Meeting corporate renewable power targets

Working Paper

Author(s):
Trivella, Alessio ; Mohseni Taheri, Danial; Nadarajah, Selvaprabu

Publication date:
2020-04

Permanent link:
https://doi.org/https://doi.org/10.3929/ethz-b-000448178

Originally published in:
TS Working Papers

Meeting Corporate Renewable Power Targets

Alessio Trivella
Institute for Transport Planning and Systems, ETH Zürich, 8092 Zürich, Switzerland
alessio.trivella@ivt.baug.ethz.ch

Danial Mohseni-Taheri, Selvaprabu Nadarajah[1]
College of Business Administration, University of Illinois at Chicago, Chicago, IL 60607
{smohse3@uic.edu, selvan@uic.edu}

First version: December 2018; This version: April 2020

Prominent companies have committed to procuring a percentage of their power demand from renewable sources by a future date. Long-term financial contracts with renewable generators, known as corporate power purchase agreements (CPPAs), are popular to meet such a renewable power purchase target (RPPT). By analyzing a simplified three-stage model, we show that the generation capacity contracted via a CPPA is more nuanced to structure optimally compared to traditional long-term power contracts due to the interplay between price and supply uncertainties as well as the RPPT. We subsequently propose a Markov decision process (MDP) to formalize rolling-power-purchase policies used in practice, that is, the construction of dynamic CPPA portfolios to meet an RPPT. The optimal MDP policy is intractable to compute but possesses the following key properties: (i) its decisions account for stochastic prices and supply, (ii) it captures the timing flexibility to enter CPPAs, and (iii) it can sign CPPAs with different tenures. We develop forecast-based reoptimization heuristics and a novel information-relaxation based reoptimization approach that sacrifice and approximate, respectively, the first property of the MDP policy and capture the remaining properties. We perform an extensive computational study on realistic procurement instances to uncover managerial insights related to procurement costs, the control of risks arising from supply uncertainty, the relevance of CPPAs as markets evolve, and the near-optimality of rolling power purchases from our information-relaxation based procurement heuristic.

1. Introduction

Corporations are playing an increasing leadership role in promoting sustainability and social responsibility around the globe. Over half of the Fortune 500 companies have publicly announced commitments to meet sustainability and climate goals, which includes greenhouse gas emissions reduction, energy efficiency improvements, and renewable power procurement (CDP et al. 2017). We focus on companies that have committed to renewable power purchase targets (RPPT), that

[1] corresponding author

is, they procure 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, Procter & Gamble and Intel have committed to RPPTs of 30% and 75%, respectively, by 2020 and 2025. In 2019, nearly 400 companies had committed to an RPPT and 221 global companies have pledged an RPPT of 100% as part of the RE100 initiative (BNEF 2020). To meet these targets, power 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 power.

Companies have increasingly resorted to corporate power purchase agreements (CPPAs) as vehicles to procure renewable power directly from the generator instead of going through a utility. A CPPA is a long-term bilateral contract between the company and a renewable generator to receive a fixed quantity of power and RECs at a predetermined *strike price* for each year of the contract's tenor. The purchase of power using CPPAs has increased 40% in 2019 from its previous record of 7.2 gigawatts in 2018 (Baker McKenzie 2018, BNEF 2018, 2020). In this paper, we consider synthetic CPPAs (henceforth, CPPAs) which have driven the recent wave of corporate renewable procurement and account for 80% of the clean energy procurement in 2019. They are long-term financial agreements on generation capacity where the producer sells the power generated to the grid, the firm buys power at a short-term (i.e., uncertain) price[1] from the grid or a utility, and payments for differences between the fixed strike price and a variable reference price are made. Typically, the CPPA reference price and the firm's short-term power purchase price are equal (or at least close to each other) to ensure a price hedge (see, e.g., RMI 2016, WBCSD 2018). Despite this price hedge, the amount of power that is generated from the contracted capacity is uncertain due to operations (e.g., maintenance) beyond the control of the buyer and the intermittency of renewable resources such as wind and solar. In other words, CPPAs ensure that the buyer pays a fixed price per MWh on an uncertain quantity of power. This feature is a blessing, as it avoids CPPAs being classified as a derivative and adhering to financial regulation (see, e.g., Davies et al. 2018, page 5), as well as a curse because it does not eliminate the companies exposure to power price and supply uncertainties.

Since CPPAs are quite different from traditional long-term contracts that ensure certainty in price and supply, there has been significant recent experimentation in corporate renewable procurement policies. Traditionally, long-term power contract lengths were twenty to twenty five years. Recent CPPAs include contracts with tenors ranging from five to fifteen years and companies have begun actively managing CPPA portfolios over time. Despite these trends, a large proportion of

[1]A short-term power purchase may be from the wholesale market at the day-ahead price or through a utility at a variable monthly retail rate.

companies have not yet committed to any renewable energy target due to a lack of strategic knowledge on corporate renewable procurement (PWC 2016). Moreover, the procurement problem faced by a firm that plans to setup and meet an RPPT has not been formally studied in the extant academic literature to the best of our knowledge. Our goal is to take a meaningful step towards reducing this knowledge gap by studying two dominant power procurement options to satisfy an RPPT (BNEF 2018): (i) enter into CPPAs, and (ii) buy short-term power, as needed, and supplement it with unbundled REC purchases[2]. To this end, we (i) analyze a simplified procurement model to highlight the structure of CPPAs and the impact of price and supply uncertainties on CPPA capacities and costs to meet an RPPT, and (ii) formulate a realistic tactical planning model and use it to investigate how practice-based policies as well as new dynamic procurement policies can be beneficial to meet an RPPT.

Our simplified procurement model uses as its long-term option a single CPPA that can be entered today with a contract length extending to the RPPT date. We analyze the behavior of the optimal expected procurement cost and CPPA capacity with respect to market parameters as well as the RPPT level. The procurement cost is greater in markets with high renewable energy penetration and high (short-term) power price volatility. Moreover, buyers procure less using CPPAs in such markets where practitioners would traditionally expect long-term agreements to be effective at reducing energy costs. Uncertainty in power supply drives this non-traditional behavior. In particular, the balancing of sustainability goals and energy costs is more nuanced: the capacity contracted via a CPPA (in the presence of power supply uncertainty) is less than and greater than, respectively, the analogous capacity in a deterministic supply setting with small and large RPPTs. Therefore, CPPAs become more effective at reducing procurement costs for companies that have aggressive RPPTs and are accessing power markets with limited renewable power penetration, which is the status quo in many US and global regions. In other words, a key insight is that combining aggressive RPPTs and CPPAs helps address climate change goals and energy costs – thus tying the knot between social responsibility and financial performance.

The intricacies of structuring procurement via a single CPPA underscore the challenge faced by companies that dynamically choose among multiple procurement options. We propose a realistic tactical procurement model that provides a unifying framework to view and benchmark multi-stage and multi-contract policies. Our model is a Markov decision process (MDP) that minimizes the expected procurement cost to meet and sustain an RPPT. The planning horizon in this MDP is divided into a reach period where the target does not have to be fulfilled (but contracts can be

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

signed) and a sustain period where the target must be satisfied. At each stage, the company decides whether to enter into new CPPAs. The set of available CPPAs depends on the contracts offered by generators, which is unpredictable over time. Moreover, when entering a CPPA, 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 *stochastic prices and supply*, (ii) capturing the *timing flexibility* to sign CPPAs, and (iii) potentially signing a *variety of CPPAs* with different tenures. Nevertheless, computing this policy is intractable owing to our MDP having a non-convex action set and a high-dimensional state space. We thus develop heuristics motivated by practice as well as a novel approach to approximate our MDP.

Practice-based procurement policies use forecasts and take a rolling power purchase approach. In its traditional form, such a policy signs CPPAs of a single fixed tenure (e.g., 20 years), maintains only one CPPA at any time, and determines the capacity of the new CPPA for the next block of the planning horizon (e.g., years 21-40) once the incumbent CPPA expires. The CPPA 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 CPPA tenure. We refer to this strategy with low timing flexibility and no CPPA variety as the forecast-based block heuristic (FBH). We also consider two extensions that both increase the timing flexibility to allow the signing of new CPPAs each year and use forecasts that span the full planning horizon. These extensions, dubbed forecast-based reoptimization heuristic (FRH) and portfolio FRH (PFRH), differ in the variety of candidate CPPAs they can access, with the former and latter using single- and multiple-CPPA types.

Since the aforementioned policies use forecasts, in contrast to an optimal policy, they do not directly capture the impact of the future evolution of uncertainty on the current procurement decision, which is conceptually undesirable. Moreover, understanding if this issue results in poor empirical policy performance requires a lower bound on the optimal policy value. We develop a novel multiple-contract procurement heuristic based on information-relaxation (PIRH) that provides both a lower bound and decisions that account for future uncertainty. PIRH 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 non-anticipative decision from a distribution of anticipative actions 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 CPPA contract lengths ranging from 5 to 25 years and a planning horizon of 40 years. The strike prices of CPPAs are specified by a model that factors in the effects of the expected power price, the expected supply, and the return on investment required by the generator in a manner that is consistent with publicly available software from the National Renewable Energy Laboratory (NREL; NREL 2017). We calibrate stochastic processes for the evolution of power/REC prices as well as supply using market data and the practitioner literature. We uncover the following insights:

- The additional cost of certifying power as being renewable increases by roughly 6% as the correlation between the power price and supply changes from being perfectly negatively correlated to being perfectly positively correlated. This suggests that companies need to account for co-variance between power price and supply uncertainties to accurately estimate the cost of meeting an RPPT using CPPAs.

- The shape of power supplied by a renewable generator exposes the firm to two types of shape risks associated with (i) meeting demand and (ii) satisfying the RPPT. Mitigating the former risk is not in the firm's direct control but the latter risk can be controlled by satisfying the RPPT less frequently (e.g., yearly). Reducing this frequency, e.g., from monthly to yearly causes the expected procurement cost to decrease by 3–4%.

- Rolling power purchases with a single long-term CPPA leads to near-optimal procurement costs and optimality gaps of 3–5%, provided that new CPPAs can be signed each year as done in FRH. Using shorter-tenure CPPAs alone in both FBH and FRH leads to larger optimality gaps of 7–15%. This finding motivates the use of portfolios with multiple CPPA tenures. Such portfolios computed by forecast-based heuristics lead to larger optimality gaps than our policies employing a single-CPPA type, which suggests that these heuristics do not sufficiently optimize CPPA portfolios. In contrast, PIRH is effective at computing near-optimal dynamic CPPA portfolios, with an average optimality gap of 1.6%. The use of CPPA portfolios, in addition to short-term purchases, can thus significantly reduce procurement costs, especially for companies with aggressive RPPTs.

- CPPAs remain valuable as a firm's operating environment evolves. In particular, corporate procurement portfolios with multiple CPPA types lead to procurement costs that are stable when contract availability and the market dynamics of REC prices change. Moreover, the removal of production tax credits has a significant impact on the cost of CPPAs, which results in a procurement cost increase of 9%. Although the presence of production tax credits directly affects the use of CPPAs in the short to medium term, they remain relevant in the long term.

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, Boyabatlı et al. 2011, Secomandi and Kekre 2014) as well as procurement in supply chains via short-term and long-term contracts, including dual- and multi-sourcing 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 an RPPT, related insights, and the dynamic policies are new to this literature. Moreover, the long-term contracts that we consider, that is CPPAs, have unique structure. For instance, CPPAs deliver power at each period over the tenure 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 CPPAs (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 CPPAs nor do we use an equilibrium model for this purpose. Rather, we obtain CPPA strike prices in our MDP by employing a modified net present value calculation consistent with the procedure in the publicly available software SAM from NREL (2017).

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 (see, e.g., 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 power procurement problem faced by a corporation with an RPPT is new to the renewable energy literature. Moreover, our investigation of the factors affecting the effectiveness of CPPAs and comparison of both traditional and contemporary procurement policies from the unifying lens of our MDP model 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 development of PIRH 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 that PIRH can outperform PFRH.

In addition to the reoptimization literature, PIRH contributes 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 because estimating a value function approximation and computing decisions using the Bellman operator are both challenging due to the large controllable part of the state space in our MDP arising from tracking the generation capacity contracted via CPPAs. Our development of PIRH thus adds a direct way to obtain non-anticipative controls when using the information relaxation and duality approach.

More broadly, PIRH adds to approximate dynamic programming (ADP; Bertsekas 2005), an area of stochastic optimization dealing with the solution of high-dimensional MDPs. Several ADP 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 (see, e.g., Salas and Powell 2017, Nadarajah and Secomandi 2018) and approaches that handle non-convex action sets are even more scarce. PIRH has potential value for solving MDPs with these complicating features, which arise beyond our specific procurement application.

1.2 Paper structure

In §2, we analyze CPPAs in a simplified procurement setting. In §3, we formulate an MDP to reach and sustain an RPPT using rolling power purchases. We present procurement heuristics that approximate this model in §§4–5. We conduct an extensive numerical study using models calibrated to data and discuss our findings in §6. We conclude in §7. All proofs and additional material related to our models and numerical study can be found in an online supplement.

2. Corporate power purchase agreements

We consider a firm that faces deterministic demand over time and commits now (time 0) to satisfy an RPPT in the future (time T), in addition to meeting its power demand at intermediate points in time. For simplicity, we assume there is only one intermediate time[3]. Thus, we have a three period problem with stage 0 representing time 0 and two future periods 1 and 2 corresponding to times $T/2$ and T, respectively, with power demands D_1 and D_2. We represent the RPPT at stage 2 as a fraction $\alpha \in [0, 1]$ of total demand $D := D_1 + D_2$, that is, the target equals αD. To meet this target, the company can (i) enter into a CPPA at stage 0 to receive power and RECs from a renewable generator at periods 1 and 2; (ii) procure any unmet power demand in these periods using short-term power purchases; and (iii) purchase unbundled RECs to satisfy any shortfall in the RPPT at period 2.

The main decision in our model is the megawatts (MWs) of capacity that the buyer contracts via the CPPA, which we denote by z. The CPPA has an associated strike price of K USD/MWh. Due to uncertainty, the megawatt hours (MWhs) of power generated in a period i may be less than z multiplied by the maximum number of hours H of production in each period. We model stochastic power supply in period i by $\eta_i z$, where η_i is a random variable with support on the interval $[0, H]$. The total renewable power supplied by the CPPA equals $\eta z = \sum_{i\in\{1,2\}} \eta_i z$. The prevalent CPPA format is an as-generated financial contract that ensures the company pays a fixed strike price K (USD/MWh) for $\eta_i z$ MWh of power generated at each period $i \in \{1, 2\}$ and receives RECs associated with ηz MWh. The cash flows when using a CPPA are as follows. The company purchases its period i power demand of D_i MWh from the short-term market at price P_i (USD/MWh). In addition, there is a financial settlement between the generator and the buyer. Specifically, if the short-term power price P_i is greater than the strike price K, the generator pays $P_i - K$ per MWh for the $\eta_i z$ MWh of power; otherwise, the company pays the generator $K - P_i$ per MWh of power

[3] Our results extend to the case with a general number of such times but this does not lead to new insights.

generated. Although there is no physical power delivery, the generator does transfer RECs for the total power generation of ηz MWh. If this amount is smaller than the target αD, the shortfall is met using unbundled REC purchases at price R_i (USD/MWh).

To represent the aforementioned cash flows, we define $w := (P_i, R_i, \eta_i, i \in \{0, 1, 2\})$ as a sample path of known and uncertain quantities, and use $\Theta := (\sigma_{P_i}, \sigma_{R_i}, \sigma_{\eta_i}, \rho_{P_i,R_i}, \rho_{P_i,\eta_i}, \rho_{R_i,\eta_i}, i \in \{1, 2\})$ to denote a vector containing the volatilities and correlations of the random components of w, where $\sigma_{(\cdot)}$ and $\rho_{(\cdot,\cdot)}$ represent the volatility and the correlation, respectively. Given a sample path realization w, the cost of choosing a CPPA with capacity z is

$$\widetilde{C}(z, w; \alpha, K, \Theta) := \sum_{i \in \{1,2\}} [P_i D_i + (K - P_i)\eta_i z] + R_2(\alpha D - \eta z)_+. \tag{1}$$

In this definition, the two terms in the summation correspond to the short-term power payment and the financial transaction between the generator and the buyer, respectively, while the third term is the cost of purchasing the REC shortfall. The cash flow in (1) can be rewritten as follows:

$$\widetilde{C}(z, w; \alpha, K, \Theta) = \sum_{i \in \{1,2\}} (K\eta_i z + P_i(D_i - \eta_i z)) + R_2(\alpha D - \eta z)_+.$$

It becomes apparent from the first term that a quantity ηz MWh is procured at the pre-agreed strike price K but the buyer is exposed to uncertain prices because supply is uncertain in the remaining terms. Specifically, supply uncertainty associated with a CPPA results in the buyer (i) being exposed to volatile power prices to varying extents dictated by the mis-match between the period i CPPA power supply $\eta_i z$ and the demand D_i in this period and (ii) the fluctuating amounts of unbundled REC purchases at uncertain REC prices when the total CPPA power supply ηz falls short of the RPPT αD. These issues are non-traditional to practitioners engaging in power procurement because they typically use long-term contracts to eliminate uncertainty. It is thus apriori unclear if and when a CPPA is an effective procurement vehicle, which prompts the following analysis of optimal CPPA procurement capacities and expected costs.

The optimal procurement cost when using a CPPA is

$$C(\alpha, K, \Theta) := \min_{z \geqslant 0} \mathbb{E}[\widetilde{C}(z, w; \alpha, K, \Theta)]. \tag{2}$$

To facilitate analysis we define two intuitive quantities. The first is the power agreement premium

$$\mathrm{PAP}(K, \Theta) := K\mathbb{E}[\eta] - \sum_{i \in \{1,2\}} \mathbb{E}[\eta_i P_i].$$

This term when positive equals the extra cost of receiving power and RECs from a unit of CPPA capacity compared to buying non-renewable power directly from the short-term market. The second quantity is the power and REC savings from using a CPPA, which we denote

$$\mathrm{PRS}(K,\Theta) := \sum_{i\in\{1,2\}} \mathbb{E}[\eta_i P_i] + \mathbb{E}[R_2\eta] - K\mathbb{E}[\eta] = \mathbb{E}[R_2\eta] - \mathrm{PAP}(K,\Theta).$$

This term equals the expected savings from receiving power and RECs from one unit of CPPA capacity relative to purchasing the same amounts of power and RECs directly from their respective short-term markets.

Proposition 2.1 characterizes the optimal CPPA capacity $z(\alpha, K, \Theta)$ and its behavior with power price related market parameters. When analyzing the behavior of functions with respect to a market parameter, we omit showing their dependence on parameters that do not play a role in this behavior. For example, we write $z(\rho_{P_i,\eta_i})$ to represent the one-dimensional function obtained by changing only ρ_{P_i,η_i} in $z(\alpha, K, \Theta)$.

Proposition 2.1. *The optimal procurement quantity satisfies*

$$z(\alpha,K,\Theta) := \begin{cases} 0 & \textit{if} \quad \mathrm{PRS}(K,\Theta) \leqslant 0 \\ \alpha D/y(K,\Theta) & \textit{otherwise,} \end{cases}$$

where $y(K,\Theta) = \inf\{y \in [0,2H] \,|\, \mathrm{PAP}(K,\Theta) - \mathbb{E}[R_2\eta\mathbf{1}(\eta \leqslant y)] \leqslant 0\}$. *Moreover, for any period i, (i) $z(\rho_{P_i,\eta_i})$ is a weakly increasing function of the power price and supply correlation ρ_{P_i,η_i}; and (ii) $z(\sigma_{P_i}, \rho_{P_i,\eta_i})$ is a weakly decreasing (weakly increasing) function of the price volatility σ_{P_i} if ρ_{η_i,P_i} is negative (positive).*

The structure of the optimal CPPA capacity in Proposition 2.1 shows that no CPPA is signed when the power and REC savings per unit of capacity is non-positive. Otherwise, it is optimal to sign a CPPA for a capacity value that balances (i) the premium paid for purchasing renewable power and RECs from the CPPA (i.e., $\mathrm{PAP}(K,\Theta)$); and (ii) the short-term REC costs avoided when meeting the RPPT as result of the RECs obtained via the CPPA (i.e., $K\mathbb{E}[\eta]$). It is also apparent that the optimal CPPA capacity increases linearly in the RPPT parameter α.

The optimal CPPA capacity increases with ρ_{P_i,η_i} because low power prices (i.e., low power price risk) tend to occur when the generator's power supply is low (i.e, high supply risk) and the buyer needs to purchase power at the short-term price. On the other hand, high power prices (i.e., high power price risk) are common when the generator's supply tends to be high (i.e., low supply risk) and most of the power needed is purchased at the pre-agreed CPPA strike price K. In other words, power price risk and supply risk work in opposing directions and thus mitigate the net effect on procurement cash flows, which in turn, makes CPPA capacity become more valuable and more capacity is procured as ρ_{P_i,η_i} increases. Following similar reasoning, the impact of power price volatility is mixed. If ρ_{P_i,η_i} is positive, high power price risk and low supply risk are likely to co-occur and entering into CPPA capacity is more beneficial as already discussed. On the other hand, if ρ_{P_i,η_i} is negative, then increasing power price volatility reduces the procured capacity as power price movements leading to high prices (i.e., high power price risk) are more likely when supply is low (i.e., high supply risk) and the buyer needs to purchase power at the short-term price. In markets with high renewable power penetration, such as ERCOT, ρ_{P_i,η_i} is negative (Woo et al. 2011) and our results indicate that procurement via a CPPA should be reduced relative to markets where correlations are zero or positive. In other words, buyers need to consider the risk arising from correlation when entering CPPAs.

Proposition 2.2 establishes the behavior of the optimal procurement cost, again focusing on the market parameters related to the power price.

Proposition 2.2. *The optimal cost $C(\alpha, K, \Theta)$ is (i) linearly increasing in α; (ii) weakly decreasing in ρ_{P_i,η_i}; and (iii) weakly increasing (decreasing) in σ_{P_i} if ρ_{P_i,η_i} is negative (positive).*

The optimal cost varies linearly with the RPPT parameter α, which shows that average spending needs to proportionally increase if the company wants to be more aggressive with its target when using a synthetic contract. The remaining findings pertaining to ρ_{P_i,η_i} and σ_{P_i} are consistent with Proposition 2.1 and the ensuing discussion underscoring the importance of correlation between power prices and supply.

Correlation and volatility risks are not limited to power price and supply but extend to the REC price in period 2, R_2, and the total supply, ηz. When the correlations ρ_{P_i,η_i} for all $i \in \{1, 2\}$ and $\rho_{R_2,\eta}$ are negative, Proposition 2.3 shows that the behavior of the optimal CPPA capacity with respect to supply volatility is more involved than its change with respect to price volatility established in Proposition 2.1. As already discussed, there is support for negative correlation between power price and supply, which we assume in Proposition 2.3. However, it is hard to rationalize a strong positive correlation between the period 2 REC price R_2 and total supply η. This correlation is likely zero or potentially negative. The former case does not change our results while that the latter, which we assume, is likely with more renewable energy penetration. We also assume (for simplicity) that σ_{η_i} equals a constant for all $i \in \{1, 2\}$ and compare cases when this constant equals zero and a strictly positive value. For $\sigma > 0$, we write $z(\sigma_{\eta_i} = \sigma)$ and $z(\sigma_{\eta_i} = 0)$ as short forms for $z(\sigma_{\eta_1} = \sigma, \sigma_{\eta_2} = \sigma)$ and $z(\sigma_{\eta_1} = 0, \sigma_{\eta_2} = 0)$, respectively. We define the lost REC value $\text{LRV}(z; \Theta) := \mathbb{E}[R_2\eta\mathbf{1}(\eta z > \alpha D)]$, which is the expected value of RECs when CPPA power supply exceeds the RPPT of αD.

Proposition 2.3. *For $\sigma > 0$, suppose ρ_{P_i,η_i} for all $i \in \{1, 2\}$ and $\rho_{R_2,\eta}$ are negative. Then $z(\sigma_{\eta_i} = \sigma) = 0$, if $z(\sigma_{\eta_i} = 0) = 0$. Instead, when $z(\sigma_{\eta_i} = 0) > 0$, we have $z(\sigma_{\eta_i} = \sigma) \leqslant z(\sigma_{\eta_i} = 0)$,* *if* $\mathrm{LRV}(\alpha D/\mathbb{E}[\eta]; \sigma_{\eta_i} = \sigma) \geqslant \mathrm{PRS}(K, \sigma_{\eta_i} = \sigma)$, *and* $z(\sigma_{\eta_i} = \sigma) \geqslant z(\sigma_{\eta_i} = 0)$, *otherwise.*

When zero CPPA capacity is purchased in the deterministic supply case ($\sigma_{\eta_i} = 0$), this proposition shows that the same is true with supply uncertainty ($\sigma_{\eta_i} > 0$). In other words, the buyer avoids using a CPPA more often under volatile supply. On the other hand, if nonzero CPPA capacity is optimal under deterministic supply, then we find that the CPPA capacity with uncertain supply can be greater than or smaller than $z(\sigma_{\eta_i} = 0)$ depending on whether the lost REC value is greater or smaller, respectively, than the REC and power price savings. In addition, this trend shows that supply variability can result in contracts with smaller CPPA capacity for firms with less aggressive RPPTs (i.e., smaller α) since the lost REC value $\mathrm{LRV}(\alpha D/\mathbb{E}[\eta]; \sigma_\eta = \sigma)$ decreases with α.

Although optimal CPPA capacities can increase or decrease when there is supply uncertainty, Proposition 2.4 establishes that the optimal procurement cost is nevertheless greater with supply uncertainty when correlations are negative. This difference is lower bounded by a quantity that increases with the magnitude of the power price, REC price, and supply volatilities. Moreover, this bound is larger for companies with more aggressive RPPTs (i.e., larger α).

Proposition 2.4. *For* $\sigma > 0$, *we have*

$$C(\sigma_{\eta_i} = \sigma) - C(\sigma_{\eta_i} = 0) \geqslant \frac{\alpha D}{y(\sigma_{\eta_i} = \sigma)} \left[- \sum_{i \in \{1,2\}} \rho_{P_i,\eta_i} \sigma_{P_i} \sigma - \rho_{R_2,\eta} \sigma_{R_2} \sigma \mathbf{1}(y(\sigma_{\eta_i} = \sigma) \geqslant \mathbb{E}[\eta]) \right].$$

Moreover, if ρ_{η_i,P_i} *for all* $i \in \{1,2\}$ *and* $\rho_{R_2,\eta}$ *are negative, then* $C(\sigma_{\eta_i} = \sigma) - C(\sigma_{\eta_i} = 0) \geqslant 0$.

Our analysis highlights that power price and supply volatilities increase procurement costs, as expected, but the magnitude of this increase is closely linked to the correlations between price and supply. The amount of capacity procured using a CPPA also depends on such correlations. It is interesting that CPPA procurement capacity increases in the presence of supply uncertainty when RPPTs are high, that is, RPPTs themselves play a role in determining whether a CPPA is an effective procurement vehicle. An insightful consequence is that companies using aggressive RPPTs have a higher potential of reducing their costs by including a CPPA as a procurement vehicle. Thus, a high RPPT in conjunction with a CPPA helps a firm balance climate goals (i.e., social responsibility) with energy costs (i.e., financial performance). Our results also shed light on the non-trivial nature of meeting an RPPT with a single CPPA. Practitioners have gone beyond a single CPPA and have started using dynamic portfolios of CPPAs with differing durations to both reach and sustain RPPTs, which we will turn to next.

3. A Markov decision process view of rolling power purchases

Two different procurement perspectives for meeting an RPPT are to (i) use short-term power and REC purchases and (ii) use a single large long-term CPPA with shortfalls procured via the short-term market. The former *short-term* policy exposes the buyer fully to price volatility. The latter long-term CPPA 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 and sustain an RPPT, which is a flexible strategy that maintains a portfolio of CPPA types with short and long tenors and signs new CPPAs, 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 §§4-5.

The planning horizon in our MDP is composed of two nested intervals as illustrated in Figure 1. The courser intervals are defined by *stages*, which are points in time when CPPA contracting decisions are made and we assume any RECs needed to meet the RPPT are purchased from the short-term market at this time. The finer intervals are demarcated by *inner-stages* at which (i) the corporation attempts to match power supply from the CPPA with demand using short-term power purchases, and (ii) exchanges cash flows with the generator to hedge the price for power generated from capacity contracted via CPPAs. In summary, the enforcement of the RPPT takes place at the decision stages, whereas CPPA settlements take place at inner-stages.

Figure 1: Illustration of planning horizon, stages, and inner-stages.

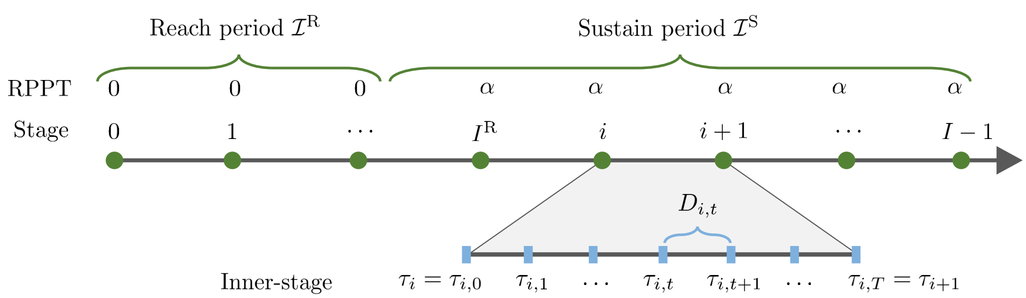

Formally, the horizon is comprised of I stages indexed by i belonging to set $\mathcal{I} := \{0, \ldots, I-1\}$. Stage i corresponds to time τ_i and time τ_0 is now. There are $T+1$ inner-stages between stages i and $i+1$ indexed by t with support in set $\mathcal{T} := \{0, \ldots, T\}$, where $\tau_{i,t}$ is the time associated with inner-stage t. Thus, $\tau_{i,0}$ and $\tau_{i,T}$ equal τ_i and τ_{i+1}, respectively. To ease exposition, we assume equi-distant annual stages and equi-distant inner-stages. An RPPT is enforced annually from year I^{R} onward, that is, a percentage $\alpha \in (0, 1]$ of the annual demand in each year $i \in \mathcal{I}^{\mathrm{S}} := \{I^{\mathrm{R}}, \ldots, I-1\}$ 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 $\mathcal{I}^{\mathrm{R}} := \{0, \ldots, I^{\mathrm{R}-1}\}$, but CPPAs can be signed. We thus refer to the years in sets $\mathcal{I}^{\mathrm{R}}$ and $\mathcal{I}^{\mathrm{S}}$, respectively, as the reach and sustain periods. The demand that needs to be satisfied in the interval $(\tau_{i,t}, \tau_{i,t+1}]$ in year $i \in \mathcal{I}$ is denoted by $D_{i,t}$ (MWh).

At each stage i, the set of potentially available CPPAs 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[4]. A CPPA of length m signed at stage i delivers power from stages $i + 1$ to $i + m$ at a fixed strike price of $K_{i,m}$ USD/MWh. The company can choose to enter into a new CPPA 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 MW[5] of the CPPA of length m years signed at year i. Given a contract availability vector a_i, the vector z_i belongs to set

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

which is non-convex when minimum purchase quantities are strictly positive (i.e., $z_m^{\min} > 0$).

The information required to make procurement decisions is described in the MDP state, which contains two components. The first component is a vector $x_i := (x_{i,l}, l \in \{0, \ldots, M-1\})$ representing the on-hand capacity from CPPA contracts, where $x_{i,l}$ is the total capacity in MW available at year $i + l$ by the on-hand CPPAs. This component is affected by the firm's decisions. The second component $w_i = (P_{i,0}, R_i, \eta_{i,0}, a_i)$ contains the stochastic factors needed to determine the Markovian evolution of the power price $P_{i,t}$ (USD/MWh), REC price R_i (USD/MWh), contract availability a_i, and stochastic power supply $\eta_{i,t} \in [0, H]$ (hours), where H is the number of hours between consecutive inner-stages, e.g., $H = 730$ for monthly inner-stages. The complete stage i MDP state is represented by the pair $(x_i, w_i) \in \mathcal{X}_i \times \mathcal{W}_i$. 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 CPPA capacity vector x_i to

$$x_{i+1,l} = f_i(x_i, z_i)_l = \begin{cases} x_{i,l+1} + \displaystyle\sum_{m \in \mathcal{M}: m > l} z_{i,m}, & \text{if } l \in \{0, \ldots, M-2\}; \\ z_{i,M}, & \text{if } l = M - 1, \end{cases} \tag{3}$$

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.

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

[5]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 energy production during this interval measured in hours.

For each stage $i \in \mathcal{I}$, the expected procurement cost accrued when entering into CPPAs is

$$c_i(x_i, w_i, z_i) = \mathbb{E}_i \Bigg[\sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t=0}^{T-1} \gamma_{\text{A}}^{l} \gamma^{t} (K_{i,m} - P_{i+l,t}) \eta_{i+l,t} z_{i,m} + \sum_{t=0}^{T-1} \gamma^{t} P_{i,t} D_{i,t} + \gamma_{\text{A}} R_{i+1} \max\left\{\alpha D_i - \eta_i x_{i,0}, 0\right\} \mathbf{1}_{\{i \in \mathcal{I}^{\text{S}}\}} \Bigg], \tag{4}$$

where $\mathbb{E}_i[\cdot] \equiv \mathbb{E}[\cdot | w_i]$, $L_{i,m} := \min\{m,\, I - i\}$ equals the number of stages of power delivery within the planning horizon, $\gamma \in [0, 1)$ denotes the discount factor over the interval defined by inner-stages, $\gamma_{\text{A}} = \gamma^T \in [0, 1)$ is the annualized discount factor, and $D_i := \sum_{t=0}^{T-1} D_{i,t}$ and $\eta_i := \sum_{t=0}^{T-1} \eta_{i,t}$ represent the aggregate (annual) power demand and power supply, respectively. The first term in (4) models the sum of the settlement cash flows between the generator and the company at each interval over the tenure of the CPPAs signed from set $\mathcal{M}$. Here $\eta_{i+l,t} z_{i,m}$ is the actual power generated in the interval $(\tau_{i,t}, \tau_{i,t+1}]$ of stage $i + l$ by the $z_{i,m}$ MW CPPA signed at stage i with tenure m. The second term is the cost of procuring the known demands $D_{i,t}$, $t = 0, \ldots, T - 1$ from the short-term market at prices $P_{i,t}$, $t = 0, \ldots, T - 1$. The third term accounts for the REC purchasing costs at a price R_{i+1} USD/MWh to cover any shortfall between the RPPT αD_i in year $i \in \mathcal{I}^{\text{S}}$ and the renewable power supply $\eta_i x_{i,0}$ in this year. We assume that at stage $I - 1$ only short-term procurement of power and RECs is possible, which means

$$c_{I-1}(x_{I-1}, w_{I-1}) = \mathbb{E}_{I-1} \left[\sum_{t=0}^{T-1} \gamma^{t} P_{I-1,t} D_{I-1,t} + \gamma_{\text{A}} R_I \max\left\{\alpha D_{I-1} - \eta_{I-1} x_{I-1,0}, 0\right\} \right]. \tag{5}$$

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 $\mathcal{I}_i := \{i, \ldots, I - 2\}$. A decision rule $Z_i^{\pi_i}$ 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 $\mathcal{Z}_i(a_i)$. 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 policy in Π_i solves

$$V_i(x_i, w_i) := \min_{\pi_i \in \Pi_i} \ \mathbb{E}_i \left[\sum_{j \in \mathcal{I}_i} \gamma_{\text{A}}^{j-i} c_j(x_j^{\pi_i}, w_j, Z_j^{\pi_i}(x_j^{\pi_i}, w_j)) + \gamma_{\text{A}}^{I-1-i} c_{I-1}(x_{I-1}^{\pi_i}, w_{I-1}) \right], \tag{6}$$

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).

An optimal policy of MDP (6) maintains a dynamically evolving portfolio of CPPAs to reach and sustain an RPPT. The MDP policy decision at a stage i accounts for (P1) the future stochasticity of prices and supply, (P2) the full flexibility to time CPPA purchases, and (P3) the access to a variety of CPPA types. While these properties are appealing, computing the optimal MDP policy is challenging due the well-known curses of dimensionality associated with MDPs (Bertsekas

2011, Powell 2011). Specifically, in MDP (6), the endogenous state x_i and decision z_i are M- and $|\mathcal{M}|$-dimensional continuous vectors, respectively, and the exogenous state w_i may also be high dimensional when using a multi-factor stochastic model for the evolution of uncertainty. In addition to dimensionality issues, another source of intractability stems from the non-convex structure of $\mathcal{Z}_i(a_i)$. 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 (see, e.g., Brown and Smith 2014, Salas and Powell 2017, Nadarajah and Secomandi 2018). The value function of MDP (6) is not convex in general as we show via an example in Online Supplement 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 3.1. (We refer the reader to Online Supplement C for further discussion.)

Proposition 3.1. *Suppose (i) $z_m^{min} = 0$ and $z_m^{max} < \infty$ for all $m \in \mathcal{M}$, and (ii) the value function satisfies $|V_i(\cdot, \cdot)| < \infty$ for each stage $i \in \mathcal{I}$. Then the value function $V_i(\cdot, w_i)$ is convex in the endogenous state x_i for each stage and exogenous state $(i, w_i) \in \mathcal{I} \times \mathcal{W}_i$.*

Owing to these difficulties we pursue simpler strategies. In §4, the policies we consider are based on forecasts. They sacrifice property P1 of the MDP policy but partially/fully capture properties P2 and P3. We then introduce in §5 policies based on information relaxation and duality (Andersen and Broadie 2004, Haugh and Kogan 2004, Brown et al. 2010) satisfying properties P1, P2, and P3.

4. Forecast-based rolling power purchase policies

The traditional use of long-tenure CPPAs in practice can be viewed as a special case of a procurement heuristic that employs a single CPPA. This CPPA is renewed in a rolling fashion (WBCSD 2018) by solving a deterministic model derived from MDP (6) by only allowing CPPAs of a single type to be signed just-in-time to meet and sustain the RPPT, replacing uncertain quantities by forecasts, and limiting these forecasts to the tenure of the CPPA type, that is, they may not extend to the end of the planning horizon. This heuristic, dubbed the *forecast-based block heuristic* (FBH), is parameterized by a fixed $m \in \mathcal{M}$ and works as follows. The first contract is entered at the last year of the reach period, $I^{\mathrm{R}} - 1$, and delivers renewable power during the first m years of the sustain period. The second contract is ordered one year before the first contract expires to ensure the continuous delivery of power from CPPAs. This process is repeated until the end of the planning horizon. If at year i contract m is not available, that is, $a_{i,m} = 0$, then the RPPT in $i+1$ is fulfilled by unbundled

RECs and the CPPA procurement decision is postponed to year $i+1$. The optimal CPPA capacity is

$$z^*_{i,m} = \operatorname*{argmin}_{z_{i,m}\in\mathcal{Z}_i(a_i)} \sum_{l=1}^{L_{i,m}} \gamma_{\mathrm{A}}^l \Bigg[\sum_{t=0}^{T-1} \gamma^t (K_{i,m} - \mathbb{E}_i[P_{i+l,t}])\, \mathbb{E}_i[\eta_{i+l,t}] z_{i,m} \;+\; \sum_{t=0}^{T-1} \gamma^t\, \mathbb{E}_i[P_{i+l,t}] D_{i+l,t}$$

$$+\, \mathbb{E}_i[R_{i+l+1}] \max\left\{\alpha D_{i+l} \,-\, \mathbb{E}_i[\eta_{i+l}] z_{i,m}, 0\right\}\Bigg]. \tag{7}$$

Compared to the optimal policy of MDP (6), FBH has low timing flexibility and no CPPA variety because new CPPAs are signed only when the incumbent contract expires and it uses a single CPPA type. We introduce a rolling-planning approach that extends FBH in terms of timing flexibility and CPPA variety. This approach, dubbed *portfolio forecast-based reoptimization heuristic* (PFRH), solves at each stage a model obtained by replacing random quantities in MDP (6) by their respective forecasts, i.e., expected values. In the case of contract availability, given the binary nature of this variable, we assign a forecast of 1 if the contract is available with probability greater than 0.5, and 0 otherwise. Formally, the stage j forecast for contract $m \in \mathcal{M}$ made at stage i, with $i \leqslant j$, 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 (4) with expectations replacing random quantities:

$$c_i^{\mathrm{DET}}(x_i, w_i, z_i) = \sum_{m\in\mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}}^l \gamma^t (K_{i,m} - \mathbb{E}_i[P_{i+l,t}])\, \mathbb{E}_i[\eta_{i+l,t}] z_{i,m} \;+\; \sum_{t=0}^{T-1} \gamma^t\, \mathbb{E}_i[P_{i,t}] D_{i,t}$$

$$+\; \gamma_{\mathrm{A}} \mathbb{E}_i[R_{i+1}] \max\left\{\alpha D_i \,-\, \mathbb{E}_i[\eta_i] x_{i,0}, 0\right\} \mathbf{1}_{\{i\in\mathcal{I}^{\mathrm{S}}\}}. \tag{8}$$

We omit the expression for c_{I-1}^{DET} as it is a simplification of the terminal cost (5). At stage i and state $(x_i, w_i) \in \mathcal{X}_i \times \mathcal{W}_i$, PFRH solves

$$\min_{y_j, z_j} \sum_{j\in\mathcal{I}_i} \gamma_{\mathrm{A}}^{j-i} c_j^{\mathrm{DET}}\big(y_j, \mathbb{E}_i[w_j], z_j\big) + \gamma_{\mathrm{A}}^{I-1-i} c_{I-1}^{\mathrm{DET}}\big(y_{I-1}, \mathbb{E}_i[w_{I-1}]\big) \tag{9a}$$

$$\text{s.t.: } y_i = x_i, \tag{9b}$$

$$y_{j+1} = f_j(y_j, z_j), \;\; \forall\, j \in \mathcal{I}_i, \tag{9c}$$

$$y_j \in \mathcal{X}_j, \;\; \forall\, j \in \mathcal{I}_i \cup \{I-1\}, \tag{9d}$$

$$z_j \in \mathcal{Z}_j(\bar{a}_{i,j}), \;\; \forall\, j \in \mathcal{I}_i. \tag{9e}$$

This math program computes decisions $z_j \in \mathcal{I}_i$ and includes auxiliary variables y_j to track the endogenous MDP state. Its objective (9a) is the sum of discounted procurement costs. Constraint (9b) initializes the stage i state to the current state x_i. Constraints (9c) ensure the feasibility of state transitions. Constraints (9d)–(9e) restrict decision variables to their respective feasible sets.

Although solving (9) at stage i and state (x_i, w_i) provides the PFRH 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 (9) 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 $I - 1$.

We also consider a variant of PFRH that can only sign a single CPPA type. We refer to this approach as FRH, which has full timing flexibility but no CPPA variety, and is thus more flexible than FBH but less flexible than PFRH. The computational burden of using the FBH, FRH, and PFRH 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 non-convexities discussed in §3. Note that both (7) and (9) are mixed integer programs that can be solved using off-the-shelf commercial optimization software. Conceptually, by replacing random quantities by their respective forecasts in our setting, FBH, FRH, and PFRH do not capture the impact of the future evolution of uncertainty on the current procurement decision. It is apriori unclear how this factor affects the policy performance of these forecast-based rolling planning approaches. In addition, the optimal objective of the PFRH 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 (see, e.g., Gallego and Van Ryzin 1997). We show in Online Supplement D that such a lower bounding property is not true for even a "convexified" version of the PFRH math program.

5. Information-relaxation based reoptimization heuristic

We propose a procurement policy 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 heuristics discussed at the end of §4. We describe the dual bound in §5.1 and then present in §5.2 the reoptimization scheme used to define the new policy.

5.1 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 (6) as well. In its most commonly used form, a dual bound is estimated in Monte Carlo simulation by solving a deterministic variant of MDP (6) endowed with full information about future uncertainty and costs adjusted for this knowledge using a dual penalty. Let $q_i\big(x_i, W_i, z_i\big)$ and $q_{I-1}\big(x_{I-1}, W_{I-1}\big)$ denote respectively the stage $i < I - 1$ and stage $I - 1$ dual penalty function, where $W_i := (w_{j,t}, j = i \ldots, I, \, t \in \mathcal{T})$ is a vector of realized stochastic factors for each stage from i to I and interval $t \in \mathcal{T}$. If $\mathbb{E}_j[q_j\big(x_j, W_j, z_j\big)] \geqslant 0$ holds for

$j \in \mathcal{I}_i$ and $\mathbb{E}_{I-1}[q_{I-1}(x_{I-1}, W_{I-1})] \geqslant 0$, 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^{IR}_{I-1})

$$\begin{aligned} c_i^{\mathrm{IR}}(x_i, W_i, z_i) = \sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}}^l \gamma^t (K_{i,m} - P_{i+l,t})\, \eta_{i+l,t} z_{i,m} \;+\; \sum_{t=0}^{T-1} \gamma^t\, P_{i,t} D_{i,t} \\ +\; \gamma_{\mathrm{A}} R_{i+1} \max\left\{\alpha D_i \;-\; \eta_i x_{i,0}, 0\right\} \mathbf{1}_{\{i \in \mathcal{I}^{\mathrm{S}}\}}, \end{aligned} \tag{10}$$

and consider the following deterministic optimization problem

$$V_i^{\mathrm{IR}}(x_i; W_i) = \min_{y_j, z_j} \sum_{j \in \mathcal{I}_i} \gamma_{\mathrm{A}}^{j-i} \Big[c_j^{\mathrm{IR}}(y_j, W_j, z_j) - q_j(y_j, W_j, z_j) \Big]$$
$$+\, \gamma_{\mathrm{A}}^{I-1-i} \Big[c^{\mathrm{IR}}_{I-1}(y_{I-1}, W_{I-1}) - q_{I-1}(y_{I-1}, W_{I-1}) \Big] \tag{11a}$$
$$\text{s.t.:}\;\; y_i = x_i, \tag{11b}$$
$$y_{j+1} = f_j(y_j, z_j), \;\; \forall\, j \in \mathcal{I}_i, \tag{11c}$$
$$y_j \in \mathcal{X}_j, \;\; \forall\, j \in \mathcal{I}_i \cup \{I-1\}, \tag{11d}$$
$$z_j \in \mathcal{Z}_j(a_j), \;\; \forall\, j \in \mathcal{I}_i. \tag{11e}$$

Constraints (11b)–(11d) are identical to constraints (9b)–(9d) in the math program solved by PFRH. Constraints (11e) differ from (9e) 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 while in the latter case we use the contract availability forecast vector described in §4. The objective (11a) can be obtained by modifying the PFRH objective (9a) by subtracting dual penalty terms and replacing the forecasted uncertainty with the elements of W_i. The expectation $\mathbb{E}_i[V_i^{\mathrm{IR}}(x_i; W_i)]$ 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 (11). Choosing this function to be zero, i.e., $q_i(\cdot,\cdot,\cdot) \equiv 0$, 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:

$$q_{I-1}(x_{I-1}, W_{I-1}) = c^{\mathrm{IR}}_{I-1}(x_{I-1}, W_{I-1}) - c_{I-1}(x_{I-1}, W_{I-1}), \tag{12a}$$
$$q_i(x_i, W_i, z_i) = \gamma_{\mathrm{A}} \Big\{ V_{i+1}\big(f_i(x_i, z_i), w_{i+1}\big) - \mathbb{E}_i\big[V_{i+1}\left(f_i(x_i, z_i), w_{i+1}\right) \big] \Big\}$$
$$+\, c_i^{\mathrm{IR}}(y_i, W_i, z_i) - c_i(y_i, w_i, z_i). \tag{12b}$$

The term $c_i^{\mathrm{IR}}(y_i, W_i, z_i) - c_i(y_i, w_i, z_i)$ is atypical in the definition of an ideal penalty but needed here since our MDP cost function (4) includes expectations over future exogenous states; hence it differs from its hindsight counterpart (10). Since the exact value function in (12b) is not available, value function approximations can be used instead or simpler dual penalties that do not rely on such approximations can be constructed (see, e.g., Brown and Smith 2014, Secomandi 2015, Nadarajah and Secomandi 2018). We take the latter approach and define the dual penalty function as

$$q_i(x_i, W_i, z_i) := \sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}}^l \gamma^t \Big[\beta_i^{\mathrm{P}} (\mathbb{E}_i[P_{i+l,t}] - P_{i+l,t}) + \beta_i^{\eta} (\mathbb{E}_i[\eta_{i+l,t}] - \eta_{i+l,t})\Big] z_{i,m} + \sum_{m \in \mathcal{M}} \sum_{l=1}^{L_{i,m}} \gamma_{\mathrm{A}}^l \beta_i^{\mathrm{R}} (\mathbb{E}_i[R_{i+l}] - R_{i+l}) z_{i,m}, \tag{13}$$

where $(\beta_i^{\mathrm{P}}, \beta_i^{\mathrm{R}}, \beta_i^{\eta})$ are stage- and factor-dependent weights. In (13), the information gained when taking a decision is approximated by spreads between the value taken by the uncertainty in future periods $(i + l, t)$ and its expectation computed at period $(i, 0)$. These spreads are multiplied by the CPPA purchase decisions. The dual penalty (13) is linear in z_i, does not depend on x_i, and is feasible because the expectation of $\mathbb{E}_i[P_{i+l,t}] - P_{i+l,t}$ equals zero (analogously for the other factors). In addition to its simple form, linear penalties ensure that the math program (11) falls into the same complexity class as the deterministic version of MDP (6) and the math program (9) solved by PFRH.

5.2 Decisions

Traditionally, an operating policy is computed independent of the dual bound computation described above (see, e.g., Desai et al. 2012 and the related discussion in §1.1). We now define a non-anticipative decision directly during the dual bound estimation process, where being non-anticipative refers to a decision that only depends on the information available at stage i (the decision resulting from solving (11) is anticipative as it relies on future information on the sample path W_i). Since the dual bound $\mathbb{E}_i[V_i^{\mathrm{IR}}(x_i; W_i)]$ involves solving the math program (11) 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 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 non-anticipative decision. We call this functional a *decision measure* $\mathcal{H}_i$ that maps a distribution $z_i(W_i)$ to a vector of $\mathbb{R}^{|\mathcal{M}|}$. Proposition 5.1 establishes some useful properties of a decision measure with respect to optimality and feasibility.

Proposition 5.1. *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. $\mathcal{Z}_i(a_i)$ is convex and $\mathcal{H}_i(z_i(W_i)) := \mathbb{E}_i[z_i(W_i)]$.

2. $\mathcal{H}_i(z_i(W_i)) = \bar{z}_i$ and $\bar{z}_{i,m} = z_{i,m}(W_i^m)$ for all $m \in \mathcal{M}$, where W_i^m is a realization of the random variable $W_i|w_i$.

Moreover, if the decision $\mathcal{H}_i(z_i(W_i))$ satisfies one of the conditions above and we use the ideal dual penalty (12) in math program (11), then $\mathcal{H}_i(z_i(W_i))$ is an optimal solution to MDP (6).

Examples of $\mathcal{H}_i$ that satisfy conditions (1) and (2) of Proposition 5.1 are $\mathcal{H}_i(z_i(W_i)) = \mathbb{E}_i[z_i(W_i)]$ and $\mathcal{H}_i(z_i(W_i)) = \mathbb{M}[z_i(W_i)|w_i]$, respectively, 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 $m \in \mathcal{M}$). This definition of median leads to a feasible decision in our application since there are no constraints on the action space $\mathcal{Z}_i(a_i)$ that tie different contracts in $\mathcal{M}$. It is important to note the feasibility of $\mathcal{H}_i(z_i(W_i))$ under condition (2) holds even when $\mathcal{Z}_i(a_i)$ is non-convex. Thus, in our application the median decision measure is more robust in terms of feasibility than the mean decision measure since the latter measure satisfies only condition (1). More broadly, the 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 §4 that employ forecasts. 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 action distribution can be useful.

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 $\{w_{j,t}^h, (j,t,h) \in \{i,\ldots,I\} \times \mathcal{T} \times \{1,\ldots,H\}\}$ which provide a discrete approximation $\hat{W}_i|w_i$ of the random variable $W_i|w_i$. Based on this approximation, we estimate both a dual bound $\sum_{h=1}^H V_i^{\mathrm{IR}}(x_i; W_i^h)/H$ and a decision $\mathcal{H}_i(z_i(\hat{W}_i))$, which requires the solution of H math programs of the type (11). We apply the aforementioned decision to move to an endogenous state $x_{i+1} = f_i(x_i, \mathcal{H}_i(z_i(\hat{W}_i)))$. 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. We call the resulting approach the *portfolio information-relaxation based reoptimization heuristic* (PIRH) and the policy computed in the process using decision measures as the PIRH policy.

6. Numerical study

In §6.1 and §6.2, respectively, we describe our realistic instances and explain the computational setup. In §§6.3-6.6, we present our results and discuss related managerial insights.

6.1 Procurement instances

Table 1 summarizes the parameter values of our baseline MDP instance. We use a 40 year planning horizon (I) and a 5 year reach period (I^{R}) to attain a stochastic RPPT that is 90% (α). Within each year, there are 12 monthly inner-stages (T). The set of contracts $\mathcal{M}$ and the minimum/maximum quantities are consistent with the CPPA portfolio of Google (Google 2017). The corresponding availability factors p_m are chosen based on Wiser and Bolinger (2017) and Baker McKenzie (2015): According to the first report, 15- to 25-year CPPAs are predominant, with 20-year contracts being the most common, while the second report indicates that CPPAs of length between 10 and 20 years are prevalent. We use a constant monthly power demand $D_{i,t}$ throughout the horizon which corresponds to the consumption of two large data centers. The annual risk-free rate is set to 3.1% and corresponds to the average 10-year United States treasury rate in November 2018 (Bloomberg 2018), which implies a discount factor γ_{A} equal to 0.97.

Table 1: Parameters defining the baseline CPPA instance.

Name	Value	Unit	Name	Value	Unit	Name	Value	Unit
I	40	years	$D_{i,t}$	$5 \cdot 10^4\ (\forall i, t \in \mathcal{I} \times \mathcal{T})$	MWh/month	$z_m^{\min}$	$20\ (\forall m \in \mathcal{M})$	MW
I^{R}	5	years	$\mathcal{M}$	$\{5, 10, 15, 20, 25\}$	years	$z_m^{\max}$	$400\ (\forall m \in \mathcal{M})$	MW
T	12	months	p_m	$\{0.3, 0.4, 0.5, 0.5, 0.4\}$	-	γ_{A}	0.97	-
α	90%	-	$\bar{\rho}$	$-0.2\ (\forall i, t \in \mathcal{I} \times \mathcal{T})$	-	r_{A}	0.93	-

We designed a realistic model to obtain CPPA strike prices. A renewable power generator typically sets a strike price to recoup its project investment and maintenance costs as well as a return on investment. We chose the generator's discount factor (r_{A}) so that the return on investment is roughly twice the risk-free interest rate. In addition, historical data and models from NREL show that the CPPA strike price is affected by several factors including the expected quantity of power produced as a fraction of installed capacity (i.e., capacity factor), tax credits, improvements in technology, the contract duration, and the expected power price over the tenor of the contract (NREL 2010, DOE 2016, Wiser and Bolinger 2017). We develop a model that accounts for these factors based on a net present value (NPV) calculation consistent with software from NREL (NREL 2017). For instance, we account for a production tax credit (PTC) per MWh that is granted to CPPAs signed in the first five years of the planning horizon. We relegate details of the strike price model to Online Supplement E.

We consider stochastic processes for the power price, REC price, and power supply. We consider monthly power prices, which can be seen as the firm purchasing power from the utility at a variable retail rate that tracks the average wholesale price (see, e.g., the CPPA signed by Amazon in RMI

2016)[6]. We model the power price using a mean-reverting stochastic process with seasonality and jumps (Lucia and Schwartz 2002, Weron 2014). The monthly power supply is also modeled using a mean-reverting process but without jumps (Loukatou et al. 2018). Power price and power supply share instantaneous correlation $\rho_{P_i\eta_i}$ equal to a constant $\bar{\rho}$ of -0.2 for all stages based on the study by Woo et al. (2011). We employ a Jacobi diffusion process to forecast yearly REC prices following Zeng et al. (2015) and assume zero correlation between REC price and power supply as discussed in §2. We employ a maximum likelihood estimation process to calibrate the power price, REC price, and power supply models using real data from PJM, New Jersey, and EIA, respectively. Availability of CPPA $m \in \mathcal{M}$ follows a Bernoulli random variable, where $p_m \in [0, 1]$ denotes the probability that this contract is available. Online Supplement F contains details of these models and the calibration.

Table 2 summarizes six instance sets S1–S6, with 34 instances in total, obtained by perturbing the parameters of our baseline instance (distinguished by a superscript B). These instances allow us to analyze the behavior of procurement policies as market parameters change.

Table 2: Extended instance sets with the baseline-instance parameter superscripted by B.

Set	Modified parameter	Values
S1	Renewable energy target α	$\{0, 0.1, \ldots, 0.9^{\mathrm{B}}, 1\}$
S2	CPPA availability change p_m for all $m \in \mathcal{M}$	$\{-0.2, -0.1, 0^{\mathrm{B}}, +0.1, +0.2, +0.3\}$
S3	Long-term mean of power price	$\{20, 25, 30, 39.7^{\mathrm{B}}\}$ USD/MWh
S4	Long-term mean of RECs price	$\{5, 9.4^{\mathrm{B}}, 15, 20\}$ USD/MWh
S5	Generator annual discount factor r_{A}	$\{0.9, 0.91, 0.92, 0.93^{\mathrm{B}}, 0.94, 0.95\}$
S6	Power price and supply correlation $\bar{\rho}$	$\{-1, -0.66, -0.33, 0, 0.33, 0.66, 1\}$

6.2 Computational setup

Our computational study is setup to shed light on how to structure effective procurement strategies to meet an RPPT. To this end, we evaluate the procurement heuristics of §§4–5 (i.e., FBH, FRH, PFRH, and PIRH) as well as estimate cost components via Monte Carlo simulation. We denote by FBH_m and FRH_m the specifications of these heuristics that use CPPA m from set $\mathcal{M}$. We implemented PIRH using linear dual penalties defined as in (13) with weights $(\beta_i^{\mathrm{P}}, \beta_i^{\mathrm{R}}, \beta_i^{\eta}) = (0.3, 0, 0)$ for each $i \in \mathcal{I}$, i.e. that include power price spreads alone, and $q_{I-1}(\cdot,\cdot) \equiv 0$. We found that this dual-penalty specification led to near-optimal procurement policies and lower bounds.

All algorithms were programmed using C++ with Gurobi 8.1 as the math programming solver. We estimated the value of the procurement policies (i.e., expected discounted total costs over the planning horizon) and the PIRH lower bound using 1000 Monte Carlo sample paths (i.e., H

[6] Our MDP and methods can be used at finer time scales at the expense of higher computational cost.

equals 1000), as this choice resulted in standard errors of lower/upper bound estimates of 0.36% on average and at most 0.53%. We followed an analogous approach to obtain low standard errors when estimating other quantities. The PIRH upper bound estimation process requires computing procurement decisions at each stage of an evaluation sample path. Therefore, for a fixed stage and evaluation sample path, we formulated and solved the math program (11) on 30 inner sample paths and then applied a decision measure to back out a non-anticipative control.

6.3 Managing the cost of renewable power certification

The company's cash flow function (4) is the sum of two components: (i) the term $\sum_{t=0}^{T-1}\gamma^t P_{i,t}D_{i,t}$ captures the cost of purchasing power from the short-term market and (ii) the remaining terms corresponding to the expected cost of "certifying" this power as being renewable through CPPAs and/or unbundled RECs. We refer to the latter component as the certification cost, which simplifies as follows for a setting with 1 MWh of power demand in each hour, an RPPT of 100%, and a 1 MW CPPA of length $m \in \mathcal{M}$ entered at stage i:

$$u_i^m := \sum_{l=1}^{L_{i,m}} \gamma_{\mathrm{A}}^l \left[\sum_{t=0}^{T-1} \gamma^t \Big(K_{i,m}\mathbb{E}_i[\eta_{i+l,t}] - \mathbb{E}_i[P_{i+l,t}]\mathbb{E}_i[\eta_{i+l,t}] - \mathrm{Cov}(P_{i+l,t}, \eta_{i+l,t}) \Big) + \mathbb{E}_i[R_{i+l+1}(TH - \eta_{i+l})] \right].$$

The certification cost terms inside the summation explicitly capture the covariance between power price and supply, where $\mathrm{Cov}(P_{i+l,t}, \eta_{i+l,t}) = \bar{\rho}\,\sigma_{P_{i+l,t}}\sigma_{\eta_{i+l,t}}$. Specifically, when $\bar{\rho}$ takes positive and negative values, respectively, the certification cost decreases and increases relative to the case with zero correlation because of the covariance term. The remaining term related to short-term REC purchases replaces the 1 MWh of demand by its implied annual demand of TH MWh. To be able to compare certification costs across CPPAs of different tenures, we annuitize u_i^m, which amounts to converting the total cash flow into a sequence of equal discounted monthly payments and then dividing this discounted cost by H (i.e., the hours per month) to obtain a 1 MWh cost. Henceforth, we denote by u_i^m the annuitized unit cost (with an abuse of notation) and by u_i^{ST} the certification cost per unit when using no CPPAs and only short-term REC purchases.

Figure 2(a) displays the impact that the signing year i has on a 20-year CPPA, that is u_i^{20}, for $\bar{\rho} \in \{-1, 0, 1\}$, as well as the behavior of u_i^{ST} with year i, which is not affected by correlation. The cost u_i^{ST} varies only marginally with stage i while u_i^{20} displays a significant jump at year 5 followed by a decrease over time. The former jump in certification cost stems from the PTC expiring five years into the planning horizon. On our baseline instance, entering into a 20-year CPPA at year 5 (i.e., u_5^{20}) rather than year 4 (i.e., u_4^{20}) translates into an annual expected increase in expenditure of about 2 mln. USD. This cost increase makes certifying power to be renewable using a 20-year CPPA more

Figure 2: Expected certification cost for 1 MWh of power.

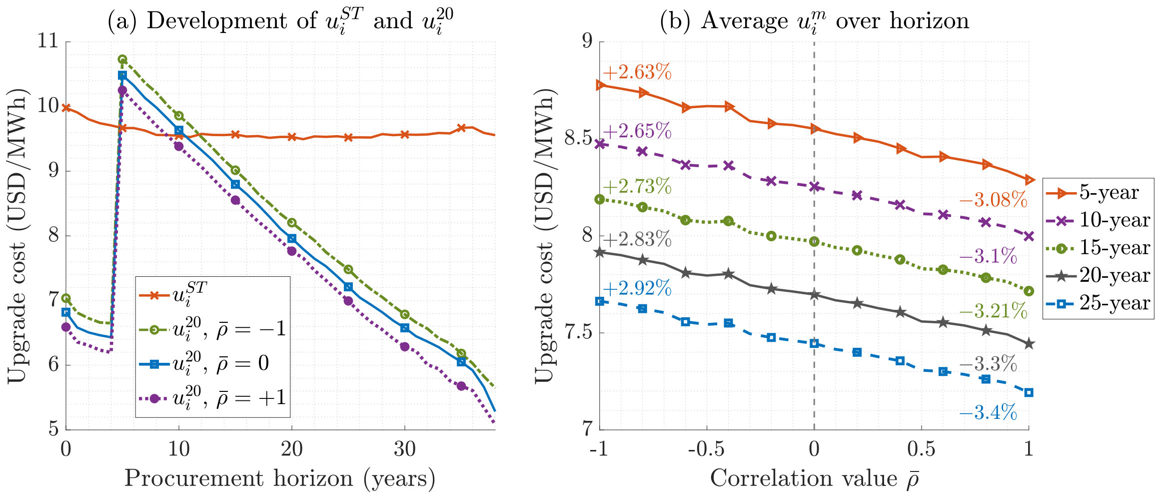

expensive than procuring RECs from the short-term market (i.e., $u_5^{20} > u_5^{\mathrm{ST}}$). The decrease in u_i^{20} after year 5 is the result of improving technology and efficiency which brings down the CPPA strike price. Despite the significant impact of subsidies on the cost of CPPAs, the use of such agreements can substantially decrease the cost of certifying power to be renewable for most stages in the planning horizon and all three correlation levels compared to buying RECs from the short-term market.

As expected from our analysis in §2, the values of u_i^{20} in Figure 2(a) for $\bar{\rho}$ equal to 1 and -1 is smaller than and greater than, respectively, this value for zero correlation. Next, we assess the certification cost change due to the correlation between power price and supply. Figure 2(b) displays the average annuitized certification cost $\sum_{i\in\mathcal{I}} u_i^m/|\mathcal{I}|$ for CPPA m as a function of $\bar{\rho}$. The averaging here smooths the impact of the signing year i and allows us to focus on the covariance effect. Varying $\bar{\rho}$ from $+1$ to -1 significantly increases the average certification cost by about 6% for all contracts. This finding suggests that companies need to account for the impact of power price and supply correlations; otherwise, their cost estimates for certifying power to meet an RPPT could be biased.

6.4 Shape risk and RPPT settlement time scale

MDP (6) has two time scales defined by inner-stages and stages. The time interval separating inner-stages depends on how often short-term power prices change, as this affects when power supplied by the CPPA is matched with power demand. Indeed the firm can negotiate with a utility on index pricing and the frequency of price changes but this is not under its direct control. In contrast, the company directly controls the time scale at which the RPPT is met. We assume that this time scale coincides with the decision stages but this need not be the case.

Using the block heuristic FBH_m, we explore in Figure 3(a) how the procurement cost in the baseline instance varies with the RPPT settlement time scale. Reducing this time scale from yearly to monthly causes the expected procurement cost to increase by 3.5%, 4.4%, and 3.2%, respectively, for 15-year, 20-year, and 25-year CPPAs. To understand such procurement cost increases, consider a 1 MW CPPA that delivers power in year 10. This CPPA delivers different amounts of power each month due to a seasonal component and a random component, both embedded in our stochastic model for power supply. Figure 3(b) shows the average delivery profile of this contract with the RPPT depicted using the dotted line. Under annual RPPT settlements, the average CPPA is sufficient to meet the target since months with supply shortfalls are balanced by ones with excess supply. Instead, if the RPPT is met every six months, then the supply from the CPPA is enough to fulfill the target in months 1–6 but 0.12 MWh of unbundled RECs are needed in months 7–12, which results in higher procurement cost. Similarly, when using a monthly RPPT commitment, 0.2 MWh of unbundled RECs need to be purchased in months 7, 8, and 9.

Figure 3: Effect of different contract lengths and RPPT settlements on FBH_m procurement.

(a) Procurement costs (b) Power supply

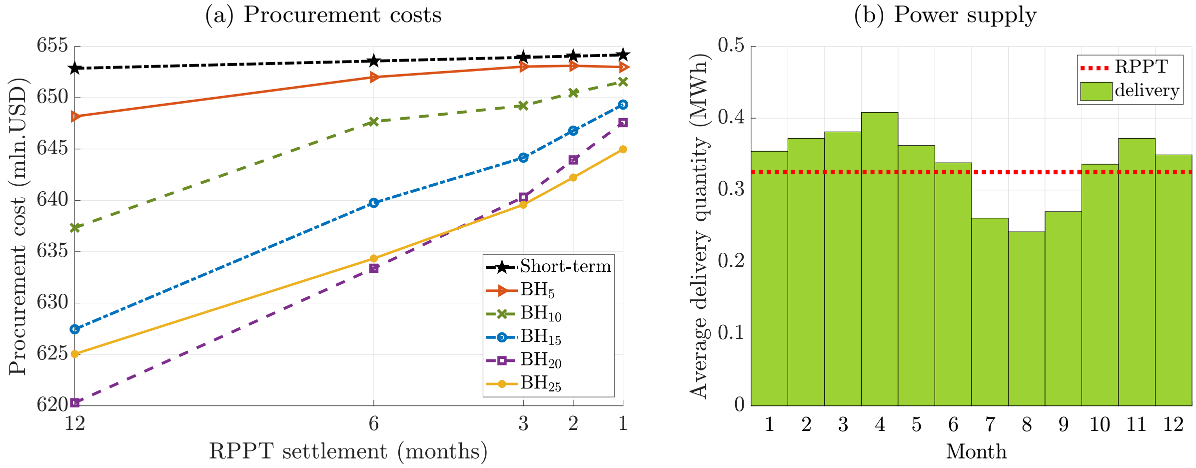

Intuitively, when using a shorter RPPT settlement, the shape of supply more heavily impacts the exposure of a firm to the volatile short-term REC price. Indeed, when using a CPPA, supply uncertainty already exposes a firm to volatile power prices (see discussion in §2). What we show above is that this shape risk arising from volatile power prices is further exacerbated by exposure to volatile REC prices in the presence of an RPPT. Fortunately, shape risk due to REC prices can be mitigated by using less frequent RPPT settlements, which motivates an annual RPPT time scale.

6.5 Value of flexibility in rolling power purchases and portfolio optimization

Rolling power purchase strategies discussed in §4 differ based on their timing flexibility and CPPA variety. We begin by investigating the value of timing flexibility in single-CPPA policies of FBH_m and FRH_m (i.e., the value of their respective policies when using CPPA type $m \in \mathcal{M}$). Recall that the former policy has less timing flexibility and less forecast information than the latter one. Table 3 reports their performance on the baseline instance as well as their average performance over the instance sets S1–S6. The performance of the policy that executes short-term power and REC purchases alone is also reported as a benchmark. Optimality gaps are computed with respect to the PIRH dual bound.

Table 3: Procurement costs in mln. USD and optimality gaps (in parenthesis) for FBH_m, FRH_m, and short-term policies.

Instance		$m = 5$	$m = 10$	$m = 15$	$m = 20$	$m = 25$	Short-term
Baseline	FBH_m	648.2 (14.9%)	637.3 (13.0%)	627.4 (11.2%)	620.3 (10.0%)	625.0 (10.8%)	652.9
	FRH_m	638.6 (13.2%)	624.5 (10.7%)	607.8 (7.8%)	591.6 (4.9%)	582.7 (3.3%)	(15.7%)
S1–S6	FBH_m	621.6 (11.2%)	614.1 (9.9%)	606.7 (8.6%)	601.0 (7.6%)	602.1 (7.8%)	626.4
	FRH_m	613.0 (9.7%)	601.8 (7.7%)	590.5 (5.7%)	579.9 (3.8%)	574.5 (2.9%)	(12.1%)

FRH_m results in a lower procurement cost than FBH_m under each contract in set $\mathcal{M}$ and these policies for longer tenure CPPAs significantly improve on the short-term policy. The difference between FRH and FBH increases with the contract length and can be up to 42.3 mln. USD on the baseline instance when using 25-year CPPAs. Further investigation showed that the average number of years between signing two consecutive CPPAs for FRH_{25} was roughly 12 years while FBH_{25} has a much longer average interval equal to 26 years. Moreover, the minimum, average, and maximum capacity of CPPAs (in MW) signed by FRH_{25} are 22.2, 114.6, and 185.1, respectively, while analogous values for FBH_{25} are 182.6, 187.9, and 193.2. These statistics show that FRH_{25} uses its additional timing flexibility to sign CPPAs more often and with less capacity compared to FBH_{25}, which results in significant cost reduction. It is also encouraging that FRH_{20} and FRH_{25} have low optimality gaps (between 3% and 5%). However, in the absence of long-term CPPAs spanning 20 to 25 years, using a shorter tenure CPPA instead can lead to a substantial increase in procurement cost.

Motivated by the above finding, we assess the value of signing multiple CPPA types, that is, the benefit of maintaining portfolios of CPPAs with different tenures. Table 4 reports the performance of PFRH, PIRH with mean and median decision measures, and a PIRH variant with zero dual penalty, labeled PIRH-0. The PFRH optimality gaps are worse than the ones for FRH_{25} (see Table 3). This worsening of procurement cost cannot occur if a CPPA portfolio is optimized (for instance using MDP (6)), which suggests that PFRH is not effective at structuring dynamic portfolios. The results in Table 4 corresponding to PIRH show this to certainly be the case. Specifically, the PIRH optimality gap when using a median decision measure is 2% and 1.6% on the baseline instance and on the other instances sets, respectively, which is a significant improvement over FRH_{25}. This finding shows that portfolios constructed by PIRH can reduce procurement costs

Table 4: Procurement costs and optimality gaps for PFRH, PIRH-0, and PIRH policies.

Instance		PFRH	PIRH-0		PIRH	
			mean	median	mean	median
Baseline	Procurement cost (mln. USD)	590.9	589.0	597.0	576.0	575.2
	Optimality gap	4.8%	4.4%	5.8%	2.1%	2.0%
S1–S6	Procurement cost (mln. USD)	577.6	584.8	584.2	568.7	567.4
	Optimality gap	3.4%	4.7%	4.6%	1.8%	1.6%

relative to a rolling planning approach such as PFRH.

Next, we provide some insight into the superior performance of PIRH compared to PFRH. The conceptual advantage of PIRH highlighted in §5.2 is that the action distribution it uses to extract a non-anticipative decision encodes information about the impact of future uncertainty on the current decision. In contrast, PFRH does not have such information as it uses a forecast. A natural question is when the information encoded in the PIRH action distribution is useful to improve on PFRH. An answer to this question can be gleaned from the optimality gaps of PIRH-0 in Table 4, which are worse than PFRH and PIRH. Specifically, if the quality of the dual penalty is poor, as in PIRH-0, then relying on an action distribution can be worse than using a forecast. On the other hand, it is encouraging that the simple linear dual penalties that we use in PIRH are sufficient to make decision measures improve upon the PFRH decisions as well as lead to near optimal decisions and a dual bound. On our instances, there is significant value in maintaining CPPA portfolios and PIRH provides an effective methodology to compute and update such portfolios.

The performance differences between PIRH and PFRH also translate to contracting differences in their respective procurement policies. Table 5 reports the contract diversity, which we represent using the mean and maximum numbers of distinct CPPAs in the policy's portfolio over the stages in the planning horizon. Since each of these statistics depends on the sample path of uncertainty in the policy simulation, we report their respective averages over the sample paths. The table also contains the fraction of sample paths (as a percentage) that use CPPAs of a particular tenure under PFRH and PIRH. Both approaches have comparable average contract diversities but PIRH

Table 5: Contract diversity and usage of PFRH and PIRH policies on the baseline instance.

	Contract diversity		Percentage of sample paths using a CPPA tenure				
	Mean	Max	5	10	15	20	25
PFRH	1.89	2.42	6%	15%	43%	80%	96%
PIRH	1.84	3.10	36%	18%	38%	81%	98%

has a larger maximum. Their policies use CPPAs of all tenures with PIRH using 5-year CPPAs significantly more often than PFRH, possibly because shorter term CPPAs avoid locking in an unfavorable strike price for a long time horizon.

The benefits of the PIRH policy comes with a higher computational burden than FRH_m and PFRH. At a given stage and state, the time taken to compute a decision using FRH_m, PFRH, and PIRH are 0.01, 0.04, and 1.9 seconds, respectively. Most of this time for PIRH goes towards solving the dual optimization models on sample paths and the time required to extract the non-anticipative decision was negligible. Estimating the value of a policy, i.e., applying it over the entire planning horizon using 1000 Monte Carlo simulations, took 2 and 9 minutes respectively for FRH_m and PFRH, and 9 hours for FIRH. Estimating the value of the PIRH policy is more expensive than PFRH because the former policy solves at each stage math programs for each of the inner sample paths (which is 30 in our simulation) whereas the latter policy solves a single math program. The solution of PIRH math programs in the inner samples can be parallelized to substantially reduce this overhead. Finally, estimating the PIRH lower bound took about 1 minute on average.

6.6 Current and future relevance of CPPAs

The results in Tables 3 and 4 show that the PIRH procurement policy reduces expected cost on the baseline instance by 13.5% compared to the short-term policy, which translates to 78 mln. USD. Therefore, entering into CPPAs appears highly beneficial for a firm in an environment resembling our baseline instance. Since regulatory and market conditions evolve, we consider below the impact of some of these changes on the benefits of a CPPA.

Support schemes have spurred the increase in renewable energy generation but these schemes are being gradually phased out in several countries. We investigate if removing the PTC[7] would significantly affect how corporations meet an RPPT. Figure 4 displays the evolution of an "average" CPPA portfolio over the procurement horizon under the PIRH policy with and without a PTC, where the averaging occurs over the CPPA portfolios associated with the sample paths used in our policy simulation. For our baseline instance, which has a PTC expiring in year five, Figure 4(a) shows that power is purchased aggressively in the first five years via CPPAs to meet the RPPT due to lower strike prices. The PTC expiration results in the strike price of CPPAs to jump in year six, causing new CPPAs to become less desirable for a few years before the strike price decreases again due to improvement in technology – a feature captured by our strike price model (see Online Supplement E). This strike price trend is supported in Figure 4(a) by the decrease in CPPA power

[7] The presence of a PTC results in additional revenue to the generator and thus a lower strike price. See our strike price model in Online Supplement E for details.

Figure 4: Average PIRH renewable portfolio composition with and without a PTC.

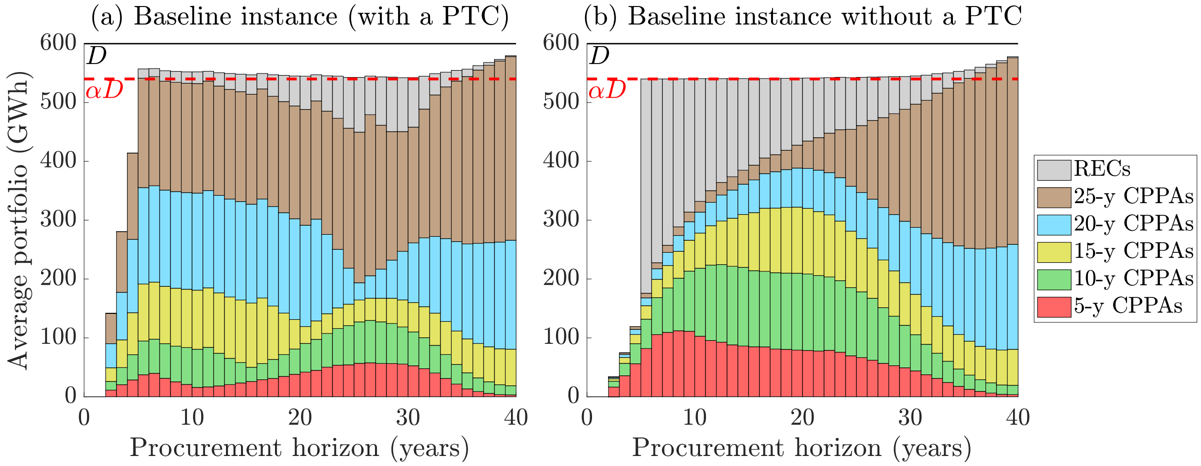

delivery from year 6 for some years and the subsequent increase in the CPPA power supply. With no PTC, procurement via CPPAs is more subdued at the beginning, with significant short-term REC purchases used to meet the RPPT as seen in Figure 4(b). The increase in the use of CPPAs in later years is driven by the decrease in the strike price due to technology improvements captured via a learning rate in our strike price model. The average contribution of CPPAs towards the RPPT in our baseline instance decreases from 93% with a PTC to 79% without a PTC. Moreover, the total procurement cost increases from 575.2 to 627.0 mln. USD when the PTC is removed, which equals a 9% increase. Thus, the presence of a PTC directly affects the use of CPPAs in the short- to medium-term but, regardless of a PTC, CPPAs are competitive procurement instruments in the long term.

Other important market trends pertain to power and REC prices, and CPPA availability. Market outlooks suggest that long-term power prices will decrease due to increasing penetration of renewable energy (Mills et al. 2017). Figure 5(a) shows the impact of reducing our calibrated long-term mean power price of 38 USD/MWh to 20 USD/MWh on the short-term and PIRH policy costs. As expected, the procurement costs of both policies decrease as the long-term mean power price drops. When the long-term mean power price falls below 25 USD/MWh, the performance of both policies become comparable. Since such a price drop is substantial, our results suggest that CPPAs will likely continue to play an important role in corporate procurement for several years to come.

In contrast to the power price, the average REC price can increase or decrease in the long-term due to regulatory changes (EPA 2018). Figure 5(b) is analogous to Figure 5(a) but considers increasing and decreasing the calibrated long-term mean of the REC price of 9.4 USD/MWh to 20 and 5 USD/MWh, respectively. When the long-term mean of the REC price increases, the procurement costs of the short-term policy is affected substantially while PIRH is stable across instances. This robustness to REC price variability is due to PIRH purchasing a considerable amount of CPPAs

Figure 5: Procurement costs in mln. USD for the short-term policy (dashed line), PIRH (continuous line), and FRH_{25} (dotted line).

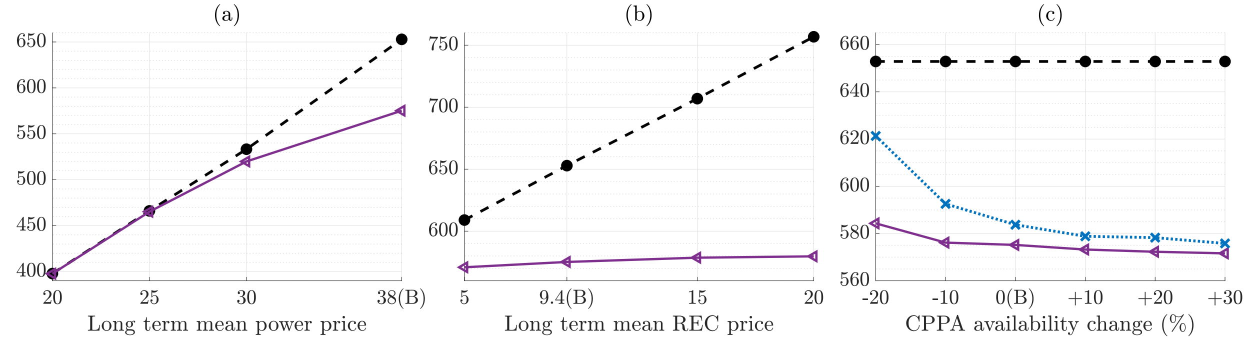

even when the REC prices are low, which insulates its procurement policy to REC price increases.

Finally, Figure 5(c) shows how procurement costs of the short-term, FRH_{25}, and PIRH policies change when varying the availability p_m of each contract $m \in \mathcal{M}$ from its base value to a value from -20% to $+30\%$. We included FRH_{25} in the comparison as it was the best single-contract policy we tested. Indeed, contract availability has no impact on the short-term policy. In contrast, but as intuitively expected, the cost of the single-contract-type heuristic FRH_{25} changes quite substantially as availability drops. Interestingly, we find that contract availability does not have a severe impact on the PIRH cost – the procurement cost increases by only 2% when decreasing contract availability from $+30\%$ to -20% relative to the baseline, which is a 50% availability cut for each contract type. This suggests that: (i) portfolios containing multiple CPPAs and optimized dynamically are rather stable under contract availability changes, and (ii) there is some level of substitutability between different subsets of CPPAs.

7. Conclusion

Motivated by the global trend in corporate energy procurement, we investigated the problem of meeting an RPPT using CPPAs in the presence of supply and price uncertainties. We analyzed a simplified model with a single CPPA and showed that it exposes a firm to price uncertainty as a result of volatile supply from the renewable energy generator. Nevertheless, signing a CPPA decreases procurement costs for companies with aggressive RPPTs when the correlations between power/REC prices and supply are not too negative, which is currently true in power markets. To inform the construction of a dynamic CPPA portfolio to meet and sustain an RPPT, as done in practice, we formulate an MDP, which is intractable. We design three forecast-based heuristics and a novel information-relaxation based technique with policies that approximate features of the idealized MDP policy. The latter technique also provides a dual bound for benchmarking policies. Through an extensive computational study on realistic procurement instances, we shed light on the cost of renewable power certification, the risks arising from supply uncertainty, the value of dynamic CPPA portfolios, and the effectiveness of our information-relaxation heuristic for constructing them.

References

Aflaki, S., S. Netessine. 2017. Strategic investment in renewable energy sources: The effect of supply intermittency. *Manufacturing & Service Operations Management* **19**(3) 489–507.

Al-Gwaiz, M., X. Chao, O. Q. Wu. 2016. Understanding how generation flexibility and renewable energy affect power market competition. *Manufacturing & Service Operations Management* **19**(1) 114–131.

Alizamir, S., F. de Véricourt, P. Sun. 2016. Efficient feed-in-tariff policies for renewable energy technologies. *Operations Research* **64**(1) 52–66.

Allon, G., J. Van Mieghem. 2010. Global dual sourcing: Tailored base-surge allocation to near- and offshore production. *Management Science* **56**(1) 110–124.

Andersen, L., M. Broadie. 2004. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science* **50**(9) 1222–1234.

Baker McKenzie. 2015. The rise of corporate PPAs - A new driver for renewables. Tech. rep.

Baker McKenzie. 2018. The rise of corporate PPAs 2.0. Tech. rep.

Bertsekas, D. P. 2005. Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control* **11**(4-5) 310–334.

Bertsekas, D. P. 2011. *Dynamic Programming and Optimal Control*. 3rd ed. Athena Scientific, Belmont, MA, USA.

Bloomberg. 2018. United States rates and bonds. URL http://pages.news/bloomberg3.

BNEF. 2018. Corporations already purchased record clean energy volumes in 2018, and it's not an anomaly. URL http://energy.financial/clean4. Bloomberg New Energy Finance (BNEF).

BNEF. 2018. Executing your corporate energy strategy. Tech. rep., Bloomberg New Energy Finance (BNEF).

BNEF. 2020. Corporate clean energy buying leapt 44% in 2019, sets new record. Tech. rep., Bloomberg New Energy Finance (BNEF).

Boyabatlı, O., P. R. Kleindorfer, S. R. Koontz. 2011. Integrating long-term and short-term contracting in beef supply chains. *Management Science* **57**(10) 1771–1787.

Brown, D. B., M. B. Haugh. 2017. Information relaxation bounds for infinite horizon Markov decision processes. *Operations Research* **65**(5) 1355–1379.

Brown, D. B., J. E. Smith. 2011. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science* **57**(10) 1752–1770.

Brown, D. B., J. E. Smith. 2014. Information relaxations, duality, and convex stochastic dynamic programs. *Operations Research* **62**(6) 1394–1415.

Brown, D. B., J. E. Smith, P. Sun. 2010. Information relaxations and duality in stochastic dynamic programs. *Operations Research* **58**(4) 785–801.

Bruck, M., P. Sandborn, N. Goudarzi. 2018. A levelized 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. Tech. rep., Carbon Disclosure Project (CDP), World Wide Fund for Nature (WWF), Coalition for Environmentally Responsible Economies (Ceres).

Chand, S., V. N. Hsu, S. Sethi. 2002. Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing & Service Operations Management* **4**(1) 25–43.

Chao, X., F. Y. Chen. 2005. An optimal production and shutdown strategy when a supplier offers an incentive program. *Manufacturing & Service Operations Management* **7**(2) 130–143.

Davies, K., G.M.J. Orrick, L. Taylor. 2018. Proxy generation ppas. Tech. rep., Microsoft, Orrick, and REsurety.

De Farias, D. P., B. Van Roy. 2003. The linear programming approach to approximate dynamic programming. *Operations Research* **51**(6) 850–865.

Denholm, P., R. Sioshansi. 2009. The value of compressed air energy storage with wind in transmission-constrained electric power systems. *Energy Policy* **37**(8) 3149–3158.

Desai, V. V., V. F. Farias, C. C. Moallemi. 2012. Pathwise optimization for optimal stopping problems. *Management Science* **58**(12) 2292–2308.

DOE. 2016. Renewable electricity production tax credit (ptc). URL http://production.credit/ptc1. Department of Energy (DOE).

Drake, D. F., P. R. Kleindorfer, L. N. Van Wassenhove. 2016. Technology choice and capacity portfolios under emissions regulation. *Production and Operations Management* **25**(6) 1006–1025.

EPA. 2018. Green power pricing. URL http://powers.supplies/green6. United States Environmental Protection Agency (EPA).

Gallego, Guillermo, Garrett Van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations research* **45**(1) 24–41.

Google. 2017. Environmental report: 2017 progress update. Tech. rep. URL http://visit.news/google45.

Haugh, M. B., L. Kogan. 2004. Pricing American options: A duality approach. *Operations Research* **52**(2) 258–270.

Haugh, M. B., O. R. Lacedelli. 2018. Information relaxation bounds for partially observed Markov decision processes. Working paper, Imperial College London.

Hu, S., G. C. Souza, M. E. Ferguson, W. Wang. 2015. Capacity investment in renewable energy technology with supply intermittency: Data granularity matters! *Manufacturing & Service Operations Management* **17**(4) 480–494.

İşlegen, Ö., S. Reichelstein. 2011. Carbon capture by fossil fuel power plants: An economic analysis. *Management Science* **57**(1) 21–39.

Jiang, D. R., W. B. Powell. 2015. Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS Journal on Computing* **27**(3) 525–543.

Kim, J. H., W. B. Powell. 2011. Optimal energy commitments with storage and intermittent supply. *Operations research* **59**(6) 1347–1360.

Kleindorfer, P. R., D. J. Wu. 2003. Integrating long-and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Science* **49**(11) 1597–1615.

Lai, G., F. Margot, N. Secomandi. 2010. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations research* **58**(3) 564–582.

Li, C. L., P. Kouvelis. 1999. Flexible and risk-sharing supply contracts under price uncertainty. *Management Science* **45**(10) 1378–1398.

Lin, Q., S. Nadarajah, N. Soheili. 2019. Revisiting approximate linear programming: Constraint-violation learning with applications to inventory control and energy storage. *Management Science* **66**(4) 1544–1562.

Longstaff, F. A., E. S. Schwartz. 2001. Valuing American options by simulation: A least-squares approach. *Review of Financial Studies* **14**(1) 113–147.

Lopuhaa, H. P., P. J. Rousseeuw. 1991. Breakdown points of affine equivariant estimators of multivariate location and covariance matrices. *The Annals of Statistics* **19**(1) 229–248.

Loukatou, A., S. Howell, P. Johnson, P. Duck. 2018. Stochastic wind speed modelling for estimation of expected wind power output. *Applied energy* **228** 1328–1340.

Lucia, J. J., E. S. Schwartz. 2002. Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Review of Derivatives Research* **5**(1) 5–50.

Martínez-de Albéniz, V., D. Simchi-Levi. 2005. A portfolio approach to procurement contracts. *Production and Operations Management* **14**(1) 90–114.

Meibom, P., H. V. Larsen, R. Barth, H. Brand, A. Tuohy, E. Ela. 2011. Advanced unit commitment strategies in the united states eastern interconnection. *Contract* **303** 275–3000.

Milligan, M., E. Ela, D. Lew, D. Corbus, Y. Wan, B. Hodge, B. Kirby. 2012. Operational analysis and methods for wind integration studies. *IEEE Transactions on Sustainable Energy* **3**(4) 612–619.

Mills, A., J. Seel, R. Wiser. 2017. As more solar and wind come onto the grid, prices go down but new questions come up. URL http://costs.solar/wind5.

Nadarajah, S., N. Secomandi. 2018. Merchant energy trading in a network. *Operations Research* **66**(5) 1304–1320.

NREL. 2010. Cost and performance assumptions for modeling electricity generation technologies. Tech. Rep. NREL/SR-6A20-48595, National Renewable Energy Laboratory (NREL).

NREL. 2017. System Advisor Model (SAM). URL https://sam.nrel.gov/. National Renewable Energy Laboratory (NREL).

Pandžić, H., Y. Wang, T. Qiu, Y. Dvorkin, D. S. Kirschen. 2015. Near-optimal method for siting and sizing of distributed storage in a transmission network. *IEEE Transactions on Power Systems* **30**(5) 2288–2300.

Pereira, M. V. F., L. M. V. G. Pinto. 1991. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* **52**(1-3) 359–375.

Powell, W. B. 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. 2nd ed. John Wiley & Sons, Hoboken, NJ, USA.

PWC. 2016. Corporate renewable energy procurement survey insights. Tech. rep., PricewaterhouseCoopers (PWC).

Ritzenhofen, I., J. R. Birge, S. Spinler. 2016. The structural impact of renewable portfolio standards and feed-in tariffs on electricity markets. *European Journal of Operational Research* **255**(1) 224–242.

RMI. 2016. Amazon and dominion virginia power reach breakthrough renewable energy agreement. Tech. rep., Rocky Mountain Institute (RMI). URL https://rmi.org/amazon-dominion-virginia-power-reach-breakthrough-renewable-energy-agreement/.

Salas, D. F., W. B. Powell. 2017. Benchmarking a scalable approximate dynamic programming algorithm for stochastic control of grid-level energy storage. *INFORMS Journal on Computing* **30**(1) 106–123.

Secomandi, N. 2015. Merchant commodity storage practice revisited. *Operations Research* **63**(5) 1131–1143.

Secomandi, N., S. Kekre. 2014. Optimal energy procurement in spot and forward markets. *Manufacturing & Service Operations Management* **16**(2) 270–282.

Shapiro, A. 2011. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research* **209**(1) 63–72.

Singh, S. P., A. A. Scheller-Wolf. 2017. That's not fair: Tariff structures for electricity markets with rooftop solar. Working paper, Carnegie Mellon University.

Sunar, N., J. R. Birge. 2019. Strategic commitment to a production schedule with uncertain supply and demand: Renewable energy in day-ahead electricity markets. *Management Science* **65**(2) 714–734.

Tomlin, B., Y. Wang. 2005. On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. *Manufacturing & Service Operations Management* **7**(1) 37–57.

Tsitsiklis, J. N., B. Van Roy. 2001. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* **12**(4) 694–703.

Veeraraghavan, S., A. A. Scheller-Wolf. 2008. Now or later: A simple policy for effective dual sourcing in capacitated systems. *Operations Research* **56**(4) 850–864.

WBCSD. 2018. Innovation in power purchase agreement structures. Tech. rep., World Business Council For Sustainable Development (WBCSD).

Webb, E., O. Q. Wu, K. Cattani. 2017. Mind the gap: Coordinating energy efficiency and demand response. Working paper, Indiana University.

Weber, C., P. Meibom, R. Barth, H. Brand. 2009. WILMAR: A stochastic programming tool to analyze the large-scale integration of wind energy. *Optimization in the energy industry*. Springer, 437–458.

Weron, R. 2014. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* **30**(4) 1030–1081.

Wiser, R., M. Bolinger. 2017. 2016 wind technologies market report. Tech. rep.

Woo, C.-K., I. Horowitz, J. Moore, A. Pacheco. 2011. The impact of wind generation on the electricity spot-market price level and variance: The texas experience. *Energy Policy* **39**(7) 3939–3944.

Wu, O. Q., V. Babich. 2012. Unit-contingent power purchase agreement and asymmetric information about plant outage. *Manufacturing & Service Operations Management* **14**(2) 245–261.

Wu, O. Q., R. Kapuscinski. 2013. Curtailing intermittent generation in electrical systems. *Manufacturing & Service Operations Management* **15**(4) 578–595.

Wu, O. Q., D. D. Wang, Z. Qin. 2012. Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management* **14**(3) 455–471.

Ye, F., H. Zhu, E. Zhou. 2018. Weakly coupled dynamic program: Information and Lagrangian relaxations. *IEEE Transactions on Automatic Control* **63**(3) 698–713.

Zeng, Y., D. Klabjan, J. Arinez. 2015. Distributed solar renewable generation: Option contracts with renewable energy credit uncertainty. *Energy Economics* **48** 295–305.

Zhou, Y., A. Scheller-Wolf, N. Secomandi, S. Smith. 2019. Managing wind-based electricity generation in the presence of storage and transmission capacity. *Production and Operations Management* **28**(4) 970–989.

Meeting Corporate Renewable Power Targets
(Online Supplement)

A. Proofs

Proof of Proposition 2.1. The derivative of the expected cost with respect to z is $\mathrm{PAP}(K,\Theta) - \mathbb{E}[R_2\eta\mathbf{1}(\eta z \leqslant \alpha D)]$, which is a non-decreasing function of z since $\mathrm{PAP}(K,\Theta)$ is a constant and $\mathbb{E}[R_2\eta\mathbf{1}(\eta z \leqslant \alpha D)]$ weakly decreases in z. Thus, the expected cost is a convex function of z and a comparison of slopes suffices to characterize the optimal procureme nt quantity.

We first consider the conditions for $z^*(\alpha,K,\Theta) = 0$. Since η_i is a random variable with support on the interval $[0,H]$, the random variable η has support on the interval $[0,2H]$. This implies that there exists a small enough value for z, denoted $\bar{z}$, such that for $0 \leqslant z \leqslant \bar{z}$ we have $\mathrm{PAP}(K,\Theta) - \mathbb{E}[R_2\eta\mathbf{1}(\eta z \leqslant \alpha D)] = \mathrm{PAP}(K,\Theta) - \mathbb{E}[R_2\eta]$. Therefore, if $\mathrm{PAP}(K,\Theta) - \mathbb{E}[R_2\eta] \geqslant 0$, which is the same as $\mathrm{PRS}(K,\Theta) \leqslant 0$, then $z^*(\alpha,K,\Theta)$ equals 0. Instead, if $\mathrm{PRS}(K,\Theta) > 0$ holds, then there is a $\infty > y^*(K,\Theta) \geqslant 0$ such that $\mathrm{PAP}(K,\Theta) - \mathbb{E}[R_2\eta\mathbf{1}(\eta \leqslant y^*(K,\Theta))] = 0$ since $\mathbb{E}[R_2\eta\mathbf{1}(\eta \leqslant y)]$ is an increasing function of y that is bounded below by 0 for $y = 0$ and upper bounded by $\mathbb{E}[R_2\eta]$ for y equal to $2H$. This implies that $z^*(\alpha,K,\Theta) = \alpha D/y^*(K,\Theta)$. We pick the smallest such $y^*(K,\Theta)$ in defining the optimal procurement quantity.

We can write $\mathrm{PAP}(K,\Theta)$ as follows

$$\begin{aligned}
\mathrm{PAP}(K,\Theta) &= K\mathbb{E}[\eta] - \sum_{i\in\{1,2\}} \mathbb{E}[\eta_i P_i] \\
&= K\mathbb{E}[\eta] - \sum_{i\in\{1,2\}} \left(\mathbb{E}[\eta_i]\mathbb{E}[P_i] + \mathrm{Cov}(P_i,\eta_i)\right) \\
&= K\mathbb{E}[\eta] - \sum_{i\in\{1,2\}} \left(\mathbb{E}[\eta_i]\mathbb{E}[P_i] + \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i}\right).
\end{aligned}$$

It is apparent from this rewriting that $\mathrm{PAP}(K,\Theta)$ is a decreasing function of each ρ_{P_i,η_i} if $\sigma_{\eta_i} > 0$ and $\sigma_{P_i} > 0$. In addition, $\mathbb{E}[R_2\eta\mathbf{1}(\eta \leqslant y)]$ does not depend on ρ_{P_i,η_i}. It follows that $y^*(\rho_{P_i,\eta_i})$ is non-increasing in ρ_{P_i,η_i} and $z^*(\rho_{P_i,\eta_i})$ is weakly increasing in this parameter.

(ii) From our rewriting of $\mathrm{PAP}(K,\Theta)$ in part (i) above, it is clear that this quantity is non-decreasing in σ_{P_i} when $\rho_{P_i,\eta_i} \leqslant 0$. Since $\mathbb{E}[R_2\eta\mathbf{1}(\eta \leqslant y)]$ does not depend on σ_{P_i}, $y^*(\sigma_{P_i},\rho_{P_i,\eta_i})$ is non-decreasing in σ_{P_i}, which implies that $z^*(\sigma_{P_i},\rho_{P_i,\eta_i})$ weakly increases in this parameter. □

Proof of Proposition 2.2. (i) If $\mathrm{PRS}(K,\Theta) \leqslant 0$, by Proposition 2.1, we have $z^*(\alpha,K,\Theta) = 0$ and $C^*(\alpha,K,\Theta) = \sum_{i\in\{1,2\}} D_i\mathbb{E}[P_i] + \alpha\mathbb{E}[R_2]D$, which is linear a function of α. Instead, if $\mathrm{PRS}(K,\Theta) > 0$, then $z^*(\alpha,K,\Theta) = \alpha D/y^*(K,\Theta)$, where $y^*(K,\Theta)$ does not depend on α. In this case,

$$C^*(\alpha,K,\Theta) = \sum_{i\in\{1,2\}} D_i\mathbb{E}[P_i] + (\alpha D/y^*(K,\Theta))\left(K\mathbb{E}[\eta] - \sum_{i\in\{1,2\}} \mathbb{E}[\eta_i P_i] + \mathbb{E}[R_2(y^*(K,\Theta)-\eta)_+]\right),$$

which is again a linear function of α.

(ii) The expected cost can be written as follows

$$
\begin{aligned}
\mathbb{E}[\widetilde{C}(z,w;\alpha,K,\Theta)] &= K\mathbb{E}[\eta]z + \sum_{i\in\{1,2\}} (D_i\mathbb{E}[P_i] - \mathbb{E}[\eta_i P_i]z) + \mathbb{E}[R_2(\alpha D - \eta z)_+], \\
&= K\mathbb{E}[\eta]z + \sum_{i\in\{1,2\}} (D_i\mathbb{E}[P_i] - \mathbb{E}[\eta_i]\mathbb{E}[P_i]z - \mathrm{Cov}(\eta_i, P_i)z) + \mathbb{E}[R_2(\alpha D - \eta z)_+] \\
&= K\mathbb{E}[\eta]z + \sum_{i\in\{1,2\}} (D_i\mathbb{E}[P_i] - \mathbb{E}[\eta_i]\mathbb{E}[P_i]z - \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i}z) + \mathbb{E}[R_2(\alpha D - \eta z)_+].
\end{aligned}
$$

This shows that expected cost function weakly decreases in ρ_{P_i,η_i} for any z. In particular, consider correlations $\rho^1_{P_i,\eta_i}$ and $\rho^2_{P_i,\eta_i}$ such that $\rho^1_{P_i,\eta_i} > \rho^2_{P_i,\eta_i}$ for all $i \in \{1,2\}$. Then we have

$$
\begin{aligned}
C^*(\rho^1_{P_i,\eta_i}) &= \min_{z\geqslant 0} \mathbb{E}[\widetilde{C}(z,w;\rho^1_{\eta_1,P_1},\rho^1_{\eta_2,P_2})] \\
&= \mathbb{E}[\widetilde{C}(z^*(\rho^1_{\eta_1,P_1},\rho^1_{\eta_2,P_2}),w;\rho^1_{\eta_1,P_1},\rho^1_{\eta_2,P_2})] \\
&\geqslant \mathbb{E}[\widetilde{C}(z^*(\rho^1_{\eta_1,P_1},\rho^1_{\eta_2,P_2}),w;\rho^2_{\eta_1,P_1},\rho^2_{\eta_2,P_2})] \\
&\geqslant \min_{z\geqslant 0} \mathbb{E}[\widetilde{C}(z,w;\rho^2_{\eta_1,P_1},\rho^2_{\eta_2,P_2})] \\
&= C^*(\rho^2_{\eta_1,P_1},\rho^2_{\eta_2,P_2}).
\end{aligned}
$$

(iii) Similarly, for negative (positive) ρ_{P_i,η_i} for all $i \in \{1,2\}$, the expected cost function weakly increases (decreases) with σ_{P_i} for each z and thus by analogous reasoning to part (i), it follows that the optimal cost increases (decreases) with σ_{P_i}. □

***Proof of Proposition** 2.3.* (i) If $z^*(\sigma_\eta = 0) = 0$ then $\mathrm{PRS}(K,\sigma_\eta = 0) \leqslant 0$. Otherwise, $\mathrm{PRS}(K,\sigma_\eta = 0) = \mathbb{E}[R_T\eta\mathbf{1}(\eta \leqslant \infty)] - \mathrm{PAP}(K,\sigma_\eta = 0) > 0$ and $0 \leqslant y^*(K,\sigma_\eta = 0) < \infty$, implying that $z^*(K,\sigma_\eta = 0) > 0$. Therefore, we can write

$$
\begin{aligned}
\mathrm{PRS}(K,\sigma_\eta = \sigma) &= \mathbb{E}[R_2\eta] - \mathrm{PAP}(K,\sigma_\eta = \sigma) \\
&= \mathbb{E}[R_2]\mathbb{E}[\eta] + \rho_{R_2\eta}\sigma_{R_2}\sigma_\eta - K\mathbb{E}[\eta] + \sum_{i\in\{1,2\}} (\mathbb{E}[\eta_i]\mathbb{E}[P_i] + \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i}) \\
&= \mathrm{PRS}(K,\sigma_\eta = 0) + \rho_{R_2\eta}\sigma_{R_2}\sigma_\eta + \sum_{i\in\{1,2\}} \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i} \\
&\leqslant \rho_{R_2\eta}\sigma_{R_2}\sigma_\eta + \sum_{i\in\{1,2\}} \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i} \\
&\leqslant 0,
\end{aligned}
$$

where the first and second inequalities follows from $\mathrm{PRS}(K,\sigma_\eta = 0) \leqslant 0$ and negative correlations, respectively. Therefore, $z^*(\sigma_\eta = \sigma) = 0$ by Part (i) of Proposition 2.1.

(i) If $z^*(\sigma_\eta = 0) > 0$ then $\mathrm{PRS}(K,\sigma_\eta = 0) > 0$ and $y^*(K,\sigma_\eta = 0) = \mathbb{E}[\eta]$. Suppose $\mathrm{LRV}(\alpha D/\mathbb{E}[\eta]) \geqslant \mathrm{PRS}(K,\sigma_\eta = \sigma)$. In this case, we have

$$
\begin{aligned}
\mathrm{PAP}(K,\sigma_\eta = \sigma) - \mathbb{E}[R_2\eta\mathbf{1}(\eta \leqslant \mathbb{E}[\eta])] &= \mathrm{PAP}(K,\sigma_\eta = \sigma) - \mathbb{E}[R_2\eta] + \mathbb{E}[R_2\eta\mathbf{1}(\eta > \mathbb{E}[\eta])] \\
&= -\mathrm{PRS}(K,\sigma_\eta = \sigma) + \mathbb{E}[R_2\eta\mathbf{1}(\eta > \mathbb{E}[\eta])]
\end{aligned}
$$

$$\geqslant 0,$$

where the first equality employs the relationship $\mathbb{E}[R_2\eta] = \mathbb{E}[R_2\eta\mathbf{1}(\eta \leqslant \mathbb{E}[\eta])] + \mathbb{E}[R_2\eta\mathbf{1}(\eta > \mathbb{E}[\eta])]$, the second equality using the definition of $\mathrm{PRS}(K, \sigma_\eta = \sigma)$, and the inequality from $\mathrm{LRV}(\alpha D/\mathbb{E}[\eta]) \geqslant \mathrm{PRS}(K, \sigma_\eta = \sigma)$ which is equivalent to $\mathbb{E}[R_2\eta\mathbf{1}(\eta > \mathbb{E}[\eta])] \geqslant \mathrm{PRS}(K, \sigma_\eta = \sigma)$. Therefore, $y^*(K, \sigma_\eta = \sigma) \geqslant \mathbb{E}[\eta] = y^*(K, \sigma_\eta = 0)$, which implies that $z^*(\sigma_\eta = \sigma) \leqslant z^*(\sigma_\eta = 0)$. Using analogous reasoning it is easy to verify that $z^*(\sigma_\eta = \sigma) \geqslant z^*(\sigma_\eta = 0)$ if $\mathrm{LRV}(\alpha D/\mathbb{E}[\eta]) \leqslant \mathrm{PRS}(K, \sigma_\eta = \sigma)$. □

***Proof of Proposition** 2.4.* The difference $C^*(\sigma_\eta = \sigma) - C^*(\sigma_\eta = 0)$

$$\begin{aligned}
&= C(z^*(K, \sigma_\eta = \sigma); \sigma_\eta = \sigma) - C^*(\sigma_\eta = 0) \\
&\geqslant C(z^*(K, \sigma_\eta = \sigma); \sigma_\eta = \sigma) - C(z^*(K, \sigma_\eta = \sigma); \sigma_\eta = 0) \\
&= C(\alpha D/y^*(K, \sigma_\eta = \sigma); \sigma_\eta = \sigma) - C(\alpha D/y^*(K, \sigma_\eta = \sigma); \sigma_\eta = 0) \\
&= -\sum_{i\in\{1,2\}} \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i}\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)} + \mathbb{E}\left[R_2\left(\alpha D - \eta\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)}\right)_+\right] \\
&\quad - \mathbb{E}[R_2]\left(\alpha D - \mathbb{E}[\eta]\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)}\right)_+ \\
&\geqslant -\sum_{i\in\{1,2\}} \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i}\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)} + \left(\mathbb{E}[R_2]\alpha D - \mathbb{E}[R_2\eta]\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)}\right)_+ \\
&\quad - \mathbb{E}[R_2]\left(\alpha D - \mathbb{E}[\eta]\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)}\right)_+ \\
&= -\sum_{i\in\{1,2\}} \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i}\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)} + \left(\mathbb{E}[R_2]\alpha D - \mathbb{E}[R_2]\mathbb{E}[\eta]\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)}\right. \\
&\quad \left. -\rho_{R_2\eta}\sigma_{R_2}\sigma\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)}\right)_+ - \mathbb{E}[R_2]\left(\alpha D - \mathbb{E}[\eta]\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)}\right)_+ \\
&\geqslant -\sum_{i\in\{1,2\}} \rho_{P_i,\eta_i}\sigma_{\eta_i}\sigma_{P_i}\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)} - \rho_{R_2\eta}\sigma_{R_2}\sigma\mathbf{1}(y^*(K, \sigma_\eta = \sigma) \geqslant \mathbb{E}[\eta])\frac{\alpha D}{y^*(K, \sigma_\eta = \sigma)} \\
&\geqslant 0.
\end{aligned}$$

The first inequality follows from $z^*(K, \sigma_\eta = \sigma)$ being a feasible solution to $C(z; \sigma_\eta = 0)$, the second inequality is a consequence of applying Jensen's inequality, the third inequality follows from noticing that $\mathbb{E}[R_2](\alpha D - \mathbb{E}[\eta]\alpha D/y^*(K, \sigma_\eta = \sigma)) \geqslant 0$ when $y^*(K, \sigma_\eta = \sigma) \geqslant \mathbb{E}[\eta]$, and the last inequality follows from the assumption on negative correlations. □

***Proof of Proposition** 3.1.* Consider the Bellman recursion associated with MDP (6):

$$\begin{aligned}
V_{I-1}(x_{I-1}, w_{I-1}) &= c_{I-1}(x_{I-1}, w_{I-1}),\ \forall (x_{I-1}, w_{I-1}) \in \mathcal{X}_{I-1} \times \mathcal{W}_{I-1}, \\
V_i(x_i, w_i) &= \min_{z_i \in \mathcal{Z}_i(a_i)} \left\{c_i(x_i, w_i, z_i) + \gamma_{\text{A}}\mathbb{E}_i\left[V_{i+1}\left(f_i(x_i, z_i), w_{i+1}\right)\right]\right\} \\
&\qquad \forall (i, x_i, w_i) \in \{0, \ldots, I-2\} \times \mathcal{X}_i \times \mathcal{W}_i.
\end{aligned}$$

At a given stage $i \in \{0, \ldots, I-2\}$, let $\mathcal{C}_i := \{(x_i, z_i) | x_i \in \mathcal{X}_i, z_i \in \mathcal{Z}_i(a_i)\}$ denote the set of actions and states at i, and also define

$$G_i(x_i, w_i, z_i) := c_i(x_i, w_i, z_i) + \gamma_{\text{A}} \mathbb{E}_i\big[V_{i+1}(f(x_i, z_i), w_{i+1})\big].$$

To prove the convexity of MDP (6), we show first that $G_{I-2}(x_{I-2}, w_{I-2}, z_{I-2})$ is jointly convex in the state x_{I-2} and action z_{I-2}, i.e. in set $\mathcal{C}_{I-2}$, at stage $I-2$. For a given w_{I-2}, we have:

$$V_{I-2}(x_{I-2}, w_{I-2}) = \min_{z_{I-2} \in \mathcal{Z}_{I-2}(a_{I-2})} \Big\{ c_{I-2}(x_{I-2}, w_{I-2}, z_{I-2}) + \gamma_{\text{A}} \mathbb{E}_{I-2}\big[V_{I-1}(f(x_{I-2}, z_{I-2}), w_{I-1})\big] \Big\}.$$

We know that $V_{I-2}(x_{I-2}, w_{I-2})$ and hence $\mathbb{E}_{I-2}\big[V_{I-1}(f(x_{I-2}, z_{I-2}), w_{I-1})\big]$ are finite quantities since by assumption $V_{I-2}(x_{I-2}, w_{I-2})$ is bounded. Moreover, $c_{I-2}(x_{I-2}, w_{I-2}, z_{I-2})$ is piece-wise linear and convex in x_{I-2} from the definition in Equation (4). Similarly, the terminal value function $V_{I-1}(x_{I-1}, w_{I-1}) = c_{I-1}(x_{I-1}, w_{I-1})$ is convex in x_{I-1} for a given w_{I-1} based on the definition in Equation (5). Thus, the continuation function $\gamma_{\text{A}} \mathbb{E}_{I-2}[V_{I-1}(f(x_{I-2}, z_{I-2}), w_{I-1})]$ is convex in x_{I-2}. In addition, the set $\mathcal{C}_{I-2}$ is convex since both $\mathcal{X}_{I-2}$ and $\mathcal{Z}_{I-2}(a_{I-2})$ are convex. As a consequence, for a realization of w_{I-2}, the function $G_{I-2}(z_{I-2}, x_{I-2}, w_{I-2})$ is convex in the set $\mathcal{C}_{I-2}$. The convexity of $V_{I-2}(x_{I-2}, w_{I-2})$ as a function of x_{I-2} now follows from Proposition B-4 in Heyman and Sobel (1982).

Suppose $V_{i+1}(\cdot, w_{i+1})$ is convex in x_{i+1}. Following the same steps above, we can show that $G_i(x_i, w_i, z_i)$ is convex in $\mathcal{C}_i$ and consequently $V_i(\cdot, w_i)$ is convex in x_i. Hence, the result follows from mathematical induction. □

***Proof of Proposition** 5.1.* We first prove the feasibility of $\mathcal{H}_i(z_i(W_i))$ under the two conditions of Proposition 5.1.

1. Since $\mathcal{Z}_i(a_i)$ is convex and bounded, for any $W_i | w_i$, the optimal decision $z_i(W_i)$ at stage i is finite. Moreover, the average of the actions for the random variable W_i given w_i, $\mathbb{E}_i[z_i(W_i)]$, belongs to the set $\mathcal{Z}_i(a_i)$; thus, $\mathbb{E}_i[z_i(W_i)]$ is feasible.

2. $\mathcal{H}_i(z_i(W_i))$ is feasible since $\bar{z}_i$ is feasible and finite based on the statement of the condition and boundedness of the set $\mathcal{Z}_i(a_i)$.

Next, assuming $\mathcal{H}_i$ satisfies one of these conditions, we establish its optimality when using ideal dual penalties in math program (11). Specifically, we show that the decision at every sample path W_i is optimal. Our proof relies on the stochastic dynamic programming (SDP) reformulation of (11) at stage i, which is

$$V^{\text{IR}}_{I-1}(x_{I-1}; W_{I-1}) = c^{\text{IR}}_{I-1}(x_{I-1}, W_{I-1}) - q_{I-1}\big(x_{I-1}, W_{I-1}\big), \;\; \forall x_{I-1} \in \mathcal{X}_{I-1}, \tag{14a}$$

$$\begin{aligned} V^{\text{IR}}_j(x_j; W_j) = \min_{z_j \in \mathcal{Z}_j(a_j)} \Big\{ & c^{\text{IR}}_j(x_j, W_j, z_j) - q_j\,(x_j, W_j, z_j) \\ & + \gamma_{\text{A}} V^{\text{IR}}_{j+1}\big(f_j(x_j, z_j)\big) \Big\}, \;\; \forall j \in \mathcal{I}_i, \forall x_j \in \mathcal{X}_i. \end{aligned} \tag{14b}$$

At stage $I-1$, it is true that for each $x_{I-1} \in \mathcal{X}_{I-1}$:

$$V^{\text{IR}}_{I-1}(x_{I-1}; W_{I-1}) = c^{\text{IR}}_{I-1}(x_{I-1}, W_{I-1}) - q_{I-1}\big(x_{I-1}, W_{I-1}\big) = c_{I-1}(x_{I-1}, w_{I-1}) = V_{I-1}(x_{I-1}, w_{I-1}).$$

By backward induction, we assume that $V_{j+1}^{\mathrm{IR}}(x_{j+1}; W_{j+1})$ in equation (11) equals $V_{j+1}(x_{j+1}, w_{j+1})$ in equation (6) and prove the equality for stage j. The following equation represents the stage-j step of SDP (14) when the ideal dual penalty (12) is used:

$$V_j^{\mathrm{IR}}(x_j; W_j) = \min_{z_j \in \mathcal{Z}_j(a_j)} \Big\{ c_j^{\mathrm{IR}}(x_j, W_j, z_j) - \gamma_{\mathrm{A}} V_{j+1}\left(f_j(x_j, z_j), w_{j+1}\right) + \gamma_{\mathrm{A}} \mathbb{E}_j\big[V_{j+1}\big(f_j(x_j, z_j), w_{j+1}\big)\big]$$
$$+ \gamma_{\mathrm{A}} V_{j+1}^{\mathrm{IR}}\left(f_j(x_j, z_j)\right) - c_j^{\mathrm{IR}}(x_j, W_j, z_j) + c_j(x_j, w_j, z_j)\Big\}, \ \ \forall x_j \in \mathcal{X}_j. \tag{15a}$$
$$= \min_{z_j \in \mathcal{Z}_j(a_j)} \Big\{ c_j(x_j, w_j, z_j) + \gamma_{\mathrm{A}} \mathbb{E}_j\big[V_{j+1}\big(f_j(x_j, z_j), w_{j+1}\big)\big]\Big\}, \ \ \forall x_j \in \mathcal{X}_j. \tag{15b}$$
$$= V_j(x_j, w_j), \ \ \forall x_j \in \mathcal{X}_j, \tag{15c}$$

where (15b) follows from the induction hypothesis, and (15c) from the definition of the SDP associated with MDP (6). The relation, $V_j^{\mathrm{IR}}(x_j; W_j) = V_j(x_j, w_j)$, thus holds at the generic stage $j \in \mathcal{I}_i$ for the principle of mathematical induction. The optimality of the action for every W_i is immediate from this equality. Since $z_i(W_i)$ is equal to the optimal decision for any sample path, a feasible decision measure $\mathcal{H}_i$ that satisfies conditions 1 and 2 leads to an optimal solution to MDP (6). $\square$

B. Non-convexity of the MDP value function

We show that the value function $V_i(\cdot, w_i)$ of MDP (6) is non-convex in the endogenous state x_i by using a simple counter-example with two stages (i.e., 0 and 1) and no intermediate inner-stages. A CPPA can be entered at stage 0 with delivery during stage 1, and an RPPT of $\alpha = 0.8$ has to be fulfilled during stage 1. Any RPPT shortfall at the end of stage 1 can be procured via short-term RECs at price R_2. We assume that the contract is available at stage 0 (i.e., $a_0 = 1$), the renewable power supply during the contract delivery period is known and equal to $\eta_1 = 1$, and that minimum and maximum procurement quantities are $z^{\min} = 6$ MW and $z^{\max} = 10$ MW, respectively. Thus, a feasible action z_0 belongs to the action set $\mathcal{Z}_0(a_0) = \{0\} \cup [6, 10] \subset \mathbb{R}_+$. Moreover, we assume there is no discounting, demand equals $D_0 = D_1 = 10$ MWh, power and REC prices have starting values $P_0 = R_0 = 10$ USD/MWh and are martingales (i.e., $\mathbb{E}_i[P_j] = P_i$ and $\mathbb{E}_i[R_j] = R_i$ for $j \geqslant i$), and the strike price K_0 is equal to 11 USD/MWh. Proceeding backward, the terminal value function (stage 1) is convex in the state x_1 as it is equal to

$$\begin{aligned} V_1(x_1, w_1) &= P_1 D_1 + \mathbb{E}_1[R_2](\alpha D_1 - \eta_1 x_1)_+ \\ &= P_1 D_1 + R_1(\alpha D_1 - x_1)_+, \end{aligned}$$

which has the slope $-R_1$ for $x_1 \in [0, \alpha D_1]$, and is a constant (i.e., zero slope) equal to $P_1 D_1$ if $x_1 \geqslant \alpha D_1$. Instead, the stage 0 value function is given by

$$\begin{aligned} V_0(x_0, w_0) &= \min_{z_0 \in \mathcal{Z}_0(a_0)} \Big\{ \mathbb{E}_0[(K_0 - P_1)] z_0 + P_0 D_0 + \mathbb{E}_0\big[P_1 D_1 + R_1(\alpha D_1 - x_0 - z_0)_+\big]\Big\} \\ &= \min_{z_0 \in \mathcal{Z}_0(a_0)} \big\{ (K_0 - P_0) z_0 + P_0 D_0 + P_0 D_1 + R_0(\alpha D_1 - x_0 - z_0)_+ \big\}, \end{aligned}$$

which is non-convex in x_0 as illustrated in Figure 6. In particular, Figure 6(a) displays the value function $V_1(\cdot, w_1)$ as a function of x_1 given $P_1 = R_1 = 10$, and Figure 6(b) shows $V_0(\cdot, w_0)$ and z_0^* as functions of x_0 for w_0, which is known. The value function $V_0(\cdot, w_0)$ takes a constant value in the interval that is roughly $[2, 7.5]$. In this interval, it is indeed optimal to enter into a CPPA of size $z^{\min} = 6$, even if this decision causes over-procurement, because the CPPA is cheaper than the expected short-term option. Instead, when $x_0 > 7.5$, the optimal action jumps from $z_0^* = z^{\min} = 6$ to $z_0^* = 0$, i.e., it becomes optimal to rely only on the short-term option.

Figure 6: Stage 1 and stage 0 MDP value functions.

(a) $V_1(\cdot, w_1)$ with $P_1 = R_1 = 10$. (b) $V_0(\cdot, w_0)$ and z_0^*.

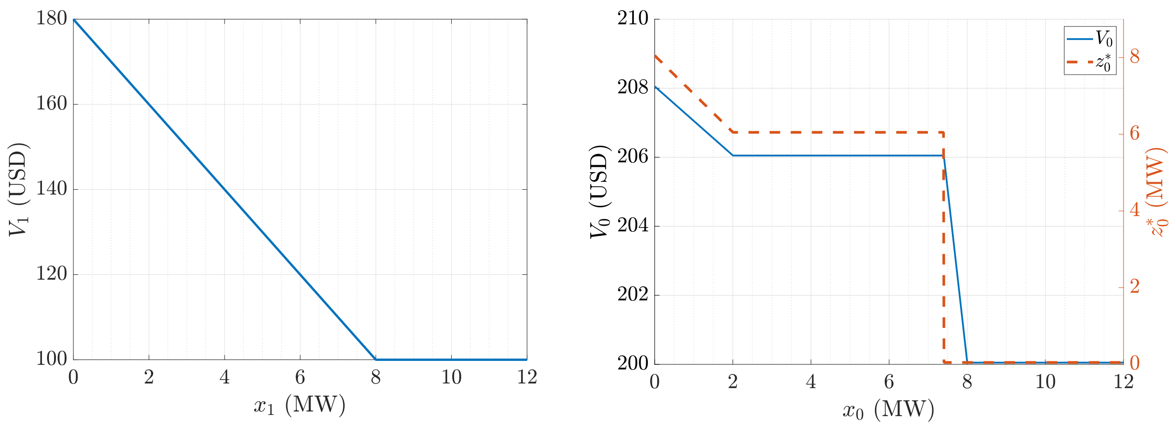

C. Finiteness of the MDP value function

In this section we introduce two sets of sufficient conditions under which the MDP value function (6) is finite, which is one of the assumptions needed in Proposition 3.1 to prove convexity of the value function. Specifically, the first set of conditions refers to the power and REC price processes having bounded support and is discussed in Proposition C.1. The second set of conditions is presented in Proposition C.2 and is less strict, only requiring expectations of price processes to be finite.

Proposition C.1. *If $z_m^{max} < \infty$ for each $m \in \mathcal{M}$ and prices $P_{i,t}$ and R_i have bounded support, respectively, for all $(i,t) \in \mathcal{I} \times \mathcal{T}$ and for all $i \in \mathcal{I} \cup \{I\}$, then $|V_i(\cdot,\cdot)| < \infty$ for each $i \in \mathcal{I}$.*

***Proof of Proposition** C.1.* For each stage and interval $(i,t) \in \mathcal{I} \times \mathcal{T}$, we know that the power supply uncertainty have bounded support $\eta_{i,t} \in [0, H]$, and that the same is true for $P_{i,t}$ and R_i by assumption. Moreover, since $z_m^{\max} < \infty$ for each $m \in \mathcal{M}$, we have that $\mathcal{Z}_i(a_i)$ is a bounded set at each stage $i \in \mathcal{I}$. As a consequence, there exists $N \in \mathbb{R}_+$ such that $|c_i(x_i, w_i, z_i)| \leqslant N$ for each $(i, x_i, w_i, z_i) \in \{0, \ldots, I-2\} \times \mathcal{X}_i \times \mathcal{W}_i \times \mathcal{Z}_i(a_i)$ and $|c_{I-1}(x_{I-1}, w_{I-1})| \leqslant N$ for each $(x_{I-1}, w_{I-1}) \in \mathcal{X}_{I-1} \times \mathcal{W}_{I-1}$. It follows immediately that $|V_i(\cdot,\cdot)| \leqslant N(I-i+2) \leqslant N \cdot I < \infty$ for each $i \in \mathcal{I}$. □

Proposition C.2. *If $z_m^{max} < \infty$ for each $m \in \mathcal{M}$ and price expectations are finite, i.e., $\mathbb{E}_i[|P_{j,t}|] < \infty$ and $\mathbb{E}_i[R_{j+1}] < \infty$ for each $(i,j,t) \in \mathcal{I} \times \{i, \ldots, I-1\} \times \mathcal{T}$, then $|V_i(\cdot,\cdot)| < \infty$ for $i \in \mathcal{I}$.*

Proof of Proposition C.2. We start by showing that $V_i(x_i, w_i) < \infty$ for $i \in \mathcal{I}$. From the definition of MDP (6) at stage i, we know that for each policy $\pi_i \in \Pi_i$, we have

$$V_i(x_i, w_i) \leqslant \mathbb{E}_i\left[\sum_{j\in\mathcal{I}_i} \gamma_{\mathrm{A}}^{j-i} c_j(x_j^{\pi_i}, w_j, Z_j^{\pi_i}(x_j^{\pi_i}, w_j)) + \gamma_{\mathrm{A}}^{I-1-i} c_{I-1}(x_{I-1}^{\pi_i}, w_{I-1})\right].$$

By choosing π_i to be the policy that only procures from the short-term market (i.e., CPPA procurement decisions $z_{i,m}$ are always zero), we obtain the following inequalities:

$$V_i(x_i, w_i) \leqslant \mathbb{E}_i\left[\sum_{j=i}^{I-1} \gamma_{\mathrm{A}}^{j-i} \mathbb{E}_j\left[\sum_{t=0}^{T-1} \gamma^t P_{j,t} D_{j,t} + \gamma_{\mathrm{A}} \alpha R_{j+1} D_j \mathbf{1}_{\{j\in\mathcal{I}^{\mathrm{S}}\}}\right]\right] \tag{16a}$$

$$= \sum_{j=i}^{I-1} \gamma_{\mathrm{A}}^{j-i} \left(\sum_{t=0}^{T-1} \gamma^t D_{j,t} \mathbb{E}_i\big[\mathbb{E}_j[P_{j,t}]\big] + \gamma_{\mathrm{A}} \alpha D_j \mathbf{1}_{\{j\in\mathcal{I}^{\mathrm{S}}\}} \mathbb{E}_i\big[\mathbb{E}_j[R_{j+1}]\big]\right), \tag{16b}$$

where (16a) is obtained by accounting for the costs of short-term procurement and (16b) follows by linearity of the expectation operator. Due to the law of iterated expectations, it holds that $\mathbb{E}_i[\mathbb{E}_j[P_{j,t}]] = \mathbb{E}_i[P_{j,t}] \leqslant \mathbb{E}_i[|P_{j,t}|]$ and $\mathbb{E}_i[\mathbb{E}_j[R_{j+1}]] = \mathbb{E}_i[R_{j+1}]$. The value function satisfying $V_i(x_i, w_i) < \infty$ thus follows by our assumption on finite price expectations.

Next we show that $V_i(x_i, w_i) > -\infty$. The cost function $c_j(x_j, w_j, z_j)$ defined in (4) contains terms of the form $-P_{j+l,t}\, \eta_{j+l,t}\, z_{j,m}$ which are generally negative (unless power prices are negative), while the remaining cost components are positive. Thus, for each $j \in \mathcal{I}_i$, it holds that:

$$c_j(x_j, w_j, z_j) \geqslant -\left|\mathbb{E}_j\left[\sum_{m\in\mathcal{M}} \sum_{l=1}^{L_{j,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}} \gamma^t P_{j+l,t}\, \eta_{j+l,t}\, z_{j,m}\right]\right| \tag{17a}$$

$$= -\left|\sum_{m\in\mathcal{M}} \sum_{l=1}^{L_{j,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}} \gamma^t \eta_{j+l,t}\, z_{j,m} \mathbb{E}_j[P_{j+l,t}]\right| \tag{17b}$$

$$\geqslant -\sum_{m\in\mathcal{M}} \sum_{l=1}^{L_{j,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}} \gamma^t \eta_{j+l,t}\, z_{j,m}\, \mathbb{E}_j[|P_{j+l,t}|] \tag{17c}$$

$$\geqslant -\sum_{m\in\mathcal{M}} \sum_{l=1}^{L_{j,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}} \gamma^t \eta_{j+l,t}\, N\, \mathbb{E}_j[|P_{j+l,t}|], \tag{17d}$$

where (17b) follows by linearity of the expectation operator, (17c) holds by bringing the absolute values inside the summations and further inside the expectation operator, and (17d) follows by the assumption on bounded action space by choosing $N := \max\{z_m^{\max} : m \in \mathcal{M}\} < \infty$. Therefore, at a given stage i, it holds that

$$V_i(x_i, w_i) \geqslant \mathbb{E}_i\left[\sum_{j\in\mathcal{I}_i} \gamma_{\mathrm{A}}^{j-i}\left[-\sum_{m\in\mathcal{M}} \sum_{l=1}^{L_{j,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}} \gamma^t \eta_{j+l,t}\, N\, \mathbb{E}_j[|P_{j+l,t}|]\right]\right] \tag{18}$$

$$= -\sum_{j\in\mathcal{I}_i} \gamma_{\mathrm{A}}^{j-i} \sum_{m\in\mathcal{M}} \sum_{l=1}^{L_{j,m}} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}} \gamma^t \eta_{j+l,t}\, N\, \mathbb{E}_i[\mathbb{E}_j[|P_{j+l,t}|]], \tag{19}$$

where (19) follows by linearity of the expectation operator. Since $\mathbb{E}_i[\mathbb{E}_j[|P_{j,t}|]] = \mathbb{E}_i[|P_{j,t}|]$ due to the law of iterated expectations, $V_i(x_i, w_i) > -\infty$ follows from our assumption on finite price expectations. Therefore, it holds that $|V_i(\cdot,\cdot)| < \infty$. $\square$

D. Convex and fluid PFRH approximation

In this section, we show that a convex and fluid approximation of the PFRH math program does not provide a lower (dual) bound to the optimal MDP cost, although this is the case with deterministic approximations in other applications (Gallego and Van Ryzin 1997), which further bolsters the need for information relaxations to define lower bounds in our procurement application. Below we: (i) introduce a "convex-fluid" variant of the PFRH decision model, and (ii) construct a simple counter-example where such a model does not provide a lower bound.

Recall that the stage-i PFRH decision model (9) relies on forecasts of the availability factors defined by $\bar{a}_{i,j,m} = 1$, if $\mathbb{E}_i[a_{j,m}] > 0.5$, and $\bar{a}_{i,j,m} = 0$ otherwise, for $i \leqslant j$ and $m \in \mathcal{M}$. These forecasts define the feasible action set $\mathcal{Z}_j(\bar{a}_{i,j})$ used in PFRH. To define the convex-fluid PFRH approximation, we first convexify the action set $\mathcal{Z}_j(\bar{a}_{i,j})$ by dropping the minimum CPPA purchase quantity $z_m^{\min}$ for each stage $j \geqslant i$ and contract $m \in \mathcal{M}$. This operation results in a relaxation of the original PFRH model (9) constraints. Second, we consider a fluid approximation with respect to the contract availability factors $a_{i,j}$, which amounts to allowing decisions z_j to vary in the set $\mathcal{Y}_j(a_{i,j})$ defined by

$$\mathcal{Y}_j(a_{i,j}) := \left\{ z_j \in \mathbb{R}_+^{|\mathcal{M}|} \,\middle|\, z_{j,m} \in \left[0,\, \mathbb{E}_i[a_{j,m}] \cdot z_m^{\max}\right],\ \forall m \in \mathcal{M} \right\}.$$

Finally, notice that the objective function (9a) of PFRH is inconsistent with the MDP objective function as it replaces individual uncertain variables with their expected value whereas the MDP uses the expected cost (4). Therefore, we use the latter expected cost as the objective. The resulting convex-fluid PFRH approximation defined at stage i thus solves:

$$\hat{V}_i(x_i, w_i) = \min_{y_j, z_j} \sum_{j \in \mathcal{I}_i} \gamma_{\mathrm{A}}^{j-i} c_j\big(y_j, w_j, z_j\big) + \gamma_{\mathrm{A}}^{I-1-i} c_{I-1}\big(y_{I-1}, w_{I-1}\big) \tag{20a}$$

$$\text{s.t.:}\ \ y_i = x_i, \tag{20b}$$

$$y_{j+1} = f_j(y_j, z_j), \ \ \forall\, j \in \mathcal{I}_i, \tag{20c}$$

$$y_j \in \mathcal{X}_j, \ \ \forall\, j \in \mathcal{I}_i \cup \{I-1\}, \tag{20d}$$

$$z_j \in \mathcal{Y}_j(a_{i,j}), \ \ \forall\, j \in \mathcal{I}_i. \tag{20e}$$

The decision variables in the math program (20) coincide with those from the original PFRH model (9), and constraints (20b)–(20d) are the same too. However, model (20) differs from (9) since the former model uses a modified objective function (20a) and the convex-fluid action set (20e).

Despite the convex and fluid approximation introduced in (20), solving this model at stage i does not provide in general a lower bound on the optimal MDP value function $V_i(x_i, w_i)$. We illustrate this by constructing a counter-example based on a three-stage model with no inner-stages and with two CPPA procurement options: a "long-term" CPPA which can be signed at stage 0 and delivers power at stages 1 and 2, and a "short-term" CPPA that can be signed at stage 1 and delivers power at stage 2. The initial CPPA inventory is empty, i.e., $x_0 = 0$. We consider a demand of 10 MWh at each stage, an RPPT of 80% in years 1 and 2, and for simplicity assume a stage is made of a single hour. We also consider deterministic and constant power and REC prices equal to 20 and 10 USD/MWh, respectively, and a deterministic strike price of 22 USD/MWh at each stage. There are two equally-likely supply scenarios for stage 1: $\eta_1 = 0.3$ (low supply) and $\eta_1 = 0.7$ (high supply) and no further branching after stage 1, which means $\eta_2|\eta_1 = \eta_1$. The action space is convex and equal to $(z_0, z_1) \in [0, 50]^2 \subset \mathbb{R}^2$ (unit: MW^2) and both CPPAs are always available, i.e., all availability factors $a_{i,m}$ are equal to 1.

Solving the convex-fluid approximation (20) in this setting results in an objective value of $\hat{V}_0(x_0, w_0) = 453.3$ USD and optimal decisions $\hat{z}_0 = 26.67$ MW and $\hat{z}_1 = 0.00$ MW. Solving the MDP to optimality (which is easy in this simplified setting) gives instead the value $V_0(x_0, w_0) = 450.3$ USD. Since $V_0(x_0, w_0) < \hat{V}_0(x_0, w_0)$, then $\hat{V}_0(x_0, w_0)$ is not a lower bound on the optimal MDP cost. The optimal MDP decisions are $z_0^* = 11.43$ MW for stage 0 and $z_1^* = 15.24$ and $z_1^* = 0.00$ for stage 1, respectively, for the low and high supply scenario. In this example, the MDP optimal solution indicates that it is preferable to enter a long-term CPPA with smaller capacity compared to the decisions from (20), observe the supply scenario at stage 1, and then decide whether or not to complement the on-hand CPPA with a second shorter CPPA. If supply is low, signing a second CPPA is economical but if supply is high, it is optimal to rely on the CPPA already signed and unbundled RECs. Intuitively, the flexibility to adapt decisions at a stage i to realizations of the uncertainty at that stage allows the optimal MDP policy to better optimize costs compared to the decisions in (20).

E. CPPA strike price

Consider a renewable power generator that begins installation at period $(i, t = 0)$, $i \in \mathcal{I}$. We assume the generator needs one year to be operational, that is, production starts at $(i+1, t = 0)$. The generator has an expected lifetime of L^{P} years and incurs a cost of C_i^{INV} capturing the one-time installation and estimated maintenance costs associated with 1 MW of production capacity as well as any applicable investment tax credit[1]. We assume that there is a production tax credit (PTC)[2] of G_i USD per MWh for the next L_i^{G} years and that future cash flows are discounted at each interval by rate $r \in (0, 1]$, which can be chosen to also account for the generator's target return on investment. The associated annual discount rate is $r_{\mathrm{A}} = r^T \in (0, 1]$. The production $\eta_{i,t} \in [0, H]$ is stochastic at each year l over the generator lifetime and each interval $t \in \mathcal{T}$. We begin by computing the fixed strike price $\hat{K}_i$ of a CPPA that spans the lifetime of the generator using the net present value (NPV) of the contract's cash flows. This approach is consistent with the System Advisor Model[3]

[1]An investment tax credit represents a one-time federal tax deduction equal to a pre-specified percentage of the installation cost of a renewable power project.

[2]A production tax credit provides a per-megawatt-hour tax credit for power generation for a fixed number of future years from the installation of a renewable power project.

[3]SAM is an open source performance and financial tool designed by NREL to access the feasibility of renewable energy projects (e.g., wind, solar, or biomass).

(SAM; NREL 2017). The NPV of 1 MW of installed capacity contracted via such a CPPA is

$$\mathrm{NPV}_i = \sum_{l=1}^{L^{\mathrm{P}}} \sum_{t=0}^{T-1} r_{\mathrm{A}}^l r^t \, \mathbb{E}_i[\eta_{i+l,t}] \, \hat{K}_i \, + \, \sum_{l=1}^{L_i^{\mathrm{G}}} \sum_{t=0}^{T-1} r_{\mathrm{A}}^l r^t \, \mathbb{E}_i[\eta_{i+l,t}] \, G_i \, - \, C_i^{\mathrm{INV}}.$$

Setting NPV_i to zero, we obtain the following strike price formula:

$$\hat{K}_i = \frac{1}{\sum_{l=1}^{L^{\mathrm{P}}} \sum_{t=0}^{T-1} r_{\mathrm{A}}^l r^t \, \mathbb{E}_i[\eta_{i+l,t}]} \left[C_i^{\mathrm{INV}} - \sum_{l=1}^{L_i^{\mathrm{G}}} \sum_{t=0}^{T-1} r_{\mathrm{A}}^l r^t \, \mathbb{E}_i[\eta_{i+l,t}] \, G_i \right]. \tag{21}$$

Expression (21) captures the dependence of the strike price on generator vintage (i.e., the year that production begins) by treating the production tax credit G_i, investment cost C_i^{INV}, and capacity factors $\eta_{i,t}$ as time-dependent quantities. Currently, renewable power generators that started construction before 2020 are eligible for productions tax credits for 10 years from the date the facility starts production (DOE 2016) but this status-quo is likely to change with government regulation. Investment costs and capacity factors typically decrease and increase, respectively, over time due to improvements in technology. The capacity factor, in addition, exhibits significant inter-region variation. For instance, in the case of wind power, capacity factors in the "internal" regions of the United States are significantly higher than those of coastal regions (Wiser and Bolinger 2017).

Next, we describe how the strike price $\hat{K}_i$ in (21) can be modified to account for shorter contract lengths and the expected power price over the tenure of the contract. Consider a CPPA with a duration of m years that is less than the lifetime L^P of the generator. Shorter contracts result in additional cash flow risk over the period of the generator's life time for which they do not generate revenue (ACORE 2016). We thus define a risk-adjusted strike price $\hat{K}_{i,m} := \hat{K}_i \cdot K_m^+$, where $K_m^+ \geqslant 1$ is a risk factor that inflates the strike price if $m < L^{\mathrm{P}}$ and equals 1 otherwise, that is, $K_m^+ = 1$ when the contract spans the life of the generator. The CPPA strike price is not solely determined by NPV but is also tied to the long-term expected power price because higher expected (future) power prices give the generator leverage to increase the CPPA price since the company's outside option is expensive (Wiser and Bolinger 2017). To account for this effect, we lower bound the CPPA strike price by the discounted average power price over the tenure of the contract, which is

$$\underline{K}_{i,m} = \frac{1}{\sum_{l=1}^{m} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}}^l \gamma^t} \sum_{l=1}^{m} \sum_{t=0}^{T-1} \gamma_{\mathrm{A}}^l \gamma^t \, \mathbb{E}_i \left[P_{i+l,t} \right].$$

Our final strike price expression is $K_{i,m} := \max \left\{ \hat{K}_{i,m}, \, \underline{K}_{i,m} \right\}$.

The parameters defining the strike price model of our baseline CPPA instance are reported in Table 6. Following NREL (2010), we use a functional form for C_i^{INV} that decreases over time by a fixed percentage ξ; specifically, it evolves according to a learning model $C_i^{\mathrm{INV}} = C_0^{\mathrm{INV}}(1-\xi)^i$. We chose the initial cost (C_0^{INV}) based on 2015–2016 wind projects in the U.S. (EIA 2018) and the learning rate $\xi = 1\%$ based on the range of values in NREL (2010). Wind turbines are usually designed to operate for 20–25 years but many remain operational for a longer period of time (Ziegler et al. 2018), thus we select the lifetime (L^{P}) to be 30 years also to account for improving technology.

Table 6: Strike price model parameters in the baseline CPPA instance.

Name	Value	Unit	Name	Value	Unit
L^{P}	30	years	γ_{A}	0.97	-
L_i^{G}	10	years	r_{A}	0.93	-
G_i	23	USD/MWh	K_5^+	1.2	-
C_0^{INV}	$1.7 \cdot 10^6$	USD/MW	K_{25}^+	1.1	-

The duration of the production tax credit (L_i^{G}) and its amount (G_i) are based on United States policy in 2016 (DOE 2016). Moreover, we assume the tax credit expires in 5 years, i.e., it is only granted to renewable energy facilities commencing construction at stages $i < 5$. The discount factors γ_{A} and r_{A} are set to 0.97 and 0.93, respectively, as explained in §6.1. We use a maximum risk factor $K_+^5 = 1.2$, i.e. a 20% premium for the 5-year CPPAs, which decreases linearly as m is increased until $K_+^{25} = 1.1$.

F. Model of market dynamics and calibration

In this section we present in detail the stochastic processes used to describe the evolution of the power price, REC price, and power supply, and discuss their calibration using market data.

The evolution of electricity prices has been studied using various processes that capture features such as seasonality (Lucia and Schwartz 2002), mean-reversion and long-term trends (Schwartz and Smith 2000 and references therein), and jumps (Weron 2014, 2007, Cartea and Figueroa 2005, Seifert and Uhrig-Homburg 2007, Escribano et al. 2011). To obtain a power price model that captures the main features of short-term electricity prices, we construct a mean-reverting stochastic process with jumps and seasonality. We use a continuous-time process for the power price $\{P_s,\ s \in \mathbb{R}_+\}$, and then consider in our decision model discrete-time values $\{P_{i,t},\ (i,t) \in \mathcal{I} \times \mathcal{T}\}$, which are the values taken by this process at the beginning of month t of year i (we do analogously for the stochastic processes of REC price and power supply). Following Weron (2014), the power price model is written as:

$$\ln(P_s) = \chi_s + g(s), \tag{22a}$$

$$d\chi_s = (\nu_P - \mathcal{K}_P\, \chi_s)\, ds + \sigma_P\, dW_s + J(\mu_J, \sigma_J)\, d\Pi(\lambda), \tag{22b}$$

$$g(s) = \phi_0 + \sum_{k=1}^{12} \phi_k \hat{P}_s^k. \tag{22c}$$

Equation (22a) describes the log power price process as the sum of a stochastic component χ_s and a deterministic component $g(s)$. The stochastic component evolves according to (22b), where $\mathcal{K}_P$ is the speed of mean reversion, ν_P models the drift, σ_P is the volatility, and W_t is a standard Brownian motion. We model spikes in monthly prices by a jump diffusion process in which the jump size follows a normal distribution $J(\mu_J, \sigma_J)$ and the jump frequency a Poisson distribution $\Pi(\lambda)$ (Cartea and Figueroa 2005). The deterministic function $g(s)$ in (22c) models the monthly price seasonality by using a constant ϕ_k for each month k, and binary values $\hat{P}_s^k$ equal to one if time s falls in month k and zero otherwise.

We calibrated the parameters of model (22) using historical monthly power price data from the Pennsylvania New Jersey Maryland Interconnection LLC (PJM) market during the period January

2010–August 2017. Analyzing power prices, we found that the jump frequency and intensity are small when considering monthly prices. We thus tested the number of jumps in the monthly power prices using the algorithm presented in Weron (2014, p. 1047), and found the jump diffusion parameters to be insignificant. Therefore, we removed jumps from the model and only focused on a mean-reverting process with seasonality. We first estimated the seasonality function $g(s)$ directly from the data using linear regression, resulting in the coefficients $\{\phi_k,\ k = 0 \ldots, 12\} = \{3.519, 0.163, 0.163, 0.078, -0.026, 0.003, 0.039, 0.174, 0.013, 0.009, -0.037, -0.038, 0.0\}$. Then, we calibrated the mean reverting coefficients using maximum likelihood estimation. The resulting estimates were $\mathcal{K}_P = 0.295$, $\sigma_P = 0.178$ (both with a p-value below 0.001), and $\nu_P = 0$. We set $P_0 = 31.5$ USD/MWh which is the average power price observed in 2017. When plotting the sample paths using the calibrated parameters, we noticed that the speed of mean reversion was too strong and so we decreased it from $\mathcal{K}_P = 0.295$ to $\mathcal{K}_P = 0.04$ for our numerical study. We illustrate this effect in Figures 7(a) and 7(b), which display 300 power price sample paths generated in Monte Carlo simulation over 20 years using $\mathcal{K}_P = 0.295$ and $\mathcal{K}_P = 0.04$, respectively. Finally, to construct the instance set S3, we varied the calibrated value for ν_P from 0 to -0.0105, -0.0178, and -0.0268 to obtain long-term mean power prices equal to 30, 25, and 20 USD/MWh, respectively.

Figure 7: 300 Monte Carlo sample paths of power prices for different mean-reversion speeds.

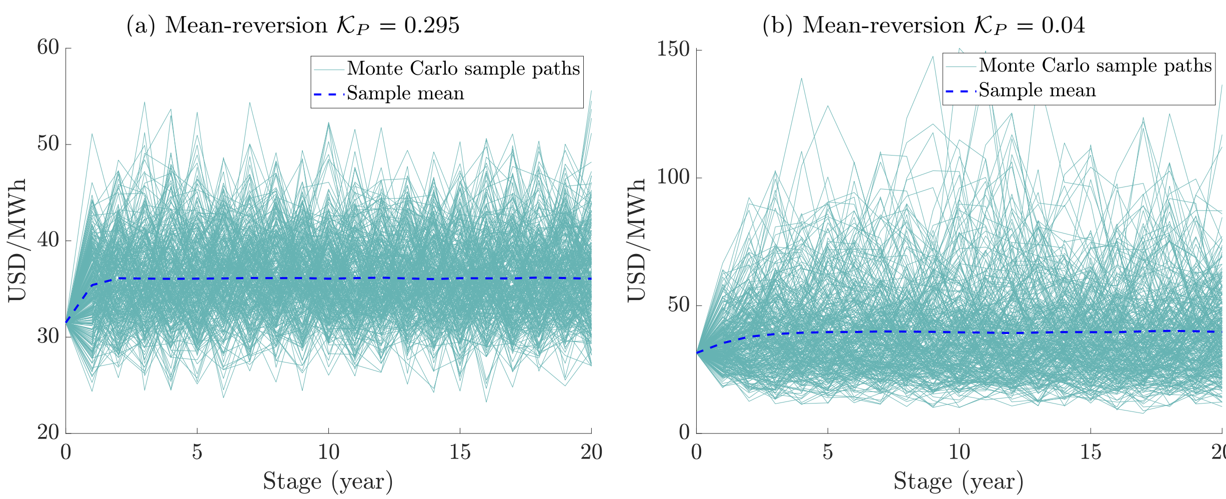

The dynamics of REC prices has been less studied in the literature. Under renewable portfolio standards, which is one of the prominent support programs for renewable energy sources, the regulator requires producers, distributors, and consumers to purchase RECs. To forecast the evolution of REC prices, following Zeng et al. (2015), we use a Jacobi diffusion process to generate values between zero and one, and obtain REC prices as the product between the output of this stochastic process and an upper bound threshold. The stochastic process is defined by

$$dr_s = (\nu_R - \mathcal{K}_R r_s)ds + \sigma_R \sqrt{r_s(1 - r_s)}\, dW_s, \tag{23a}$$

$$R_s = r_s \cdot \overline{R}. \tag{23b}$$

In equation (23a), ν_R and $\mathcal{K}_R$ are the mean-reverting parameters, σ_R is the volatility, and W_s is a standard Brownian motion. The process (23a) generates values $r_s \in [0,1]$, which are then scaled in the second equation (23b) by the threshold value $\overline{R}$ to ensure R_s belongs to the interval $[0, \overline{R}]$.

We chose an upper bound of $\overline{R} = 60$ USD/MWh that is representative of wind RECs (EPA 2018). We then estimated the parameters of the REC price model using monthly data for New Jersey REC prices between May 2015 and December 2017, and an adaptation of the maximum likelihood estimation method for Jacobi diffusion processes from Gouriéroux and Valéry (2004). We obtained the parameter estimates $\mathcal{K}_R = 0.448$, $\nu_R = 0.066$, and $\sigma_R = 0.109$ (p-values were below 0.001). We set the initial value $R_0 = 10$ USD/MWh, which is representative of the average REC price in our time series. To construct the instance set S4, we changed the calibrated value for ν_R from 0.066 to 0.0343, 0.105, and 0.141 to obtain long-term mean REC prices equal to 5, 15, and 20 USD/MWh, respectively.

We model monthly power supply using a mean reverting process with seasonality (Loukatou et al. 2018), which is similar to the power price model (22) but without jumps, as shown below.

$$\ln(\Theta_s) = \chi_s + g(s), \tag{24a}$$

$$d\chi_s = (\nu_\eta - \mathcal{K}_\eta \chi_s)\, ds + \sigma_\eta \, dW_s, \tag{24b}$$

$$g(s) = \phi_0 + \sum_{k=1}^{12} \phi_k \hat{\eta}_s^k, \tag{24c}$$

$$\eta_s = H \cdot \Theta_s. \tag{24d}$$

The stochastic process defined in (24a)–(24c) models a capacity factor $\Theta_s \in [0,1]$. In (24d), the output from this process is then multiplied by the number of hours H in an inner-stage to obtain a power supply value in $\eta_s \in [0, H]$ (hours) consistent with our MDP (6).

We calibrated model (24) using historical capacity factors in the United States from January 2011 to June 2019 (EIA 2019). We proceeded essentially using the same steps we used to calibrate the power price model (22). Specifically, we first estimated the seasonality function $g(s)$ from the data with a linear regression, which resulted in coefficients $\{\phi_k, k = 0 \ldots, 12\} = \{-1.055, 0.014, 0.064, 0.089, 0.162, 0.041, -0.030, -0.289, -0.369, -0.254, -0.035, 0.067, 0\}$. We then used maximum likelihood to estimate the remaining coefficients of the model, obtaining the estimates $\mathcal{K}_\eta = 0.814$, $\sigma_\eta = 0.0919$ (both with p-value lower than 0.001), and $\nu_\eta = 0$. We set $\Theta_0 = 0.36$, which was the average capacity factor observed in our time series.

The calibrated models allow us to generate sample paths of the uncertainty in Monte Carlo simulation, which are needed to estimate the value of the procurement policies and dual bounds.

G. Additional numerical results

In Figures 8–10 we compare the behavior of the different procurement policies on the extended instance set S1–S6. Specifically, we consider the PIRH dual bound and the following policies: Short-term, FBH_{20}, PFRH, and PIRH with median decision measure (henceforth PIRH).

In Figure 8(a), which refers to instance set S1, PIRH performs best on all the S1 instances with average optimality gap of 1.4%, followed by PFRH with average optimality gap of 3.4%. The performance of block and short-term heuristics is largely inferior to reoptimization methods (the average

Figure 8: Procurement costs for the S1 and S2 instance sets.

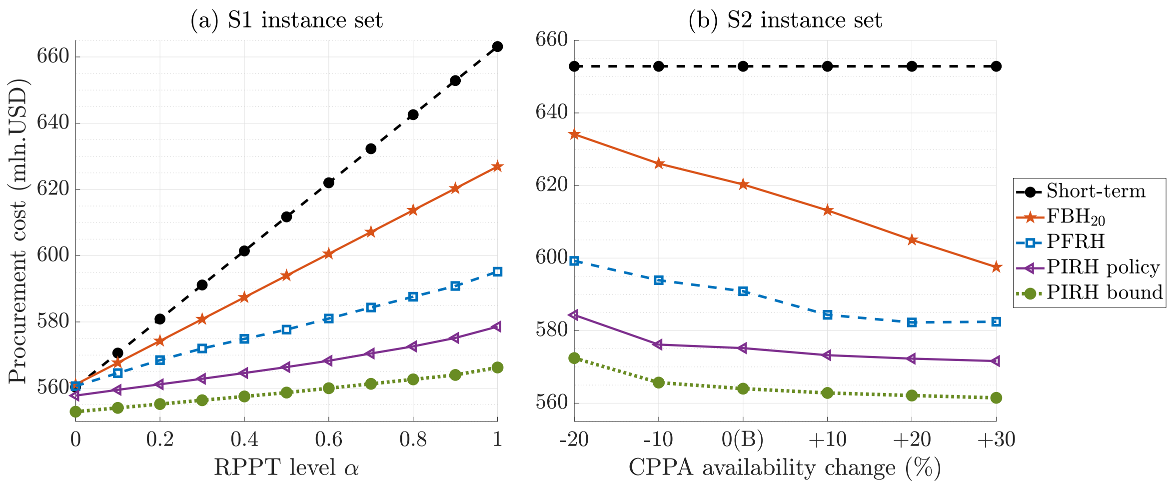

optimality gap is respectively 6.2% and 9.4%). Figure 8(b) refers to instance set S2. The relative ranking of methods in these instances is the same as in set S1, with PIRH achieving the smallest optimality gaps across instances (1.90% on average). However, contract availability impacts the block policy more heavily than PFRH/PIRH. In fact, the FBH$_{20}$ procurement cost increases by more than 6% when decreasing contract availability from +30% to −20% relative to the baseline, which corresponds to more than 36 mln. USD. In contrast, the maximum cost increase is less than 3% (16 mln. USD) under PFRH/PIRH that use dynamic portfolios containing multiple CPPAs.

Results for the S3 and S4 instance sets are displayed in Figures 9(a) and 9(b), respectively. The relative performance of methods on these instance sets is largely consistent with our prior observations on the S1–S2 instances. Interestingly, the shortsighted objective function of block policies (see the FBH model (7)) results in FBH$_{20}$ entering into some CPPAs even under low power prices as shown in left part of Figure 9(a), leading to larger procurement costs than the short-term policy.

Figure 9: Procurement costs for the S3 and S4 instance sets.

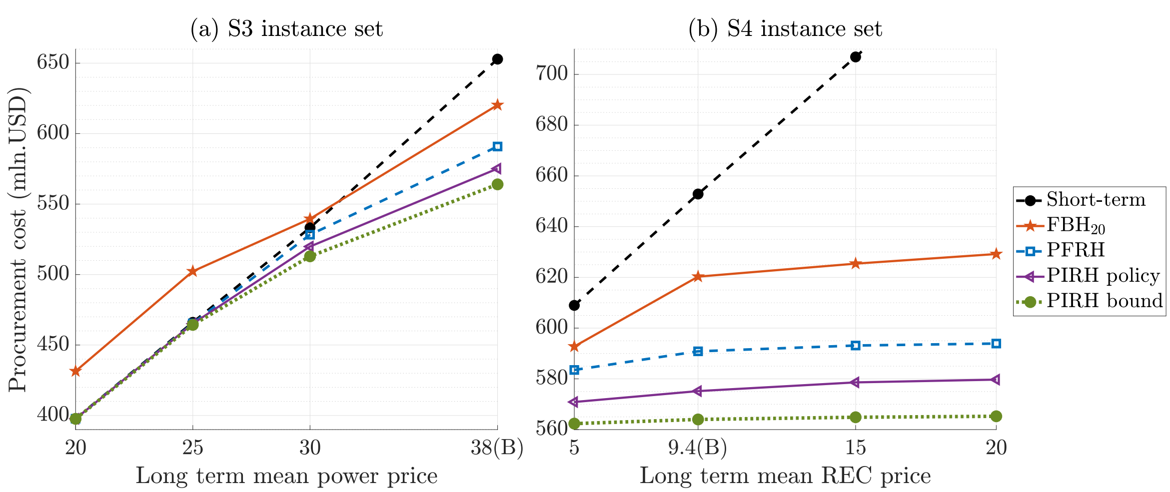

Finally, results for the S5 and S6 instance sets are displayed in Figures 10(a) and 10(b), respectively. The results in Figure 10(a) show that the procurement cost increases under all methods when the generator discount factor decreases (i.e., return on investment increases), except for the short-term policy. Similar to instances S3, Figure 10(a) shows that FBH$_{20}$ enters into CPPAs even when the short-term policy is near optimal. The impact of correlation on procurement cost is visible in Figure 10(b). Varying $\bar{\rho}$ from $+1$ to -1 results in an cost increase of 2–3% across methods.

Figure 10: Procurement costs for the S5 and S6 instance sets.

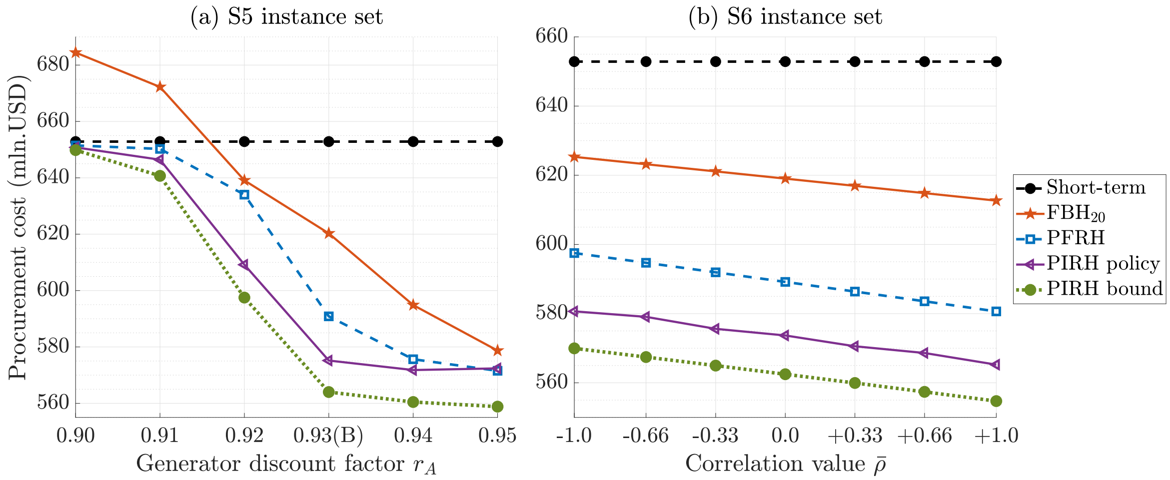

References

ACORE. 2016. Renewable energy ppa guidebook for corporate & industrial purchasers. Tech. rep., American Council On Renewable Energy (ACORE). URL http://article.media/purchasers8.

Cartea, A., M. G. Figueroa. 2005. Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* **12**(4) 313–335.

DOE. 2016. Renewable electricity production tax credit (ptc). URL http://production.credit/ptc1. Department of Energy (DOE).

EIA. 2018. Wind generators' cost declines reflect technology improvements and siting decisions. URL http://costs.energy/eia6. Energy Information Administration (EIA).

EIA. 2019. Capacity factors for utility scale generators primarily using non-fossil fuels. URL http://output.energy/eia3. Energy Information Administration (EIA).

EPA. 2018. Green power pricing. URL http://powers.supplies/green6. United States Environmental Protection Agency (EPA).

Escribano, A., J. Ignacio Pena, P. Villaplana. 2011. Modelling electricity prices: International evidence. *Oxford Bulletin of Economics and Statistics* **73**(5) 622–650.

Gallego, Guillermo, Garrett Van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations research* **45**(1) 24–41.

Gouriéroux, C., P. Valéry. 2004. Estimation of a Jacobi process. Working paper, University of Toronto and Université de Montréal.

Heyman, D. P., M. J. Sobel. 1982. *Stochastic Models in Operations Research: Stochastic Optimization*, vol. 2. Courier Corporation.

Loukatou, A., S. Howell, P. Johnson, P. Duck. 2018. Stochastic wind speed modelling for estimation of expected wind power output. *Applied energy* **228** 1328–1340.

Lucia, J. J., E. S. Schwartz. 2002. Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Review of Derivatives Research* **5**(1) 5–50.

NREL. 2010. Cost and performance assumptions for modeling electricity generation technologies. Tech. Rep. NREL/SR-6A20-48595, National Renewable Energy Laboratory (NREL).

NREL. 2017. System Advisor Model (SAM). URL https://sam.nrel.gov/. National Renewable Energy Laboratory (NREL).

Schwartz, E. S., J. E. Smith. 2000. Short-term variations and long-term dynamics in commodity prices. *Management Science* **46**(7) 893–911.

Seifert, J., M. Uhrig-Homburg. 2007. Modelling jumps in electricity prices: Theory and empirical evidence. *Review of Derivatives Research* **10**(1) 59–85.

Weron, R. 2007. *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*, vol. 403. John Wiley & Sons, West Sussex, Englend.

Weron, R. 2014. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* **30**(4) 1030–1081.

Wiser, R., M. Bolinger. 2017. 2016 wind technologies market report. Tech. rep.

Zeng, Y., D. Klabjan, J. Arinez. 2015. Distributed solar renewable generation: Option contracts with renewable energy credit uncertainty. *Energy Economics* **48** 295–305.

Ziegler, L., E. Gonzalez, T. Rubert, U. Smolka, J. J. Melero. 2018. Lifetime extension of onshore wind turbines: A review covering Germany, Spain, Denmark, and the UK. *Renewable and Sustainable Energy Reviews* **82** 1261–1271.