
Learning Quantile Functions without Quantile Crossing for Distribution-free Time Series Forecasting

Youngsuk Park[†]

Kelvin Kan * ‡

[†]Amazon Research

Danielle C. Maddix[†]

Jan Gasthaus[†]

François-Xavier Aubet[†]

Yuyang Wang[†]

[‡]Emory University

Abstract

Quantile regression is an effective technique to quantify uncertainty, fit challenging underlying distributions, and often provide full probabilistic predictions through joint learnings over multiple quantile levels. A common drawback of these joint quantile regressions, however, is *quantile crossing*, which violates the desirable monotone property of the conditional quantile function. In this work, we propose the Incremental (Spline) Quantile Functions I(S)QF, a flexible and efficient distribution-free quantile estimation framework that resolves quantile crossing with a simple neural network layer. Moreover, I(S)QF inter/extrapolate to predict arbitrary quantile levels that differ from the underlying training ones. Equipped with the analytical evaluation of the continuous ranked probability score of I(S)QF representations, we apply our methods to NN-based time series forecasting cases, where the savings of the expensive re-training costs for non-trained quantile levels is particularly significant. We also provide a generalization error analysis of our proposed approaches under the sequence-to-sequence setting. Lastly, extensive experiments demonstrate the improvement of consistency and accuracy errors over other baselines.

1 INTRODUCTION

Probabilistic time series forecasting methods are increasingly replacing point prediction techniques in practical applications, as it is crucial for downstream decision-making processes to take the uncertainties in

predictions into consideration (Park et al., 2019, 2020; Kim et al., 2020; Petropoulos et al., 2022). A popular approach to probabilistic forecasting is to combine a sequential model with a likelihood model that determines how to emit from the hidden or latent states to the observations. Examples include classical statistical models such as State-Space Models (Hyndman et al., 2008), Gaussian processes (Rasmussen and Williams, 2006), and more recently the neural network models like DeepAR / DeepVAR (Salinas et al., 2020, 2019), DeepState (de Bézenac et al., 2020; Rangapuram et al., 2018), DeepFactor (Wang et al., 2019), etc. While full probabilistic predictions can be obtained from these models, a practical conundrum of which likelihood function to choose arises. As an example, DeepAR, a RNN-based probabilistic forecaster, offers likelihood choices of normal, student- t , negative binomial, etc. Therefore, it is desirable in both theory and practice to have a method that requires no assumption of the data generating process.

Fortunately, quantile regression (Koenker and Bassett Jr, 1978; Koenker, 2005), which has been successfully used for robustly modeling probabilistic outputs, comes to rescue. The incorporation of the quantile regression component to various sequential neural network backbones has been shown to be particularly effective with recent advances in deep learning (Wen et al., 2017; Gasthaus et al., 2019; Lim et al., 2019; Eisenach et al., 2020). Obtaining a full probabilistic prediction (i.e., the ability to query a forecast at an arbitrary quantile) usually requires generating multiple quantiles at once. However, when modeling multiple quantiles simultaneously, *quantile crossing*, i.e., the failure of the estimated conditional quantile function to obey the required monotonicity constraint, is a commonly observed concern beyond mere theoretical consistency. This phenomenon particularly stands out in the case of sequence-to-sequence (Seq2Seq) predictions, where the inconsistency can propagate over a sequence of predictions. Even though various strategies have been

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

*Work done as an intern at Amazon Research.

proposed to remedy quantile crossing (See Section 1.1), these techniques are not widely adopted in combination with recent deep probabilistic forecasting models, due to their perceived complexity and lack of solid principles. In particular, Wen et al. (2017); Lim et al. (2019) do not explicitly address quantile crossing.

In addition to quantile crossing, another issue of typical quantile regression approaches in the time series context is that they require the set of predicted quantile levels to be fixed for training apriori. This restriction requires the expense of re-training the models when making a prediction at arbitrary quantile levels. This is a particular concern with the computational overhead for the class of heavy deep forecasting models with numerous training panels of time series.

In this paper, we propose a simple methodology of Incremental (Spline) Quantile Functions I(S)QF (in Sections 3, 4), both of which quantify prediction uncertainty in a distribution-free manner. I(S)QF consist of a family of conditional quantile functions that interpolates between the given quantile levels and extrapolates in the parametric tail distribution beyond the extremal training quantiles. These inter/extrapolation strategies ultimately lead to resolving both quantile crossing and re-training cost issues, along with supporting analytical Continuous Ranked Probability Score (CRPS) (Gneiting and Raftery, 2007) evaluation in the training stage. We apply our methods to Seq2Seq time series forecasting, where both aforementioned issues are raised most frequently (in Section 5). We analyze the generalization errors under multi-horizon and multi-quantile time series forecasting, characterized with quantities, such as the Rademacher complexity and discrepancy measure for stationarity (in Section 6). Under the state-of-the-art Seq2Seq MQ-CNN model (Wen et al., 2017), extensive experiments on real-world datasets demonstrate the consistency and accuracy improvement of our methodology with I(S)QF layers over various other baseline layers, e.g., default quantile (Wen et al., 2017), Gaussian (Flunkert et al., 2017), and SQF (Gasthaus et al., 2019) (in Section 7).

1.1 Related Works

Seq2Seq Quantile Forecasters. Recently, several Seq2Seq forecasting models have been proposed to jointly learn the quantile estimates over multiple quantiles without any explicit distribution assumptions. Wen et al. (2017) propose MQ-R(C)NN, which uses a RNN or dilated causal convolution (CNN) encoder, respectively, and a multilayer perceptron (MLP) decoder that outputs a set of quantile levels for the entire forecast horizon. A disadvantage is that the model needs to be re-trained if other quantile levels than the training ones are requested at inference time. Wen

and Torkkola (2019) extend the MQ-CNN model into a marginal quantile model with a generative quantile Gaussian copula. This Gaussian copula component improves the forecast at the distribution tails, but the quantile crossing drawback that can occur with MQ-CNN still remains. Chen et al. (2020) propose DeepTCN, another Seq2Seq model where the encoder is also a dilated causal convolution with residual blocks, and the decoder is simply a MLP with residual connections. Structure-wise, DeepTCN is almost the same as the basic structure of MQ-CNN, i.e., without the local MLP component that aims to model spikes and events. When the quantiles are learned jointly, all of these Seq2Seq models can still suffer from quantile crossing.

Quantile Regression. There are various approaches to resolve the issue of quantile crossing that occurs when quantile estimations over multiple quantiles are learned jointly. Most of these works are sorting-based post-processing at the end (Chernozhukov et al., 2010; Kim et al., 2021) or or expensive constrained optimization (Liu and Wu, 2009). Schmidt and Zhu (2016) propose learning on the non-negative increment between quantile estimates on pre-determined quantile levels, and then stacking them. In addition, some linear/non-linear interpolation between quantile estimates allows for a form of quantile function in some range of quantile levels. SQF (Gasthaus et al., 2019) is another method that gives a functional form based on the CRPS and its analytical evaluation, and results in no quantile crossing. However, this method is too flexible in choosing knot positions to optimize efficiently. Therefore, interpolation through linear splines leads to general performance degradation, especially on the tail regions.

Theoretical Analysis on Forecasting. There are several recent theoretical analyses of time series forecasting with the tools from learning theory. Kuznetsov and Mohri (2015) investigate the theoretical analysis on the general scenario of non-stationary and non-mixing stochastic processes in terms of a data-dependent measure of sequential complexity and a discrepancy measure. Zimin and Lampert (2017) study the learnability of stochastic processes with respect to the conditional risk, focusing on analyzing scenarios, where the pairwise discrepancy is controllable. Mariet and Kuznetsov (2019) examine theoretical studies on general multivariate sequence-to-sequence settings in terms of discrepancy measure and mixing coefficient of the underlying stochastic process. Albeit the powerful intuitions and mathematical machinery provided, these works focus on general one-step prediction models, lacking more fruitful intuition for our use case of global multi-horizon and multi-quantile deep time series forecasting.

2 PRELIMINARIES

For a random variable $Z \in \mathbb{R}$, we denote $F_Z(z)$ as its cumulative distribution function (CDF). Then, the α -quantile of Z is given as:

$$q_Z(\alpha) := F_Z^{-1}(\alpha) = \inf\{z \in \mathbb{R} : \alpha \leq F_Z(z)\},$$

where $\alpha \in (0, 1)$ denotes a quantile level. The quantile function $q(\cdot)$ is also called the percent-point function or inverse cumulative distribution function.

2.1 Quantile Regression

Let $(X, Z) \sim F_{(X, Z)}$ for a regression setting. Quantile regression seeks to estimate the α -quantile of $Z \in \mathbb{R}$ conditioned on $X = x$ for some $\alpha \in [0, 1]$, i.e., $q(\alpha | x) = F_{Z|X=x}^{-1}(\alpha)$. Equivalently, the α -quantile is the solution of minimizing the expected quantile loss (Koenker and Bassett Jr, 1978):

$$q(\alpha | x) = \operatorname{argmin}_{q \in \mathbb{R}} \mathbb{E}_{Z|X=x}(\rho_\alpha(Z - q)), \quad (1)$$

where $\rho_\alpha(u) = u \times (\alpha - \mathbf{1}\{u < 0\})$ with $u = z - q$ denotes the quantile loss of $q \in \mathbb{R}$ w.r.t. any $z \in \mathbb{R}$.

Continuous Ranked Probability Score (CRPS). The CRPS (Gneiting and Raftery, 2007) is a *proper scoring rule*¹ that averages the quantile losses over all quantiles, rather than optimizing a single quantile loss. Formally, given a fixed target $z \in \mathbb{R}$ and a family of quantile functions \mathcal{Q} , the CRPS $L : \mathcal{Q}, \mathbb{R} \rightarrow \mathbb{R}$ is given as:

$$L(q, z) = \int_{\alpha=0}^1 2\rho_\alpha(z - q(\alpha))d\alpha. \quad (2)$$

2.2 Quantile Regression in Time Series Forecasting

Quantile regression settings can be extended to the common case of multiple horizon and multiple quantile estimations in time series forecasting as follows. Suppose we have m related time series data, each of which consists of observation $z_{i,t} \in \mathbb{R}$ with (optional) input covariates $\xi_{i,t} \in \mathbb{R}^d$ at time t . In i -th time series forecasting, given T past target observations $z_{i,1:T}$ and all (optional) *future* covariates $\xi_{i,1:T+\tau}$, we wish to make τ future quantile predictions at time T :

$$\{\hat{z}_{i,T+1}^\alpha, \dots, \hat{z}_{i,T+\tau}^\alpha\}_{\alpha \in \{\alpha\}} = F_\theta(z_{i,1:T}, \xi_{i,1:T+\tau}),$$

where $\{\alpha\}$ denotes a set of quantile levels with $\alpha \in [0, 1]$, and F_θ denotes a global² sequence-to-sequence quantile

¹ $\mathbb{E}_{Z \sim F} L(F^{-1}, Z) \leq \mathbb{E}_{Z \sim F} L(G^{-1}, Z)$.

²Forecast models are the same across all time series i .

prediction. In short, we can express the α -quantile estimate as:

$$\hat{z}_{i,t}^\alpha = q_\theta^t(\alpha | x_i),$$

where $x_i = (z_{i,1:T}, \xi_{i,1:T+\tau})$ denotes the input to the quantile function.

Finally, Empirical Risk Minimization (ERM) with L in (2) is formulated to find the best-fit parameters:

$$\operatorname{minimize}_\theta \frac{1}{m\tau} \sum_{i=1}^m \sum_{t=T+1}^{T+\tau} L(q_\theta^t(\cdot | x_i), z_{i,t}). \quad (3)$$

2.3 Quantile Crossing

When tackling the problem of quantile learning naively, inconsistencies that violate the non-decreasing property of quantile estimate on α , called *quantile crossing* may arise. In other words, for $\alpha_1 > \alpha_2$, there exists x s.t.

$$q(\alpha_1 | x) < q(\alpha_2 | x),$$

where $q(\cdot | x)$ is the quantile estimate of the target conditioned on input x . The quantile crossing issue occurs in many cases when quantile regression is applied separately over multiple quantiles $\{\alpha\}$, and/or ERM with CRPS in (3) are solved under naive parameterizations on quantile functions.

3 INCREMENTAL QUANTILE FUNCTIONS (IQF)

Most forecast models output quantile estimates on a fixed finite set $\{\alpha_k\}_{k=1}^K$, called *quantile knots*. In this section, we propose the Incremental Quantile Functions (IQF), a family of conditional quantile functions. IQF estimates targets on any quantile $\alpha \in (0, 1)$, and overcomes quantile crossing. We first learn some basis quantile estimates on quantile knots, and then apply interpolation and extrapolation strategies around the knots to make it accessible on any query quantile.

For notational simplicity, we omit index i, t . Let h be the last hidden variable with $h = h(x)$ where h denotes a function that maps an input x to a hidden variable h . We denote the last output layer as q_ϕ with parameter ϕ into which the last hidden variable h is fed into, i.e., $q_\phi(\alpha | h) := q_\theta(\alpha | x)$. We often use $q(\alpha | h)$ by omitting parameters ϕ, θ when obvious.

3.1 Basis Estimates on Quantile Knots

On the quantile knots $\{\alpha_k\}_{k=1}^K$ with $\alpha_k < \alpha_{k+1}$, IQF first gives basis quantile estimates \hat{q} by cumulatively adding non-negative increments as the quantile level




Figure 1: Neural network design of I(S)QF representation. I(S)QF architecture never allows quantile crossing to happen to whichever (hidden) input h is given.

α_k increases:

$$\hat{q}_\phi(\alpha_k | h) = \begin{cases} \pi^{\text{Linear}}(h; \phi_1), & k = 1, \\ \hat{q}_\phi(\alpha_{k-1} | h) + \pi^{\text{MLP}}(h; \phi_k), & k \neq 1, \end{cases} \quad (4)$$

where π^{Linear} denotes a linear function, and π^{MLP} denotes a multilayer perceptron (MLP) with a non-negative non-linear component, e.g., ReLu, Sigmoid. Here, $\phi = \{\phi_i\}_{i=1}^K$ are the associated parameters to learn.

3.2 Inter/Extrapolation beyond Knots

In order to provide full quantile estimates beyond basis estimates in (4), i.e., $q(\alpha | h)$ on $\alpha \notin \{\alpha_k\}$, IQF takes the inter/extrapolation strategies into account: non-parametric interpolation in the middle and parametric extrapolation in the tails. It adopts a simple yet efficient (deterministic) linear interpolation and extrapolation with exponential tail distributions.

Assume the knot size $K \geq 2$, and use extremal knots to define tail regions³, i.e., $\alpha_{\text{tail}_L} = \alpha_1$ and $\alpha_{\text{tail}_R} = \alpha_K$.

Linear Interpolation. For the non-tail region $\alpha \in [\alpha_{\text{tail}_L}, \alpha_{\text{tail}_R}]$, each linear interpolation on $\alpha_k \leq \alpha \leq \alpha_{k+1}$ interval gives:

$$q(\alpha | h) = w(\alpha)\hat{q}(\alpha_k | h) + (1 - w(\alpha))\hat{q}(\alpha_{k+1} | h), \quad (5)$$

where $w(\alpha) = (\alpha_{k+1} - \alpha)/(\alpha_{k+1} - \alpha_k)$. Note that interpolation q and basis \hat{q} match at every knot α_k .

Extrapolation with Exponential Tails. For the tail regions, we extrapolate the quantile estimate through exponential distribution at both ends:

$$\alpha = \begin{cases} \exp(\beta_L(q(\alpha | h) - \gamma_L)), & \alpha \leq \alpha_{\text{tail}_L}, \\ 1 - \exp(-\beta_R(q(\alpha | h) - \gamma_R)), & \alpha \geq \alpha_{\text{tail}_R}. \end{cases}$$

³The tail region can be selected more arbitrarily.

Here, parameters (β_L, γ_L) and (β_R, γ_R) are uniquely chosen so that extrapolation q and basis \hat{q} coincide on the two leftmost and rightmost quantile knots α_1, α_2 and α_{K-1}, α_K , respectively (See the supplementary materials).


Then, this is equivalent to extrapolating:

$$q(\alpha | h) = \begin{cases} \frac{1}{\beta_L} \log \frac{\alpha}{\alpha_2} + \hat{q}(\alpha_2 | h), & \alpha \leq \alpha_{\text{tail}_L}, \\ \frac{1}{\beta_R} \log \frac{1 - \alpha_{K-1}}{1 - \alpha} + \hat{q}(\alpha_{K-1} | h), & \alpha \geq \alpha_{\text{tail}_R}. \end{cases} \quad (6)$$


We add numeric safeguards to the denominators and log operations that enhance numerical stability while maintaining monotonicity. Finally, aggregating inter/extrapolation completes IQF:

$$q(\alpha | h) = \begin{cases} (5), & \alpha \in [\alpha_{\text{tail}_L}, \alpha_{\text{tail}_R}], \\ (6), & \alpha \in (0, \alpha_{\text{tail}_L}] \cup [\alpha_{\text{tail}_R}, 1). \end{cases}$$

Figure 1 displays the corresponding neural network architecture.



(a) IQF with 5 knots.



(b) IQF with 20 knots.

Figure 2: IQF recovers the underlying multimodal distribution via inter/extrapolation beyond training knots with higher accuracy as the knot size K increases.

Distribution Recovery. Figure 7 shows the behaviors of IQF for a multi-modal distribution over different number of training knots without any distributional assumption. IQF also fits heavy tail distributions, e.g., Exponential and even Cauchy (although CRPS under true Cauchy distributions may not be well-defined). See the Section H for the additional experiments. Note that naive distributional assumptions, e.g. Gaussianity fail in these examples. Even though IQF is in a distribution-free class, any additional information regarding the tail

distribution for IQF can be enforced based on prior knowledge on the data, e.g., non-negative targets in count data, and the choice of tail regions.

4 INCREMENTAL SPLINE QUANTILE FUNCTIONS (ISQF)

We extend IQF into a more flexible family, Incremental Spline Quantile Functions (ISQF). In particular, ISQF adopts learnable inter/extrapolation strategies.

4.1 Learnable Inter/Extrapolation

While IQF takes fixed inter/extrapolation strategies into account, ISQF uses learnable versions of these strategies: linear spline interpolation in the middle and parametric extrapolation with exponential or Generalized Pareto Distribution (GPD) at the tails.

Piecewise-linear Interpolation. Given two fixed knot positions α_k and α_{k+1} , we additionally adopt learnable linear splines, similar to SQF, to interpolate between them:

$$q(\alpha | h) = \sum_{s=0}^{S-1} \max\left\{\min\left\{\frac{\alpha - d_s}{d_{s+1} - d_s}, 1\right\}, 0\right\}(p_{s+1} - p_s) + \hat{q}(\alpha_k | h). \quad (7)$$

Here, the S number of spline knots $\{d_s\}$ and associated quantile estimates $\{p_s\}$ are parameterized to be non-decreasing, and match on the basis points $\{\alpha_k\}_{k=1}^K$, i.e., $d_0 = \alpha_k$, $d_S = \alpha_{k+1}$ and $p_0 = \hat{q}(\alpha_k | h)$, $p_S = \hat{q}(\alpha_{k+1} | h)$. (7) is equivalent to SQF by a change of variables, and simpler to parameterize and learn. See more discussion in the supplementary materials.

Extrapolation with Parametric Tails. We generalize the exponential in IQF by requiring it to only pass through the tail quantile knot, and letting the parameters β_L, β_R in (6) be trainable. Another class of extrapolation other than exponential is Generalized Pareto Distribution (GPD). Similar to Ehrlich et al. (2021), we model the tail distribution with GPD as:

$$\alpha = \begin{cases} \alpha_{\text{tail}_L} (1 - \eta \psi_\mu(\alpha, \alpha_{\text{tail}_L}))^{-1/\eta}, & \alpha \leq \alpha_{\text{tail}_L}, \\ \alpha_{\text{tail}_R} (1 - (1 + \eta \psi_\mu(\alpha, \alpha_{\text{tail}_R}))^{-1/\eta}), & \alpha \geq \alpha_{\text{tail}_R}, \end{cases} \quad (8)$$

where $\psi_\mu(\alpha, \alpha') := [q(\alpha | h) - q(\alpha' | h)]/\mu$, and the shape and scale (η, μ) are learnable parameters. Note that the quantiles from extrapolation match on the extremal knots. Lastly, re-arranging the CDF (8) into quantile form gives the desirable extrapolation.




Figure 3: An example of applying I(S)QF as the output layers in a Seq2Seq framework.

4.2 Analytical CRPS Evaluation

To this end, an analytic representation of the CRPS L is preferable to an integral approximation because it saves computational cost, and is more accurate. We derive the analytical expression of CRPS integral for I(S)QF with both exponential and generalized Pareto tail distribution (Friederichs and Thorarinsdottir, 2012). See Section C for our derivations and final expressions.

5 FORECASTING WITH I(S)QF

Architecture Design for Seq2Seq Forecasting.

We now walk through how to use I(S)QF for the conditional quantile function in the Seq2Seq time series forecasting setting. Under the general encoder and decoder architecture of NN-based Seq2Seq models, we replace the last layers of the decoder with I(S)QF, as depicted in Figure 3. Specifically, an input x passes through an encoder and decoder to generate the last hidden variable $h = \mathbf{h}(x)$, which is fed into the I(S)QF layers, resulting in the desirable conditional quantile estimates. Finally, we learn the forecast parameters by minimizing the ERM with the CRPS in (3) with some training data. I(S)QF is also applicable to autoregressive models (See the supplementary materials for details.)

Sample Path Generation. Based on the full conditional quantile function from I(S)QF, we draw $\alpha \sim U[0, 1]$, and then generate a sample path:

$$[\hat{z}_{T+1}, \dots, \hat{z}_{T+\tau}] = [q^{T+1}(\alpha | x), \dots, q^{T+\tau}(\alpha | x)],$$

for any input x .

Comparison of Various Output Layers. IQF, ISQF and SQF support a monotonicity-preserving conditional quantile function on arbitrary quantiles and sample path generation. On the other hand, the naive

	Eliminates quantile crossing	Eliminates re-training for arbitrary quantile query	Supports sample path generation	Inter- polation	Extra- polation	Knots	Number of parameters
QF	No	No	No	No	No	Fixed	K
IQF	Yes	Yes	Yes	Yes	Yes (Exp)	Fixed	K
ISQF	Yes	Yes	Yes	Yes	Yes (Exp, GPD)	Mixed	$K+2S$
SQF	Yes	Yes	Yes	Yes	No	Non-fixed	$2(K+S)$

Table 1: Comparison of various quantile output layers. K and S denote the total number of quantile and spline knots for I(S)QF, respectively. In case of SQF, all the knots used for ISQF are treated as its spline knots to fit.

MLP-based quantile function (QF) used in MQ-CNN (Wen et al., 2017) is not guaranteed to support these properties. In particular, its inability to support sample paths hinders QF from its applicability to autoregressive models, e.g., DeepAR (Flunkert et al., 2017). The additional flexibility in choosing the knots with SQF results in more parameters to optimize over, which can be more difficult in the multi-horizon forecasting. Table 1 summarizes these differences.

6 GENERALIZATION ERRORS

To analyze generalization errors, we first transform the available data $\mathcal{D} = \{z_{i,1:T}, \xi_{i,1:T+\tau}\}_{i=1}^m$ into some training data $\mathcal{D}_{\text{train}}$ and test data $\mathcal{D}_{\text{test}}$ suitable for defining empirical and generalized losses under Seq2Seq framework. First, in the training stage, we assume to forecast at $T - \tau$ with the next τ targets $\{z_{i,t}\}_{t=T-\tau+1}^T$ accessible. Formally, the *conditional empirical loss* $\hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}})$ under the CRPS L is given as:

$$\hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}}) = \frac{1}{m\tau} \sum_{i=1}^m \sum_{t=T-\tau+1}^T L(z_{i,t}, q_{\theta}^t(\cdot | x_i^{-\tau})).$$

Here, $\mathcal{D}_{\text{train}} = \{z_{i,T-\tau+1:T}, x_i^{-\tau}\}_{i=1}^m$ denotes the training data with τ -shifted input $x_i^{-\tau} = (z_{i,1:T-\tau}, \xi_{i,1:T})$.

Next, the *conditional generalized loss* $\mathcal{L}(\theta | \mathcal{D}_{\text{test}})$ takes the expected loss over the next unknown τ targets $\{z_{i,t}\}_{t=T+1}^{T+\tau}$ from the actual forecast time T as follows:

$$\mathcal{L}(\theta | \mathcal{D}_{\text{test}}) = \frac{1}{m\tau} \sum_{i=1}^m \sum_{t=T+1}^{T+\tau} \mathbb{E}[L(Z_{i,t}, q_{\theta}^t(\cdot | X_i)) | X_i = x_i],$$

where $Z_{i,t}$ denotes a random variable for the future predictions. Here, $\mathcal{D}_{\text{test}} = \{x_i\}_{i=1}^m$ denotes the test data with input $x_i = (z_{i,1:T}, \xi_{i,1:T+\tau})$ to the quantile function. Then, optimal solution is denoted as:

$$\hat{\theta} = \operatorname{argmin} \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}}), \theta^* = \operatorname{argmin} \mathcal{L}(\theta | \mathcal{D}_{\text{test}}),$$

for conditional empirical loss and generalized loss, respectively.

To begin our analysis, we make the following assumption.

Assumption 1. *The target time series $Z_{i,t}$ are bounded, i.e., $\|Z_{i,t}\| \leq D$ for some $D \geq 0$.*

6.1 Conditional Generalization Error

Let \mathcal{F} be the set of functions f where $f = L \circ q$ is the composition of CRPS loss L in (3) and any quantile function q .

Definition 1. *The Rademacher complexity of \mathcal{F} is defined as:*

$$\mathfrak{R}_N(\mathcal{F}) = \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left(\frac{1}{N} \sum_{i=1}^N \sigma_i f(Z_i) \right) \right],$$

where $\sigma_1, \dots, \sigma_N$ are independent random variables uniformly chosen from $\{-1, 1\}$. Here the expectation is taken over both σ_i and data Z_i for $i = 1, \dots, N$.

Lemma 1. *For a class of quantile functions \mathcal{Q} and CRPS loss L , the Rademacher complexity of a composited class $\mathcal{F} = L \circ \mathcal{Q}$ with sample size N is upper bounded by that of \mathcal{Q} , i.e.,*

$$\mathfrak{R}_N(\mathcal{F}) \leq \mathfrak{R}_N(\mathcal{Q}).$$

Definition 2. *The temporal discrepancy over τ time difference is defined as:*

$$\Delta_{\text{dis}}^{\tau}(\mathcal{D}_{\text{test}}) = \sup_{\theta} [\mathcal{L}(\theta | \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau})].$$

Here, $\mathcal{D}_{\text{test}}^{-\tau} = \{x_i^{-\tau}\}_{i=1}^m$ denotes the backtest data with τ -shifted input $x_i^{-\tau} = (z_{i,1:T-\tau}, \xi_{i,1:T})$.

Theorem 1. *For any $\delta > 0$, a conditional generalization error*

$$\begin{aligned} \mathcal{L}(\hat{\theta} | \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta^* | \mathcal{D}_{\text{test}}) &\leq 2\Delta_{\text{dis}}^{\tau}(\mathcal{D}_{\text{test}}) \\ &+ 2\mathfrak{R}_{m\tau}(\mathcal{Q} | \mathcal{D}_{\text{test}}^{-\tau}) + \frac{4D}{\sqrt{m\tau}} \sqrt{\log \left(\frac{1}{\delta} \right)}, \end{aligned}$$

holds with at least $1 - \delta$ probability. Here, $\mathfrak{R}_{m\tau}(\mathcal{Q} | \mathcal{D}_{\text{test}}^{-\tau})$ denotes the Rademacher complexity of \mathcal{Q} , where the expectation on target data is conditioned on the backtest data $\mathcal{D}_{\text{test}}^{-\tau}$.

Theorem 2 says that the error consists of three terms: First, the discrepancy Δ^{τ} term captures errors from non-stationarity, which becomes zero if the data is stationary or periodic, and possibly nonzero otherwise. Note that this stationarity is not required in particular as a data generating process since the corresponding error would appear in the generalization errors. Second,

Rademacher complexity $\mathfrak{R}_{m\tau}$ gets smaller under larger sample sizes or richer quantile function classes \mathcal{Q} . The last one is a high probability term but only with the logarithmic dependency. All of the terms except the discrepancy diminish to zero as the sample size $m\tau$ and knot size K increase.

Remark. Even though a seemingly similar analysis exists (Kuznetsov and Mohri, 2015; Zimin and Lampert, 2017; Mariet and Kuznetsov, 2019), our theoretical results are arguably more sophisticated and informative, specifically targeted for the global Seq2Seq quantile forecasting setting. Note that the generalization errors are directly expressed with the complexity of quantile functionals \mathcal{Q} by eliminating the CRPS loss dependency (in Lemma 1) and potential quantization errors through the analytical CRPS integral supported (in Section 4.2). Moreover, we analyze the multi-horizon $\tau > 1$ cases with exogenous variables $\xi_{i,t}$, applying τ -shifted backtest, adopting τ -shifted discrepancy with sophisticated telescoping and etc. On the other hand, existing works mainly depend on the complexity of both \mathcal{F} , analyze one-step prediction without exogenous variables, and can involve quantization errors.

6.2 Unconditional Generalization Error

Recall input time series x_i to the quantile function consists of the i th observed target and input covariates. We assume that all the targets $Z_{i,1:T+\tau}$ and the input X_i are independent time series over i . The unconditional generalization loss \mathcal{L} is defined as:

$$\mathcal{L}(\theta) = \frac{1}{\tau} \sum_{t=T+1}^{T+\tau} \mathbb{E}[L(Z_{i,t}, q_{\theta}^t(\cdot | X_i))].$$

The following Lemma explicitly connects the discrepancy quantity to stationarity and beyond.

Lemma 2. *Assume each input time series is stationary or periodic. Then the expected temporal discrepancy $\Delta_{\text{dis}}^{\tau} = \sup_{\theta} \mathbb{E}[\mathcal{L}(\theta | \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau})] = 0$.*

Finally, we have a generalization error below.

Corollary 1. *For any $\delta > 0$, a generalization error*

$$\mathcal{L}(\hat{\theta}) - \mathcal{L}(\theta^*) \leq 2\Delta_{\text{dis}}^{\tau} + 2\mathfrak{R}_{m\tau}(\mathcal{Q}) + \frac{2D}{\sqrt{m\tau}} \sqrt{\log\left(\frac{1}{\delta}\right)},$$

holds with at least $1 - \delta$ probability, where $\Delta_{\text{dis}}^{\tau}$ denotes the expected temporal discrepancy.

7 EXPERIMENTS

In the experiments, we show the impact of I(S)QF on the state-of-the-art Seq2Seq forecasting model, MQ-CNN (Wen et al., 2017). We examine and compare several properties of I(S)QF including no quantile crossing,

dynamic quantile query, and accuracy improvements. We implement the model in the open-source GluonTS⁴ (Alexandrov et al., 2020) library. The experiments are done using AWS SageMaker (Liberty et al., 2020).

7.1 Experimental Setup

We conduct benchmark experiments on the **Elec** and **Traf** from the UCI data repository (Dheeru and Karra Taniskidou, 2017), **Wiki** from Kaggle (Lai, 2017), and 6 different **M4** competition datasets (Makridakis et al., 2018). We report the weighted quantile losses to measure the accuracy at various trained, interpolated and extrapolated quantile knots, the mean weighted quantile loss to approximate the CRPS, the mean scaled interval score (MSIS) (Gneiting and Raftery, 2007) to measure the sharpness and coverage of the distribution, and the crossing error to measure the severity of quantile crossing for our I(S)QF and various baselines: MLP-based QF, SQF, and Gaussian output layers. See the supplementary materials for more details on the setup and hyperparameters.

7.2 Experimental Results

Removal of Quantile Crossing Errors. The crossing errors in Table 2 show that QF suffers from quantile crossing, often severely with 42% occurrence and 0.5 of averaged crossed distance on $[0, 1]$ target domain on **Traf** dataset. On the other hand, all the other layers including our I(S)QF do not suffer from any quantile crossing. Figure 4 visualizes the quantile crossing for the QF layers, and no quantile crossing for the IQF layers on the **Elec** data.

Supporting any Quantile Query and Saving Re-training Costs. The original MQ-CNN model equipped with default QF can make inferences only on pre-determined quantile knots, resulting in N/A for other arbitrary quantile queries shown in Table 2. Therefore, inferences on different quantiles require an expensive re-training. I(S)QF shown in Figures 4b-4c and Table 2 provides the full conditional quantile function queryable on any quantile, only in seconds through inter/extrapolation. This results in re-training costs savings of approximately 10-60 minutes for each run depending on the dataset for the default QF layers.

Accuracy Improvements over Baselines. Table 2 shows that I(S)QF, especially IQF, do not compromise accuracy against QF, and often improve accuracy over other baselines. First, I(S)QF show often notable improvement compared to SQF. IQF with 5 basis knots performed far better than SQF up to 10 times and 20

⁴available in <https://github.com/awsmlabs/gluon-ts/blob/master/src/gluonts/model/seq2seq/>.

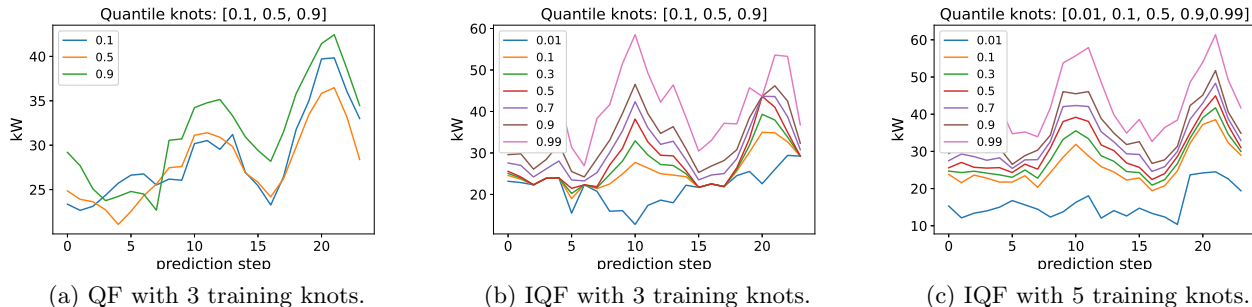


Figure 4: Experimental results for the Elec dataset under the MQ-CNN forecaster with $\tau = 24$ predicted time steps, and different quantile output layers and training quantile knots. The default quantile function (QF) is limited to query only on training quantiles, together with severe quantile crossing. On the other hand, IQF shows more consistent predictions with no quantile crossing over any quantile query.

Dataset	Strategy	mean_wQL	crossing %	wQL[0.5]	wQL[0.7] (I)	wQL[0.9]	wQL[0.995] (E)	MSIS[0.1]	MSIS[0.02]
Elec	Gaussian	0.129 \pm 0.03	0.0 \pm 0.0	0.234 \pm 0.05	0.224 \pm 0.04	0.159 \pm 0.04	0.068 \pm 0.04	42.8 \pm 12.0	79.8 \pm 22.9
	SQF	0.379 \pm 0.18	0.0 \pm 0.0	0.506 \pm 0.2	0.556 \pm 0.25	0.542 \pm 0.28	0.457 \pm 0.27	157.8 \pm 84.1	369.1 \pm 204.1
	QF	0.047 \pm 0.0	1.271 \pm 0.21	0.104 \pm 0.0	N/A	0.054 \pm 0.0	N/A	N/A	17.3 \pm 0.2
	IQF	0.055 \pm 0.01	0.0 \pm 0.0	0.12 \pm 0.01	0.101 \pm 0.01	0.056 \pm 0.0	0.009 \pm 0.0	12.2 \pm 0.3	17.6 \pm 0.5
	ISQF	0.198 \pm 0.06	0.0 \pm 0.0	0.378 \pm 0.06	0.37 \pm 0.08	0.253 \pm 0.11	0.069 \pm 0.05	66.9 \pm 34.8	109.1 \pm 71.3
Traf	Gaussian	0.154 \pm 0.01	0.0 \pm 0.0	0.299 \pm 0.02	0.3 \pm 0.02	0.205 \pm 0.03	0.075 \pm 0.03	14.5 \pm 1.3	25.1 \pm 4.0
	SQF	0.193 \pm 0.01	0.0 \pm 0.0	0.331 \pm 0.02	0.339 \pm 0.02	0.274 \pm 0.05	0.152 \pm 0.05	22.2 \pm 2.1	45.4 \pm 4.2
	QF	1.483 \pm 1.44	42.227 \pm 5.39	1.257 \pm 1.07	N/A	1.035 \pm 0.76	N/A	N/A	1128.7 \pm 1569.1
	IQF	1.31 \pm 0.74	0.0 \pm 0.0	2.046 \pm 1.11	1.705 \pm 1.23	1.132 \pm 1.07	0.917 \pm 1.36	166.5 \pm 150.9	373.6 \pm 354.0
	ISQF	0.187 \pm 0.06	0.0 \pm 0.0	0.355 \pm 0.08	0.377 \pm 0.12	0.295 \pm 0.12	0.043 \pm 0.01	18.8 \pm 7.5	29.4 \pm 18.1
Wiki	Gaussian	0.36 \pm 0.2	0.0 \pm 0.0	0.575 \pm 0.23	0.634 \pm 0.33	0.58 \pm 0.38	0.38 \pm 0.32	52.2 \pm 35.3	109.1 \pm 81.8
	SQF	0.176 \pm 0.01	0.0 \pm 0.0	0.323 \pm 0.02	0.355 \pm 0.04	0.291 \pm 0.04	0.122 \pm 0.01	21.6 \pm 2.2	40.0 \pm 3.0
	QF	0.136 \pm 0.0	0.052 \pm 0.02	0.247 \pm 0.01	N/A	0.214 \pm 0.0	N/A	N/A	34.5 \pm 0.3
	IQF	0.135 \pm 0.0	0.0 \pm 0.0	0.246 \pm 0.0	0.278 \pm 0.0	0.213 \pm 0.0	0.087 \pm 0.0	19.0 \pm 0.3	34.7 \pm 0.3
	ISQF	0.051 \pm 0.09	0.0 \pm 0.0	0.098 \pm 0.17	0.087 \pm 0.15	0.064 \pm 0.11	0.026 \pm 0.04	21.2 \pm 0.0	41.9 \pm 0.0
M4-daily	Gaussian	0.017 \pm 0.0	0.0 \pm 0.0	0.034 \pm 0.0	0.029 \pm 0.0	0.017 \pm 0.0	0.005 \pm 0.0	43.8 \pm 6.6	74.0 \pm 15.3
	SQF	0.017 \pm 0.0	0.0 \pm 0.0	0.031 \pm 0.0	0.028 \pm 0.0	0.017 \pm 0.0	0.009 \pm 0.0	48.9 \pm 5.7	95.2 \pm 17.0
	QF	0.015 \pm 0.0	0.329 \pm 0.27	0.029 \pm 0.0	N/A	0.013 \pm 0.0	N/A	N/A	52.2 \pm 1.6
	IQF	0.017 \pm 0.0	0.0 \pm 0.0	0.033 \pm 0.0	0.026 \pm 0.0	0.017 \pm 0.0	0.006 \pm 0.0	54.3 \pm 17.7	93.1 \pm 34.8
	ISQF	0.014 \pm 0.01	0.0 \pm 0.0	0.025 \pm 0.01	0.023 \pm 0.01	0.013 \pm 0.01	0.003 \pm 0.0	42.5 \pm 0.9	154.1 \pm 59.0
M4-weekly	Gaussian	0.043 \pm 0.0	0.0 \pm 0.0	0.086 \pm 0.0	0.087 \pm 0.0	0.059 \pm 0.0	0.017 \pm 0.0	60.9 \pm 6.8	81.3 \pm 9.0
	SQF	0.046 \pm 0.0	0.0 \pm 0.0	0.085 \pm 0.0	0.084 \pm 0.0	0.063 \pm 0.0	0.022 \pm 0.0	64.6 \pm 4.7	96.8 \pm 18.0
	QF	0.038 \pm 0.0	0.048 \pm 0.0	0.067 \pm 0.0	N/A	0.057 \pm 0.0	N/A	N/A	86.8 \pm 2.4
	IQF	0.042 \pm 0.0	0.0 \pm 0.0	0.069 \pm 0.0	0.072 \pm 0.0	0.058 \pm 0.0	0.019 \pm 0.01	70.1 \pm 4.3	103.7 \pm 8.3
	ISQF	0.047 \pm 0.0	0.0 \pm 0.0	0.084 \pm 0.0	0.084 \pm 0.0	0.063 \pm 0.0	0.012 \pm 0.0	65.8 \pm 5.3	157.6 \pm 31.2
M4-monthly	Gaussian	0.149 \pm 0.01	0.0 \pm 0.0	0.188 \pm 0.01	0.179 \pm 0.01	0.153 \pm 0.01	0.115 \pm 0.01	37.0 \pm 3.0	80.4 \pm 7.2
	SQF	0.12 \pm 0.01	0.0 \pm 0.0	0.179 \pm 0.01	0.168 \pm 0.01	0.128 \pm 0.01	0.079 \pm 0.01	25.3 \pm 2.0	50.8 \pm 5.8
	QF	0.087 \pm 0.0	0.239 \pm 0.04	0.131 \pm 0.0	N/A	0.09 \pm 0.0	N/A	N/A	31.5 \pm 0.7
	IQF	0.081 \pm 0.0	0.0 \pm 0.0	0.131 \pm 0.0	0.121 \pm 0.0	0.093 \pm 0.0	0.038 \pm 0.02	19.3 \pm 1.2	33.4 \pm 5.4
	ISQF	0.115 \pm 0.0	0.0 \pm 0.0	0.18 \pm 0.0	0.174 \pm 0.0	0.14 \pm 0.01	0.042 \pm 0.0	28.0 \pm 1.0	51.3 \pm 3.3

Table 2: Comparison of the accuracy for our I(S)QF and various baselines with 5 training quantile knots [0.01, 0.1, 0.5, 0.9, 0.99]. The mean and standard deviation are computed over 4 runs, and the winning method is shown in bold. (I) and (E) indicate the quantiles, where interpolation and extrapolation are performed, respectively. The MSIS[0.1] measures the 90% prediction interval using the interpolated 95th and 5th quantiles, if defined. Similarly, the MSIS[0.02] measures the 98% prediction interval using the 99th and 1st training quantiles.

times in terms of meanQuantile and MSIS, respectively on most of the datasets. Clearly, SQF has difficulty in fitting the tail due to the lack of extrapolation based on the fact that the wQL[0.995] 10 of SQF in Tables 2 is often much larger than that for our I(S)QF. This demonstrates that too many flexible knots are difficult to fit, and a moderate amount of fixed knot positions are helpful in practice. Second, the performance of the Gaussian output varies. It performs poorly on Elec, Wiki, and M4-monthly, and well on Traf, arguably due to whether the Gaussianity assumption matches

the underlying distribution. Third, IQF performs similarly to QF on 4-6 out of 9 datasets, depending on the metric. For the remaining datasets, IQF generally outperforms QF on the mean weighted quantile loss, but underperforms on the MSIS. We argue that QF sacrifices consistency, and instead gets tighter coverage. Note that some MSIS is not accessible for QF unless it is trained on the corresponding quantile. Fourth, ISQF outperforms SQF and Gaussian in general. Compared to QF/IQF, ISQF improves the performance on Traf and Wiki by large margins. It performs similarly to

IQF on 4 out of 9 datasets, and worse especially on the remaining 3 datasets. Lastly, the $wQL[0.995]$ metric shows the benefit of the extrapolation methods in I(S)QF. See Section G for the extensive results.

Effects of Training Knots and Other Parameters.

Through experiments under different number of training knots reported in Section G, we first observe that the improvement gain for IQF becomes more significant under a larger number of quantile knots. This result implies that a larger number of training knots tends to improve the accuracy since the CRPS of IQF better matches on that of the true underlying distribution. In addition, as long as the extremal knots are beyond 0.1 and 0.9, I(S)QF behaves reasonably but the most efficient choices may depend on each dataset. Table 6 in Section G.2 demonstrates these results, where 3 vs. 5 training knots are compared.

Second, IQF and QF have the same number of parameters to fit, but ISQF has more parameters due to the additional spline knots. Therefore, ISQF generally requires more hyperparameter tuning, e.g., different number of epochs and flexible spline knots. The performance of ISQF improves with more epochs (200 vs. 100), and in general we observe a (one sided part of) U-shaped curve depending on the level of difficulty of fitting on each dataset as the number of knots increases. See additional discussion and figures in Section G.2.

8 DISCUSSION

In this paper, we propose a distribution-free methodology that infers quantile estimates without quantile crossing. Our approach, the Incremental (Spline) Quantile Functions I(S)QF, is capable of inter/extrapolating to form a conditional quantile function that is accessible at any quantile query, and saves re-training costs accordingly. We apply our method to Seq2Seq time series forecasting, analyze generalization bounds, and validate the superiority of our methods through empirical studies. While Seq2Seq time series forecasting is mainly covered, our contributions open up our framework to new potential applications including autoregressive probabilistic time series forecasting or other applications like domain adaptation (Jin et al., 2022) where multiple quantile regressions are widely-used, providing additional benefits to future research. We leave for the future work the development of a variant of I(S)QF for multi-quantile prediction on multivariate time series alongside its theoretical analysis and suitable faster training schemes (Lu et al., 2021).

Acknowledgements

The authors would like to thank Syama Sundar Rangapuram, Lorenzo Stella, Hilaf Hasson, and Anoop

Deoras in Amazon research for providing their fruitful feedback at the early stage of this work.

References

- Alexandrov, A., Benidis, K., Bohlke-Schneider, M., Flunkert, V., Gasthaus, J., Januschowski, T., Maddix, D. C., Rangapuram, S., Salinas, D., Schulz, J., Stella, L., Türkmen, A. C., and Wang, Y. (2020). GluonTS: Probabilistic and Neural Time Series Modeling in Python. *Journal of Machine Learning Research*, 21(116):1–6.
- Chen, Y., Kang, Y., Chen, Y., and Wang, Z. (2020). Probabilistic forecasting with temporal convolutional neural network. *Neurocomputing*, 399:491–501.
- Chernozhukov, V., Fernández-Val, I., and Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica*, 78(3):1093–1125.
- de Bézenac, E., Rangapuram, S. S., Benidis, K., Bohlke-Schneider, M., Kurle, R., Stella, L., Hasson, H., Galinari, P., and Januschowski, T. (2020). Normalizing kalman filters for multivariate time series analysis. In *NeurIPS*.
- Dheeru, D. and Karra Taniskidou, E. (2017). UCI machine learning repository.
- Ehrlich, E., Callot, L., and Aubet, F.-X. (2021). Spliced binned-pareto distribution for robust modeling of heavy-tailed time series. *arXiv preprint arXiv:2106.10952*.
- Eisenach, C., Patel, Y., and Madeka, D. (2020). Mq-transformer: Multi-horizon forecasts with context dependent and feedback-aware attention. *arXiv preprint arXiv:2009.14799*.
- Flunkert, V., Salinas, D., and Gasthaus, J. (2017). Deepar: Probabilistic forecasting with autoregressive recurrent networks. *arXiv preprint arXiv:1704.04110*.
- Friederichs, P. and Thorarinsdottir, T. L. (2012). Forecast verification for extreme value distributions with an application to probabilistic peak wind prediction. *Environmetrics*, 23(7):579–594.
- Gasthaus, J., Benidis, K., Wang, Y., Rangapuram, S. S., Salinas, D., Flunkert, V., and Januschowski, T. (2019). Probabilistic forecasting with spline quantile function rnns. In Chaudhuri, K. and Sugiyama, M., editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1901–1910. PMLR.
- Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.

- Hyndman, R., Koehler, A. B., Ord, J. K., and Snyder, R. D. (2008). *Forecasting with exponential smoothing: the state space approach*. Springer Science & Business Media.
- Jin, X., Park, Y., Maddix, D. C., Wang, H., and Wang, Y. (2022). Domain adaptation for time series forecasting via attention sharing.
- Kim, J., Park, Y., Fox, J. D., Boyd, S. P., and Dally, W. (2020). Optimal operation of a plug-in hybrid vehicle with battery thermal and degradation model. In *2020 American Control Conference (ACC)*, pages 3083–3090. IEEE.
- Kim, T., Fakoor, R., Mueller, J., Smola, A. J., and Tibshirani, R. J. (2021). Deep quantile aggregation. *arXiv preprint arXiv:2103.00083*.
- Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs. Cambridge University Press.
- Koenker, R. and Bassett Jr, G. (1978). Regression quantiles. *Econometrica: Journal of the Econometric Society*, pages 33–50.
- Kuznetsov, V. and Mohri, M. (2015). Learning theory and algorithms for forecasting non-stationary time series. In *NIPS*, pages 541–549. Citeseer.
- Lai (2017). Dataset of Kaggle Competition Web Traffic Time Series Forecasting, Version 3.
- Liberty, E., Karnin, Z., Xiang, B., Rouesnel, L., Coskun, B., Nallapati, R., Delgado, J., Sadoughi, A., Astashonok, Y., Das, P., et al. (2020). Elastic machine learning algorithms in amazon sagemaker. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 731–737.
- Lim, B., Arik, S. O., Loeff, N., and Pfister, T. (2019). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *arXiv preprint arXiv:1912.09363*.
- Liu, Y. and Wu, Y. (2009). Stepwise multiple quantile regression estimation using non-crossing constraints. *Statistics and its Interface*, 2(3):299–310.
- Lu, Y., Park, Y., Chen, L., Wang, Y., De Sa, C., and Foster, D. (2021). Variance reduced training with stratified sampling for forecasting models. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 7145–7155. PMLR.
- Makridakis, S. et al. (2018). The M4 competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 34(4):802 – 808.
- Mariet, Z. and Kuznetsov, V. (2019). Foundations of sequence-to-sequence modeling for time series. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 408–417. PMLR.
- Park, Y., Mahadik, K., Rossi, R. A., Wu, G., and Zhao, H. (2019). Linear quadratic regulator for resource-efficient cloud services. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 488–489.
- Park, Y., Rossi, R., Wen, Z., Wu, G., and Zhao, H. (2020). Structured policy iteration for linear quadratic regulator. In *International Conference on Machine Learning*, pages 7521–7531. PMLR.
- Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Taieb, S. B., Bergmeir, C., Bessa, R. J., Bijak, J., Boylan, J. E., et al. (2022). Forecasting: theory and practice. *International Journal of Forecasting*.
- Rangapuram, S. S., Seeger, M. W., Gasthaus, J., Stella, L., Wang, Y., and Januschowski, T. (2018). Deep state space models for time series forecasting. *Advances in neural information processing systems*, 31:7785–7794.
- Rasmussen, C. and Williams, C. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- Salinas, D., Bohlke-Schneider, M., Callot, L., Medico, R., and Gasthaus, J. (2019). High-dimensional multivariate forecasting with low-rank gaussian copula processes. *Advances in Neural Information Processing Systems*, 32:6827–6837.
- Salinas, D., Flunkert, V., Gasthaus, J., and Januschowski, T. (2020). Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191.
- Schmidt, L. and Zhu, Y. (2016). Quantile spacings: A simple method for the joint estimation of multiple quantiles without crossing. *Available at SSRN 2220901*.
- Wang, Y., Smola, A., Maddix, D., Gasthaus, J., Foster, D., and Januschowski, T. (2019). Deep factors for forecasting. In *International conference on machine learning*, pages 6607–6617. PMLR.
- Wen, R. and Torkkola, K. (2019). Deep generative quantile-copula models for probabilistic forecasting. *arXiv preprint arXiv:1907.10697*.
- Wen, R., Torkkola, K., and Narayanaswamy, B. (2017). A multi-horizon quantile recurrent forecaster. *arXiv preprint arXiv:1711.11053*.
- Zimin, A. and Lampert, C. (2017). Learning theory for conditional risk minimization. In *Artificial Intelligence and Statistics*, pages 213–222. PMLR.

Supplementary Material for “Learning Quantile Functions without Quantile Crossing for Distribution-free Time Series Forecasting”

A Extension to Auto-Regressive (AR) Models

A similar procedure done for Seq2Seq models can be applied to autoregressive models. There is an additional intermediate stage, where the prediction at each time to be fed to the next time-step is sampled from I(S)QF for the next prediction in the sequence. In case of sample path generation, we sample $\hat{z}_t = q^t(\alpha_t | x)$ by drawing α_t uniformly at random at prediction time $t = T + 1, \dots, T + \tau$, recursively.

B Extrapolation

B.1 Extrapolation with Exponential Tails for IQF

The fixed parameters β_L, β_R in the exponential extrapolation for IQF are given as:

$$\beta_L = \frac{\log[(\alpha_2 + \epsilon)/(\alpha_1 + \epsilon) + \epsilon]}{\hat{q}(\alpha_2 | h) - \hat{q}(\alpha_1 | h)}, \quad \beta_R = \frac{\log[(1 - \alpha_{K-1} + \epsilon)/(1 - \alpha_K + \epsilon) + \epsilon]}{\hat{q}(\alpha_K | h) - \hat{q}(\alpha_{K-1} | h)}, \quad (9)$$

where ϵ is a tolerance parameter, set to half of machine precision, to prevent numerical errors when dividing by these parameters. Note that this choice also still enforces the desired monotonicity.

B.2 Extrapolation with Exponential Tails for ISQF

In IQF, the left and right exponential tails are required to coincide with the two leftmost and rightmost quantile knots, respectively. In ISQF, this is generalized by restricting the tails to pass through only the leftmost/rightmost quantile knots and letting the free variables to be trainable. In particular, the exponential tails are given as

$$q(\alpha | h) = \begin{cases} a_L \log(\alpha) + b_L & =: q_{\text{tail}_L}(\alpha | h), & 0 < \alpha \leq \alpha_{\text{tail}_L}, \\ a_R \log(1 - \alpha) + b_R & =: q_{\text{tail}_R}(\alpha | h), & \alpha_{\text{tail}_R} \leq \alpha < 1, \end{cases} \quad (10)$$

where $a_L = 1/\beta_L, a_R = 1/\beta_R, b_L = -a_L \log(\alpha_{\text{tail}_L}) + \hat{q}(\alpha_{\text{tail}_L} | h), b_R = a_R \log(1 - \alpha_{\text{tail}_R}) + \hat{q}(\alpha_{\text{tail}_R} | h)$.

C Analytical CRPS for I(S)QF

In this section, we provide the derivation for the analytical CRPS of I(S)QF. For notational simplicity, we denote $q(\alpha) := q(\alpha | h)$ and $\hat{q}(\alpha) := \hat{q}(\alpha | h)$. The CRPS L of I(S)QF is given by

$$L(q, z) = \int_0^1 2\rho_\alpha(z - q(\alpha))d\alpha = \left(\int_0^{\alpha_{\text{tail}_L}} + \int_{\alpha_{\text{tail}_L}}^{\alpha_{\text{tail}_R}} + \int_{\alpha_{\text{tail}_R}}^1 \right) 2\rho_\alpha(z - q(\alpha))d\alpha. \quad (11)$$

C.1 Exponential Tails

Left Tail CRPS The first term of (11) which corresponds to the left tail can be evaluated as

$$\begin{aligned} \int_0^{\alpha_{\text{tail}_L}} 2\rho_\alpha(z - q_{\text{tail}_L}(\alpha))d\alpha &= \int_0^{\alpha_{\text{tail}_L}} 2(\alpha - \mathbf{1}\{z \leq q_{\text{tail}_L}(\alpha)\})(z - q_{\text{tail}_L}(\alpha))d\alpha \\ &= \int_0^{\alpha_{\text{tail}_L}} 2\alpha(z - q_{\text{tail}_L}(\alpha))d\alpha - \int_{\hat{\alpha}_{\text{tail}_L}}^{\alpha_{\text{tail}_L}} 2(z - q_{\text{tail}_L}(\alpha))d\alpha. \end{aligned}$$

Here, the quantile level $\tilde{\alpha}_{\text{tail}_L}$ denotes the quantile, where the indicator function changes value and is given by

$$\tilde{\alpha}_{\text{tail}_L} = \begin{cases} \exp((z - b_L)/a_L), & \text{if } z < \hat{q}(\alpha_{\text{tail}_L}), \\ \alpha_{\text{tail}_L}, & \text{otherwise.} \end{cases}$$

Evaluating the integrals using (10), we obtain the following

$$\begin{aligned} \int_0^{\alpha_{\text{tail}_L}} 2\rho_\alpha(z - q_{\text{tail}_L}(\alpha))d\alpha &= (z - b_L)(\alpha_{\text{tail}_L}^2 - 2\alpha_{\text{tail}_L} + 2\tilde{\alpha}_{\text{tail}_L}) \\ &+ a_L \left[\alpha_{\text{tail}_L}^2 (-\log(\alpha_{\text{tail}_L}) + \frac{1}{2}) + 2\alpha_{\text{tail}_L}(\log(\alpha_{\text{tail}_L}) - 1) - 2\tilde{\alpha}_{\text{tail}_L}(\log(\tilde{\alpha}_{\text{tail}_L}) - 1) \right]. \end{aligned} \quad (12)$$

Right Tail CRPS By using the symmetry of the two tails, we compute the right tail CRPS using the results of the left tail in (12) as

$$\begin{aligned} \int_{\alpha_{\text{tail}_R}}^1 2\rho_\alpha(z - q_{\text{tail}_R}(\alpha))d\alpha &= - \int_0^{\alpha_{\text{tail}_L}} 2\rho_{\alpha'}(z - q_{\text{tail}_L}(\alpha'))d\alpha' \\ &= -(z - b_R)(1 + \alpha_{\text{tail}_R}^2 - 2\tilde{\alpha}_{\text{tail}_R}) - a_R \left[(1 - \alpha_{\text{tail}_R}^2) \log(1 - \alpha_{\text{tail}_R}) \right. \\ &\quad \left. + \frac{1}{2} + \frac{\alpha_{\text{tail}_R}^2}{2} + \alpha_{\text{tail}_R} - 2(1 - \tilde{\alpha}_{\text{tail}_R}) \log(1 - \tilde{\alpha}_{\text{tail}_R}) - 2\tilde{\alpha}_{\text{tail}_R} \right], \end{aligned}$$

with the change of variables: $\alpha' = 1 - \alpha$, $\alpha_{\text{tail}_R} = 1 - \alpha_{\text{tail}_L}$, $\beta_R = -\beta_L$, $\hat{q}(\alpha_{\text{tail}_R}) = \hat{q}(\alpha_{\text{tail}_L})$, $a_R = a_L$, $b_R = b_L$, $\tilde{\alpha}_{\text{tail}_R} = 1 - \tilde{\alpha}_{\text{tail}_L}$, and $q_{\text{tail}_R}(\alpha) = q_{\text{tail}_L}(\alpha')$.

C.2 GPD Tails

The analytical CPRS for the Generalized Pareto (GPD) right tail is given in Friederichs and Thorarinsdottir (2012) as:

$$\begin{aligned} \text{CRPS}_{\text{GPD}_{\text{tail}_R}}(\alpha) &= \frac{\mu}{\eta} [1 + \eta\psi_\mu(\alpha, \alpha_{\text{tail}_R})] [2\text{GPD}_{\text{tail}_R}(\alpha) - 1] \\ &\quad - \frac{2\mu}{\eta(\eta - 1)} \left[\frac{1}{\eta - 2} + [1 - \text{GPD}_{\text{tail}_R}(\alpha)] [1 + \eta\psi_\mu(\alpha, \alpha_{\text{tail}_R})] \right], \end{aligned}$$

where $\eta \neq 0$, and $\text{GPD}_{\text{tail}_R}(\alpha)$ is the expression in Equation 8 in the main body for $\alpha \geq \alpha_{\text{tail}_R}$ divided by α_{tail_R} . The $\text{CRPS}_{\text{GPD}_{\text{tail}_L}}(\alpha)$ has a similar form with α_{tail_L} .

C.3 Spline CRPS

The middle term of (11) corresponds to the non-tail region (spline), and the integral is given by

$$\int_{\alpha_k}^{\alpha_{k+1}} 2\rho_\alpha(z - q_{\text{spline}}(\alpha))d\alpha = \int_{\alpha_k}^{\alpha_{k+1}} 2\alpha(z - q_{\text{spline}}(\alpha))d\alpha - \int_{\tilde{\alpha}}^{\alpha_{k+1}} 2(z - q_{\text{spline}}(\alpha))d\alpha, \quad (13)$$

where q_{spline} denotes a linear spline with S knots defined in $[\alpha_k, \alpha_{k+1}]$ as:

$$q_{\text{spline}}(\alpha) = \sum_{s=0}^{S-1} \max\left\{ \min\left\{ \frac{\alpha - d_s}{d_{s+1} - d_s}, 1 \right\}, 0 \right\} (p_{s+1} - p_s) + \hat{q}(\alpha_k), \quad (14)$$

with $\{d_s\}$ and $\{p_s\}$ spline knots and associated quantile estimates, respectively. Note that since $p_0 = \hat{q}(\alpha_k)$ and $p_S = \hat{q}(\alpha_{k+1})$, (14) reduces to the linear interpolation of IQF when $S = 1$. Here, the quantile level $\tilde{\alpha}$ is where the indicator function changes value and given by

$$\tilde{\alpha} = \begin{cases} \alpha_k, & \text{if } z \leq \hat{q}(\alpha_k), \\ d_{s_0} + (z - q(\alpha_k) - \sum_{s=0}^{s_0-1} (p_{s+1} - p_s)) \left(\frac{d_{s_0+1} - d_{s_0}}{p_{s_0+1} - p_{s_0}} \right), & \text{if } \hat{q}(\alpha_k) < z < \hat{q}(\alpha_{k+1}), \\ \alpha_{k+1}, & \text{if } \hat{q}(\alpha_{k+1}) \leq z, \end{cases}$$

where $s_0 = \max\{s \mid q_{\text{spline}}(d_s) < z, 0 \leq s \leq S\}$. The quantile level $\tilde{\alpha}$ can be computed in $\mathcal{O}(S)$ time by evaluating the spline at the knots. Since the knots are in increasing order, and q_{spline} is a non-decreasing function, the time complexity can be reduced to $\mathcal{O}(\log(S))$ if binary search is used.

We compute the first term in (13) as follows:

$$\int_{\alpha_k}^{\alpha_{k+1}} 2\alpha(z - q_{\text{spline}}(\alpha))d\alpha = (\alpha_{k+1}^2 - \alpha_k^2)(z - \hat{q}(\alpha_k)) - 2 \sum_{s=0}^{S-1} (p_{s+1} - p_s) \int_{\alpha_k}^{\alpha_{k+1}} \max\{\min\{\frac{\alpha - d_s}{d_{s+1} - d_s}, 1\}, 0\}d\alpha. \quad (15)$$

To evaluate the integral in the second term of (15), we break up the integral at the spline knots, where $\alpha_k \leq d_s \leq d_{s+1} \leq \alpha_{k+1}$:

$$\begin{aligned} \int_{\alpha_k}^{\alpha_{k+1}} 2\alpha \max\{\min\{\frac{\alpha - d_s}{d_{s+1} - d_s}, 1\}, 0\}d\alpha &= \int_{\alpha_k}^{d_s} 0d\alpha + 2 \int_{d_s}^{d_{s+1}} \alpha \frac{\alpha - d_s}{d_{s+1} - d_s} d\alpha + 2 \int_{d_{s+1}}^{\alpha_{k+1}} \alpha d\alpha \\ &= \frac{1}{d_{s+1} - d_s} \left(\frac{2d_{s+1}^3}{3} - d_s d_{s+1}^2 + \frac{d_s^3}{3} \right) + \alpha_{k+1}^2 - d_{s+1}^2. \end{aligned}$$

Substituting this into (15) gives:

$$\int_{\alpha_k}^{\alpha_{k+1}} 2\alpha(z - q_{\text{spline}}(\alpha))d\alpha = (\alpha_{k+1}^2 - \alpha_k^2)(z - \hat{q}(\alpha_k)) - \sum_{s=0}^{S-1} (p_{s+1} - p_s) \left[\frac{1}{d_{s+1} - d_s} \left(\frac{2d_{s+1}^3}{3} - d_s d_{s+1}^2 + \frac{d_s^3}{3} \right) + \alpha_{k+1}^2 - d_{s+1}^2 \right]. \quad (16)$$

Similarly, we compute the second term in (13) as follows:

$$\int_{\tilde{\alpha}}^{\alpha_{k+1}} 2(z - q_{\text{spline}}(\alpha))d\alpha = 2(\alpha_{k+1} - \tilde{\alpha})(z - \hat{q}(\alpha_k)) - 2 \sum_{s=0}^{S-1} (p_{s+1} - p_s) \int_{\tilde{\alpha}}^{\alpha_{k+1}} \max\{\min\{\frac{\alpha - d_s}{d_{s+1} - d_s}, 1\}, 0\}d\alpha. \quad (17)$$

Here, we have the following three cases: 1. $\tilde{\alpha} \leq d_s \leq d_{s+1}$, 2. $d_s \leq \tilde{\alpha} \leq d_{s+1}$, 3. $d_s \leq d_{s+1} \leq \tilde{\alpha}$. Hence, we can succinctly write the integral in the second term of (17) as:

$$\begin{aligned} \int_{\tilde{\alpha}}^{\alpha_{k+1}} 2 \max\{\min\{\frac{\alpha - d_s}{d_{s+1} - d_s}, 1\}, 0\}d\alpha &= 2 \int_{r_s}^{d_{s+1}} \frac{\alpha - d_s}{d_{s+1} - d_s} d\alpha + 2 \int_{\max(\tilde{\alpha}, d_{s+1})}^{\alpha_{k+1}} d\alpha \\ &= d_{s+1}^2 - 2d_s d_{s+1} - r_s^2 - 2d_s r_s + 2(\alpha_{k+1} - \max(\tilde{\alpha}, d_{s+1})), \end{aligned}$$

where $r_s = \max\{\min\{\tilde{\alpha}, d_{s+1}\}, d_s\}$.

Substituting this into (17) gives:

$$\begin{aligned} \int_{\tilde{\alpha}}^{\alpha_{k+1}} 2(z - q_{\text{spline}}(\alpha))d\alpha &= 2(\alpha_{k+1} - \tilde{\alpha})(z - \hat{q}(\alpha_k)) - \sum_{s=0}^{S-1} (p_{s+1} - p_s) \left[\frac{1}{d_{s+1} - d_s} \left(d_{s+1}^2 - 2d_s d_{s+1} - r_s^2 + 2d_s r_s \right) \right. \\ &\quad \left. + 2(\alpha_{k+1} - \max(\tilde{\alpha}, d_{s+1})) \right]. \end{aligned} \quad (18)$$

In summary, substituting the integrals in (16) and (18) into (13) gives the following final expression:

$$\begin{aligned} \int_{\alpha_k}^{\alpha_{k+1}} 2\rho_\alpha(z - q_{\text{spline}}(\alpha))d\alpha &= (\alpha_{k+1}^2 - \alpha_k^2 - 2(\alpha_{k+1} - \tilde{\alpha}))(z - \hat{q}(\alpha_k)) \\ &\quad + \sum_{s=0}^{S-1} (p_{s+1} - p_s) \left[\frac{1}{d_{s+1} - d_s} \left(d_{s+1}^2 \left(-\frac{2}{3}d_{s+1} + d_s + 1 \right) - d_s \left(\frac{d_s^2}{3} + 2d_{s+1} \right) \right. \right. \\ &\quad \left. \left. - r_s(r_s - 2d_s) \right) - \alpha_{k+1}^2 + d_{s+1}^2 + 2\alpha_{k+1} - 2\max(\tilde{\alpha}, d_{s+1}) \right]. \end{aligned}$$

D Equivalence of the Splines of ISQF and SQF

We show that the spline of ISQF is equivalent to SQF. We first recall for $\alpha \in [\alpha_k, \alpha_{k+1}]$, SQF is given by

$$\text{SQF}(\alpha) = \hat{q}(\alpha_k) + \sum_{s=0}^{S-1} c_s \max\{\alpha - d_s, 0\}, \quad (19)$$

where $\{d_s\}$ denote the x -coordinate of the spline knots as in ISQF, and $\{c_s\}$ denote the parameters determining the slopes of the pieces. Since $\{d_s\}$ is in an increasing order, the following holds:

$$\sum_{s=0}^{S-1} c_s \max\{\alpha - d_s, 0\} = \sum_{s=0}^{S-1} m_s \max\{\min\{\alpha - d_s, d_{s+1} - d_s\}, 0\}, \quad (20)$$

where $m_s = \sum_{i=0}^s c_i$ denotes the slope of the $(s+1)$ -th piece. Substituting (20) into (19), we have

$$\text{SQF}(\alpha) = \hat{q}(\alpha_k) + \sum_{s=0}^{S-1} m_s \max\{\min\{\alpha - d_s, d_{s+1} - d_s\}, 0\}. \quad (21)$$

Letting $m_s = (p_{s+1} - p_s)/(d_{s+1} - d_s)$ and substituting it into (21), we obtain

$$\begin{aligned} \text{SQF}(\alpha) &= \hat{q}(\alpha_k) + \sum_{s=0}^{S-1} \frac{p_{s+1} - p_s}{d_{s+1} - d_s} \max\{\min\{\alpha - d_s, d_{s+1} - d_s\}, 0\} \\ &= \hat{q}(\alpha_k) + \sum_{s=0}^{S-1} (p_{s+1} - p_s) \max\left\{\min\left\{\frac{\alpha - d_s}{d_{s+1} - d_s}, 1\right\}, 0\right\}, \end{aligned}$$

which is equivalent to the spline of ISQF (14).

E Proofs of Generalization Errors

E.1 Proof of Lemma 1

Lemma 3. *Let $z_1, z_2 \in \mathcal{Z}_D$ where $\mathcal{Z}_D \subset \mathbb{B}_D := \{z \mid \|z\| \leq D\}$, then*

$$|L(z_1, q) - L(z_2, q)| \leq 2D.$$

Proof. Without loss of generality, let $z_1 \leq z_2$. Then, for a fixed α with $q := q(\alpha)$,

$$\rho_\alpha(q - z_1) - \rho_\alpha(q - z_2) = \begin{cases} \alpha(z_2 - z_1), & q \geq z_2, \\ q - z_2 + \alpha(z_2 - z_1), & z_2 \leq q \leq z_1, \\ (1 - \alpha)(z_2 - z_1), & q \leq z_1. \end{cases}$$

Since

$$(1 - \alpha)(z_1 - z_2) \leq q - z_2 + \alpha(z_2 - z_1) \leq \alpha(z_2 - z_1),$$

under $z_2 \leq q \leq z_1$,

$$|\rho_\alpha(q - z_1) - \rho_\alpha(q - z_2)| \leq \max(\alpha, 1 - \alpha)|z_2 - z_1| \leq |z_2 - z_1|,$$

holds. Thus

$$\begin{aligned} \int \rho_\alpha(q(\alpha) - z_1) - \rho_\alpha(q(\alpha) - z_2) d\alpha &\leq \int |\rho_\alpha(q(\alpha) - z_1) - \rho_\alpha(q(\alpha) - z_2)| d\alpha \\ &\leq |z_2 - z_1| \\ &\leq 2D. \end{aligned}$$

□

Lemma 4. CRPS $L(q, z)$ is 1-Lipschitz continuous w.r.t any quantile function $q \in \mathcal{Q}$ in L_1 norm, i.e., for any $z \in \mathcal{Z}_D$ and $q_1, q_2 \in \mathcal{Q}$,

$$|L(z, q_1) - L(z, q_2)| \leq \|q_1 - q_2\|_1.$$

Proof. Similar to the proof of Lemma 3, for a fixed α and z ,

$$|\rho_\alpha(q_1(\alpha) - z) - \rho_\alpha(q_2(\alpha) - z)| \leq |q_2(\alpha) - q_1(\alpha)|,$$

holds. Thus

$$\begin{aligned} \left| \int \rho_\alpha(q_1(\alpha) - z) - \rho_\alpha(q_2(\alpha) - z) d\alpha \right| &\leq \int |\rho_\alpha(q_1(\alpha) - z) - \rho_\alpha(q_2(\alpha) - z)| d\alpha \\ &\leq \int |q_2(\alpha) - q_1(\alpha)| d\alpha \\ &= \|q_2 - q_1\|_1. \end{aligned}$$

□

Lemma 5 (Lemma 1 in main body). The Rademacher complexity on $L \circ q \in \mathcal{F}$ is upper-bounded by

$$\mathfrak{R}_N(\mathcal{F}) \leq \mathfrak{R}_N(\mathcal{Q}),$$

where $q \in \mathcal{Q}$.

Proof. The proof is immediate from Lemma 4 and the property of Rademacher complexity. □

E.2 Proof of Theorem 1

Recall that $\mathcal{D}_{\text{test}} = \{x_i\}$ with $x_i = (z_{i,1:T}, \xi_{i,1:T+\tau})$, $\mathcal{D}_{\text{train}} = \{z_{i,T-\tau+1:T}, x_i^{-\tau}\}$, and $\mathcal{D}_{\text{test}}^{-\tau} = \{x_i^{-\tau}\}_{i=1}^m$ with $x_i^{-\tau} = (z_{i,1:T-\tau}, \xi_{i,1:T})$.

Lemma 6. Let $\{\bar{z}_{i,t}\}_{i,t}$ be the same as $\{z_{i,t}\}_{i,t}$ for all $1 \leq t \leq T$ and $i \in [m]$ except for one element, i.e., $\bar{z}_{i,t} \neq z_{i,t}$ for some $T - \tau + 1 \leq t \leq T$ and some $i \in [m]$, and $\bar{z}_{i,t} = z_{i,t}$ otherwise. Then, for $\phi(\{z_{i,T-\tau+1:T}, x_i^{-\tau}\}) := \sup_\theta [\mathcal{L}(\theta | \{x_i^{-\tau}\}) - \hat{\mathcal{L}}(\theta; \{z_{i,T-\tau+1:T}, x_i^{-\tau}\})]$,

$$|\phi(\{z_{i,T-\tau+1:T}, x_i^{-\tau}\}) - \phi(\{\bar{z}_{i,T-\tau+1:T}, x_i^{-\tau}\})| \leq \frac{2D}{m\tau},$$

holds.

Proof.

$$\begin{aligned} |\phi(\{z_{i,T-\tau+1:T}, x_i^{-\tau}\}) - \phi(\{\bar{z}_{i,T-\tau+1:T}, x_i^{-\tau}\})| &\stackrel{(a)}{\leq} \sup_\theta |\hat{\mathcal{L}}(\theta; \{\bar{z}_{i,T-\tau+1:T}, x_i^{-\tau}\}) - \hat{\mathcal{L}}(\theta; \{z_{i,T-\tau+1:T}, x_i^{-\tau}\})| \\ &\stackrel{(b)}{\leq} \frac{1}{m\tau} \sum_{i=1}^m \sum_{t=T-\tau+1}^T |L(\bar{z}_{i,t}, q_\theta^t(\cdot | x_i^{-\tau})) - L(z_{i,t}, q_\theta^t(\cdot | x_i^{-\tau}))| \\ &\stackrel{(c)}{\leq} \frac{2D}{m\tau}, \end{aligned}$$

where (a) holds due to $|\sup A - \sup B| \leq \sup |A - B|$, (b) holds due to sub-additivity of the supremum and the definition of the supremum attainer $\tilde{\theta}$, and (c) holds due to the assumption on $\{z_{i,t}\}$, and $\{\bar{z}_{i,t}\}$ and Lemma 3. □

Lemma 7. For $\phi(\mathcal{D}_{\text{train}}) = \phi(\{z_{i,T-\tau+1:T}, x_i^{-\tau}\}) := \sup_\theta [\mathcal{L}(\theta | \{x_i^{-\tau}\}) - \hat{\mathcal{L}}(\theta; \{z_{i,T-\tau+1:T}, x_i^{-\tau}\})] = \sup_\theta [\mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau}) - \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}})]$,

$$\Pr(\phi(\mathcal{D}_{\text{train}}) - \mathbb{E}\phi(\mathcal{D}_{\text{train}}) \geq \epsilon) \leq e^{-m\tau\epsilon^2/2D^2},$$

holds where \mathbb{E} is the expectation over $Z_{i,T-\tau+1:T}$ conditioned on $x_i^{-\tau}$ for $i \in [m]$ appeared in $\mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau})$.

Proof. The result is immediate through applying Lemma 6 into McDiarmid's inequality. \square

Assumption 2. The target time series $Z_{i,t}$ are bounded, i.e., $\|Z_{i,t}\| \leq D$ for some $D \geq 0$.

Definition 3. The temporal discrepancy over τ time difference is defined as:

$$\Delta_{\text{dis}}^\tau(\mathcal{D}_{\text{test}}) = \sup_{\theta} [\mathcal{L}(\theta | \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau})].$$

Here, $\mathcal{D}_{\text{test}}^{-\tau} = \{x_i^{-\tau}\}_{i=1}^m$ denotes the backtest data with τ -shifted input $x_i^{-\tau} = (z_{i,1:T-\tau}, \xi_{i,1:T})$.

Theorem 2 (Theorem 1 in the main body). For any $\delta > 0$, a conditional generalization error

$$\mathcal{L}(\hat{\theta} | \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta^* | \mathcal{D}_{\text{test}}) \leq 2\Delta_{\text{dis}}^\tau(\mathcal{D}_{\text{test}}) + 2\mathfrak{R}_{m\tau}(\mathcal{Q} | \mathcal{D}_{\text{test}}^{-\tau}) + \frac{4D}{\sqrt{m\tau}} \sqrt{\log\left(\frac{1}{\delta}\right)},$$

holds with at least $1 - \delta$ probability. Here, $\mathfrak{R}_{m\tau}(\mathcal{Q} | \mathcal{D}_{\text{test}}^{-\tau})$ denotes the Rademacher complexity of \mathcal{Q} , where the expectation on target data is conditioned on the backtest data.

Proof. Let $\mathcal{L}(\theta) := \mathcal{L}(\theta | \mathcal{D}_{\text{test}})$ and $\hat{\mathcal{L}}(\theta) = \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}})$, together with $\hat{\theta} = \text{argmin } \hat{\mathcal{L}}(\theta)$ and $\theta^* = \text{argmin } \mathcal{L}(\theta)$. Then

$$\begin{aligned} \mathcal{L}(\hat{\theta}) - \mathcal{L}(\theta^*) &= \mathcal{L}(\hat{\theta}) - \hat{\mathcal{L}}(\hat{\theta}) + \hat{\mathcal{L}}(\hat{\theta}) - \hat{\mathcal{L}}(\theta^*) + \hat{\mathcal{L}}(\theta^*) - \mathcal{L}(\theta^*) \\ &\leq \mathcal{L}(\hat{\theta}) - \hat{\mathcal{L}}(\hat{\theta}) + \hat{\mathcal{L}}(\theta^*) - \mathcal{L}(\theta^*) \\ &\leq 2 \sup_{\theta} |\mathcal{L}(\theta) - \hat{\mathcal{L}}(\theta)|. \end{aligned} \quad (22)$$

The sub-additivity of the supremum gives

$$\sup[\mathcal{L}(\theta | \mathcal{D}_{\text{test}}) - \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}})] \leq \underbrace{\sup[\mathcal{L}(\theta | \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau})]}_{\Delta_{\text{dis}}^\tau(\mathcal{D}_{\text{test}})} + \underbrace{\sup[\mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau}) - \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}})]}_{\phi(\mathcal{D}_{\text{train}})}, \quad (23)$$

where

$$\begin{aligned} L(\theta | \mathcal{D}_{\text{test}}) &= \frac{1}{m\tau} \sum_{i=1}^m \left[\sum_{t=T+1}^{T+\tau} \mathbb{E} [L(Z_{i,t}, q_\theta^t(\cdot | X_i)) | X_i = x_i] \right], \\ L(\theta | \mathcal{D}_{\text{test}}^{-\tau}) &= \frac{1}{m\tau} \sum_{i=1}^m \left[\sum_{t=T-\tau+1}^T \mathbb{E} [L(Z_{i,t}, q_\theta^t(\cdot | X_i^{-\tau})) | X_i^{-\tau} = x_i^{-\tau}] \right], \\ \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}}) &= \frac{1}{m\tau} \sum_{i=1}^m \sum_{t=T-\tau+1}^T L(z_{i,t}, q_\theta^t(\cdot | x_i^{-\tau})). \end{aligned}$$

The first term on RHS follows from the definition of $\Delta_{\text{dis}}^\tau(\mathcal{D}_{\text{test}})$.

For the second term $\phi(\mathcal{D}_{\text{train}})$, let's define

$$L^t(h) = \frac{1}{m\tau} \sum_{i=1}^m \left[\sum_{\gamma=T-\tau+1}^t \mathbb{E} [L(Z_{i,t}, q_\theta^\gamma(\cdot | X_i^{-\tau})) | X_i^{-\tau} = x_i^{-\tau}] + \sum_{\gamma=t+1}^T L(z_{i,t}, q_\theta^\gamma(\cdot | x_i^{-\tau})) \right],$$

for any $t = T - \tau + 1, \dots, T$. Then telescoping the term gives

$$\begin{aligned} \mathcal{L}(\theta | \mathcal{D}_{\text{test}}^{-\tau}) - \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}}) &= \mathcal{L}(h | \{x_i^{-\tau}\}) - \hat{\mathcal{L}}(h; \{z_{i,T-\tau+1:T}, x_i^{-\tau}\}) \\ &= L^T(h) - L^{T-\tau}(h) \\ &= \sum_{t=T-\tau+1}^T L^t(h) - L^{t-1}(h) \\ &= \sum_{i=1}^m \frac{1}{m\tau} \left[\sum_{t=T-\tau+1}^T \mathbb{E} [L(Z_{i,t}, q_\theta^t(\cdot | X_i^{-\tau})) | X_i^{-\tau} = x_i^{-\tau}] - L(z_{i,t}, q_\theta^t(\cdot | x_i^{-\tau})) \right]. \end{aligned}$$

Now the expectation of ϕ conditioned on $\{x_i^{-\tau}\}$ is

$$\begin{aligned} \mathbb{E}\phi(\mathcal{D}_{\text{train}}) &= \mathbb{E} \sup[\mathcal{L}(\theta \mid \mathcal{D}_{\text{test}}^{-\tau}) - \hat{\mathcal{L}}(\theta; \mathcal{D}_{\text{train}})] \\ &\stackrel{(a)}{\leq} \mathbb{E} \sup \frac{1}{m\tau} \sum_{i,t} |\mathbb{E}f(\theta; i, t) - f(\theta; i, t)| \\ &\stackrel{(b)}{\leq} \mathfrak{R}_{m\tau}(\mathcal{F} \mid \{x_i^{-\tau}\}), \end{aligned} \tag{24}$$

where (a) holds by defining

$$\begin{aligned} \mathbb{E}f(\theta; i, t) &= \mathbb{E}_{Z_{i,t} \mid X_i^{-\tau}} [L(Z_{i,t}, q_{\theta}^t(\cdot \mid X_i^{-\tau})) \mid X_i^{-\tau} = x_i^{-\tau}], \\ f(\theta; i, t) &= L(z_{i,t}, q_{\theta}^t(\cdot \mid x_i^{-\tau})), \end{aligned}$$

and (b) holds due to the standard symmetrization trick followed by the definition of Rademacher complexity on $f = L \circ q \in \mathcal{F}$.

Aggregating (23) and (24) gives

$$G_{m\tau} := \sup_{\theta} |\mathcal{L}(\theta) - \hat{\mathcal{L}}(\theta)| - (\Delta_{\text{dis}}^{\tau}(\mathcal{D}_{\text{test}}) + \mathbb{E}\phi(\mathcal{D}_{\text{train}})) \leq \phi(\mathcal{D}_{\text{train}}) - \mathbb{E}\phi(\mathcal{D}_{\text{train}}).$$

Writing $\phi := \phi(\mathcal{D}_{\text{train}})$ in shorthand, we have

$$\begin{aligned} \Pr(G_{m\tau} \geq \epsilon) &\leq \Pr(\phi - \mathbb{E}\phi \geq \epsilon) \\ &\stackrel{(c)}{\leq} e^{-m\tau\epsilon^2/2D^2}, \end{aligned}$$

where (c) hold due to Lemma 7. Choosing $\delta > 0$ and setting

$$\epsilon = \sqrt{\frac{2D^2}{m\tau} \log\left(\frac{1}{\delta}\right)},$$

guarantees the following w.p. $1 - \delta$:

$$\begin{aligned} \sup |\mathcal{L}(\theta) - \hat{\mathcal{L}}(\theta)| &\leq \Delta_{\text{dis}}^{\tau}(\mathcal{D}_{\text{test}}) + \mathbb{E}\phi + \sqrt{\frac{2D^2}{m\tau} \log\left(\frac{1}{\delta}\right)} \\ &\stackrel{(d)}{\leq} \Delta_{\text{dis}}^{\tau}(\mathcal{D}_{\text{test}}) + \mathfrak{R}_{m\tau}(\mathcal{Q} \mid \mathcal{D}_{\text{test}}^{-\tau}) + \sqrt{\frac{2D^2}{m\tau} \log\left(\frac{1}{\delta}\right)}, \end{aligned}$$

where (d) holds due to Equation (24) and Lemma 5. Finally, combining with (22) gives the desired result. \square

E.3 Proof of Lemma 2

Lemma 8 (Lemma 2 in the main body). *Assume each input time series is stationary or periodic. Then the expected temporal discrepancy $\Delta_{\text{dis}}^{\tau} = \sup_{\theta} \mathbb{E}[\mathcal{L}(\theta \mid \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta \mid \mathcal{D}_{\text{test}}^{-\tau})] = 0$.*

Proof. Recall the temporal discrepancy over τ time difference $\Delta_{\text{dis}}^{\tau}(\mathcal{D}_{\text{test}}) = \sup_{\theta} [\mathcal{L}(\theta \mid \mathcal{D}_{\text{test}}) - \mathcal{L}(\theta \mid \mathcal{D}_{\text{test}}^{-\tau})]$.

Equivalently,

$$\begin{aligned} \Delta_{\text{dis}}^{\tau}(\mathcal{D}_{\text{test}}) &= \sup \left[\frac{1}{m\tau} \sum_{i=1}^m \sum_{t=T+1}^{T+\tau} \mathbb{E} [L(Z_{i,t}, q_{\theta}^t(\cdot \mid x_i))] - \mathbb{E} [L(Z_{i,t-\tau}, q_{\theta}^{t-\tau}(\cdot \mid x_i^{-\tau}))] \right] \\ &= \sup \left[\frac{1}{m\tau} \sum_{i=1}^m \sum_{t=T+1}^{T+\tau} \sum_{\gamma=t-\tau+1}^t \mathbb{E} [L(Z_{i,\gamma}, q_{\theta}^{\gamma}(\cdot \mid x_i^{\gamma}))] - \mathbb{E} [L(Z_{i,\gamma-1}, q_{\theta}^{\gamma-1}(\cdot \mid x_i^{\gamma-1}))] \right], \end{aligned}$$

where we denote $x_i^\gamma = (z_{i,1:\gamma-\tau}, \xi_{1:\gamma})$, $x_i^{T+\tau} := x_i = (z_{i,1:T}, \xi_{1:T+\tau})$ and $x_i^T := x_i^{-\tau} = (z_{i,1:T-\tau}, \xi_{1:T})$.

Note that, after swapping second and third summation and taking the expectation, applying sub-additivity of supremum gives

$$\Delta_{\text{dis}}^\tau \leq \sum_{\gamma=T+1}^{T+\tau} \Delta_{\text{dis}}^\gamma,$$

where $\Delta_{\text{dis}}^\gamma = \sup[\mathbb{E}\mathcal{L}(\theta \mid \{x_i^\gamma\}) - \mathbb{E}\mathcal{L}(\theta \mid \{x_i^{\gamma-1}\})]$.

The rest of the proof is similar to Mariet and Kuznetsov (2019). It is immediate to see $\Delta_{\text{dis}}^\gamma = 0$ as $\Pr(Z_{i,\gamma-\tau+1:\gamma} \mid X_i^\gamma) = \Pr(Z_{i,\gamma-\tau:\gamma-1} \mid X_i^{\gamma-1})$ under stationarity or periodicity, which gives the desired result. \square

E.4 Proof of Corollary 1

Corollary 2 (Corollary 1 in the main body). *For any $\delta > 0$, a generalization error*

$$\mathcal{L}(\hat{\theta}) - \mathcal{L}(\theta^*) \leq 2\Delta_{\text{dis}}^\tau + 2\mathfrak{R}_{m\tau}(\mathcal{Q}) + \frac{4D}{\sqrt{m\tau}} \sqrt{\log\left(\frac{1}{\delta}\right)},$$

holds with at least $1 - \delta$ probability, where Δ_{dis}^τ denotes the expected temporal discrepancy.

Proof. We begin with expected version $L(\theta)$ on Equation (22) and (23) and its empirical version under i.i.d. setup. Then the rest of the proof is basically the same. \square

F Metrics

We use the `Evaluator` class from `GluonTS`⁵ to report the following probabilistic error metrics.

F.1 Weighted Quantile Losses (`wQL`[α])

For a given quantile $\alpha \in (0, 1)$, a target value z_t , input x , and α -quantile prediction $q^t(\alpha \mid x)$, the α -quantile loss is defined as:

$$\rho_\alpha(z_t - q^t(\alpha \mid x)) = (z_t - q^t(\alpha \mid x))(\alpha - \mathbf{1}\{z_t - q^t(\alpha \mid x) < 0\}).$$

We report the normalized sum of quantile losses,

$$\text{wQL}[\alpha] = 2 \frac{\sum_{i,t} \rho_\alpha(z_{i,t} - q^t(\alpha \mid x_i))}{\sum_{i,t} |z_{i,t}|},$$

to compute the weighted quantile losses for a given time span $t = T, \dots, T + \tau$ across all time series. We include results for the training quantiles $\alpha = 0.5, 0.9$ and interpolated quantile $\alpha = 0.7$, and the extrapolated quantile $\alpha = 0.995$.

F.2 Mean Weighted Quantile Loss (`mean_wQL`)

We also report the mean weighted quantile loss by averaging the weighted quantile losses over all of the $k = 1, \dots, K$ quantiles α_k as:

$$\text{mean_wQL} = \frac{1}{K} \sum_{k=1}^K \text{wQL}[\alpha_k],$$

as an approximation to the CRPS.

⁵<https://github.com/awsmlabs/gluon-ts/blob/master/src/gluonts/evaluation/>.

F.3 Percent Quantile Crossing

We quantify the effect of quantile crossing by introducing the percent quantile crossing metric as the proportion of adjacent quantiles that are crossing for time series at every predicted time step as:

$$\frac{1}{m\tau(K-1)} \sum_{i,t,k} \mathbf{1}\{q^t(\alpha_k | x_i) > q^t(\alpha_{k+1} | x_i), \alpha_k < \alpha_{k+1}\} \times 100\%,$$

for some input x_i , where $i = 1, \dots, m$, $t = T + 1, \dots, T + \tau$, and $k = 1, \dots, K - 1$. Note that this metric only measures the number of quantile crossing between adjacent quantiles, and is not reflective of the potential total amount of quantile crossings. For example in Figure 4a, we observe extreme crossing with the QF output layer between the 0.1 and 0.9 non-adjacent quantile levels on the `Elec` dataset, which is not counted under this metric. Hence, this metric is a lower bound on the total number of quantile crossings exhibited by QF.

F.4 Mean Scaled Interval Score (MSIS)

The mean scaled interval score (MSIS) (Gneiting and Raftery, 2007) is another probabilistic metric to measure the sharpness and coverage of the prediction interval. It is also reported in Gasthaus et al. (2019); Kim et al. (2021), and used as a metric in the M4 forecasting competition⁶. The $\text{MSIS}[\zeta]$ is defined as follows:

$$\begin{aligned} \text{MSIS}[\zeta] = \frac{1}{\text{SE}(z)} \left(\frac{1}{m\tau} \sum_{i,t} (q^t(\alpha^U | x_i) - q^t(\alpha^L | x_i)) + \frac{2}{\zeta} \left[(q^t(\alpha^L | x_i) - z_{i,t}) \mathbf{1}\{z_{i,t} < q^t(\alpha^L | x_i)\} \right. \right. \\ \left. \left. + (z_{i,t} - q^t(\alpha^U | x_i)) \mathbf{1}\{z_{i,t} > q^t(\alpha^U | x_i)\} \right] \right), \end{aligned}$$

where $i = 1, \dots, m$, $t = T + 1, \dots, T + \tau$, the upper quantile $\alpha^U = 1 - \zeta/2$, and the lower quantile $\alpha^L = \zeta/2$. The seasonal error SE for time series frequency f is given as:

$$\text{SE}(z) = \frac{1}{m(T-f)} \sum_{i,t'} |z_{i,t'} - z_{i,t'+f}|,$$

where $t' = 1, \dots, T - f$.

G Additional Experiments

G.1 Setup

Datasets. We use the open-source versions of the benchmarking datasets available in the GluonTS dataset repository⁷ as referenced in Table 3.

Forecast Models and Parameters. We select the MQ-CNN (Wen et al., 2017) model, which is a state-of-the-art time series quantile forecasting model in the sequence-to-sequence framework. The `MQCNNEstimator`⁸ is available in the open-source GluonTS (Alexandrov et al., 2020) package. For each dataset, we use the default time features $\xi_t \in \mathbb{R}^d$, e.g. day of the week and hour of the day, as covariates, where d is determined by the time series frequency.

We use the default hyperparameters optimized for the default QF layer from GluonTS. In our experiments, we only vary the training quantiles from 3 to 5 knots for all layers, and the number of spline knots and epochs for ISQF. In our experiments, we can simply toggle between the QF and IQF output layers using the `is_iqf` hyper-parameter. Similarly, to test the other three output layers, e.g. ISQF, SQF and Gaussian, we pass the corresponding `DistributionOutput` class to the `distr_output` hyperparameter. Table 4 summarizes the hyper-parameter settings for the various output layers.

⁶<https://www.m4.unic.ac.cy/wp-content/uploads/2018/03/M4-Competitors-Guide.pdf>.

⁷<https://github.com/awsmlabs/gluon-ts/blob/master/src/gluonts/dataset/repository/datasets.py>.

⁸https://github.com/awsmlabs/gluon-ts/blob/master/src/gluonts/model/seq2seq/_mq_dnn_estimator.py.

Learning Quantile Function without Quantile Crossing

DOMAIN	NAME	SUPPORT	FREQ	NO. TS	AVG. LEN	PRED. LEN	NO. COVARIATES
electrical load	Elec	\mathbb{R}^+	H	321	21044	24	4
road traffic	Traf	$[0, 1]$	H	862	14036	24	4
visit counts of wikipedia pages	Wiki	N	D	9535	762	30	3
M4 forecasting competition	M4-daily	\mathbb{R}^+	D	4227	2357	14	3
	M4-weekly	\mathbb{R}^+	W	359	1022	13	2
	M4-monthly	\mathbb{R}^+	M	48000	216	18	1
	M4-quarterly	\mathbb{R}^+	Q	24000	92	8	1
	M4-yearly	\mathbb{R}^+	Y	23000	31	6	0

Table 3: Summary of dataset statistics, where **Elec** and **Traf** are derived from the UCI data repository (Dheeru and Karra Taniskidou, 2017), **Wiki** from Kaggle (Lai, 2017), and 6 different M4 competition datasets (Makridakis et al., 2018).

Strategy	is_iqf	distr_output	epochs	no. spline knots (S)
Gaussian	N/A	GaussianOutput	100	N/A
SQF	N/A	PiecewiseLinearOutput	100	10
QF	False	N/A	100	N/A
IQF	True	N/A	100	N/A
ISQF	N/A	ISQFOutput	200	3

Table 4: Hyperparameter settings for the MQCNEstimator for our I(S)QF and the various other output layer baselines with the training quantiles set to $[0.1, 0.5, 0.9]$ (3) and $[0.01, 0.1, 0.5, 0.9, 0.99]$ (5) in the experiments.

G.2 All Benchmark Results

Table 6 shows the complete benchmarking results on the 9 real-world datasets.

ISQF spline knots experiments Figure 5 shows the mean weighted quantile loss of ISQF when varying the number of spline knots, both for 3 and 5 quantile knots. We observe that the added flexibility and complexity of using spline knots can improve the performance drastically up to a certain point at which it starts to slowly deteriorate, forming a U-shaped curve where the left side is much steeper than the right one. The point at which the performance starts to slowly deteriorate depends on the dataset. On some datasets, the added flexibility is needed as we see for the **Traf** dataset on the left side of the Figures 5c-5d. On the **Elec** dataset in Figures 5a-5b, it appears that the added flexibility is hurting the performance, and we see that the mean weighted quantile loss slowly increases as the number of spline knots increase. This is a special case where the center of the U-shaped curve is at 0, which corresponds to IQF. Similarly, in Figure 5c, the performance slowly degrades after 5 spline knots, showing the beginning of the right part of a U-shaped curve.

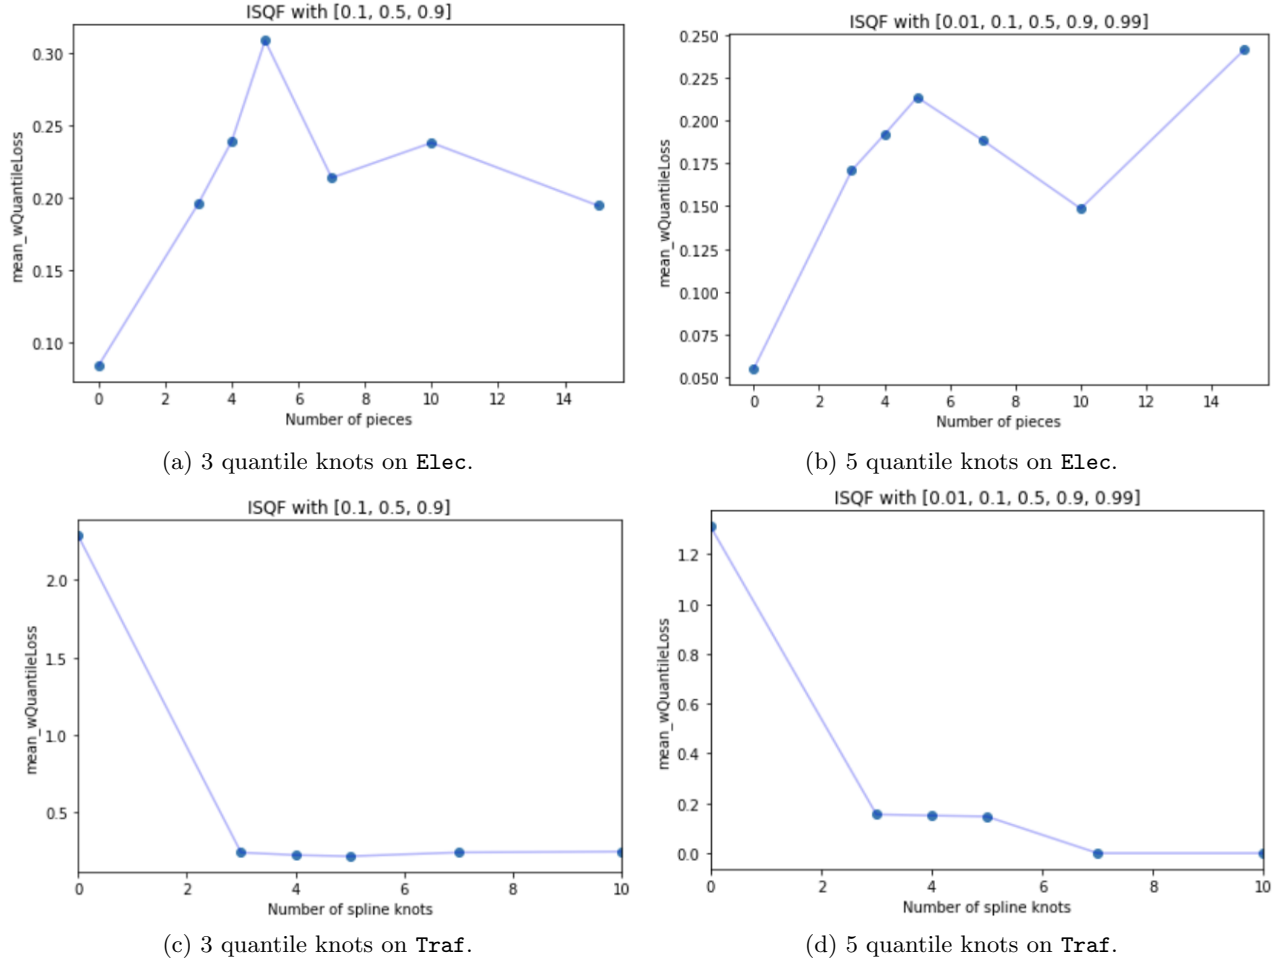


Figure 5: Experiments that vary the number of spline knots for ISQF on Elec (top) and Traf (bottom) datasets with various quantile knots [0.1, 0.5, 0.9] (3) (left) and [0.01, 0.1, 0.5, 0.9, 0.99] (5) (right).

DeepAR ISQF results Table 5 shows the results of applying the auto-regressive probabilistic DeepAR model with different benchmark layers and our ISQF. We observe our I(S)QF performing strongly on these datasets. The experiments are ran using the PyTorch version of the GluonTS (Alexandrov et al., 2020) package.

Dataset	Strategy	mean_wQL	crossing %	wQL[0.5]	wQL[0.7] (I)	wQL[0.9]	wQL[0.995] (E)	MSIS[0.1]	MSIS[0.02]
Elec	Gaussian	0.053 ± 0.006	0.0 ± 0.0	0.082 ± 0.008	0.076 ± 0.008	0.047 ± 0.006	0.010 ± 0.003	5.685 ± 0.430	10.254 ± 1.198
	SQF	0.054 ± 0.008	0.0 ± 0.0	0.084 ± 0.015	0.078 ± 0.011	0.046 ± 0.009	0.008 ± 0.004	6.424 ± 1.488	12.485 ± 3.211
	IQF	0.048 ± 0.005	0.0 ± 0.0	0.073 ± 0.006	0.067 ± 0.007	0.037 ± 0.004	0.015 ± 0.006	6.177 ± 0.607	15.887 ± 2.533
	ISQF	0.055 ± 0.010	0.0 ± 0.0	0.081 ± 0.013	0.079 ± 0.016	0.049 ± 0.013	0.012 ± 0.002	6.299 ± 0.656	14.413 ± 1.617
Traf	Gaussian	0.126 ± 0.004	0.0 ± 0.0	0.176 ± 0.004	0.179 ± 0.008	0.123 ± 0.005	0.025 ± 0.000	5.599 ± 0.244	10.870 ± 0.247
	SQF	0.111 ± 0.002	0.0 ± 0.0	0.157 ± 0.004	0.152 ± 0.003	0.105 ± 0.001	0.032 ± 0.002	4.889 ± 0.071	12.020 ± 0.353
	IQF	0.119 ± 0.014	0.0 ± 0.0	0.171 ± 0.026	0.164 ± 0.022	0.107 ± 0.008	0.034 ± 0.002	5.174 ± 0.443	12.569 ± 0.458
	ISQF	0.111 ± 0.003	0.0 ± 0.0	0.154 ± 0.005	0.150 ± 0.005	0.103 ± 0.003	0.035 ± 0.002	5.059 ± 0.225	12.608 ± 0.713

Table 5: Comparison of the accuracy for using DeepAR with Gaussian, SQF and our ISQF with 5 training quantile knots [0.01, 0.1, 0.5, 0.9, 0.99]. The mean and standard deviation are computed over 4 runs, and the winning method is shown in bold. (I) and (E) indicate the quantiles, where interpolation and extrapolation are performed, respectively. The MSIS[0.1] measures the 90% prediction interval using the interpolated 95th and 5th quantiles, if defined. Similarly, the MSIS[0.02] measures the 98% prediction interval using the 99th and 1st training quantiles.

Sample path generation Figure 6 visualizes the sample paths of applying the MQ-CNN model with different benchmark layers and our I(S)QF.

Learning Quantile Function without Quantile Crossing




Figure 6: Sample paths generated by using MQ-CNN with our I(S)QF and various benchmark layers. Three sample paths are generated for each of the 6 time series from the ELeC dataset. The dashed vertical lines represent the start of the prediction horizon.

H Distribution Recovery with IQF

H.1 Multimodal Distribution




Figure 7: Toy Multimodal distribution with 3 peaks: IQF interpolates and extrapolates around training knots to approximate a Multimodal distribution (c) without any distributional assumption, even more accurately as the number of knots increases from 5 ((a), (d)) to 20 ((b), (e)).

H.2 Cauchy Distribution




Figure 8: IQF fits a Cauchy distribution as a distribution-free approach with a various number of training knots.