



HAL
open science

Statistical Learning of Value-at-Risk and Expected Shortfall

D Barrera, S Crépey, E Gobet, Hoang-Dung Nguyen, B Saadeddine

► **To cite this version:**

D Barrera, S Crépey, E Gobet, Hoang-Dung Nguyen, B Saadeddine. Statistical Learning of Value-at-Risk and Expected Shortfall. 2024. hal-03775901v4

HAL Id: hal-03775901

<https://hal.science/hal-03775901v4>

Preprint submitted on 18 Sep 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.

Statistical Learning of Value-at-Risk and Expected Shortfall*

D. Barrera[†], S. Crépey[‡], E. Gobet[§], Hoang-Dung Nguyen[¶], B. Saadeddine^{||}

September 6, 2024

Abstract

We propose a non-asymptotic convergence analysis of a two-step approach to learn a conditional value-at-risk (VaR) and a conditional expected shortfall (ES) using Rademacher bounds, in a non-parametric setup allowing for heavy-tails on the financial loss. Our approach for the VaR is extended to the problem of learning at once multiple VaRs corresponding to different quantile levels. This results in efficient learning schemes based on neural network quantile and least-squares regressions. An a posteriori Monte Carlo procedure is introduced to estimate distances to the ground-truth VaR and ES. This is illustrated by numerical experiments in a Student- t toy model and a financial case study where the objective is to learn a dynamic initial margin.

Keywords: value-at-risk, expected shortfall, quantile regression, quantile crossings, neural networks, numerical finance.

AMS Subject Classification: 62G32, 62L20, 62M45, 91G60, 91G70.

* Python notebooks reproducing the results of this paper are available on <https://github.com/BouazzaSE/Learning-VaR-and-ES>.

[†]Email: j.barrerac@uniandes.edu.co. Departamento de Matemáticas, Universidad de los Andes, Cra 1 # 18a-12, Edificio H. Bogotá, Colombia. Postal code: 111711. The research of D. Barrera benefited from the support of the *Chair Capital Markets Tomorrow: Modeling and Computational Issues* under the aegis of the Institut Europlace de Finance, a joint initiative of Laboratoire de Probabilités, Statistique et Modélisation (LPSM) / Université Paris Cité and Crédit Agricole CIB, and from the support of the *Chair Stress Test, RISK Management and Financial Steering of the Foundation Ecole Polytechnique*.

[‡]Email: stephane.crepey@lpsm.paris. LPSM, Université Paris Cité, France. The research of S. Crépey benefited from the support of the *Chair Stress Test, RISK Management and Financial Steering*, led by the French Ecole polytechnique and its Foundation and sponsored by BNP Paribas.

[§]Email: emmanuel.gobet@polytechnique.edu. Centre de Mathématiques Appliquées (CMAP), CNRS, Ecole Polytechnique, Institut Polytechnique de Paris. Route de Saclay, 91128 Palaiseau Cedex, France. The research of E. Gobet is supported by the *Chair Stress Test, RISK Management and Financial Steering of the Foundation Ecole Polytechnique*.

[¶]Email: hdnguyen@lpsm.paris. PhD student, Université Paris Cité, France. The research of H.-D. Nguyen is funded by a CIFRE grant from Natixis.

^{||}Email: bouazza.saadeddine2@ca-cib.com. Quantitative research GMD, Credit Agricole CIB, Paris.

1 Introduction

Motivation. Quantile regression is a classical statistical problem that has received attention since the 1750s. According to Koenker (2017), the least absolute criterion (or pinball loss function) for the median even preceded the least squares for the mean (introduced by Legendre in 1805). Quantile regression is commonly performed in the context of linear models, where the ensuing minimization problem can be cast as a linear program solved by the simplex method. Alternative approaches include nonlinear quantile regression based on interior point methods (Koenker and Park, 1996) or non-parametric quantile regression implemented by stochastic gradient descent (Rodrigues and Pereira, 2020).

In harmony with the numerous financial applications, we refer to quantile as value-at-risk (VaR) and to superquantile (Rockafellar and Royset, 2013), i.e. the expected loss given the loss exceeds the VaR, as expected shortfall (ES, initially named as conditional value-at-risk in the literature, however this terminology makes it confusing when considering additional conditioning as we do). In this paper we learn conditional versions of VaR and ES, accounting for some available information in the conditioning, represented by a random variable X . Conditional VaR and ES appear naturally in various financial applications. Learning the VaR and ES provides a way to shortcut nested Monte Carlo simulations by regression (cf. Broadie, Du, and Moallemi (2015) for learning or regressing a conditional expectation). Our initial motivation was the simulation of dynamic initial margin or economic capital processes in the context of refined FVA or KVA computations as per Albanese, Crépey, Hoskinson, and Saadeddine (2021, Section A.4) or Abbas-Turki, Crépey, Li, and Saadeddine (2024, Section 5). As a second application, one may consider a forward-looking risk management exercise, whereby a bank samples several possible scenarios along which it must compute all its risk metrics in order to assess the amount of regulatory capital required to secure its activities. A third application could be related to stress testing exercises required by regulators, whereby the risk metrics need be evaluated along various stressed scenarios. Importantly, our approach is readily extendible, both in theory and practice, to the problem of learning multiple quantiles at the same time, furnishing in particular efficient approaches to the so called crossing quantiles problem (see Section 5.1 for a discussion of the related literature). A fourth application of this work thus consists in parameterizing the VaR and ES as a function of their risk level $\alpha \in (0, 1)$, so as to have an evaluation of risk metric functions conditionally on α (and other informative variables).

Our contributions in relation to the literature. Dimitriadis and Bayer (2019) introduced a joint linear regression estimator for VaR and ES based on their joint elicibility properties (Fissler and Ziegel, 2016; Fissler, Ziegel, and Gneiting, 2015), implemented numerically using the nonlinear simplex optimization algorithm. They developed an asymptotic convergence analysis for their estimators, establishing their consistency and asymptotic normality under somewhat strong semiparametric assumptions and regularity conditions. Instead, we propose a non-asymptotic convergence analysis of a learning algorithm for VaR and ES using a two-step approach in a non-parametric setup. We then specify our results to practical schemes for learning the

conditional VaR and ES using neural networks as the function approximators.

Padilla, Tansey, and Chen (2022) also consider quantile regression with ReLU networks, but only in the case of deterministic covariates X_i (in our notation). In a setup similar to our Assumption 2.2 (see Remark 2.3), they provide qualitative non-asymptotic estimates for such networks, of which our corresponding results can be considered quantitative versions (i.e. with explicit constants); they also establish minimax rates for quantile functions with Hölder-related regularity or improved rates under Besov regularity on the target function.

Assuming that the quantile function has a compositional structure in terms of Hölder-continuous functions, Shen, Jiao, Lin, Horowitz, and Huang (2021) derive non-asymptotic error bounds that depend only on the dimension of the composed functions (as opposed to the dimension d of the inputs usually in the literature), but a statistical error term requiring an order of integrability $p > 1$ of the response variable Y .

As opposed to the previous references who mainly follow Vapnik-Chervonenkis (VC) based approaches, we mainly follow a Rademacher based approach. Our value-at-risk error bounds do not require any finite moments beyond integrability of Y . In the neural network case, depending on the nature of the regularization that is used, our results may not depend on the dimension d of the inputs either. Also accounting for the values of the constants, fully explicit in our case, which appear in our error bounds, we show in the discussion following Theorem 3.3 how Rademacher-based bounds are better than VC bounds for data in the small to medium size regime that matters in finance, while VC only dominates Rademacher for really big data that may be relevant in other application fields. Notice that the two papers discussed above only consider value-at-risk (quantile), as opposed to expected shortfall also in our case. Since the first arXiv version of this work was published, other works appeared on the two-step approach for learning the ES, notably He, Tan, and Zhou (2023), sometimes with a non-asymptotic error analysis, but always restricted to a parametric or semiparametric linear setup.

Beyond theory, we contribute by several algorithmic tricks. In the context of machine learning on simulated data, which is very relevant for quantitative finance modulo validity, Proposition 3.4 underpins an a posteriori error estimation method in order to compute errors against ground-truth values of the conditional VaR and ES (even without access to the latter). In the neural net setup, our two-step methodology enables the reuse of the VaR neural network’s hidden layers in the training of the neural network approximating the ES, reducing the additional learning of the ES to a linear regression against the already learned regression basis for the VaR. The multi-quantile learning approaches of Section 5 (with the multi- α (III) approach of Section 5.4 often found the best alternative in our numerics) not only diminish the computational burden with respect to several single-quantile learnings that would be run separately, but are also found to better learn the value-at-risks of high confidence levels α .

Outline of the paper. Section 2 introduces the setup and our learning algorithm. Section 3 delivers the corresponding convergence analysis. Section 4 discusses specializations of this scheme and its errors to the case of inference via neural networks. We introduce multi-quantile extensions of the above in Section 5. Sections 6 and 7 discuss

numerical experiments. Section A reviews classical properties of the unconditional VaR and ES. Technical proofs are deferred to Section B.

2 A Learning Algorithm for VaR and ES

We denote by $(\Omega, \mathcal{A}, \mathbb{P})$ a fixed probability space, which admits all the random variables appearing below, with corresponding expectation operator denoted by $\mathbb{E}[\cdot]$. The notation $L^p_{\mathbb{P}}$ (with $p \in [1, +\infty]$) stands for the usual \mathbb{P} -integrability spaces and we denote by $\|\cdot\|_{\mathbb{P}, p}$ the related norms. For a general Polish space \mathcal{Z} with Borel sigma algebra $\mathcal{B}_{\mathcal{Z}}$, by a random element Z of \mathcal{Z} we mean an $\mathcal{A}/\mathcal{B}_{\mathcal{Z}}$ measurable function $Z : \Omega \rightarrow \mathcal{Z}$; \mathbb{P}_Z denotes the law induced by Z on \mathcal{Z} , i.e., for every $A \in \mathcal{B}_{\mathcal{Z}}$, $\mathbb{P}_Z(A) = \mathbb{P}[Z \in A]$.

From now on, \mathcal{X} denotes a fixed Polish space and

$$(X, Y) : \Omega \rightarrow \mathcal{X} \times \mathbb{R} \text{ is a fixed random element of } \mathcal{X} \times \mathbb{R} \text{ with } Y \in L^1_{\mathbb{P}}.$$

We fix a conditional distribution function $\mu : \mathcal{X} \times \mathcal{B}_{\mathbb{R}} \rightarrow [0, 1]$ of Y given X (Kallenberg, 2006, Theorem 5.3 p.84) and we assume that the function $\mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}$ defined by $(x, y) \mapsto \mu(x, (-\infty, y])$ is $(\mathcal{B}_{\mathcal{X}} \times \mathcal{B}_{\mathbb{R}})/\mathcal{B}_{\mathbb{R}}$ (i.e. Borel) measurable (as for instance if $\mathcal{X} = \mathbb{R}^d$ and (X, Y) admits a density with respect to Lebesgue measure). We use the corresponding version $\mathbb{P}[Y \in \cdot | X] = \mu(X, \cdot)$ of the conditional probability of Y given X and the conditional cdf of Y given X , $F_{Y|X}(y) = \mathbb{P}[Y \leq y | X] = \mu(X, (-\infty, y])$. We assume, without loss of generality, that $F_{Y|X(\omega)}(\cdot)$ is integrable for every $\omega \in \Omega$ (since $Y \in L^1_{\mathbb{P}}$, we have that $\infty > \mathbb{E}[|Y|] = \mathbb{E}[\mathbb{E}[|Y| | X]] = \mathbb{E}[\int_{\mathbb{R}} |y| F_{Y|X}(dy)]$, thus $F_{Y|X(\omega)}$ is integrable for \mathbb{P} a.e. ω : it suffices to change the version of X to guarantee integrability for every ω). In particular, if $g : \mathbb{R} \rightarrow \mathbb{R}$ is such that $g(Y)$ is \mathbb{P} integrable, then

$$\mathbb{E}[g(Y)|X](\omega) = \int_{\mathbb{R}} g(y) F_{Y|X(\omega)}(dy), \quad \mathbb{P} \text{ a.s.} \quad (2.1)$$

For a function $F : \mathbb{R} \rightarrow \mathbb{R}$ and $q \in \mathbb{R}$, $F(q-) := \lim_{z \uparrow q} F(z)$.

Definition 2.1. The conditional value-at-risk (VaR) and expected shortfall (ES) of Y given X at the confidence level $\alpha \in (0, 1)$ are (cf. (A.2))

$$\begin{aligned} \text{VaR}(Y|X) &= \text{VaR}(F_{Y|X}) = \min_{F_{Y|X}^{-1}([\alpha, 1])} = \min\{y \in \mathbb{R} : F_{Y|X}(y) \geq \alpha\}, \\ \text{ES}(Y|X) &= \frac{1}{1 - F_{Y|X}(\text{VaR}(Y|X)-)} \int_{[\text{VaR}(Y|X), \infty)} y F_{Y|X}(dy). \end{aligned}$$

Lemma 2.1. *There exist Borel measurable functions $q : \mathcal{X} \rightarrow \mathbb{R}$ and $s : \mathcal{X} \rightarrow \mathbb{R}$ such that*

$$q(X) = \text{VaR}(Y|X), \quad s(X) = \text{ES}(Y|X), \quad \mathbb{P} \text{ a.s.} \quad (2.2)$$

Proof. See Section B.1. \square

Notation 2.1. Besides (2.2), we also introduce $r = s - q$, which is nonnegative, \mathbb{P}_X a.s..

Lemma 2.2. (i) *With probability one, we have*

$$-\alpha^{-1}\mathbf{E}[Y^-|X] \leq q(X) \leq (1-\alpha)^{-1}\mathbf{E}[Y^+|X], \quad (2.3)$$

$$|s(X)| \leq (1-\alpha)^{-1}\mathbf{E}[|Y||X]. \quad (2.4)$$

(ii) *The (resp. square) integrability of $q(X)$ and $s(X)$ follows from that of Y .*

Proof. See Section B.2. \square

2.1 VaR and ES as Minimizers

For our VaR error control, we assume the knowledge of a tube of bounded width containing the graph of the target function q . The more a priori knowledge one has on q , i.e. the narrower this tube, the better will be the statistical error of our a priori estimates (see e.g. the second line in (3.8)).

Assumption 2.2. There exist functions $a, b : \mathcal{X} \rightarrow \mathbb{R}$ such that (i) the law $F_{Y|X}(\cdot)$ has a density $\dot{F}_{Y|X}(\cdot)$ on $[a(X), b(X)]$, P a.s., (ii)

$$F_{Y|X}(a(X)) < \alpha < F_{Y|X}(b(X)) \text{ holds P a.s.}, \quad (2.5)$$

and (iii) $0 < \|b - a\|_{P_X, \infty} < \infty$.

Remark 2.2. The absolute continuity of $F_{Y|X}$ around the quantile of level α postulated in Assumption 2.2(i) is a common assumption in the literature, see e.g. Reiss (2012, Chapter 3). It can be ensured (with a quantifiable impact on the quantiles) by adding a small independent Gaussian noise to Y . Assumptions 2.2(i)-(ii) imply, in particular, that $F_{Y|X}(q(X)) = \alpha$ and $a(X) < q(X) < b(X)$ P a.s. hold, but the converse is not true as there can be α -quantile functions b greater than the left-quantile (VaR) q . Assumption 2.2(iii) holds in particular if Y is bounded or if \mathcal{X} is compact and q can be chosen continuous, case in which q is also bounded, but it may also hold when neither X nor Y (or Y given X) nor q are bounded. The \mathcal{X} compact and q continuous case may seem restrictive but even this very special sub-case of Assumption 2.2 covers most use-cases of interest, as data practitioners typically restrict their learnings to compact sets of the state space, even if this means truncating it and doing several learnings on different subdomains. Also note that one can always reduce the problem to the case of bounded Y by applying a bijective increasing, hence VaR preserving (modulo an application of the inverse bijection to the estimated quantile function), bounded transformation to the data Y —with the caveat that applying such transformation and its inverse may not be innocuous numerically. In any case, the key point, which we see as a significant contribution of this work, is that Assumption 2.2 makes mainly (even though implicitly) assumptions on X , and no tail assumptions on Y or $Y|X$ (beyond the integrability and square integrability of Y that are postulated in the VaR and ES respective parts of the paper). This is important for applications in finance as Cont, Cucuringu, Xu, and Zhang (2023) demonstrate that heavy tail losses frequently arise from trading strategies, even in Gaussian models.

Remark 2.3. Padilla et al. (2022, Assumption 2) can be seen as a slightly stronger version of Assumption 2.2, in the sense that the width of their guarantee tube on q does not depend on x and they assume positive lower and upper bounds on the density $\dot{F}_{Y|X}(\cdot)$ along their tube. Shen et al. (2021) even assumes q (f_0 in their notation) bounded throughout the paper.

Notation 2.4. Under Assumption 2.2, $Y^{a,b}(\omega) := a(X(\omega)) \vee Y(\omega) \wedge b(X(\omega))$, $\omega \in \Omega$, and $\mathcal{L}_{a,b}$ denotes the set of functions f on \mathcal{X} such that $f(x) \in [a(x), b(x)]$ (resp. $f(x) \geq 0$) holds for all x ; \mathcal{L}_+^2 denotes the set of P_X square integrable functions f on \mathcal{X} such that $f(x) \in [0, +\infty)$ holds for all x .

Lemma 2.3. *Under Assumption 2.2, the (resp. square) integrability of Y implies that any function $f \in \mathcal{L}_{a,b}$ is P_X (resp. square) integrable.*

Proof. Assumption 2.2(i)-(ii) implies that

$$\begin{aligned} q(X) < b(X) &= b(X) - a(X) + a(X) < b(X) - a(X) + q(X) \\ q(X) > a(X) &= a(X) - b(X) + b(X) > a(X) - b(X) + q(X) \end{aligned} \quad (2.6)$$

hold P a.s.. By Lemma 2.2, the (resp. square) integrability of Y implies the (resp. square) integrability of q . Combining this with the bounded $b - a$ in Assumption 2.2(iii), (2.6) yields the (resp. square) integrability of a and b , which implies the one of any $f \in \mathcal{L}_{a,b}$. \square

Definition 2.3. Given $\alpha \in (0, 1)$, the respective pointwise loss functions for VaR and ES–VaR given the hypothesis u for VaR, both at the level α , are given respectively by

$$\begin{aligned} \text{the pinball loss, i.e. } \mathbb{R}^2 \ni (y, u) &\mapsto \phi(y, u) = (1 - \alpha)^{-1}(y - u)^+ + u \in \mathbb{R}, \text{ resp.} \\ \mathbb{R}^2 \times [0, +\infty) \ni (y, u, v) &\mapsto ((1 - \alpha)^{-1}(y - u)^+ - v)^2 = (\phi(y, u) - u - v)^2 \in [0, +\infty). \end{aligned} \quad (2.7)$$

By application of Lemma 2.2(ii), the P integrability of Y also implies that of $\phi(Y, q(X))$; if $Y \in L_P^2$, then $(\phi(Y, q(X)) - q(X) - r(X))^2$ is P integrable. The loss functions $\phi(y, u)$ and $(\phi(y, u) - u - v)^2$ underlie the following representations of the functions q and $r = s - q$ in Lemma 2.1 and Notation 2.1 as solutions to optimization problems.

Lemma 2.4. *Under Assumption 2.2, we have*

$$q \in \underset{f \in \mathcal{L}_{a,b}}{\operatorname{argmin}} \mathbb{E} [\phi(Y, f(X))] = \underset{f \in \mathcal{L}_{a,b}}{\operatorname{argmin}} \mathbb{E} [\phi(Y^{a,b}, f(X))]. \quad (2.8)$$

Moreover, any function belonging to either argmin in (2.8) is an α -quantile function of Y or, equivalently, of $Y^{a,b}$. In addition, for any function q' in either argmin in (2.8):

$$s(X) = q'(X) + (1 - \alpha)^{-1} \mathbb{E} [(Y - q'(X))^+ | X] \text{ holds } P \text{ a.s.}; \quad (2.9)$$

if $Y \in L_P^2$ and $q' \in \mathcal{L}_{a,b}$, then

$$s - q' \in \underset{g \in \mathcal{L}_+^2}{\operatorname{argmin}} \mathbb{E} [(\phi(Y, q'(X)) - q'(X) - g(X))^2] \quad (2.10)$$

and, for any function r' belonging to the argmin in (2.10),

$$s(X) = q'(X) + r'(X) \text{ holds } P \text{ a.s..}$$

Proof. Using (A.4) and (2.1), for every $f \in \mathcal{L}_{a,b} \subset L^1_{\mathbb{P}_X}$ (by Lemma 2.3), we obtain $\mathbb{E}[\phi(Y, q(X))|X] \leq \mathbb{E}[\phi(Y, f(X))|X]$ \mathbb{P} a.s., hence $\mathbb{E}[\phi(Y, q(X))] \leq \mathbb{E}[\phi(Y, f(X))]$. This implies that $q \in \operatorname{argmin}_{f \in \mathcal{L}_{a,b}} \mathbb{E}[\phi(Y, f(X))]$. Moreover, by (2.5), the conditional α -quantiles of $Y^{a,b}$ and of Y are the same; in particular

$$q(X) = \operatorname{VaR}(Y^{a,b}|X), \quad \mathbb{P} \text{ a.s.},$$

and the second equality of (A.4) yields (2.8).

Conversely, if $g \in \mathcal{L}_{a,b}$ is not an α -quantile function of Y , i.e. if $F_{Y|X}(g(X))$ differs from α on a set of positive \mathbb{P} probability, then, according to (A.4) and (2.1), the random variable $Z = \mathbb{E}[\phi(Y, g(X))|X] - \mathbb{E}[\phi(Y, q(X))|X]$ satisfies

$$\mathbb{P}[Z \geq 0] = 1 \text{ and } \mathbb{P}[Z > 0] > 0,$$

whence

$$\mathbb{E}[\phi(Y, g(X))] - \mathbb{E}[\phi(Y, q(X))] = \mathbb{E}[Z] > 0,$$

showing that g does not belong to $\operatorname{argmin}_{f \in \mathcal{L}_{a,b}} \mathbb{E}[\phi(Y, f(X))]$. Likewise,

$$\operatorname{argmin}_{f \in \mathcal{L}_{a,b}} \mathbb{E}[\phi(Y^{a,b}, f(X))]$$

is included in the set of the α -quantile functions of $Y^{a,b}$ or, equivalently, of Y .

Finally, (A.5) and (2.1) yield (2.9), from which (2.10) follows by the representation of the conditional expectation as an L^2 projection in the square integrable case. \square

Remark 2.5. The minimizers in (2.8) do not need to be unique: any function $q' \in \mathcal{L}_{a,b}$ satisfying $F_{Y|X}(q'(X)) = \alpha$ \mathbb{P} a.s. is a minimizer of $f \mapsto \mathbb{E}[\phi(Y, f(X))]$ and there are infinitely many such functions q' (and random variables $q'(X)$) in the “degenerate case” where $F_{Y|X}^{-1}(\alpha)$ is an interval of positive length on a set of positive \mathbb{P}_X measure.

Remark 2.6. There exist whole families, generalizing (2.8) and (2.10), of representations of the functions q and $r = s - q$ (or q and s) as minimizers of suitable functionals, including joint (but non globally convex) representations of the pair (q, s) based on the joint elicibility properties of value-at-risk and expected shortfall (Fissler and Ziegel, 2016; Fissler et al., 2015): see Theorem 2.3 in the arXiv v1 version of this work. Using such a joint representation of (q, s) was actually our original choice for practical computations in Albanese et al. (2021), in the footsteps of Dimitriadis and Bayer (2019). Deriving the VaR and ES at once indeed looks an attractive idea, but after more empirical investigation reported in the paper’s GitHub the best turned out to be the simplest, i.e. the two-step algorithm that first produces an approximation of the (conditional) VaR and then learns the ES by least-squares using the VaR approximation.

The joint approach in fact suffers from two different scales present in the same problem, to the effect that either the VaR or the ES is badly handled. Moreover, the mathematical analysis of the joint approach poses difficulties of its own, even under locally convex parametrizations such as the ones considered (for tail risk scenario generation) in Cont et al. (2023); also note that (local or global) convexity at the functional

level does not necessarily imply convexity with respect to the weights of a neural net approximator.

In view of these practical and theoretical considerations, we focus on the two-step algorithm hereafter. On this two-step approach (in a semiparametric linear setup), see also He, Tan, and Zhou (2023).

The two-step learning scheme for VaR and ES. The functional representations (2.8) and (2.10) give rise to associated approximation algorithms for q and s . The numerical recipe is just to replace the minimization problems in (2.8) or (2.10) by empirical versions: instead of $\mathcal{L}_{a,b}$ and \mathcal{L}_+^2 , we use convenient hypotheses spaces (families of functions represented by neural nets in our numerics) inside the previous ones; instead of \mathbb{P} expectation, we use a Monte Carlo approximation based on i.i.d. samples $(X, Y)_{1:n} = ((X_1, Y_1), \dots, (X_n, Y_n))$ of (X, Y) in $\mathcal{X} \times \mathbb{R}$, with (X, Y) independent of the sample.

3 Convergence Analysis

The a priori error analysis of Sections 3.1–3.4 corresponds essentially to the above scheme based on the data $(X, Y)_{1:n}$, ignoring the numerical optimization error (i.e. assuming one has access to global minimizers of the empirical risk functions). A more practical (and exhaustive, but only a posteriori) error control is then provided in Section 3.5.

3.1 Estimation of VaR: General Setting

For a possibly data dependent $f \in \mathcal{L}_{a,b} \subset L_{\mathbb{P}_X}^1$ (by Lemma 2.3), i.e. a function of $\mathcal{L}_{a,b}$ possibly parameterized by the data $(X, Y)_{1:n}$, we define

$$\Phi(f) = \int_{\mathcal{X} \times \mathbb{R}} \phi\left(a(x) \vee y \wedge b(x), f(x)\right) \mathbb{P}_{X,Y}(dx, dy) = \mathbb{E} \left[\phi\left(Y^{a,b}, f(X)\right) \middle| (X, Y)_{1:n} \right],$$

(cf. Notation 2.4). We also fix a hypothesis space $\mathcal{F} \subseteq \mathcal{L}_{a,b}$ ensuring the existence of

$$\tilde{q} \in \underset{f \in \mathcal{F}}{\operatorname{argmin}} \Phi(f) \text{ and } \hat{q} \in \underset{f \in \mathcal{F}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \phi\left(Y_i^{a,b}, f(X_i)\right), \quad (3.1)$$

dubbed best mean and best empirical hypothesis for VaR within \mathcal{F} .

3.2 A Priori Error Bounds for the Estimation of VaR

Under Assumption 2.2, q as in (2.2) defines an α -quantile function (the smallest one for \mathbb{P}_X a.e. x). For such q , define

$$\gamma_f(x) = \Gamma_{F_{Y|X=x}}(f(x), q(x)) \quad (3.2)$$

via (A.7), so that $\gamma_f(x)$ multiplied by $2(1-\alpha)$ coincides with $\dot{F}_{Y|X=x}$ evaluated at some middle point between $f(x)$ and $q(x)$ (see (A.11)). The following lemma interprets the approximation error $\Phi(\tilde{q}) - \Phi(q)$ as a weighted distance between \tilde{q} and q in the L^2

norm; by (3.3), with P_X probability one, the weights $\gamma_f(x)$ are equal to zero at some x only if $f(x)$ is already an α -quantile of $F_{Y|X=x}$. Note that the inequality in (3.4) is similar to Lemma 3 in Shen et al. (2021).

Lemma 3.1. *Under Assumption 2.2, with γ_f as in (3.2),*

$$P\left(\{\gamma_f(X) = 0\} \setminus \{F_{Y|X}(f(X)) = \alpha\}\right) = 0 \quad (3.3)$$

holds for any $f \in \mathcal{L}_{a,b}$. Moreover, for \tilde{q} defined by (3.1),

$$\Phi(\tilde{q}) - \Phi(q) = \|(\tilde{q} - q)\gamma_{\tilde{q}}\|_{P_{X,2}}^2 = \inf_{f \in \mathcal{F}} \|(f - q)\gamma_f\|_{P_{X,2}}^2 \leq \frac{2 - \alpha}{1 - \alpha} \inf_{f \in \mathcal{F}} \|f - q\|_{P_{X,1}}. \quad (3.4)$$

Proof. See Section B.3. \square

Example 3.1. Assume that

$$0 < c \leq \gamma_{\tilde{q}}(X) \quad (3.5)$$

holds P a.s. for some constant c (according to (A.10) this is true for instance if

$$c \leq \frac{1}{2(1 - \alpha)} \dot{F}_{Y|X}(tf(X) + (1 - t)q(X)) \leq C \quad (3.6)$$

holds P a.s. for every $t \in (0, 1)$ and $f \in \mathcal{F}$, for some positive constants $c \leq C$). Suppose additionally that $\inf_{f \in \mathcal{F}} \|f - q\|_{P_{X,1}} < \delta$ holds for some positive constant δ . Then (3.4) yields

$$\sqrt{c} \|\tilde{q} - q\|_{P_{X,2}} \leq \sqrt{\frac{2 - \alpha}{1 - \alpha}} \delta. \quad (3.7)$$

Example 3.2. Assume that (i) (3.5) holds, (ii) (2.5) holds for some constants $a(X) \equiv A$ and $b(X) \equiv B$, i.e. $q(X) \in (A, B)$, P a.s., and (iii) there exists (e.g. because \mathcal{X} is compact and q Hölder continuous) an enumerable partition $\{\mathcal{X}_j\}_j$ of measurable subsets of \mathcal{X} such that $\sup_j \sup_{(x,x') \in \mathcal{X}_j \times \mathcal{X}_j} |q(x) - q(x')| < \delta$. Then (3.7) holds for the following hypothesis space \mathcal{F} (over which \tilde{q} minimizes Φ , cf. (3.2))

$$\mathcal{F} = \left\{ x \mapsto \sum_j c_j \mathbf{1}_{\mathcal{X}_j}(x); c_j \in [A, B], \forall j \right\}.$$

We now give an upper bound for an error in probability associated with the empirical estimator \hat{q} of q . For this, we need to introduce the following measures of complexity applicable to our hypothesis spaces.

Definition 3.1. Let \mathcal{Z} be a Polish space and \mathcal{H} be a set of measurable real valued functions on \mathcal{Z} . For any random sequence $Z_{1:n}$ in \mathcal{Z} , the empirical Rademacher complexity $\mathcal{R}_{emp}(\mathcal{H}, Z_{1:n})$ and the Rademacher complexity $\mathcal{R}_{ave}(\mathcal{H}, Z_{1:n})$ of \mathcal{H} at $Z_{1:n}$ are defined as

$$\mathcal{R}_{emp}(\mathcal{H}, Z_{1:n}) = E \left[\frac{1}{n} \sup_{h \in \mathcal{H}} \sum_{k=1}^n U_k h(Z_k) \middle| Z_{1:n} \right], \quad \mathcal{R}_{ave}(\mathcal{H}, Z_{1:n}) = E [\mathcal{R}_{emp}(\mathcal{H}, Z_{1:n})],$$

where $U_{1:n}$ is an i.i.d. Rademacher sequence $P[U_k = 1] = P[U_k = -1] = 1/2$ independent of $Z_{1:n}$.

Theorem 3.2. Under the assumptions of Lemma 3.1, for any $\delta \in (0, 1)$,

$$(1 - \alpha)^{1/2} \|(\widehat{q} - q)\gamma_{\widehat{q}}\|_{\mathbb{P}_{X,2}} \leq (1 - \alpha)^{1/2} \inf_{f \in \mathcal{F}} \|(f - q)\gamma_f\|_{\mathbb{P}_{X,2}} + \left(2(2 - \alpha)\mathcal{R}_{ave}(\mathcal{F}, X_{1:n}) + \|b - a\|_{\mathbb{P}_{X,\infty}} \sqrt{\frac{2 \log(2/\delta)}{n}} \right)^{1/2} \quad (3.8)$$

holds with probability at least $1 - \delta$.

Proof. See Section B.4. \square

Remark 3.3. From Theorem 3.2, if there exist some constants c, C such that

$$0 < c \leq \gamma_f(X) \leq C < \infty \text{ holds P a.s.,}$$

for every $f \in \mathcal{F}$ (e.g. if (3.6) holds), then it can be deduced from (3.8) and (3.4) that

$$c(1 - \alpha)^{1/2} \|\widehat{q} - q\|_{\mathbb{P}_{X,2}} \leq \left((2 - \alpha) \inf_{f \in \mathcal{F}} \|f - q\|_{\mathbb{P}_{X,1}} \right)^{1/2} \wedge \left(C(1 - \alpha)^{1/2} \inf_{f \in \mathcal{F}} \|f - q\|_{\mathbb{P}_{X,2}} \right) + \left(2(2 - \alpha)\mathcal{R}_{ave}(\mathcal{F}, X_{1:n}) + \|b - a\|_{\mathbb{P}_{X,\infty}} \sqrt{\frac{2 \log(2/\delta)}{n}} \right)^{1/2} \quad (3.9)$$

holds with probability at least $1 - \delta$. A better a priori knowledge on q , i.e. a smaller $\|b - a\|_{\mathbb{P}_{X,\infty}}$ (see above Assumption 2.2), also results into larger c and smaller C , hence better constants in (3.9), but always with the same rate $n^{-\frac{1}{4}}$ for the statistical error assuming $\mathcal{R}_{ave}(\mathcal{F}, X_{1:n}) = O(n^{-\frac{1}{2}})$ as satisfied in our neural network application below.

Remark 3.4. In the proof of Lemma 2.3, $\mathcal{L}_{a,b} \subset L^1_{\mathbb{P}_X}$ follows from Assumption 2.2 and the integrability of q . Consequently, all the above results in this section remain valid even if the integrability assumption on Y is relaxed to that of q .

Remark 3.5. For $\varepsilon \geq 0$ and \mathcal{Z}, \mathcal{H} and $Z_{1:n}$ as in Definition 3.1, the ε covering number of \mathcal{H} with respect to the empirical L^1 norm at $Z_{1:n}$ is

$$\mathcal{N}_1(\mathcal{H}, Z_{1:n}, \varepsilon) = \min \left\{ m \in \mathbb{N} : \exists g_{1:m} \in (\mathcal{L}(\mathcal{Z}))^m : \sup_{h \in \mathcal{H}} \min_{l \in 1..m} \frac{1}{n} \sum_{k=1}^n |h(Z_k) - g_l(Z_k)| < \varepsilon \right\}$$

(with the convention $\inf \emptyset = \infty$). The quantity $\log(\mathcal{N}_1(\mathcal{H}, Z_{1:n}, \varepsilon))$ is called the L^1 ε -entropy of \mathcal{H} at $Z_{1:n}$. The interplay between the entropy and the Rademacher complexity is explained in Massart's Lemma B.2, which illuminates the usefulness of having upper bounds on the entropy for applications of the bounds above.

3.3 Estimation of ES-VaR: General Setting

Note that two random variables whose cdf differ only beyond their (assumed common) VaR can have arbitrarily far away ES. The truncated loss $\psi^B(y, u, v) := \left(B \wedge (\phi(y, u) - u) - v \right)^2$ for some positive constant B (compare with the second line in (2.7))

is therefore introduced in view of establishing ES related concentrations. We assume Y square integrable, so that $\mathcal{L}_{a,b} \subset L^2_{\mathbb{P}_X}$, by Lemma 2.3. We fix a hypothesis space \mathcal{G} of $[0, B]$ valued $\mathcal{B}_X/\mathcal{B}_{\mathbb{R}}$ measurable, square integrable functions and, given $f \in \mathcal{L}_{a,b}$, we fix some function

$$\hat{r}_f \in \operatorname{argmin}_{g \in \mathcal{G}} \frac{1}{n} \sum_{i=1}^n \psi^B(Y_i, f(X_i), g(X_i)), \quad (3.10)$$

assumed to exist, interpreted as the best empirical hypothesis for ES – VaR (after truncation by B) within \mathcal{G} , conditioned to the hypothesis f for VaR.

3.4 A Priori Error Bounds for the Estimation of ES–VaR

We are now in a position to establish our a priori bound on the error of the theoretical estimate \hat{r}_f with respect to r given an hypothesis f for q .

Theorem 3.3. *In the above setting, given $\delta \in (0, 1)$, the inequality*

$$\begin{aligned} \|\hat{r}_f - r\|_{\mathbb{P}_X, 2} &\leq \inf_{g \in \mathcal{G}} \|g - r\|_{\mathbb{P}_X, 2} \\ &\quad + 2 \left((1 - \alpha)^{-1} \|f - q\|_{\mathbb{P}_X, 2} + \|((1 - \alpha)^{-1}(Y - q(X))^+ - B)^+\|_{\mathbb{P}, 2} \right) \\ &\quad + B \left(\frac{4\mathcal{R}_{ave}(\mathcal{G}, X_{1:n})}{B} + \sqrt{\frac{2 \log(2/\delta)}{n}} \right)^{1/2} \end{aligned} \quad (3.11)$$

holds with probability at least $1 - \delta$.

Proof. See Section B.6. \square

Rademacher Versus Vapnik-Chervonenkis (VC) Regimes. An application of Massart’s Lemma B.2 for $\varepsilon = B\sqrt{n}$ to (3.11) yields the upper bound

$$\begin{aligned} \|\hat{r}_f - r\|_{\mathbb{P}_X, 2} &\leq \inf_{g \in \mathcal{G}} \|g - r\|_{\mathbb{P}_X, 2} \\ &\quad + 2 \left((1 - \alpha)^{-1} \|f - q\|_{\mathbb{P}_X, 2} + \|((1 - \alpha)^{-1}(Y - q(X))^+ - B)^+\|_{\mathbb{P}, 2} \right) \\ &\quad + \frac{B}{n^{1/4}} \left(4 \left(1 + \mathbb{E} \left[\sqrt{2 \log(\mathcal{N}_1(\mathcal{G}, X_{1:n}, B/\sqrt{n}))} \right] \right) + \sqrt{2 \log(2/\delta)} \right)^{1/2} \end{aligned} \quad (3.12)$$

with probability at least $1 - \delta$.

According to well-known facts, both (3.11) and (3.12) show a statistical error (the third line in both displays) with a suboptimal rate for the case at hand, namely $O(n^{-1/4})$ instead of $O(n^{-1/2})$. An upper bound with “the right rate” $O(n^{-1/2})$ can be achieved via a separate Vapnik-Chervonenkis (VC) analysis that can be found in the arXiv v1 version of this work (based upon Barrera and Gobet (2019, Theorem 2.2), itself refining Györfi, Kohler, Krzyżak, and Walk (2002, Theorem 11.4 p.201)),

according to which the inequality

$$\begin{aligned}
\|\widehat{r}_f - r\|_{\mathbb{P}_{X,2}} &\leq \sqrt{(6\lambda - 5)} \inf_{g \in \mathcal{G}} \|g - r\|_{\mathbb{P}_{X,2}} \\
&+ (1 + \sqrt{(6\lambda - 5)})((1 - \alpha)^{-1} \|f - q\|_{\mathbb{P}_{X,2}} + \|((1 - \alpha)^{-1}(Y - q(X))^+ - B)^+\|_{\mathbb{P},2}) \\
&+ \frac{2^3 3^{1/2} B}{\sqrt{(\lambda - 1)n}} \left(2 \log(21 \mathbb{E}[\mathcal{N}_1(\mathcal{G}, X_{1:n}, B/(24n))]) + 2 \log(2/\delta) \right)^{1/2},
\end{aligned} \tag{3.13}$$

holds with probability at least $1 - \delta$ and for every $1 < \lambda \leq 13/12$ (see also Barrera and Gobet (2019, Remark 3.2)).

But the size of the constants involved in (3.13) makes it less accurate than (3.12) in a relatively wide range of sample sizes. To give a crude comparison between the bounds (3.12) and (3.13), first note that, since $\sqrt{6\lambda - 5} \approx 1$, it is reasonable to limit the discussion to a comparison between the terms in the third line of the inequalities (3.12) and (3.13). If we focus on the dependence of the error bounds w.r.t. the data size n , and if we neglect the complexity terms by considering $\mathcal{N}_1(\dots) = 1$, a consideration of the usual case $\delta = 1/20$ and the choice $\lambda - 1 = 1/12$ shows that the ratio between the two terms in the third line of the inequalities (3.12) and (3.13) is (crudely) approximated by

$$\frac{\left(\frac{1}{\sqrt{n}} \left(4 + \sqrt{2 \log(2/\delta)} \right) \right)^{1/2}}{2^3 3^{1/2} \left(\frac{12}{n} \left[2 \log(21) + 2 \log(2/\delta) \right] \right)^{1/2}} \stackrel{\delta=1/20}{\approx} 1.47 \times 10^{-2} n^{1/4}. \tag{3.14}$$

This shows that the Rademacher-approach inequality (3.11) is better than the VC-approach inequality (3.13), namely the right-hand side of (3.14) is smaller than 1, whenever

$$n \gtrsim 2.14 \times 10^7, \tag{3.15}$$

which can be understood as a heuristic boundary between “small-medium” and “big” data, where we pass from the Rademacher to the VC regime.

Remark 3.6. The bounds (3.11) as well as (3.12) or (3.13) above can be combined with the bound (3.9) for giving rise to bounds on the estimate $\widehat{r}_{\widehat{q}}$ of r based on the estimation \widehat{q} of q if $\widehat{q}, q \in \mathcal{L}_{a,b}$, e.g.

$$\begin{aligned}
\|\widehat{r}_f - r\|_{\mathbb{P}_{X,2}} &\leq \inf_{g \in \mathcal{G}} \|g - r\|_{\mathbb{P}_{X,2}} + 2 \|((1 - \alpha)^{-1}(Y - q(X))^+ - B)^+\|_{\mathbb{P},2} \\
&+ \frac{2}{c(1 - \alpha)^{3/2}} \left[\left((2 - \alpha) \inf_{f \in \mathcal{F}} \|f - q\|_{\mathbb{P}_{X,1}} \right)^{1/2} \wedge \left(C(1 - \alpha)^{1/2} \inf_{f \in \mathcal{F}} \|(f - q)\|_{\mathbb{P}_{X,2}} \right) \right. \\
&\quad \left. + \left(2(2 - \alpha) \mathcal{R}_{ave}(\mathcal{F}, X_{1:n}) + \|b - a\|_{\mathbb{P}_{X,\infty}} \sqrt{\frac{2 \log(2/\delta)}{n}} \right)^{1/2} \right] \\
&+ \left(4B \mathcal{R}_{ave}(\mathcal{G}, X_{1:n}) + B^2 \sqrt{\frac{2 \log(2/\delta)}{n}} \right)^{1/2}.
\end{aligned}$$

3.5 A Posteriori Monte Carlo Validation of VaR and ES learners

The previous a priori error analysis is specific to the estimates \hat{q} and $\hat{s}_{\hat{q}} = \hat{q} + \hat{r}_{\hat{q}}$ of q and s . It is also theoretical in the sense that it supposes that global minimization is reached through training.

By contrast, the following a posteriori error analysis can be applied to any tentative approximations $f(X)$ of $q(X) = \text{VaR}(Y|X)$ and $h(X)$ of $s(X) = \text{ES}(Y|X)$ at the confidence level α , including the ones that will arise from the practical training schemes of Sections 4.2 and 5.3-5.4, namely numerical stochastic gradient descent (SGD) approximations of \hat{q} and $\hat{q} + \hat{r}_{\hat{q}}$. In fact, assuming one has access to the data generating process, as it is the case in most quantitative finance problems, one can estimate distances of any guesses to the groundtruth (conditional) VaR and ES without directly computing the latter, using an companion out-of-sample (dubbed “twin” in reference to $Y^{(1)}$ and $Y^{(2)}$ below) Monte Carlo procedure relying on the following result.

Proposition 3.4. *Let $Y^{(1)}$ and $Y^{(2)}$ denote two conditionally independent copies of Y given X , meaning that for any bounded, $\mathcal{B}_{\mathbb{R}}$ measurable functions ρ and ϱ ,*

$$\mathbb{E} \left[\rho(Y^{(j)}) | X \right] = \mathbb{E} [\rho(Y) | X], \quad \mathbb{E} \left[\rho(Y^{(1)}) \varrho(Y^{(2)}) | X \right] = \mathbb{E} \left[\rho(Y^{(1)}) | X \right] \mathbb{E} \left[\varrho(Y^{(2)}) | X \right].$$

(i) *For any $\mathcal{B}_X / \mathcal{B}_{\mathbb{R}}$ measurable function f ,*

$$\begin{aligned} \|\mathbb{P}[Y > f(X) | X] - 1 + \alpha\|_{\mathbb{P},2}^2 &= (1 - \alpha) \times \\ &\left(1 - \alpha - \left(\mathbb{P}[Y^{(1)} > f(X)] + \mathbb{P}[Y^{(2)} > f(X)] \right) \right) + \mathbb{P}[Y^{(1)} \wedge Y^{(2)} > f(X)]. \end{aligned} \quad (3.16)$$

(ii) *For any f and h in $\mathcal{L}_{a,b}$ under Assumption 2.2, if Y is square integrable and if, for some positive constant c , $\inf_{y \in (a(x), b(x))} \tilde{F}_{Y|X=x}(y) \geq c$ holds \mathbb{P}_X a.s., then*

$$\begin{aligned} \|h(X) - s(X)\|_{\mathbb{P},2} &\leq \left\| h(X) - f(X) - \mathbb{E} \left[(1 - \alpha)^{-1} (Y - f(X))^+ | X \right] \right\|_{\mathbb{P},2} \\ &+ \frac{2 - \alpha}{c(1 - \alpha)} \|\mathbb{P}[Y > f(X) | X] - 1 + \alpha\|_{\mathbb{P},2}, \end{aligned} \quad (3.17)$$

where

$$\begin{aligned} \left\| h(X) - f(X) - \mathbb{E} \left[(1 - \alpha)^{-1} (Y - f(X))^+ | X \right] \right\|_{\mathbb{P},2}^2 &= \|h(X) - f(X)\|_{\mathbb{P},2}^2 \\ &+ \frac{1}{(1 - \alpha)^2} \mathbb{E} \left[(Y^{(1)} - f(X))^+ (Y^{(2)} - f(X))^+ \right] \\ &- \frac{2}{1 - \alpha} \mathbb{E} \left[(h(X) - f(X)) (Y - f(X))^+ \right] \end{aligned} \quad (3.18)$$

and $\frac{2 - \alpha}{c(1 - \alpha)} \|\mathbb{P}[Y \geq f(X) | X] - 1 + \alpha\|_{\mathbb{P},2}$ is given by (3.16).

Proof. See Section B.7. \square

The expectations and probabilities in (3.16) and (3.18) can be estimated by Monte Carlo simulation (see Algorithm 1), as opposed to nested Monte Carlo that would be

required to explicitly attempt to approximate the conditional expectations involved in the left hand sides of (3.16) and (3.18). Moreover the accuracy of the twin Monte-Carlo estimates can be controlled by computing confidence intervals.

As $1 - \alpha = \mathbb{P}[Y \geq q(X)|X]$ holds \mathbb{P} almost surely, the distance in (3.16) can be interpreted as a distance in p -values between the quantile estimate $f(X)$ and the true quantile $q(X)$. Note that because (3.17) is only an inequality and due to the $\frac{1}{c}$ factor in (3.18), the control on $\|h(X) - s(X)\|_{\mathbb{P},2}$ provided by (3.17) can be quite crude. Hence the control on $\|h(X) - s(X)\|_{\mathbb{P},2}$ provided by twin Monte Carlo based on (3.16)-(3.17)-(3.18) can be quite conservative. In particular, such a procedure is appropriate for ensuring that $\|h(X) - s(X)\|_{\mathbb{P},2}$ is good enough for a given estimator h of s , but it cannot be used to compare two estimators of s . For comparing different estimators, however (of the expected shortfall or the quantile), twin Monte Carlo for e.g. $\mathbb{E} [(h(X) - \mathbb{E}[\phi(Y, q(X))|X])^2]$ is not needed anyway, it is enough to compare the corresponding out-of-sample losses e.g.

$$\begin{aligned} \mathbb{E} [(h(X) - \phi(Y, q(X)))^2] &= \mathbb{E} [(h(X) - \mathbb{E}[\phi(Y, q(X))|X])^2] + \\ &\mathbb{E} [(\phi(Y, q(X)) - \mathbb{E}[\phi(Y, q(X))|X])^2]. \end{aligned}$$

It may happen that the empirical version of the right-hand side in (3.16), or at least the lower bound of the corresponding Monte Carlo confidence interval, is negative, in which case the corresponding Monte Carlo estimate or lower bound cannot be used directly for the left-hand side. However, one could mitigate that by using a confidence upper-bound for the twin Monte-Carlo estimator, where a high enough confidence level can help get a more robust upper-bound for the a posteriori L^2 error. Also, the fact that we chose $\mathbb{P} [Y^{(1)} > f(X)] + \mathbb{P} [Y^{(2)} > f(X)]$ rather than $2\mathbb{P} [Y^{(1)} > f(X)]$ in (3.16) is arbitrary in this respect (and just motivated by the formal symmetry of the ensuing formula).

In the case where the twin Monte Carlo estimates for the right-hand-sides in (3.16) and (3.18), after having been confirmed to be accurate by drawing enough samples, are not good enough, one can improve the numerical optimization, in first attempt, and then act on the hypothesis space. For instance, in the case of the next sections of the paper where hypothesis spaces of neural networks are used, one can improve the corresponding stochastic gradient descent by changing the optimizer (e.g. switching from the basic SGD of Algorithm 2 to a more sophisticated Adam SGD as effectively done in our numerics), in first attempt, and then try to train with more layers/units or better architectures.

<p>name : TwinVal</p> <p>input : out-of-sample $\{(X_i, Y_i^{(1)}, Y_i^{(2)})\}_{i=1}^n$ with $Y_i^{(1)}, Y_i^{(2)}$ independent copies of Y given $X = X_i$, a confidence level α, corresponding estimates f and h of q and s, tolerance levels δ^{var} and δ^{es}</p> <p>output: Quality of f and h</p> <p>1 Compute $(\epsilon^{\text{var}})^2 = \frac{1}{n} \sum_{i=1}^n ((1 - \alpha)(1 - \alpha - \mathbb{1}_{Y_i^{(1)} > f(X_i)} - \mathbb{1}_{Y_i^{(2)} > f(X_i)}) + \mathbb{1}_{Y_i^{(1)} \wedge Y_i^{(2)} > f(X_i)})$</p> <p>2 if $\epsilon^{\text{var}} > \delta^{\text{var}}$ then</p> <p>3 Reply already f is bad</p> <p>4 else</p> <p>5 Compute $(\epsilon^{\text{es}})^2 = \frac{1}{n} \sum_{i=1}^n \left[(h(X_i) - f(X_i))^2 + \frac{1}{(1-\alpha)^2} (Y_i^{(1)} - f(X_i))^+ (Y_i^{(2)} - f(X_i))^+ - \frac{2}{1-\alpha} (h(X_i) - f(X_i))(Y_i^{(1)} - f(X_i))^+ \right]$</p> <p>6 if $\epsilon^{\text{es}} > \delta^{\text{es}}$ then</p> <p>7 Reply f is good but h is bad</p> <p>8 else</p> <p>9 Reply f and h are good</p> <p>10 end</p> <p>11 end</p>

Algorithm 1: Twin Monte Carlo validation for VaR and ES.

4 Learning Using Neural Networks

In this section we follow up on the theory of Section 3 in the case of learning with the following class of neural networks that will be used as hypothesis spaces.

Definition 4.1. Let σ be a 1-Lipschitz, positive-homogeneous activation function (such as ReLU), applied element-wise when supplied with a vector as input, and let $(d, o, l, m, B_{1:(l+1)}) \in \mathbb{N}^4 \times (0, \infty)^{l+1}$. We consider the family of neural networks on $\mathcal{X} \subseteq \mathbb{R}^d$ with o outputs, l hidden layers, m (or less) hidden units, activation function σ , and Euclidean norm regularisation on the weights and bias, i.e., with $\tilde{x} := \begin{bmatrix} x \\ 1 \end{bmatrix}$ introduced to take into account the bias,

$$\mathcal{NN}(d, o, B_{1:(l+1)}, l, m, \sigma) = \left\{ \mathbb{R}^d \ni x \mapsto W_{l+1} \sigma(W_l \sigma(\dots \sigma(W_1 \tilde{x}))) \in \mathbb{R}^o; \right. \\ \left. W_1 \in \mathbb{R}^{m \times (d+1)}, W_{2:l} \in (\mathbb{R}^{m \times m})^{l-1}, W_{l+1} \in \mathbb{R}^{o \times m}, |W_k|_2 \leq B_k \text{ for } k = 1 \dots l + 1 \right\}.$$

The norm $|\cdot|_2$, defined as the square root of the sum of the squares of all entries of \cdot , is called the Frobenius norm when \cdot is a matrix. It is not difficult to see that, for matrices (or vectors) W and V such that the product WV is well defined,

$$|WV|_2 \leq |W|_2 |V|_2. \quad (4.1)$$

4.1 Error Bound of the Single- α Learning Algorithm With Neural Networks

Lemma 4.1. *For any $f \in \mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma)$, we have*

$$|f(X)| \leq (|X|_2 + 1) \prod_{k=1}^{l+1} B_k. \quad (4.2)$$

If \mathcal{X} is compact, then $\mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma)$ is uniformly bounded and

$$\mathcal{R}_{ave}(\mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma), X_{1:n}) \leq \frac{\sqrt{2(l+1)\log 2} + 1}{\sqrt{n}} (\|X\|_2 \|P_{X,\infty}\| + 1) \prod_{k=1}^{l+1} B_k. \quad (4.3)$$

Proof. See Section B.8. \square

The estimate (4.3) can be combined with Theorem 3.2 to provide a more explicit error control on \hat{q} .

Theorem 4.2. *In the setup of Theorem 3.2 with $-a(X) = b(X) = (|X|_2 + 1) \prod_{k=1}^{l+1} B_k$ and $\mathcal{F} = \mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma)$, assume that*

- (i) \mathcal{X} is a compact subset of \mathbb{R}^d ,
- (ii) $c \leq \gamma_f(X) \leq C$ holds P a.s. for some constants $0 < c \leq C < \infty$ independent of $f \in \mathcal{F}$.

Then there exist \tilde{q} and \hat{q} satisfying (3.1). For any $\delta \in (0, 1)$,

$$\begin{aligned} c(1-\alpha)^{1/2} \|\hat{q} - q\|_{P_{X,2}} &\leq \left((2-\alpha) \inf_{f \in \mathcal{F}} \|f - q\|_{P_{X,1}} \right)^{1/2} \wedge \left(C(1-\alpha)^{1/2} \inf_{f \in \mathcal{F}} \|f - q\|_{P_{X,2}} \right) \\ &+ \sqrt{2} \left(\frac{(2-\alpha)(\sqrt{2(l+1)\log 2} + 1)}{\sqrt{n}} + \sqrt{\frac{2\log(2/\delta)}{n}} \right)^{1/2} \left((\|X\|_2 \|P_{X,\infty}\| + 1) \prod_{k=1}^{l+1} B_k \right)^{1/2} \end{aligned} \quad (4.4)$$

holds with probability at least $1 - \delta$.

Proof. See Section B.9. \square

Remark 4.1. Under the regularization embedded in $\mathcal{F} = \mathcal{NN}(d, \sigma, B_{1:(l+1)}, l, m, \sigma)$, the estimation error term (second line) in (4.4) does not depend on the number of neurons m nor on the dimension d of the input space. The bias term in the first line is decreasing in m (via \mathcal{F}) and potentially increasing in d (via \mathcal{F} and q).

Remark 4.2. Our convergence rate $n^{-1/4}$ for the statistical error is consistent with Padilla et al. (2022, Theorem 2), which however only deals with the case of deterministic X_i and the corresponding empirical measure. Shen, Jiao, Lin, Horowitz, and Huang (2021) requires integrability of Y of order $p > 1$, with rates better than $n^{-1/4}$ for $p > 2$. On the other side, both papers rely on VC arguments leading to unspecified constants in their statistical error bounds, see also our discussion below Theorem 3.3.

An analogous reasoning, using this time Theorem 3.3, leads to the following error control on \widehat{r}_f (with $(\mathcal{H})^+ = \{h^+ : h \in \mathcal{H}\}$ for any set of functions \mathcal{H}):

Theorem 4.3. *In the setup of Theorem 3.3 with \mathcal{X} compact, $B = (\|X\|_2 \|P_{X,\infty} + 1) \prod_{k=1}^{l+1} B_k$, and $\mathcal{G} = (\mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma))^+$, there exists a function \widehat{r}_f satisfying (3.10). For any $\delta \in (0, 1)$ and $f \in \mathcal{L}_{a,b}$,*

$$\begin{aligned} & \|\widehat{r}_f - r\|_{P_{X,2}} \leq \inf_{g \in \mathcal{G}} \|g - r\|_{P_{X,2}} \\ & + 2 \left((1 - \alpha)^{-1} \|f - q\|_{P_{X,2}} + \left\| \left(\frac{(Y - q(X))^+}{1 - \alpha} - (\|X\|_2 \|P_{X,\infty} + 1) \prod_{k=1}^{l+1} B_k \right)^+ \right\|_{P,2} \right) \\ & + \left(\frac{4(\sqrt{2(l+1)\log 2} + 1)}{\sqrt{n}} + \sqrt{\frac{2\log(2/\delta)}{n}} \right)^{1/2} (\|X\|_2 \|P_{X,\infty} + 1) \prod_{k=1}^{l+1} B_k \end{aligned}$$

holds with probability at least $1 - \delta$.

Proof. See Section B.10. \square

4.2 Algorithms

<p>name : SGDOpt</p> <p>input : $(X, Y)_{1:n}$, a partition Π of $\{1 \dots n\}$, a number of epochs $E \in \mathbb{N}^*$, a learning rate $\eta > 0$, initial weight (matrix) \widehat{W} parameters, a loss function $\ell = \ell(W, \text{batch})$ and a regularisation weight κ (set at 0.01 by default)</p> <p>output: Trained parameters \widehat{W}</p> <pre> 1 Set $\ell(W, \text{batch}) = \ell(W, \text{batch}) + \kappa W _2$ 2 for $epoch = 1, \dots, E$ do // loop over epochs 3 for $batch \in \Pi$ do // loop over batches 4 $\widehat{W} \leftarrow \widehat{W} - \eta \nabla_W \ell(\widehat{W}, \text{batch})$ 5 end 6 end</pre>
--

Algorithm 2: Mini-batch stochastic gradient descent in a neural net hypothesis space.

In practice the hard constraints $B_{1:(l+1)}$ of the neural network family are “softly” handled by penalization: see Algorithm 2. Accordingly, we drop $B_{1:(l+1)}$ and simplify the notation for the network family to $\mathcal{NN}(d, o, l, m, \sigma)$ in what follows. Let $\zeta_{l+1}^{d,o}(x, W)$ denote a function in the family $\mathcal{NN}(d, o, l, m, \sigma)$ with the Softplus activation function, i.e. $\sigma(x) = \log(1 + \exp(x))$, where W represents the set of network parameters. We observed numerically similar performances between ReLU and Softplus networks. Despite our theoretical bounds in Section 4.1 being built upon the ReLU network, we use the Softplus networks in our numerics below for their ability to provide analytical derivatives of their outputs with respect to their inputs, as required in Section 5.

Given an i.i.d sample $(X, Y)_{1:n}$ of (X, Y) as before, instead of considering global argminima \widehat{q} and (given \widehat{q}) $\widehat{r}_{\widehat{q}}$ of the related empirical losses as in the theoretical analysis

of Sections 3.1–3.4, we consider various mini-batch SGD estimates (cf. Algorithm 2) for \widehat{q} and $\widehat{s}_{\widehat{q}}$. These SGD estimates are then amenable to the a posteriori (twin Monte Carlo) error analysis of Section 3.5.

name : VaRAlg
input : $(X, Y)_{1:n}$, a partition Π of $\{1 \dots n\}$, a quantile level α , a number of epochs $E \in \mathbb{N}^*$, a learning rate $\eta > 0$, initial values for the network parameters \widehat{W} , and neural network output function $\zeta_{l+1}^{d,1}(x, W)$
output: Trained parameters of VaR network
1 define $\ell^{\text{var}}(W, \text{batch}) = \frac{1}{|\text{batch}|} \sum_{i \in \text{batch}} [(Y_i - \zeta_{l+1}^{d,1}(X_i, W))^+ + (1 - \alpha)\zeta_{l+1}^{d,1}(X_i, W)]$
2 $\widehat{W}^{\text{var}} \leftarrow \text{SGDOpt}((X, Y)_{1:n}, \Pi, E, \eta, \widehat{W}, \ell^{\text{var}})$

Algorithm 3: Neural network regression for learning the VaR.

Algorithm 3 thus produces a VaR predictor

$$\zeta_{l+1}^{d,1}(X, \widehat{W}^{\text{var}}) \approx \widehat{q}(X)$$

in the form a function of X represented by a neural network from $\mathcal{NN}(d, 1, l, m, \sigma)$, for given m and l .

name : ESAlg
input : $(X, Y)_{1:n}$, a partition Π of $\{1 \dots n\}$, a quantile level α , a number of epochs $E \in \mathbb{N}^*$, a learning rate $\eta > 0$, initial values for the network parameters \widehat{W} and neural network output function $\zeta_{l+1}^{d,1}(x, W)$
output: Trained parameters of ES network \widehat{W}^{es}
1 // Learn the corresponding VaR
2 $\widehat{W}^{\text{var}} \leftarrow \text{VaRAlg}((X, Y)_{1:n}, \Pi, \alpha, E, \eta, \widehat{W})$
3 **if** *linear regression* **then**
4 | // Remind W_k denote the weight of k th layer
5 |
$$\{\widehat{W}_k^{\text{es}}\}_{k=1}^l \leftarrow \{\widehat{W}_k^{\text{var}}\}_{k=1}^l$$

6 |
$$\widehat{W}_{l+1}^{\text{es}} \leftarrow \underset{W_{l+1}}{\text{argmin}} \frac{1}{n} \sum_{i=1}^n \left[(1 - \alpha)^{-1} (Y_i - \zeta_{l+1}^{d,1}(X_i, \widehat{W}^{\text{var}}))^+ + \zeta_{l+1}^{d,1}(X_i, \widehat{W}^{\text{var}}) - \zeta_{l+1}^{d,1}(X_i, (\{\widehat{W}_k^{\text{var}}\}_{k=1}^l, W_{l+1})) \right]^2$$

7 **else**
8 | define $\ell^{\text{es}}(W, b, \text{batch}) = \frac{1}{|\text{batch}|} \sum_{i \in \text{batch}} [((1 - \alpha)^{-1} (Y_i - \zeta_{l+1}^{d,1}(X_i, \widehat{W}^{\text{var}}))^+ + \zeta_{l+1}^{d,1}(X_i, \widehat{W}^{\text{var}}) - \zeta_{l+1}^{d,1}(X_i, W))^2]$
9 | $\widehat{W}^{\text{es}} \leftarrow \text{SGDOpt}((X, Y)_{1:n}, \Pi, E, \eta, \widehat{W}^{\text{var}}, \ell^{\text{es}})$
9 **end**

Algorithm 4: Neural network regressions for learning the ES in two steps, using least-squares regression of the linear readout map or full NN training or for deducing the ES from the VaR.

Given this VaR predictor, Algorithm 4 produces

$$\zeta_{l+1}^{d,1}(X, \widehat{W}^{\text{es}}) \approx \widehat{s}_{\hat{q}}(X),$$

in two possible ways: either by training a dedicated neural network from scratch, or by using the same architecture as the one used for the VaR, freezing the weights of all hidden layers as those of the VaR network and least-square regressing the linear-readout map. We show in Section 6 that such a transfer learning scheme is enough to obtain good approximations, while also being very fast (a fraction of a second in our experiments) using highly optimized linear algebra routines such as the ones implemented by cuBLAS for Nvidia GPUs.

5 Multi- α Learning for VaR

In this part we are interested in learning $\text{VaR}(Y|X)$ for multiple confidence levels $\alpha \in (0, 1)$ using a single empirical error minimization. This can help give insights into the sensitivity of $\text{VaR}(Y|X)$ with respect to the confidence level, or into the full distribution of Y given X (e.g. approximated by a histogram representation). Given that the previous trainings were done for a single fixed confidence level α , we refer to them as the single- α learning (or single- α for brevity in the numerics): under this approach, if one is interested in finding the conditional VaR for another confidence level, one has to repeat the training procedure using the new confidence level.

By contrast, the multi- α learning approaches below allow learning $\text{VaR}(Y|X)$ for multiple confidence levels within a single simulation run. Regarding the learning of $\text{ES}(Y|X)$ for multiple confidence levels, the transfer learning trick (cf. the linear regression case in Algorithm 4) in Section 4.2 was found to provide the most valuable alternative, whether done α by α (our choice in the numerics below), as each run of it is very fast, or globally across α 's based on either of the multi- α VaR approaches below. Hence we focus on the multi- α learning of VaR in what follows.

5.1 The Crossing Quantile Issue

When several quantile levels α are considered, a flaw inherent to linear quantile regression is the problem of crossing quantile curves, i.e. the violation of the monotonicity with respect to α . The simultaneous learning of conditional quantiles for multiple confidence levels and the problem of quantile crossing were early addressed in He (1997), Koenker (2004) and Takeuchi, Le, Sears, and Smola (2006), see also Moon, Jeon, Lee, and Kim (2021) for a review of more recent references. To deal with the quantile crossing problem, two strategies for constraints can be considered.

The first strategy is to consider explicitly the non-crossing constraints during the learning phase of the model in form of either hard constraints (that the model must strictly satisfy) or soft constraints (i.e. penalization). Once the non-crossing hard constraints are employed, the model is usually learned using primal-dual optimization algorithms. The latter are applicable in a wide class of models, e.g. support vector regression (Takeuchi, Le, Sears, and Smola, 2006; Sangnier, Fercoq, and d'Alché Buc, 2016) and spline regression (Bondell, Reich, and Wang, 2010), but notably not in the

case of the family of (deep) neural networks, because of the computational cost and the poor scalability of projected gradient descent. Therefore, the non-crossing constraints are more preferably embedded in the training of neural networks via a penalty term, based in Moon et al. (2021) on a finite difference of the output of the neural network (that approximates the value-at-risk) for two confidence levels.

In Section 5.3 we use a similar penalization strategy, where, instead of penalizing the negative part of a finite difference, we penalize the negative part of the partial derivative of the network with respect to the confidence level. The partial derivative gives more information about the local behavior around training points and we can penalize its negative part at every α that appears at the training stage, e.g. for several thousands values of α in our numerics below, as opposed to penalizing negative increments at a few fixed values of α as in Moon et al. (2021). Our approach also spares one hyperparameter, namely the size of the discrete increment in confidence levels used for the finite differences.

The second strategy is to use hypothesis spaces of functions nondecreasing with respect to the confidence level. Meinshausen and Ridgeway (2006) introduce quantile regression forests. In this model the predicted quantile of a new point is based on the empirical percentile of the group (i.e. the terminal leaf of each tree) where this point belongs, hence the monotonicity of the quantile estimates is satisfied by construction. Regarding neural networks, Hatalis, Lamadrid, Scheinberg, and Kishore (2017) propose a specific initialization scheme for the weights of the output layer, which does not prevent quantile crossings, but appears to reduce them significantly in their experiments. Cannon (2018) considers the confidence level as an additional explanatory variable and then explores a network such that the estimate is monotone with a defined covariate (confidence level), imposing the non-crossing. Gasthaus, Benidis, Wang, Rangapuram, Salinas, Flunkert, and Januschowski (2019) and Padilla, Tansey, and Chen (2022) use a (deep) network with multiple outputs, constrained by design to be positive, which are expected to approximate quantile increments.

The latter resembles our approach in Section 5.4, but we sample the confidence level uniformly on a given interval and we further interpolate linearly with respect to the confidence level before insertion of the output of the neural network in the training loss (cf. (5.2)-(5.3)), in order to have a conditional quantile function that is valid for all quantile levels in the interval.

5.2 Extension of the Bounds to Multi- α Learning

The various proofs and bounds presented in this paper for a fixed $\alpha \in [0, 1]$ can be extended to the multi- α learning framework where α is now a random variable supported by $\mathcal{I} = [\underline{\alpha}, \bar{\alpha}] \subseteq (0, 1)$, treated as a new covariate alongside X . Hereafter we randomize α assuming $\alpha \sim \mathcal{U}([\underline{\alpha}, \bar{\alpha}])$. We then consider a finite i.i.d sample $\alpha_1, \dots, \alpha_n$ of α , independent of covariates (α, X) and of the sample $(X, Y)_{1:n}$, and the loss functions that appear, together with corresponding changes to the results above, in Table 1. The implementation of this approach using neural networks is discussed below.

Single- α	Multi- α
X valued in \mathcal{X}	(α, X) valued in $\mathcal{I} \times \mathcal{X}$
$F_{Y X}(a(X)) < \alpha < F_{Y X}(b(X))$, P-a.s.	$F_{Y X}(a(X)) < \underline{\alpha} \leq \bar{\alpha} < F_{Y X}(b(X))$, P-a.s.
$\phi(y, u) = (1 - \alpha)^{-1}(y - u)^+ + u$	$\phi(\alpha, y, u) = (1 - \alpha)^{-1}(y - u)^+ + u$
$\Phi(f) = \mathbb{E} [\phi(Y^{a,b}, f(X)) (X, Y)_{1:n}]$	$\Phi(f) = \mathbb{E} [\phi(\alpha, Y^{a,b}, f(\alpha, X)) (\alpha, X, Y)_{1:n}]$
spaces $\mathcal{F} \subseteq \mathcal{L}_{a,b}$ of functions on \mathcal{X}	spaces $\mathcal{F} \subseteq \mathcal{L}_{a,b}$ of functions on $\mathcal{I} \times \mathcal{X}$
$\tilde{q} \in \operatorname{argmin}_{f \in \mathcal{F}} \Phi(f)$	$\tilde{q} \in \operatorname{argmin}_{f \in \mathcal{F}} \Phi(f)$
$\hat{q} \in \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \phi(Y_i^{a,b}, f(X_i))$	$\hat{q} \in \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \phi(\alpha_i, Y_i^{a,b}, f(\alpha_i, X_i))$
$\gamma_f(x) = \Gamma_{F_{Y X=x}}(f(x), q(x))^*$	$\gamma_f(\alpha, x) = \Gamma_{F_{Y X=x}}(f(\alpha, x), q(\alpha, x))^*$
$\Phi(\tilde{q}) - \Phi(q) = \ \gamma_{\tilde{q}}(\tilde{q} - q)\ _{\mathbb{P}_{X,2}}^2 =$ $\inf_{f \in \mathcal{F}} \ \gamma_f(f - q)\ _{\mathbb{P}_{X,2}}^2 \leq \frac{2 - \alpha}{1 - \alpha} \inf_{f \in \mathcal{F}} \ f - q\ _{\mathbb{P}_{X,1}}$	$\Phi(\tilde{q}) - \Phi(q) = \ \gamma_{\tilde{q}}(\tilde{q} - q)\ _{\mathbb{P}_{\alpha,X,2}}^2 =$ $\inf_{f \in \mathcal{F}} \ \gamma_f(f - q)\ _{\mathbb{P}_{\alpha,X,2}}^2 \leq \frac{2 - \bar{\alpha}}{1 - \bar{\alpha}} \inf_{f \in \mathcal{F}} \ f - q\ _{\mathbb{P}_{\alpha,X,1}}$
$(1 - \alpha)^{-1} \left(2(2 - \alpha) \mathcal{R}_{ave}(\mathcal{F}, X_{1:n}) \right.$ $\left. + \frac{\ b - a\ _{\mathbb{P}_{X,\infty}} \sqrt{2 \log(2/\delta)}}{\sqrt{n}} \right)$	$(1 - \bar{\alpha})^{-1} \left(2(2 - \bar{\alpha}) \mathcal{R}_{ave}(\mathcal{F}, (\alpha, X)_{1:n}) \right.$ $\left. + \frac{\ b - a\ _{\mathbb{P}_{X,\infty}} \sqrt{2 \log(2/\delta)}}{\sqrt{n}} \right)$
$\mathcal{R}_{ave}(\mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma), X_{1:n})$ $\leq \frac{\sqrt{2(l+1) \log 2} + 1}{\sqrt{n}} (\ X\ _2 \ X\ _{\mathbb{P}_{X,\infty}} + 1) \prod_{k=1}^{l+1} B_k$	$\mathcal{R}_{ave}(\mathcal{NN}(d+1, 1, B_{1:(l+1)}, l, m, \sigma), (\alpha, X)_{1:n})$ $\leq \frac{\sqrt{2(l+1) \log 2} + 1}{\sqrt{n}} (\bar{\alpha} + \ X\ _2 \ X\ _{\mathbb{P}_{X,\infty}} + 1) \prod_{k=1}^{l+1} B_k$
$es(u) = (1 - \alpha)^{-1} \int_u^\infty (y - u)^+ F(dy) + u,$ $\Delta_F(u, q) = es(u) - es(q)$ $\Gamma_F(u, q) = \frac{\Delta_F(u, q)}{(u - q)^2} \mathbf{1}_{(0,\infty)}(u - q)$ $+ \frac{\dot{F}(q)}{2(1 - \alpha)} \mathbf{1}_{\{0\}}(u - q)$	$es(\alpha, u) = (1 - \alpha)^{-1} \int_u^\infty (y - u)^+ F(dy) + u,$ $\Delta_F(\alpha, u, q(\alpha)) = es(\alpha, u) - es(\alpha, q(\alpha))$ $\Gamma_F(\alpha, u, q(\alpha)) = \frac{\Delta_F(\alpha, u, q(\alpha))}{(u - q(\alpha))^2} \mathbf{1}_{(0,\infty)}(u - q(\alpha))$ $+ \frac{\dot{F}(q(\alpha))}{2(1 - \alpha)} \mathbf{1}_{\{0\}}(u - q(\alpha))$

Table 1: Main changes required to adapt the previous results and proofs from a single-quantile to a multi-quantile regression setup.

* defined from the unconditional notation displayed after the double line.

5.3 Learning With a Continuum of α

The finite-sample training problem for this approach can be stated as follows:

$$\operatorname{argmin}_W \frac{1}{n} \sum_{i=1}^n [(Y_i - \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W))^+ + (1 - \alpha_i) \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W)],$$

where $[a, x]$ denotes the vector obtained by concatenating a vector x to a real a . One can also approximately impose the non-crossing of the quantiles by penalizing the sample average of the negative part of the partial derivative $\partial_\alpha \zeta_{l+1}^{d+1,1}([\alpha, X], W)$, as per

$$\operatorname{argmin}_W \frac{1}{n} \sum_{i=1}^n \left[(Y_i - \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W))^+ + (1 - \alpha_i) \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W) + \lambda \left(\partial_\alpha \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W) \right)^- \right], \quad (5.1)$$

where $\lambda > 0$ determines the strength of the penalization. Algorithm 5 thus produces

<pre> name : MultiContinuousVaRAlg input : $(X, Y)_{1:n}$, a partition Π of $\{1 \dots n\}$, a quantile upper bound level $\bar{\alpha}$, and lower bound level $\underline{\alpha}$, a number of epochs $E \in \mathbb{N}^*$, a learning rate $\eta > 0$, a regularisation parameter $\lambda \geq 0$, initial values for the network parameters \widehat{W} and neural network output function $\zeta_{l+1}^{d+1,1}([a, x], W)$ output: Trained parameters of multi-VaR network \widehat{W} 1 // Sample quantile levels α 2 $\alpha_i \sim \text{Uniform}(\underline{\alpha}, \bar{\alpha})$ for $i = 1 \dots n$ 3 // Define a loss function 4 if <i>non-crossing quantile regularisation</i> then 5 // multi-α(I) 6 define $\ell^{\text{vars}}(W, \text{batch}) = \frac{1}{ \text{batch} } \sum_{i \in \text{batch}} [(Y_i - \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W))^+ + (1 -$ 7 $\alpha_i) \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W) + \lambda (\partial_\alpha \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W))^-]$ 8 where $\partial_\alpha \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W)^-$ can be quickly computed 9 else 10 // multi-α(II) 11 define $\ell^{\text{vars}}(W, \text{batch}) =$ 12 $\frac{1}{ \text{batch} } \sum_{i \in \text{batch}} [(Y_i - \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W))^+ + (1 - \alpha_i) \zeta_{l+1}^{d+1,1}([\alpha_i, X_i], W)]$ 13 end 14 $\widehat{W}^{\text{vars}} \leftarrow \text{SGDOpt}((X, Y)_{1:n}, \Pi, E, \eta, \widehat{W}, \ell^{\text{vars}})$ </pre>
--

Algorithm 5: Learning multi continuous VaR.

$$\zeta_{l+1}^{d+1,1}(\alpha, X, \widehat{W}^{\text{vars}}) \approx \widehat{q}_\alpha(X),$$

where $\widehat{q}_\alpha(X)$ denotes the theoretical estimator (multi- α analog of $\widehat{q}(X)$ before, ignoring the SGD numerical optimization error) of the VaR($Y|X$) quantile function at any quantile level in $[\underline{\alpha}, \bar{\alpha}]$. Note that for non-zero λ , one uses a penalized pinball loss function

that is beyond the scope of Theorems 3.2-4.2. The a posteriori Monte Carlo validation technique of Section 3.5, though, is still applicable to this (as to any) estimator. Importantly, one can compute the derivative in (5.1) fast in closed-form given our neural network parametrization, as $\partial_\alpha \zeta_{l+1}^{d+1,1}([\alpha, X], W) = W_{l+1} \partial_\alpha \zeta_n^{d+1,1}([\alpha, X], W)$, where

$$\begin{aligned} \partial_\alpha \zeta_0^{d+1,1}([\alpha, X], W) &= [1, 0_d] \text{ and, for } k = 1, \dots, l, \\ \partial_\alpha \zeta_k^{d+1,1}([\alpha, X], W) &= (W_k \partial_\alpha \zeta_{k-1}^{d+1,1}([\alpha, X], W)) \odot \sigma'(W_k \zeta_{k-1}^{d+1,1}([\alpha, X], W)). \end{aligned}$$

Here \odot is an element-wise product and σ' is the derivative of σ (applied element-wise). Given the computations of $\zeta_{l+1}^{d+1,1}([\alpha, X], W)$ and $\partial_\alpha \zeta_{l+1}^{d+1,1}([\alpha, X], W)$ share many common sub-expressions, the recursions can be done at the same time, i.e. at each $k \in \{0, \dots, l+1\}$, compute $\zeta_k^{d+1,1}([\alpha, X], W)$ and then reuse the common sub-expressions to compute also $\partial_\alpha \zeta_k^{d+1,1}([\alpha, X], W)$. In the numerics, we refer to this approach with multi- α (I) if we use a non-zero λ , and multi- α (II) otherwise.

5.4 Learning Via a Discrete Set of α 's and Linear Interpolation

Another approach for multi- α learning is to use a finite set of confidence levels $\alpha^{(1)} < \dots < \alpha^{(o)}$ in $[\underline{\alpha}, \bar{\alpha}]$ in conjunction with linear interpolation. More precisely, we solve

$$\operatorname{argmin}_W \frac{1}{n} \sum_{i=1}^n \left[\left(Y_i - \Sigma(\alpha_i, \zeta_{l+1}^{d,o}(X_i, W)) \right)^2 + (1 - \alpha_i) \Sigma(\zeta_{l+1}^{d,o}(\alpha_i, X_i, W)) \right], \quad (5.2)$$

where, for $y = (y_0, \dots, y_{o-1})^\top$ and $a \in [\underline{\alpha}, \bar{\alpha}]$,

$$\Sigma(a, y) = y_0 + \sum_{j=1}^{o-1} \mathbf{1}_{\alpha^{(j)} \leq a} \frac{(\alpha^{(j+1)} \wedge a - \alpha^{(j)})}{\alpha^{(j+1)} - \alpha^{(j)}} y_j. \quad (5.3)$$

Algorithm 6 thus produces

$$\Sigma(\alpha, \zeta_{l+1}^{d,o}(X, \widehat{W}^{vars})) \approx \widehat{q}_\alpha(X),$$

where $[\zeta_{l+1}^{d,o}(X, W)]_0$ can be interpreted as a predictor of the value-at-risk of lowest grid level $\alpha^{(1)}$, whereas, for each $j \geq 1$, $[\zeta_{l+1}^{d,o}(X, W)]_j$ is a predictor of the increment between the value-at-risks of levels $\alpha^{(j)}$ and $\alpha^{(j+1)}$. Notice that one could impose the monotonicity by design by adding a positive activation function σ to each neuron in the output layer of $\zeta_{l+1}^{d+1,o}$, except for the first neuron, e.g. by replacing

$$y_j \text{ with } \sigma(y_j), \text{ for all } j \in 1, \dots, o-1,$$

in (5.3). However we have not found doing so to be satisfactory numerically and thus we keep the formulation in (5.3) as it is. In the numerics, we refer to this approach as multi- α (III).

We now test the proposed procedures on a Student toy example and a dynamic initial margin (DIM) case study. Any minimization of loss functions over $\mathcal{NN}(d, o, l, m, \sigma)$

<p>name : MultiDiscreteVaRAlg// multi-α(III)</p> <p>input : $(X, Y)_{1:n}$, a partition Π of $\{1 \dots n\}$, an increasing quantile level sequence $\alpha^{(1)} < \dots < \alpha^{(o)}$, a number of epochs $E \in \mathbb{N}^*$, a learning rate $\eta > 0$, initial values for the network parameters \widehat{W}, neural network output function $\zeta_{l+1}^{d,o}(x, W)$</p> <p>output: Trained parameters of multi-VaR network \widehat{W}</p> <p>1 // Sample quantile levels α</p> <p>2 $\alpha_i \sim \text{Uniform}(\underline{\alpha}, \bar{\alpha})$ for $i = 1 \dots n$</p> <p>3 // Define a loss function</p> <p>4 define $\Sigma(y, a) = y_0 + \sum_{j=1}^{o-1} \mathbf{1}_{\alpha^{(j)} \leq a} \frac{(\alpha^{(j+1)} \wedge a - \alpha^{(j)})}{\alpha^{(j+1)} - \alpha^{(j)}} y_j$</p> <p>5 define $\ell^{vars}(W, \text{batch}) =$</p> $\frac{1}{ \text{batch} } \sum_{i \in \text{batch}} \left[\left(Y_i - \Sigma(\zeta_{l+1}^{d,o}(X_i, W), \alpha_i) \right)^+ + (1 - \alpha_i) \Sigma(\zeta_{l+1}^{d,o}(X_i, W), \alpha_i) \right]$ <p>6 $\widehat{W}^{vars} \leftarrow \text{SGDOpt}(\{(X_i, Y_i)\}_{i=1}^n, \Pi, E, \eta, \widehat{W}, \ell^{vars})$</p>
--

Algorithm 6: Learning multi discrete VaR.

or similar sets of neural networks is done using the Adam algorithm of Kingma and Ba (2014) over the parameters W together with mini-batching.

All of our neural networks have 3 hidden layers, and twice their input dimensionality as the number of neurons per hidden layer. In both examples below, for the multi- α (I) and multi- α (II) learning approaches, we use the bounds $(1 - \bar{\alpha}, 1 - \underline{\alpha}) = (10^{-4}, 0.15)$. For the multi- α (III) approach, we use a uniform interpolation grid $1 - \alpha^{(k)} = 10^{-3} + k \frac{0.15 - 10^{-3}}{20}$, with $k \in \{0, \dots, 20\}$. The different runs and the ensuing RMSE errors referred to in our numerics share a common dataset. What is randomized from one run to the next is only the initialization of the network, in order to make our numerical conclusions robust to the training noise related to the random initialization of the SGD.

6 Conditionally Student- t Toy Model

We first apply the above algorithms to the data generating process (X, Y) such that X is a standard multivariate normal vector

$$X \sim \mathcal{N}(0, I_d), \text{ for some } d \in \mathbb{N}^*$$

and, for given functions P, Q, S of x ,

$$Y|X = T|X + UQ(X),$$

where U is an independent Rademacher variable (worth ± 1 with probabilities $1/2$), while the conditional distribution $T|X$ is Student- t with degree $\nu > 1$, location $P(X)$ and scale $S(X)$. According to Khokhlov (2016),

$$\begin{aligned} \text{VaR}(T|X) &= F_\nu^{-1}(\alpha, P(X), S(X)) = P(X) + S(X)F_\nu^{-1}(\alpha), \\ \text{ES}(T|X) &= P(X) + S(X) \frac{\dot{F}_\nu(F_\nu^{-1}(\alpha)) \nu + (F_\nu^{-1}(\alpha))^2}{1 - \alpha} \frac{1}{\nu - 1}, \end{aligned}$$

where $F_\nu(\cdot, \mu, \sigma)$ and $\dot{F}_\nu(\cdot, \mu, \sigma)$, shortened for $\mu = 0$ and $\sigma = 1$ as $F_\nu(\cdot)$ and $\dot{F}_\nu(\cdot)$, are the cdf and pdf of the Student- t distribution with degree ν , location μ and scale σ . The cdf of $Y|X$ is given by

$$\begin{aligned} F_{Y|X}(y) &= \text{P}(Y < y|X) = \frac{1}{2} \left[\text{P}(Y < y|X, U = 1) + \text{P}(Y < y|X, U = -1) \right] \\ &= \frac{1}{2} \left[\text{P}(T < y + Q(X)) + \text{P}(T < y - Q(X)) \right] \\ &= \frac{1}{2} \left[F_\nu(y, P(X) - Q(X), S(X)) + F_\nu(y, P(X) + Q(X), S(X)) \right], \end{aligned} \quad (6.1)$$

with corresponding conditional pdf

$$\dot{F}_{Y|X}(y) = \frac{1}{2} \left[\dot{F}_\nu(y, P(X) - Q(X), S(X)) + \dot{F}_\nu(y, P(X) + Q(X), S(X)) \right]. \quad (6.2)$$

The conditional VaR of Y given X , at level α , by

$$\text{VaR}(Y|X) = F_{Y|X}^{-1}(\alpha),$$

cannot be computed analytically, but one can efficiently approximate it numerically (by monotonicity of quantile functions). The formula for the ES of $Y|X$ is deduced by the following result, whose proof is deferred to Section B.11:

$$\begin{aligned} \text{ES}(Y|X) &= P(X) + \frac{1}{2(1-\alpha)} Q(X) \left[F_\nu \left(F_{Y|X}^{-1}(\alpha), P(X) - Q(X), S(X) \right) \right. \\ &\quad \left. - F_\nu \left(F_{Y|X}^{-1}(\alpha), P(X) + Q(X), S(X) \right) \right] \\ &\quad + \frac{\nu S(X)^2 + (F_{Y|X}^{-1}(\alpha) - P(X))^2 + Q(X)^2}{(\nu - 1)(1 - \alpha)} \dot{F}_{Y|X} \left(F_{Y|X}^{-1}(\alpha) \right) \\ &\quad + \frac{\left(F_{Y|X}^{-1}(\alpha) - P(X) \right) Q(X)}{(\nu - 1)(1 - \alpha)} \left[\dot{F}_\nu \left(F_{Y|X}^{-1}(\alpha), P(X) - Q(X), S(X) \right) \right. \\ &\quad \left. - \dot{F}_\nu \left(F_{Y|X}^{-1}(\alpha), P(X) + Q(X), S(X) \right) \right]. \end{aligned} \quad (6.3)$$

The corresponding values of $\text{VaR}(Y|X)$ and $\text{ES}(Y|X)$ will serve us as ground-truth values. This provides an heavy-tailed setup in which the solution can be computed quasi analytically, for benchmarking purposes, but sufficiently rich so that the data are not parametrically determined by a few scaling parameters (as it would be the case with conditionally Gaussian or even elliptical distributions).

6.1 Numerical Results

We take P, Q , and S as quadratic functions of x and set $\nu = 3$. We use a dimension of $d = 25$ for the state space of X . The nonzero coefficients of these polynomials are drawn independently from a standard normal distribution. For this example, we use $n = 2^{19} = 524288$ training points and the same number of testing points for computing the errors. For the Adam algorithm, we used 2000 epochs, mini-batching with batches

of size $2^{15} = 32768$, a learning rate $\eta = 0.01$, and the rest of the parameters kept at their default values as per (Kingma and Ba, 2014).

Figure 1 shows the convergence of the learnings at rate close to $n^{-\frac{1}{4}}$, consistent with the controls of Theorem 4.2 for the VaR and of Theorem 4.3 for the ES. As our sample size $n = 2^{19}$ satisfies (3.15), this is in line with the discussion following Theorem 3.3. Notice that setting $\nu = 3$ yields the integrability of order $p = 2$ of the t -distributed response Y , making the statistical bound in Shen et al. (2021) have a lower rate than ours (see Remark 4.2).

Tables 2, as also 6, 7 and 8 below in the DIM case, suggest that the multi- α approaches are competitive compared to the single- α approach by yielding acceptable errors for confidence levels below 99%, while requiring only one single training, as opposed to the single- α approach which requires one training per target confidence level. For very extreme confidence levels, like 99.9%, the multi- α (III) approach outperforms all the other approaches. This can be explained by the fact that, even if the target confidence level is hard to reach given a limited training set, the lower confidence levels in the interpolation grid contribute to inferring the VaR for the target confidence level. Table 3 confirms that one can rely on the twin-simulation trick of Section 3.5 to draw mostly similar conclusions as in Table 2, without the need to have access to the groundtruth estimators. Note that we computed upper-bounds of 95% confidence intervals for (3.16), instead of the estimates directly in order to be conservative and take into account the potentially high variance in the indicator functions that need to be simulated in order to estimate (3.16). Table 4 demonstrates the effectiveness of the penalization term (for λ simply set to 1) in the multi- α (I) approach to mitigate the quantiles crossing problem. Table 4 also shows that the other multi- α learning approaches, even without directly penalizing the crossing of the quantiles, behave better than a single- α learning in terms of the crossing of the quantiles.

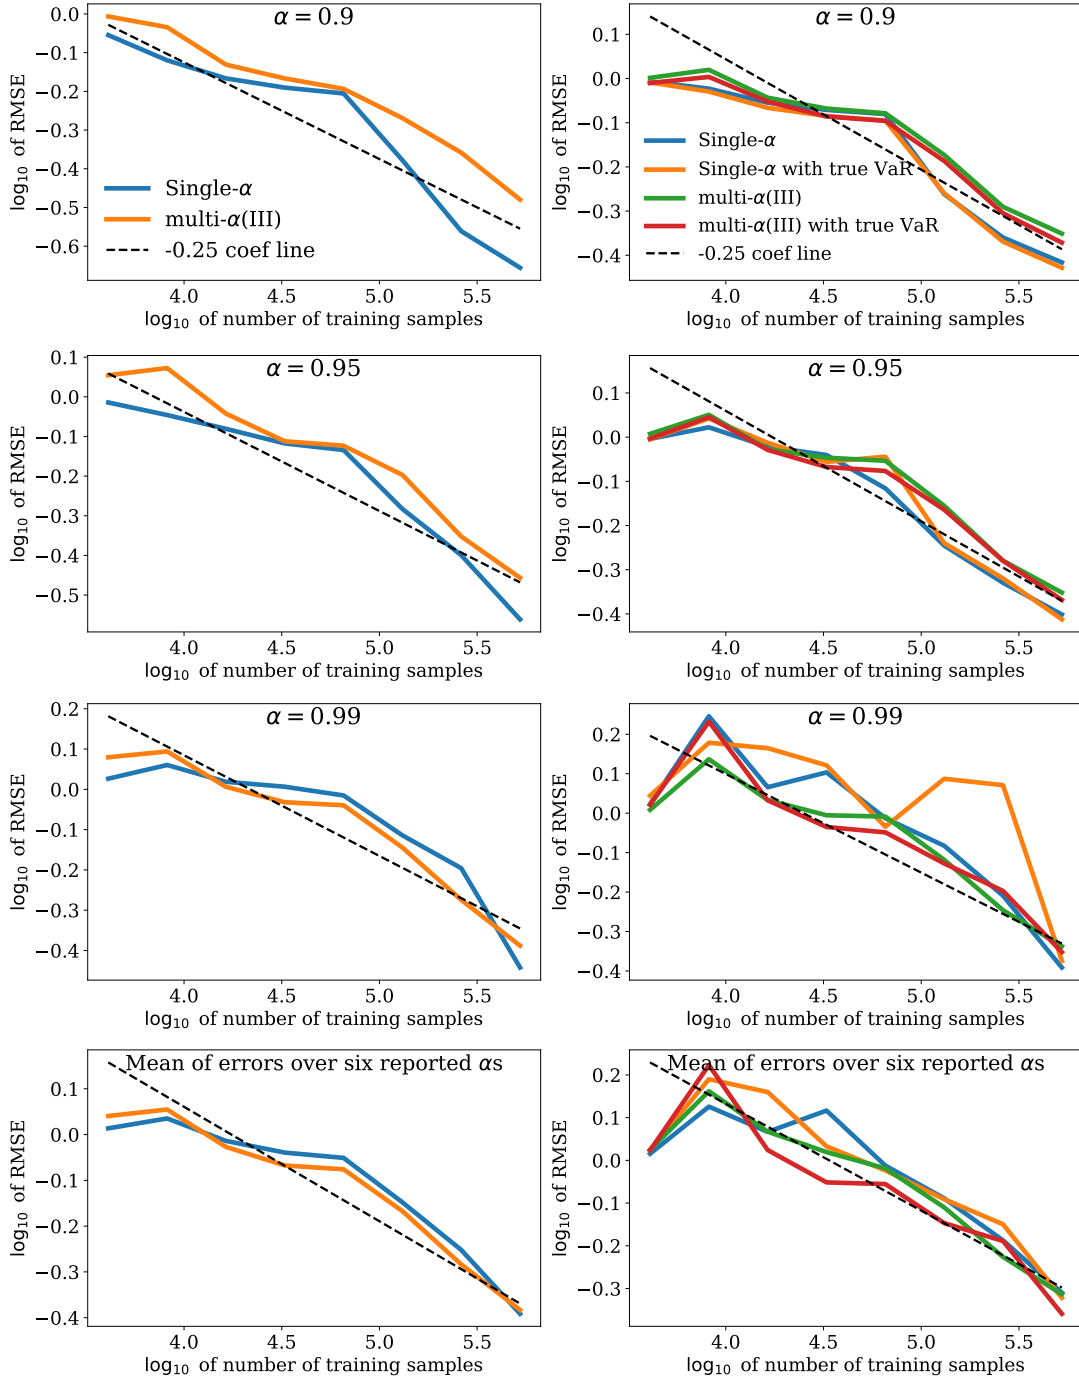


Figure 1: Convergence curves of the log-RMSE errors against their groundtruth values, in the Student- t toy-example when increasing (the log of) the number of training samples, of: (Left) the conditional VaR learned by the single and multi- α (III) models; (Right) the conditional ES learned by the single and multi- α (III) models, using the true or learned conditional VaR as VaR hypothesis for the ES. The slope of the black dash lines correspond to the theoretical convergence rate of Theorems 4.2 and 4.3.

α	0.999	0.995	0.99
single- α	0.882 (0.037)	0.548 (0.035)	0.429 (0.016)
multi- α (I)	0.821 (0.048)	0.438 (0.043)	0.342 (0.035)
multi- α (II)	0.81 (0.05)	0.438 (0.035)	0.352 (0.034)
multi- α (III)	0.549 (0.037)	0.422 (0.028)	0.408 (0.022)
α	0.98	0.95	0.9
single- α	0.363 (0.014)	0.25 (0.013)	0.209 (0.013)
multi- α (I)	0.319 (0.035)	0.263 (0.02)	0.242 (0.014)
multi- α (II)	0.325 (0.036)	0.26 (0.019)	0.237 (0.014)
multi- α (III)	0.394 (0.029)	0.346 (0.031)	0.34 (0.034)

Table 2: Means (standard deviations) of RMSE errors of learned conditional VaR estimators against groundtruth values in the Student- t toy-example across 32 runs. The RMSE errors are normalized by division by the standard deviation of the groundtruth VaR.

α	0.999	0.995	0.99
single- α	-1.989 (-2.842)	-1.641 (-2.756)	-1.528 (-2.716)
multi- α (I)	-2.029 (-2.828)	-2.0 (-2.846)	-1.903 (-2.906)
multi- α (II)	-2.019 (-2.846)	-1.98 (-2.893)	-1.888 (-2.834)
multi- α (III)	-2.36 (-2.684)	-2.11 (-2.699)	-1.713 (-2.478)
α	0.98	0.95	0.9
single- α	-1.399 (-2.83)	-1.311 (-2.549)	-1.199 (-2.68)
multi- α (I)	-1.691 (-2.908)	-1.351 (-2.622)	-1.159 (-2.538)
multi- α (II)	-1.676 (-2.762)	-1.349 (-2.662)	-1.159 (-2.525)
multi- α (III)	-1.584 (-2.412)	-1.273 (-2.309)	-1.051 (-2.222)

Table 3: Means (standard deviations) \log_{10} of the means (standard deviations) across 32 runs of a posteriori twin Monte Carlo p -value error estimates, i.e. of the right-hand side of (3.16), of learned conditional VaR estimators in the Student- t toy-example.

For the ES learning in the Student- t toy-example, for brevity, we denote by “LR using M VaR” an ES learning using linear regression only for the output layer, corresponding to the linear regression case in Algorithm 4, and a VaR learned using the method M as the candidate VaR. For example, LR using single- α VaR refers to the linear regression approach for learning the ES, by using a VaR that is learned with the single- α approach as the VaR candidate. To demonstrate the effectiveness of this linear regression approach, we also introduce an ES that is learned by neural regression, by using a neural network corresponding to the second (else) case in Algorithm 4, without freezing any weights and using the groundtruth VaR as the VaR candidate. Table 5 shows that our linear regression approach for the ES outperforms the neural regression, no matter which approach is used for learning the embedded VaR candidate. The rela-

E	$\{\widehat{q}_{0.999}(X) < \widehat{q}_{0.995}(X)\}$	$\{\widehat{q}_{0.995}(X) < \widehat{q}_{0.99}(X)\}$	$\{\widehat{q}_{0.99}(X) < \widehat{q}_{0.98}(X)\}$
single- α	-0.817 (-1.871)	-0.601 (-1.85)	-0.706 (-1.861)
multi- α (I)	-4.692 (-4.311)	-4.412 (-4.156)	-3.943 (-3.869)
multi- α (II)	-4.449 (-4.094)	-4.209 (-3.97)	-3.79 (-3.72)
multi- α (III)	-3.445 (-2.986)	-3.346 (-2.959)	-1.674 (-1.852)
E	$\{\widehat{q}_{0.98}(X) < \widehat{q}_{0.97}(X)\}$	$\{\widehat{q}_{0.97}(X) < \widehat{q}_{0.96}(X)\}$	$\{\widehat{q}_{0.96}(X) < \widehat{q}_{0.95}(X)\}$
single- α	-0.589 (-1.687)	-0.519 (-1.511)	-0.488 (-1.601)
multi- α (I)	-3.219 (-3.315)	-2.516 (-2.626)	-1.929 (-2.076)
multi- α (II)	-3.132 (-3.189)	-2.488 (-2.648)	-1.936 (-2.139)
multi- α (III)	-1.824 (-1.994)	-1.39 (-1.585)	-1.721 (-1.766)

Table 4: \log_{10} of the empirical estimates (and of the corresponding standard deviations) of $P(E)$, for the events E listed in the first row, for learned conditional VaR estimators in the Student- t toy-example across 32 runs.

α	0.999	0.995	0.99
NNR using true VaR	1.598 (0.196)	1.063 (0.045)	1.006 (0.022)
LR using single- α VaR	0.955 (0.015)	2.231 (1.149)	0.804 (0.429)
LR using multi- α (I) VaR	2.581 (0.505)	0.799 (0.169)	0.55 (0.098)
LR using multi- α (II) VaR	2.394 (0.396)	0.74 (0.137)	0.518 (0.081)
LR using multi- α (III) VaR	1.276 (0.355)	0.574 (0.085)	0.491 (0.042)
α	0.98	0.95	0.9
NNR using true VaR	0.968 (0.014)	0.887 (0.029)	0.816 (0.079)
LR using single- α VaR	0.492 (0.091)	0.371 (0.021)	0.383 (0.027)
LR using multi- α (I) VaR	0.43 (0.053)	0.38 (0.027)	0.403 (0.02)
LR using multi- α (II) VaR	0.416 (0.046)	0.374 (0.024)	0.396 (0.017)
LR using multi- α (III) VaR	0.45 (0.029)	0.432 (0.025)	0.432 (0.02)

Table 5: Means (standard deviations) of RMSE errors of learned conditional ES estimators against groundtruth values in the Student- t toy-example across 32 runs. The RMSE errors are normalized by division by the standard deviation of the groundtruth ES.

tive performance of the different linear regression approaches in Table 5 is explained by the relative performance of the VaR learning approaches, given that the VaR learning error contributes to the ES learning error.

7 Dynamic Initial Margin Case Study

A financial application of the quantile learning framework is the learning of a path-wise, dynamic initial margin (DIM) in the context of XVA computations (see e.g. Crépey (2022, Table 4.1) and Albanese et al. (2021, Section A.4)). Let there be given

respectively \mathbb{R}^d valued and real valued stochastic processes $S = (S_t)_{t \geq 0}$ and $\text{MtM} = (\text{MtM}_t)_{t \geq 0}$, where S is Markov and S_t represents the state of the market at time t (e.g. diffused market risk factors), whereas MtM_t represents the mark-to-market (price) of the portfolio of the bank at time t . We include in this price the cash flows cumulated up to time t , so that $\text{MtM}_{t+\delta} - \text{MtM}_t$ is $\sigma(X_s, t \leq s \leq t + \delta)$ measurable. We ignore risk-free discounting in the notation (while preserving it in the numerical experiments). The initial margin of the bank at time t at the confidence level α , denoted by IM_t , is defined as

$$\text{VaR}(\text{MtM}_{t+\delta} - \text{MtM}_t | S_t) =: \text{IM}(t, S_t).$$

Hence, having n Euler simulated paths of S and MtM , one can estimate the DIM process $(\text{IM}(t, S_t))$ at grid times, using one quantile regression with data $(X, Y)_{1:n} \equiv (S_t, \text{MtM}_{t+\delta} - \text{MtM}_t)_{1:n}$ for each grid time t . Alternatively, this DIM process can be estimated by a brute force nested Monte Carlo method detailed in the arXiv v1 preprint version of this work, taking several nights of computation time on a workstation, as opposed to a few minutes by regressions.

7.1 Numerical Results

We consider a portfolio composed of 100 interest rate swaps with randomly drawn characteristics and final maturity 10 years, assessed in the market model of Abbas-Turki, Crépey, and Saadeddine (2023, Appendix B), i.e. a multi-factor market model with 10 short-rate processes representing 10 economies and 9 cross-currency rate processes. Given that swap coupons can depend on short-rates at previous fixing dates, we also include in the regression basis the same short-rates but observed at the latest previous fixing date, which leads in total to a dimensionality of $d = 29$ for the state vector S_t at a given time $t > 0$, with 100 time steps uniformly spread between time 0 and the final maturity of the portfolio equal to 10 years. We use $2^{22} = 4194304$ simulated paths (generated in 25 seconds using the code developed in Abbas-Turki, Crépey, and Saadeddine (2023, Appendix B)) for training and 2^{14} simulated paths, independent of the former, for evaluating the nested Monte Carlo benchmark and computing the errors. We leverage the transfer learning trick used in Abbas-Turki et al. (2023, Appendix B), which consists in doing the training starting from the latest time-step and then proceeding backwards by reusing the solution obtained at each successive time-step t_{k+1} as an initialization for the learning to be done at time t_k . This allows us to use only 16 training epochs. As in the Student toy-example, we use mini-batching. The batch size is taken to be $2^{17} = 131072$, we use a learning rate of 0.001, and the rest of the Adam parameters are kept at their default values.

To illustrate that the quantile learning approach allows one to learn an entire stochastic process (dynamic initial margin), we plot the mean and 5 th/95 th percentiles of the learned IM process at each time-step for the different quantile learning schemes in Figure 2. At $t = 0$ the IM is deterministic because there is no randomness in the model yet, at $t = 10$ (last time step, i.e. final maturity of the portfolio) the IM vanishes because there are no later cash flows. The sawtooth-like behaviour in the paths of $(\text{IM}_t)_{t \geq 0}$ that is visible in the plots in Figure 2 is expected, due to the recurring cash-flows inherent to interest rate swaps (Andersen, Pykhtin, and Sokol, 2017).

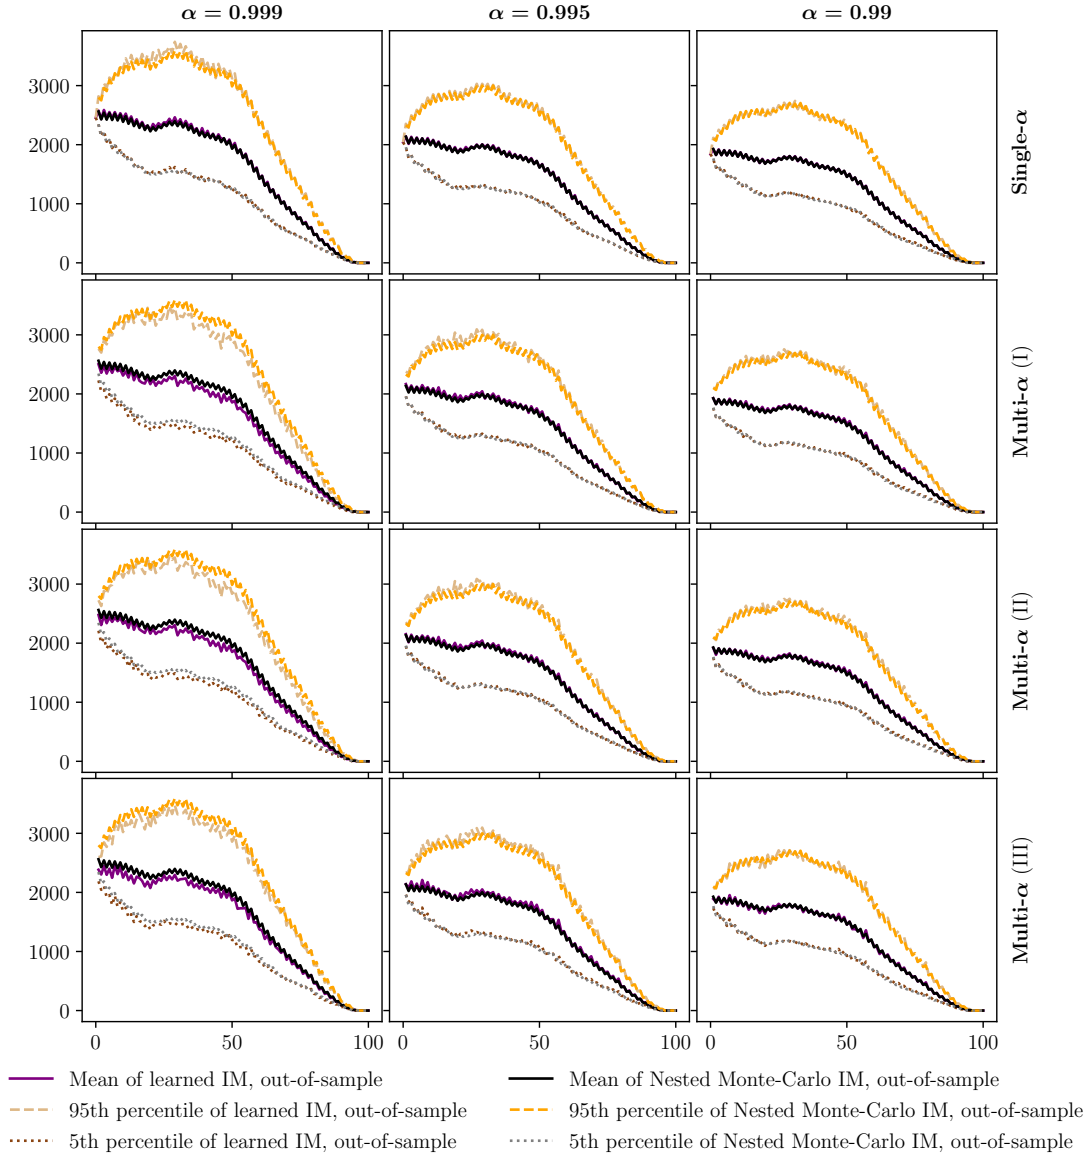


Figure 2: Mean and 5 th/95 th percentiles of both the learned and the nested Monte Carlo IM at different time steps and for different values of α and learning approaches. The learning approach used for the plots in each row is indicated on the right, and each column corresponds to one value of α which is indicated at the top of each column. Statistics are computed using out-of-sample trajectories of the diffused risk-factors, and the time steps are on the x axis.

α	0.999	0.995	0.99	0.98	0.95	0.9
multi- α (I)	0.265	0.160	0.109	0.065	0.058	0.056
multi- α (II)	0.261	0.155	0.107	0.066	0.057	0.056
multi- α (III)	0.128	0.185	0.102	0.133	0.116	0.074
Single- α	0.134	0.074	0.070	0.056	0.066	0.065

Table 6: RMSE errors of learned IM_t estimators against nested Monte Carlo estimators, for $t = 2.5$ years. Errors are normalized by division by the standard deviation of the nested Monte Carlo benchmark.

α	0.999	0.995	0.99	0.98	0.95	0.9
multi- α (I)	0.204	0.166	0.131	0.072	0.061	0.069
multi- α (II)	0.212	0.162	0.127	0.072	0.062	0.069
multi- α (III)	0.150	0.123	0.067	0.065	0.066	0.068
Single- α	0.165	0.095	0.070	0.057	0.060	0.066

Table 7: RMSE errors of learned IM_t estimators against nested Monte Carlo estimators, for $t = 5$ years. Errors are normalized by division by the standard deviation of the nested Monte Carlo benchmark.

α	0.999	0.995	0.99	0.98	0.95	0.9
multi- α (I)	0.292	0.119	0.122	0.095	0.073	0.072
multi- α (II)	0.296	0.118	0.118	0.091	0.071	0.070
multi- α (III)	0.157	0.118	0.090	0.089	0.079	0.086
Single- α	0.119	0.088	0.082	0.068	0.061	0.061

Table 8: RMSE errors of learned IM_t estimators against nested Monte Carlo estimators, for $t = 7.5$ years. Errors are normalized by division by the standard deviation of the nested Monte Carlo benchmark.

Tables 6, 7 and 8 (using the nested Monte Carlo as a benchmark) confirm the conclusions of Table 2 regarding the competitiveness of the multi- α approaches.

Conclusion of the Numerical Experiments

The numerical experiments of Sections 6 and 7 suggest that learning multiple quantiles (multi- α (I), multi- α (II) or multi- α (III)), although counter-intuitive at first, can help better target extreme quantiles than a standard single quantile learning approach. This can be explained by the fact that multiple quantile approaches leverage the information given by nearby quantiles and thus are better at extrapolating at the extremes. The multi- α (I) approach is remarkably good at ensuring, via soft-constraints on the derivative with respect to the quantile level, monotonicity (avoiding quantile crossings), in cases where consistency among different quantile levels is desired. Our experiments

also show that one can successfully use these quantile estimation methods in an XVA or dynamic risk calculation setting, where the computation burden may be greatly diminished by replacing nested Monte Carlo estimations by VaR and ES learnings.

As **practical take-away messages** to the reader, we would like to emphasize three algorithmic breakthroughs of the paper, namely **(i)** the a posteriori twin Monte Carlo validation trick of Algorithm 1, **(ii)** the neural network VaR to ES transfer learning trick corresponding to the linear regression case in Algorithm 4, and **(iii)** the multi- α learning schemes Algorithms 5-6.

A Value-at-Risk and Expected Shortfall Representations

In this section we recall the elicibility results behind our learning algorithms.

A cumulative distribution function (cdf) $F : \mathbb{R} \rightarrow [0, 1]$ is by definition (Stieltjes) integrable if

$$\int_{\mathbb{R}} |y| F(dy) < \infty. \quad (\text{A.1})$$

If Y is a random variable with cdf F (i.e. $\mathbb{P}[Y \leq t] = F(t)$, $t \in \mathbb{R}$), then (A.1) holds if and only if Y is \mathbb{P} integrable (the left-hand side of (A.1) is then $\mathbb{E}|Y|$).

Definition A.1. Let $F : \mathbb{R} \rightarrow [0, 1]$ be an integrable cdf and let $\alpha \in (0, 1)$. The value-at-risk (VaR) and expected shortfall (ES) of F at the confidence level α are defined respectively by

$$\text{VaR}(F) = \min F^{-1}([\alpha, 1]), \quad \text{ES}(F) = \frac{1}{1 - F(\text{VaR}(F)-)} \int_{[\text{VaR}(F), \infty)} y F(dy). \quad (\text{A.2})$$

If Y is an integrable random variable on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$, we write

$$\text{VaR}(Y) = \text{VaR}(F_Y), \quad \text{ES}(Y) = \text{ES}(F_Y),$$

where $F_Y(t) = \mathbb{P}[Y \leq t]$ is the cdf of Y .

Remark A.1. If Y is an integrable random variable, then it is easy to see that

$$\text{VaR}(Y) = \min\{t : \mathbb{P}[Y \leq t] \geq \alpha\}, \quad \text{ES}(Y) = \mathbb{E}[Y|Y \geq \text{VaR}(Y)] \quad (\text{A.3})$$

The version of (A.3) for abstract distribution functions F is clear *mutatis mutandis*.

Assumption A.2. There exists an interval $[a, b]$ where F is continuous and

$$F(a) < \alpha < F(b).$$

Modulo the additional truncation feature needed for this work, The following result corresponds to Rockafellar and Uryasev (2000).

Lemma A.1. *Let ϕ be given by (2.7). If F is an integrable distribution function satisfying Assumption A.2, then*

$$\begin{aligned} F^{-1}(\{\alpha\}) &= \operatorname{argmin}_{u \in [a, b]} \underbrace{(1 - \alpha)^{-1} \int_u^\infty (y - u)^+ F(dy) + u}_{=: es(u)} \\ &= \operatorname{argmin}_{u \in [a, b]} \underbrace{(1 - \alpha)^{-1} \int_u^\infty ((a \vee y \wedge b) - u)^+ F(dy) + u}_{=: es^{a, b}(u)} \end{aligned} \quad (\text{A.4})$$

$$\{\text{ES}(F)\} = es(F^{-1}(\{\alpha\})). \quad (\text{A.5})$$

Proof. Under Assumption A.2, $\emptyset \subsetneq F^{-1}(\{\alpha\}) \subseteq (a, b)$. Since F is continuous in $[a, b]$, the identity $\partial_u es(u) = 1 - (1 - \alpha)^{-1}(1 - F(u))$ holds for $u \in [a, b]$. It follows in particular that the continuously differentiable function es has critical points as the set of α -quantiles of F . Since es is convex, these critical points are the minimizers of es over $[a, b]$, which proves the first equality in (A.4). Likewise,

$$\begin{aligned} \partial_u es^{a, b}(u) &= 1 - (1 - \alpha)^{-1} \left(((a \vee u \wedge b) - u)^+ + \int_u^\infty \mathbf{1}_{(a \vee y \wedge b) \geq u} F(dy) \right) \\ &= 1 - (1 - \alpha)^{-1}(1 - F(u)), \text{ for } u \in [a, b]. \end{aligned}$$

Consequently, $F^{-1}(\{\alpha\})$ are the critical points of the function under the second argmin in (A.4), which are also the minimizers over $[a, b]$ due to the convexity of $es^{a, b}$. This completes the proof of (A.4).

Moreover, for any $q \in F^{-1}(\{\alpha\})$, F is constant on $[\text{VaR}(F), q)$, hence (A.2) yields

$$\text{ES}(F) = (1 - \alpha)^{-1} \int_q^\infty y F(dy) = es(u).$$

This proves (A.5). \square

We now introduce the functions Δ_F and $\Gamma_F : [a, b] \times F^{-1}(\{\alpha\}) \rightarrow [0, \infty)$ defined by

$$\Delta_F(u, q) = es(u) - es(q), \quad (\text{A.6})$$

$$\Gamma_F(u, q) = \frac{\Delta_F(u, q)}{(u - q)^2} \mathbf{1}_{(0, \infty)}(|u - q|) + \frac{\dot{F}(q)}{2(1 - \alpha)} \mathbf{1}_{\{0\}}(u - q), \quad (\text{A.7})$$

where $\dot{F}(q)$ is the density of Y at q (assumed to exist).

Lemma A.2. *Under Assumption A.2, if F is absolutely continuous in (a, b) with associated density denoted by \dot{F} , then*

$$\Delta_F(u, q) = (u - q)^2 \Gamma_F(u, q), \quad (u, q) \in [a, b] \times F^{-1}(\{\alpha\}) \quad (\text{A.8})$$

$$\Gamma_F^{-1}(\{0\}) \subseteq F^{-1}(\{\alpha\}) \times F^{-1}(\{\alpha\}) = \Delta_F^{-1}(\{0\}) \quad (\text{A.9})$$

$$\inf_{t \in (0, 1)} \dot{F}(tu + (1 - t)q) \leq 2(1 - \alpha) \Gamma_F(u, q) \leq \sup_{t \in (0, 1)} \dot{F}(tu + (1 - t)q). \quad (\text{A.10})$$

Proof. (A.8) is immediate from the definition of Γ_F . It is also clear from (A.4) that Δ_F as in (A.6) is nonnegative, hence so is in turn (A.7). (A.4) implies at once the equality of sets at the right in (A.9), while the inclusion of sets at the left in (A.9) is an immediate consequence of the definition of Γ_F and of the first equality in (A.4).

If F is differentiable in $[a, b]$, then (A.4) implies that for every $(u, q) \in [a, b] \times F^{-1}(\{\alpha\})$,

$$\Delta_F(q, q) = 0 = \partial_1 \Delta_F(q, q), \quad \partial_{1,1}^2 \Delta_F(tu + (1-t)q, q) = (1-\alpha)^{-1} \dot{F}(tu + (1-t)q),$$

where $\partial_1 \Delta_F$ and $\partial_{1,1}^2 \Delta_F$ refer to the first and second derivatives of Δ_F with respect to its first argument. Hence, by an application of Taylor's theorem with the Lagrange form of the remainder, given $(u, q) \in [a, b] \times F^{-1}(\{\alpha\})$ there exists $t_{u,q} \in (0, 1)$ for which

$$\Delta_F(u, q) = \frac{(u-q)^2}{2(1-\alpha)} \dot{F}(t_{u,q}q + (1-t_{u,q})u). \quad (\text{A.11})$$

(A.10) is then clear from (A.11). \square

In view of (A.6)-(A.7), (A.10) reflects the connection between the value of \dot{F} and the accuracy of u as an approximation of q as measured by Δ_F .

B Proofs of Main Results

B.1 Proof of Lemma 2.1

The functions $\omega \mapsto \text{VaR}(Y|X(\omega))$ and $\omega \mapsto \text{ES}(Y|X(\omega))$ are $\sigma(X)$ measurable. In fact, given $t \in \mathbb{R}$, we have

$$\{\text{VaR}(Y|X) < t\} = \cup_{n \in \mathbb{N}} \{F_{Y|X}(t - 1/n) \geq \alpha\},$$

which is a countable union of $\sigma(X)$ measurable sets (as $F_{Y|X}(y)$ is $\sigma(X)$ measurable for every fixed y). This shows the claim for $\text{VaR}(Y|X)$.

As for the $\sigma(X)$ measurability of $\text{ES}(Y|X)$, notice that the function $es : \mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$es(x, u) = \frac{1}{1 - \mu(x, (-\infty, u))} \int y \mathbf{1}_{[u, \infty)}(y) \mu(x, dy)$$

is Borel measurable (on the set where $\mu(x, (-\infty, u)) < 1$) and that

$$\text{ES}(Y|X) = es(X, \text{VaR}(Y|X)).$$

The result then follows by an application of the Doob-Dynkin lemma. \square

B.2 Proof of Lemma 2.2

Combining $\mathbb{P}[Y < q(X)|X] \leq \alpha \leq \mathbb{P}[Y \leq q(X)|X]$ with easy inequalities, we get

$$\begin{aligned} q(X) \geq 0 &\implies (1 - \alpha)q(X) \leq q(X)\mathbb{P}[Y \geq q(X)|X] \leq \mathbb{E}[Y\mathbf{1}_{Y \geq q(X)}|X] \leq \mathbb{E}[Y^+|X], \\ q(X) \leq 0 &\implies \alpha q(X) \geq q(X)\mathbb{P}[Y \leq q(X)|X] \geq \mathbb{E}[Y\mathbf{1}_{Y \leq q(X)}|X] \geq -\mathbb{E}[Y^-|X], \end{aligned}$$

which implies (2.3).

The inequality (2.4) stems from the bound $|\mathbb{E}S(Y|X)| \leq (1 - \alpha)^{-1}\mathbb{E}[|Y||X]$ that follows from Definition 2.1. These inequalities (resp. and the conditional Jensen inequality) show that the (resp. square) integrability of $q(X)$ and $s(X)$ follows from that of Y . \square

B.3 Proof of Lemma 3.1

For any $f \in \mathcal{L}_{a,b}$, (3.3) is a consequence of the definition of γ_f in (3.2) and (A.9). Now notice that, according to (2.1), the definition (3.2) of γ_f , (A.6) and (2.7),

$$\Phi(f) - \Phi(q) = \mathbb{E}[\mathbb{E}[\phi(Y, f(X)) - \phi(Y, q(X))|X]] = \mathbb{E}[\Delta_{F_Y|X}(f(X), q(X))]$$

for $f \in \mathcal{F}$, which leads via (A.8) to the equality

$$\Phi(f) - \Phi(q) = \mathbb{E}[\gamma_f(X)(f(X) - q(X))^2]. \quad (\text{B.1})$$

Together with

$$\Phi(\tilde{q}) - \Phi(q) \leq \Phi(f) - \Phi(q) \quad (\text{B.2})$$

valid for every $f \in \mathcal{F}$, this gives rise to

$$\|\gamma_{\tilde{q}}(\tilde{q} - q)\|_{\mathbb{P}_X, 2}^2 = \Phi(\tilde{q}) - \Phi(q) = \inf_{f \in \mathcal{F}} \|\gamma_f(f - q)\|_{\mathbb{P}_X, 2}^2,$$

which implies the equalities in (3.4). In addition, the upper bound

$$\begin{aligned} \Phi(\tilde{q}) - \Phi(q) &= \frac{1}{1 - \alpha} \mathbb{E}[(Y - \tilde{q}(X))^+ - (Y - q(X))^+] + \mathbb{E}[\tilde{q}(X) - q(X)] \\ &\leq \left(\frac{2 - \alpha}{1 - \alpha}\right) \inf_{f \in \mathcal{F}} \|f - q\|_{\mathbb{P}_X, 1} \end{aligned}$$

follows via an elementary estimation using

$$|w^+ - v^+| \leq |w - v| \quad (\text{B.3})$$

and the triangle inequality, together with a new application of (B.2). This shows the inequality at the end in (3.4). \square

B.4 Proof of Theorem 3.2

Before proving the theorem we introduce the following Rademacher bound.

Lemma B.1. *Let $Z_{1:n}$ be an i.i.d. sample of a random element Z of \mathcal{Z} , with Z independent of $Z_{1:n}$, and let \mathcal{H} be a family of functions $\mathcal{Z} \rightarrow \mathbb{R}$. Define, for $z \in \mathcal{Z}$, $A_{\mathcal{H}}(z) = \inf_{h \in \mathcal{H}} h(z)$, $B_{\mathcal{H}}(z) = \sup_{h \in \mathcal{H}} h(z)$, and assume that $0 < \|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}} < \infty$. For a possibly data dependent $h \in \mathcal{H}$, let $\Theta(h) = \mathbb{E}[h(Z)|Z_{1:n}]$, $\hat{\Theta}(h) = \frac{1}{n} \sum_{i=1}^n h(Z_k)$, $\hat{p} \in \operatorname{argmin}_{h \in \mathcal{H}} \Theta(h)$, $\tilde{p} \in \operatorname{argmin}_{h \in \mathcal{H}} \hat{\Theta}(h)$. Then for any $\delta \in (0, 1)$*

$$\Theta(\hat{p}) - \Theta(\tilde{p}) \leq 2\mathcal{R}_{ave}(\mathcal{H}, Z_{1:n}) + \|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}} \sqrt{\frac{2 \log \frac{2}{\delta}}{n}} \quad (\text{B.4})$$

holds with probability at least $1 - \delta$.

Proof. First,

$$\begin{aligned} \Theta(\hat{p}) - \Theta(\tilde{p}) &= \left(\Theta(\hat{p}) - \hat{\Theta}(\hat{p}) \right) + \left(\hat{\Theta}(\hat{p}) - \hat{\Theta}(\tilde{p}) \right) + \left(\hat{\Theta}(\tilde{p}) - \Theta(\tilde{p}) \right) \\ &\leq \left(\Theta(\hat{p}) - \hat{\Theta}(\hat{p}) \right) + \left(\hat{\Theta}(\tilde{p}) - \Theta(\tilde{p}) \right). \end{aligned} \quad (\text{B.5})$$

To bound from above the first difference at the second line in (B.5), we define the family \mathcal{H}' by

$$\mathcal{H}' = \frac{\mathcal{H} - A_{\mathcal{H}}}{\|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}}} = \left\{ \mathcal{Z} \ni z \mapsto \frac{h(z) - A_{\mathcal{H}}(z)}{\|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}}}; h \in \mathcal{H} \right\}.$$

It then follows from Definition 3.1 that

$$\begin{aligned} \mathcal{R}_{emp}(\mathcal{H}', Z_{1:n}) &= \mathcal{R}_{emp} \left(\frac{\mathcal{H} - A_{\mathcal{H}}}{\|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}}}, Z_{1:n} \right) \\ &= \frac{\mathcal{R}_{emp}(\mathcal{H}, Z_{1:n}) + \mathcal{R}_{emp}(A_{\mathcal{H}}, Z_{1:n})}{\|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}}} = \frac{\mathcal{R}_{emp}(\mathcal{H}, Z_{1:n})}{\|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}}}, \end{aligned} \quad (\text{B.6})$$

where the last equality follows from the fact that the Rademacher complexity of a single function set is null. By Mohri, Rostamizadeh, and Talwalkar (2018, Eqn. (3.7) and (3.13) p.31–32),

$$\sup_{h \in \mathcal{H}'} \left(\Theta(h) - \hat{\Theta}(h) \right) \leq 2\mathcal{R}_{ave}(\mathcal{H}', Z_{1:n}) + \sqrt{\frac{\log \frac{2}{\delta}}{2n}}$$

holds with probability at least $1 - \delta/2$. Multiplying both sides of the above inequality by $\|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}}$, changing the set of the sup term in the left-hand side to \mathcal{H}' and using (B.6) yields

$$\Theta(\hat{p}) - \hat{\Theta}(\hat{p}) \leq \sup_{h \in \mathcal{H}'} \left(\Theta(h) - \hat{\Theta}(h) \right) \leq 2\mathcal{R}_{ave}(\mathcal{H}, Z_{1:n}) + \|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}} \sqrt{\frac{\log \frac{2}{\delta}}{2n}} \quad (\text{B.7})$$

with probability at least $1 - \delta/2$. Since \tilde{p} does not depend on $Z_{1:n}$, the difference $\widehat{\Theta}(\tilde{p}) - \Theta(\tilde{p})$ in (B.5) can be bounded using the same argument with $\mathcal{H} = \{\tilde{p}\}$ to obtain

$$\widehat{\Theta}(\tilde{p}) - \Theta(\tilde{p}) \leq \|B_{\mathcal{H}} - A_{\mathcal{H}}\|_{P_{Z,\infty}} \sqrt{\frac{\log \frac{2}{\delta}}{2n}} \quad (\text{B.8})$$

with probability at least $1 - \delta/2$. Combining (B.5), (B.7) and (B.8) yields (B.4). \square

Back to the proof of the theorem, applying (B.1) (valid for every $f \in \mathcal{F}$) to $f = \widehat{q} \in \mathcal{F}$ yields

$$\|(\widehat{q} - q)\gamma_{\widehat{q}}\|_{P_{X,2}}^2 = \Phi(\widehat{q}) - \Phi(q) = (\Phi(\widehat{q}) - \Phi(\tilde{q})) + (\Phi(\tilde{q}) - \Phi(q)). \quad (\text{B.9})$$

The term $(\Phi(\tilde{q}) - \Phi(q))$ is handled by (3.4). To upper bound $\Phi(\widehat{q}) - \Phi(\tilde{q})$ in probability, notice that

$$(1 - \alpha)\phi(y, u) = (y - \alpha u) \vee ((1 - \alpha)u). \quad (\text{B.10})$$

If both y and u are in the range (v, w) for $\infty < v \leq w < \infty$, then a little algebra shows, via (B.10), that

$$(1 - \alpha)v \leq (1 - \alpha)\phi(y, u) \leq w - \alpha v,$$

which implies that

$$a(X) \leq \phi(Y^{a,b}, f(X)) \leq (1 - \alpha)^{-1}(b(X) - \alpha a(X)), \text{ P a.s.} \quad (\text{B.11})$$

Hence, an application of Lemma B.1 with $Z_{1:n} = (X, Y^{a,b})_{1:n}$, $\mathcal{H} = \phi(\mathcal{F}) = \{\mathcal{X} \times \mathbb{R} \ni (x, y) \mapsto \phi(y, f(x)); f \in \mathcal{F}\}$, $A_{\mathcal{H}} = a$ and $B_{\mathcal{H}} = (1 - \alpha)^{-1}(b - \alpha a)$ results in the inequality

$$\Phi(\widehat{q}) - \Phi(\tilde{q}) \leq 2\mathcal{R}_{ave}(\phi(\mathcal{F}), (X, Y^{a,b})_{1:n}) + \frac{\|b - a\|_{P_{X,\infty}} \sqrt{2 \log(2/\delta)}}{(1 - \alpha)\sqrt{n}} \quad (\text{B.12})$$

with probability at least $1 - \delta$.

Note now that, owing to Talagrand contraction lemma (Mohri et al. (2018, Lemma 5.7 p.93)), since $u \mapsto (1 - \alpha)^{-1}u^+$ is $(1 - \alpha)^{-1}$ Lipschitz, we have for any $(x, y)_{1:n} \subset (\mathcal{X} \times \mathbb{R})^n$:

$$\begin{aligned} \mathcal{R}_{emp}(\phi(\mathcal{F}), (x, y)_{1:n}) &\leq \mathcal{R}_{emp}(\{(1 - \alpha)^{-1}(y - f)^+ : f \in \mathcal{F}\}, (x, y)_{1:n}) + \mathcal{R}_{emp}(\mathcal{F}, x_{1:n}) \\ &\leq (1 - \alpha)^{-1} \mathcal{R}_{emp}(\{y - f : f \in \mathcal{F}\}, (x, y)_{1:n}) + \mathcal{R}_{emp}(\mathcal{F}, x_{1:n}) \\ &\leq (1 - \alpha)^{-1} \mathcal{R}_{emp}(\{y\}, y_{1:n}) + \left(\frac{2 - \alpha}{1 - \alpha}\right) \mathcal{R}_{emp}(\mathcal{F}, x_{1:n}) \\ &= \left(\frac{2 - \alpha}{1 - \alpha}\right) \mathcal{R}_{emp}(\mathcal{F}, x_{1:n}), \end{aligned}$$

where the last equality follows from Definition 3.1. The corresponding inequality for average complexities,

$$\mathcal{R}_{ave}(\phi(\mathcal{F}), (X, Y^{a,b})_{1:n}) \leq \left(\frac{2 - \alpha}{1 - \alpha}\right) \mathcal{R}_{ave}(\mathcal{F}, X_{1:n}), \quad (\text{B.13})$$

follows by integration with respect to the law of $(X, Y)_{1:n}$. A combination of (B.9) with (B.12), (B.13), (3.4) and the inequality $\sqrt{w^2 + v^2} \leq |w| + |v|$ yields (3.8). \square

B.5 Massart's Lemma

Lemma B.2. *Let $Z_{1:n}$ be a sequence of random variables, and consider a family of functions \mathcal{H} mapping \mathcal{Z} to \mathbb{R} , then*

$$\begin{aligned} n\mathcal{R}_{ave}(\mathcal{H}, Z_{1:n}) &\leq \varepsilon + \mathbb{E} \left[\sup_{h \in \mathcal{H}} \sqrt{2 \log(\mathcal{N}_1(\mathcal{H}, Z_{1:n}, \varepsilon/n)) \sum_{i=1}^n h^2(Z_i)} \right] \\ &\leq \varepsilon + \sqrt{2n} \|\sup_{h \in \mathcal{H}} |h(Z_1)|\|_{\mathbb{P}, \infty} \mathbb{E} \left[\sqrt{\log(\mathcal{N}_1(\mathcal{H}, Z_{1:n}, \varepsilon/n))} \right]. \end{aligned}$$

holds for any $\varepsilon > 0$.

Proof. Integrating the estimate in Rebeschini (2021, Proposition 5.2) yields the first inequality, from which the second one is deduced by upper bounding the Euclidean norm of $(h(Z_1), \dots, h(Z_n))$ with $\sqrt{n} \|\sup_{h \in \mathcal{H}} |h(Z_1)|\|_{\mathbb{P}, \infty}$. \square

B.6 Proof of Theorem 3.3

We consider the family of functions $\psi_f^B(\mathcal{G})$ defined (for fixed f) by

$$\psi_f^B(\mathcal{G}) = \{\mathcal{X} \times \mathbb{R} \ni (x, y) \mapsto \psi^B(y, f(x), g(x)) \in \mathbb{R}; g \in \mathcal{G}\}.$$

We also denote by r_f any function

$$r_f \in \operatorname{argmin}_{g \in \mathcal{L}_+^2} \mathbb{E} [(\phi(Y, f(X)) - f(X) - g(X))^2],$$

i.e., as $(Y - f(X))^+$ is square integrable, any function $r_f \in \mathcal{L}_+^2$ such that

$$r_f(X) = \mathbb{E} [(1 - \alpha)^{-1}(Y - f(X))^+ | X], \quad \mathbb{P} \text{ a.s.},$$

and we let $r_f^B : \mathcal{X} \rightarrow [0, B]$ be one of its truncated companions, in the sense that

$$r_f^B(X) = \mathbb{E} [B \wedge ((1 - \alpha)^{-1}(Y - f(X))^+) | X], \quad \mathbb{P} \text{ a.s.}$$

Let now $\|\cdot\|$ be a shorthand for the $L_{\mathbb{P}, X}^2$ seminorm on the space of square integrable measurable functions on \mathcal{X} .

Lemma B.3. *For every square integrable functions f, f' , we have*

$$\begin{aligned} \|r_f - r_f^B\| &\leq \|((1 - \alpha)^{-1}(Y - f(X))^+ - B)^+\|_{\mathbb{P}, 2}, \\ \|r_f^B - r_{f'}^B\| &\leq (1 - \alpha)^{-2} \|f - f'\|. \end{aligned}$$

Proof. The first inequality is a direct consequence of Jensen's inequality

$$\mathbb{E} [\mathbb{E} [(W - B \wedge W) | X]^2] \leq \mathbb{E} [|W - B \wedge W|^2] = \mathbb{E} [((W - B)^+)^2],$$

valid for any square integrable positive random variable W . As for the second one, notice that $B \wedge \cdot$ is 1-Lipschitz. Combining this property with (B.3) and with Jensen's inequality we obtain:

$$\begin{aligned} \|r_f^B - r_{f'}^B\|^2 &\leq \mathbb{E} \left[\mathbb{E} \left[\left| B \wedge \left(\frac{(Y - f(X))^+}{(1 - \alpha)} \right) - B \wedge \left(\frac{(Y - f'(X))^+}{(1 - \alpha)} \right) \right|^2 \middle| X \right] \right] \\ &\leq (1 - \alpha)^{-2} \|f - f'\|^2. \quad \square \end{aligned}$$

Back to the proof of Theorem 3.3, for $f \in \mathcal{L}_{a,b}$, the triangle inequality gives

$$\begin{aligned} \|\widehat{r}_f - r\| &\leq \|\widehat{r}_f - r_f^B\| + \|r_f^B - r_q^B\| + \|r_q^B - r\| \\ &\leq \|\widehat{r}_f - r_f^B\| + (1-\alpha)^{-1}\|f - q\| + \|((1-\alpha)^{-1}(Y - q(X))^+ - B)^+\|_{\mathbb{P},2}, \end{aligned} \quad (\text{B.14})$$

by Lemma B.3.

For a fixed f , applying Lemma B.1 to $Z_{1:n} = (X, Y)_{1:n}$, $\mathcal{H} = \psi_f^B(\mathcal{G})$, and to the constant functions $A_{\mathcal{H}} = 0$ and $B_{\mathcal{H}} = \sup_{g \in \mathcal{G}} \|\psi^B(Y, f(X), g(X))\|_{\mathbb{P},\infty}$ yields

$$\begin{aligned} &\|\widehat{r}_f - r_f^B\|^2 - \inf_{g \in \mathcal{G}} \|g - r_f^B\|^2 \\ &\leq 2\mathcal{R}_{ave}(\psi_f^B(\mathcal{G}), (X, Y)_{1:n}) + \frac{\sup_{g \in \mathcal{G}} \|\psi^B(Y, f(X), g(X))\|_{\mathbb{P},\infty} \sqrt{2 \log(2/\delta)}}{\sqrt{n}} \end{aligned} \quad (\text{B.15})$$

with probability at least $1 - \delta$, for any given $\delta \in (0, 1)$. In addition, Lemma B.3 together with the triangle inequality implies that

$$\inf_{g \in \mathcal{G}} \|g - r_f^B\| \leq \inf_{g \in \mathcal{G}} \|g - r\| + (1-\alpha)^{-1}\|f - q\| + \|((1-\alpha)^{-1}(Y - q(X))^+ - B)^+\|_{\mathbb{P},2}. \quad (\text{B.16})$$

In virtue of the triangle inequality and the inequality $\sqrt{w^2 + v^2} \leq |w| + |v|$, replacing (B.16) in (B.15) and combining the resulting inequality with (B.14) implies

$$\begin{aligned} \|\widehat{r}_f - r\| &\leq \inf_{g \in \mathcal{G}} \|g - r\| \\ &\quad + 2 \left((1-\alpha)^{-1}\|f - q\| + \|((1-\alpha)^{-1}(Y - q(X))^+ - B)^+\|_{\mathbb{P},2} \right) \\ &\quad + \left(2\mathcal{R}_{ave}(\psi_f^B(\mathcal{G}), (X, Y)_{1:n}) + \frac{\sup_{g \in \mathcal{G}} \|\psi^B(Y, f(X), g(X))\|_{\mathbb{P},\infty} \sqrt{2 \log(2/\delta)}}{\sqrt{n}} \right)^{1/2}. \end{aligned} \quad (\text{B.17})$$

Let us now upper bound the Rademacher complexity in the above equation. Since the square function on $[-B, B]$ has Lipschitz constant equal to $2B$, Talagrand's contraction lemma gives

$$\mathcal{R}_{ave}(\psi_f^B(\mathcal{G}), (X, Y)_{1:n}) \leq 2B\mathcal{R}_{ave}(\mathcal{G}, X_{1:n}). \quad (\text{B.18})$$

Replacing (B.18) and the bound $\sup_{g \in \mathcal{G}} \|\psi_f^B(g)(X, Y)\|_{\mathbb{P},\infty} \leq B^2$ in (B.17) implies (3.11). \square

B.7 Proof of Proposition 3.4

We have

$$\|\mathbb{P}[Y > f(X)|X] - 1 + \alpha\|_{\mathbb{P},2}^2 = \mathbb{E} \left[\mathbb{P}[Y > f(X)|X]^2 \right] + (1-\alpha)^2 - 2(1-\alpha)\mathbb{P}[Y > f(X)],$$

where

$$\mathbb{P}[Y > f(X)|X]^2 = \mathbb{P} \left[Y^{(1)} \wedge Y^{(2)} > f(X)|X \right].$$

Thus (3.16) follows. For the ES, we have $E[\phi(Y, q(X))|X] = s(X)$ (see (2.7) and (2.9)), hence

$$\|h(X) - s(X)\|_{\mathbb{P},2}^2 = \|E[Z|X]\|_{\mathbb{P},2}^2,$$

where $Z = h(X) - \phi(Y, q(X))$ satisfies by the conditional Jensen inequality:

$$\|E[Z|X]\|_{\mathbb{P},2}^2 = E[(E[Z|X])^2] \leq E[E[Z^2|X]] = E[Z^2] = \|h(X) - \phi(Y, q(X))\|_{\mathbb{P},2}^2.$$

An application of the triangular inequality yields

$$\|h(X) - \phi(Y, q(X))\|_{\mathbb{P},2} \leq \|h(X) - \phi(Y, f(X))\|_{\mathbb{P},2} + \|\phi(Y, q(X)) - \phi(Y, f(X))\|_{\mathbb{P},2}.$$

By $\frac{2-\alpha}{1-\alpha}$ Lipschitz regularity of ϕ with respect to its second argument,

$$\|\phi(Y, q(X)) - \phi(Y, f(X))\|_{\mathbb{P},2} \leq \frac{2-\alpha}{1-\alpha} \|q(X) - f(X)\|_{\mathbb{P},2}. \quad (\text{B.19})$$

Given our assumption that $\inf_{y \in (a(X), b(X))} \dot{F}_{Y|X}(y) \geq c$ holds \mathbb{P} a.s., it follows that, \mathbb{P} a.s.,

$$\sup_{t \in (\alpha \wedge (1 - \mathbb{P}[Y > f(X)|X]), \alpha \vee (1 - \mathbb{P}[Y > f(X)|X]))} \widehat{F}_{Y|X}^{-1}(t) \leq \sup_{t \in (F^{-1}(a(X)), F^{-1}(b(X)))} \widehat{F}_{Y|X}^{-1}(t) \leq \frac{1}{c}.$$

By writing $q(X) = F_{Y|X}^{-1}(\alpha)$ and $f(X) = F_{Y|X}^{-1}(1 - \mathbb{P}[Y > f(X)|X])$, applying the mean value theorem, we obtain

$$|q(X) - f(X)| = |\alpha - 1 + \mathbb{P}[Y > f(X)|X]| \widehat{F}_{Y|X}^{-1}(t)$$

for some $t \in (\alpha \wedge (1 - \mathbb{P}[Y > f(X)|X]), \alpha \vee (1 - \mathbb{P}[Y > f(X)|X]))$, \mathbb{P} a.s.. Hence

$$\|q(X) - f(X)\|_{\mathbb{P},2} \leq \frac{1}{c} \|1 - \alpha - \mathbb{P}[Y > f(X)|X]\|_{\mathbb{P},2}.$$

Replacing the above bound to (B.19) yields (3.17). By conditional independence and tower law, we get (3.18). \square

B.8 Proof of Lemma 4.1

For any $x \in \mathcal{X}$, $\tilde{x} = \begin{bmatrix} x \\ 1 \end{bmatrix}$, and $f \in \mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma)$, an application of (4.1) yields

$$|f(x)| = |W_{l+1} \sigma(W_l \sigma(\dots \sigma(W_1 \tilde{x})))| \leq |W_{l+1}|_2 |\sigma(W_l \sigma(\dots \sigma(W_1 \tilde{x})))|_2.$$

Since σ is 1-Lipschitz and positive homogeneous, $|\sigma(V)|_2 \leq |V|_2$ holds for any vector V . Successive applications of this and of (4.1) imply (4.2).

Given (4.2), integrating both sides of the inequality in Golowich et al. (2018, Theorem 1) w.r.t. the law of $X_{1:n}$ yields (4.3). \square

B.9 Proof of Theorem 4.2

By Lemma 4.1, $\mathcal{F} = \mathcal{NN}(d, 1, B_{1:(l+1)}, l, m, \sigma)$ is uniformly bounded and $\mathcal{F} \subseteq \mathcal{L}_{a,b}$. Similar computations as in Section B.8 show that for $f \in \mathcal{F}$, $x, x' \in \mathcal{X}$ and $\delta > 0$, if $|x - x'|_2 < \delta$, then

$$|f(x) - f(x')|_2 \leq \prod_{k=1}^{l+1} B_k \delta, \quad f \in \mathcal{F},$$

which implies the equicontinuity of \mathcal{F} . By the Ascoli-Arzelà theorem that then applies under the assumption (i) (see e.g. Brezis (2011, Theorem 4.25 p.111)), \mathcal{F} is compact in the space of continuous functions on \mathcal{X} . Now let $\{f_j\}$ be a sequence in \mathcal{F} that converges pointwise to $f \in \mathcal{F}$. Then the same calculations as for (B.11) yield, for any j ,

$$|\phi(Y^{a,b}, f_j(X))| \leq a(X) \wedge (1 - \alpha)^{-1} (b(X) - \alpha a(X)),$$

which is integrable by assumption (i). Hence, by dominated convergence, we have

$$\lim_{j \rightarrow \infty} \Phi(f_j) = \Phi(f), \quad \text{and} \quad \lim_{j \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \phi(Y_i^{a,b}, f_j(X_i)) = \frac{1}{n} \sum_{i=1}^n \phi(Y_i^{a,b}, f(X_i)),$$

which implies the continuity of the risk functions. This and the compactness of \mathcal{F} results in the existence of \tilde{q} and \hat{q} by an application of the Weierstrass extreme value theorem.

By combining (3.9) with the assumption (ii) and (4.3), we obtain (4.4). \square

B.10 Proof of Theorem 4.3

By Lemma 4.1, \mathcal{G} is uniformly bounded. Analogous reasoning used to prove Theorem 4.2 leads to the compactness of \mathcal{G} in the space of continuous functions on \mathcal{X} . Let g_j functions in \mathcal{G} converge pointwise to $g \in \mathcal{G}$. For any fixed $f \in \mathcal{L}$, by continuity of $z \mapsto (B \wedge ((1 - \alpha)^{-1}(y - f(x))^+ - z))^2$, we then have that

$$\lim_{j \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \psi^B(Y_i, f(X_i), g_j(X_i)) = \frac{1}{n} \sum_{i=1}^n \psi^B(Y_i, f(X_i), g(X_i)),$$

which proves the continuity of the risk $g \mapsto \hat{\Psi}_f(g)$. The existence of \hat{r}_f follows by an application of the Weierstrass extreme value theorem. \square

B.11 Proof of the Toy Model ES Formula (6.3)

By Khokhlov (2016, Eqn. (5c)), one has

$$y \dot{F}_\nu(y, \mu, \sigma) = \partial_y \left(-\frac{\nu \sigma^2 + (y - \mu)^2}{\nu - 1} \dot{F}_\nu(y, \mu, \sigma) + \mu F_\nu(y, \mu, \sigma) \right).$$

When y tends to ∞ , the cdf $F_\nu(y, \mu, \sigma)$ tends to 1 and, for $\nu > 1$, $\frac{\nu \sigma^2 + (y - \mu)^2}{\nu - 1} \dot{F}_\nu(y, \mu, \sigma)$ tends to 0, because the pdf converges to 0 faster than the polynomial of order 2 goes

to ∞ . Hence

$$\int_{y_0}^{\infty} y \dot{F}_{\nu}(y, \mu, \sigma) dy = \mu(1 - F_{\nu}(y_0, \mu, \sigma)) + \frac{\nu\sigma^2 + (y_0 - \mu)^2}{\nu - 1} \dot{F}_{\nu}(y_0, \mu, \sigma). \quad (\text{B.20})$$

For conciseness, in the following calculus, we use $y_0 = F_{Y|X}^{-1}(\alpha)$ and functions P, Q, S without the input. By (6.2) and (B.20),

$$\begin{aligned} \text{ES}(Y|X) &= \frac{1}{1 - \alpha} \int_{y_0}^{\infty} y \dot{F}_{Y|X}(y) dy \\ &= \frac{1}{(1 - \alpha)} \int_{y_0}^{\infty} \frac{y}{2} \left[\dot{F}_{\nu}(y, P - Q, S) + \dot{F}_{\nu}(y, P + Q, S) \right] dy \\ &= \frac{1}{2(1 - \alpha)} \left[(P - Q)(1 - F_{\nu}(y_0, P - Q, S)) + \frac{\nu S^2 + (y_0 - P + Q)^2}{\nu - 1} \dot{F}_{\nu}(y_0, P - Q, S) \right. \\ &\quad \left. + (P + Q)(1 - F_{\nu}(y_0, P + Q, S)) + \frac{\nu S^2 + (y_0 - P - Q)^2}{\nu - 1} \dot{F}_{\nu}(y_0, P + Q, S) \right] \\ &= \frac{1}{2(1 - \alpha)} \left[2P - 2PF_{Y|X}(y_0) + Q \times (F_{\nu}(y_0, P - Q, S) - F_{\nu}(y_0, P + Q, S)) + \right. \\ &\quad \left. 2 \frac{\nu S^2 + (y_0 - P)^2 + Q^2}{\nu - 1} \dot{F}_{Y|X}(y_0) + 2 \frac{(y_0 - P)Q}{\nu - 1} (\dot{F}_{\nu}(y_0, P - Q, S) - \dot{F}_{\nu}(y_0, P + Q, S)) \right], \end{aligned}$$

by (6.1) and (6.2) again. Reminding that $y_0 = F_{Y|X}^{-1}(\alpha)$, we obtain (6.3). \square

References

- Abbas-Turki, L., S. Crépey, B. Li, and B. Saadeddine (2024). An explicit scheme for pathwise XVA computations. arXiv:2401.13314.
- Abbas-Turki, L., S. Crépey, and B. Saadeddine (2023). Pathwise CVA regressions with oversimulated defaults. *Mathematical Finance* 33(2), 274–307.
- Albanese, C., S. Crépey, R. Hoskinson, and B. Saadeddine (2021). XVA analysis from the balance sheet. *Quantitative Finance* 21(1), 99–123.
- Andersen, L., M. Pykhtin, and A. Sokol (2017). Rethinking the margin period of risk. *Journal of Credit Risk* 13(1), 1–45.
- Barrera, D. and E. Gobet (2019). Quantitative bounds for concentration of measure inequalities and empirical regression: the independent case. *Journal of Complexity* 52, 45–81.
- Bondell, H., B. Reich, and H. Wang (2010). Noncrossing quantile regression curve estimation. *Biometrika* 97(4), 825–838.
- Brezis, H. (2011). *Functional analysis, Sobolev spaces and partial differential equations*. Springer.

- Broadie, M., Y. Du, and C. Moallemi (2015). Risk estimation via regression. *Operations Research* 63(5), 1077–1097.
- Cannon, A. J. (2018). Non-crossing nonlinear regression quantiles by monotone composite quantile regression neural network, with application to rainfall extremes. *Stochastic environmental research and risk assessment* 32(11), 3207–3225.
- Cont, R., M. Cucuringu, R. Xu, and C. Zhang (2023). Tail-gan: Learning to simulate tail risk scenarios. arXiv:2203.01664v3.
- Crépey, S. (2022). Positive XVAs. *Frontiers of Mathematical Finance* 1(3), 425–465.
- Dimitriadis, T. and S. Bayer (2019). A joint quantile and expected shortfall regression framework. *Electronic Journal of Statistics* 13(1), 1823–1871.
- Fissler, T. and J. Ziegel (2016). Higher order elicibility and Osband’s principle. *The Annals of Statistics* 44(4), 1680–1707.
- Fissler, T., J. Ziegel, and T. Gneiting (2015, December). Expected shortfall is jointly elicitable with value at risk—implications for backtesting. *Risk Magazine*, January.
- Gasthaus, J., K. Benidis, Y. Wang, S. S. Rangapuram, D. Salinas, V. Flunkert, and T. Januschowski (2019). Probabilistic forecasting with spline quantile function rnns. In *The 22nd international conference on artificial intelligence and statistics*, pp. 1901–1910. Proceedings of Machine Learning Research.
- Golowich, N., A. Rakhlin, and O. Shamir (2018, 06–09 Jul). Size-independent sample complexity of neural networks. In *Proceedings of the 31st Conference On Learning Theory*, Volume 75, pp. 297–299. Proceedings of Machine Learning Research.
- Györfi, L., M. Kohler, A. Krzyżak, and H. Walk (2002). *A distribution-free theory of nonparametric regression*. Springer Series in Statistics.
- Hatalis, K., A. J. Lamadrid, K. Scheinberg, and S. Kishore (2017). Smooth pinball neural network for probabilistic forecasting of wind power. arXiv:1710.01720.
- He, X. (1997). Quantile curves without crossing. *The American Statistician* 51(2), 186–192.
- He, X., K. M. Tan, and W.-X. Zhou (2023). Robust estimation and inference for expected shortfall regression with many regressors. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 85(4), 1223–1246.
- Kallenberg, O. (2006). *Foundations of modern probability*. Springer.
- Khokhlov, V. (2016). Conditional value-at-risk for elliptical distributions. *Evropskij časopis ekonomiky a managementu* 2(6), 70–79.
- Kingma, D. P. and J. Ba (2014). Adam: A method for stochastic optimization. arXiv:1412.6980.

- Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis* 91(1), 74–89.
- Koenker, R. (2017). Quantile regression: 40 years on. *Annual Review of Economics* 9, 155–176.
- Koenker, R. and B. J. Park (1996). An interior point algorithm for nonlinear quantile regression. *Journal of Econometrics* 71(1-2), 265–283.
- Meinshausen, N. and G. Ridgeway (2006). Quantile regression forests. *Journal of Machine Learning Research* 7(6), 983–999.
- Mohri, M., A. Rostamizadeh, and A. Talwalkar (2018). *Foundations of machine learning* (second ed.). MIT press.
- Moon, S. J., J.-J. Jeon, J. S. H. Lee, and Y. Kim (2021). Learning multiple quantiles with neural networks. *Journal of Computational and Graphical Statistics* 30(4), 1238–1248.
- Padilla, O. H. M., W. Tansey, and Y. Chen (2022). Quantile regression with ReLU networks: Estimators and minimax rates. *Journal of Machine Learning Research* 23(247), 1–42.
- Rebeschini, P. (2021). Covering numbers bounds for Rademacher complexity. chaining. Lecture note. Available at <https://www.stats.ox.ac.uk/~rebeschi/teaching/AFoL/22/material/lecture05.pdf>.
- Reiss, R.-D. (2012). *Approximate distributions of order statistics: with applications to nonparametric statistics*. Springer science & business media.
- Rockafellar, R. and S. Uryasev (2000). Optimization of conditional value-at-risk. *Journal of risk* 2, 21–42.
- Rockafellar, R. T. and J. O. Royset (2013). Superquantiles and their applications to risk, random variables, and regression. In *Theory Driven by Influential Applications*, pp. 151–167. Informs.
- Rodrigues, F. and F. C. Pereira (2020). Beyond expectation: Deep joint mean and quantile regression for spatiotemporal problems. *IEEE transactions on neural networks and learning systems* 31(12), 5377–5389.
- Sangnier, M., O. Fercoq, and F. d’Alché Buc (2016). Joint quantile regression in vector-valued RKHSs. In *Neural Information Processing Systems*, Volume 29.
- Shen, G., Y. Jiao, Y. Lin, J. L. Horowitz, and J. Huang (2021). Deep quantile regression: Mitigating the curse of dimensionality through composition. arXiv:2107.04907.
- Takeuchi, I., Q. Le, T. Sears, and A. Smola (2006). Nonparametric quantile estimation. *Journal of machine learning research* 7(7), 1231–1264.