Forecast-Based Reoptimization Heuristic

In this section, we introduce the forecast-based policies. We first present the models used to make procurement decisions in Section 3.1 and provide guidelines on how to implement such policies in Section 3.2.

3.1. Decision Models

The traditional use of long-tenor VPPAs in practice can be viewed as a specific case of a reoptimization heuristic that uses a single VPPA. This VPPA is renewed in a rolling fashion (WBCSD 2018) by solving a deterministic model derived from MDP (4) by only allowing VPPAs of a single type to be signed just-in-

time to meet the target, replacing uncertain quantities by forecasts, and limiting these forecasts to the tenor of the VPPA type; that is, they may not extend to the end of the planning horizon. This heuristic, dubbed the *forecast-based block heuristic* (FBH_m), is parameterized by a fixed $m \in \mathcal{M}$ and solves the following model to determine the optimal VPPA capacity at stage $i < I-1$:

$$z_{i,m}^* = \arg \min_{z_{i,m} \in \mathcal{Z}_i(a_i)} \sum_{l=1}^{L_{i,m}} \gamma_a^l \left[\sum_{t \in \mathcal{T}} \gamma^t (K_{i,m} \mathbb{E}_i[\eta_{i+l,t}] - \mathbb{E}_i[P_{i+l,t} \eta_{i+l,t}]) z_{i,m} + \sum_{t \in \mathcal{T}} \gamma^t \mathbb{E}_i[P_{i+l,t}] d_{i+l,t} + \mathbb{E}_i[R_{i+l+1}] \max\{\alpha d_{i+l} - \mathbb{E}_i[\eta_{i+l}] z_{i,m}, 0\} \mathbf{1}_{\{i+l \in \mathcal{T}^S\}} \right]. \quad (5)$$

Compared with the optimal policy of MDP (4), FBH_m has low timing flexibility and zero sourcing flexibility (i.e., it uses VPPAs of a single type) because new VPPAs are signed only when the incumbent contract expires, and it uses a single VPPA type. We introduce a rolling planning approach that has full timing and sourcing flexibilities and thus extends FBH_m. This approach, dubbed *forecast-based reoptimization heuristic* (FRH), solves at each stage a model obtained by replacing random quantities in MDP (4) by their respective forecasts, that is, expected values. In the case of contract availability, which is a binary variable, we assign a forecast of one if the contract is available with probability greater than 0.5 and zero otherwise. Formally, the stage j forecast for contract $m \in \mathcal{M}$ made at stage i , with $i \leq j$, is defined as $\bar{a}_{i,j,m} = 1$, if $\mathbb{E}_i[a_{j,m}] > 0.5$, and $\bar{a}_{i,j,m} = 0$ otherwise. Consider the following version of the cost function (2) with random quantities $\eta_{i+l,t}$, $P_{i+l,t} \eta_{i+l,t}$, R_{i+1} , and η_i replaced by their respective expectations:

$$c_i^{\text{DET}}(x_i, w_i, z_i) = \sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \gamma_a^l \gamma^t (K_{i,m} \mathbb{E}_i[\eta_{i+l,t}] - \mathbb{E}_i[P_{i+l,t} \eta_{i+l,t}]) z_{i,m} + \sum_{t \in \mathcal{T}} \gamma^t \mathbb{E}_i[P_{i,t}] d_{i,t} + \gamma_a \mathbb{E}_i[R_{i+1}] \max\{\alpha d_i - \mathbb{E}_i[\eta_i] x_{i,0}, 0\} \mathbf{1}_{\{i \in \mathcal{T}^S\}}.$$

We omit the expression for c_{I-1}^{DET} as it is an analogous simplification of the terminal cost (3).

At stage i and state $(x_i, w_i) \in \mathcal{X}_i \times \mathcal{W}_i$, FRH solves

$$\min_{y_j, z_j} \sum_{j \in \mathcal{I}_i} \gamma_a^{j-i} c_j^{\text{DET}}(y_j, \mathbb{E}_i[w_j], z_j) + \gamma_a^{I-1-i} c_{I-1}^{\text{DET}}(y_{I-1}, \mathbb{E}_i[w_{I-1}]) \quad (6a)$$

$$\text{s.t.} : y_i = x_i, \quad (6b)$$

$$y_{j+1} = f_j(y_j, z_j), \quad \forall j \in \mathcal{I}_i, \quad (6c)$$

$$y_j \in \mathcal{X}_j, \quad \forall j \in \mathcal{I}_i \cup \{I-1\}, \quad (6d)$$

$$z_j \in \mathcal{Z}_j(\bar{a}_{i,j}), \quad \forall j \in \mathcal{I}_i. \quad (6e)$$

This math program computes decisions $z_j \in \mathcal{I}_i$ and includes auxiliary variables y_j to track the

endogenous MDP state. The decision at stage j is static, that is, it does not depend on the state at this future stage. This is in contrast to the decision rule at stage j in the MDP, which is a function of the state. Objective (6a) can be viewed as a forecast of the procurement costs under static decisions. Constraint (6b) initializes the stage i state to the current state x_i . Constraints (6c) ensure the feasibility of state transitions. Constraints (6d)–(6e) restrict decision variables to their respective feasible sets.

We also consider a variant FRH_m of FRH that uses a single VPPA type $m \in \mathcal{M}$, which has full timing flexibility but no sourcing flexibility. The decision model corresponding to FRH_m is analogous to (6) but involves only the variables and the objective function terms corresponding to contract m .

3.2. Implementation Guidelines

The objective functions of the decision models associated with FBH_m , FRH_m , and FRH involve expectations. In general, they can be estimated using Monte Carlo simulation, which is the strategy we use in our numerical experiments. Under certain stochastic processes for the evolution of uncertainties, expectations may also be available in closed form.

The FBH_m model (5) is not solved each year. To elaborate, it is first solved at the year preceding the target date to potentially enter into a VPPA that delivers renewable electricity for m years from this date. The decision model is needed again in the year before the first contract expires to potentially sign a second contract. This process is repeated until the end of the planning horizon. If contract m is not available at year i , that is, $a_{i,m} = 0$, then the target in $i + 1$ is fulfilled by unbundled RECs and the VPPA procurement decision is postponed to year $i + 1$.

In contrast to FBH_m , the FRH_m and FRH decision models are solved each year. We elaborate on how the FRH policy is constructed next. Although solving (6) at a given stage i and state (x_i, w_i) provides procurement decisions $\{z_j^*, j \in \mathcal{I}_i\}$, we only implement z_i^* corresponding to the current stage. Implementing this decision results in a transition to a new inventory of power $x_{i+1} = f_i(x_i, z_i^*)$. Once new information w_{i+1} becomes available at stage $i + 1$, we recompute the expectations of uncertain quantities and solve an analogue of math program (6) formulated using these updated expectations at state (x_{i+1}, w_{i+1}) to obtain z_{i+1}^* . We repeat this procedure until we reach stage $I - 1$. The construction of the FRH_m follows the same logic but using its decision model.

The computational burden of using forecast-based policies is tied to the difficulty of solving the

deterministic form of our MDP or a simplification thereof, which is a favorable algorithmic property, especially in the presence of the nonconvexities discussed in Section 2. Both (5) and (6) are mixed-integer programs that can be solved using off-the-shelf commercial optimization software. Conceptually, the FRH_m and FRH decisions, although reoptimized, are based on static decision models that precludes future decisions to dynamically adapt to new information. It is a priori unclear how this factor affects the policy performance of these forecast-based rolling planning approaches. In addition, the optimal objective of the FRH math program does not provide a lower bound on the optimal policy cost whereas deterministic (convex/fluid) approximations in other applications are known to provide such an optimistic bound (Gallego and Van Ryzin 1997). We show in Online Appendix C that such a lower bounding property is not true for even a “convexified” version of the FRH math program.

4. Information-Relaxation–Based Reoptimization Heuristic

We propose a heuristic based on the information relaxation and duality framework (Andersen and Broadie 2004, Haugh and Kogan 2004, Brown et al. 2010) that addresses the shortcomings of forecast-based policies discussed in Section 3 while retaining the favorable implementation property of solving deterministic math programs. We describe the dual bound in Section 4.1, present the decision model used to define the new policy in Section 4.2, and provide implementation guidelines in Section 4.3.

4.1. Dual Bound

Information relaxation and duality is a useful framework to obtain dual bounds on the optimal policy cost of intractable MDPs and is applicable to MDP (4) as well. In its most commonly used form, a dual bound is estimated in Monte Carlo simulation by solving a deterministic variant of MDP (4) endowed with full information about future uncertainty and costs adjusted for this knowledge using a dual penalty. Let $q_i(x_i, W_i, z_i)$ and $q_{I-1}(x_{I-1}, W_{I-1})$ denote, respectively, the stage $i < I - 1$ and stage $I - 1$ dual penalty function, where $W_i := (w_{j,t}, j = i \dots, I, t \in \mathcal{T})$ is a vector of realized stochastic factors for each stage from i to I and interval $t \in \mathcal{T}$. If the inequalities $\mathbb{E}_j[q_j(x_j, W_j, z_j)] \geq 0$, for $j \in \mathcal{I}_i$, and $\mathbb{E}_{I-1}[q_{I-1}(x_{I-1}, W_{I-1})] \geq 0$ hold for all feasible (i.e., Markovian/nonanticipative) policies, then the dual penalty function is feasible. Given knowledge of W_i , define the following hindsight cost function (we omit the analogous definition for c_{I-1}^{IR})

$$\begin{aligned}
 c_i^{\text{IR}}(x_i, W_i, z_i) &= \sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \gamma_a^l \gamma^t (K_{i,m} - P_{i+l,t}) \eta_{i+l,t} z_{i,m} \\
 &\quad + \sum_{t \in \mathcal{T}} \gamma^t P_{i,t} d_{i,t} \\
 &\quad + \gamma_a R_{i+1} \max\{\alpha d_i - \eta_i x_{i,0}, 0\} \mathbf{1}_{\{i \in \mathcal{T}^S\}},
 \end{aligned} \tag{7}$$

and consider the following deterministic optimization problem:

$$\begin{aligned}
 V_i^{\text{IR}}(x_i; W_i) &= \min_{y_j, z_j} \sum_{j \in \mathcal{I}_i} \gamma_a^{j-i} [c_j^{\text{IR}}(y_j, W_j, z_j) - q_j(y_j, W_j, z_j)] \\
 &\quad + \gamma_a^{I-1-i} [c_{I-1}^{\text{IR}}(y_{I-1}, W_{I-1}) - q_{I-1}(y_{I-1}, W_{I-1})]
 \end{aligned} \tag{8a}$$

$$\text{s.t. :} \quad y_i = x_i, \tag{8b}$$

$$y_{j+1} = f_j(y_j, z_j), \quad \forall j \in \mathcal{I}_i, \tag{8c}$$

$$y_j \in \mathcal{X}_j, \quad \forall j \in \mathcal{I}_i \cup \{I-1\}, \tag{8d}$$

$$z_j \in \mathcal{Z}_j(a_j), \quad \forall j \in \mathcal{I}_i. \tag{8e}$$

Constraints (8b)–(8d) are identical to Constraints (6b)–(6d) in the math program solved by FRH. Constraints (8e) differ from (6e) in the availability vector used to define $\mathcal{Z}_j(\cdot)$. In the former case, we use the realization of the random contract availability vector on a given sample path, whereas in the latter case we use the contract availability forecast vector described in Section 3. Objective (8a) can be obtained by modifying FRH Objective (6a) by subtracting dual penalty terms and replacing the forecasted uncertainty with the elements of W_i . The expectation $\mathbb{E}_i[V_i^{\text{IR}}(x_i; W_i)]$ taken with respect to the random variable $W_i|w_i$ defines a dual bound on the value function $V_i(x_i, w_i)$, that is, the optimal policy value starting from stage i and state (x_i, w_i) .

The quality of the dual bound depends on the choice of the dual penalty function in math program (8). Choosing this function to be zero, that is, $q_i(\cdot, \cdot, \cdot) \equiv 0$, results in the dual bound being equivalent to the well-known perfect information bound, which can be weak. Brown et al. (2010) show that the dual bound is instead equal to the optimal policy value when using the following ideal dual penalty based on the MDP value function:

$$q_{I-1}(x_{I-1}, W_{I-1}) = c_{I-1}^{\text{IR}}(x_{I-1}, W_{I-1}) - c_{I-1}(x_{I-1}, W_{I-1}), \tag{9a}$$

$$\begin{aligned}
 q_i(x_i, W_i, z_i) &= \gamma_a \{V_{i+1}(f_i(x_i, z_i), w_{i+1}) \\
 &\quad - \mathbb{E}_i[V_{i+1}(f_i(x_i, z_i), w_{i+1})]\} \\
 &\quad + c_i^{\text{IR}}(y_i, W_i, z_i) - c_i(y_i, w_i, z_i).
 \end{aligned} \tag{9b}$$

The term $c_i^{\text{IR}}(y_i, W_i, z_i) - c_i(y_i, w_i, z_i)$ is atypical in the definition of an ideal penalty but needed here because our MDP Cost Function (2) includes expectations over future exogenous states; hence, it differs from its hindsight counterpart (7). Because the exact value function

in (9b) is not available, value function approximations can be used instead or simpler dual penalties that do not rely on such approximations can be constructed (Brown and Smith 2014, Secomandi 2015, Nadarajah and Secomandi 2018).

4.2. Decision Model

Traditionally, a decision is computed independent of the dual bound computation described previously (Desai et al. 2012 and the related discussion in Section 1.1). We now define a nonanticipative decision directly during the dual bound estimation process, where being nonanticipative refers to a decision that only depends on the information available at stage i (the decision resulting from solving (8) is anticipative as it relies on future information on the sample path W_i). Because the dual bound $\mathbb{E}_i[V_i^{\text{IR}}(x_i; W_i)]$ involves solving the math program (8) over multiple realizations of the random variable $W_i|w_i$, we also have a distribution of optimal solutions of this math program. This decision distribution encodes information from the evolution of future prices and supply uncertainties. We focus on stage i decisions obtained during the dual bound estimation process and represent them via a random decision $z_i(W_i)$ that is a function of the random variable $W_i|w_i$. Our key idea is to define a functional that operates on the distribution of the random variable $z_i(W_i)$ and returns a single nonanticipative decision. We call this functional a *decision measure* \mathcal{H}_i as it maps a distribution $z_i(W_i)$ to a vector of $\mathbb{R}^{|\mathcal{M}|}$. Moreover, we refer to $\mathcal{H}_i(z_i(W_i))$ as the *information-relaxation based reoptimization heuristic* (IRH) decision. Proposition 2 establishes some useful properties of the IRH decision with respect to optimality and feasibility.

Proposition 2. *The decision $\mathcal{H}_i(z_i(W_i))$ is guaranteed to be feasible, that is, $\mathcal{H}_i(z_i(W_i)) \in \mathcal{Z}_i(a_i)$, if any of the following conditions hold:*

(1) *The feasible decision set $\mathcal{Z}_i(a_i)$ is convex and the decision vector $\mathcal{H}_i(z_i(W_i))$ is the mean of $z_i(W_i)$.*

(2) *For each $m \in \mathcal{M}$, the m -th element of the decision vector $\mathcal{H}_i(z_i(W_i))$ equals the m th element of $z_i(W_i)$ for some realization W_i of the random variable $W_i|w_i$.*

Moreover, if the decision $\mathcal{H}_i(z_i(W_i))$ satisfies one of the previous conditions and we use the ideal dual penalty (9) in math program (8), then $\mathcal{H}_i(z_i(W_i))$ is an optimal solution to MDP (4).

The convexity of $\mathcal{Z}_i(a_i)$ in Condition (1) is important to ensure that the mean of $z_i(W_i)$ is a feasible decision. If $\mathcal{Z}_i(a_i)$ is nonconvex, which is the case for a strictly positive procurement minimum $z_m^{\text{min}} > 0$ in (1), feasibility of the mean can no longer be guaranteed. To see this, the realizations of $z_i(W_i)$ are either zero or lie in the interval $[z_m^{\text{min}}, z_m^{\text{max}}]$. Although each realization belongs to $\mathcal{Z}_i(a_i)$, averaging over the aforementioned

zero and strictly positive values can lead to a mean of $z_i(W_i)$ that is strictly less than z_m^{\min} . An extra step is needed to recover feasibility in this case as discussed in Section 4.3. In contrast, a decision measure satisfying condition (2) in Proposition 2 gives a feasible decision even when $\mathcal{Z}_i(a_i)$ is nonconvex. This is because each component of the decision vector is directly selected from a realization of $z_i(W_i)$, which is feasible. Such component-wise construction leads to feasibility in our application since there are no constraints on the action space $\mathcal{Z}_i(a_i)$ that tie different contracts together. An example of \mathcal{H}_i that satisfies condition (2) is $\mathcal{H}_i(z_i(W_i)) = \mathbb{M}[z_i(W_i)|w_i]$, where \mathbb{M} denotes the component-wise median of a multivariate distribution (Lopuhaa and Rousseeuw 1991); that is, the standard univariate median applied to each dimension of the distribution (in our case, each contract type $m \in \mathcal{M}$). Regardless of which condition in Proposition 2 is true, the IRH decision $\mathcal{H}_i(z_i(W_i))$ is impacted by the unfolding of uncertainty because it is a function of the distribution of $z_i(W_i)$: a favorable property shared by the MDP optimal policy but not by heuristics in Section 3 that use forecasts, where decisions are based on static models. The optimality of $\mathcal{H}_i(z_i(W_i))$ under an ideal dual penalty shows that the information about future uncertainty encoded in the IRH decision distribution can be useful.

A single contract variant IRH_m of IRH can be defined in a similar manner using a decision model analogous to (8) that involves only the variables and the objective function terms corresponding to contract m .

4.3. Implementation Guidelines

Implementing the IRH and IRH_m policies requires (i) a dual penalty function q_i and (ii) a decision measure \mathcal{H}_i . We provide some guidelines regarding the two choices and the resulting rolling planning approach, which are adopted in our numerical experiments.

Because computing a value function approximation is complex in our MDP due to a high-dimensional state space and nonconvexities, we propose defining a dual penalty function of the form

$$q_i(x_i, W_i, z_i) := \sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \sum_{b=1}^B \gamma_a^l \gamma^t \beta_{i,b} \times (\mathbb{E}_i[w_{i+l,t,b}] - w_{i+l,t,b}) z_{i,m}, \quad (10)$$

where B is the number of stochastic factors (i.e., the number of elements in vector $w_{i,t}$), $w_{i,t,b}$ denotes the b th component of $w_{i,t}$, and $\beta_{i,b}$ are stage- and factor-dependent weights.⁸ Intuitively, Dual Penalty (10) accounts for the information gained when taking a decision by spreads between the values taken by stochastic factors in future periods and their expectations computed at period i . These spreads are multiplied by

the VPPA purchase decisions. Representation (10) allows the dual penalty to depend on the factors in the stochastic processes driving the evolution of prices and supply. For our numerical experiments, we use a specification of (10) detailed in Online Appendix E that involves the power and REC prices, uncertain supply, a factor associated with the strike price stochastic process, and weights that are only factor dependent.

Dual Penalty (10) is linear in z_i , does not depend on x_i , and is feasible because the stage i expectation of $\mathbb{E}_i[w_{i+l,t,b}] - w_{i+l,t,b}$ equals zero. This linear form ensures that the math program (8) falls into the same complexity class as the deterministic version of MDP (4) and the math program (6) solved by FRH.⁹ The weights in the penalty can be defined using a grid search or a simple local search heuristic guided by the information relaxation dual bound objective. The latter approach, which we adopt in our numerical study, is described in Online Appendix E.

To compute a dual bound and decision at stage i and state (x_i, w_i) , we generate H Monte Carlo sample paths of uncertainty $\{w_{j,t}^h, (j, t, h) \in \{i, \dots, I\} \times \mathcal{T} \times \{1, \dots, H\}\}$, which provide a discrete approximation $\widehat{W}_i|w_i$ of the random variable $W_i|w_i$. Based on this approximation, we estimate both a dual bound $\sum_{h=1}^H V_i^{\text{IR}}(x_i; W_i^h)/H$ and a decision $\mathcal{H}_i(z_i(\widehat{W}_i))$. The mean and the median introduced in Section 4.2 are natural choices for the decision measure. Implementing either of them requires the solution of H math programs of Type (8), which are mixed-integer programs under a linear dual penalty. Applying the median yields a feasible decision by Proposition 2. Instead, the mean may lead to $z_{i,m} \in (0, z_m^{\min})$ for some $m \in \mathcal{M}$, which is infeasible. Feasibility can be restored by projecting the infeasible decision z_i onto the feasible set $\mathcal{Z}_i(a_i)$ in different ways. A simple approach is to replace each component $z_{i,m}$ by the closest value between zero and z_m^{\min} . Alternatively, one can consider the capacity profile of a decision. Given an infeasible decision vector z_i , the resulting capacity available l periods in the future is $\sum_{m \in \mathcal{M}: m > l} z_{i,m}$. The capacity profile associated with z_i is defined as the vector of such capacities: $(\sum_{m \in \mathcal{M}: m > l} z_{i,m}, l = 1, 2, \dots, M)$. A feasible solution z'_i can be defined such that its capacity profile $(\sum_{m \in \mathcal{M}: m > l} z'_{i,m}, l = 1, 2, \dots, M)$ is close to the one of z_i , where closeness is measured using a norm. We take this approach in our numerical study and use a one norm to measure distance as it results in solving a small mixed-integer program to find z'_i .

The IRH (or IRH_m) policy applies the IRH (or IRH_m) decision from our decision measure to move to an endogenous state $x_{i+1} = f_i(x_i, \mathcal{H}_i(z_i(\widehat{W}_i)))$. Then we observe the stage $i + 1$ uncertainty w_{i+1} and repeat the same process at state (x_{i+1}, w_{i+1}) and keep moving forward in time until we reach the terminal stage.

5. Strike Price Dynamics

The VPPA strike price, the prices of electricity and RECs, and electricity supply evolve stochastically over time. The literature contains reduced-form stochastic processes for the last three quantities (Weron 2014, Zeng et al. 2015, Loukatou et al. 2018), but we are not aware of stochastic processes for the strike price. An open-source tool to help with determining the strike price is the system advisory model (SAM), designed by the National Renewable Energy Laboratory. SAM computes a VPPA strike price based on the generator’s net present value (NREL 2017) and requires generator-specific inputs (e.g., expected capacity factor, tax credits, cash incentives, expected return on investments, technology improvements, degradation rates). Because this information is not available to a corporate buyer, especially across heterogeneous generators, SAM is not easy to leverage in our procurement model. We therefore propose a stochastic model for VPPA strike prices.

Intuitively, the strike price of a VPPA $m \in \mathcal{M}$ of length $L_{i,m}$ signed in stage $i \in \mathcal{I}$ would account for the following discounted sum of expected prices of electricity and RECs over the tenor of the agreement as a baseline price:

$$\text{BP}_{i,m} = \frac{1}{\sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \gamma_a^l \gamma^t} \sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \gamma_a^l (\gamma^t \mathbb{E}_i [P_{i+l,t}] + \mathbb{E}_i [R_{i+l+1}]).$$

Here the constant $1/[\sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \gamma_a^l \gamma^t]$ annualizes the sum $\sum_{l=1}^{L_{i,m}} \sum_{t \in \mathcal{T}} \gamma_a^l (\gamma^t \mathbb{E}_i [P_{i+l,t}] + \mathbb{E}_i [R_{i+l+1}])$ so that the baseline can be directly compared with the strike price $K_{i,m}$ in USD per MWh. The baseline $\text{BP}_{i,m}$ is the expected cost of the outside option to procure electricity and RECs from short-term markets without using a VPPA. Proposition 3 shows that it is necessary for the strike price $K_{i,m}$ to be lower than $\text{BP}_{i,m}$ for the corresponding VPPA to be a part of an optimal procurement policy.

Proposition 3. *Suppose $K_{i,m} > \text{BP}_{i,m}$ for some $i \in \mathcal{I}$ and $m \in \mathcal{M}$. Then $z_{i,m}$ equals zero in an optimal policy to MDP (4).*

The intuition behind this proposition is the following. Compared with short-term procurement, signing a VPPA entails a supply disadvantage as the shape and uncertainty of renewable energy supply may result in overprocurement of RECs. To be attractive, a VPPA should compensate for this supply disadvantage with a price advantage. In general, $K_{i,m}$ could be above or below $B_{i,m}$.

Using $\text{BP}_{i,m}$ as a baseline to model tenor-dependent strike prices is currently challenging because of the lack of public strike price data. We leverage a recent

data set from BerkeleyLab (2020a, b), which does not specify the tenor of each contract, but reports execution date, power market, capacity, and a levelized strike price of roughly 800 wind and solar VPPAs signed in the United States. The levelized VPPA strike price accounts for the full revenue available over the contract tenor, including subsidies and other project specific features, and then converts the revenue stream to a USD per MWh value. Considering the available data, we develop a stochastic model for the levelized strike price, denoted \widehat{K}_i , that can be calibrated to public data and then scale it by a tenor-dependent premium ξ_m to obtain a tenor-dependent strike price:

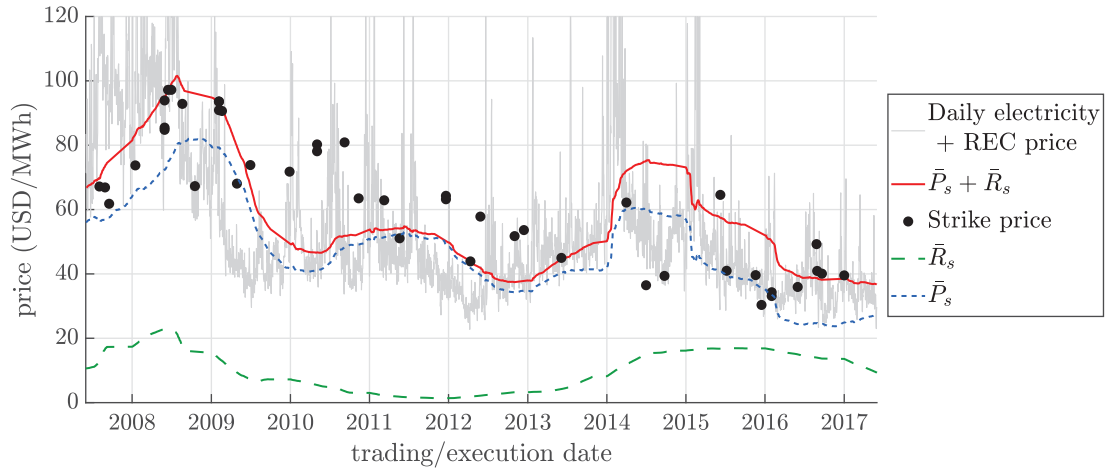
$$K_{i,m} := \widehat{K}_i \cdot (1 + \xi_m). \quad (11)$$

The structure of our levelized strike price model is motivated by an analysis of data on the VPPA price, electricity price, and REC price from BerkeleyLab (2020a, b), PJM (2020), and MarexSpectron (2020), respectively. Figure 2 displays the levelized strike price of VPPAs signed in Pennsylvania New Jersey Maryland Interconnection (PJM) by their execution date s , average of monthly PJM electricity prices for the year preceding this date (\bar{P}_s), an analogous average of New Jersey Class I RECs (\bar{R}_s), the sum of these two averages ($\bar{P}_s + \bar{R}_s$), and the significantly more volatile sum of daily prices. As this figure shows, $\bar{P}_s + \bar{R}_s$ tracks the behavior of the strike price data well. In addition, we observe that the strike price can be higher or lower than this baseline but appears to mean-revert around this level, that is, periods where the VPPA strike prices are greater than the baseline are followed by a period where these prices are more likely to be below the baseline. Consistent with these model-free observations, we specify the levelized strike price as

$$\widehat{K}_i := \bar{P}_i + \bar{R}_i + q_i U_i, \quad (12)$$

where $\bar{P}_i = (1/T) \sum_{t=1}^T P_{i-1,t}$ is the average of electricity prices over the past T periods, $\bar{R}_i = (1/T) \sum_{t=1}^T R_{i-1,t}$ an analogous average of REC prices, q_i a scalar, and U_i the time $\tau_{i,0}$ value of a mean-reverting latent variable in $[-1, 1]$, which we model as a Jacobi diffusion. The term $q_i U_i$ determines the quantity by which the levelized strike price is above or below the baseline. We define $q_i := \zeta^p \bar{P}_i + \zeta^r \bar{R}_i$ and require $\zeta^p \leq 1$ and $\zeta^r \leq 1$, which are sufficient conditions for \widehat{K}_i to be nonnegative when the electricity and REC prices are nonnegative.

We now overview the calibration of our strike price model (11) at a monthly time scale with T equal to 12, which we use in our numerical study, relegating details to Online Appendix D. In addition to the scalars ζ^p and ζ^r , and the parameters of the Jacobi

Figure 2. (Color online) Historical Data from PJM

diffusion underlying U_i , we need to specify stochastic processes for the electricity and REC prices that appear in the baseline term $\bar{P}_i + \bar{R}_i$. We model the monthly electricity price using a mean-reverting stochastic process with seasonality and jumps based on Lucia and Schwartz (2002) and Weron (2014). The evolution of monthly REC prices follows the Jacobi diffusion process in Zeng et al. (2015). We allow for correlations between the power price, the REC price, the latent variable U_i , and uncertain monthly supply η_i , which we model using mean-reverting stochastic process without jumps in Loukatou et al. (2018). We use maximum likelihood estimation to determine parameters, except for correlations, using the data underlying Figure 2 for prices and EIA (2019) for supply. The resulting p values of the estimates equal 0.01 or lower. We estimated correlations directly from data. The tenor-dependent premiums ξ_m need to be specified, perhaps based on guidance from practitioner reports, because data to calibrate this term are unavailable.¹⁰

6. Numerical Study

We numerically investigate the benefits of constructing dynamic VPPA portfolios using rolling power purchases, which involves comparing the policies summarized in Table 1. This table indicates for each

policy, the section in which it was previously introduced (column 2), how many VPPAs it considers if any (column 3), and whether its decision at a given stage and state accounts for future stochasticity (column 4), timing flexibility (column 5), and sourcing flexibility (column 6). We follow the implementation guidelines for these policies discussed in Sections 3.2 and 4.3.

We describe the procurement setting considered for our experiments in Section 6.1. We study the value of the advanced reoptimization underlying IRH in Section 6.2. We then use the IRH policy to assess the value of dynamic VPPA portfolios in Section 6.3 and the importance of timing and sourcing flexibilities in constructing such portfolios in Section 6.4. We perform sensitivity analysis in Section 6.5.

6.1. Procurement Setting

Table 2 summarizes the parameter values defining our procurement setting. We consider a 40-year planning horizon (I) and a target to reach in 5 years (I^R). Within each year, there are 12 monthly inner stages (T). The annual risk-free rate is set to 3.1% and corresponds to the average 10-year U.S. treasury rate in November 2018 (Bloomberg 2018), which implies a yearly discount factor γ_A equal to 0.97 and a monthly discount factor γ of 0.9975. We use a constant monthly

Table 1. Summary of Procurement Policies and Their Features

Policy	Section	VPPA	Uncertainty-aware decisions	Timing flexibility	Sourcing flexibility
Short-term	2	None	No	No	No
FBH _m	3	Single	No	No	No
FRH _m	3	Single	No	Yes	No
FRH	3	Portfolio	No	Yes	Yes
IRH _m	4	Single	Yes	Yes	No
IRH	4	Portfolio	Yes	Yes	Yes

Table 2. Parameters Defining the Procurement Setting

Name	Value	Unit
I	40	Years
I^R	5	Years
T	12	Months
γ_A	0.97	—
γ	0.9975	—
$d_{i,t}$	$50(\forall i, t \in \mathcal{I} \setminus \{I\} \times T)$	GWh/month
\mathcal{M}	$\{5, 10, 15, 20, 25\}$	Years
p_m	$\{0.4, 0.5, 0.6, 0.6, 0.5\}$	—
ξ_m	$\{5, 2.5, 0, -2.5, -5\}$	%
z_m^{\min}	$20(\forall m \in \mathcal{M})$	MW
z_m^{\max}	$400(\forall m \in \mathcal{M})$	MW

energy demand $d_{i,t}$ throughout the horizon which represents the consumption of two large data centers (Kamiya and Kvarnström 2019). The set of contracts \mathcal{M} and the minimum/maximum quantities are consistent with the VPPA portfolio of Google (Google 2017). We consider the availability of VPPA $m \in \mathcal{M}$ following a Bernoulli random variable, where $p_m \in [0, 1]$ denotes the probability that this contract is available at a given stage. The availability factors p_m are chosen based on Wisser and Bolinger (2017) and Baker McKenzie (2015): According to the first report, 15- to 25-year VPPAs are predominant, with 20-year contracts being the most common, whereas the second report indicates that VPPAs of length between 10 and 20 years are prevalent. We calibrate stochastic models for the evolution of prices and supply following the description at the end of Section 5 (see Online Appendix D for details). We specify VPPA risk factors ξ_m that decrease with m to reflect the following benefit to a producer: long-term (typically at least 15 years) agreements provide guaranteed revenue stream for considerable time, which the renewable power producer can use to secure financing. In particular, we use a maximum risk factor equal to $\xi_5 = 5\%$ for 5-year VPPAs, which decreases linearly as m is increased until $\xi_{25} = -5\%$ for 25-year VPPAs.

6.2. Value of Advanced Reoptimization

We compare FRH, which is based on a known forecast-based rolling power purchase strategy, with IRH, which uses more advanced reoptimization.

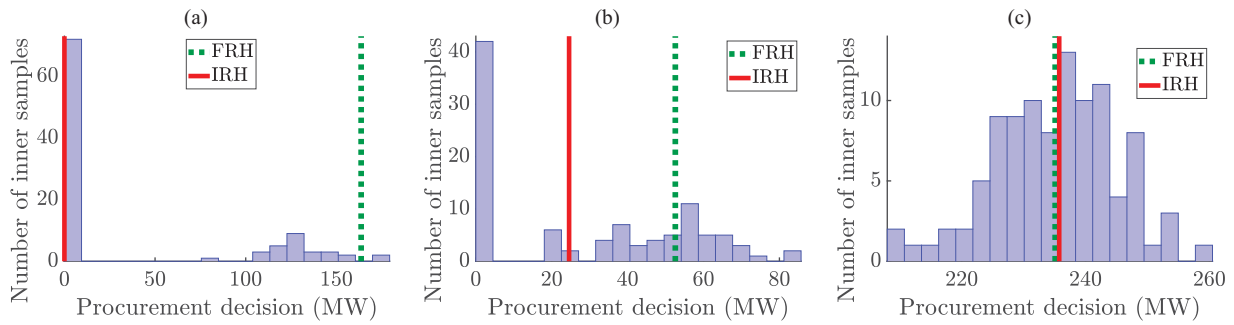
Table 3 reports the procurement cost and the optimality gap of the FRH policy and four specifications of IRH that differ in their choice of decision measure and dual penalty. Statistics are provided for instances with target levels (α) equal to 50% and 100%. We also report average statistics across 11 instances that vary this target from 0% to 100% in increments of 10%. Optimality gaps are computed using the dual bound with penalties (9), which is tighter than an analogous bound when using a zero penalty on all instances and, in particular, by 2.2% on average.

The best IRH configuration uses dual penalties and the median decision measure. Its optimality gap under a target of 50%/100% is 5.6%/7.0% compared with 9.8%/12% for FRH. The average IRH and FRH gaps are 5.6% and 9.6%, respectively. These improvements translate to average savings of 18.6 million USD when using IRH instead of FRH, which reach 24.3 million USD for a 100% target. This suggests significant value in the more advanced reoptimization underlying IRH, where a nonanticipative decision is extracted from a decision distribution that encodes information about the impact of future uncertainty on the current decision. In contrast, FRH solves a model based on forecasts and thus does not have such information. To understand the value of distributional information used by IRH, it helps to note that the FRH optimization model (6) is equivalent to a hindsight optimization model (8) formulated with a zero dual penalty on a sample path that equals the forecast. If the decision computed by this hindsight optimization is sensitive to changes in sample path information (i.e., prices and other uncertainties), then hindsight decisions on likely sample paths will be quite different from the FRH decision. In this case, the FRH decision can be highly suboptimal in hindsight on likely sample paths. In contrast, if one uses IRH with a median decision measure and zero dual penalty, its decision is “centered” on the distribution of decisions across sample paths, which mitigates the deviation of this decision from the hindsight decision on likely sample paths. Put differently, FRH can suffer from overfitting on the forecast (an “average” sample path), whereas IRH reduces this issue by having access to the decision distribution across sample paths. Figure 3

Table 3. Procurement Cost (Million USD) and Optimality Gap (%) of FRH and IRH Policies

Policy	Dual penalty	Decision measure	Procurement cost			Optimality gap		
			$\alpha = 0.5$	$\alpha = 1$	Average	$\alpha = 0.5$	$\alpha = 1$	Average
FRH	not applicable	not applicable	512.6	545.0	512.8	9.8	12.0	9.6
IRH	(9)	Average	496.8	529.0	499.0	6.4	8.7	6.7
IRH	(9)	Median	493.0	520.7	494.2	5.6	7.0	5.6
IRH	Zero	Average	499.5	535.4	502.4	7.0	10.0	7.4
IRH	Zero	Median	497.9	529.0	499.5	6.6	8.7	6.8

Figure 3. (Color online) Examples of FRH and IRH Median Decisions Taken from the Same State



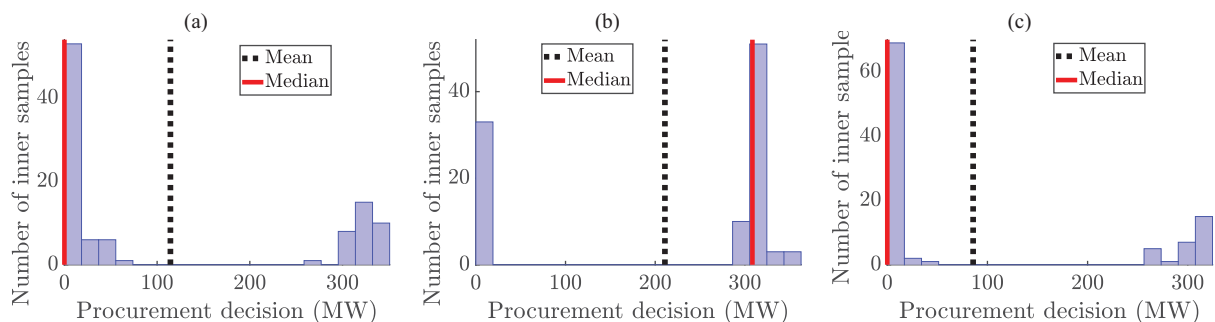
illustrates this behavior by simulating the FRH and IRH decisions just discussed by superimposing them on the decision distribution used by IRH. Each panel corresponds to a different state. In panels (a) and (b) of this figure, the hindsight decision is very sensitive to the sample path. For instance, in the former panel, the FRH decision is to enter a VPPA of roughly 160 MW, but such a large purchase only occurs very rarely. Most decisions in the distribution are either zero or smaller than this value and hence deviate substantially from the FRH decision that is overfitted to the forecast. Such overfitting need not always hurt as shown in panel (c). Here, all anticipative decisions lie in the interval [208,260] MW with more than half in the narrow range of [228,242] MW. In this case, the procurement decision taken by FRH and IRH are almost identical.

Another encouraging aspect of the results in Table 3 is that all four versions of IRH outperform FRH. Nevertheless, dual penalties and decision measures matter. Using Dual Penalty (9) in conjunction with a mean/median decision measure improves policy performance compared with the zero dual penalty case by 0.7%/1.1% on average. These penalties appear to have a favorable effect on the decision distribution used by the IRH policy, and this observation is in sync with Proposition 2, which shows the optimality of the decision distribution when using ideal penalties. The

median decision measure improves on the mean decision measure by 1.0% on average when using dual penalties. Further numerical investigation suggests that this improvement is because of the former measure being more robust to outliers than the latter one when the decision distribution is skewed. For example, a left-skewed distribution for a given contract type m arises at stages and states where most IRH sample path decisions correspond to either signing no VPPA of this tenor (i.e., zero capacity) or signing the minimum capacity, and only a few sample paths correspond to signing a large capacity or the maximum possible capacity. In this case, the median of the decisions is close to zero or the minimum capacity, whereas the mean gets considerably affected by the small number of samples with large capacity values. We observe a similar behavior when the decision distribution is right skewed. Figure 4 provides illustrative examples of such skewed distributions encountered in our numerical experiments.

The aforementioned policy improvements associated with the advanced reoptimization in IRH come at a computational cost because it involves solving at each stage math programs for each of the inner sample paths, whereas FRH solves a single math program. We used 30 inner samples to construct the IRH decision distribution and used a sample average approximation with 100 samples to estimate expectations in

Figure 4. (Color online) Illustrative Examples of Skewed IRH Decision Distributions



the FRH objective. For these choices, the time taken to compute a decision at a given stage and state using FRH and IRH is about 0.1 and 1 seconds, respectively.¹¹ We estimated the value of a policy over the planning horizon using 1,000 sample paths (i.e., H equals 1,000), which resulted in lower and upper bound estimates with standard errors that were at most 0.47% of their respective upper bound estimates. This policy evaluation required 46 minutes and 3.7 hours for FRH and IRH, respectively, with the latter CPU times being substantially larger than the former but still practical for our application.¹² These times also highlight the value of IRH's advanced reoptimization in solving large-scale MDPs with nonconvexities, which are difficult to tackle for common value function approximation strategies used in energy applications such as least squares Monte Carlo (Nadarajah et al. 2017). As discussed in Section 2, MDP (4) has a 51-dimensional continuous state space and a 5-dimensional action space on our instances.

6.3. Cost of Target and Relevance of VPPAs

Equipped with a near-optimal dynamic VPPA portfolio policy from the best IRH variant, we investigate the cost of a target and the role of VPPAs in managing this cost. We benchmark against the short-term policy (i.e., short-term power and unbundled REC purchases), which does not use VPPAs and is thus a simple strategy one can use in practice with no need for optimization.

Figure 5 displays the procurement costs as the target is varied from 0% to 100% in our setting. The cost of the target depends on the procurement policy. When using IRH, the cost of a 100% target is 521 million USD compared with a cost of 470 million USD with no target (i.e., α equals zero). The cost increases linearly with the target in this case at the rate of 0.51

million USD for each percentage. The short-term policy with VPPAs instead starts at 501 million USD with no target and incurs 603 million USD for a 100% target, with costs increasing linearly by 1.02 million USD for each percentage. These results show significant value in using VPPAs: Savings of roughly 50% (or 0.51 million USD) for each percentage increase in the target and an overall cost reduction of 82 million USD for a 100% target.

The risks associated with the short-term and IRH strategies help understand their relative performance. The former policy is not affected by supply uncertainty, but it does face price risk, whereas the latter policy can manage price risk by signing VPPAs but takes on supply risk. The substantial cost savings when using IRH over the short-term strategy suggests that price risk is mitigated using VPPAs and this benefit overshadows the risk because of power supply. Consistent with this reasoning, we find that the average total price of electricity and unbundled RECs is 45.5 USD/MWh for a 90% target, whereas the average (weighted on contract size and length) strike price of VPPAs signed by IRH is 33.5 USD/MWh, a substantial cost advantage indeed. Interestingly, the average of the strike prices over all available VPPAs (i.e., signed and unsigned) is 46.9 USD/MWh, which is slightly higher than the short-term power and REC price average. This indicates that the cost advantage of using dynamic VPPA portfolios is linked to the value of timing and sourcing flexibilities and the ability to optimize them in IRH, which we investigate in Section 6.4. We also considered an additional setting that allows the firm to sell RECs back to a secondary market when their quantity exceeds the target. The value of VPPAs changes by less than 1% on average in this case compared with the model we consider where REC sales are not possible, which is

Figure 5. (Color online) Procurement Costs of Short-Term and IRH Policies for Different Target Levels α

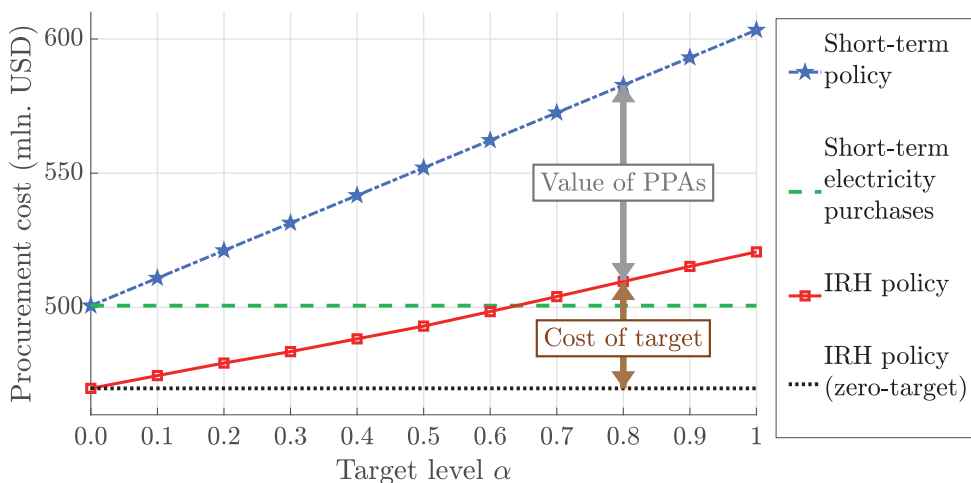
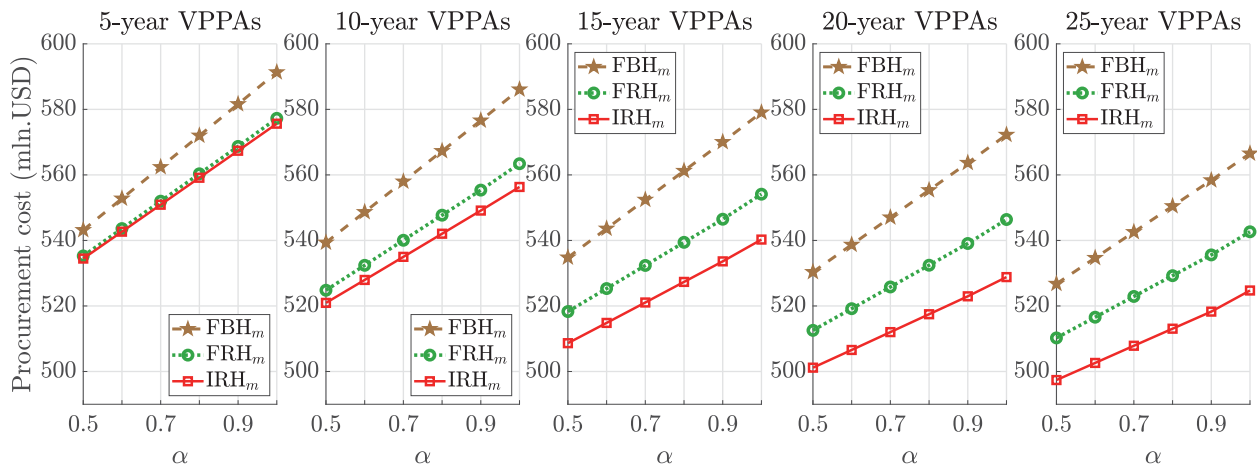


Figure 6. (Color online) Procurement Cost Under Single-Contract Policies from FBH_m , FRH_m , and IRH_m



encouraging. We provide further details of these experiments in Online Appendix F.

6.4. Value of Timing and Sourcing Flexibilities

To isolate the value of timing flexibility from sourcing flexibility, we assess the former by considering three policies that can only sign contracts of a single tenor: FBH_m , FRH_m , and IRH_m . Here FBH_m serves as a benchmark with limited timing flexibility as it can only sign a new contract once the incumbent expires, while FRH_m and IRH_m have full timing flexibility and can sign new contracts each year but differ in how they optimize this flexibility.

Figure 6 reports the procurement cost of the aforementioned methods for different contract tenors in separate subfigures. Each subfigure displays results across target levels from 50% to 100% (which are more likely in practice), in increments of 10%. FRH_m outperforms FBH_m for all contract tenors, with this difference increasing for larger ones. For a five-year VPPA, the value of timing flexibility (i.e., FRH_5 improvement over FBH_5) ranges between 1.5% and 2.4%, and the value of optimizing this flexibility (i.e., IRH_5 improvement over FRH_5) varies from 0.2% to 0.3%. The analogous ranges for a 25-year VPPA are 3.2%–4.4% and 2.6%–3.4%, respectively. The reason

for the value of timing flexibility increasing with the tenor is because FBH_m can sign fewer new deals for longer-term agreements. For example, FBH_{20} can sign a maximum of two contracts over a 40-year planning horizon, whereas FRH_{20} has 40 opportunities to sign new VPPAs. Figure 6 also shows that long-term VPPAs lead to lower procurement costs under all policies, which can be attributed to the lower premiums associated with these contracts and suggests that the timing flexibility reduction from using long-term contracts in FBH_m is outweighed by the lower price contracts.

On average, across instances and contracts, the optimality gaps of FBH_m and FRH_m are 16.9% and 13.0%, respectively. Therefore, the average value of timing flexibility is roughly 4% when using forecast-based methods. The optimality of IRH_m is on average 10.9%. These optimality gaps indicate that IRH_m is better at optimizing timing flexibility than FRH_m , and such optimization decreases procurement costs by 2.1% (i.e., 13.0%–10.9%) on average. To understand how the use of timing flexibility affects VPPA procurement capacities, we report the following statistics in Table 4 for the instance with a 90% target ($\alpha = 0.9$): (i) the VPPA purchase frequency, defined as the average number of years between signing two consecutive

Table 4. Average Frequency and Size of the VPPAs Signed by Single-Contract Policies ($\alpha = 0.9$)

	$m \in \mathcal{M}$					$m \in \mathcal{M}$				
	Average frequency (years)					Average contract size (MW)				
	5	10	15	20	25	5	10	15	20	25
FBH_m	8.6	13.0	17.1	20.7	25.9	186.3	186.5	186.8	187.3	187.5
FRH_m	8.4	11.1	11.9	14.0	15.7	167.5	151.9	139.4	138.4	141.8
IRH_m	8.6	11.2	12.3	13.3	14.5	159.4	146.9	142.2	141.3	148.7

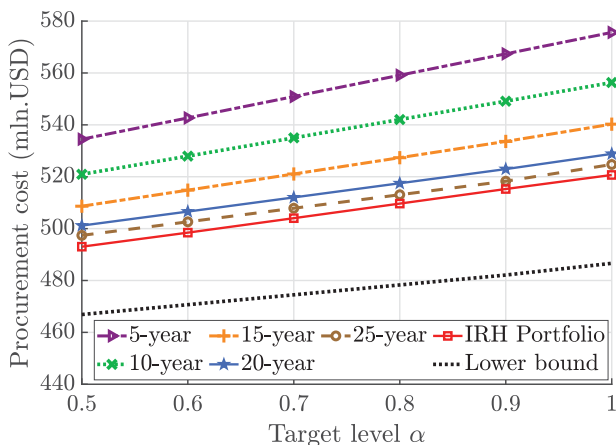
Table 5. Average VPPA Strike Price (USD/MWh) Signed by Single-Contract Policies ($\alpha = 0.9$)

	VPPA tenor				
	5	10	15	20	25
FBH _m	43.7	43.5	43.2	41.9	41.7
FRH _m	38.2	38.5	38.5	38.2	37.6
IRH _m	36.6	35.9	35.5	34.7	34.1

contracts, and (ii) the average contracted capacity in MW of the signed VPPAs. Although these statistics are quite similar for FRH_m and IRH_m, both these policies use their additional timing flexibility to sign VPPAs more often and with less capacity compared with FBH_m. For example, the average interval between contracts signed by FBH₂₅ is roughly 26 years, whereas it is less than 16 years for FRH_m and IRH_m. Moreover, the average size of VPPAs signed by FBH₂₅ is 187.5 MW, which is 26%–32% higher than that of FRH₂₅ and IRH₂₅ (141.8–148.7 MW). By signing less capacity at a given point in time, the policies with timing flexibility “reserve” capacity for even better deals arising in the future, whereas a policy with limited timing flexibility does not have the opportunity to do this and thus locks in deals with larger capacity. This is also confirmed by additional statistics in Table 5 reporting the average strike price (weighted on contract size) of VPPAs signed by single-contract policies on the same instance with $\alpha = 0.9$. On average, FRH_m and IRH_m sign cheaper contracts than FBH_m, saving 3.7–5.5 and 7.0–7.7 USD/MWh, respectively.

To isolate the value of sourcing flexibility, we compare IRH_m and IRH for target levels from 50% to 100% in Figure 7, also showing the information relaxation lower bound. The value of sourcing flexibility (i.e., improvement of IRH over IRH_m) depends on the length of the VPPAs and ranges between 0.7% and

Figure 7. (Color online) Procurement Costs of the IRH_m and IRH Policies



9.5%, which corresponds to a 3–55 million USD decrease in procurement cost. Specifically, the value of sourcing flexibility is between 5.5% and 9.5% when computed relative to IRH₅ and IRH₁₀ and between 0.7% and 1.6% relative to IRH₂₀ and IRH₂₅. For $\alpha = 0.9$, Table 6 reports procurement statistics for the IRH policy disaggregated by contract type. Comparing this table with Tables 4 and 5 shows that IRH signs VPPAs of smaller size and at a cheaper strike price on average than IRH_m for every tenor. In addition, new agreements are signed every 7.9 years on average by IRH, which is more frequent than the IRH_m policies (Table 4). We also verified that the IRH VPPA portfolio contains contracts with 2.1 different tenors on average. Thus, IRH’s allocation of capacity among VPPAs is more price efficient as it can exploit the availability of heterogenous VPPA types (i.e., sourcing flexibility), in addition to its ability to time the signing of these contracts (i.e., timing flexibility).

6.5. Sensitivity to Strike Price Specification

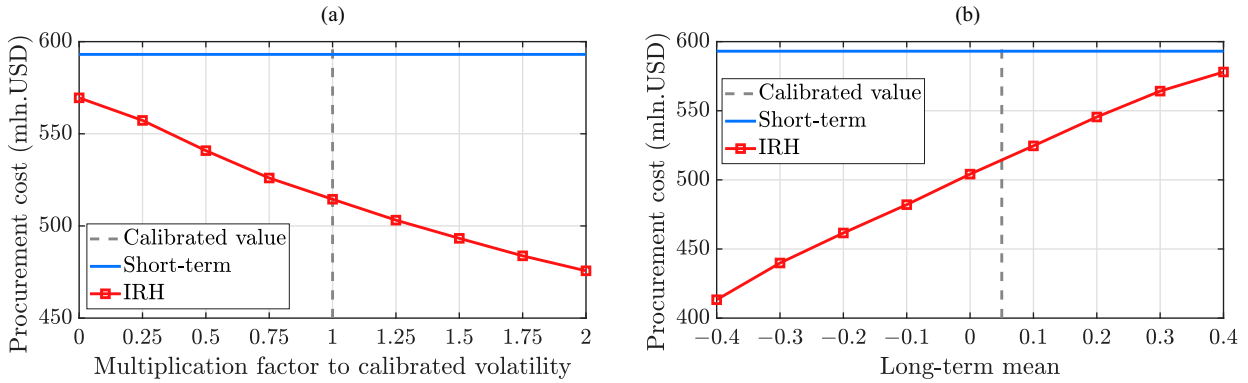
Our results regarding the value of dynamic VPPA portfolios indeed depend on the strike price specification discussed in Section 5. Because the renewable energy market is evolving rapidly, it is possible for the mean and the volatility of VPPA strike prices to change. To capture the former and latter effects, we vary respectively the long-term mean and the volatility of the stochastic process U_t , which captures deviations from the baseline in (12). Specifically, we increase and decrease the calibrated volatility by 25%, 50%, 75%, and 100% and vary the long-term mean from -0.4 to 0.4 , which has a calibrated value of 0.0468.

Figure 8(a) plots the IRH procurement costs as a function of the volatility and uses the short-term policy procurement cost, which only depends on power and REC prices, as a reference. This figure reveals that procurement costs under IRH reduce with higher strike price volatility: Costs decrease substantially from 514 million USD at the calibrated volatility to 476 million USD when the volatility is double this value. The cost improvement from using the IRH policy over the short-term policy reaches 25% for the latter volatility value. Intuitively, this happens because higher strike price volatility implies that it is likely for our model to have strike prices that are cheaper and

Table 6. Strike Price (USD/MWh) and Size (MW) of the VPPAs Signed by IRH ($\alpha = 0.9$)

	VPPA tenor				
	5	10	15	20	25
Average strike price	36.2	35.5	34.0	33.4	33.1
Average size	46.5	87.1	99.9	119.0	131.0

Figure 8. (Color online) Sensitivity of Procurement Cost to Changes in the U_i Parameters ($\alpha = 0.9$)



Notes. (a) Changes to the volatility. (b) Changes to the long-term mean.

more expensive than when this volatility is low. IRH can use its timing and sourcing flexibilities to asymmetrically benefit from this volatility by signing cheaper contracts and avoid more expensive contracts. As expected, Figure 8(b) shows that decreasing the long-term mean of U_i , which translates to lowering the average strike price, reduces procurement costs, whereas the opposite happens when this mean is increased.

Next, we study parameters of the strike price model that depend on the contract tenor because the public data set we used for calibration in Section 5 did not contain tenor information. We first study the effect of varying the premiums ξ_m for $m \in \mathcal{M}$ in (11). These premiums were chosen in Table 2 based on the generator’s preference for longer-term contracts. The competing force in negotiations is that a long-term commitment involves the buyer agreeing to strike price terms over a more uncertain future, where power prices may drop. Buyers thus prefer shorter-term VPPAs. For instance, Walmart typically signs VPPAs of 10 or 15 years (Ozment 2014). We define seven instances labeled A through G where the contract premiums are varied across tenors in \mathcal{M} as shown in Table 7 (instance F corresponds to the setting in Table 2 analyzed in Sections 6.2–6.4). We increase and decrease the premiums for shorter- and longer-tenor contracts in a symmetric

fashion around a 15-year premium ξ_{15} of zero. The IRH and IRH_m procurement costs on these instances and their information relaxation dual bound are shown in Figure 9 for a 90% target.

IRH has the lowest procurement cost on all instances. Short-duration (i.e., IRH_5 and IRH_{10}) and long-duration (i.e., IRH_{20} and IRH_{25}) single-contract policies increase the IRH procurement costs by an average of 6.9% when premiums decrease (i.e., instances E–G) and increase (i.e., instances A–C) with tenor length, respectively. Therefore, the sourcing flexibility in the VPPA portfolios managed by IRH allows it to maintain low costs across different variations in the contract premiums, whereas the performance of single-contract policies worsens. Further investigation strengthens this intuition. Specifically, IRH_{20} and IRH_{25} have optimality gaps of 8.5% and 7.6%, respectively, on instance F considered in Section 6.4, but the analogous gaps on instance A are 14.5% and 17.5%, that is, much larger. In contrast, the optimality gaps of IRH vary between 6.2% and 8.4% across the seven instances, with the

Table 7. Definition of Strike Price Premiums ξ_m for Instances A–G

Instance	ξ_5	ξ_{10}	ξ_{15}	ξ_{20}	ξ_{25}
A	-10	-5	0	+5	+10
B	-5	-2.5	0	+2.5	+5
C	-2	-1	0	+1	+2
D	0	0	0	0	0
E	+2	+1	0	-1	-2
F	+5	+2.5	0	-2.5	-5
G	+10	+5	0	-5	-10

Figure 9. (Color online) Procurement Costs of the IRH_m and IRH Policies at Varying Strike Price Premiums ($\alpha = 0.9$)

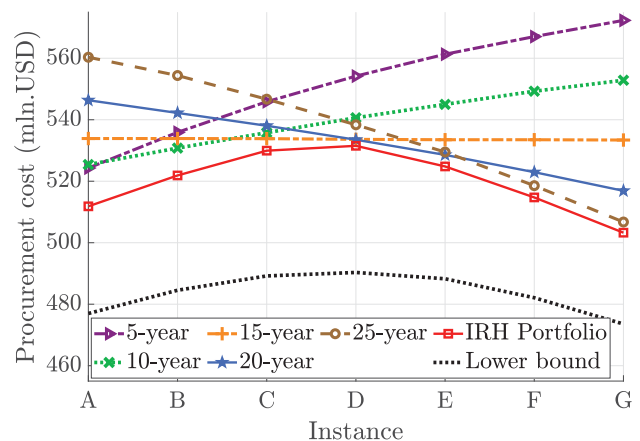
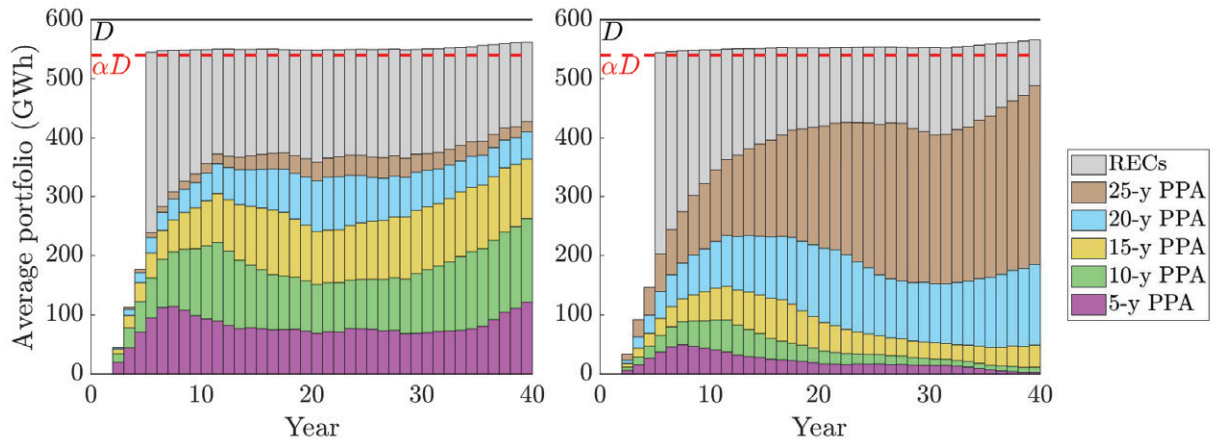


Figure 10. (Color online) Average VPPA Portfolios for Instances B (Left) and F (Right)



amount of capacity allocated to different VPPA contracts adjusting to the variation in premiums as illustrated in Figure 10 for instances B and F.

We also performed experiments that vary the availability p_m of each VPPA. We find that the procurement costs of each single-contract policy worsen as its availability decreases because there is exposure to the short-term power and REC markets when prices in these markets are high. The IRH procurement cost as expected is more insulated from the change in the availability of a particular VPPA because of its sourcing flexibility to sign other VPPAs. Online Appendix F contains more details on the sensitivity analysis regarding availability factors.

Overall, the sensitivity results show that key insights from Sections 6.2–6.4 continue to hold when parameters of our strike price model are changed: (i) significant procurement cost reduction can be achieved by leveraging the sourcing flexibility of VPPA portfolios to meet a target rather than using contracts of a single tenor, and (ii) the procurement portfolios constructed by IRH are near optimal. Finally, the sensitivity of the procurement portfolios highlights the value of tenor-dependent strike price data (as opposed to leveled strike price data) to corporate buyers as it would allow them to estimate tenor-dependent premiums ξ_m for procurement planning purposes.

7. Conclusion

Motivated by the global trend in corporate energy procurement, we investigated the problem of meeting a renewable power target using VPPAs. To inform the construction of a dynamic VPPA portfolio to meet the target, we formulated an MDP that accounts of price and supply uncertainties. Because the MDP policy is

challenging to obtain, we proposed forecast-based heuristics that extend current practice-based policies and developed a novel information-relaxation based technique. The latter technique is unique in its ability to account for the impact of uncertainty when computing procurement decisions. It also provides a dual bound for benchmarking policies. We also proposed a new latent variable model that governs the evolution of strike prices. Through an extensive computational study on realistic procurement instances, we find that using VPPAs to meet a target reduces costs substantially compared with procurement in the short-term markets without these agreements. Dynamically updating a VPPA portfolio with multiple tenors results in the least procurement cost, that is, timing and sourcing flexibilities have value. Fully exploiting these flexibilities requires using our information-relaxation based reoptimization heuristic and leads to near optimal VPPA portfolios, whereas forecast-based reoptimization can result in significantly higher procurement costs despite also constructing VPPA portfolios.

The research in this paper takes a meaningful step toward formally studying corporate power procurement to meet a renewable procurement target, and in particular, suggests significant value in using dynamic VPPA procurement portfolios and provides related decision analytics. Given the rapidly evolving nature of corporate renewable power procurement and energy markets, several forward-looking extensions are possible, of which, we state a few. Assessing the value of virtual storage in the context of VPPAs would be quite valuable, as it would partly manage the financial risk associated with uncertain renewable power supply in these agreements and even provide an additional source of merchant trading revenue from price arbitrage in the power market. Another

possibility is to consider the forward purchase of RECs or swap more expensive RECs for cheaper ones of a different type from another region. Finally, recent corporate renewable procurement deals involve third-party insurance companies to alleviate supply, price, and other operational risks. Understanding the value of this insurance and its relationship to the risk aversion of corporations would be valuable.

Acknowledgments

The authors thank the review team for valuable feedback that helped us significantly improve the original version of this manuscript. The first author thanks the University of Illinois Chicago, ETH Zurich, and Denmark Technical University for providing resources to conduct this research as a visiting scholar, PhD student, and postdoctoral scholar, respectively.

Endnotes

¹ See <https://www.there100.org/re100-members>.

² See <https://www.climateutraldatacentre.net/>.

³ A short-term electricity purchase may be from the wholesale market at the day-ahead price or through a utility or retail energy supplier at a variable monthly or quarterly retail rate.

⁴ WBCSD, World Business Council for Sustainable Development.

⁵ Unbundled RECs represent the short-term option to buy RECs at an uncertain price from a secondary market.

⁶ The definition of the set of contracts in our formulation can be easily extended to differentiate contracts based on features other than length.

⁷ The amount of energy in MWh delivered in a specific interval is given by the product of the contracted capacity in MW and the stochastic renewable capacity availability during this interval measured in hours.

⁸ Weights β could be more general and depend on other time indexes (i.e., t and l) and on the VPPA tenor m .

⁹ This is instead not true when considering more general dual penalties that are functions of x_i and z_i (i.e., bilinear).

¹⁰ If tenor-dependent price data are available, one could either calibrate ξ_m or, alternatively, adapt our model to use $BP_{i,m}$ as the baseline and eliminate the need for specifying a tenor-dependent premium ξ_m .

¹¹ All methods were implemented in C++ and use Gurobi 9.1 as math programming solver. The tests were run on a laptop equipped with an Intel Core i7-8650U processor and consumed at most 200 MB RAM.

¹² The solution of IRH math programs in the inner samples can be parallelized to substantially reduce this overhead, which we did not do.

References

- Aflaki S, Netessine S (2017) Strategic investment in renewable energy sources: The effect of supply intermittency. *Manufacturing Service Oper. Management* 19(3):489–507.
- Al-Gwaiz M, Chao X, Wu OQ (2016) Understanding how generation flexibility and renewable energy affect power market competition. *Manufacturing Service Oper. Management* 19(1):114–131.
- Alizamir S, de Véricourt F, Sun P (2016) Efficient feed-in-tariff policies for renewable energy technologies. *Oper. Res.* 64(1):52–66.
- Allon G, Van Mieghem J (2010) Global dual sourcing: Tailored base-surge allocation to near- and offshore production. *Management Sci.* 56(1):110–124.
- Andersen L, Broadie M (2004) Primal-dual simulation algorithm for pricing multidimensional American options. *Management Sci.* 50(9):1222–1234.
- Baker McKenzie (2015) The rise of corporate PPAs: A new driver for renewables. Technical report, Baker McKenzie, Chicago, USA.
- Baker McKenzie (2018) The rise of corporate PPAs 2.0. Technical report, Baker McKenzie, Chicago, USA.
- BerkeleyLab (2020a) Utility-scale solar. Accessed January 10, 2022, <https://rmp.lbl.gov/utility-scale-solar>.
- BerkeleyLab (2020b) Utility-scale solar. Accessed January 10, 2022, <https://emp.lbl.gov/wind-technologies-market-report/>.
- Bertsekas DP (2005) Dynamic programming and suboptimal control: A survey from ADP to MPC. *Eur. J. Control* 11(4-5):310–334.
- Bertsekas DP (2011) *Dynamic Programming and Optimal Control*, 3rd ed. (Athena Scientific, Belmont, MA).
- Bloomberg (2018) United States rates and bonds. Accessed January 10, 2022, <http://pages.news/bloomberg3>.
- BNEF (2018a) Corporations already purchased record clean energy volumes in 2018, and it's not an anomaly. Accessed January 10, 2022, <http://energy.financial/clean4>.
- BNEF (2018b) Executing your corporate energy strategy. Technical report, BNEF, New York, USA.
- BNEF (2020) *Corporate clean energy buying leapt 44% in 2019, sets new record*. Technical report, BNEF, New York, USA.
- Boyabatli O, Kleindorfer PR, Koontz SR (2011) Integrating long-term and short-term contracting in beef supply chains. *Management Sci.* 57(10):1771–1787.
- Brown DB, Haugh MB (2017) Information relaxation bounds for infinite horizon Markov decision processes. *Oper. Res.* 65(5):1355–1379.
- Brown DB, Smith JE (2011) Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Sci.* 57(10):1752–1770.
- Brown DB, Smith JE (2014) Information relaxations, duality, and convex stochastic dynamic programs. *Oper. Res.* 62(6):1394–1415.
- Brown DB, Smith JE, Sun P (2010) Information relaxations and duality in stochastic dynamic programs. *Oper. Res.* 58(4):785–801.
- Bruck M, Sandborn P, Goudarzi N (2018) A leveled cost of energy (LCOE) model for wind farms that include power purchase agreements (PPAs). *Renewable Energy* 122:131–139.
- CDP, WWF, Calvert Investments, Ceres (2017) How the largest US companies are capturing business value while addressing climate change. Technical report, Carbon Disclosure Project (CDP), World Wide Fund for Nature (WWF), Coalition for Environmentally Responsible Economies (Ceres), Washington DC, USA.
- Chand S, Hsu VN, Sethi S (2002) Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing Service Oper. Management* 4(1):25–43.
- Chao X, Chen FY (2005) An optimal production and shutdown strategy when a supplier offers an incentive program. *Manufacturing Service Oper. Management* 7(2):130–143.
- Davies K, Orrick GMJ, Taylor L (2018) Proxy generation ppas. Technical report, Microsoft, Orrick, and REsurety, Boston, USA.
- De Farias DP, Van Roy B (2003) The linear programming approach to approximate dynamic programming. *Oper. Res.* 51(6):850–865.
- Denholm P, Sioshansi R (2009) The value of compressed air energy storage with wind in transmission-constrained electric power systems. *Energy Policy* 37(8):3149–3158.
- Desai VV, Farias VF, Moallemi CC (2012) Pathwise optimization for optimal stopping problems. *Management Sci.* 58(12):2292–2308.

- Drake DF, Kleindorfer PR, Van Wassenhove LN (2016) Technology choice and capacity portfolios under emissions regulation. *Production Oper. Management* 25(6):1006–1025.
- EIA (2019) Capacity factors for utility scale generators primarily using non-fossil fuels. Accessed January 10, 2022, <http://output.energy/eia3>.
- Gallego G, Van Ryzin G (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* 45(1):24–41.
- Google (2017) Environmental report: 2017 progress update. Accessed January 10, 2022, <https://sustainability.google/reports/environmental-report-2017/#data-centers>.
- Haugh MB, Kogan L (2004) Pricing American options: A duality approach. *Oper. Res.* 52(2):258–270.
- Haugh MB, Lacedelli OR (2018) Information relaxation bounds for partially observed Markov decision processes. Working paper, Imperial College London, London.
- Hu S, Souza GC, Ferguson ME, Wang W (2015) Capacity investment in renewable energy technology with supply intermittency: Data granularity matters! *Manufacturing Service Oper. Management* 17(4):480–494.
- İşlegen Ö, Reichelstein S (2011) Carbon capture by fossil fuel power plants: An economic analysis. *Management Sci.* 57(1):21–39.
- Jiang DR, Powell WB (2015) Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS J. Comput.* 27(3):525–543.
- Kamiya Y, Kvarnström A (2019) International Energy Agency: Data centres and energy: From global headlines to local headaches? Technical report, International Energy Agency, Paris, France.
- Kim JH, Powell WB (2011) Optimal energy commitments with storage and intermittent supply. *Oper. Res.* 59(6):1347–1360.
- Kleindorfer PR, Wu DJ (2003) Integrating long-and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Sci.* 49(11):1597–1615.
- Lai G, Margot F, Secomandi N (2010) An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Oper. Res.* 58(3):564–582.
- Li CL, Kouvelis P (1999) Flexible and risk-sharing supply contracts under price uncertainty. *Management Sci.* 45(10):1378–1398.
- Lin Q, Nadarajah S, Sohelli N (2019) Revisiting approximate linear programming: Constraint-violation learning with applications to inventory control and energy storage. *Management Sci.* 66(4):1544–1562.
- Longstaff FA, Schwartz ES (2001) Valuing American options by simulation: A least-squares approach. *Rev. Financial Stud.* 14(1):113–147.
- Lopuhaa HP, Rousseeuw PJ (1991) Breakdown points of affine equivariant estimators of multivariate location and covariance matrices. *Ann. Statist.* 19(1):229–248.
- Loukatou A, Howell S, Johnson P, Duck P (2018) Stochastic wind speed modelling for estimation of expected wind power output. *Appl. Energy* 228:1328–1340.
- Lucia JJ, Schwartz ES (2002) Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Rev. Derivative Res.* 5(1):5–50.
- MarexSpectron (2020) Renewable energy credits. Technical report.
- Martínez-de Albéniz V, Simchi-Levi D (2005) A portfolio approach to procurement contracts. *Production Oper. Management* 14(1):90–114.
- Meibom P, Larsen HV, Barth R, Brand H, Tuohy A, Ela E (2011) Advanced unit commitment strategies in the United States Eastern Interconnection. *Contract* 303:275–300.
- Milligan M, Ela E, Lew D, Corbus D, Wan Y, Hodge B, Kirby B (2012) Operational analysis and methods for wind integration studies. *IEEE Trans. Sustainable Energy* 3(4):612–619.
- Min S, Maglaras C, Moallemi CC (2021) Thompson sampling with information relaxation penalties. Working paper, Columbia University, New York.
- Nadarajah S, Secomandi N (2018) Merchant energy trading in a network. *Oper. Res.* 66(5):1304–1320.
- Nadarajah S, Margot F, Secomandi N (2017) Comparison of least squares Monte Carlo methods with applications to energy real options. *Eur. J. Oper. Res.* 256(1):196–204.
- NREL (2017) System advisor model (SAM). Accessed January 10, 2022, <https://sam.nrel.gov/>.
- Ozment D (2014) Walmart’s approach to renewable energy. Technical report, Walmart.
- Pandžić H, Wang Y, Qiu T, Dvorkin Y, Kirschen DS (2015) Near-optimal method for siting and sizing of distributed storage in a transmission network. *IEEE Trans. Power Systems* 30(5):2288–2300.
- Pereira MVF, Pinto LMVG (1991) Multi-stage stochastic optimization applied to energy planning. *Math. Programming* 52(1-3):359–375.
- PJM (2020) Day-ahead hourly Imps. Technical report.
- Powell WB (2011) *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd ed. (John Wiley & Sons, Hoboken, NJ).
- Ritzenhofen I, Birge JR, Spinler S (2016) The structural impact of renewable portfolio standards and feed-in tariffs on electricity markets. *Eur. J. Oper. Res.* 255(1):224–242.
- RMI (2016) Amazon and dominion virginia power reach breakthrough renewable energy agreement. Technical report.
- Salas DF, Powell WB (2017) Benchmarking a scalable approximate dynamic programming algorithm for stochastic control of grid-level energy storage. *INFORMS J. Comput.* 30(1):106–123.
- Secomandi N (2015) Merchant commodity storage practice revisited. *Oper. Res.* 63(5):1131–1143.
- Secomandi N, Kekre S (2014) Optimal energy procurement in spot and forward markets. *Manufacturing Service Oper. Management* 16(2):270–282.
- Shankleman J (2020) Google targets carbon-free power for data centers by 2030. Accessed January 10, 2022, <https://www.bloomberg.com/news/articles/2020-09-14/google-ramps-up-clean-energy-ambition-as-californian-fires-rage>.
- Shapiro A (2011) Analysis of stochastic dual dynamic programming method. *Eur. J. Oper. Res.* 209(1):63–72.
- Singh SP, Scheller-Wolf AA (2017) That’s not fair: Tariff structures for electricity markets with rooftop solar. Working paper, Carnegie Mellon University, Pittsburgh, PA.
- Sunar N, Birge JR (2019) Strategic commitment to a production schedule with uncertain supply and demand: Renewable energy in day-ahead electricity markets. *Management Sci.* 65(2):714–734.
- Tomlin B, Wang Y (2005) On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. *Manufacturing Service Oper. Management* 7(1):37–57.
- Tsitsiklis JN, Van Roy B (2001) Regression methods for pricing complex American-style options. *IEEE Trans. Neural Networks* 12(4):694–703.
- Veeraraghavan S, Scheller-Wolf AA (2008) Now or later: A simple policy for effective dual sourcing in capacitated systems. *Oper. Res.* 56(4):850–864.
- WBCSD (2018) *Innovation in power purchase agreement structures. Technical report*, WBCSD.
- Webb E, Wu OQ, Cattani K (2017) Mind the gap: Coordinating energy efficiency and demand response. Working paper, Indiana University, Indiana, PA.
- Weber C, Meibom P, Barth R, Brand H (2009) WILMAR: A stochastic programming tool to analyze the large-scale integration of wind energy. *Optimization in the Energy Industry* (Springer, Berlin), 437–458.

- Weron R (2014) Electricity price forecasting: A review of the state-of-the-art with a look into the future. *Internat. J. Forecasting* 30(4):1030–1081.
- Wiser R, Bolinger M (2017) 2016 wind technologies market report. Accessed January 10, 2022, <https://www.energy.gov/eere/wind/downloads/2016-wind-technologies-market-report>.
- Wu OQ, Babich V (2012) Unit-contingent power purchase agreement and asymmetric information about plant outage. *Manufacturing Service Oper. Management* 14(2):245–261.
- Wu OQ, Kapuscinski R (2013) Curtailing intermittent generation in electrical systems. *Manufacturing Service Oper. Management* 15(4):578–595.
- Wu OQ, Wang DD, Qin Z (2012) Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing Service Oper. Management* 14(3):455–471.
- Ye F, Zhu H, Zhou E (2018) Weakly coupled dynamic program: Information and Lagrangian relaxations. *IEEE Trans. Automated Control* 63(3):698–713.
- Zeng Y, Klabjan D, Arinez J (2015) Distributed solar renewable generation: Option contracts with renewable energy credit uncertainty. *Energy Econom.* 48:295–305.
- Zhou Y, Scheller-Wolf A, Secomandi N, Smith S (2019) Managing wind-based electricity generation in the presence of storage and transmission capacity. *Production Oper. Management* 28(4):970–989.