



HAL
open science

Adaptive probabilistic forecasting of French electricity spot prices

Grégoire Dutot, Margaux Zaffran, Olivier Féron, Yannig Goude

► **To cite this version:**

Grégoire Dutot, Margaux Zaffran, Olivier Féron, Yannig Goude. Adaptive probabilistic forecasting of French electricity spot prices. 2024. hal-04586509

HAL Id: hal-04586509

<https://hal.science/hal-04586509v1>

Preprint submitted on 24 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Adaptive probabilistic forecasting of French electricity spot prices

Grégoire Dutot*, Margaux Zaffran^{*1,2,3}, Olivier Féron^{1,4}, and Yannig Goude^{1,5}

¹Electricité De France R&D, Palaiseau, France

²PreMeDICAL project team, INRIA Sophia-Antipolis, Montpellier, France

³CMAP, École polytechnique, Institut Polytechnique de Paris, Palaiseau, France

⁴FiME, Université Paris-Dauphine, Paris, France

⁵LMO, Université Paris-Saclay, Orsay, France

Abstract

Electricity price forecasting (EPF) plays a major role for electricity companies as a fundamental entry for trading decisions or energy management operations. As electricity can not be stored, electricity prices are highly volatile which make EPF a particularly difficult task. This is all the more true when dramatic fortuitous events disrupt the markets. Trading and more generally energy management decisions require risk management tools which are based on probabilistic EPF (PEPF). In this challenging context, we argue in favor of the deployment of highly adaptive black-boxes strategies allowing to turn any forecasts into a robust adaptive predictive interval, such as conformal prediction and online aggregation, as a fundamental last layer of any operational pipeline.

We propose to investigate a novel data set containing the French electricity spot prices during the turbulent 2020-2021 years, and build a new explanatory feature revealing high predictive power, namely the nuclear availability. Benchmarking state-of-the-art PEPF on this data set highlights the difficulty of choosing a given model, as they all behave very differently in practice, and none of them is reliable. However, we propose an adequate conformalisation, `OSSCP-horizon`, that improves the performances of PEPF methods, even in the most hazardous period of late 2021. Finally, we emphasize that combining it with online aggregation significantly outperforms any other approaches, and should be the preferred pipeline, as it provides trustworthy probabilistic forecasts.

Keywords: electricity prices, probabilistic forecasting, adaptive forecasting, conformal prediction, online aggregation

1 Introduction

Electricity price forecasting (EPF) plays a major role for electricity companies as a fundamental entry for trading decisions or energy management operations. As electricity can not be stored, electricity prices are highly volatile which make EPF a particularly difficult task (Weron, 2014; Lago et al., 2021).

The increase of renewable production in many countries (RTE, 2022; IEA, 2022a), the development of storage devices or more generally demand response programs (e.g., electrical vehicle smart charging (Nassar et al., 2022), electric water heater management (Amabile et al., 2021; Marin Moreno et al., 2023)) simultaneously entails a need for good EPF and generates more complexity for price modelling. Furthermore, prices can be affected by fortuitous events such as Covid-19 pandemic in 2020-2021 (IEA, 2021), the stress corrosion issue which affected French nuclear power plants in 2022 or the crisis of the gas markets triggered by Russia’s invasion of Ukraine (IEA, 2022b). Trading and more generally energy management decisions require risk management tools which are based on probabilistic EPF (Bunn et al., 2016). This supports the advancement of adaptive probabilistic approaches for forecasting prices, which can continuously learn and adjust to the evolving behaviors of EP, resulting in accurate and reliable probabilistic forecasts.

The literature on EPF is growing rapidly and most papers deals with point forecasts (Weron, 2014; Lago et al., 2021). We focus on short term (day-ahead) EPF as the mainstay of short-term power trading in Europe is the day-ahead market. As proposed in (Lago et al., 2021), models used for forecasting electricity prices can be categorized as either statistical, machine learning or hybrid models.

Statistical models are dominated by auto-regressive models and their variants, in particular the state of the art Lasso Estimated AutoRegressive (LEAR) model proposed by Uniejewski et al. (2016) and recently used as state of the art benchmark in (Lago et al., 2021; Tschora et al., 2022). It consists in a high dimensional ARX model where the fitting process is done by minimizing an elastic net regularization. The high dimension (around 250 parameters) comes from a large number of lags of prices and forecasts of variable of interests (generation, zonal prices, consumption). As highlighted by Lago et al. (2021) pre-processing of EP such as log transformations or more generally variance stabilizing transformations (Uniejewski et al., 2018) are a common practice to deal with heavy tailed distribution. Regarding

*These authors contributed equally to this work.

non-stationarity of the prices, regime switching ARX models are proposed in (Nitka et al., 2021). Marcjasz et al. (2018) propose to average a set of point forecasts obtained from learning with different time windows to derive probabilistic forecasts.

The utilisation of **machine learning** tools including deep learning approaches for electricity price forecasting (EPF) has grown over the past decade. Recent studies (Tschora et al., 2022; Jędrzejewski et al., 2022) reveal that complex ML methods such as deep neural networks can achieve better forecasting performances than the LEAR model at the cost of significantly higher computational cost. The relatively important dimension of these models require a significant amount of data for their calibration, making them poor candidate to adapt to abrupt changes in price distribution (Çağatay Berke Bozlak and Yaşar, 2024). Yang et al. (2023) show how graphical neural network could be used to model spatial dependency to forecast the day-ahead electricity prices of the Nord Pool market.

Probabilistic price forecasting is progressively becoming more popular in the forecasting literature following the GEFCom2014 energy forecasting competition (Hong et al., 2016). This is a natural goal as the final objective EPF is to optimize a financial risk criteria (Bjorgan et al., 1999; Deschatre et al., 2021). Most of the previous parametric statistical models are based on statistical assumptions and could be naturally extended to produce probabilistic forecast (more or less accurate as we will explore in this paper). Relaxing distributional assumption, non parametric regression models such as quantile regression have been investigated (Uniejewski and Weron, 2021). In Loizidis et al. (2024), machine learning models coupled with bootstrap methods are compared with classical time series models for German and Finnish day-ahead market. Marcjasz et al. (2023) recently proposed a distributional network that outperforms state-of-the-art benchmarks. Nickelsen and Müller (2024) present a Bayesian forecasting framework for the German continuous intraday market and show that orthogonal matching pursuit methods can outperform LEAR. Cornell et al. (2024) propose quantile regression with varying training-length periods and model averaging to forecast prices of the South Australia region of the Australian National Electricity Market.

PEPF models face many pitfalls: extreme price spikes, non-stationarity due to exogenous factors inducing time-varying mean and/or volatility. Conformal methods (Vovk et al., 1999; Papadopoulos et al., 2002; Vovk et al., 2005) and more specifically adaptive conformal methods, proposed for example by Gibbs and Candès (2021); Zaffran et al. (2022), are a way to adapt PEPF models in a very general way. It can be applied to any of the previously cited PEPFs to improve them. We propose to extend the work of Zaffran et al. (2022) to forecast electricity prices in France during the turbulent period 2020-2022. Another framework allowing to adapt PEPF models is online aggregation under expert advice

(Cesa-Bianchi and Lugosi, 2006), which was successfully used in financial non-stationary environments (Remlinger et al., 2023; Berrisch and Ziel, 2024a). Our aim is to investigate if and how it is possible to make adaptive an existing probabilistic forecasting algorithm. This approach is driven by an operational concern: proposing a plug-in tool that can be applied to any underlying model eases its integration in the current pipeline.

Contributions We list below our main contributions:

- **New data:** we study the recent turbulent period 2020-2022 and we add a new feature, the nuclear availability
- **Benchmark:** we consider state-of-the-art PEPF methods, their windowed versions (rolling window estimation) and benchmark them on this new dataset
- **Analysis** of the improvements (or not) of existing **online conformal methods**
- Suggestion of **novel online conformal strategy** coined `OSSCP-horizon`
- Unified framework of **sequential aggregation** of all these probabilistic forecasting
- **Understanding the benefits** of these 2 frameworks of probabilistic post-processing (i.e. CP and aggregation) and how they can help each other: *sequential aggregation with conformalized expert is the best*

2 Data presentation and insightful new explanatory variables

2.1 Dataset’s description

The considered dataset spans approximately 6 years of observations at a hourly frequency, from January 11th, 2016 to December 31st, 2021, and is decomposed of a training set (from January 11th, 2016 to December 31st, 2018) to estimate the parameters of the models, a validation test (year 2019) to estimate the hyperparameters, and a test set (years 2020 and 2021) to evaluate the performances (see Figure 1). We consider the task of forecasting day-ahead (DAH) prices on the French EPEX market. As the 24 hours of day d are fixed from EUPHEMIA¹’s market clearing at 12:00pm of day $d - 1$, the features considered to predict each of them are selected so that they are available before 12:00pm of day $d - 1$. More precisely the dataset contains the following features, for a target at day d , hour h :

- the 24 French DAH prices at days $d - 1$ and $d - 7$;

¹EUPHEMIA is the algorithm that solves the market coupling problem for the Central West European region, used by EPEX to compute the day-head power prices




Figure 1: Evolution of the Spot prices (first panel), Residual Load (second panel), Nuclear availability (third panel) and commodity prices (last panel) from 2016 to 2021 (x-axis).

- the observed daily price of Gas on the French PEG market at $d - 1$ and the month-ahead futures prices for Oil (Brent) and Coal (CIF ARA Argus-McCloskey);
- the forecasted residual load signal built with data available before 12pm at $d - 1$: the load forecasts for the 24 hours of day d , estimated on day $d - 2$, minus the renewable production forecasts (i.e., wind and solar forecasts estimated on day $d - 2$, and the observed run-of-river electricity on $d-2$);
- the availability of French nuclear electricity on day d , i.e. the announced available capacity of nuclear generation;
- the observed electricity generation from all production types at $d - 2$ and $d - 7$ (in the case of nuclear energy, the production is divided by the nuclear availability);
- the EUR vs. GBP and EUR vs. USD exchange rate (last observed at $d - 1$);
- the total electricity volume exchanges between France and all its neighbors (observed at $d - 2$);
- the specific electricity volume exchanges between France and Germany (observed at $d - 2$);
- dummy variables, including dummy variables for French holidays (as a percentage of the total population concerned), holiday bridges, weekends, and weekdays;
- the time of year as a sine and cosine function, as well as a clock variable to capture a possible trend.

2.2 First point forecast and feature importance

The proposed dataset comprises features classically used to forecast electricity prices, and also a new feature, the nuclear availability, for we intuit that nuclear availability has a significant impact on DAH prices due to the French energy mix.

At first we proceed a point forecast exercise, with Lasso CV and Random forest models, to detect the most important features and highlight the relevance of the proposed new variables. Here, the meaning of the term “feature importance” varies according to the model: in the case of Lasso CV, it refers to the value of the coefficient associated to a given feature, whereas for Random Forest it refers to the Mean Decrease in Impurity (MDI).

In Figure 2, we observe the top 20 mean feature importances over both models trained in 2020. Spot price at H-23 of the previous day is the “most important” feature for the Lasso CV model. This is coherent with what is found in

(Maciejowska et al., 2022; Ziel and Weron, 2018). The MDI-based importances computed for the Random Forest suggests the same conclusion, even though high correlation between all $d-1$ spot prices makes the interpretation harder. The Lasso CV model, which allows for a better modelisation with highly correlated features, suggests that gas prices and nuclear availability have a high explanatory power. This speaks in favour of an inclusion of these features in EPF prediction models, at least in the case of the French market.

We also compute the feature importance of both model over every days in the test period and observe the evolution in the predominance of the various feature groups. To do so, we first aggregate features into groups: “Change” for all exchange rates, “Commodity price” for gas, coal and oil prices, “Exchange” for all hourly power volumes exchanges, and the rest of features groups are hourly features aggregated at a daily level. The group aggregation consists in summing up the absolute importance value of all features belonging to this group, then normalize these values by the total sum over all groups. Figure 3 represents the evolution we obtain. We observe a considerable change in the relative group’s explanatory power: for both the Random Forest and Lasso model, we observe a significant increase in the aggregated explanatory power of the commodity prices, at the expense of the residual load forecast. This indicates an important

distribution shift in the relationships between the times series by September 2021.

3 Probabilistic forecasting methods

Notations Given the nature of the data and in particular the hourly patterns, we will build one model per hour, as explained in Section 5.1. From now on, the temporal index t is used and it elapses at a daily rate (i.e., for a given hour h). $t = 1$ corresponds to the beginning of the training data, $t = T_0$ marks the end of the training data and $t = T_1$ refers to the last test observation to be predicted. In other words, we aim at predicting the French spot prices between $T_0 + 1$ and T_1 , corresponding to the years 2020 and 2021 (see Figure 1).

3.1 Framework

One objective of probabilistic forecast is to build *Prediction Intervals* (PIs) for a variable Y_t depending on the covariates X_t . Let $\alpha \in [0, 1]$ be a *miscoverage rate*. A PI at the $1 - \alpha$ level is expected to contain at least $1 - \alpha$ of the realisations: $\mathbb{P}(Y_t \in \text{PI}_{1-\alpha}(X_t)) \geq 1 - \alpha$, while being as small as possible. In order to retrieve as much information as possible about the distribution of Y_t , one can consider multiple values of the miscoverage rate α .

A PI can be characterized by two “point forecasts”: its lower




Figure 2: Feature (y -axis) importance (x -axis) for Lasso CV (left panel) and Random Forest (right panel) models. The colors are associated with a type of feature.




Figure 3: Evolution of normalized feature importance (y -axis) for Lasso CV (left panel) and Random Forest (right panel) models over the whole test period (x -axis). The colors are associated with the features.

($\ell(X)$) and upper ($u(X)$) bounds. A natural choice for the PI is $\ell(X) = Q_{\alpha/2}(X)$ and $u(X) = Q_{1-\alpha/2}(X)$, where Q_β is the β -th quantile of the cumulative function distribution (c.d.f.) of the price conditionally to the covariates used to forecast.

However, in practice, these true Q are never known and we have to estimate them, e.g., using quantile regression (Koenker, 2005). This approach is detailed in Section 3.2.

Another path is to post-process individual predictors (see Section 3.3). The individual predictors can either estimate the mean as in point forecasting and the post-processing step will turn them into PI, or directly estimate a conditional quantile (as described in Section 3.2).

3.2 Quantile regression methods

We present here the quantile regression methods that we retained for our benchmark study. These methods were chosen for their good performance on time series data, and in particular on electricity related data. They are all quite easy to fit automatically and have a relatively low computational cost (this is a key asset due to the intensive benchmark including rolling window estimation).

3.2.1 DESCRIPTION OF THE METHODS

Basics on Quantile Regression (QR) QR (Koenker, 2005) replaces the usual quadratic loss by the *pinball loss* to forecast a conditional quantile of the distribution of Y (i.e. the price) given the features X :

$$\min_{g \in \mathcal{G}} \mathbb{E} [\rho_\beta(Y - g(X)) | X = x],$$

for any x , with ρ_β the *pinball loss* of level β : $\rho_\beta(y - \hat{y}) = (1 - \beta)|y - \hat{y}| \mathbb{1}\{y \leq \hat{y}\} + \beta|y - \hat{y}| \mathbb{1}\{y \geq \hat{y}\}$, and \mathcal{G} the class of regressors considered, e.g. linear models, Lasso (QLR-Lasso), additive non-linear models (QGAM) or gradient boosting regressors (QGB).

Quantile Linear Regression (Linear QR) and Quantile Lasso (Lasso QR) The class of regressors \mathcal{G} is restricted to linear models. For Lasso QR, We perform a Lasso selection process (Tibshirani, 1996) to deal with the pretty high number of covariates, the class of regressors is thus the linear models on all possible subsets of covariates.

Quantile Generalized Additive Models (QGAM) Generalized Additive Models (GAMs) (Hastie and Tibshirani, 1986) consists in explaining the conditional expectation $\mu(X)$ of Y over X with a semi-parametric additive structure. The estimation of GAMs is based on a (regularized) mean squared error (MSE) criterion. Our objective is to use GAMs for a QR problem. One could replace the MSE by the pinball loss function in the estimation process as described in the previous paragraph. However, Fasiolo et al. (2020) demonstrate that the pinball loss is statistically sub-optimal in this framework and propose a procedure based on the smooth Extended Log-F loss instead.

Quantile Random Forests (QRF) Meinshausen (2006) adapts Random Forests to the QR task. The same forest is built than for mean-regression, that is a forest grown in order to minimize the mean squared error. However, to adapt to the quantile task at hand, the final decision rule for prediction now corresponds to evaluating an empirical conditional quantile (conditional on the fact that the features of the test point belongs to the corresponding leaves).

Quantile (tree based) Gradient Boosting (QGB) Gradient boosting machine (Friedman, 2001) are widely used in the forecasting community where it has demonstrated excellent performance for different applications on tabular data (Grinsztajn et al., 2022) or time series (Makridakis et al., 2022). As for the Random Forests, the regressors are here regression trees. The boosting algorithm consists in adding a sequence of simple models (called weak learners and trained on a subsample randomly selected of the training set) obtained by sequentially fitting a quantile regression tree to the residuals by minimizing the pinball loss, which is a key difference with QRF.

3.2.2 OPERATIONAL PIPELINE

We explore these prediction methods through their implementation in the Python package `scikit-learn` package (Pedregosa et al., 2011) for linear quantile regression, Lasso and QGB. QRF are implemented through `scikit-garden`. The QGAM are implemented in the R package (Fasiolo et al., 2021).

All of these models depend on hyper-parameters, and QGAM additionally requires an exact formula. In particular, we optimized for the regularizer (Lasso), the number of trees and their maximum depth (QRF and QGB), as well as the learning rate and fraction of samples (QGB), and the formula (QGAM). Their estimation is based on grid-searching on the validation set after estimation of mean-regression models on the training set, as illustrated in Figure 1. Therefore, the formula of the QGAM is the same for all quantiles. It includes:

- *linear effects*: for the indicator of the week days;
- *univariate non-linear terms*: the announced French nuclear availability, the lagged 2 days of the fossil hard coal and observed nuclear productions, the square root of the lagged one day of the Gaz prices, cosin and sin of the time of year;
- *functional smooth effects*: as proposed in Amara-Ouali et al. (2023) in the context of electricity load forecasting, we model the lagged (one day and one week) prices and the load forecast effects via a functional smooth effect. It allows to capture the effect of these functional (in function of time) covariates over the price at a given instant of the day.

In this paper we do not consider online re-estimation of the hyperparameters, which in practice is very time consuming and statistically challenging. We study the performance of operational fixed prediction models that can be made adaptive through a plugged-in layer, useful when facing non-stationarity without completely retraining them.

Also, as illustrated in the preliminary results of Figure 4, before September 2021, only QRF and QGAM achieved *validity*. We explore strategies to recover validity in Section 3.3. What is more, none of the probabilistic methods attain the target coverage level after September 2021. Indeed, the high explosion of the prices after this date, both in average and in variability, calls for more adaptive strategies, that we discuss in Section 4. Note that the standard rolling training procedure did adapt to this change as illustrated by the lengths of the PIs after September 2021, but more adaptiveness is required given the strength of the shift and variability.

3.3 Conformal methods: add-on to traditional probabilistic approaches

Conformal Prediction (CP) (Vovk et al., 1999; Papadopoulos et al., 2002; Vovk et al., 2005) builds PI around any kind of prediction models. These intervals are valid (achieving marginal nominal coverage) in finite samples under the only assumption of exchangeability of the data. Therefore, CP has to be seen as an add-on protective layer to existing probabilistic (or not) forecasts, that is able to robustify them in terms of validity but whose efficiency and shape will

always rely on the quality of the underlying forecast.

Suppose that we have T_0 random variables $\{(X_t, Y_t)\}_{t=1}^{T_0}$. For a given miscoverage rate $\alpha \in [0, 1]$, we aim at building a *marginally valid* PI \hat{C}_α of Y_{T_0+1} , i.e. \hat{C}_α should satisfy:

$$\mathbb{P}(Y_{T_0+1} \in \hat{C}_\alpha(X_{T_0+1})) \geq 1 - \alpha. \quad (1)$$

To achieve this, Split Conformal Prediction (SCP) (Papadopoulos et al., 2002; Lei et al., 2018) randomly splits the T_0 data points into a training set Tr and a calibration set Cal . A regression model $\hat{\mu}$ is then fitted on Tr and used to predict on Cal to obtain a set of conformity scores $\mathcal{S}_{\text{Cal}} = \{S_t := s(X_t, Y_t; \hat{\mu}), t \in \text{Cal}\}$. These scores assess the conformity between the calibration’s observed values and the predicted ones: the smaller the better. In the case of regression, they are usually computed using the absolute value of the residuals, i.e. $S_t := s(X_t, Y_t; \hat{\mu}) = |\hat{\mu}(X_t) - Y_t|$. A corrected² $(1 - \tilde{\alpha})$ -th empirical quantile of the conformity scores $Q_{1-\tilde{\alpha}}(\mathcal{S}_{\text{Cal}})$ is obtained, to finally build the prediction interval $\hat{C}_\alpha := \{y : s(X_{T_0+1}, y; \hat{\mu}) \leq Q_{1-\tilde{\alpha}}(\mathcal{S}_{\text{Cal}})\}$. In the standard regression case, it boils down to $\hat{C}_\alpha(X_{T_0+1}) = [\hat{\mu}(X_{T_0+1}) \pm Q_{1-\tilde{\alpha}}(\mathcal{S}_{\text{Cal}})]$. This procedure is guaranteed theoretically to satisfy Equation (1) for any model $\hat{\mu}$, any sample size T_0 , as long as the calibration and test data are exchangeable.

Proposed by Romano et al. (2019), Conformalized Quantile Regression (CQR) benefits simultaneously from the adap-

²The correction $1 - \tilde{\alpha} = (1 - \alpha)(1 + \frac{1}{\#\text{Cal}})$ is needed to ensure finite sample validity, because of the inflation of the quantiles.

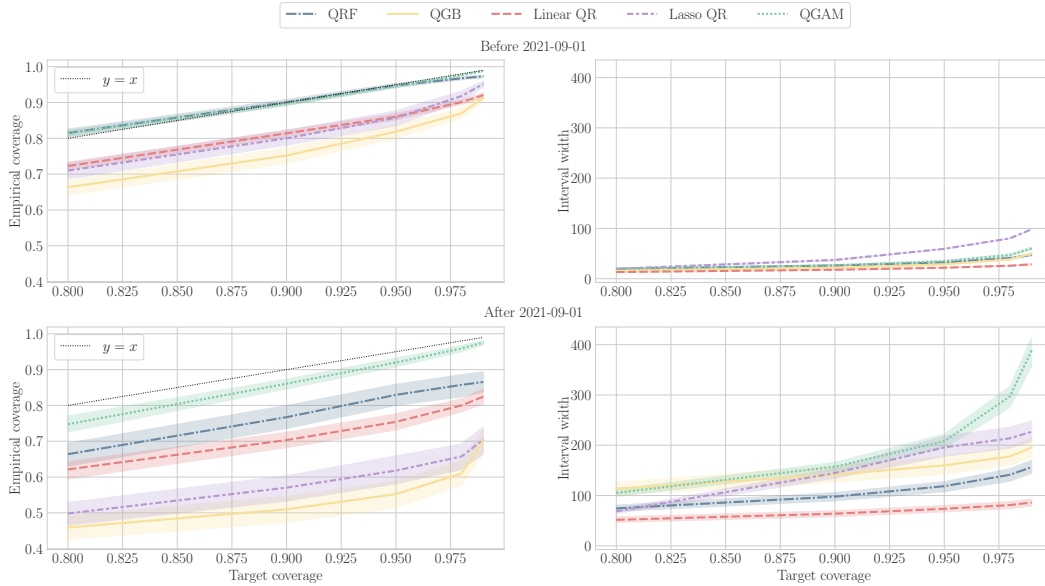


Figure 4: PIs’s performance of individual probabilistic forecasts at test time, before September 2021 (top row) and after September 2021 (bottom row), for various target coverage levels (x -axis). The left column represents the average empirical coverage: the closest to the $y = x$ line the better, and above it is best. The right column represents the average interval width: the lower the better. The colors and shapes are associated with the models. The shaded regions correspond to the 5% and 95% empirical quantiles after bootstrapping 500 times the test time series, see Section 5.1 for details.

tiveness of classical QR methods and from the theoretical guarantees ensured by CP. Instead of training a mean regression model on the training set Tr , CQR requires to fit two conditional quantile regression models $\hat{q}_\ell(\cdot), \hat{q}_u(\cdot)$ ³. In this context, the conformity scores now quantify the error made by the fitted PI $\hat{C}(x) := [\hat{q}_\ell(x), \hat{q}_u(x)]$. Precisely, $S_t := s(X_t, Y_t; \hat{q}_\ell, \hat{q}_u) = \max\{\hat{q}_\ell(X_t) - Y_t; Y_t - \hat{q}_u(X_t)\}$. Accordingly, the PI becomes $\hat{C}_\alpha(X_{T_0+1}) = [\hat{q}_\ell(X_{T_0+1}) - Q_{1-\alpha}(\mathcal{S}_{\text{Cal}}), \hat{q}_u(X_{T_0+1}) + Q_{1-\alpha}(\mathcal{S}_{\text{Cal}})]$.

To account for the temporal aspect of time series, an online and sequential version of SCP is usually considered, in which the split leading to Tr and Cal is not random, but constrained so that any point in Tr occurs before any point in Cal (Wisniewski et al., 2020; Zaffran et al., 2022). See Figure 5 for an illustration.

4 Adaptiveness as a wrapper around individual forecasts

The online setting—in which the environment reveals the true value before the next prediction—allows to post-process individual predictors to adapt to previous errors (e.g., as done in CP). This approach demonstrates all its interest when stationarity – and consequently neither exchangeability – does not hold, as in our case study. One way to implement such a post-processing, coming from the online literature, is online aggregation of predictors, as described in Section 4.1⁴. Another strategy, within the CP framework, is to modify the calibration step of CP (see Section 4.2) and make it adaptive.

4.1 Online aggregation based strategies

Adaptive aggregation of *experts* (Cesa-Bianchi and Lugosi, 2006), with $K \in \mathbb{N}^*$ experts denoted $(\hat{f}_t^{(k)}(\cdot))_{k \in [1, K]}$ being various individual forecasters for the prices at time t (that is a corresponding day d on a given hour h) such as the ones introduced in Section 3.2, computes an optimal weighted mean of the experts. At each time t (i.e., day d , for a given hour h), the weights $\omega_t^{(k)}$ assigned to expert k depend on all experts' suffered *losses*, i.e. their performances on the previous time steps until $t-1$. In our case, these performances are evaluated through the pinball loss ρ_β , standard in quantile regression, with the pinball parameter β being the target quantile level. These losses are plugged in the *aggregation rule* Φ , outputting the aggregation weights. Finally, the aggregation rule can include the computation of the gradients of the loss (*gradient trick*, see (Cesa-Bianchi and Lugosi, 2006) for more details). As aggregation rules require bounded experts, a thresholding step is added. Concretely, the aggregated predictor at time t , $\hat{f}_t^\Phi(\cdot)$, is defined

³Usually $\ell = \alpha/2$ and $u = 1 - \alpha/2$, but this is not necessary. Romano et al. (2019) suggest to choose these values by cross-validation, to improve PI's efficiency.

⁴This does not include Quantile Regression Averaging (QRA) (Nowotarski and Weron, 2014) as it is an offline averaging, thus non-adaptive.

by

$$\hat{f}_t^\Phi(X_t) = \sum_{k=1}^K \omega_t^{(k)} f_t^{(k)}(X_t).$$

In our experiments, the different forecasts obtained are aggregated quantile by quantile, using the appropriate pinball loss as a score. The aggregation rule Φ is set to be the Bernstein Online Aggregation (BOA) (Wintenberger, 2017) algorithm, along with the gradient trick. We use the R package OPERA (Gaillard and Goude, 2016) to perform such an aggregation, and reorder the quantiles predicted by the aggregation models to avoid quantile crossing.

Recently, Berrisch and Ziel (2021) proposed an approach that jointly aggregates every quantile forecasting model together and gives directly a probabilistic prediction as an output, instead of performing independent aggregation for each quantile level. Berrisch and Ziel (2021)'s method reduces the number of aggregation parameters to be computed, while yielding preferable probabilistic performances. It is available in the R-Package `profoc` (Berrisch and Ziel, 2024b), compatible with the BOA method with the gradient trick and automatically reordering the predicted quantiles. It has to be noted that we did not explore the full range of tuning possibilities allowed by this method. In our experiments, both approaches performed similarly. Therefore, to avoid overloading the analysis, we present in this paper only the first method.

4.2 Adaptive conformal approaches

In addition to online aggregation, we consider another post-processing of individual forecasters which consists in adding a conformal layer on top of them, adaptively. As explained in Section 3.3, CP requires exchangeable data, an assumption clearly not satisfied in a time series setting, and even less in our highly non-stationary case study.

The first theoretically grounded result on CP for dependent data is given by Chernozhukov et al. (2018): it shows that when the data is strongly mixing and the learned model is close “enough” to the underlying data generation process then CP guarantees still hold, along with proposing an extension for full CP⁵ under which the previous theorem holds. Again, this is not sufficient to encapsulate our setting.

In practice, Online Sequential Split Conformal Prediction (OSSCP) is often used to take into account the temporal structure, introduced in Wisniewski et al. (2020); Zaffran et al. (2022). The idea is (i) to enforce a sequential split where all the training observations are temporally consecutive, and preceding the ones of the calibration set and (ii) to update this split in order to incorporate the newly observed

⁵Full CP is a version of CP that does not require to split the data, at the cost of a bigger computational burden. This is the reason why we do not consider it in this work, along with the fact that full CP can be plugged in on an existing pipeline, making it particularly appealing for operational purposes. The interested reader on full CP can have a look at (Vovk et al., 2005)

data points at each prediction step $t + 1$, forgoing the oldest ones, leading to adaptive sets Tr_t and Cal_t . See Figure 5 (a) for an illustration. Note that OSSCP does not enjoy any form of theoretical guarantees beyond the exchangeable setting, despite its good empirical performances in the time series framework, as highlighted in (Zaffran et al., 2022).

4.2.1 IMPROVING CP ONLINE ADAPTIVENESS: OSSCP-HORIZON

One drawback of OSSCP is that the set on which the models were fitted can be far from the points on which it will be applied (either calibration or test points). If the temporal data suffers from a strong distribution shift, this may hinder the accuracy of the base learner, and therefore the performances of the PI, both in terms of coverage (the exchangeability assumption is not satisfied anymore) and in terms of efficiency, i.e. interval's length (as large errors cause large intervals).

In order to avoid high errors on the calibration and test points, we propose a new approach, coined OSSCP-horizon. The idea is to ensure that the underlying model is trained on the data just preceding each calibration point: in other words, to only compute test errors of horizon one, as is the forecast horizon. More generally, for any forecasting task at horizon h , OSSCP-horizon computes calibration errors of horizon h . See Figure 5 (b) for an illustration. Formally, at prediction time $T + 1$, OSSCP-horizon thus builds the calibration set as follows:

- For each $X_t \in \text{Cal}_T$, fit quantile regression estimators $\hat{q}_\ell^{-(t)}$, $\hat{q}_u^{-(t)}$ on $\{(X_{t-|\text{Tr}|}, Y_{t-|\text{Tr}|}), \dots, (X_{t-1}, Y_{t-1})\}$ ⁶;
- Compute the calibration score $S_t = s(X_t, Y_t; \hat{q}_\ell^{-(t)}, \hat{q}_u^{-(t)})$ and add it to the set of scores $\mathcal{S}_{\text{Cal}_T}$.

After having built $\mathcal{S}_{\text{Cal}_T} = \{S_{T-|\text{Cal}|+1}, \dots, S_T\}$, OSSCP-horizon computes the PI for the test point X_{T+1} :

$$\hat{C}_\alpha(X_{T+1}) := \left[\hat{q}_\ell^{-(T+1)}(X_{T+1}) - Q_{1-\tilde{\alpha}}(\mathcal{S}_{\text{Cal}_T}); \hat{q}_u^{-(T+1)}(X_{T+1}) + Q_{1-\tilde{\alpha}}(\mathcal{S}_{\text{Cal}_T}) \right].$$

Again, while demonstrating empirical improvements upon standard OSSCP in the temporal setting, OSSCP-horizon does not enjoy any form of theoretical guarantees. To theoretically account for the online setting, a popular method is Adaptive Conformal Inference (ACI) (Gibbs and Candès, 2021).

4.2.2 ADAPTIVE CONFORMAL INFERENCE (ACI)

Proposed in (Gibbs and Candès, 2021), ACI adapts CP to an arbitrary online setting, including temporal distribution shifts.

⁶For a horizon $h \neq 1$, then $\hat{q}_\ell^{-(t)}$, $\hat{q}_u^{-(t)}$ are fitted on $\{(X_{t-|\text{Tr}|}, Y_{t-|\text{Tr}|}), \dots, (X_{t-h}, Y_{t-h})\}$.

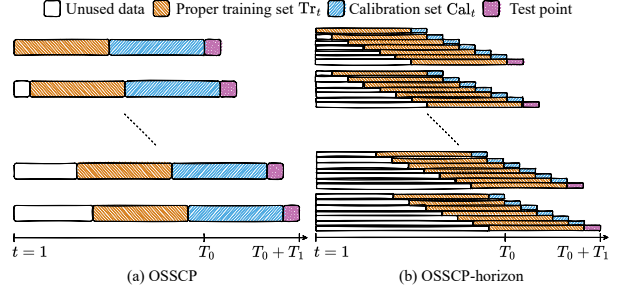


Figure 5: Scheme of OSSCP (a) and our proposal (b), OSSCP-horizon, when the horizon is 1.

To do so, ACI recursively updates the *effective* miscoverage rate $\tilde{\alpha} := \alpha_t$ used in the computation of the PI. Set $\alpha_1 = \alpha$. For $t \geq T_0$, and for a chosen $\gamma \geq 0$ the ACI update formula is:

$$\begin{cases} \hat{C}_{\alpha_t}(X_t) := [\hat{q}_\ell(X_t) - Q_{1-\alpha_t}(\mathcal{S}_{\text{Cal}_t}), \hat{q}_u(X_t) + Q_{1-\alpha_t}(\mathcal{S}_{\text{Cal}_t})] \\ \alpha_{t+1} = \alpha_t + \gamma \left(\alpha - \mathbb{1}_{\{Y_t \notin \hat{C}_{\alpha_t}(X_t)\}} \right) \end{cases}$$

The underlying idea is the following. If the PI does not cover at time t , then $\alpha_{t+1} \leq \alpha_t$ which increases the size of the PI. Conversely, the size of the interval decreases gently at time $t + 1$ when it covers at time t . As noted in (Zaffran et al., 2022), it is possible to have $\alpha_t \geq 1$ or $\alpha_t \leq 0$: the former case is quite rare and produces by convention $\hat{C}_{\alpha_t} = [\hat{q}_\ell(\cdot), \hat{q}_u(\cdot)]$; however, the latter can happen frequently, especially for a high γ , giving a prediction interval of infinite size ($\hat{C}_{\alpha_t} \equiv \mathbb{R}$).

The main theoretical result on ACI is that for any sequence $(X_t, Y_t)_t$, $\left| \frac{1}{T_1 - T_0} \sum_{t=T_0+1}^{T_1} \mathbb{1}_{\{Y_t \in \hat{C}_{\alpha_t}(X_t)\}} - (1 - \alpha) \right| \leq \frac{2}{\gamma(T_1 - T_0)}$. It shows the asymptotically valid frequency of ACI intervals for any arbitrary (possibly adversarial) distribution.

Note that the convergence rate is in γ^{-1} , hence favoring large γ which are the ones leading to more variability and in the extreme case to infinite PIs (discussed previously). This illustrates the need for guidance on how to choose properly γ , and even avoid having to choose it and being able to switch between different γ depending on the current data distribution's evolution.

4.2.3 AGACI

The goal of AgACI, proposed in (Zaffran et al., 2022), is precisely to provide a parameter-free method based on ACI, that can adapt to temporal changes in the data distribution adaptively. Given a list of K γ values $\{\gamma_k\}_{k=1}^K$, AgACI works as an adaptive aggregation of experts (Cesa-Bianchi and Lugosi, 2006) (see also Section 4.1), with expert k being ACI with parameter γ_k . At each prediction step t , it performs two independent aggregations of the K ACI intervals $\hat{C}_{\alpha_{t,k}}(\cdot) \stackrel{\text{not.}}{=} [\hat{b}_{t,k}^{(\ell)}(\cdot), \hat{b}_{t,k}^{(u)}(\cdot)]$, one for each bound, and

outputs $\tilde{C}_t(\cdot) \stackrel{\text{not}}{=} [\tilde{b}_t^{(\ell)}(\cdot), \tilde{b}_t^{(u)}(\cdot)]$. According to Zaffran et al. (2022), the standard different aggregation rules gave similar results. In this work, we restrict ourselves to the setting of (Zaffran et al., 2022), that is BOA, with the gradient trick.

4.2.4 LATEST RELATED WORKS

Since the analysis presented in this paper was performed, the line of research on adaptive and online conformal approaches has been expanding fast. Recent developments include: Gibbs and Candès (2023) improving on ACI by online aggregation on a grid of different γ , similarly to AgACI, at the crucial difference that the aggregation is on the value of α_t and not on the lower and upper bounds independently (Section 5.2 highlights why we argue in favor of different aggregations); Bastani et al. (2022) who achieve stronger coverage guarantees (conditional on the effective level, and conditional on specified subsets of the explanatory variables); Bhatnagar et al. (2023) enjoy anytime regret bound, by leveraging tools from the strongly adaptive regret minimization literature; Angelopoulos et al. (2023) who extend upon ACI ideas by relying on control theory to add more information on the temporal structure; Angelopoulos et al. (2024) proposing to use adaptive learning rates γ_t in ACI.

Our goal in this analysis is to deeply investigate the improvements, or not, brought by conformal as one of the layers for probabilistic forecasts with an operational lens. Therefore, we restricted the study to OSSCP, OSSCP-horizon, and AgACI as it has already shown benefits on electricity prices and does not require to select any hyper-parameter (Zaffran et al., 2022). Indeed, it allows us to easily understand what is the cause of the improved or declined performance. Furthermore, the most recent works are either complex structures (thus less interpretable) or depend on hyper-parameter tuning, making them more costly to implement in operational use.

5 Application and results

5.1 Setting and evaluation

Experimental details In order to span a wide range of the price distribution function, we vary the PIs' miscoverage level $1 - \alpha > 0.6$. For the final probabilistic forecasts, the overall training set comprises 4 years of data, from 2016 to 2019 included (i.e. merging the training and validation sets).

Due to training time constraints, we trained and evaluated the considered models on hours 3, 8, 13, 18, and 23 of every day. These 5 hours encompass best the different phases of hourly electricity prices in a given day, while uniformly covering the 24 hours of the day.

Finally, due to the high non-stationarity, we trained each of the base models presented in Section 3.2 on different window sizes: approximately 4 years, 3 years, 2 years, 1 year, 270 days, 180 days, and 90 days. For the sake of clarity,

for each analysis performed, the largest window size will be selected and presented in this paper. In the same vein, the calibration size of the conformal approaches (Sections 3.3 and 4.2) varies among 25%, 50% and 75% of the overall windowed training set. Again, to ease interpretation of our results, we present here only the results for a calibration set of proportion 50% (except if stated otherwise) as it allows for an intermediary adaptation speed, hence being a good trade-off between up-to-date quantile regression models and calibration set large enough to perform the estimation of the highly non-stationary conformal correction. We recall that in the i.i.d. setting a general rule of thumb for the calibration size is around 25% (Sesia and Candès, 2020). In our study, the impact of non-stationarity induces a need for a trade-off between adaptivity and the calibration window length.

Evaluation procedure The main challenge of evaluating a probabilistic forecast is that the true distribution of the underlying process cannot be observed. Hence, it is impossible to compare the estimated distribution with the actual distribution of the true spot prices. This is not the case for a sequence of PIs $([\hat{b}^{(\ell)}(\cdot), \hat{b}^{(u)}(\cdot)])_t$ that can be evaluated through:

- **empirical average coverage**,

$$\frac{1}{T_1 - T_0} \sum_{t=T_0+1}^{T_1} \mathbb{1} \{y_t \in [\hat{b}^{(\ell)}(x_t), \hat{b}^{(u)}(x_t)]\}, \text{ that should be close and above to the target level } 1 - \alpha \text{ for } \textit{validity} \text{ (also known as } \textit{reliability}),$$

- **empirical average length**, $\frac{1}{T_1 - T_0} \sum_{t=T_0+1}^{T_1} \hat{b}^{(u)}(x_t) - \hat{b}^{(\ell)}(x_t)$, for *efficiency*⁷ (also known as *sharpness*).

For each of these metrics, confidence intervals are constructed by time series bootstrapping (non-overlapping moving block bootstrap) (Kunsch, 1989; Politis and Romano, 1994).

Results on the CRPS are provided in A. Indeed, our goal is really to compare PIs and not predictive distributions. Therefore, the forecasts' objective is truly to **be as sharp as possible while satisfying validity**.

5.2 Results

Impact of the conformalisations In Figures 6 and 7 we represent the performance of Linear Quantile Regression and Quantile Random Forest respectively, with various layers of conformalisation. The display choice of these two base models is motivated by the fact that they represent a diverse range of modelisation.

⁷Indeed, achieving exactly $1 - \alpha$ coverage can be trivially done by outputting $1 - \alpha$ of the time \mathbb{R} and the empty set otherwise, which is critically uninformative. Thus, one wants to attain *validity* while minimizing the size of the resulting intervals, that is maximizing *efficiency*.

In both cases, we observe that a naive conformalisation – in the form of OSSCP – does not allow to achieve the nominal coverage level, neither before nor after September 2021.

Yet, our proposal OSSCP-horizon does improve drastically the coverage level: before September 2021 it manages to reach the target level while improving the lengths of the PIs, and after September 2021 it allows to reduce the gap with the target considerably (linear model), while recovering the approximately satisfactory performances of the

individual QRF that was deteriorated by OSSCP.

Finally, making the conformalisation even more adaptive through the use of AgACI especially enhances validity after September 2021. Yet, it has to be noted that it seems to be insufficiently adaptive to perfectly reach the target level.

Analysis of various aggregations Therefore, we go further and add another adaptive post-processing layer by performing online aggregation. In Figure 8 we compare

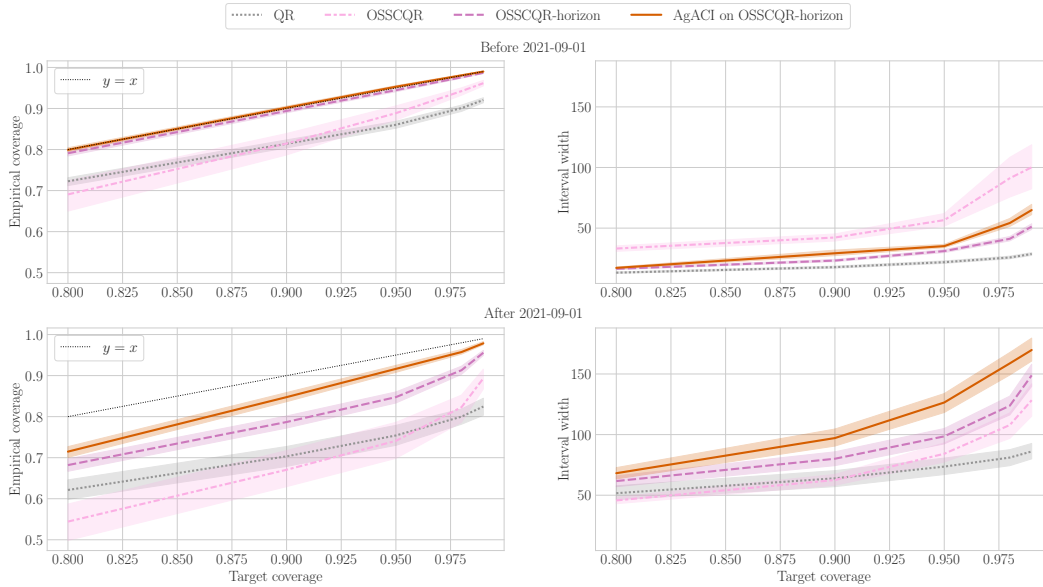


Figure 6: PIs’s performances with different levels of conformalisation on the **quantile linear model**, before September 2021 (top row) and after September 2021 (bottom row), for various target coverage levels (x -axis). The colors and shapes are associated with the conformalisation layers. The shaded regions correspond to the 5% and 95% empirical quantiles after bootstrapping 500 times the test time series.

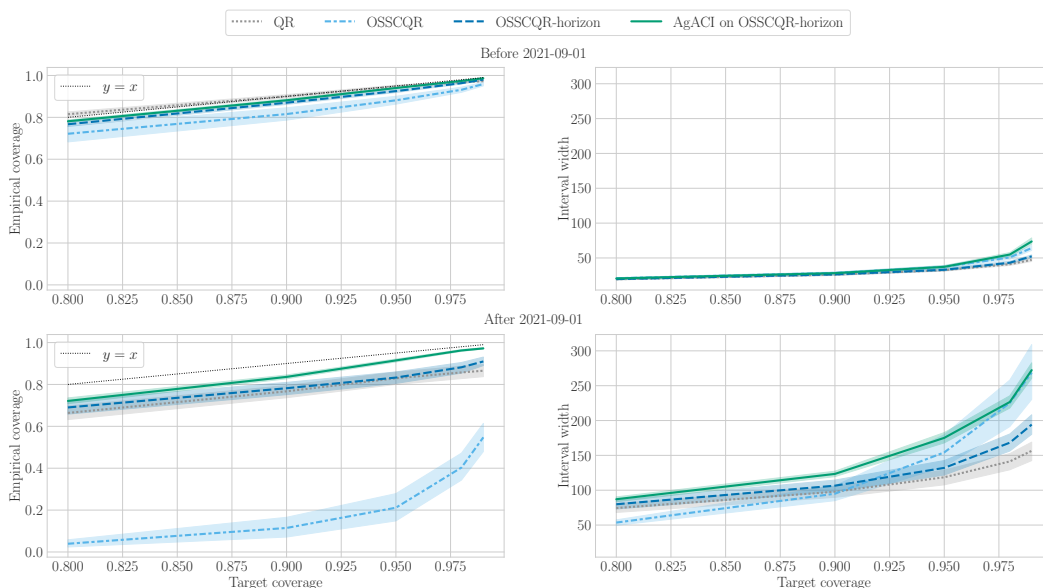


Figure 7: Same caption than Figure 6 but for the **quantile random forest model**.

the performances of various aggregations, each of them considering a different set of experts (individual forecasts, OSSSCP-horizon forecasts, AgACI forecasts, and all of them). As a baseline, we add the uniform average of all of these experts. For each of the aggregation, we compared aggregating forecasts with a unique window size for training with aggregating forecasts with multiple training window size (hence augmenting the number of experts in the set). This latter strategy is usually referred to as *windowing* (Marjasz et al., 2018). We selected the best aggregation (namely aggregating AgACI forecasts with windowing) and, for the sake of readability and for coherence, we displayed in Figure 8 all the aggregations with windowing. It has to be noted that there is a lot of variability, as it can be seen in Figure 8, and that for some aggregation the best choice was in fact without windowing.

Figure 8 highlights that online aggregation improves considerably the robustness to non-stationarity in terms of validity. Furthermore, after September 2021, online aggregation on AgACI forecasts enhances the sharpness of the forecasts with respect to the uniform average, that has similar coverage. This can be explained by the fact that the individual performances degrade in this non-stationary environment, leading to aggregation’s weights close to uniform so as to minimise the risk (as we will also see in the next analysis).

Analysis of aggregation of various AgACI: applying the best conformalisation possible (AgACI) on each model and then aggregating them In Figure 9 we represent the evolution of the weights associated to each of the AgACI (the color representing the base model, and the shade of

it indicating the calibration percentage) with time x -axis, for various coverage level (columns). To improve readability, we display these weights for the aggregation without windowing.

The first striking observation is the presence of temporal ruptures in the weights’ distribution. They are informative as they are associated with domain phenomena, which depend on the considered bound (lower or upper). Particularly, the first one happening is the big negative spike in Easter 2020 (April 13, 2020, see top row of Figure 1) due to both the public holiday and the Covid-19 lockdown. This especially affects the lower bound. The second one occurs in the second fortnight of September 2020 when the first extreme positive peaks take place, impacting the upper bound. These positive spikes are mainly due to a very low wind generation in France (less than 1 GW) and more generally in Europe, along with a French nuclear production well below its level of previous years at the same time. The last significant rupture is around October 2021, when spot prices start to rise drastically and get more and more volatile, corresponding to the increase in level and volatility of gas and carbon emission prices. This one affects both the lower and upper bounds. In particular, the weights’ distribution becomes uniform after this rupture, which is expected in a setting where the aggregation tries to minimize the risk with experts performing poorly.

The second observation is that the methods on which the aggregation places the most of the weights is different depending on the bound: remarkably, at the levels 0.95 and 0.98, the lower bound places high mass on quantile random forests, while the upper bound relies more on qgam. This

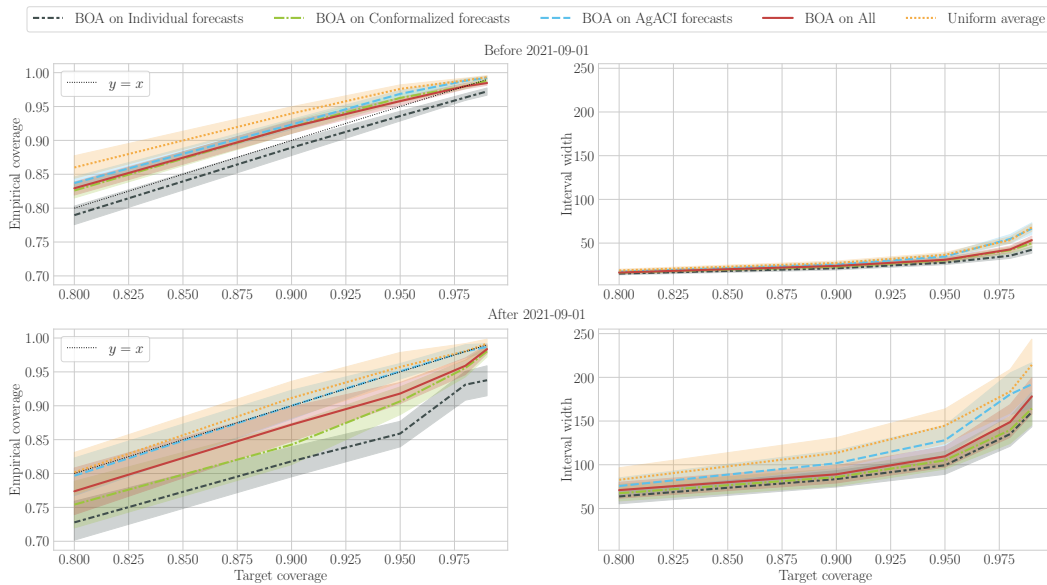


Figure 8: PIs’s performances of online aggregation on multiple set of experts with windowing, before September 2021 (top row) and after September 2021 (bottom row), for various target coverage levels (x -axis). The colors and shapes are associated with the set of experts. The shaded regions correspond to the 5% and 95% empirical quantiles after bootstrapping 500 times the test time series.

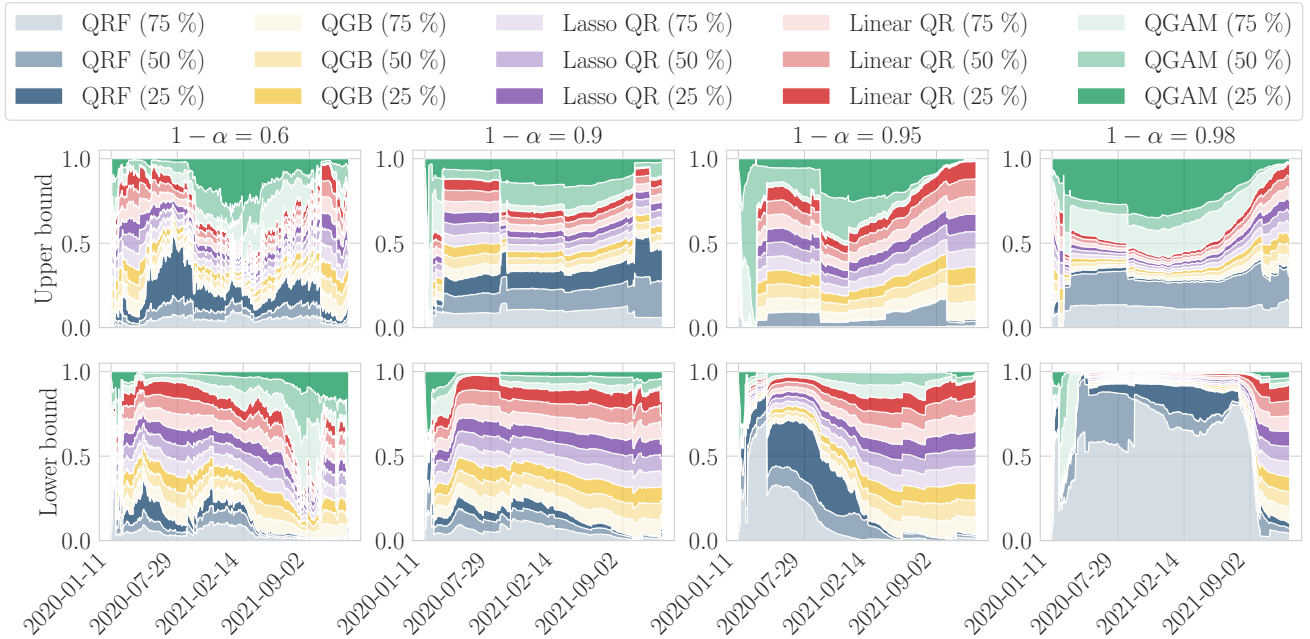


Figure 9: Temporal evolution (x -axis) of the weights associated with each expert in the online aggregation, for different values of (columns). The top row (resp. bottom row) shows the weights assigned for the upper (resp. lower) bound forecast. The colors correspond to the base model on which AgACI is applied to, and the transparency to the proportion of training data kept for actually fitting these base models.

can be explained by the fact that the various methods depend differently on the provided features: additive models such as qgam or linear ones have a great extrapolation ability, while random forests and gradient boosting benefit from more flexibility on features' interaction modeling. This idea is also reflected in Figures 2 and 3 comparing the feature importance in Lasso with the one of Random forest.

Lastly, for high levels of coverage such as 0.95 and 0.98, the aggregation also places weights on different training size depending on the bound. While the upper bound favors small training size, the lower bound encourages large training size. This might be due to the effective sample size which is required to appropriately learn the lower quantiles of the prices, which are less impacted by the non-stationarity; while the upper bound is particularly complex to model, and having more data points correct the predictive model through conformalisation might be a better usage of the available data.

These three key observations argue in favor aggregating independently the upper and lower bounds.

6 Conclusion and perspectives

In this study, we have analysed the performances of a wide range of probabilistic methods in a particularly challenging task: forecasting electricity spot prices in France in 2020 and 2021. On the design, we have highlighted the importance of including the new explanatory variable corresponding to

the nuclear plants' availability. We were also able to bring new insights into the post-processing of individual forecasts, such as conformalisation or aggregation. Indeed, our extensive experiments demonstrate that *i*) conformalisation, when appropriately done as through `OSSCP-horizon`, considerably improves PI's quality despite the non-stationarity, *ii*) online aggregation of experts is extremely powerful in terms of adaptiveness bringing enhanced PI's performances and taking advantage of windowing, *iii*) combining both conformalisation and online aggregation appears on this data set to be the best strategy, and most importantly sheds light on many domain phenomena thanks to great interpretability.

There are many avenues for future works. From the electricity lens, the prices have continued to evolve significantly since 2022 and pursuing the study on newer data would undoubtedly yield new knowledge. Speaking of which, our study did not investigate the crucial question of peaks and extreme forecasts, dominant in electricity prices. Works on online procedure tailored for extremes have already been deployed (Himych et al., 2024), and it might be relevant to see how it can be paired with conformal approaches. Another natural perspective that would deepen our understanding on the benefits of conformalisation is to conformalize the aggregated models as suggested in Susmann et al. (2024), as opposed to aggregating the conformalized models which is what we performed. It would also be interesting to assess the performances of the most recent online conformal algorithms (listed in Section 4.2.4), that might be better suited for non-stationarity. Finally, our angle of approach is

to showcase the advantages of black-box plugs-in such as CP and aggregation. It is attractive to couple it with recent developments that enhance the interpretability of complex statistical models, such as Wood et al. (2022).

References

- Loris Amabile, Delphine Bresch-Pietri, Gilbert El Hajje, Sébastien Labbé, and Nicolas Petit. Optimizing the self-consumption of residential photovoltaic energy and quantification of the impact of production forecast uncertainties. *Advances in Applied Energy*, 2:100020, 2021.
- Yvenn Amara-Ouali, Matteo Fasiolo, Yannig Goude, and Hui Yan. Daily peak electrical load forecasting with a multi-resolution approach. *International Journal of Forecasting*, 39(3):1272–1286, 2023.
- Anastasios Angelopoulos, Emmanuel Candes, and Ryan Tibshirani. Conformal pid control for time series prediction. In *Advances in Neural Information Processing Systems*, 2023.
- Anastasios N. Angelopoulos, Rina Foygel Barber, and Stephen Bates. Online conformal prediction with decaying step sizes, 2024.
- Osbert Bastani, Varun Gupta, Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Practical adversarial multivalid conformal prediction. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- Jonathan Berrisch and Florian Ziel. CRPS learning. *Journal of Econometrics*, dec 2021.
- Jonathan Berrisch and Florian Ziel. Multivariate probabilistic crps learning with an application to day-ahead electricity prices. *International Journal of Forecasting*, 2024a.
- Jonathan Berrisch and Florian Ziel. *The profoc Package: An R package for probabilistic forecast combination using CRPS Learning*, 2024b. R package version 1.3.1.
- Aadyot Bhatnagar, Huan Wang, Caiming Xiong, and Yu Bai. Improved online conformal prediction via strongly adaptive online learning. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pages 2337–2363. PMLR, 23–29 Jul 2023.
- Roger Bjorgan, Chen-Ching Liu, and Jacques Lawarree. Financial risk management in a competitive electricity market. *IEEE Transactions on power systems*, 14(4): 1285–1291, 1999.
- Derek Bunn, Arne Andresen, Dipeng Chen, and Sjur Westgaard. Analysis and forecasting of electricity price risks with quantile factor models. *The Energy Journal*, 37(1), 2016.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- Victor Chernozhukov, Kaspar Wüthrich, and Zhu Yinchu. Exact and Robust Conformal Inference Methods for Predictive Machine Learning with Dependent Data. In *Conference On Learning Theory*, pages 732–749. PMLR, July 2018. ISSN: 2640-3498.
- Cameron Cornell, Nam Trong Dinh, and S. Ali Pourmousavi. A probabilistic forecast methodology for volatile electricity prices in the Australian national electricity market. *International Journal of Forecasting*, 2024.
- Thomas Deschatre, Olivier Féron, and Pierre Gruet. A survey of electricity spot and futures price models for risk management applications. *Energy Economics*, 102: 105504, 2021.
- Matteo Fasiolo, Simon N. Wood, Margaux Zaffran, Raphaël Nedellec, and Yannig Goude. Fast calibrated additive quantile regression. *Journal of the American Statistical Association*, 116(535):1402–1412, 2020.
- Matteo Fasiolo, Simon N. Wood, Margaux Zaffran, Raphaël Nedellec, and Yannig Goude. qgam: Bayesian nonparametric quantile regression modeling in R. *Journal of Statistical Software*, 100(9):1–31, 2021.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- Pierre Gaillard and Yannig Goude. *OPERA, a R package for online aggregation of experts*, 2016. R package version 1.2.0.
- Isaac Gibbs and Emmanuel Candes. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.
- Isaac Gibbs and Emmanuel Candès. Adaptive conformal inference under distribution shift. In *Advances in Neural Information Processing Systems*, 2021.
- Isaac Gibbs and Emmanuel Candès. Conformal inference for online prediction with arbitrary distribution shifts. preprint, 2023.
- Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems*, 35:507–520, 2022.
- Trevor Hastie and Robert Tibshirani. Generalized additive models. *Statistical Science*, 1(3):297–310, 1986.
- Omar Himych, Amaury Durand, and Yannig Goude. Adaptive Forecasting of Extreme Electricity Load. preprint, 2024.
- Tao Hong, Pierre Pinson, Shu Fan, Hamidreza Zareipour, Alberto Troccoli, and Rob J Hyndman. Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond, 2016.

- Paris IEA. Covid-19 impact on electricity. Technical report, Technical report, 2021.
- Paris IEA. Renewable electricity. Technical report, Technical report, 2022a.
- Paris IEA. World energy outlook 2022. Technical report, Technical report, 2022b.
- Arkadiusz Jędrzejewski, Jesus Lago, Grzegorz Marcjasz, and Rafał Weron. Electricity price forecasting: The dawn of machine learning. *IEEE Power and Energy Magazine*, 20(3):24–31, 2022.
- Roger Koenker. *Quantile Regression*. Econometric Society Monographs. Cambridge University Press, 2005. ISBN 978-0-521-84573-1.
- Hans R Kunsch. The jackknife and the bootstrap for general stationary observations. *The annals of Statistics*, pages 1217–1241, 1989.
- Jesus Lago, Grzegorz Marcjasz, Bart De Schutter, and Rafał Weron. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293:116983, 2021.
- Jing Lei, Max G’Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-Free Predictive Inference for Regression. *Journal of the American Statistical Association*, 113(523):1094–1111, July 2018.
- Stylianios Loizidis, Andreas Kyprianou, and George E. Georghiou. Electricity market price forecasting using elm and bootstrap analysis: A case study of the german and finnish day-ahead markets. *Applied Energy*, 363:123058, 2024.
- Katarzyna Maciejowska, Bartosz Uniejewski, and Rafał Weron. Forecasting electricity prices. preprint, 2022.
- S. Makridakis, E. Spiliotis, and V. Assimakopoulos. M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, 38(4):1346–1364, 2022.
- Grzegorz Marcjasz, Tomasz Serafin, and Rafał Weron. Selection of calibration windows for day-ahead electricity price forecasting. *Energies*, 11(9), 2018.
- Grzegorz Marcjasz, Michał Narajewski, Rafał Weron, and Florian Ziel. Distributional neural networks for electricity price forecasting. *Energy Economics*, 125:106843, 2023.
- Bianca Marin Moreno, Margaux Brégère, Pierre Gaillard, and Nadia Oudjane. A mirror descent approach for mean field control applied to demande-side management. preprint, 2023.
- Nicolai Meinshausen. Quantile regression forests. *Journal of Machine Learning Research*, 7(35):983–999, 2006.
- N. Nassar, D. Silva, and H. Morais. Hierarchical energy management solution for smart charging. In *CIREP Porto Workshop 2022: E-mobility and power distribution systems*, volume 2022, pages 721–725, 2022.
- Daniel Nickelsen and Gernot Müller. Bayesian hierarchical probabilistic forecasting of intraday electricity prices. preprint, 2024.
- Weronika Nitka, Tomasz Serafin, and Dimitrios Sotiros. Forecasting electricity prices: Autoregressive hybrid nearest neighbors (arhnn) method. In Maciej Paszynski, Dieter Kranzlmüller, Valeria V. Krzhizhanovskaya, Jack J. Dongarra, and Peter M.A. Sloom, editors, *Computational Science – ICCS 2021*, pages 312–325, 2021.
- Jakub Nowotarski and Rafał Weron. Computing electricity spot price prediction intervals using quantile regression and forecast averaging. *Computational Statistics*, 30(3):791–803, 2014.
- Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive Confidence Machines for Regression. In Tapio Elomaa, Heikki Mannila, and Hannu Toivonen, editors, *Machine Learning: ECML 2002*, pages 345–356, 2002.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Dimitris N Politis and Joseph P Romano. The stationary bootstrap. *Journal of the American Statistical association*, 89(428):1303–1313, 1994.
- Carl Remlinger, Clémence Alasseur, Marie Brière, and Joseph Mikael. Expert aggregation for financial forecasting. *The Journal of Finance and Data Science*, 9:100108, 2023.
- Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized Quantile Regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- RTE. Bilan électrique 2022, 2022.
- Matteo Sesia and Emmanuel J. Candès. A comparison of some conformal quantile regression methods. *Stat*, 9(1), 2020.
- Herbert Susmann, Antoine Chambaz, Julie Josse, Mathias Wargon, Philippe Aegerter, and Emmanuel Bacry. Probabilistic Prediction of Arrivals and Hospitalizations in Emergency Departments in Île-de-France. preprint, 2024.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- Léonard Tschora, Erwan Pierre, Marc Plantevit, and Céline Robardet. Electricity price forecasting on the day-ahead market using machine learning. *Applied Energy*, 313:118752, 2022.
- Bartosz Uniejewski and Rafał Weron. Regularized quantile regression averaging for probabilistic electricity price forecasting. *Energy Economics*, 95:105121, 2021.

- Bartosz Uniejewski, Jakub Nowotarski, and Rafał Weron. Automated variable selection and shrinkage for day-ahead electricity price forecasting. *Energies*, 9(8), 2016.
- Bartosz Uniejewski, Rafał Weron, and Florian Ziel. Variance stabilizing transformations for electricity spot price forecasting. *IEEE Transactions on Power Systems*, 33(2): 2219–2229, 2018.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer US, 2005.
- Volodya Vovk, Alexander Gammerman, and Craig Saunders. Machine-Learning Applications of Algorithmic Randomness. In *Proceedings of the Sixteenth International Conference on Machine Learning, ICML '99*, pages 444–453, San Francisco, CA, USA, 1999.
- Rafał Weron. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International journal of forecasting*, 30(4):1030–1081, 2014.
- Olivier Wintenberger. Optimal learning with bernstein online aggregation. *Machine Learning*, 106:119–141, 2017.
- Wojciech Wisniewski, David Lindsay, and Sian Lindsay. Application of conformal prediction interval estimations to market makers’ net positions. In Alexander Gammerman, Vladimir Vovk, Zhiyuan Luo, Evgeni Smirnov, and Giovanni Cherubin, editors, *Proceedings of the Ninth Symposium on Conformal and Probabilistic Prediction and Applications*, volume 128 of *Proceedings of Machine Learning Research*, pages 285–301. PMLR, 2020.
- S. N. Wood, Y. Goude, and M. Fasiolo. Interpretability in generalized additive models. In *Interpretability for Industry 4.0: Statistical and Machine Learning Approaches*, page 85–123. Springer International Publishing, 2022.
- Yifan Yang, Ju’e Guo, Yi Li, and Jiandong Zhou. Forecasting day-ahead electricity prices with spatial dependence. *International Journal of Forecasting*, 2023.
- Margaux Zaffran, Olivier Feron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 25834–25866. PMLR, 2022.
- Florian Ziel and Rafał Weron. Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics*, 70:396–420, 2018.
- Çağatay Berke Bozlak and Claudia Fernanda Yaşar. An optimized deep learning approach for forecasting day-ahead electricity prices. *Electric Power Systems Research*, 229:110129, 2024.

A Results on the CRPS

To assess the performance of a probabilistic method on the overall range of quantiles, one can use the Continuous Ranked Probability Score (CRPS). This score is originally described in terms of the predictive CDS $\hat{F}_{d,h}$:

$$CRPS(\hat{F}_{d,h}, y_{d,h}) = \int_{-\infty}^{\infty} (\hat{F}_{d,h}(y|x_{d,h}) - \mathbb{1}_{\{y_{d,h} \leq y\}})^2 dy.$$

Interestingly, the CRPS can be reformulated (to a multiplicative constant) as :

$$CRPS(\hat{F}_{d,h}, y_{d,h}) = \int_0^1 \rho_{\alpha}(y_{d,h}, \hat{F}_{d,h}^{-1}(\alpha)) d\alpha,$$

where $\hat{F}_{d,h}^{-1}(\alpha)$ actually corresponds to the predicted value at quantile α . By approximating this integral as a Riemann sum, we can transform pinball scores over multiple quantiles into one single metric.




Figure 10: PIs’s CRPS with different levels of conformalisation on the **quantile linear model**, depending on the time. The colors and shapes are associated with the conformalisation layers.




Figure 11: Same caption than Figure 10 but for the **quantile random forest model**.

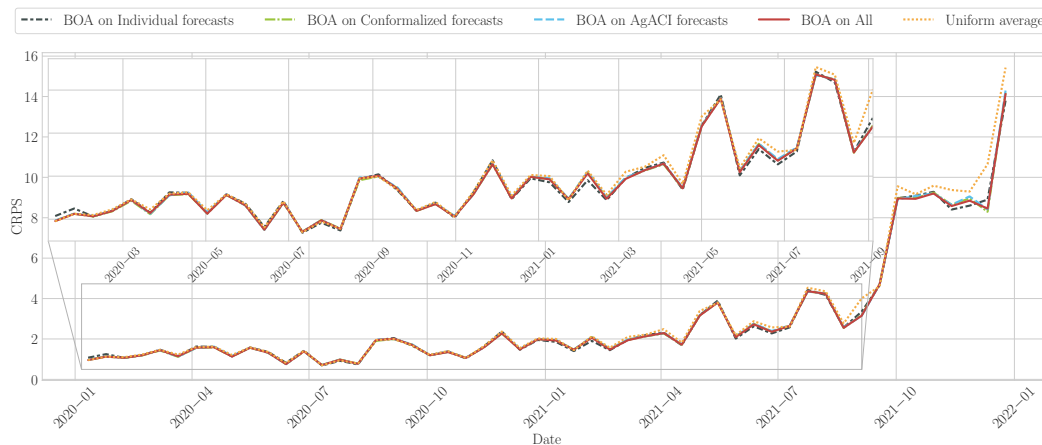


Figure 12: PIs’s CRPS of online aggregation on multiple set of experts with windowing, depending on the time. The colors and shapes are associated with the set of experts.