

Optimal Renewable Power Purchase Agreements for Data Centers

Kekun Gao*, Yuejun Yan*, Yuke Zhou*, Endong Liu, Zhaoyang Wang, Pengcheng You

Abstract—Data centers have become integral to contemporary infrastructure. However, their electricity demand significantly adds to global energy consumption, resulting in substantial electricity costs and carbon footprints. Power purchase agreements, when used as bilateral long-term contracts for procuring renewable energy, not only hedge against volatility in an electricity market to secure locked-in prices but also play a crucial role in reducing carbon emissions. This paper studies the strategy of a data center to sign such power purchase agreements, which encompasses an optimal agreement design problem (static) and an optimal signing timing problem (dynamic) to maximize long-term expected electricity cost savings. In particular, we develop a continuous-time stochastic process model for long-term market price evolution in an electricity market, alongside a power purchase agreement model comprising the starting time, contract term, contractual power supply, and locked-in price. These models enable a novel problem formulation for PPA signing, which further allows for a tractable solution via decomposition. We first fix an arbitrary starting time and explicitly analyze how to set the other three variables with the most expected total saving. We then propose a dynamic threshold policy that online identifies an optimal starting time based on real-time observations of market prices. The proposed strategy is extensively tested with simulation experiments to validate our theoretical analysis. Numerical results also provide additional insights – over 35% reduction in long-term total electricity expenditure can be made possible for data centers, as long as power purchase agreements are properly signed.

Index Terms—Power purchase agreement, Bilateral contract, Electricity market, Data center, Renewable energy.

I. INTRODUCTION

Data centers, as critical infrastructure, have been growing explosively in the past decade, leading to significant energy consumption and carbon emissions that negatively impact climate change. Renewable energy sources, e.g., solar and wind, emit almost zero greenhouse gases, and are widely recognized as sustainable alternatives to fossil energy. A soaring number of data centers are committed, e.g., via the RE100 initiative [1], to sustainability and climate goals [2], by purchasing renewable energy to meet demand, reducing emissions, and improving energy efficiency. However, the intermittent nature of renewable energy supply makes it unreliable for the uninterrupted operation of data centers. In this

This work was supported by NSFC through grants 72201007, 72431001, and CCF-Alibaba Innovative Research Fund. (Corresponding author: Pengcheng You)

K. Gao, Y. Zhou, and P. You are with the Department of Control Science and Systems Engineering, Peking University, Beijing, China (e-mail: pcyou@pku.edu.cn).

Y. Yan and Z. Wang are with Alibaba Group, Hangzhou, China.

E. Liu is with the School of Electrical and Mechanical Engineering, Heze University, Heze, China.

*These authors contributed equally to this work.

regard, renewable Power Purchase Agreements (PPAs) provide a promising solution.

A PPA is a long-term bilateral contract between a (commercial or industrial) consumer and a renewable generator. PPAs may last, say, between 5 and 20 years, during which the consumer purchases fixed amounts of electricity from the renewable generator at pre-negotiated prices [3]. On the one hand, by signing a PPA, the consumer locks in long-term energy costs, thereby mitigating risks associated with fluctuating electricity market prices. It also fulfills the commitment to decarbonization by procuring specifically renewable energy, such as wind and solar. On the other hand, PPAs will also be favored by the renewable generator as they provide financial certainty – a catalyst to boost investment in renewable facilities [4]. In 2023, the global volume of PPAs reached 46GW – a 12% increase compared with the total in 2022. From 2008 to 2023, the cumulative PPA capacity signed by corporations exceeded the total installed power generation capacity of countries such as France, the United Kingdom, and South Korea [5]. Notably, technology companies – particularly Google, Microsoft, and Meta – have been at the forefront of using PPAs to directly procure renewable energy. For example, Microsoft has committed over 10 billion US dollars to renewable energy investments to power its data centers [6]. As corporate adoption of PPAs continues to grow, especially within the technology sector, developing effective strategies to guide data centers in negotiating and signing PPAs has become increasingly important.

Recent studies have devoted increasing attention to signing PPAs, particularly with respect to managing key risk, price risk, and quantity risk in particular. One stream of research addresses price risk, which is especially pronounced in long-term fixed-price PPAs. When electricity market prices fall consistently below the agreed-upon PPA price, the power purchaser bears substantial financial exposure. To avoid this issue, [4], [7], [8] study optimal PPA prices and primarily adopt cost-based models to more accurately assess the Levelized Cost of Energy (LCOE) in renewable PPAs. These models account for production constraints and purchase limitations to set a fair price, but they do not explicitly consider electricity market price fluctuations. In contrast, other works such as [9] and [10] apply Nash Bargaining Theory to determine PPA pricing, with a particular emphasis on economic equilibrium between buyers and sellers to ensure mutual benefits through negotiation. Unlike LCOE-based models, these works account for both generation costs and strategic interactions to determine prices. However, existing literature typically restricts the focus to the internalities of PPAs and ignores the external impacts,

e.g., from alternative electricity markets. Our work advances a different approach by explicitly modeling long-term electricity market price behavior and capturing its influence on PPA performance.

Another major concern in PPA design is quantity risk, which arises from the intermittent nature of renewable energy generation. To mitigate such risk, a body of research has focused on portfolio-based strategies that diversify power generation assets across technologies and geographies. For instance, [3] demonstrates that optimizing multi-technology and multi-location PPA portfolios effectively reduces energy supply volatility. Similarly, [11] and [12] show that combining complementary energy sources can mitigate economic risk and enhance the financial stability of renewable power portfolios. Dynamic PPAs have also been proposed as a means to manage quantity risk. These contracts essentially use model predictive control to adjust decision variables over time based on predicted market conditions, facilitating flexible responses from both buyers and sellers [13]. However, the main focus of these studies remains at the portfolio level, emphasizing diversification rather than optimization for an individual PPA. By contrast, our work takes a contract-level perspective and aims to exploit the flexibility in PPA design to improve PPA performance.

A third line of research explores investment in renewable generation capacity under PPA frameworks. These studies typically examine how fixed-price contracts can incentivize early investment and ensure long-term access to renewable electricity. For example, [2] investigates how utilities determine optimal investment levels in response to fixed-price PPAs, while [14] analyzes investment decisions using a newsvendor-style model. Further advancing this line, [15] considers both investment timing and transfer payment structure, employing an optimal stopping framework to determine when a firm should sign a PPA into effect. Although our work is similar in terms of addressing design and timing of PPAs, it differs in focus and perspective. Specifically, we approach the problem from the standpoint of an electricity consumer, i.e., a data center, and consider not only signing timing for a PPA but also a more complex contract design, including finer-grained specifications of the contract term, locked-in price, and contractual power supply. By incorporating volatile electricity market prices into our analysis, we provide a novel framework for consumer-side optimization of renewable energy procurement via PPAs.

This paper investigates an optimal renewable PPA signing strategy for a data center, focusing specifically on both (static) PPA design and (dynamic) optimal signing timing to maximize long-term electricity expenditure savings. We establish a novel continuous-time stochastic programming formulation for the PPA sign problem, incorporating a Geometric Brownian Motion (GBM) model to characterize the realistic evolution of electricity market prices over long time horizons. The PPA is explicitly modeled as a bilateral contract specified by the starting time, contract term, contractual power supply, and locked-in price. Leveraging a decomposition approach, we first derive analytical insights into optimal PPA parameters for a given starting time, and subsequently develop a dynamic

threshold policy to determine the optimal timing to sign a PPA into effect. Extensive simulations validate the analytical results, providing robust insights into the substantial economic benefits achievable through strategically signing renewable PPAs.

A comparison with related works is summarized in Table I to highlight the gaps in the existing literature. Our work bridges these gaps by simultaneously addressing the optimal design of a PPA (contract term, contractual power supply, and locked-in price) and the optimal timing for starting the PPA based on dynamic electricity market conditions, modeled as a GBM. We further explicitly demonstrate the advantages of this integrated approach through extensive simulations that use real-world electricity market price data from New York ISO, providing robust insights and actionable decision support for the PPA signing of a data center.

Contributions of our work. The contributions are mainly threefold.

- We model the long-term evolution of electricity market prices using a GBM to capture general tendencies and fluctuations. This allows us to formulate the PPA signing problem as a continuous-time stochastic program. To the best of our knowledge, it is the first work that investigates the analytical impact of the PPA-market interplay on PPA decisions.
- We propose an integrated solution to PPA signing that jointly optimizes static PPA design (contract term, contractual power supply, and locked-in price) given an arbitrary starting time, and dynamically determines the optimal starting time by explicitly accounting for evolving electricity market prices. Our strategy is analytically sound, straightforward to implement, and adaptive to market dynamics.
- Extensive numerical simulations based on real-world electricity market price data illustrate that our PPA signing strategy can lead to substantial economic benefits with a possible 35% reduction in long-term electricity expenditure for a data center. Further sensitivity analyses provide additional insights into performance robustness and potential risk.

Organization of this paper. The rest of this paper is organized as follows: Section II describes our PPA signing model and problem formulation. Section III summarizes the main analytical results of the PPA model. Section IV shows the numerical results. Section V concludes the paper.

NOMENCLATURE

Acronyms

GBM Geometric Brownian Motion

LCOE Levelized Cost of Energy

PPA Power Purchase Agreement

SDE Stochastic Differential Equation

Parameters

λ Currency discount rate

μ Drift constant of a GBM

$\pi(t)$ Electricity market price at time t subject to a GBM

σ Volatility constant of a GBM

TABLE I
COMPARISON WITH RELATED WORKS.

References	Contract Design	Signing Timing	Dynamic Prices	Decision Support	Realistic Simulation
[4], [7], [8] (LCOE Models)	✓	✗	✗	✗	✓
[9], [10] (Bargaining Models)	✓	✗	✗	✓	✗
[3], [11], [12] (Portfolio Optimization)	✓ (multi-asset)	✗	✗	✓	✓
[13] (Dynamic PPA)	✓	✗	✗	✓	✓
[14], [15] (Investment)	✗	✓	✗	✓	✗
Our Work	✓	✓	✓ (partially)	✓	✓

$D(t)$	Electricity demand of a data center at time t
m	Constant
$W(t)$	The Wiener process
Symbols	
T_{critical}	Critical threshold for contract terms
\bar{p}	Upper-bound market price to guarantee nonnegative expected total saving from a PPA
\mathbb{E}	Expectation operator
$\mathbb{R}_{\geq 0}$	Set of nonnegative real numbers
Var	Variance operator
ω	Realization of market prices
ϕ	Optimal value function of PPA signing
S	Conditional expected total saving of a PPA (evaluated at its starting time)
t	Continuous time
$X(t)$	Monetary value at time t
Variables	
τ	Starting time of a PPA
k	Contractual power supply of a PPA
p	Locked-in price of a PPA
T	Contract term of a PPA

II. PROBLEM FORMULATION

Consider a setting where a data center procures renewable energy to meet its electricity demand by signing a PPA with a renewable generator. Since both the supply and demand are time-varying, the data center meanwhile participates in an electricity market – purchasing electricity when the renewable energy secured by the PPA is insufficient, while selling electricity when there is surplus [16]. A diagram of this setting is illustrated in Fig. 1. In this section, we first describe the models for the PPA and the electricity market, respectively, before explicitly formulating our problem of how to optimally sign the PPA.

A. Power Purchase Agreement

We particularly consider a PPA specified by a quadruple (τ, T, k, p) :

- starting time $\tau \in \mathbb{R}_{\geq 0}$: the PPA takes effect starting from time τ ;
- contract term $T \in \mathbb{R}_{\geq 0}$: the PPA is valid in the time window $[\tau, \tau + T]$;
- contractual power supply $k \in \mathbb{R}_{\geq 0}$: the PPA stipulates that the renewable generator must provide a fixed amount k of power supply;
- locked-in price $p \in \mathbb{R}_{\geq 0}$: the contractual power supply is charged at a fixed unit price p .




Fig. 1. Electricity and money flows of a data center. The data center signs a PPA with a renewable generator to procure renewable energy. Meanwhile, it trades electricity in an electricity market for any surplus or shortfall of the PPA.

The total cost of signing the PPA is then simply $p k T$ for the entire contract term. While such a PPA model is stylized, it captures key components of practical PPAs and will offer useful insights for PPA signing. Given a signed PPA, the data center will use the contractual power supply first to meet its electricity demand since (i) it is already paid for and reserved; (ii) it is from carbon-free renewable energy. There are two possibilities:

- 1) The contractual power supply k falls short of the real-time electricity demand. In this case, the data center has to make up for the shortfall from an electricity market, subject to time-varying and volatile market prices.
- 2) The contractual power supply k exceeds the real-time electricity demand. In this case, the data center can sell the surplus to the electricity market also at market prices.¹

To focus on the standpoint of the data center, we assume that the contractual power supply is reliable, i.e., the renewable generator is always able to provide the fixed amount k of power supply specified in the PPA. In practice, there is usually a penalty for failing to fulfill the commitment. For instance, the renewable generator may be asked to compensate for the undersupply from the electricity market. Therefore, approximately, the contractual power supply can be viewed as guaranteed for the data center. The extension to a potentially unreliable power supply from the renewable generator falls out of the scope of this work and will be left for future studies.

¹In wholesale electricity markets, clearing prices are identical for suppliers and consumers at the same locations.

B. Electricity Market

The data center can participate in an electricity market to purchase or sell electricity. While on-site small-scale distributed renewable generation for the data center is possibly available, we will assume it is all locally consumed and will not be fed back to power grids, in order to restrict our focus on the impact of the electricity market on PPA signing. Therefore, the data center serves as a supplier in the electricity market only when there is surplus from the PPA. The electricity market is considered a huge pool with adequate supply and demand. Therefore, the data center is always able to purchase or sell any amount of electricity at time-varying market prices that are treated as an external stochastic input.²

Inspired by the Black–Scholes model in mathematical finance [17], [18], we adopt a GBM to characterize long-term stochastic market price behavior in the electricity market, e.g., in a time span of 1-2 decades. In particular, denote the market price at time t as $\pi(t)$, which is governed by a continuous-time stochastic process, i.e.,

$$d\pi(t) = \mu\pi(t)dt + \sigma\pi(t)dW(t), \quad \forall t. \quad (1)$$

(1) is a stochastic differential equation (SDE) parameterized by a drift constant $\mu > 0$ and a volatility constant $\sigma > 0$ to respectively capture general growth trends and unpredictable fluctuations of market prices $\pi(t)$. The randomness is captured in $W(t)$ subject to the Wiener process [19], [20], [21].

To tackle the SDE (1), we apply Ito's Lemma to the particular natural logarithmic function of the market price $\ln(\pi(t))$ to attain its derivative [22], [23]:

$$d[\ln(\pi(t))] = \frac{1}{\pi(t)}d\pi(t) - \frac{1}{2}\frac{1}{\pi^2(t)}(d\pi(t))^2,$$

which, with (1) plugged in, reduces to

$$d[\ln(\pi(t))] = \left(\mu - \frac{\sigma^2}{2}\right)dt + \sigma dW(t).$$

Integrating both sides from time 0 to time t yields

$$\ln(\pi(t)) = \ln(\pi(0)) + \left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W(t).$$

By taking the exponential of both sides, we can explicitly derive the market price $\pi(t)$ as

$$\pi(t) = \pi(0) \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W(t)\right). \quad (2)$$

(2) is the solution of the SDE (1) governing the market price $\pi(t)$ and implies that given a positive initial market price $\pi(0)$, if $\mu < \frac{\sigma^2}{2}$ holds, we have $\pi(t) \rightarrow 0$ as $t \rightarrow \infty$; instead if $\mu > \frac{\sigma^2}{2}$ holds, we have $\pi(t) \rightarrow \infty$ as $t \rightarrow \infty$. Based on empirical observations, obviously, the former case does not reflect the long-term evolution of real-world electricity market prices. Therefore, we will base all our discussions on the parameterization $\pi(0) > 0$ and $\mu > \frac{\sigma^2}{2}$.

²While the electricity demand of a data center is enormous, it still only accounts for a tiny portion of the total demand in a national-level or state-level electricity market. Therefore, we implicitly assume data centers are price takers.

Further, the expectation of the market price at any time t can be readily computed from (2):

$$\begin{aligned} \mathbb{E}[\pi(t)] &= \mathbb{E}\left[\pi(0) \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W(t)\right)\right] \\ &= \pi(0) \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)t\right) \mathbb{E}[\exp(\sigma W(t))] \\ &= \pi(0) \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)t\right) \exp\left(\frac{\sigma^2 t}{2}\right) \\ &= \pi(0)e^{\mu t}, \end{aligned} \quad (3)$$

where the third equality follows from the property of the Wiener process $W(t)$ that is normally distributed with $\mathbb{E}[W(t)] = 0$ and $\text{Var}[W(t)] = t$.

Denote the electricity demand of the data center at time t as $D(t)$. Suppose the data center signs a PPA specified by the quadruple (τ, T, k, p) . During the time window $[\tau, \tau + T]$, whenever $D(t) > k$ holds, the data center needs to purchase the shortfall $D(t) - k$ from the electricity market at the market price $\pi(t)$, leading to an instantaneous cost of $\pi(t)(D(t) - k)$. Whenever $D(t) < k$ holds, the data center can sell the surplus $k - D(t)$ to the electricity market also at the market price $\pi(t)$. This earns the data center an instantaneous revenue of $\pi(t)(k - D(t))$. Note that outside the time window $[\tau, \tau + T]$ or if no PPA is signed, all the electricity demand of the data center has to be met with the supply from the electricity market, which implies an instantaneous cost $\pi(t)D(t)$ at time t .

C. Problem Formulation

The goal of optimizing PPA signing is to minimize the expected total electricity expenditure of the data center over the entire time horizon. However, as the time span considered in our work could span decades, it is important to account for currency depreciation due to inflation. To this end, we introduce a constant discount rate $\lambda \in (0, 1)$ to capture the change in monetary values across time:

$$\frac{dX(t)}{dt} = \lambda X(t), \quad (4)$$

where $X(t)$ denotes the value of a certain amount of money at time t . Consider the starting time τ of a PPA, it can be readily verified that the solution of (4) satisfies

$$X(\tau) = X(t)e^{-\lambda(t-\tau)},$$

which basically reflects the equivalent value at time τ for the amount of money that is worth $X(t)$ at any time t . On this basis, we are ready to formulate the problem of signing an optimal PPA.

In the presence of a PPA (τ, T, k, p) , the total electricity expenditure of the data center during the time window $[\tau, \tau + T]$ when the PPA is valid, adjusted to the equivalent value at time τ , can be explicitly given by

Expenditure with PPA

$$\begin{aligned}
 & \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) \max\{0, D(t) - k\} dt \\
 & - \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) \max\{0, k - D(t)\} dt \\
 & + pkT \\
 & = \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) (D(t) - k) dt + pkT,
 \end{aligned} \tag{5}$$

which consists of the (discounted) total cost of market purchase in case of shortfall, the (discounted) total revenue of market selling in case of surplus, and the PPA signing cost. The equality in (5) holds since market purchase and selling do not occur simultaneously for the data center.

In the absence of the PPA, the electricity demand of the data center is all fulfilled in the electricity market. The corresponding (discounted) total electricity expenditure is

Expenditure without PPA

$$\int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) D(t) dt. \tag{6}$$

As we explore the benefits of signing the PPA, it is more convenient to focus on the total saving, defined as

Expenditure without PPA – Expenditure with PPA

$$\begin{aligned}
 & = \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) D(t) dt \\
 & - \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) (D(t) - k) dt - pkT \\
 & = k \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) dt - pkT.
 \end{aligned} \tag{7}$$

As far as the PPA is concerned, the total saving (7) is an equivalent indicator for the total electricity expenditure since the benchmark expenditure without the PPA is for sure independent of PPA signing. While the PPA is not valid outside the time window $[\tau, \tau + T]$, when to sign it into effect, i.e., τ , still plays a key role in improving the total saving.

Given the uncertainty in future market prices $\pi(t)$, the problem of signing the PPA can be explicitly formulated as a stochastic program:

$$\max_{(\tau, T, k, p)} \mathbb{E} \left[e^{-\lambda\tau} \cdot \left(k \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) dt - pkT \right) \right], \tag{8}$$

where the discount $e^{-\lambda\tau}$ is included to fairly evaluate the saving at any possible starting time τ .

The PPA signing problem (8) is still, in general, difficult to solve. However, in the next section, we propose to decompose it in a tractable manner such that the data center tackles a static PPA design problem first by optimizing (T, k, p) only for any given τ , and then determines the starting time τ from a dynamic online perspective.

III. ANALYSIS AND RESULTS

In this section, we first analyze for a given starting time τ of the PPA how to set the contract term T , the contractual

power supply k , and the locked-in price p that lead to the most expected total saving. We then exploit this insight to develop a dynamic threshold policy that online identifies an optimal starting time τ using real-time observations of market prices.

A. How to Design a PPA

Suppose at a given starting time τ , the data center signs a PPA into effect. Define

$$S(\pi(\tau)) := \mathbb{E} \left[k \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} \pi(t) dt - pkT \mid \pi(\tau) \right]$$

to be the expected total saving of the PPA (evaluated at time τ) conditional on the observed market price $\pi(\tau)$. The following lemma characterizes its expected total saving.

Lemma 1. *The expected total saving of a PPA starting from time τ , conditional on the observed market price $\pi(\tau)$, can be explicitly expressed as*

$$S(\pi(\tau)) = \frac{k\pi(\tau)}{\mu - \lambda} (e^{(\mu - \lambda)T} - 1) - pkT. \tag{9}$$

Proof. A similar argument to (2) will lead to

$$\pi(t) = \pi(\tau) \exp \left(\left(\mu - \frac{\sigma^2}{2} \right) (t - \tau) + \sigma(W(t) - W(\tau)) \right).$$

When $\pi(\tau)$ is observed, for $t \in [\tau, \tau + T]$ we have

$$\mathbb{E}[\pi(t) \mid \pi(\tau)] = \pi(\tau) e^{\mu(t-\tau)}, \tag{10}$$

following (3). Then it immediately leads to

$$\begin{aligned}
 S(\pi(\tau)) & = k\pi(\tau) \int_{\tau}^{\tau+T} e^{-\lambda(t-\tau)} e^{\mu(t-\tau)} dt - pkT \\
 & = k\pi(\tau) \int_{\tau}^{\tau+T} e^{(\mu - \lambda)(t-\tau)} dt - pkT \\
 & = \frac{k\pi(\tau)}{\mu - \lambda} (e^{(\mu - \lambda)T} - 1) - pkT.
 \end{aligned} \tag{11}$$

□

Given Lemma 1, we first analyze particularly the impact of the contract term T on the expected total saving S under two scenarios: $\mu > \lambda$ and $\mu < \lambda$. Note that $T = 0$ immediately implies $S = 0$ and $\frac{\partial S}{\partial T} = k(\pi(\tau) - p)$ from (9). The following discussion is predicated on $k > 0$.

1) $\mu > \lambda$, i.e., the growth rate of market prices is larger than the currency discount rate.

- $p < \pi(\tau)$: Note that $\frac{\partial S}{\partial T}$ is positive for any $T \geq 0$. Therefore, the PPA always yields a positive expected total saving that is increasing in the contract term $T > 0$.

- $p > \pi(\tau)$: Note that $\frac{\partial S}{\partial T}$ increases in T , starting negative but growing to be positive with $T > \frac{\ln p - \ln \pi(\tau)}{\mu - \lambda}$. Therefore, the expected total saving initially decreases in T and becomes negative. It reaches the minimum value with $T = \frac{\ln p - \ln \pi(\tau)}{\mu - \lambda}$. Then the expected total saving changes to increase in T and become positive with $T > T_{\text{critical}}$, where $T_{\text{critical}} > 0$ is the solution to $S(\pi(\tau)) = 0$.

In case of $\mu > \lambda$, a short message from above is to pursue the longest possible contract term for the PPA. This is consistent with the intuition when electricity market prices are growing at a fast pace.

2) $\mu < \lambda$, i.e., the currency discount rate is larger than the growth rate of market prices.

- $p < \pi(\tau)$: Note that $\frac{\partial S}{\partial T}$ decreases in T , starting positive but growing to be negative with $T > \frac{\ln p - \ln \pi(\tau)}{\mu - \lambda}$. Therefore, the expected total saving initially increases in T and becomes positive. It reaches the maximum value with $T = \frac{\ln \pi(\tau) - \ln p}{\lambda - \mu}$. Then the expected total saving changes to decrease in T and become negative with $T > T_{\text{critical}}$, where $T_{\text{critical}} > 0$ is the solution to $S(\pi(\tau)) = 0$. In this case, it is risky to sign the PPA with excessively long contract terms, despite $p < \pi(\tau)$.
- $p > \pi(\tau)$: Note that $\frac{\partial S}{\partial T}$ is negative for any $T \geq 0$. Therefore, the PPA always yields a negative expected total saving that is decreasing in the contract term $T > 0$. In this case, it is better off without the PPA.

The above analysis can be summarized into the following proposition:

Proposition 1. Suppose the data center signs a PPA (τ, T, k, p) at time τ upon observing the market price $\pi(\tau)$.

- In case of $\mu > \lambda$, the conditional expected total saving of the PPA is positive if either (i) $p < \pi(\tau)$ or (ii) $p \geq \pi(\tau)$ and $T > T_{\text{critical}}$. Moreover, it increases in T and k , while decreasing in p .
- In case of $\mu < \lambda$, the optimal contract term is given by

$$T^* = \begin{cases} \frac{\ln \pi(\tau) - \ln p}{\lambda - \mu}, & \text{if } p < \pi(\tau), \\ 0, & \text{if } p \geq \pi(\tau). \end{cases} \quad (12)$$

The corresponding conditional expected total saving is

$$S(\pi(\tau)) = \begin{cases} \frac{kp - k\pi(\tau) - pk(\ln p - \ln \pi(\tau))}{\mu - \lambda}, & \text{if } p < \pi(\tau), \\ 0, & \text{if } p \geq \pi(\tau). \end{cases} \quad (13)$$

which, in case of $p < \pi(\tau)$, increases in k , while decreasing in p .

Note that the locked-in price p also plays a key role. First, T^* decreases in the locked-in price p in case of $p < \pi(\tau)$. Intuitively, it shall be worth extending the contract term of a PPA with a lower locked-in price. Second, from (9) there exists an upper-bound market price \bar{p} such that as long as $p \leq \bar{p}$, the PPA is beneficial, i.e., $S(\pi(\tau)) \geq 0$. \bar{p} can be noted down explicitly as

$$\bar{p} = \frac{\pi(\tau)[e^{(\mu-\lambda)T} - 1]}{(\mu - \lambda)T}, \quad (14)$$

which leads to more insights:

- 1) $\mu > \lambda$: \bar{p} increases in T . This also explains why in the case of $p \geq \pi(\tau)$, the PPA becomes beneficial to cut expenditure only when the contract term T exceeds the critical threshold T_{critical} .
- 2) $\mu < \lambda$: \bar{p} decreases in T . In particular, in the case of $p \geq \pi(\tau) > \bar{p}$, the PPA cannot achieve any saving.

Therefore, the locked-in price p has to be set below the market price $\pi(\tau)$ when the PPA is signed into effect.

B. When to Sign a PPA

We now turn to the second part of the PPA signing problem (8), where in practice the data center has to identify online the optimal timing to sign and start a PPA, as market prices are observed sequentially.

We are inspired by the classical problem of optimal stopping to weigh signing a PPA immediately against taking a wait-and-see stance. In particular, given the initial market price $\pi(0) = \omega$, denote the optimal expected total saving of a PPA in (8) by

$$\phi(\omega) := \max_{\tau} \mathbb{E} [e^{-\lambda\tau} S(\pi(\tau)) \mid \pi(0) = \omega], \quad (15)$$

where we assume $S(\pi(\tau))$ has been optimized with proper choice of (T, k, p) , as discussed in Section III-A. $\phi(\omega)$ is essentially the optimal value function of PPA signing and corresponds to an optimal starting time τ^* of the PPA. Moreover, τ^* can be more explicitly characterized by the following lemma.

Lemma 2. There is a unique ω^* such that the optimal starting time of the PPA is $\tau^* = \inf\{\tau \geq 0 : \pi(\tau) \geq \omega^*\}$.

See [24] for the proof. Lemma 2 basically suggests that there exists a unique threshold ω^* such that whenever the market price $\pi(t)$ first goes beyond this threshold, it is the optimal timing to sign and start the PPA. Next, we focus on solving for ω^* .

Since the market price behavior follows a GBM, Ito's Lemma requires the optimal value function $\phi(\omega)$ to satisfy the Hamilton–Jacobi–Bellman (HJB) equation:

$$\phi(\omega) = \max\{S(\omega), \frac{1}{\lambda} \mu \omega \phi'(\omega) + \frac{1}{2\lambda} \sigma^2 \omega^2 \phi''(\omega)\}. \quad (16)$$

A more detailed derivation of (16) can be found in Appendix A. It implies that the optimal value function weighs the immediate payoff (the expected total saving of the PPA) from signing the PPA now against the expected payoff from taking a wait-and-see stance. Based on (16), it can be inferred that in the case of $\phi(\omega) > S(\omega)$, i.e., the expected wait-and-see payoff is higher than the immediate payoff, the data center is better off skipping the current signing opportunity. Instead, in case of $\phi(\omega) = S(\omega)$, it indicates the optimal timing to start the PPA immediately.

Recall that $S(\omega)$ is an increasing function of ω from (9). Therefore, the threshold ω^* should satisfy

$$S(\omega^*) = \frac{1}{\lambda} \mu \omega^* \phi'(\omega^*) + \frac{1}{2\lambda} \sigma^2 \omega^{*2} \phi''(\omega^*), \quad (17a)$$

along with the regularity conditions

$$S(\omega^*) = \phi(\omega^*), \quad (17b)$$

$$S'(\omega^*) = \phi'(\omega^*). \quad (17c)$$

Notably, (17b) enforces value matching while (17c) enforces smooth pasting at the threshold ω^* . More specifically, the threshold ω^* is characterized in the following proposition.

Proposition 2. The threshold ω^* satisfying (17) is given by

$$\omega^* = \frac{pmT(\mu - \lambda)}{[e^{(\mu - \lambda)T} - 1](m - 1)}, \quad (18)$$

with

$$m := \frac{(\sigma^2 - 2\mu) + \sqrt{(2\mu - \sigma^2)^2 + 8\lambda\sigma^2}}{2\sigma^2}. \quad (19)$$

The proof is provided in Appendix III-B. Combining Proposition 1 and Proposition 2, we can further pin down the threshold ω^* in the following corollary:

Corollary 1. In case of $\mu < \lambda$ and $p < \pi(\tau)$, the threshold ω^* satisfying (12) and (18) is the solution to

$$(p - \omega^*)(m - 1) - pm(\ln p - \ln \omega^*) = 0. \quad (20)$$

We summarize the timing strategy to start a PPA in the following proposition:

Proposition 3. Suppose the data center observes the market price $\pi(t)$ sequentially.

- In case of $\mu > \lambda$, it is optimal to start the PPA immediately, i.e., $\phi(\omega) = S(\omega)$.
- In case of $\mu < \lambda$, the data center should sign and start the PPA whenever the market price $\pi(t)$ first exceeds the threshold ω^* given by (18). In this case, the optimal value function $\phi(\omega)$ depends on $S(\omega)$ via

$$\phi(\omega) = \begin{cases} e^{-\lambda\tau^*} S(\omega^*), & \text{if } \pi(0) < \omega^*, \\ S(\omega), & \text{if } \pi(0) \geq \omega^*. \end{cases} \quad (21)$$

When the growth rate μ of market prices is larger than the currency discount rate λ , we have $m < 1$ in (19), which guarantees that $\omega^* \leq 0$ always stays below $\pi(t)$ for $\forall t \geq 0$. Therefore, it is optimal to start the PPA immediately. This is consistent with Proposition 1 that suggests the longest possible contract term is preferred in this case. On the contrary, $\mu < \lambda$ leads to $m > 1$, and the data center should only proceed to start the PPA when $\pi(t) \geq \omega^*$ first occurs. Due to the continuity of $\phi(\omega)$ and $S(\omega)$ as required in (17), the optimal starting time τ^* in Lemma 2 reduces to $\tau^* = \inf\{t \geq 0 : \pi(t) = \omega^*\}$, meaning that it is optimal to start the PPA whenever the threshold ω^* is reached.

Corollary 2. If the data center starts the PPA at the optimal starting time τ^* , the conditional expected total saving $S(\pi(\tau^*))$ is given by

$$S(\omega^*) = \frac{k(p - \omega^*)}{(\mu - \lambda)m}, \quad (22)$$

with the corresponding optimal value function

$$\phi(\omega) = \frac{k(p - \omega^*)}{(\mu - \lambda)m\omega^{*m}} \omega^m. \quad (23)$$

Finally, we outline the key steps to implement the proposed PPA signing strategy in Algorithm 1 below.³

³While we adopt a continuous-time model for analysis and design, in practice, decisions are made in discrete time.

Algorithm 1: PPA Signing Strategy

```

Data:  $\mu, \sigma, \lambda$ 
Result:  $(\tau, T, k, p)$ 
 $t \leftarrow 0$ ;
 $\tau \leftarrow -1$ ; /* PPA is not signed yet */
while  $\tau \neq t$  do
   $k \leftarrow$  maximum possible ;
   $p \leftarrow$  minimum possible ;
  if  $\mu > \lambda$  then
     $\tau \leftarrow t$ ; /* Start PPA now */
     $T \leftarrow$  maximum possible ;
  else
    Observe  $\pi(t)$  ;
    if  $\pi(t) \leq p$  then
       $t \leftarrow t + 1$ ; /* Defer PPA */
    else
      Compute  $\omega^*$  by solving (20) ;
      if  $\pi(t) \geq \omega^*$  then
         $\tau \leftarrow t$ ;
         $T \leftarrow T^*$  based on (12) ;
      else
         $t \leftarrow t + 1$  ;
      end
    end
  end
end

```

IV. SIMULATION

In this section, we demonstrate the proposed strategy for PPA signing by assessing total saving using real-world electricity market prices from New York ISO spanning 2015 to 2025 [25]. Based on the data, we estimate for market prices an annual growth rate $\mu = 0.15$ and an annual volatility constant $\sigma = 0.48$. The initial market price is set to $\pi(0) = 38\$/MWh$. The currency discount rate λ is fixed at 0.05. Monte Carlo experiments of 3000 simulation runs are carried out for each test result to show its statistical pattern. Further, a series of sensitivity analyses is conducted to illustrate the robustness of the test results by varying key parameters, including the GBM constants μ and σ , the PPA locked-in price p , and the currency discount rate λ .

Base case with $\lambda = 0.05$ and $p = 36\$/MWh$, i.e., $\mu > \lambda$ and $p < \pi(0)$. We show particularly how the total saving of the PPA changes with respect to the contract term T , along with the impact of μ and σ in Fig. 2 and Fig. 3, respectively. The solid lines represent the sample averages of the total saving due to PPA signing over 3000 simulation runs, while the shaded areas represent the interquartile range (IQR) between the 25th and 75th percentiles. In both figures, the expected total saving exhibits approximately exponential growth in terms of the contract term T , which is consistent with our analysis that suggests starting the PPA immediately with the longest possible contract term. Moreover, it is observed that both an increase in μ and σ leads to a larger expected total saving. The former is straightforward due to the faster growth of market prices. The latter is empirical due to limited samples but illustrates the




Fig 2. Total saving with respect to contract term T under different μ 's ($\sigma = 0.48$, $p = 36\$/MWh$, $\lambda = 0.05$).




Fig 4. Total saving with respect to contract term T under different μ 's ($\sigma = 0.48$, $p = 55\$/MWh$, $\lambda = 0.05$).




Fig 3. Total saving with respect to contract term T under different σ 's ($\mu = 0.15$, $p = 36\$/MWh$, $\lambda = 0.05$).




Fig 5. Total saving with respect to contract term T under different σ 's ($\mu = 0.15$, $p = 55\$/MWh$, $\lambda = 0.05$).

benefit of the PPA in hedging against price volatility. Notably, the IQR of the total saving of the PPA grows rapidly over time, reflecting the fact that the uncertainty (variance) in the logarithm of electricity market prices increases linearly in time for a GBM, as shown in (2).

Base case with $\lambda = 0.05$ and $p = 55\$/MWh$, i.e., $\mu > \lambda$ and $p > \pi(0)$. For comparison, Fig. 4 and Fig. 5 display the impact of the relative relation between p and $\pi(0)$. In this case, the expected total saving starts with a decline but soon switches to grow exponentially after the contract term T exceeds $\frac{\ln p - \ln \pi(0)}{\mu - \lambda} \approx 3.7$ years (Fig. 5). This implies that signing a PPA with short contract terms ($T < T_{\text{critical}}$) may not be economically beneficial. Additionally, it can be observed that T_{critical} decreases in μ and σ (empirical observation only), which is aligned with the previous analysis for the expected total saving.

Artificial cases with $\lambda = 0.16$, i.e., $\mu < \lambda$. To further show the impact of the relative relation between μ and λ , we consider a set of artificial cases where λ is increased to 0.16 while the rest parameters remain the same. In case of $p = 36\$/MWh < \pi(0)$, as shown in Fig. 6 and Fig. 7, the PPA achieves the expected total saving that first grows and then declines, and could even become negative when the contract term T is too long. It can be similarly observed that the maximum is achieved by the optimal contact term $T^* = \frac{\ln \pi(0) - \ln p}{\lambda - \mu} \approx 5.4$ years (Fig. 7). Further, Fig. 6 suggests that T^* increases in μ , which is aligned with both the analysis and the intuition that faster-growing market prices would incentivize the data

center to pursue longer contract terms. For completeness, we also test the case with $p = 55\$/MWh > \pi(0)$. Fig. 8 and Fig. 9 basically echo our analysis that it is better off without the PPA in this case.

A summary of the sensitivity analysis for the base cases is presented in Table II, where the contract term is set to $T = 8$ years – the longest considered in the experiments. In general, the expected total saving seems more sensitive to μ than to σ , particularly reflected by the relative change with respect to the benchmark case of $\mu = 0.15$ and $\sigma = 0.48$. For instance, in the $p < \pi(0)$ case, a modest increase in μ by 0.005 leads to a 9.09% increase in the expected total saving. The relative saving provides a more straightforward measure of cost reduction with the PPA. Obviously, if $p < \pi(0)$ can be achieved in negotiation, the benefit of the PPA will be tremendously enhanced. The median of total savings provides additional insights into its positively skewed distribution under uncertainty, i.e., the distribution is biased towards small or even negative total savings. It suggests that while the expected total saving could be promising, there are hidden risks unaccounted for in our work. We would like to leave risk-averse PPA signing strategies for future studies.

V. CONCLUSION

This paper studies the problem of how to sign a PPA for a data center with the most expected total savings. We first establish a novel stochastic programming formulation for such a problem by modeling PPA decisions using a quadruple

TABLE II
SENSITIVITY ANALYSIS FOR TOTAL SAVING OF PPA

Case	μ	σ	Mean (Million\$)	Relative Change	Relative Saving	Median (Million\$)
$p < \pi(0)$	0.15	0.48	172.48	—	37.46%	24.10
	0.145	0.48	159.00	-7.81%	35.57%	24.80
	0.155	0.48	188.16	9.09%	39.52%	36.58
	0.15	0.475	169.41	-1.78%	37.04%	32.81
	0.15	0.485	177.39	2.85%	38.12%	28.28
$p > \pi(0)$	0.15	0.48	20.48	—	4.45%	-127.90
	0.145	0.48	7.00	-65.79%	1.57%	-127.20
	0.155	0.48	36.16	76.61%	7.59%	-115.42
	0.15	0.475	17.41	-14.97%	3.81%	-119.19
	0.15	0.485	25.39	23.98%	5.46%	-123.72




Fig 6. Total saving with respect to contract term T under different μ 's ($\sigma = 0.48$, $p = 36\$/MWh$, $\lambda = 0.16$).




Fig 8. Total saving with respect to contract term T under different μ 's ($\sigma = 0.48$, $p = 55\$/MWh$, $\lambda = 0.16$).




Fig 7. Total saving with respect to contract term T under different σ 's ($\mu = 0.15$, $p = 36\$/MWh$, $\lambda = 0.16$).




Fig 9. Total saving with respect to contract term T under different σ 's ($\mu = 0.15$, $p = 55\$/MWh$, $\lambda = 0.16$).

and long-term evolution of electricity market prices using a GBM. We then propose an integrated PPA signing strategy to tackle the problem: (i) given an arbitrary starting time, we analytically set the corresponding optimal contract term, contractual power supply, and locked-in price; (ii) we develop a threshold policy to dynamically pick an optimal starting time based on real-time observations of market prices. Our analysis leads to several novel results that offer analytical guidelines for practical PPA signing. In particular, we identify scenarios in which a PPA is not profitable – an important message to avoid signing inadvisable PPAs. Our dynamic threshold policy is friendly to implement – computing a threshold online and comparing it with the real-time market price. Our numerical results provide additional insights – different parameter sensi-

tivity and potential risk due to the skewed distribution of PPA performance.

Limitations and future directions. There are several limitations in our work that open avenues for future studies. First, a more general PPA model beyond the quadruple (τ, T, k, p) would be closer to contractual practices. Second, it becomes technically more challenging in the presence of an unreliable power supply from a PPA. Third, a risk-averse strategy to sign PPAs would be more attractive under asymmetric uncertainty. Last, an extension to managing multiple PPAs along with other asset options via portfolio optimization would have practical implications for a return-risk tradeoff.

REFERENCES

[1] South Pole. Going 100% renewable: How companies are demanding a

faster market response. [Online] Available: <https://www.southpole.com/publications/re100-annual-report-2019>.

[2] S. Afkaki and S. Netessine, "Strategic investment in renewable energy sources: The effect of supply intermittency," *Manufacturing & Service Operations Management*, vol. 19, no. 3, pp. 489–507, 2017.

[3] P. Gabrielli, R. Aboutalebi, and G. Sansavini, "Mitigating financial risk of corporate power purchase agreements via portfolio optimization," *Energy Economics*, vol. 109, pp. 980–991, 2022.

[4] M. Bruck, P. Sandborn, and N. Goudarzi, "A levelized cost of energy model for wind farms that include power purchase agreements," *Renewable Energy*, vol. 122, pp. 131–139, 2018.

[5] Bloomberg NEF. Corporate clean power buying grew 12% to new record in 2023. [Online] Available: <https://about.bnef.com/blog/corporate-clean-power-buying-grew>.

[6] District Energy. Microsoft signs deal to invest more than \$10 billion on renewable energy capacity to power data centers. [Online] Available: <https://www.districtenergy.org/blogs/district-energy/2024/05/08/microsoft-signs-deal-to-invest-more-than-10-billio>.

[7] L. Mendicino, D. Menniti, A. Pinnarelli, and N. Sorrentino, "Corporate power purchase agreement: Formulation of the related levelized cost of energy and its application to a real life case study," *Applied Energy*, vol. 253, pp. 113577–113589, 2019.

[8] M. Bruck and P. Sandborn, "Pricing bundled renewable energy credits using a modified LCOE for power purchase agreements," *Renewable Energy*, vol. 170, pp. 224–235, 2021.

[9] B. Kandpal, S. Backe, and P. C. del Granado, "Power purchase agreements for plus energy neighbourhoods: Financial risk mitigation through predictive modelling and bargaining theory," *Applied Energy*, vol. 358, pp. 122589–122592, 2024.

[10] S. Chen and X. Hou, "Collective renewable power purchase agreement: Negotiation and risk-sharing mechanisms," *SSRN 4903084*, 2024.

[11] D. P. Neto, E. G. Domingues, A. P. Coimbra, A. T. de Almeida, A. J. Alves, and W. P. Calixto, "Portfolio optimization of renewable energy assets: Hydro, wind, and photovoltaic energy in the regulated market in Brazil," *Energy Economics*, vol. 64, pp. 238–250, 2017.

[12] J. Schmidt, R. Cancella, and A. O. Pereira Jr, "An optimal mix of solar PV, wind and hydro power for a low-carbon electricity supply in Brazil," *Renewable Energy*, vol. 85, pp. 137–147, 2016.

[13] K. Fedorová, T. Ábelová, and M. Kvasnica, "Dynamic power purchase agreement," in *International Conference on Process Control (PC)*, 2023, pp. 156–161.

[14] S. Hu, G. C. Souza, M. E. Ferguson, and W. Wang, "Capacity investment in renewable energy technology with supply intermittency: Data granularity matters!" *Manufacturing & Service Operations Management*, vol. 17, no. 4, pp. 480–494, 2015.

[15] Z. Gao, N. Sunar, and J. R. Birge, "Power purchase agreements with renewables: Optimal timing and design," in *IEEE Conference on Decision and Control (CDC)*, 2023, pp. 7580–7585.

[16] Greentech Media. It's official: Google can sell power like a utility. [Online] Available: <https://www.fastcompany.com/1555267/its-official-google-can-buy-and-sell-energy>.

[17] L. Andreis, M. Flora, F. Fontini, and T. Vargiu, "Pricing reliability options under different electricity price regimes," *Energy Economics*, vol. 87, pp. 104–115, 2020.

[18] S. Borovkova and M. D. Schmeck, "Electricity price modeling with stochastic time change," *Energy Economics*, vol. 63, pp. 51–65, 2017.

[19] R. R. Marathe and S. M. Ryan, "On the validity of the geometric Brownian motion assumption," *The Engineering Economist*, vol. 50, no. 2, pp. 159–192, 2005.

[20] K. Reddy and V. Clinton, "Simulating stock prices using geometric Brownian motion: Evidence from Australian companies," *Australasian Accounting, Business and Finance Journal*, vol. 10, no. 3, pp. 23–47, 2016.

[21] V. Stojkoski, T. Sandev, L. Basnarkov, L. Kocarev, and R. Metzler, "Generalised geometric Brownian motion: Theory and applications to option pricing," *Entropy*, vol. 22, no. 12, pp. 1432–1443, 2020.

[22] C. Chiarella, X. He, and C. S. Nikitopoulos, *Derivative Security Pricing*. Springer, 2016.

[23] R. Bhattacharya and E. C. Waymire, *Continuous Parameter Markov Processes and Stochastic Differential Equations*. Springer Nature, 2023.

[24] E. Mordecki and Y. Mishura, "Optimal stopping for Lévy processes with one-sided solutions," *SIAM Journal on Control and Optimization*, vol. 54, no. 5, pp. 2553–2567, 2016.

[25] New York ISO. Energy market & operational data. [Online] Available: <https://www.nyiso.com>, accessed 2025-07.

APPENDIX

A. Optimality of the HJB Equation

Over a sufficiently small time interval Δt , the data center faces two options: starting a PPA or continuing waiting. Assuming the data center chooses to wait during Δt , and given that $\pi(t)$ follows a GBM, we apply Ito's Lemma to derive the following equation:

$$\phi(\pi(\Delta t)) - \phi(\pi(0)) = \int_0^{\Delta t} \phi'(\pi(t)) d\pi(t) + \int_0^{\Delta t} \frac{1}{2} \phi''(\pi(t)) (d\pi(t))^2.$$

Because $d\pi(t) = \mu\pi(t)dt + \sigma\pi(t)dW(t)$, we have

$$\begin{aligned} \phi(\pi(\Delta t)) &= \phi(\omega) + \int_0^{\Delta t} \phi'(\omega)[\mu\omega dt + \sigma\omega dW(t)] \\ &\quad + \frac{1}{2} \int_0^{\Delta t} \phi''(\omega)[\mu\omega dt + \sigma\omega dW(t)]^2 \\ &= \phi(\omega) + \mu\omega\phi'(\omega)\Delta t + \int_0^{\Delta t} \phi'(\omega)\sigma\omega dW(t) \\ &\quad + \frac{1}{2}\phi''(\omega) \int_0^{\Delta t} [\mu\omega dt + \sigma\omega dW(t)]^2 + o(\Delta t) \\ &= \phi(\omega) + \mu\omega\phi'(\omega)\Delta t + \int_0^{\Delta t} \phi'(\omega)\sigma\omega dW(t) \\ &\quad + \frac{1}{2}\phi''(\omega) \int_0^{\Delta t} \sigma^2\omega^2(dW(t))^2 + o(\Delta t). \end{aligned}$$

Using $(dW(t))^2 = dt$ and

$$\phi(\omega) = \mathbb{E}[e^{-\lambda\Delta t} \phi(\pi(\Delta t)) | \pi(0) = \omega],$$

we have

$$\begin{aligned} \phi(\omega) &= \mathbb{E}[e^{-\lambda\Delta t} (\phi(\omega) + \mu\omega\phi'(\omega)\Delta t + \int_0^{\Delta t} \phi'(\omega)\sigma\omega dW(t))] \\ &\quad + \frac{1}{2}\phi''(\omega) \int_0^{\Delta t} \sigma^2\omega^2(dW(t))^2 + o(\Delta t) \\ &= e^{-\lambda\Delta t} (\phi(\omega) + \mu\omega\phi'(\omega)\Delta t + \mathbb{E}[\int_0^{\Delta t} \phi'(\omega)\sigma\omega dW(t)]) \\ &\quad + \frac{1}{2}\phi''(\omega)\sigma^2\omega^2\Delta t + o(\Delta t)) \\ &= e^{-\lambda\Delta t} (\phi(\omega) + \mu\omega\phi'(\omega)\Delta t + \frac{1}{2}\phi''(\omega)\sigma^2\omega^2\Delta t + o(\Delta t)), \end{aligned}$$

where the last equality follows from

$$\mathbb{E} \left[\int_0^{\Delta t} \phi'(\omega)\sigma\omega dW(t) \right] = 0.$$

We then simplify the expression by noting that $e^{-\lambda\Delta t} = 1 - \lambda\Delta t + o(\Delta t)$ for sufficiently small Δt , and conduct simple algebraic transformation to derive:

$$\phi(\omega) = \phi(\omega) + \left(\mu\omega\phi'(\omega) + \frac{1}{2}\phi''(\omega)\sigma^2\omega^2 - \lambda\phi(\omega) \right) \Delta t + o(\Delta t).$$

Since $\lim_{\Delta t \rightarrow 0} \frac{o(\Delta t)}{\Delta t} = 0$, we can obtain the following differential equation:

$$\phi(\omega) = \frac{\mu\omega}{\lambda} \phi'(\omega) + \frac{\sigma^2\omega^2}{2\lambda} \phi''(\omega).$$

On the other hand, if the data center chooses to stop waiting and initiate a PPA contract immediately, we have

$$\phi(\omega) = S(\omega).$$

Since $\phi(\omega)$ captures the maximum value between the immediate payoff from signing the PPA and the expected payoff from continuing to wait, the HJB equation (16) holds.

B. Proof of Proposition 2

Before determining the threshold ω^* , it is essential to first establish the form of $\phi(\omega)$. Assuming $\phi(\omega) = c \cdot \omega^x$ is a reasonable approach to solving the HJB equation. When $\phi(\omega)$ satisfies the equation $\phi(\omega) = \frac{1}{\lambda} \mu \omega \phi'(\omega) + \frac{1}{2\lambda} \sigma^2 \omega^2 \phi''(\omega)$, differentiating $\phi(\omega)$ yields

$$c\omega^x = \frac{\mu}{\lambda} c x \omega^x + \frac{\sigma^2}{2\lambda} c(x^2 - x) \omega^x.$$

Solving the equation above gives

$$m = \frac{(\sigma^2 - 2\mu) + \sqrt{(2\mu - \sigma^2)^2 + 8\lambda\sigma^2}}{2\sigma^2} > 0,$$

$$n = \frac{(\sigma^2 - 2\mu) - \sqrt{(2\mu - \sigma^2)^2 + 8\lambda\sigma^2}}{2\sigma^2} < 0.$$

Therefore, we obtain two linearly independent particular solutions, i.e. $c\omega^m$ and $c\omega^n$ and we can conclude that

$$\phi(\omega) = c_1 \omega^m + c_2 \omega^n.$$

Since the market price $\pi(\tau)$ follows a GBM, according to (2), we know that $\pi(t) \equiv 0$ holds at any time in case of $\omega = 0$, which implies $\phi(\omega) = 0$. With $w \rightarrow 0$, we have $w^n \rightarrow +\infty$ with $n < 0$, which implies $c_2 = 0$. Therefore, we can obtain

$$\phi(\omega) = c_1 \omega^m.$$

According to (17b) and (17c), we have $mS(\omega) = \omega S'(\omega)$, leading to

$$\omega^* = \frac{pmT(\mu - \lambda)}{[e^{(\mu - \lambda)T} - 1](m - 1)}.$$

In case of $\mu < \lambda$ and $p < \pi(\tau)$, there exists the optimal PPA contract term $T^* = \frac{\ln p - \ln \pi(\tau)}{\mu - \lambda}$. Combining the above yields

$$(p - \omega^*)(m - 1) - pm(\ln p - \ln \omega^*) = 0.$$



Kekun Gao earned his B.S. degree in Computer Science and Technology from Yantai University in 2022 and his M.S. degree in Mechanical Engineering from Peking University in 2025. His research focuses on prediction and optimization.



Yuejun Yan received the Ph.D. degree from the University of Pennsylvania. She is currently the Technical Lead of Alibaba Cloud Global Data Center Sustainability Innovation. She is also the Lead Principal Investigator for Green Cloud Computing Project in ANGEL (Alibaba-NTU Global e-Sustainability CorpLab). Her research interests are in green AI and green cloud computing.



Yuke Zhou received her B.S. degree in theoretical and applied mechanics from Peking University, Beijing, China, in 2022. She is currently pursuing a Ph.D. degree in Industrial and Systems Engineering with Peking University. Her research interests lie in optimization and its applications to power and energy systems and market mechanisms.



Endong Liu received the B.E. degree in Automation from Shandong University, Jinan, China, in 2012. He obtained the Ph.D. degree from the College of Control Science and Engineering, Zhejiang University, Hangzhou, China, where he was also a Member of the Networked Sensing and Control Group. He was a visiting scholar at the College of Engineering, Peking University, from 2022 to 2024. Currently, he is a faculty with the College of Mechanical and Electrical Engineering, Heze University. His research interests include privacy protection and security in smart grids, autonomous navigation of unmanned aerial robots, as well as optimization theory and methods.



Zhaoyang Wang is currently the General Manager of Global Data Center, Alibaba Cloud. He leads global data center planning, delivery, research and development, and operations.



Pengcheng You is an Assistant Professor at the Department of Control Science and Systems Engineering, Peking University. He also holds a joint appointment at the National Engineering Laboratory for Big Data Analysis and Applications, Peking University. Prior to joining PKU, he was a Postdoctoral Fellow at the ECE and ME Departments, Johns Hopkins University. He earned his Ph.D. and B.S. degrees from Zhejiang University, China. During the graduate studies, he was a visiting student at Caltech and a research intern at PNNL. His research interests lie in control, optimization, reinforcement learning, and market mechanisms, with application to power and energy systems.