

Bayesian inference for estimating generation costs in electricity markets

Matthias Pirlet ^{a,*}, Adrien Bolland ^a, Alexandre Huynen ^b, Quentin Louveaux ^a,
Gilles Louppe ^a, Damien Ernst ^a

^a University of Liège, Belgium

^b Engie, Brussels, Belgium

ARTICLE INFO

Keywords:

Electricity markets
Machine learning
Optimization
Statistics
Unit commitment

ABSTRACT

Estimating generation costs from observed electricity market data is essential for market simulation, strategic bidding, and system planning. To that end, we model the relationship between generation costs and production schedules with a latent variable model. Estimating generation costs from observed schedules is then formulated as Bayesian inference. A prior distribution encodes an initial belief on parameters, and the inference consists of updating the belief with the posterior distribution given observations. We use balanced neural posterior estimation (BNPE) to learn this posterior. Validation on the IEEE RTS-96 test system shows that marginal costs are recovered with narrow credible intervals, while start-up costs remain largely unidentifiable from schedules alone. The method is benchmarked against an inverse-optimization algorithm that exhibits larger parameter errors without uncertainty quantification.

1. Introduction

Electricity is exchanged across multiple markets by many agents, from years to minutes before actual delivery [1,2]. Agents make decisions in this complex process based on incomplete or imperfect information. Estimation of generation costs (marginal and start-up) from observed market data (generation schedule) enables more accurate price forecasts, risk assessment, and long-term revenue estimation, benefiting analysts in utilities and hedge funds [3,4]. Such estimates further allow regulators and TSOs to monitor market behavior and evaluate capacity retention in zonal markets where localized price signals are limited.

Producers aim to maximize profit by selling production through electricity markets. In a perfect market, participants bid at their marginal costs for all future periods available on that market, i.e., they propose to sell electricity at the cost of producing each additional MWh. In theory, as time advances, bids evolve as new information becomes available until the delivery period. In practice, several electricity markets exist to approximate this ideal behavior. We focus on an ideal day-ahead market, where production schedules for the day are fixed one day in advance for each participant. We assume that each participant submits independent bids for each production unit. Equivalently, each unit corresponds to a unique participant.

A unit commitment (UC) problem [5] is standardly solved to approximate the market clearing process. This mixed-integer optimization problem minimizes total generation cost while satisfying physical con-

straints such as capacity limits, ramp rates, minimum up/down periods, and transmission limits. Large-scale UC problems can be solved efficiently using modern formulations [6,7] and commercial solvers. Using UC to approximate market clearing is justified by the assumption that in a competitive environment, the collective result of participant bidding aligns with the cost-optimal physical dispatch of the system's assets [8–10]. Consequently, market participants and system operators use UC models as a proxy to simulate market outcomes and forecast prices and volumes.

Literature on UC divides into two main categories. Forward approaches treat cost parameters as known and seek to determine generation schedules. Operational uncertainties are addressed through stochastic programming [11–14] or robust optimization [15]. Inverse approaches treat cost parameters as unknown and seek to infer them from observed schedules. They typically consist of solving an inverse optimization problem to recover parameters [16,17]. For example, Liang and Dvorkin [18] apply such a method to estimate bid prices in US nodal markets from locational marginal prices and generator schedules. This approach relies on bus-level price information unavailable in zonal markets such as those in Europe [19]. Inverse optimization methods often return point estimates and provide limited uncertainty quantification on the recovered parameters. Preliminary work addresses this limitation by applying simulation-based inference [20] to generation cost parameter estimation in UC problems [21]. It demonstrates the feasibility of using neural posterior estimation (NPE) [22–25] to estimate generation cost

* Corresponding author.

E-mail address: matthias.pirlet@uliege.be (M. Pirlet).

in a simplified 9-unit system without network constraints. The present work extends this foundation to a realistic network and compares to a traditional inversion approach.

In this work, we extend the UC formulation to model stochastic demand, line and generation outages, and strategic bidding, all of which are reflected in market schedules. To that end, we use a latent variable model and formulate cost estimation as a Bayesian inference problem. This allows quantifying uncertainty about generation costs and updating estimates as data is observed. We use balanced neural posterior estimation (BNPE) to approximate posterior distributions, enabling fast repeated inference with uncertainty quantification. We also introduce an inverse-optimization baseline that provides fast point estimates when full uncertainty quantification is not required. We validate both approaches on the IEEE RTS-96 test system [26] with its 2016 update [27]. Methods are compared based on estimation accuracy, computational efficiency, and uncertainty characterization.

The remainder of the paper is organized as follows. In Section 2, the market model and the Bayesian inference problem for recovering generation costs are formulated. The two practical algorithms are described in Section 3 and validated in Section 4. Finally, conclusions are highlighted in Section 6.

2. Problem statement

In this section, we first define a latent variable model [28] that describes how electricity market schedules are obtained depending on generation costs. We then formalize the problem of estimating costs as a Bayesian inference problem with the previous model.

The latent variable model describing the market interactions is illustrated in Fig. 1. Formally, let us describe the electrical network with a graph composed of nodes $\mathcal{N}$ and edges $\mathcal{E} \subset \mathcal{N} \times \mathcal{N}$. Let $\mathcal{T} = \{1, \dots, T\}$ index time periods and $\mathcal{J} = \{1, \dots, J\}$ index generators. The model parameters include the marginal generation cost c_{tj} and the start-up cost s_{tj} for each generator $j \in \mathcal{J}$ at time $t \in \mathcal{T}$. We consider uniform priors on these parameters $c_{tj} \sim \mathcal{U}(c_{\min}, c_{\max})$ and $s_{tj} \sim \mathcal{U}(s_{\min}, s_{\max})$, which model uncertainty over a range of feasible values. The price offered for producing energy depends on the marginal costs and accounts for trading




Fig. 1. Latent variable model for cost inference. White circles represent latent variables, gray circles represent observed variables. The two rectangular boxes denote the T time periods over which the model operates, with variables influencing each other across time. In the upper box are marginal costs c (latent, uniform prior) and start-up costs s (latent, uniform prior). Marginal costs pass through a stochastic bidding model $\tilde{c} \sim \mathcal{N}(c, \sigma^2)$ (latent), while start-up costs remain unchanged. In the lower part are generator availability α (latent, Bernoulli), line availability β (latent, Bernoulli), observed system load δ (gray circle). The deterministic UC optimization problem combines bidding costs $\tilde{c}$, start-up costs s , availabilities α, β , and observed system load δ (gray circle) sampled from a distribution capturing seasonal, weekly, and diurnal demand patterns, to produce the schedule g (observed, gray circle).

strategies, which is modeled with the bidding costs $\tilde{c}_{tj} \sim \mathcal{N}(c_{tj}, \sigma^2)$. The availability of generator $j \in \mathcal{J}$ is modeled with the Bernoulli random variable $\alpha_j \sim \text{Bernoulli}(1 - p_{\text{gen}})$ and the availability of the line $e \in \mathcal{E}$ with $\beta_e \sim \text{Bernoulli}(1 - p_{\text{line}})$. The demand at bus n and time t is modeled as $\delta_{tn} = L_t \cdot s_{tn}$, where L_t is the total system load following seasonal and diurnal profiles specified in the IEEE RTS-96 test system and s_{tn} are the bus-level load shares, normalized such that $\sum_n s_{tn} = 1$. Both variables are disturbed by a Gaussian noise before scaling such that the demand vector δ follows the distribution $p(\delta)$. Generation schedules g are the solution of the UC problem, which minimizes total generation cost subject to physical and operational constraints such as capacity limits, ramp rates, minimum up/down times, and transmission limits. The generation schedule depends on the realizations of the previous events

$$g = UC(\tilde{c}, s, \alpha, \beta, \delta). \quad (1)$$

We denote the market outcome by $x = (g, \delta)$, the cost parameters by $\theta = (c, s)$ and the latent variables by $z = (\tilde{c}, \alpha, \beta)$. Finally, the forward model is defined by the joint distribution of these variables

$$p(x, z, \theta) = p(g | z, \theta, \delta) p(z | \theta) p(\theta) p(\delta), \quad (2)$$

where $p(z | \theta) = p(\tilde{c} | c) p(\alpha) p(\beta)$ and $p(g | z, \theta, \delta)$ is a Dirac centered in the solution of the UC.

The problem of finding cost parameters θ from observed market outcome $x^{\text{obs}} = (g^{\text{obs}}, \delta^{\text{obs}})$ is formulated as a Bayesian inference problem. We represent our knowledge on plausible cost parameters before observing market outcome using the prior distribution $p(\theta)$, and apply Bayes' theorem to represent plausible cost parameters after observation using the posterior distribution

$$p(\theta | x^{\text{obs}}) \propto p(\theta) p(x^{\text{obs}} | \theta). \quad (3)$$

3. Methodology

To solve the Bayesian inference problem defined in Eq. (3), we present balanced neural posterior estimation (BNPE) [29], a simulation-based inference method, to approximate the posterior distribution from simulations. We compare this approach with a baseline method using inverse optimization, which provides fast point estimates without quantifying uncertainty. These two methods offer different trade-offs between estimation accuracy, computational efficiency, and uncertainty characterization.

3.1. Simulation-based posterior inference with BNPE

Computing the posterior distribution from Bayes' theorem would require integrating over latent variables, which is intractable in practice. Instead, we use BNPE, in which we train a neural network $q_\phi(\theta | x)$, where ϕ denotes the weights of the network, to approximate the posterior $p(\theta | x)$. This network is trained by minimizing the expected Kullback-Leibler (KL) divergence between the true posterior and the neural network approximation

$$\min_{\phi} \mathbb{E}_{p(x)} [\text{KL}(p(\theta | x) \| q_\phi(\theta | x))],$$

where the expectation is taken over the marginal distribution of simulated market outcomes $p(x) = \int p(x | \theta) p(\theta) d\theta$.

The KL divergence measures the information loss when approximating $p(\theta | x)$ with $q_\phi(\theta | x)$. As shown in Appendix A.1, this is equivalent to minimizing the expected negative log-posterior

$$\min_{\phi} \mathbb{E}_{p(\theta, x)} [-\log q_\phi(\theta | x)],$$

where $p(\theta, x) = p(x | \theta) p(\theta)$ denotes the joint distribution induced by the latent variable model.

Training consists in generating samples (θ, x) by sampling θ from the prior $p(\theta)$, sampling δ from the load distribution $p(\delta)$, and computing the schedule g through the UC optimization, resulting in $x = (g, \delta)$. These samples are used to estimate the expected log-posterior, which is

minimized via stochastic gradient descent on the network parameters ϕ . Once trained, the density $q_\phi(\theta \mid x^{\text{obs}})$ serves as a surrogate for $p(\theta \mid x^{\text{obs}})$ when observing market outcome x^{obs} .

3.2. Inverse optimization with polar cone method

We compare our approach with inverse optimization [17]. This method seeks cost parameters for which the observed schedule is optimal under a deterministic model. We therefore neglect stochastic bidding and random outages to obtain a deterministic model, namely the UC problem and

$$g = \text{UC}(c, s, \mathbf{1}, \mathbf{1}, \delta). \quad (4)$$

Let us first note that for a given observed load δ^{obs} , the observed schedule g^{obs} is optimal when solving the UC problem with any parameter θ that satisfies the condition

$$\theta^\top (g - g^{\text{obs}}) \geq 0 \quad \forall g \in \mathcal{G}(\delta^{\text{obs}}), \quad (5)$$

where $\mathcal{G}(\delta^{\text{obs}})$ is the set of feasible schedules in the UC problem given the observed load δ^{obs} . The set of parameters that satisfy this condition is denoted by Θ^* , which is usually called the polar cone [30].

The inverse optimization problem refines iteratively an approximation of the polar cone. The initial approximation is the hypercube $\Theta^{(0)} = [\theta^{\min}, \theta^{\max}]$ representing physically plausible cost ranges. At each iteration k , we sample a reference parameter θ^{ref} from the prior that we project onto the current approximation $\Theta^{(k)}$ solving

$$\min_{\theta \in \Theta^{(k)}} \|\theta - \theta^{\text{ref}}\|. \quad (6)$$

Let $\theta^{(k)}$ denote the solution of the projection. We solve the UC problem from Eq. (4) with these costs to obtain the schedule $g^{(k)}$ and verify the optimality condition (5). If $\theta^{(k)\top}(g^{(k)} - g^{\text{obs}}) = 0$, then the observed schedule is optimal for the current cost parameters and the algorithm has converged. If $\theta^{(k)\top}(g^{(k)} - g^{\text{obs}}) < 0$, then $\theta^{(k)}$ violates condition (5), and we refine the polar cone approximation $\Theta^{(k+1)} = \{\theta \in \Theta^{(k)} : \theta^\top (g^{(k)} - g^{\text{obs}}) \geq 0\}$.

As the number of iterations increases, $\Theta^{(k)}$ converges to the polar cone Θ^* , and g^{obs} is an optimal schedule for any solution of the projection (6). In practice, we stop when $\theta^{(k)\top}(g^{(k)} - g^{\text{obs}}) \geq -\epsilon$, meaning that the observed schedule is ϵ -optimal, and the parameter $\theta^{(k)}$ serves as the approximate solution.

This method is only applicable for inverting deterministic models. We therefore ignore stochastic random bidding and model outages, which introduces model misspecification. Despite this limitation, inverse optimization provides fast point estimates for time-critical applications.

4. Experiments

4.1. Experimental setup

We apply our two algorithms on the IEEE RTS-96 test system [26] with its 2016 update [27]. This system includes 24 buses $n \in \mathcal{N}$ (17 buses with loads, and $J = 12$ generators connected to 10 buses), and 38 transmission lines $e \in \mathcal{E}$. This problem covers a time horizon of $T = 24$ hours.

We estimate cost parameters for each generator, i.e., marginal costs c_{ij} and start-up costs s_{ij} . We assume these costs remain constant over the 24-hour horizon such that

$$c_{ij} = c_j, \quad \forall t \in \mathcal{T}$$

$$s_{ij} = s_j, \quad \forall t \in \mathcal{T}$$

We define uniform priors for marginal costs $c_j \sim \mathcal{U}(10, 50)\text{€/MWh}$ and start-up costs $s_j \sim \mathcal{U}(500, 10000)\text{€}$. These ranges reflect typical values observed in electricity markets while encoding minimal prior knowledge.

We model bidding costs with the transition $\tilde{c}_j \sim \mathcal{N}(c_j, \sigma^2)$ with $\sigma = 2.5 \text{ €/MWh}$. We model generator and transmission line availability with $\alpha_j \sim \text{Bernoulli}(0.95)$ and $\beta_{ij} \sim \text{Bernoulli}(0.99)$, respectively.

Daily demand scenarios are generated by applying the modeling described in Section 2 to the IEEE RTS-96 system. Specifically, we set the total load noise to $\sigma_L = 5\%$ and the bus-share perturbation to $\sigma_s = 0.5\%$, using the seasonal and diurnal profiles from [27] as the base temporal patterns.

The UC model is a security-constrained unit commitment (SCUC) [31] model, which minimizes total generation cost subject to generator operational constraints (capacity limits, ramp rates, minimum up/down times) and transmission network constraints (thermal limits of line under DC power flow). The complete mathematical formulation is as follows.

Technical parameters ψ

$$\begin{aligned} G_j^{\min}, G_j^{\max} & \quad \text{min/max generation of unit } j, \\ ru_j, rd_j & \quad \text{ramp-up / ramp-down limits for unit } j, \\ mu_j, md_j & \quad \text{minimum up / down times (integer) for unit } j, \\ B_{ij} & \quad \text{susceptance of line } (i, j), \\ \bar{F}_{ij} & \quad \text{thermal capacity of line } (i, j), \\ v_j^{\text{init}}, g_j^{\text{init}} & \quad \text{initial on/off status and power of unit } j, \\ \alpha_j & \quad \text{availability indicators for generator } j, \\ \beta_{ij} & \quad \text{availability indicators for line } (i, j), \\ n_{\text{slack}} \in \mathcal{N} & \quad \text{index of the slack (reference) bus.} \end{aligned} \quad (7)$$

System load

$$\delta_{in} \quad \text{demand at bus } n \text{ and time } t.$$

Cost parameters $\theta = [c, s]$ (unknown in the inverse problem)

$$c_j \quad \text{marginal cost of unit } j,$$

$$s_j \quad \text{start-up cost of unit } j.$$

Decision variables

$$g_{tj} \in \mathbb{R}_+ \quad \text{dispatch of unit } j \text{ at time } t,$$

$$v_{tj} \in \{0, 1\} \quad \text{commitment (on/off) of unit } j \text{ at time } t,$$

$$y_{tj} \in \{0, 1\} \quad \text{start-up indicator of unit } j \text{ at time } t,$$

$$z_{tj} \in \{0, 1\} \quad \text{shut-down indicator of unit } j \text{ at time } t,$$

$$\vartheta_{in} \in \mathbb{R} \quad \text{voltage angle at bus } n \text{ and time } t,$$

$$f_{tij} \in \mathbb{R} \quad \text{directed flow on line } (i, j) \text{ at time } t.$$

Constraints

Slack bus

$$\vartheta_{tn_{\text{slack}}} = 0 \quad \forall t \in \mathcal{T}.$$

DC power flow and thermal limits:

$$f_{tij} = B_{ij}(\vartheta_{ti} - \vartheta_{tj}),$$

$$-\beta_{ij} \bar{F}_{ij} \leq f_{tij} \leq \beta_{ij} \bar{F}_{ij}, \quad \forall (i, j) \in \mathcal{E}, \forall t \in \mathcal{T}.$$

Generation bounds:

$$\alpha_j G_j^{\min} v_{tj} \leq g_{tj} \leq \alpha_j G_j^{\max} v_{tj} \quad \forall j \in \mathcal{J}, \forall t \in \mathcal{T}.$$

State-transition (start/stop):

$$y_{tj} - z_{tj} = \begin{cases} v_{1j} - v_j^{\text{init}} & t = 1, \\ v_{tj} - v_{t-1j} & t > 1, \end{cases} \quad \forall j \in \mathcal{J}.$$

Ramp constraints for $t = 1$:

$$g_{1j} - g_j^{\text{init}} \leq ru_j v_j^{\text{init}} + G_j^{\min} y_{1j}, \quad \forall j \in \mathcal{J},$$

$$g_j^{\text{init}} - g_{1j} \leq rd_j v_j^{\text{init}} + G_j^{\min} z_{1j}, \quad \forall j \in \mathcal{J}.$$

Ramp constraints for $t > 1$:

$$g_{tj} - g_{t-1j} \leq r u_j v_{t-1j} + G_j^{\min} y_{tj}, \quad \forall j \in \mathcal{J},$$

$$g_{t-1j} - g_{tj} \leq r d_j v_{tj} + G_j^{\min} z_{tj}, \quad \forall j \in \mathcal{J}.$$

Minimum up/down time constraints:

For $t \geq mu_j$

$$\sum_{k=t-mu_j+1}^t y_{kj} \leq v_{tj}, \quad \forall j \in \mathcal{J},$$

and for $t \geq md_j$

$$\sum_{k=t-md_j+1}^t z_{kj} \leq 1 - v_{tj} \quad \forall j \in \mathcal{J}.$$

Nodal power balance: For each bus $n \in \mathcal{N}$ and time $t \in \mathcal{T}$

$$\sum_{j \in \mathcal{J}_n} g_{tj} - \delta_{tn} = \sum_{m: (n,m) \in \mathcal{L}} f_{tnm} - \sum_{m: (m,n) \in \mathcal{L}} f_{tmn},$$

where $\mathcal{J}_n \subseteq \mathcal{J}$ is the set of generators connected to bus n . *Objective (cost minimization)*

$$\min_{g, v, y, z, \delta, f} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} (c_j g_{tj} + s_j y_{tj}).$$

For BNPE training, we generate $N = 2^{18} \approx 262,000$ simulation pairs (θ, x) where $x = (g, \delta)$ for training and additional 2^{18} pairs for validation. The neural density estimator is a neural spline flow [32] with 5 coupling transformations, each parameterized by a masked autoregressive network with 5 hidden layers of 256 units and ReLU activation. All cost parameters, generation schedules, and loads are standardized to zero mean and unit variance before training.

For inverse optimization, we set a maximum number of 100k iterations and an accuracy threshold at $10e^{-6}$, but the algorithm always achieves this threshold before reaching the maximum number of iterations.

4.2. Results and diagnostics

A fundamental challenge in simulation-based inference is that the true posterior distribution $p(\theta|x)$ is unknown; therefore, we cannot assess whether the learned approximation $q_\phi(\theta|x)$ has converged to the target. The quality of the learned posterior depends on the modeling choices that include the prior specification, neural architecture and hyperparameters, the distribution of training data, and the strength of the regularization. To address this lack of ground truth, multiple diagnostic tools can be used to detect potential failures and validate the reliability of these approximations.

A first standard diagnostic is to verify that the posterior conditioned on a given observation does not reproduce the prior and places a high density around the parameter vector that generated the observation. The observed market outcome x^{obs} is produced in a controlled setup by sampling a nominal parameter vector θ^* from the prior and running the forward model. If $q_\phi(\theta|x^{\text{obs}})$ matches the prior, it means that the surrogate model does not encode any information about the parameters. Figs. 2 and 3 show the distribution of $N = 2^{12}$ samples drawn from the learned posterior $q_\phi(\theta|x^{\text{obs}})$ together with the nominal parameter θ^* (marked in blue) and the solution of the inverse optimization problem (marked in red).

For marginal costs (Fig. 2), the nominal parameter values always lie in high-density regions. The narrow, peaked marginal distributions represent learning beyond the flat uniform prior $\mathcal{U}(10, 50)$ €/MWh, indicating that the posterior reflects information extracted from the observed schedule. Off-diagonal plots show correlations between marginal costs, reflecting how schedules couple the inference of cost parameters. In contrast, start-up cost posteriors (Fig. 3) remain close to the prior $\mathcal{U}(500, 10000)$ €. This diffuse distribution accurately represents the




Fig. 2. Marginal (diagonal) and joint (off-diagonal) posterior distributions for marginal cost parameters c inferred by BNPE from an observed schedule, solution of the stochastic process with nominal parameters c^* . The concentrated, non-uniform posteriors demonstrate successful learning beyond the flat prior (uniform on $[10, 50]$ €/MWh), with nominal cost parameters consistently lying in high-density regions. The inverse optimization solution shows larger deviations, as it assumes deterministic optimality while the generating process is stochastic.




Fig. 3. Marginal (diagonal) and joint (off-diagonal) posterior distributions for start-up cost parameters s inferred by BNPE from an observed schedule, solution of the stochastic process with nominal parameters s^* . The nearly uniform posteriors indicate limited learning beyond the prior (uniform on $[500, 10000]$ €), reflecting the challenge of inferring start-up costs from generation schedules. Nominal parameters s^* often lie in higher-density regions, though posterior concentration remains weak compared to marginal costs.

parameter's lack of observability in the forward model; since the schedule g is largely insensitive to s_j once a unit is committed. The fact that the model recovers the prior distribution instead of collapsing to an arbitrary point suggests that BNPE has successfully captured the physics of the problem rather than experiencing a learning failure.

The inverse optimization solution shows larger deviations from the nominal parameters c^* and often falls into low-density regions of the posterior surrogate. Unlike BNPE, inverse optimization only uses the deterministic part of the model, assuming that the observed schedule x^{obs} is a perfectly optimal response to the cost parameters θ^* . When the observation is instead generated by a stochastic process, inverse optimization incorrectly interprets random noise as a change in the UC model.

A second diagnostic assesses whether the surrogate posterior distribution provides the correct level of uncertainty about the parameters that generated an observation. If so, the model is said to be well-calibrated. We validate this by sampling nominal parameters θ^* from the prior and generate observations x^* from the forward model. For a given observation, a $(1 - \alpha)$ credible region is defined as a subset of the parameter space containing $(1 - \alpha)$ of the posterior mass. For each pair (θ^*, x^*) , we construct a $(1 - \alpha)$ credible region from $q_\phi(\theta|x^*)$ and check whether θ^* lies inside. The expected coverage is the proportion of pairs where the nominal parameter falls within the credible region. A model is well-calibrated if the expected coverage equals $(1 - \alpha)$. If the coverage is lower, the model is overconfident, which implies that the learned posterior is too narrow. If the coverage is higher, it is underconfident meaning that credible regions are unnecessarily large. Mathematical details are provided in Appendix A.2.

Our results (Fig. 4) show minor overconfidence; empirical coverage falls slightly below the nominal level across all credibility levels. The miscalibration likely arises from the high-dimensional parameter space (24 parameters), the stochastic forward model, and complex interactions




Fig. 4. Empirical coverage probability versus nominal credibility level for the learned BNPE posterior. The diagonal line represents perfect calibration, where empirical coverage equals the nominal level $1 - \alpha$. The observed curve falls slightly below the diagonal across most credibility levels, indicating modest overconfidence, i.e., credible regions are narrower than ideal. This miscalibration persists despite balanced regularization ($\lambda = 50$), likely due to the high-dimensional parameter space (24 parameters) and complex stochastic forward model.

between cost parameters and schedules that challenge the neural density estimator's capacity. This uncertainty quantification remains valuable for typical market analysis tasks, where approximate probabilistic bounds inform decision-making even if not perfectly calibrated. In contrast, overconfidence would have been problematic for safety-critical applications where credible intervals directly inform operational margins and underestimating uncertainty could lead to insufficient safety buffers. For electricity market monitoring, the learned posterior provides a reasonable balance between narrowing the prior and avoiding the complete absence of uncertainty inherent to point estimation methods.

A final diagnostic called posterior predictive checks is typically used in real-data applications to assess model adequacy. The idea is to sample parameters from the learned posterior given an observation $q_\phi(\theta|x^{\text{obs}})$, pass them through the forward model, and check whether the resulting predicted schedules x^{pred} are consistent with the observation x^{obs} . This produces samples from the posterior predictive distribution

$$p(x^{\text{pred}}|x^{\text{obs}}) = \int p(x^{\text{pred}}|\theta) q_\phi(\theta|x^{\text{obs}}) d\theta.$$

If the model is misspecified or the posterior is poor, the prediction will deviate from the observation.

In our simulated setting, this diagnostic serves a different purpose. Because the observed market outcome x^{obs} is itself generated by the same forward model that we use in the predictive check, the goal is not to detect model misspecification but to verify self-consistency. Specifically, we check whether sampling parameters from the learned posterior and pushing them forward through the model reproduces schedules compatible with x^{obs} . If this loop is coherent, then the posterior has successfully captured the inverse mapping.

Fig. 5 shows the posterior predictive distributions for all generators over the 24-hour horizon. The learned posterior $q_\phi(\theta|x^{\text{obs}})$ appears to capture the inverse mapping accurately, as parameters sampled from it generate schedules that are consistent with the observed schedule. This confirms the high quality of the posterior approximation by validating that the entire forward-inverse-forward loop is consistent with the observed data.

5. Discussion

Correctly quantifying model uncertainties is important for avoiding erroneous reasoning and enabling risk-aware decision-making, among other things. As illustrated in the experiments, the most standard approach based on inverse optimization does not quantify uncertainty and suffers from model misspecification; it is therefore prone to such failures. Approximating the posterior distribution may, however, be a difficult task. Markov Chain Monte Carlo (MCMC) methods are widely considered the gold standard for Bayesian inference, as they provide asymptotically exact samples from the posterior distribution under regularity conditions. However, these methods require evaluating the likelihood $p(x|\theta)$. In our setting, the likelihood cannot be computed efficiently. Likelihood-free MCMC variants, such as pseudo-marginal MCMC [33] or ABC-MCMC [34], bypass explicit likelihood evaluation by approximating acceptance probabilities through Monte Carlo sampling. However, these methods require solving the UC problem many times per MCMC step. Additionally, these methods suffer from exponentially vanishing acceptance rates in high dimensions. Together, these limitations make likelihood-free MCMC computationally prohibitive for this application.

Amortized optimization methods and in particular simulation-based inference (SBI) partially address previous limitations. SBI requires a one-time training phase but then provides fast posterior evaluation for any new observation, making it ideal for operational use where the same model is inverted repeatedly (e.g., daily cost estimation). For applications making many inferences between model updates, SBI's retraining cost is small compared to the cumulative cost of running inverse optimization or MCMC for each query. In contrast, inverse optimization and




Fig. 5. Posterior predictive check. For a single observation x^{obs} generated from nominal parameters θ^* , we sample $N = 2^{12}$ parameters from the learned posterior $q_{\theta}(\theta|x^{\text{obs}})$ given this observed market outcome and pass them through our forward model. Each subplot shows one generator's dispatch over 24 hours. The observed schedule consistently lies within high-probability regions across all generators and time periods, validating that the learned posterior correctly captures the stochastic forward process.

MCMC adapt immediately without requiring retraining when the model changes because of regulatory updates, network modifications, or improved modeling. This makes them better suited for research settings with evolving infrastructure.

6. Conclusion

We formulated generation cost estimation from market data as a Bayesian inference problem with a latent variable model accounting for opportunity costs and operational uncertainty. The Bayesian approach allows for encoding prior knowledge about cost parameters and systematically updating beliefs given observed data. Two complementary inference methods were discussed. Balanced neural posterior estimation (BNPE) that uses the full latent variable model to learn amortized posterior approximations with full uncertainty quantification. Second, feasibility-based inverse optimization using only a unit commitment problem provides fast point estimates suitable for time-critical decisions.

Empirical validation on the IEEE RTS-96 test system demonstrates successful parameter estimation. Marginal costs are accurately inferred with concentrated posteriors, while start-up cost posteriors appropri-

ately remain diffuse, correctly reflecting their limited observability from schedule observations alone. BNPE provides well-calibrated approximate posteriors with fast amortized inference, meaning that once trained, posterior evaluation for new observations is nearly instantaneous, enabling real-time deployment for operational market monitoring. This approach successfully combines uncertainty quantification with computational efficiency, offering a practical solution for cost parameter estimation in electricity markets.

Several promising directions extend this work. First, inferring physical parameters (e.g., ramp rates, capacity limits) jointly with cost parameters would account for model misspecification by modeling all uncertain quantities as random variables. For high-dimensional parameter spaces exceeding tens of parameters, recent advances in flow matching [35] offer more scalable alternatives to normalizing flows used here. Second, extending inference across multiple days could strengthen identification since generator marginal costs vary with publicly available fuel prices (gas, coal, oil). This would substantially increase the effective sample size, tightening the posterior and potentially identifying start-up costs from temporal patterns in consecutive start-up decisions. Finally, while Gaussian noise currently approximates unknown oppor-

tunity costs, future work should incorporate more complex distributions to better capture real-world trading strategies.

CRedit authorship contribution statement

Matthias Pirlet: Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Adrien Bolland:** Writing – review & editing, Validation, Methodology, Formal analysis, Conceptualization; **Alexandre Huynen:** Writing – review & editing, Formal analysis, Conceptualization; **Quentin Louveaux:** Methodology; **Gilles Louppe:** Writing – review & editing, Visualization, Validation, Methodology, Conceptualization; **Damien Ernst:** Writing – review & editing, Validation, Supervision, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Appendix

A.1. Derivation of the BNPE objective

We derived the BNPE objective by minimizing the expected Kullback-Leibler (KL) divergence between the true posterior $p(\theta | x)$ and its approximation $q_\phi(\theta | x)$

$$\min_{\phi} \mathbb{E}_{p(x)} [\text{KL}(p(\theta | x) \| q_\phi(\theta | x))]. \quad (\text{A.1})$$

The KL divergence is defined as

$$\begin{aligned} \text{KL}(p(\theta | x) \| q_\phi(\theta | x)) &= \int p(\theta | x) \log \frac{p(\theta | x)}{q_\phi(\theta | x)} d\theta, \\ &= \int p(\theta | x) [\log p(\theta | x) - \log q_\phi(\theta | x)] d\theta. \end{aligned} \quad (\text{A.2})$$

Using Bayes' rule, $p(\theta | x) = \frac{p(\theta)p(x|\theta)}{p(x)}$

$$\begin{aligned} &\int p(\theta | x) [\log p(\theta | x) - \log q_\phi(\theta | x)] d\theta, \\ &= \int \frac{p(x | \theta) p(\theta)}{p(x)} [\log p(\theta | x) - \log q_\phi(\theta | x)] d\theta. \end{aligned} \quad (\text{A.3})$$

The full expectation becomes

$$\begin{aligned} &\mathbb{E}_{p(x)} [\text{KL}(p(\theta | x) \| q_\phi(\theta | x))] \\ &= \int p(x) \int \frac{p(x | \theta) p(\theta)}{p(x)} [\log p(\theta | x) - \log q_\phi(\theta | x)] d\theta dx, \\ &= \int \int p(x | \theta) p(\theta) [\log p(\theta | x) - \log q_\phi(\theta | x)] d\theta dx, \\ &= \mathbb{E}_{p(\theta)p(x|\theta)} [\log p(\theta | x) - \log q_\phi(\theta | x)]. \end{aligned}$$

Minimizing this is equivalent to maximizing the expected log-probability of the approximate posterior

$$\mathbb{E}_{p(\theta)p(x|\theta)} [\log q_\phi(\theta | x)],$$

since the entropy term $\log p(\theta | x)$ is independent of the network parameters ϕ .

A.2. Coverage

Let $\Theta_{q_\phi(\theta|x)}(1 - \alpha)$ be the set of regions in the parameter space containing at least $1 - \alpha$ probability mass of the learned posterior $q_\phi(\theta|x)$

$$\Theta_{q_\phi(\theta|x)}(1 - \alpha) = \left\{ \Omega \mid \int_{\Omega} q_\phi(\theta|x) d\theta \geq 1 - \alpha \right\}.$$

Let $\Theta_{q_\phi(\theta|x)}^*(1 - \alpha) \in \Theta_{q_\phi(\theta|x)}(1 - \alpha)$ be the $(1 - \alpha)$ -highest posterior density region

$$\Theta_{q_\phi(\theta|x)}^*(1 - \alpha) = \arg \sup_{\Omega \in \Theta_{q_\phi(\theta|x)}(1 - \alpha)} \inf_{\theta \in \Omega} q_\phi(\theta|x).$$

The expected coverage is

$$\text{Cov}(1 - \alpha) = \mathbb{E}_{p(\theta,x)} [\mathbb{1}\{\theta \in \Theta_{q_\phi(\theta|x)}^*(1 - \alpha)\}]. \quad (\text{A.4})$$

In practice, we estimate this empirically on a test set of $N = 2^{12}$ samples $(\theta^{(i)}, x^{(i)})$ by

$$\widehat{\text{Cov}}(1 - \alpha) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\theta^{(i)} \in \Theta_{q_\phi(\theta|x^{(i)})}^*(1 - \alpha)\}. \quad (\text{A.5})$$

Perfect calibration yields $\widehat{\text{Cov}}(1 - \alpha) = 1 - \alpha$ for all credibility levels $(1 - \alpha)$.

References

- [1] A.J. Wood, B.F. Wollenberg, G.B. Sheblé, *Power Generation, Operation, and Control*, John Wiley & sons, 2013.
- [2] D.S. Kirschen, G. Strbac, *Fundamentals of Power System Economics*, John Wiley & Sons, 2018.
- [3] H. Zareipour, C.A. Canizares, K. Bhattacharya, Economic impact of electricity market price forecasting errors: a demand-side analysis, *IEEE Trans. Power Syst.* 25 (1) (2009) 254–262.
- [4] European Commission Joint Research Centre, Tackling energy price volatility: a smarter approach to price forecasting, 2025, Accessed: 2025-09-10, https://joint-research-centre.ec.europa.eu/jrc-news-and-updates/tackling-energy-price-volatility-smarter-approach-price-forecasting-2025-03-13_en.
- [5] N.P. Padhy, Unit commitment-a bibliographical survey, *IEEE Trans. Power Syst.* 19 (2) (2004) 1196–1205.
- [6] M. Carrión, J.M. Arroyo, A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem, *IEEE Trans. Power Syst.* 21 (3) (2006) 1371–1378.
- [7] C. Gentile, G. Morales-Espana, A. Ramos, A tight MIP formulation of the unit commitment problem with start-up and shut-down constraints, *EURO J. Comput. Optim.* 5 (1) (2017) 177–201.
- [8] A. Mas-Colell, M.D. Whinston, J.R. Green, et al., *Microeconomic Theory*, 1, Oxford university press New York, 1995.
- [9] B.E. Hobbs, Linear complementarity models of Nash-Cournot competition in bilateral and POOLCO power markets, *IEEE Trans. Power Syst.* 16 (2) (2001) 194–202.
- [10] R. Wilson, Architecture of power markets, *Econometrica* 70 (4) (2002) 1299–1340.
- [11] J.R. Birge, F. Louveaux, *Introduction to Stochastic Programming*, Springer, 1997.
- [12] S. Takriti, J.R. Birge, E. Long, A stochastic model for the unit commitment problem, *IEEE Trans. Power Syst.* 11 (3) (1996) 1497–1508.
- [13] Q.P. Zheng, J. Wang, A.L. Liu, Stochastic optimization for unit commitment-a review, *IEEE Trans. Power Syst.* 30 (4) (2014) 1913–1924.
- [14] M. Häberg, Fundamentals and recent developments in stochastic unit commitment, *Int. J. Electr. Power Energy Syst.* 109 (2019) 38–48.
- [15] D. Bertsimas, E. Litvinov, X.A. Sun, J. Zhao, T. Zheng, Adaptive robust optimization for the security constrained unit commitment problem, *IEEE Trans. Power Syst.* 28 (1) (2012) 52–63.
- [16] R.K. Ahuja, J.B. Orlin, Inverse optimization, *Oper. Res.* 49 (5) (2001) 771–783.
- [17] J.R. Birge, A. Hortaçsu, J.M. Pavlin, Inverse optimization for the recovery of market structure from market outcomes: an application to the MISO electricity market, *Oper. Res.* 65 (4) (2017) 837–855.
- [18] Z. Liang, Y. Dvorkin, Data-driven inverse optimization for marginal offer price recovery in electricity markets, in: *Proceedings of the 14th ACM International Conference on Future Energy Systems*, 2023, pp. 497–509.
- [19] L. Meeus, K. Purchala, R. Belmans, Development of the internal electricity market in Europe, *Electr. J.* 18 (6) (2005) 25–35.
- [20] K. Cranmer, J. Brehmer, G. Louppe, The frontier of simulation-based inference, *Proc. Natl. Acad. Sci.* 117 (48) (2020) 30055–30062.
- [21] M. Pirlet, A. Bolland, G. Louppe, D. Ernst, Cost estimation in unit commitment problems using simulation-based inference, in: *NeurIPS 2024 Workshop on Data-driven and Differentiable Simulations, Surrogates, and Solvers*, 2024.
- [22] D. Rezende, S. Mohamed, Variational inference with normalizing flows, in: *International Conference on Machine Learning*, PMLR, 2015, pp. 1530–1538.
- [23] G. Papamakarios, I. Murray, Fast ϵ -free inference of simulation models with bayesian conditional density estimation, *Adv. Neural Inf. Process. Syst.* 29 (2016), pp. 1036–1044.

- [24] J.-M. Lueckmann, P.J. Goncalves, G. Bassetto, K. Öcal, M. Nonnenmacher, J.H. Macke, Flexible statistical inference for mechanistic models of neural dynamics, *Adv. Neural Inf. Process. Syst.* 30 (2017), pp. 1289–1299.
- [25] D. Greenberg, M. Nonnenmacher, J. Macke, Automatic posterior transformation for likelihood-free inference, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 2404–2414.
- [26] C. Grigg, P. Wong, P. Albrecht, R. Allan, M. Bhavaraju, R. Billinton, Q. Chen, C. Fong, S. Haddad, S. Kuruganty, et al., The IEEE reliability test system-1996. a report prepared by the reliability test system task force of the application of probability methods subcommittee, *IEEE Trans. Power Syst.* 14 (3) (1999) 1010–1020.
- [27] C. Ordoudis, P. Pinson, J.M.M. González, M. Zugno, An updated version of the IEEE RTS 24-bus system for electricity market and power system operation studies (2016).
- [28] D.M. Blei, Build, compute, critique, repeat: data analysis with latent variable models, *Annu. Rev. Stat. Appl.* 1 (1) (2014) 203–232.
- [29] A. Delaunoy, B.K. Miller, P. Forré, C. Weniger, G. Louppe, Balancing simulation-based inference for conservative posteriors, [arXiv preprint arXiv:2304.10978](#). *Lingua::EN::Titlecase* (2023).
- [30] R.T. Rockafellar, *Convex Analysis*, Princeton University Press, 1970.
- [31] N. Yang, Z. Dong, L. Wu, L. Zhang, X. Shen, D. Chen, B. Zhu, Y. Liu, A comprehensive review of security-constrained unit commitment, *J. Mod. Power Syst. Clean Energy* 10 (3) (2021) 562–576.
- [32] C. Durkan, A. Bekasov, I. Murray, G. Papamakarios, Neural spline flows, *Adv. Neural Inf. Process. Syst.* 32 (2019) 7511–7522.
- [33] C. Andrieu, G.O. Roberts, The pseudo-marginal approach for efficient Monte Carlo computations *The Annals of Statistics* 37 (2) (2009) 697–725.
- [34] P. Marjoram, J. Molitor, V. Plagnol, S. Tavaré, Markov chain Monte Carlo without likelihoods, *Proc. Natl. Acad. Sci.* 100 (26) (2003) 15324–15328.
- [35] M. Dax, J. Wildberger, S. Buchholz, S.R. Green, J.H. Macke, B. Schölkopf, Flow matching for scalable simulation-based inference, [arXiv preprint arXiv:2305.17161](#). (2023).