Published: 31 May 2024

e-Energy '24: The 15th ACM International Conference on Future and Sustainable Energy Systems
June 4 - 7, 2024
Singapore, Singapore

RESEARCH-ARTICLE

Incentives for Accurate Energy Predictions: How to Reduce Epistemic Uncertainty

ROLAND SAUR, Delft University of Technology, Delft, Zuid-Holland, Netherlands

HAN LA POUTRÉ, Delft University of Technology, Delft, Zuid-Holland, Netherlands

NEIL YORKE-SMITH, Delft University of Technology, Delft, Zuid-Holland, Netherlands

e-Energy '24: Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems (June 2024)
https://doi.org/10.1145/3632775.3661956
ISBN: 9798400704802

Incentives for Accurate Energy Predictions: How to Reduce Epistemic Uncertainty

Roland Saur
Centrum Wiskunde & Informatica
IEPG, TU Delft
Amsterdam, The Netherlands
roland@cwi.nl

Han La Poutré
Centrum Wiskunde & Informatica
IEPG, TU Delft
Amsterdam, The Netherlands
Han.La.Poutre@cwi.nl

Neil Yorke-Smith
STAR Lab, TU Delft
Delft, The Netherlands
n.yorke-smith@tudelft.nl

ABSTRACT

Accurate predictions of power fluctuations are pivotal to the operation of flexibility markets. While the design of flexibility markets is an active and ongoing field of research, the question of how to elicit high quality predictions in a non-cooperative setting is often overlooked. Conceptually, we contribute the concept of best prediction incentivizing contracts. Under such contracts the best response of an agent is to report the true distribution of its power fluctuation. This concept differs from Incentive Compatibility by explicitly taking epistemic uncertainty into account: while Incentive Compatible mechanisms often assume the agent possess perfect knowledge of their own valuation, our concept incentivizes agents to reduce their epistemic uncertainty about the world. In practical terms, we present generic closed form solutions for polynomial distributions and show they can be used to approximate realistic Gaussian distributions. Lastly, placing our work in a larger context, we show that third party agents can profit from providing improved predictions via arbitrage.

CCS CONCEPTS

• **Theory of computation** → **Computational pricing and auctions**; • **Hardware** → *Renewable energy*; • **Computing methodologies** → *Multi-agent systems.*

KEYWORDS

Mechanism Design, Pricing, Epistemic Uncertainty, Energy Flexibility, Renewable Energy Integration

ACM Reference Format:
Roland Saur, Han La Poutré, and Neil Yorke-Smith. 2024. Incentives for Accurate Energy Predictions: How to Reduce Epistemic Uncertainty. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24), June 04–07, 2024, Singapore, Singapore.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3632775.3661956

1 INTRODUCTION

Renewable energy sources (RES) are generally seen as the essential solution to global warming. Excessive use of fossil fuels and the accompanying output of carbon dioxide has caused global temperatures to rise [21]. RES have become a carbon-neutral alternative to fossil fuels through advancements in technology as well as investments in mass production [15, 37].

However, large scale integration of RES has brought with it some challenges for electricity systems. In contrast to conventional power plants, most types of RES cannot control their power output. As the share of fluctuating RES increases so does the magnitude and frequency of power imbalances, which need to be addressed by grid operators [18, 24].

To effectively deal with these power imbalances, flexibility has to be integrated into electricity grids [2, 18]. While there is no agreed upon definition of flexibility, it is generally seen as the ability of power producing or consuming devices to quickly adjust their energy output/input [20].

Flexibility can be provided via several different technologies. Conventional gas turbine generators are well established and can provide flexibility by adjusting their operating point [28]. Besides adjustments on the production side, demand response can provide flexibility by adjusting electricity demand [17, 18, 29]. However, flexibility can not only be provided by adjusting production or consumption, but also by shifting energy in different dimensions. Such flexibility can be provided by for instance electrical storage devices [7, 25], which effectively shift energy from times of excess to when energy production is falling short. Given the high cost of electrical storage [4], the idea of shifting energy by utilizing several energy carriers in multi-energy systems has emerged. Flexibility providing multi-energy systems range from combined heat and power generators [42] and the latent heat in residential buildings [9, 17, 23] to tightly integrated industrial processes [3, 43].

While there exists a variety of potential flexibility providers as well as a clear need for it, efficiently integrating this flexibility is still a topic of active research.

Flexibility market mechanisms have emerged as a viable approach for large scale flexibility integration. Direct integration, i.e. equipping fluctuating RES with flexible assets [30] has allowed RES to better participate in conventional electricity markets [1, 32]. However, such integration ties the benefit to one particular fluctuating energy source limiting the overall usefulness of the flexible asset. This limitation can somewhat be addressed by integration of flexibility in larger multi-energy systems [16]. While this adds to the usefulness of flexible assets, it still requires rather complex control systems [16], which cannot be scaled to larger systems, e.g. national electricity markets. Market mechanisms on the other hand are able to coordinate between disparate energy assets leading to a more efficient use of available flexibility.

E-Energy '24, June 04–07, 2024, Singapore, Singapore

ACM ISBN 979-8-4007-0480-2/24/06
https://doi.org/10.1145/3632775.3661956

The operation of flexibility markets rely on accurate predictions. Flexibility trading has been approached both via centralized [39] as well as peer-to-peer markets [39, 47]. It is noteworthy that all of these approaches implicitly assume accurate probability models. This assumption is complicated two-fold: by the difficulty of the task itself as well as the non-cooperative nature of market settings.

The task of forecasting power fluctuations is non-trivial involving different kinds of uncertainty. The need for accurate predictions in RES integration has been noted early on [41] and steady progress has been made over the years on providing improved probabilistic models [12]. However, only in recent years, the power forecasting literature has started addressing the issue of epistemic uncertainty [6, 22, 44] i.e. the uncertainty that exists about the probabilistic models.

Uncertainty can be grouped into aleatoric and epistemic uncertainty. Aleatoric uncertainty is the known uncertainty that can be modeled by probabilistic models. Epistemic uncertainty, on the other hand, represents the unknown uncertainty or lack of knowledge [5, 26]. The epistemic component of the uncertainty can be reduced by learning from additional data [14]. However, since, epistemic uncertainty represents a lack of knowledge it is very difficult to capture in mathematical models, as has been discussed in [14]. We will circumvent this issue by creating a mechanism that incentivizes the reduction of epistemic uncertainty rather than attempting to model it.

The Forecasting issue is further complicated in flexibility markets by the non-cooperative nature of market settings. Within a market setting the question is not just how to generate accurate forecasts but why a self-interested agent should reveal them to the system. This has been explicitly noted by [47] as an open challenge. Recent years have seen the development of data markets [10, 11, 36] as well as the use of prediction markets[33–35] in the energy sector to address this challenge. However, such information markets are disjoint from the physical operation of the energy system and can therefore be more susceptible to speculation.

To address these issues, we consider a peer-to-peer flexibility trading setting that integrates the forecasting problem into the energy contracts by specifically incentivizing the reduction of epistemic uncertainty in the contract generation.

The state of the art is extended in the following way.
We

- properly define our peer-to-peer flexibility trading setting and provide the concept of best prediction incentivizing contracts as a conceptual means to incentivizing the reduction of epistemic uncertainty (Section 2).
- provide best prediction incentivizing contract functions for uniform (Theorem 3.1) and single term polynomial (Theorem 3.2) probability distributions.
- show that our single-term polynomial solution is suitable for modeling real world applications (Section 3.3).
- demonstrate how the arbitrage potential in our solution incentivizes third party agents to contribute improved predictions (Section 4).
- put our work into context by comparing our approach to data markets (Section 5.1) as well as prediction markets (Section 5.2).

2 PROBLEM FORMULATION

We consider a wind farm (WF) with actual power deviation Δq described by true probability density function (pdf) f_T.

To deal with the fluctuation, the WF can obtain an insurance contract $\Theta = (p_0, p_\Delta)$, which will protect it from costly balancing markets. The insurance contract Θ is composed of an upfront cost p_0 and a payment $p_\Delta(\Delta q)$, which depends on the actual power deviation Δq. The balancing cost to the WF of power deviation Δq given a contract Θ can be expressed as:

$$c_B = p_0 + p_\Delta(\Delta q)$$

However, the contract Θ has to be provided by an insurance provider. The contracts this insurance provider is willing to provide depend on the event that the insurance is for, i.e. Θ will depend on f_T. To make this connection explicit, we say the insurance provider uses a contract creating function Ω that accepts a probability distribution f and returns a contract Θ. Let $p_0(.|\Omega(f))$ be the up-front cost p_0 of the contract returned by $\Omega(f)$, while $p_\Delta(\Delta q|\Omega(f))$ is the cost of power deviation Δq given the contract returned by $\Omega(f)$.

Given that the insurance contract Ω depends on f_T, the main research question is how to elicit best approximation of f_T from the agent.

2.1 Best Prediction Incentivizing

In a first step we assume that the WF knows the true probability distribution f_T, while the insurance provider has no information on f_T. The WF reports probability distribution f to contract creating function Ω to generate a contract Θ.

To answer our research question we are looking for a contract creating function Ω to make reporting f_T the best approach for the WF.

The expected cost to the WF of the deviation given that the insurance provider uses contract creating function Ω

$$C_B(f|\Omega, f_T) = p_0(.|\Omega(f)) + \mathbb{E}_{f_T}\left[p_\Delta(\Delta q|\Omega(f))\right] \tag{1}$$

should be minimized when the WF provides the true distribution f_T, i.e.

$$C_B(f_T|\Omega, f_T) \leq C_B(f|\Omega, f_T) \quad \forall f$$

Definition 1 (Incentivizing Best Predictions). *We call a contract creating function Ω or the related pricing functions $(p_0(.|\Omega(f)), p_\Delta(\Delta q|\Omega(f)))$ best prediction incentivizing if*

$$C_B(f_T|\Omega, f_T) \leq C_B(f|\Omega, f_T) \quad \forall f$$

2.2 Best Prediction Incentizing under epistemic uncertainty

So far we assumed that the WF knows the true distribution f_T. However, in reality the WF only has an approximation f of the true distribution f_T. This is because there are two kinds of uncertainty at play here. The true uncertainty described by f_T that is inherent in the event itself. The second kind of uncertainty is called epistemic [26] and is caused not by the event but by a lack of knowledge that is reflected in the discrepancy between f and f_T.

Given that reporting f_T is optimal, there is a certain cost associated with reporting an approximate distribution f. We will refer to this cost as the cost of epistemic uncertainty.

$$CEU(f|\Omega, f_T) = C_B(f|\Omega, f_T) - C_B(f_T|\Omega, f_T)$$

This epistemic uncertainty about the true distribution f_T and its associated cost complicate our earlier definition of best prediction incentivizing contract functions. The earlier definition specifically incentivizes reporting the true distribution f_T. However, it is impossible to incentivize a WF to reliably report something it does not know.

We therefore extend our definition of best prediction incentivizing contract functions to settings with epistemic uncertainty. To do so we will not only require that reporting f_T is optimal, but that reporting a strictly better approximation always results in a lower cost of epistemic uncertainty. This way improving the approximation will always result in lower costs, in essence incentivizing the WF to report better approximations of f_T.

First we need to define what it means for one approximation to be better than the other. To do so we introduce an ordering $\succ_{f_T}$ with respect to the true distribution f_T, where f_T is the dominant element, i.e. $f_T \succ_{f_T} f_1 \quad \forall f_1 \neq f_T$.

Definition 2. *We say probability density function f_1 is a better approximation of f_T than probability density function f_2 if $f_1 \succ_{f_T} f_2$. Several different orderings can reasonable be considered for $\succ_{f_T}$. We will provide two here:*

a) Given an interval $[-d_T, d_T]$ and a class $\mathcal{F}$ of probability density functions on said interval, we say that $f_1 \succ_{f_T} f_2$ if

$$f_1(\Delta q) \in [min(f_T(\Delta q), f_2(\Delta q)), max(f_T(\Delta q), f_2(\Delta q))] \ \forall \Delta q$$

as well as

$$\exists \Delta q \ s.t. \ f_1(\Delta q) \neq f_2(\Delta q)$$

b) Given a parameter d, which defines a unique function in a class $\mathcal{F}$ on the interval $[-d, d]$, we say that $f(d_1) \succ_{f_T} f(d_2)$ if

$$d_1 \in (d_2, d_T] \quad or \quad d_1 \in [d_T, d_2)$$

With this definition, we can extend our definition of Incentivizing Best Predictions to settings with epistemic uncertainty.

Definition 3 (Incentivizing Best Predictions under Epistemic Uncertainty). *We call a contract creating function Ω or the related pricing functions $(p_0(.|\Omega(f)), p_\Delta(\Delta q|\Omega(f)))$ best prediction incentivizing under epistemic uncertainty if reporting a better approximation of f_T always results in a lower epistemic cost, i.e.*

$$CEU(f_1|\Omega, f_T) < CEU(f_2|\Omega, f_T) \qquad \forall f_1 \succ_{f_T} f_2$$

Corollary 1. *Given that Ω is Incentivizing Best Predictions under Epistemic Uncertainty, a contract Θ_1 created by providing f_1 will have a lower expected cost than a contract Θ_2 created by providing f_2 if f_1 is a better approximation of the true distribution f_T, i.e.*

$$C_B(f_1|\Omega, f_T) < C_B(f_2|\Omega, f_T) \qquad \forall f_1 \succ_{f_T} f_2$$

Note 1. *Any pricing function that is Incentivizing Best Predictions under Epistemic Uncertainty is also Incentivizing Best Predictions under perfect knowledge.*

Going forward, we will be referring to best prediction incentivizing contracts/pricing functions (Definition 3) as *BPI.*

2.3 Comparison to Incentive Compatibility

In this section we explore the difference between our concept and the related concept of Incentive Compatibility (IC). In an IC mechanism an agent's best response is to reveal their true valuation to the market [27]. Regarding the information revelation part, IC is similar to BPI contracts. However, in our setting the true valuation of the agent is obscured by epistemic uncertainty. The concept presented here differs from IC by explicitly accounting for the epistemic uncertainty that the agent has about the world.

The epistemic uncertainty stems from the kind of information we are interested in revealing: internal versus external information.

In settings with perfectly known valuations, IC mechanisms incentivizes the revelation of private/internal information to value a good or service. For instance, the value of a vase being sold at an art auction is entirely determined by the private appreciation of the participants. This holds similarly for any kind of commodity or service contract. The assumption, here, being that the agent knows their own true valuation.

Pricing in BPI contracts explicitly aims at revealing external information about the world itself, i.e. the true future power fluctuation. Instead of valuing an object or service, BPI aims to arrive at the true probability distribution of a wind farm. While the argument can be made that the agent knows their true internal valuation, it is very unlikely that they possess perfect knowledge about the world. This inability to perfectly model the future is what we refer to as epistemic uncertainty.

This difference in revealing external instead of internal information has implications for the purpose of pricing. In both concepts, pricing is used as a means to an end. In IC mechanisms without epistemic uncertainty, the end is to value a good or service. In the concept of BPI, pricing is directly linked to finding the correct predictions of an uncertain future event. This connection between pricing and prediction can also be seen in Corollary 1, which states that better predictions result in lower costs. The prices in BPI contracts is therefore directly connected to the amount of epistemic uncertainty.

However, taking epistemic uncertainty into account also comes with challenges for the agent. Honest reports are easy to compute. Therefore, when the true valuation is known to the agent, truth-telling is an easy strategy to implement. However, correctly predicting an external event is a non-trivial task. We will return to this challenge in Section 4, when discussing the arbitrage potential of BPI contracts.

3 APPLICATIONS OF BPI CONTRACTS

So far we have provided BPI as a novel concept and discussed how it relevantly differs from IC. However, a concept is only useful if it can feasibly be applied. To this end, we provide BPI contracts for uniform (Section 3.1) and single term polynomial distributions (Section 3.2) and show that the latter can be used in more real-world settings (Section 3.3). For the particular cases in this section, we will be using a parameter d instead of the entire function f as the input to the contract generating function Ω. Given this input, we will consider the particular ordering of Definition 2b.

3.1 Solutions for uniform distribution

To begin with, we consider a very simple setting.

Setting 1 (Uniform). *The true power deviation of a fluctuating energy source is described by a uniform probability density distribution on the interval $[-d_T, d_T]$. The task of the energy source operator is to provide the correct parameter d_T.*

We want to design prices, $p_0(.|d)$ and $p_\Delta(\Delta q|d)$ for Setting 1 such that the operator minimizes its cost by submitting $d = d_T$, i.e. they are best predictions incentivizing.

Theorem 3.1. *Contract generating function $\Omega_u(d)$ with pricing functions ($p_0(.|d) = 0.5ln(d)$, $p_\Delta(\Delta q|d) = |\Delta q|\frac{1}{d}$) are best BPI functions under Epistemic Uncertainty for Setting 1.*

Proof. To show that the epistemic cost always decreases by moving from a worse approximation d_2 to a better one d_1, we only need to show that epistemic costs are falling for d less than d_T and are rising for d larger than d_T, i.e.:

$$\frac{\partial CEU(d|\Omega_u, d_T)}{\partial d} < 0 \quad \text{if } d \in (0, d_T)$$
$$\frac{\partial CEU(d|\Omega_u, d_T)}{\partial d} > 0 \quad \text{if } d \in (d_T, \infty).$$

This can easily be shown:

$$\begin{aligned}\frac{\partial CEU(d|\Omega_u, d_T)}{\partial d} &= \frac{\partial}{\partial d}C_B(d|d_T, \Omega_u) - \underbrace{\frac{\partial}{\partial d}C_B(d_T|d_T, \Omega_u)}_{=0} \\ &= \frac{\partial}{\partial d}\left[0.5ln(d) + \frac{1}{2}d^{-1}d_T\right] \\ &= 0.5d^{-1}\left[1 - \frac{d_T}{d}\right]\end{aligned}$$

Note that $\left[1 - \frac{d_T}{d}\right] > 0$ for $d > d_T$ and $\left[1 - \frac{d_T}{d}\right] < 0$ for $d < d_T$, which concludes the proof.

□

3.2 Solutions for single term polynomial distributions

In this section, we will be deriving a family of contract generating functions $\Omega(d|k)$ parameterized on a shape parameter k. We consider a single term polynomial as the pdf $f(\Delta q|d)$ on the symmetric interval $[-d, d]$. The general form of f on the positive interval is:

$$f(\Delta q) = -m(\Delta q)^k + t$$

Furthermore, we assume that $f(\Delta q|d)$ is symmetric around $\Delta q = 0$ and $f(d) = 0$. With these two assumptions $f(\Delta q|d)$ is uniquely defined given an agent's report d.

From $f(d) = 0$ it follows that

$$t = md^k$$

Because f is a probability density function and symmetric, we know that:

$$\begin{aligned}\int_0^d f(\Delta q)\partial\Delta q = \frac{1}{2} &= \left[-\frac{m}{k+1}(\Delta q)^{k+1} + t\Delta Q\right]_0^d \\ &= -\frac{m}{k+1}d^{k+1} + t \cdot d\end{aligned}$$

Replacing $t = md^k$ we obtain:

$$\frac{1}{2} = -\frac{m}{k+1}d^{k+1} + md^{k+1}$$

From which follows:

$$m = \frac{1}{2}\frac{k+1}{k}d^{-(k+1)} \wedge t = \frac{1}{2}\frac{k+1}{k}\frac{1}{d}$$

This results in the general from of

$$\begin{aligned}&f_p(\Delta q|d, k) = \frac{1}{2}\frac{k+1}{k}\left[-d^{-(k+1)}(\Delta q)^k + \frac{1}{d}\right] \quad \text{if } \Delta q \in [0, d] \\ &f_p(\Delta q|d, k) = f(-\Delta q) \quad \text{for } \Delta q \in [-d, 0]\end{aligned} \tag{2}$$

Setting 2 (Single Term Polynomial). *The true power deviation of a fluctuating energy source is described by a probability density distribution $f_p(\Delta q|d, k)$ of the shape of equation 2 on the interval $[-d_T, d_T]$. The task of the energy source operator is to provide the correct parameter d_T.*

For Setting 2, we derive the following BPI pricing functions.

Theorem 3.2. *Given a power deviation that can be described by a single term polynomial $f_p(\Delta q|d, k)$ of specified degree k as described in equation (2), contract function Ω_p with pricing functions $p_0(.|d) = \frac{k+1}{2(k+2)}log(d)$ and $p_\Delta(\Delta q) = |\Delta q|\frac{1}{d}$ are BPI pricing functions according to Definition 3.*

Proof. To show that p_0 and p_Δ is BPI pricing functions, we need to show that the expected cost C_B contains a global minimum at $d = d_T$, which we will do by showing that the expected cost increases for $d > d_T$ and decreases for $d < d_T$. We denote f_T as the true distribution $f(\Delta q|d_T, k)$.

First, we calculate the expected cost:

$$
\begin{aligned}
C_B(f,|\Omega f_T) &= \frac{k+1}{2(k+2)} log(d) + \int_{-\infty}^{\infty} f_T(\Delta q)|\Delta q|\frac{1}{d}\partial\Delta q \\
&= \frac{k+1}{2(k+2)} log(d) + 2\int_{0}^{d_T} f_T(\Delta q)\Delta q\frac{1}{d}\partial\Delta q \\
&= \frac{k+1}{2(k+2)} log(d) \\
&\quad + \frac{k+1}{k}\frac{1}{d}\int_{0}^{d_T} -\frac{1}{d_T^{k+t}}(\Delta q)^{k+1} + \frac{1}{d_T}(\Delta q)\partial\Delta q \\
&= \frac{k+1}{2(k+2)} log(d) \\
&\quad + \frac{k+1}{k}\frac{1}{d}\left[-\frac{1}{k+2}\frac{1}{d_T^{k+1}}(\Delta q)^{k+2} + \frac{1}{2d_T}(\Delta q)^2\right]_0^{d_T} \\
&= \frac{k+1}{2(k+2)} log(d) + \frac{k+1}{k}\frac{1}{d}\left[-\frac{1}{k+2}d_T + \frac{1}{2}d_T\right] \\
&= \frac{k+1}{2(k+2)} log(d) + \frac{k+1}{k}\frac{d_T}{d}\left[-\frac{2}{2(k+2)} + \frac{k+2}{2(k+2)}\right] \\
&= \frac{k+1}{2(k+2)} log(d) + \frac{k+1}{2(k+2)}\frac{d_T}{d} \\
&= \frac{k+1}{2(k+2)}\left[log(d) + \frac{d_T}{d}\right]
\end{aligned}
$$

Taking the first derivative of the expected cost C_B with respect to the reported interval d results in:

$$
\begin{aligned}
\frac{\partial}{\partial d}C_B &= \frac{k+1}{2(k+2)}\left[\frac{1}{d} - \frac{d_T}{d^2}\right] \\
&= \underbrace{\frac{k+1}{2(k+2)}\frac{1}{d}}_{>0}\left[1 - \frac{d_T}{d}\right]
\end{aligned}
$$

It is easy to see that $\left[1 - \frac{d_T}{d}\right] > 0$ for all $d > d_T$, while also $\left[1 - \frac{d_T}{d}\right] < 0$ for all $d < d_T$ and therefore the cost increases while $d > d_T$ and decreases while $d < d_T$. This results in a global minimum at $d = d_T$. □

3.3 Real World Applications

While Theorem 3.2 provides us with a solution for a wide range of polynomial distributions, on first glance it seems to be limited when it comes to real-world applications. First, Theorem 3.2 assumes that any market operator knows the particular degree k of the polynomial up front. Secondly, not many problems regarding uncertainty in the energy domain are modeled as single term polynomial distributions.

However, we can fine tune our theoretical solution to be applicable to real world problems. Uncertainty in the energy domain is often modeled as Gaussian distributions[8, 38, 46]. We will, therefore, consider in Setting 3 a truncated Gaussian distribution as the true model of reality. Fortunately, Theorem 3.2 is parameterized on k and, thus, provides an entire family of solutions. Using parameter k, we can choose the solution that best approximates a truncated Gaussian distribution, i.e. we are looking for k^* such that $\Omega_P(d|k^*)$ is best suited for the following Gaussian setting.

Setting 3 (Gaussian). *We will consider reality to be described by a truncated Gaussian distribution $\mathcal{N}(\Delta q|\mu,\sigma)$ [31] with mean $\mu = 0$ and standard deviation σ on a symmetric interval $[-d, d]$.*

This truncated Gaussian is, therefore, completely described by the input parameters d and σ. We scale the standard deviation σ by $0.5d$ to reduce the number of input parameters. The Gaussian setting contains not only the interval length d but also the standard deviation σ as an input parameter. This input parameter does not appear in Setting 2, which we will use for approximation. By scaling $\sigma = 0.5d$, the distribution can be described solely by the input parameter d. Therefore, the input for both this setting as well as Setting 2 is the same, i.e. the interval length d.

This results in a truncated normal probability distribution $\mathcal{N}(\Delta q|d)$ that can be described solely by the interval size d in the following way:

$$
\frac{1}{0.5d}\frac{1}{\Phi^*(d|0,0.5d) - \Phi^*(-d|0,0.5d)}\mathcal{N}^*(\Delta q|0,0.5d), \qquad (3)
$$

where $\mathcal{N}^*(\Delta q|0,0.5d)$ is the untruncated normal distribution of mean $\mu = 0$ and standard deviation $\sigma = 0.5d$ with its corresponding cumulative $\Phi^*(\Delta q|0,0.5d)$ and $\Delta q \in [-d, d]$.

Note, that reducing the number of input parameters does not mean that this setting can only describe a single shape of Gaussian distributions. Figure 1 shows Gaussians from this setting for different values of d.

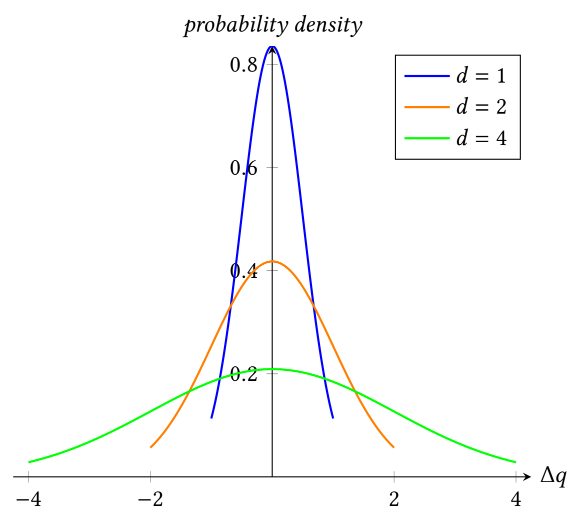

Figure 1: Truncated Gaussians for different reports of d

In the following sub-sections, we will show that the solution from Theorem 3.2 is appropriate for Setting 3. In particular, we will show that our solution can provide a good approximation of a Gaussian distribution (Section 3.3.1) and that the introduced modeling error is especially small when comparing it to the error caused by epistemic uncertainty (Section 3.3.2). Lastly, we consider the possibility of the agent to gain a profit by reporting a value other than d_T and conclude that both the ability and the effect of such misreporting is very limited (Section 3.3.3).

3.3.1 Approximating Gaussian Distributions. While we do not have BPI pricing functions for Setting 3 in closed form, we will leverage our previous result in Theorem 3.2 to derive pricing functions that specifically aim to approximate truncated Gaussian distribution as described in Setting 3. Specifically, we are looking for a value of k such that the probability distribution described in equation (2) best approximates the truncated Gaussian. We measure the quality of our approximation by the root mean-square error [19] over the interval $[-d_T, d_T]$. In this subsection, our task is to find k^* to minimize the error between $\mathcal{N}(\Delta q|d_T)$ and the single term polynomial $f_p(\Delta q|d_T, k)$ as defined in equation (2).

$$Error(\mathcal{N}, f) = \sqrt{\frac{1}{2d_T}\int_{-d_T}^{d_T} [\mathcal{N}(\Delta q|d_T) - f(\Delta q|d_T, k)]^2 \, \partial\Delta q} \quad (4)$$

Figure 2 shows the average error between our approximate solution and the Gaussian distribution over k. The error appears to be minimal for approximately $k = 1.4$. Note that this value is independent of the actual size of the interval d_T as both the Gaussian $\mathcal{N}(\Delta q|d_T)$ and $f_p(\Delta q|d_T, k)$ are being scaled by d_T.

$$k^* \approx \arg\min_k \{Error(\mathcal{N}, f)\} \quad (5)$$

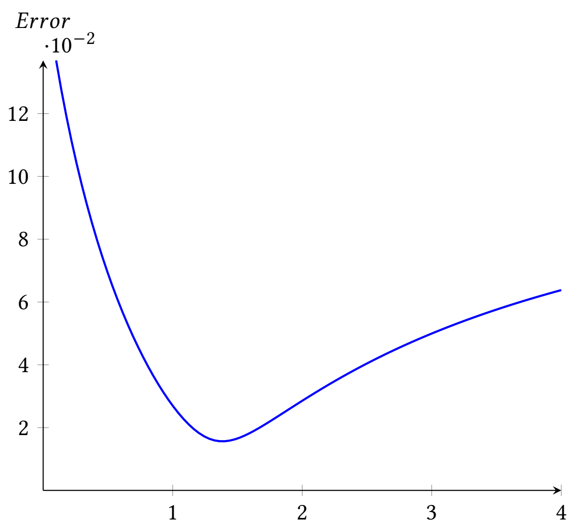

Figure 2: Average Error over k

Figure 3 shows both Gaussian distributions describing the original problem as well as the approximation f_p provided by our approach. It appears that our approach approximates the truncated Gaussian well for a range of values of d.

In subsequent sections, we will use the solution from Theorem 3.2 with $k^* = 1.4$.

3.3.2 Approximation under Epistemic Uncertainty. In this section, we explore the effect of epistemic uncertainty on our approximation and conclude that the error we introduce by using an approximate model of the problem is often insignificant compared to the error caused by epistemic uncertainty. We again consider Setting 3 and the approximation from Theorem 3.2 with $k^* = 1.4$.

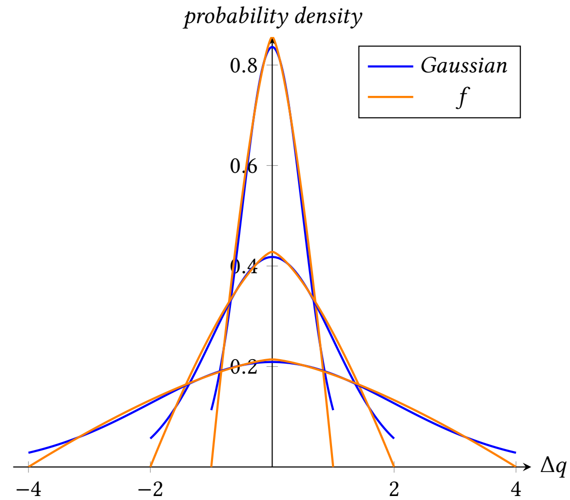

Figure 3: Gaussian and its approximate distribution for $k^* = 1.4$

In this section we will deal with two kinds of errors: the modeling error and the epistemic error. The modeling error is caused by using an approximate model, i.e. single term polynomial instead of a Gaussian. The epistemic error is due to an inaccurate report by the agent. We assume that the agent only knows an approximate value $d \approx d_T$ instead of the true value d_T.

To more clearly highlight the different (modeling and epistemic) contributions to the overall approximation error, we introduce several notions of errors. The total error $Error_{total}$ is simply the overall error caused by providing an approximate value d to function $f_p(\Delta q|d, k^*)$, while the true distribution is the Gaussian $\mathcal{N}(\Delta q|d_T)$.

$$Error_{total} = Error(\mathcal{N}(\Delta q|d_T), f_p(\Delta q|d, k^*)) \quad (6)$$

We define the purely epistemic error as the result of providing an approximate value d to the true model $\mathcal{N}(\Delta q|d)$).

$$Error_{epistemic} = Error(\mathcal{N}(\Delta q|d_T), \mathcal{N}(\Delta q|d)) \quad (7)$$

Given the total error and the purely epistemic error, we introduce the *relative modeling error* as an indication of how much of the total error is due to incorrect modeling. While the true distribution is Gaussian, we are using the polynomial f_p, see equation (2), to approximate it. The error caused by this discrepancy we refer to as the modeling error. We, then, define the relative modeling error $Error_{model}$ as the relative difference between the total and the epistemic error.

$$Error_{model} = \frac{Error_{total} - Error_{epistemic}}{Error_{total}} \quad (8)$$

Figure 4 shows the relative modeling error over reports d, while Figure 5 displays the total error. As is to be expected, the relative modeling error is 1 (i.e. the error is entirely caused by the incorrect model) when $d = d_T$ because then the epistemic error is zero and the entire error is due to the polynomial approximation. However, the relative modeling error quickly drops off as d deviates from the true value of d_T. It is worth noting that this narrow band around d_T, where the modeling error makes a significant contribution, also mostly contains the values of d for which the overall error is small, see Figure 5. Therefore, whenever there is a significant difference between the true distribution and the approximate distribution provided by our solution (Figure 5), most of the error is caused by epistemic uncertainty rather than incorrect model assumptions (Figure 4).

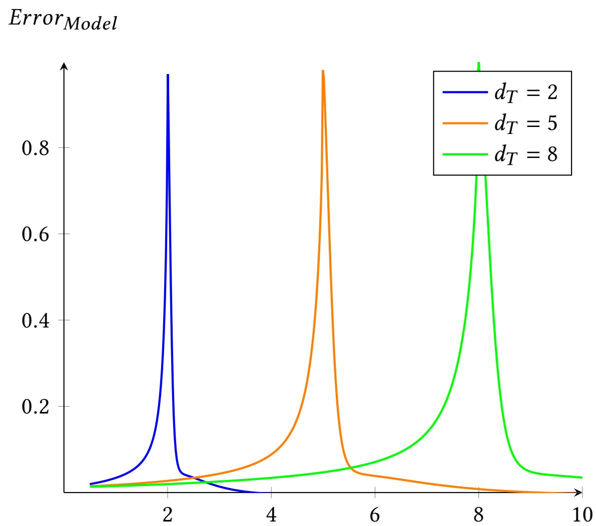

Figure 4: Relative Modeling Error over report d for $k^* = 1.4$

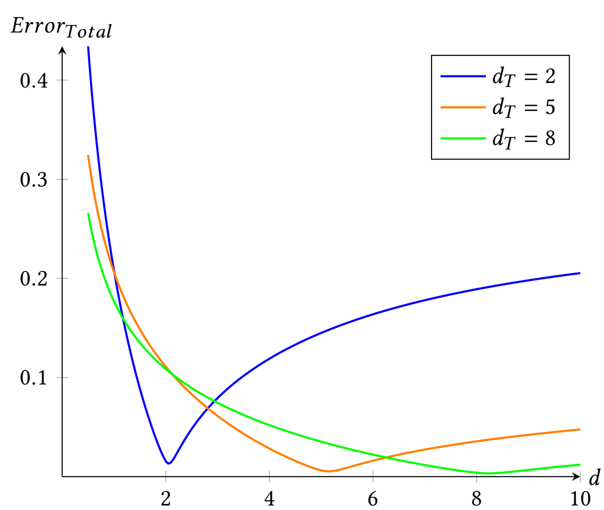

Figure 5: Total Error over Report d

In this section, we showed that our solution not only provides a good approximation of Gaussian distributions, but also that the introduced modeling error is relatively small once epistemic uncertainty is present. However, so far we assumed that the agent is still incentivized to provide accurate predictions, which we will address in the next section.

3.3.3 *Incentives under Epistemic Uncertainty.*

While Section 3.3.2 showed that the modeling error is small compared to the epistemic error, we, implicitly, assumed that the incentive to provide the best prediction of d_T still holds when the underlying true distribution is Gaussian. In this section, we consider Setting 3 and the approximation, i.e. Ω_P and f_p, from Theorem 3.2 with $k^* = 1.4$.

Specifically, we will explore to which extend the agent is still incentivized to provide an accurate estimation of d_T. Theorem 3.2 only specifies pricing functions that are BPI when the true distribution is a single term polynomial of specified degree k. This raises the question of incentives when these pricing functions are used in Setting 3. Reporting the true interval length d_T is no longer proven to be the optimal strategy. Therefore, there is a cost minimizing d^* that may not be equal to d_T.

$$d^* = \arg\min_d C_B\left(d \middle| \Omega_p(d|k^*), \mathcal{N}(\Delta q|d_T)\right) \tag{9}$$

We explore the ramifications of this deviation from truthfully reporting d_T by looking at three metrics related to the strategy, agent cost and the resulting distribution. Regarding the strategy, we consider the absolute value of the relative distance between d^* and d_T, $\left|\frac{d^*-d_T}{d_T}\right|$. For the cost of the agent, we look at the relative cost reduction gained by misreporting, i.e. $\left|\frac{\Delta Cost}{Cost_T}\right|$. Finally, we look at the quality of the resulting distribution by calculating the total error caused by optimal misreporting d^* when using the approximation $f_p(\Delta q|d, k^*)$ as defined in equation (2). This error we will refer to as the error of misreporting $Error_{mis}$.

$$Error_{mis} = Error[\mathcal{N}(\Delta q|d_T), f_p(\Delta q|d^*, k^*)] \tag{10}$$

However, any exploration of these effects is contingent on the ability to obtain the optimal report d^*. Unfortunately, an analytical solution to d^* is difficult to obtain as there is no closed form of the integral of a Gaussian. Luckily, the expected balancing cost $C_B\left(d \middle| \Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ of reporting d to the polynomial contract function $\Omega_p(d|k)$ in Gaussian Setting 3 has a single local minimum with regard to the reported parameter d, see Theorem 3.3. This means that approaches such as gradient descent and bisection method converge quickly to the global minimum d^*.

Theorem 3.3. *When applying contract function $\Omega_p(d|k)$ (Theorem 3.2), to Setting 3, the expected balancing cost $C_B\left(d \middle| \Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ has a single local minimum with regard to the report d.*

Proof. To show that $C_B\left(d \middle| \Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ has a single local minimum in $d > 0$, it suffices to show that its derivative $\frac{\partial}{\partial d} C_B\left(d \middle| \Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ changes sign from negative to positive at a single point. Recall the general cost function, equation (1) as well as the explicit pricing functions in Theorem 3.2.

$$\begin{aligned}
&\frac{\partial}{\partial d} C_B(d|\Omega_P, \mathcal{N}_T) \\
&\quad = \frac{\partial}{\partial d}\left[\frac{k+1}{2(k+2)} log(d) + \int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \frac{1}{d} \partial \Delta q\right] \\
&\quad = \frac{k+1}{2(k+2)} \frac{1}{d} + \frac{\partial}{\partial d} \int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \frac{1}{d} \partial \Delta q \\
&\quad = \frac{k+1}{2(k+2)} \frac{1}{d} + \frac{\partial}{\partial d} \frac{1}{d} \underbrace{\int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \partial \Delta q}_{const>0} \\
&\quad = \frac{k+1}{2(k+2)} \frac{1}{d} - \frac{1}{d^2} \int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \partial \Delta q \\
&\quad = \frac{1}{d}\left[\frac{k+1}{2(k+2)} - \frac{1}{d} \underbrace{\int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \partial \Delta q}_{>0}\right]
\end{aligned}$$

It is easy to see that $\left[\frac{k+1}{2(k+2)} - \frac{1}{d} \underbrace{\int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \partial \Delta q}_{>0}\right]$ changes sign at a single point, while $\frac{1}{d} > 0$ for all $d > 0$, which concludes the proof. □

We numerically explore the effect of reporting optimal d^* as in equation (10) instead of d_T in the context of using $\Omega_p(d|k^*)$ (Theorem 3.2) in Setting 3. Similar to Section 3.3.2, we assume a value of $k^* = 1.4$. We use the scipy optimize package [40] to calculate d^*.

First we consider the relative distance between the optimal report d^* and the true value d_T. Figure 6 displays the relative distance between optimal d^* and true d_T over different values of d_T. We see that the deviation from the true value d_T is generally low and less than three percentage points for a range of values. Note, that in order to reliably report d such that the cost is less than when reporting d_T, the agent needs to know d^* within this three percentage point window. Given that we can assume that the reporting agent does not have perfect knowledge of the future, it stands to reason that they will not be able to take advantage of this difference.

Next, we consider the effect that reporting an untrue but optimal value $d^* \neq d_T$ has on both the agent's cost and the quality of the resulting prediction. In Figure 7a, we see that the relative difference in cost for reporting d^* instead of d_T is very small. This indicates that any effort invested in optimizing for d^* over d_T might ultimately be larger than its gain. Figure 7b depicts the misreporting error as defined in equation 10 when assuming a single term polynomial and optimal report d^* compared to a Gaussian distribution on the true interval $[-d_T, d_T]$. As we can see, the effect on the prediction is less than $3 \cdot 10^{-2}$, indicating that even if the agent reports d^*, the flexibility provider buying the contract is still left with a virtually unchanged prediction.

Therefore, we believe our approach to still incentivize honest behavior: First, effectively deviating from the true report is difficult in practical settings, see Figure 6; Second, the gains of deviating are small, see Figure 7a; And third, the error-effects on the prediction are similarly small, see Figure 7b.

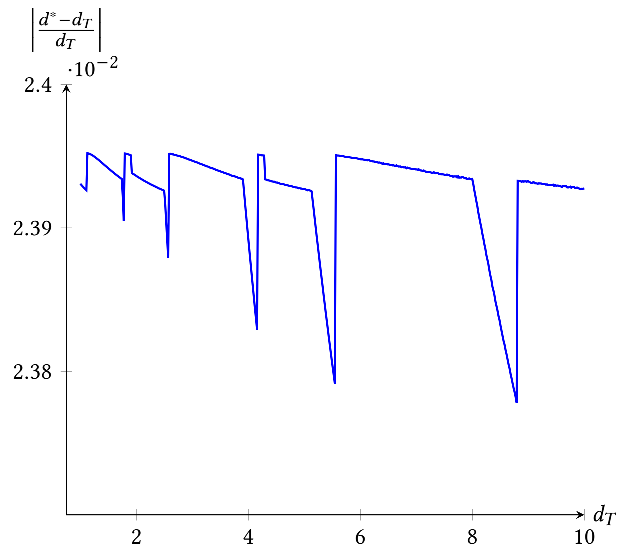

Figure 6: Relative distance between optimal and true report

3.3.4 Real World Applications: Concluding Remarks. In conclusion, we demonstrated in this section the applicability of our theoretical results to real world problems. We showed the versatility of the results from Theorem 3.2. From the family of BPI contracts provided by Theorem 3.2, we produced a solution that is suitable for a setting with Gaussian distributions.

It is important to note that the solution from Theorem 3.2 does not only fit a Gaussian well, as can be seen in Section 3.3.1. Beyond a good fit, the modeling error is often less significant than the error caused by epistemic uncertainty, i.e. the agent only knowing an approximate value of d, see Section 3.3.2. And even under perfect knowledge, the effect of misreporting is small both for the cost of the agent as well as the resulting prediction, see Section 3.3.3. All together this produces a solution that incentivizes agents to provide improved predictions most of the time.

4 ARBITRAGE: RESELLING IMPROVED PREDICTIONS

In this section, we show how reselling contracts can both benefit the reseller and lead to improved predictions.

In particular, we revisit the issue that the agents do not know their true valuation, see Section 2.3. But, we will frame the issue in the context of arbitrage. So far we have shown how our solution can incentivize a fluctuating energy source to produce accurate predictions. However, this incentive might not actually translate into better predictions, because, it is not just a matter of willingness but also ability to report the true forecast. Unlike providing private information, the task of providing accurate predictions is non-trivial and might not be within the expertise of the operator of a fluctuating energy source. However, providing such predictions might very well be the expertise of other agents.

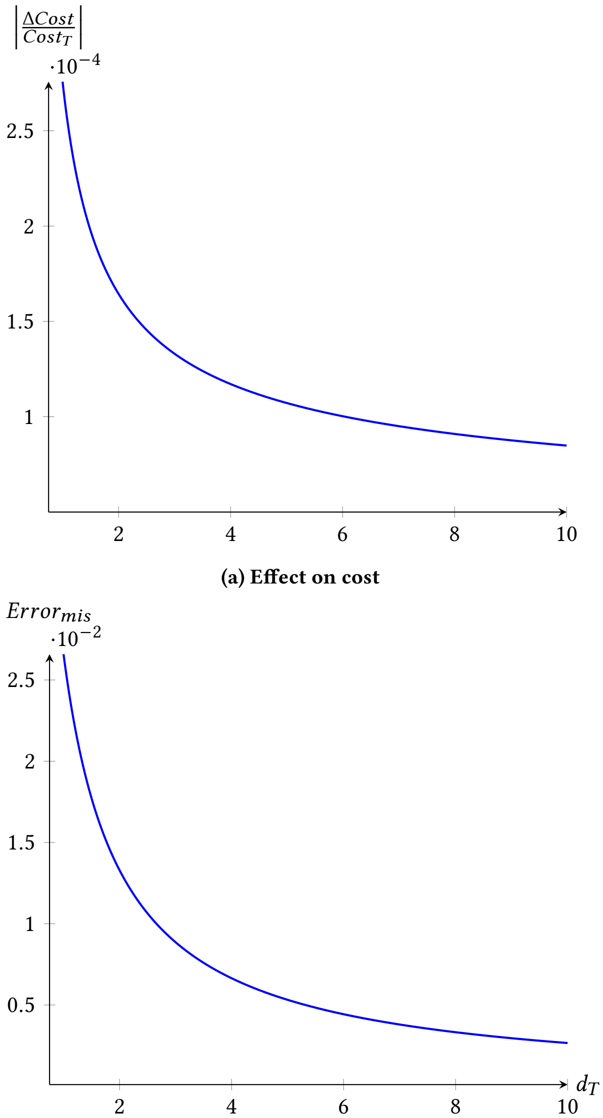

(a) Effect on cost

(b) Effect on Error of misreporting

Figure 7: Effect of reporting optimal instead of true intervals

The challenge of enabling third party agents to participate in the prediction process is not necessarily a technical one. In most cases, the sources of uncertainty in renewable energies are events with publicly available data, such as weather forecasts. This means from a data availability standpoint, the task of providing accurate predictions could be done by any third party. The challenge lies more with how to incentivize a third party to participate in the first place.

This is where arbitrage becomes relevant. Arbitrage is the act of generating a profit by reselling a good or service. The arbitrage of BPI contracts addresses the above mentioned challenge, because the contract creating function connects prices directly to the quality of the prediction, as we have argued in Section 2.3. This direct connection between prices and quality of prediction means that arbitrage of such contracts rewards agents that provide improved predictions with favorable prices.

We will illustrate this potential in the following example.

4.1 Arbitrage Example: Middleman providing an improved Prediction

To illustrate how a third party agent can benefit from arbitraging contracts created from BPI pricing functions, we will consider an example setting consisting of a Wind Farm (WF), flexibility provider (FP) and middleman (M) as depicted in Figure 8. For this example we assume the true day-ahead wind power fluctuating is described by a truncated Gaussian on the interval $[-d_T, d_T]$. For the contract generating function $\Omega_p(d|k^*)$, we use the cost function as outlined in Theorem 3.2 with parameter $k = 1.4$, as this has been shown to be a good approximation for a truncated Gaussian, see Section 3.3.2. Figure 9 depicts the expected cost the Wind farm incurs when providing an estimation d.

Initially, there only exists the rather inaccurate estimate d_1 that the WF produced – indicated in Figure 9. The WF uses this estimate to generate a contract Θ_1 from contract generating function $\Omega_p(d|k^*)$. It offers the contract Θ_1 to the market. Contract Θ_1, the input parameter d_1 as well as the wind farm that the contract refers to are public knowledge.

A third party agent M, observes contract Θ_1 as well as input d_1 and believes it has a better estimate d_2 – indicated in Figure 9. Therefore M decides to buy contract Θ_1 and generates a new contract using its improved estimate d_2. It then offers this improved contract Θ_2 to the market, where it gets sold to the flexibility provider FP.

It is important to note that middleman M has no energy obligation and only engages in arbitrage. On a physical level both contract Θ_1 and Θ_2 refer to the same uncertain future energy production. This means that both contracts insure the same event and therefore the energy Δq that is exchanged is the same in both contracts, leaving M with no obligation.

With regard to prices, however, Θ_1 and Θ_2 differ in favor of M. Recall, Corollary 1 from Section 2.2, which states that better predictions lead to a lower expected cost when using BPI contract functions. Since d_2 is closer to the true interval size d_T than d_1, the cost paid for Θ_1 is higher than the cost of Θ_2. This way the third party agent M has an expected profit of $C_B(\Theta_1) - C_B(\Theta_2) > 0$ by providing an improved prediction to the flexibility provider. This is visually indicated in Figure 9 as the difference between $C_1 = C_B(\Theta_1)$ and $C_2 = C_B(\Theta_2)$.

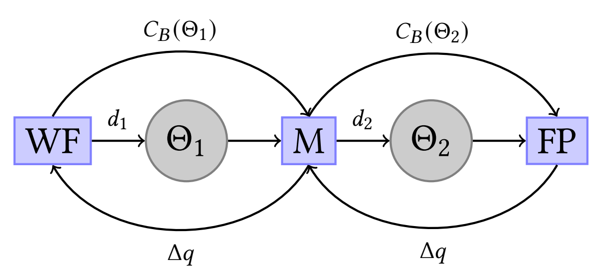

Figure 8: Arbitrage Setup

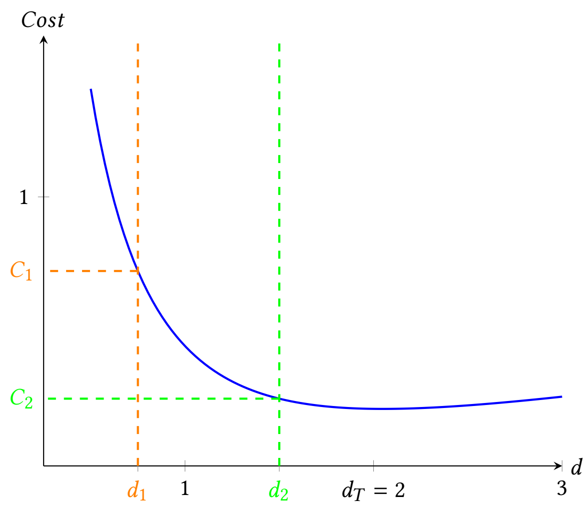

Figure 9: Arbitrage by Cost Difference

Ultimately, middleman M generated a profit and provided an improved prediction to the flexibility provider by reselling contracts.

5 DISCUSSION

Having shown the efficacy of our solution in incentivizing good predictions (Section 3) as well as the ability of our approach to accommodate third party prediction providers (Section 4) we compare our solution to two other information theoretical approaches: data markets (Section 5.1) and prediction markets (Section 5.2)

5.1 Comparison to Data Markets

Data markets are mechanisms for the sharing of data between non-cooperative agents. For the use in renewable energy forecasting, data markets enable different RES owners to purchase each other's data/forecasts in order to improve their own forecasts. A market system for a single buyer and several data sellers has been proposed in [11, 36], while a two sided mechanism with several data buyers and sellers has been designed in [10].

Both data markets and our approach have the aim of improving forecasts. However, they differ conceptually in two distinct ways.

First data markets are designed as secondary markets that operate exclusively on the informational layer. The concept of BPI contracts as presented here, on the other hand, are power markets that come with balancing responsibilities.

Secondly, data markets unlike BPI contracts lack a clear incentive structure to reduce epistemic uncertainty. Purchasing additional data may lead to an improved prediction, effectively reducing epistemic uncertainty. However, this is not guaranteed and more importantly, there is no clear incentive structure to provide the most useful data. In contrast, BPI contracts directly incentivize agents to provide improved forecasts by reducing epistemic uncertainty.

5.2 Comparison to Prediction Markets

Predictions markets aim to to predict outcomes of future events [45] and have been applied in a wide range of applications [13]. Contracts, traded on prediction markets, yield a pay-off depending on the outcome of a pre-defined event. From the discrepancy between trading price and pay-off, the probabilities of specific outcomes can be derived. Through the continuous trading process the distributed information of the participants is being aggregated.

While prediction markets generally have been around for several decades, their use in the energy literature has been a more recent development. Initial research has focused on utilizing prediction markets for hedging purposes [33, 35]. Recently, a prediction market for complete probabilistic forecasts has been proposed by Shamsi and Cuffe [34].

Prediction markets are a suitable tool for providing energy forecasts, because their aim is to reveal information about a future event. Such markets, however, differ from BPI in the way they seek to achieve this goal. While prediction markets try to leverage the wisdom of the crowds to provide an aggregate prediction, BPI contracts directly incentivize accurate predictions within the contract structure.

Similar to data markets, prediction markets differ from the concepts presented here by operating only on the informational layer. In the case of prediction market, this disconnect from the operation of the grid, makes them more susceptible to speculation. BPI contracts on the other hand are power contracts first and come with balancing obligations.

6 CONCLUSION

The work presented here contributes to the energy flexibility literature by providing the concept of Best Prediction Incentivizing contracts. The concept addresses the often overlooked issue of how to incentivize agents to accurately predict their fluctuating future energy production. The concept differs from Incentive Compatibility by focusing on reducing epistemic uncertainty about the real world rather than just incentivizing the revelation of private information.

Furthermore, we showed the applicability of this concept to the real world. Starting by providing BPI pricing functions for simple cases – uniform and single term polynomial distributions – we then showed that these simple solutions can be tuned to approximate more realistic Gaussian distributions.

Finally, we put our results into a larger context. We explored how third party agents are incentivized by the arbitrage potential to contribute improved predictions and compared our solution to other information theoretical approaches.

Going forward, there still remain some open questions around the integration of BPI contracts into larger energy systems. For instance, it is unclear how wide-spread existence of such contracts would affect balancing markets. Further research also has to be done on how inter-dependencies of different events affect their corresponding contracts.

ACKNOWLEDGMENTS

The authors thank the anonymous ACM E-Energy reviewers. This work is part of the research project Heat and Power Systems at

Industrial Sites and Harbours (HaPSISH) with project number OND1363719, and partly financed by the Dutch Research Council (NWO).
This work was also partially supported by Epistemic AI, a project funded by EU Horizon 2020 research and innovation programme under grant number 964505.

REFERENCES

[1] Mehrdad Aghamohamadi and Amin Mahmoudi. 2019. From bidding strategy in smart grid toward integrated bidding strategy in smart multi-energy systems, an adaptive robust solution approach. *Energy* 183 (2019), 75–91.
[2] Hans Auer and Reinhard Haas. 2016. On integrating large shares of variable renewables into the electricity system. *Energy* 115 (2016), 1592–1601.
[3] Zhengqi Chen, Yingyun Sun, Xin Ai, Sarmad Majeed Malik, and Liping Yang. 2019. Integrated demand response characteristics of industrial park: a review. *Journal of Modern Power Systems and Clean Energy* 8, 1 (2019), 15–26.
[4] Paul L Denholm, Erik Ela, Brendan Kirby, and Michael R. Milligan. 2010. Role of Energy Storage with Renewable Electricity Generation. https://www.osti.gov/biblio/972169
[5] Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural Safety* 31, 2 (2009), 105–112.
[6] Jinfeng Ding, Kaigui Xie, Bo Hu, Changzheng Shao, Tao Niu, Chunyan Li, and Congcong Pan. 2022. Mixed aleatory-epistemic uncertainty modeling of wind power forecast errors in operation reliability evaluation of power systems. *Journal of Modern Power Systems and Clean Energy* 10, 5 (2022), 1174–1183.
[7] Jonas Engels. 2020. Integration of flexibility from battery storage in the electricity market. *IEEE Transactions on Smart Grid* 11, 4 (2020), 3270–3279.
[8] Fuchao Ge, Yuntao Ju, Zhinan Qi, and Yi Lin. 2018. Parameter estimation of a Gaussian mixture model for wind power forecast error by Riemann L-BFGS optimization. *IEEE Access* 6 (2018), 38892–38899.
[9] Hessam Golmohamadi. 2022. Demand-side flexibility in power systems: A survey of residential, industrial, commercial, and agricultural sectors. *Sustainability* 14, 13 (2022), 7916.
[10] Carla Goncalves, Pierre Pinson, and Ricardo J Bessa. 2020. Towards data markets in renewable energy forecasting. *IEEE Transactions on Sustainable Energy* 12, 1 (2020), 533–542.
[11] Liyang Han, Pierre Pinson, and Jalal Kazempour. 2022. Trading data for wind power forecasting: A regression market with lasso regularization. *Electric Power Systems Research* 212 (2022), 108442.
[12] Shahram Hanifi, Xiaolei Liu, Zi Lin, and Saeid Lotfian. 2020. A critical review of wind power forecasting methods—past, present and future. *Energies* 13, 15 (2020), 3764.
[13] Christian Franz Horn, Bjoern Sven Ivens, Michael Ohneberg, and Alexander Brem. 2014. Prediction markets: A literature review 2014. *The Journal of Prediction Markets* 8, 2 (2014), 89–126.
[14] Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning* 110 (2021), 457–506.
[15] IRENA. 2023. Renewable power generation costs in 2022. (2023).
[16] Verena Kleinschmidt, Thomas Hamacher, Vedran S. Perić, and Mohammad Reza Hesamzadeh. 2020. Unlocking Flexibility in Multi-Energy Systems: A Literature Review. *2020 17th International Conference on the European Energy Market (EEM)* (2020), 1–6. https://doi.org/10.1109/EEM49802.2020.9221927
[17] Christos D Korkas, Simone Baldi, Iakovos Michailidis, and Elias B Kosmatopoulos. 2016. Occupancy-based demand response and thermal comfort optimization in microgrids with renewable energy sources and energy storage. *Applied Energy* 163 (2016), 93–104.
[18] Dirk Kuiken and Heyd F Más. 2019. Integrating demand side management into EU electricity distribution system operation: A Dutch example. *Energy Policy* 129 (2019), 153–160.
[19] David Lane, David Scott, Mikki Hebl, Rudy Guerra, Dan Osherson, and Heidi Zimmer. 2003. *Introduction to statistics*. David Lane.
[20] Eamonn Lannoye, Damian Flynn, and Mark O'Malley. 2012. Evaluation of power system flexibility. *IEEE Transactions on Power Systems* 27, 2 (2012), 922–931.
[21] Hoesung Lee, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, Ko Barret, et al. 2023. IPCC, 2023: Climate Change 2023: Synthesis Report, Summary for Policymakers. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland. (2023).
[22] Lei Liu, Jicheng Liu, Yu Ye, Hui Liu, Kun Chen, Dong Li, Xue Dong, and Mingzhai Sun. 2023. Ultra-short-term wind power forecasting based on deep Bayesian model with uncertainty. *Renewable Energy* 205 (2023), 598–607.
[23] Jesus Lizana, Daniel Friedrich, Renaldi Renaldi, and Ricardo Chacartegui. 2018. Energy flexible building through smart demand-side management and latent heat storage. *Applied Energy* 230 (2018), 471–485.
[24] Peter D Lund, John Byrne, Reinhard Haas, and Damian Flynn. 2019. *Advances in energy systems: The large-scale renewable energy integration challenge*. John Wiley & Sons.
[25] Alexandra Lüth, Jan Martin Zepter, Pedro Crespo Del Granado, and Ruud Egging. 2018. Local electricity market designs for peer-to-peer trading: The role of battery flexibility. *Applied Energy* 229 (2018), 1233–1243.
[26] Kevin P Murphy. 2022. *Probabilistic machine learning: an introduction*. MIT press.
[27] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. 2007. Algorithmic Game Theory. https://doi.org/10.1017/CBO9780511800481
[28] Vishwamitra Oree and Sayed Z Sayed Hassen. 2016. A composite metric for assessing flexibility available in conventional generators of power systems. *Applied Energy* 177 (2016), 683–691.
[29] Peter Palensky and Dietmar Dietrich. 2011. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics* 7, 3 (2011), 381–388.
[30] Anthony Papavasiliou and Shmuel S Oren. 2013. Large-scale integration of deferrable demand and renewable energy sources. *IEEE Transactions on Power Systems* 29, 1 (2013), 489–499.
[31] Christian P Robert. 1995. Simulation of truncated normal variables. *Statistics and Computing* 5 (1995), 121–125.
[32] Roland Saur, Han La Poutré, and Neil Yorke-Smith. 2020. Towards Optimal Demand-Side Bidding in Parallel Auctions for Time-Shiftable Electrical Loads. In *2020 IEEE 18th International Conference on Industrial Informatics (INDIN)*, Vol. 1. IEEE, 340–347.
[33] Mahdieh Shamsi and Paul Cuffe. 2021. A prediction market trading strategy to hedge financial risks of wind power producers in electricity markets. *IEEE Transactions on Power Systems* 36, 5 (2021), 4513–4523.
[34] Mahdieh Shamsi and Paul Cuffe. 2021. Prediction markets for probabilistic forecasting of renewable energy sources. *IEEE Transactions on Sustainable Energy* 13, 2 (2021), 1244–1253.
[35] Mahdieh Shamsi and Paul Cuffe. 2021. Using binary prediction markets as hedging instruments: Strategies for renewable generators. *IEEE Transactions on Sustainable Energy* 13, 2 (2021), 1160–1163.
[36] Zelin Sun, Leandro Von Krannichfeldt, and Yi Wang. 2023. Trading and Valuation of Day-Ahead Load Forecasts in an Ensemble Model. *IEEE Transactions on Industry Applications* (2023).
[37] Govinda R Timilsina. 2021. Are renewable energy technologies cost competitive for electricity generation? *Renewable Energy* 180 (2021), 658–672.
[38] Pierre-Julien Trombe, Pierre Pinson, and Henrik Madsen. 2012. A general probabilistic forecasting framework for offshore wind power fluctuations. *Energies* 5, 3 (2012), 621–657.
[39] Orlando Valarezo, Tomás Gómez, José Pablo Chaves-Ávila, Leandro Lind, Mauricio Correa, David Ulrich Ziegler, and Rodrigo A. Escobar. 2021. Analysis of New Flexibility Market Models in Europe. *Energies* (2021). https://doi.org/10.3390/en14123521
[40] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. https://doi.org/10.1038/s41592-019-0686-2
[41] Yih-huei Wan and Brian K Parsons. 1993. *Factors relevant to utility integration of intermittent renewable technologies*. Technical Report. National Renewable Energy Lab.(NREL), Golden, CO (United States).
[42] Jiawei Wang, Shi You, Yi Zong, Chresten Træholt, Zhao Yang Dong, and You Zhou. 2019. Flexibility of combined heat and power plants: A review of technologies and operation strategies. *Applied Energy* 252 (2019), 113445.
[43] Jianxiao Wang, Haiwang Zhong, Ziming Ma, Qing Xia, and Chongqing Kang. 2017. Review and prospect of integrated demand response in the multi-energy system. *Applied Energy* 202 (2017), 772–782.
[44] Honglin Wen, Jinghuan Ma, Jie Gu, Lyuzerui Yuan, and Zhijian Jin. 2022. Sparse variational gaussian process based day-ahead probabilistic wind power forecasting. *IEEE Transactions on Sustainable Energy* 13, 2 (2022), 957–970.
[45] Justin Wolfers and Eric Zitzewitz. 2004. Prediction markets. *Journal of Economic Perspectives* 18, 2 (2004), 107–126.
[46] Jinhua Zhang, Jie Yan, David Infield, Yongqian Liu, and Fue-sang Lien. 2019. Short-term forecasting and uncertainty analysis of wind turbine power based on long short-term memory network and Gaussian mixture model. *Applied Energy* 241 (2019), 229–244.
[47] Yue Zhao, Junjie Qin, Ram Rajagopal, Andrea J. Goldsmith, and H. Vincent Poor. 2015. Wind Aggregation Via Risky Power Markets. *IEEE Transactions on Power Systems* 30 (2015), 1571–1581. https://doi.org/10.1109/TPWRS.2014.2359671