



Prediction of quantiles by statistical learning and application to GDP forecasting

Pierre Alquier, Xiaoyin Li

► To cite this version:

Pierre Alquier, Xiaoyin Li. Prediction of quantiles by statistical learning and application to GDP forecasting. The Fifteenth International Conference on Discovery Science (DS 2012), Oct 2012, Lyon, France. pp.22-36. hal-00671982v3

HAL Id: hal-00671982

<https://hal.science/hal-00671982v3>

Submitted on 6 Aug 2012

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.

Prediction of quantiles by statistical learning and application to GDP forecasting

Pierre Alquier^{1,3} and Xiaoyin Li²

¹ LPMA (Université Paris 7)
175, rue du Chevaleret
75013 Paris FRANCE
alquier@math.jussieu.fr
<http://alquier.ensae.net/>

² Laboratoire de Mathématiques (Université de Cergy-Pontoise)
UCP site Saint-Martin, 2 boulevard Adolphe Chauvin
95000 Cergy-Pontoise, FRANCE

xiaoyin.li@u-cergy.fr
³ CREST (ENSAE)

Abstract. In this paper, we tackle the problem of prediction and confidence intervals for time series using a statistical learning approach and quantile loss functions. In a first time, we show that the Gibbs estimator is able to predict as well as the best predictor in a given family for a wide set of loss functions. In particular, using the quantile loss function of [1], this allows to build confidence intervals. We apply these results to the problem of prediction and confidence regions for the French Gross Domestic Product (GDP) growth, with promising results.

Keywords: Statistical learning theory, time series, quantile regression, GDP forecasting, PAC-Bayesian bounds, oracle inequalities, weak dependence, confidence intervals, business surveys.

1 Introduction

Motivated by economics problems, the prediction of time series is one of the most emblematic problems of statistics. Various methodologies are used that come from such various fields as parametric statistics, statistical learning, computer science or game theory.

In the parametric approach, one assumes that the time series is generated from a parametric model, e.g. ARMA or ARIMA, see [2, 3]. It is then possible to estimate the parameters of the model and to build confidence intervals on the prevision. However, such an assumption is unrealistic in most applications.

In the statistical learning point of view, one usually tries to avoid such restrictive parametric assumptions - see, e.g., [4, 5] for the online approach dedicated to the prediction of individual sequences, and [6–8] for the batch approach. However, in this setting, a few attention has been paid to the construction of confidence intervals or to any quantification of the precision of the prediction. This

is a major drawback in many applications. Notice however that [9] proposed to minimize the cumulative risk corresponding to the quantile loss function defined by [1]. This led to asymptotically correct confidence intervals.

In this paper, we propose to adapt this approach to the batch setting and provide nonasymptotic results. We also apply these results to build quarterly prediction and confidence regions for the French Gross Domestic Product (GDP) growth. Our approach is the following. We assume that we are given a set of basic predictors - this is a usual approach in statistical learning, the predictors are sometimes referred as “experts”, e.g. [4]. Following [8], we describe a procedure of aggregation, usually referred as Exponentially Weighed Aggregate (EWA), [10, 11], or Gibbs estimator, [12, 13]. It is interesting to note that this procedure is also related to aggregations procedure in online learning as the weighted majority algorithm of [14], see also [15]. We give a PAC-Bayesian inequality that ensures optimality properties for this procedure. In a few words, this inequality claims that our predictor performs as well as the best basic predictor up to a remainder of the order \mathcal{K}/\sqrt{n} where n is the number of observations and \mathcal{K} measures the complexity of the set of basic predictors. This result is very general, two conditions will be required: the time series must be weakly dependent in a sense that we will make more precise in Section 4, and loss function must be Lipschitz. This includes, in particular, the quantile loss functions. This allows us to apply this result to our problem of economic forecasting.

The paper is organized as follows: Section 2 provides notations used in the whole paper. We give a definition of the Gibbs estimator in Section 3. The PAC-Bayesian inequality is given in Section 4, and the application to GDP forecasting in Section 5. Finally, the proof of all the theorems are given in Section 7.

2 The context

Let us assume that we observe X_1, \dots, X_n from a \mathbb{R}^p -valued stationary time series $X = (X_t)_{t \in \mathbb{Z}}$ defined on $(\Omega, \mathcal{A}, \mathbb{P})$. From now, $\|\cdot\|$ will denote the Euclidian norm on \mathbb{R}^p . Fix an integer k and let us assume that we are given a family of predictors $\{f_\theta : (\mathbb{R}^p)^k \rightarrow \mathbb{R}^p, \theta \in \Theta\}$: for any θ and any t , f_θ applied to the last past values $(X_{t-1}, \dots, X_{t-k})$ is a possible prediction of X_t . For the sake of simplicity, let us put for any $t \in \mathbb{Z}$ and any $\theta \in \Theta$,

$$\hat{X}_t^\theta = f_\theta(X_{t-1}, \dots, X_{t-k}).$$

We also assume that $\theta \mapsto f_\theta$ is linear. Note that we may want to include parametric models as well as non-parametric prediction. In order to deal with various family of predictors, we propose a model-selection type approach: $\Theta = \cup_{j=1}^m \Theta_j$.

Example 1 *We deal with only one model, $m = 1$ and $\Theta = \Theta_1$. We put $\theta = (\theta_0, \theta_1, \dots, \theta_k) \in \Theta = \mathbb{R}^{k+1}$ and define the linear autoregressive predictors*

$$f_\theta(X_{t-1}, \dots, X_{t-k}) = \theta_0 + \sum_{j=1}^k \theta_j X_{t-j}.$$

Example 2 We may generalize the previous example to non-parametric auto-regression, for example using a dictionary of functions $(\mathbb{R}^p)^k \rightarrow \mathbb{R}^p$, say $(\varphi_i)_{i=0}^\infty$. Then we can fix $m = n$, and take $\theta = (\theta_1, \dots, \theta_\ell) \in \Theta_j = \mathbb{R}^j$ and

$$f_\theta(X_{t-1}, \dots, X_{t-k}) = \sum_{i=1}^j \theta_i \varphi_i(X_{t-1}, \dots, X_{t-k}).$$

We now define a quantitative criterion to evaluate the quality of the predictions. Let ℓ be a loss function, we will assume that ℓ satisfies the following assumption.

Assumption LipLoss: ℓ is given by: $\ell(x, x') = g(x - x')$ for some convex function g satisfying $g \geq 0$, $g(0) = 0$ and g is K -Lipschitz.

Definition 1 We put, for any $\theta \in \Theta$, $R(\theta) = \mathbb{E} \left[\ell \left(\hat{X}_t^\theta, X_t \right) \right]$.

Note that because of the stationarity, $R(\theta)$ does not depend on t .

Example 3 A first example is $\ell(x, x') = \|x - x'\|$. In this case, the Lipschitz constant K is 1. This example was studied in detail in [8]. In [6, 7], the loss function is the quadratic loss $\ell(x, x') = \|x - x'\|^2$. Note that it also satisfies our Lipschitz condition, but only if we assume that the time series is bounded.

Example 4 When the time-series is real-valued, we can use a quantile loss function. The class of quantile loss functions is defined as

$$\ell_\tau(x, y) = \begin{cases} \tau(x - y), & \text{if } x - y > 0 \\ -(1 - \tau)(x - y), & \text{otherwise} \end{cases}$$

where $\tau \in (0, 1)$. It is motivated by the following remark: if U is a real-valued random variable, then any value t^* satisfying $\mathbb{P}(U \leq t^*) = \tau$ is a minimizer of $t \mapsto \mathbb{E}(\ell_\tau(X - t))$; such a value is called quantile of order τ of U . This loss function was introduced by [1], see [16] for a survey. Recently, [17] used it in the context of high-dimensional regression, and [9] in learning problems.

3 Gibbs estimator

In order to introduce the Gibbs estimator, we first define the empirical risk.

Definition 2 For any $\theta \in \Theta$, $r_n(\theta) = \frac{1}{n-k} \sum_{i=k+1}^n \ell \left(\hat{X}_i^\theta, X_i \right)$.

Let \mathcal{T} be a σ -algebra on Θ and \mathcal{T}_ℓ be its restriction to Θ_ℓ . Let $\mathcal{M}_+^1(\Theta)$ denote the set of all probability measures on (Θ, \mathcal{T}) , and $\pi \in \mathcal{M}_+^1(\Theta)$. This probability measure is usually called the *prior*. It will be used to control the complexity of the set of predictors Θ .

Remark 1 In the case where $\Theta = \cup_j \Theta_j$ and the Θ_j are disjoint, we can write $\pi(d\theta) = \sum_{j=1}^m \mu_j \pi_j(d\theta)$ where $\mu_j := \pi(\Theta_j)$ and $\pi_j(d\theta) := \pi(d\theta) \mathbf{1}_{\Theta_j}(\theta) / \mu_j$. Here π_j can be interpreted as a prior probability measure inside the model Θ_j and that the μ_j as a prior probability measure between the models.

Definition 3 (Gibbs estimator) We put, for any $\lambda > 0$,

$$\hat{\theta}_\lambda = \int_{\Theta} \theta \hat{\rho}_\lambda(d\theta), \text{ where } \hat{\rho}_\lambda(d\theta) = \frac{e^{-\lambda r_n(\theta)} \pi(d\theta)}{\int e^{-\lambda r_n(\theta')} \pi(d\theta')}.$$

The choice of the parameter λ is discussed in the next sections.

4 Theoretical results

In this section, we provide a PAC-Bayesian oracle inequality for the Gibbs estimator. PAC-Bayesian bounds were introduced in [18, 19], see [12, 13, 20–22] for more recent advances. The idea is that the risk of the Gibbs estimator will be close to $\inf_{\theta} R(\theta)$ up to a small remainder. More precisely, we upper-bound it by

$$\inf_{\rho \in \mathcal{M}_+^1(\Theta)} \left\{ \int R(\theta) \rho(d\theta) + \text{remainder}(\rho, \pi) \right\}.$$

To establish such a result, we need some hypothesis. The first hypothesis concerns the type of dependence of the process, it uses the $\theta_{\infty, n}(1)$ -coefficients of [23]. Such a condition is also used in [8], and is more general than the mixing conditions used in [7, 6].

Assumption WeakDep: we assume that the distribution \mathbb{P} is such that almost surely, $\|X_0\|_\infty \leq \mathcal{B} < \infty$, and that there is a constant \mathcal{C} with $\theta_{\infty, k}(1) \leq \mathcal{C} < \infty$ for any k . We remind that for any σ -algebra $\mathfrak{S} \subset \mathcal{A}$, for any $q \in \mathbb{N}$, for any $(\mathbb{R}^p)^q$ -valued random variable Z defined on $(\Omega, \mathcal{A}, \mathbb{P})$, we define

$$\theta_\infty(\mathfrak{S}, Z) = \sup_{f \in A_1^q} \left\| \mathbb{E}[f(Z)|\mathfrak{S}] - \mathbb{E}[f(Z)] \right\|_\infty$$

where

$$A_1^q = \left\{ f : (\mathbb{R}^p)^q \rightarrow \mathbb{R}, \quad \frac{|f(z_1, \dots, z_q) - f(z'_1, \dots, z'_q)|}{\sum_{j=1}^q \|z_j - z'_j\|} \leq 1 \right\},$$

and

$$\theta_{\infty, k}(1) := \sup \{ \theta_\infty(\sigma(X_t, t \leq p), (X_{j_1}, \dots, X_{j_\ell})), \quad p < j_1 < \dots < j_\ell, 1 \leq \ell \leq k \}.$$

Intuitively, these coefficients provide a quantification of the dependence between the past and the future of the time series. The sequence $\theta_{\infty, k}(1)$ is growing with k , but, when X_k behaves almost like a random variable independent from X_0, X_{-1}, \dots , the sequence is bounded. Examples of processes satisfying **WeakDep** are provided in [8, 23]. It includes processes of the form $X_t = H(\xi_t, \xi_{t-1}, \xi_{t-2}, \dots)$ where the ξ_t are iid and bounded and H satisfies a Lipschitz condition, in particular, ARMA processes with bounded innovations. It also includes uniform φ -mixing processes (see [24]) and some dynamical systems.

Assumption Lip: we assume that there is a constant $L > 0$ such that for any $\theta \in \Theta$, there are coefficients $a_j(\theta)$ for $1 \leq j \leq k$ satisfying, for any x_1, \dots, x_k and y_1, \dots, y_k ,

$$\|f_\theta(x_1, \dots, x_k) - f_\theta(y_1, \dots, y_k)\| \leq \sum_{j=1}^k a_j(\theta) \|x_j - y_j\|, \text{ with } \sum_{j=1}^k a_j(\theta) \leq L.$$

Theorem 1 (PAC-Bayesian Oracle Inequality) *Let us assume that assumptions **LipLoss**, **WeakDep** and **Lip** are satisfied. Then, for any $\lambda > 0$, for any $\varepsilon > 0$, with probability at least $1 - \varepsilon$,*

$$R(\hat{\theta}_\lambda) \leq \inf_{\rho \in \mathcal{M}_+^1(\Theta)} \left[\int R(\theta) \rho(d\theta) + \frac{2\lambda\kappa^2}{n(1 - \frac{k}{n})^2} + \frac{2\mathcal{K}(\rho, \pi) + 2\log(\frac{2}{\varepsilon})}{\lambda} \right]$$

where $\kappa = \kappa(K, L, \mathcal{B}, \mathcal{C}) := K(1+L)(\mathcal{B}+\mathcal{C})/\sqrt{2}$ and \mathcal{K} is the Kullback divergence, given by $\mathcal{K}(\rho, \pi) = \int \log \left[\frac{d\rho}{d\pi}(\theta) \right] \rho(d\theta)$ if $\rho \ll \pi$ and $+\infty$ otherwise.

The choice of λ in practice is a hard problem. In [25, 13] a general method is proposed to optimize the bound with respect to λ . However, while adapted in the iid case, this method is more difficult to use in the context of time series as it would require the knowledge of κ . However, some empirical calibration seems to give good results, as shown in Section 5.

At the price of a more technical analysis, this result can be extended to the case where the X_t are not assumed to be bounded: the results in [8] require subGaussian tails for X_t , but suffer a $\log(n)$ loss in the learning rate.

5 Application to French GDP and quantile prediction

5.1 Uncertainty in GDP forecasting

Every quarter t , economic forecasters at INSEE⁴ are asked a prediction for the quarterly growth rate of the French Gross Domestic Product (GDP). Since it involve a lot of information, the “true value” of the growth rate $\log(\text{GDP}_t/\text{GDP}_{t-1})$ is only known after two years, but *flash estimates* of the growth rate, say ΔGDP_t , are published 45 days after the end of the current quarter t . One of the most relevant economic pieces of information available at time t to the forecaster, apart from past GDP observations, are *business surveys*. Indeed, they are a rich source of information, for at least two reasons. First, they are rapidly available, on a monthly basis. Moreover, they provide information coming directly from the true economic decision makers.

A business survey is traditionally a fixed questionnaire of ten questions sent monthly to a panel of companies. This process is described in [26]. INSEE publishes a composite indicator called the *French business climate indicator*: it summarises information of the whole survey. This indicator is defined in [27], see also

⁴ *Institut National de la Statistique et des Etudes Economiques*, the French national bureau of statistics, <http://www.insee.fr/>

[28]. All these values are available from the INSEE website. Note that a quite similar approach is used in other countries, see also [29] for a prediction of the European Union GDP based on EUROSTATS data.

It is however well known among economic forecasters that interval confidence or density forecasts are to be given with the prediction, in order to provide an idea of the uncertainty of the prediction. The ASA and the NBER started using density forecasts in 1968, see [30, 31] for historical surveys on density forecasting. The Central Bank of England and INSEE, among others, provide their prediction with a “fan chart”, [32]. However, it is interesting to note that the methodology used is often very crude, see the criticism in [33, 34]. For example, until 2012, the fan chart provided by the INSEE led to the construction of confidence intervals with constant length. But there is an empirical evidence that it is more difficult to forecast economic quantities during crisis (e.g. the subprime crisis in 2008). The Central Bank of England fan chart is not reproducible as it includes subjective information. Recently, [33] proposed a clever density forecasting method based on quantile regressions that gives satisfying results in practice. However, this method did not receive any theoretical support up to our knowledge.

Here, we use the Gibbs estimator to build a forecasting of ΔGDP_t , using the quantile loss function. This allows to return a prediction: the forecasted median, for $\tau = 0.5$, that is theoretically supported. This also allows to provide confidence intervals corresponding to various quantiles.

5.2 Application of Theorem 1

At each quarter t , the objective is to predict the flash estimate of GDP growth, ΔGDP_t . As described previously, the available information is $\Delta\text{GDP}_{t'}$ for $t' < t$ and $I_{t'}$ for $t' < t$, where for notational convenience, I_{t-1} is the climate indicator available to the INSEE at time t (it is the mean of the climate indicator at month 3 of quarter $t-1$ and at month 1 and 2 of quarter t). The observation period is 1988-Q1 (1st quarter of 1988) to 2011-Q3.

We define $X_t = (\Delta\text{GDP}_t, I_t)' \in \mathbb{R}^2$. As we are not interested by the prevision of I_t but only by the prediction of the GDP growth, the loss function will only take into account ΔGDP_t . We use the quantile loss function of Example 4:

$$\begin{aligned} \ell_\tau((\Delta\text{GDP}_t, I_t), (\Delta'\text{GDP}_t, I'_t)) \\ = \begin{cases} \tau (\Delta\text{GDP}_t - \Delta'\text{GDP}_t), & \text{if } \Delta\text{GDP}_t - \Delta'\text{GDP}_t > 0 \\ -(1 - \tau) (\Delta\text{GDP}_t - \Delta'\text{GDP}_t), & \text{otherwise.} \end{cases} \end{aligned}$$

To keep in mind that the risk depends on τ , we add a subscript τ in the notation $R^\tau(\theta) := \mathbb{E}[\ell_\tau(\Delta\text{GDP}_t, f_\theta(X_{t-1}, X_{t-2}))]$. We also let r_n^τ denote the associated empirical risk. Following [33, 35] we consider predictors of the form:

$$f_\theta(X_{t-1}, X_{t-2}) = \theta_0 + \theta_1 \Delta\text{GDP}_{t-1} + \theta_2 I_{t-1} + \theta_3 (I_{t-1} - I_{t-2}) | I_{t-1} - I_{t-2}| \quad (1)$$

where $\theta = (\theta_0, \theta_1, \theta_2, \theta_3) \in \Theta(B)$. For any $B > 0$ we define

$$\Theta(B) = \left\{ \theta = (\theta_0, \theta_1, \theta_2, \theta_3) \in \mathbb{R}^4, \|\theta\|_1 = \sum_{i=0}^3 |\theta_i| \leq B \right\}.$$

These predictors of Equation 1 correspond to the model used in [33] for forecasting, one of the conclusions of [33, 35] is that these family of predictors allow to obtain a forecasting as precise as the INSEE one.

For technical reason that will become clear in the proofs, if one wants to achieve a prediction performance comparable to the best $\theta \in \Theta(B)$, it is more convenient to define the prior π as the uniform probability distribution on some slightly larger set, e.g. $\Theta(B+1)$. We will let Π_B denote this distribution. We let $\hat{\rho}_{B,\lambda}^\tau$ and $\hat{\theta}_{B,\lambda}^\tau$ denote respectively the associated aggregation distribution and the associated estimator, defined in Definition 3.

In this framework, Assumption **Lip** is satisfied with $L = B + 1$, and the loss function is K -Lipschitz with $K = 1$ so Assumption **LipLoss** is also satisfied.

Theorem 2 *Let us fix $\tau \in (0, 1)$. Let us assume that Assumption **WeakDep** is satisfied, and that $n \geq \max(10, \kappa^2/(3\mathcal{B}^2))$. Let us fix $\lambda = \sqrt{3n}/\kappa$. Then, with probability at least $1 - \varepsilon$ we have*

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\theta \in \Theta(B)} \left\{ R^\tau(\theta) + \frac{2\sqrt{3}\kappa}{\sqrt{n}} \left[2.25 + \log \left(\frac{(B+1)\mathcal{B}\sqrt{n}}{\kappa} \right) + \frac{\log(\frac{1}{\varepsilon})}{3} \right] \right\}.$$

The choice of λ proposed in the theorem may be a problem as in practice we will not know κ . Note that from the proof, it is obvious that in any case, for n large enough, when $\lambda = \sqrt{n}$ we still have a bound

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\theta \in \Theta(B)} \left\{ R^\tau(\theta) + \frac{C(B, \mathcal{B}, \kappa, \varepsilon)}{\sqrt{n}} \right\}.$$

However, in practice, we will work in an online setting: at each date t we compute the Gibbs estimator based on the observations from 1 to t and use it to predict the GDP and its quantiles at time $t + 1$. Let $\hat{\theta}_{B,\lambda}^\tau[t]$ denote this estimator. We propose the following empirical approach: we define a set of values $\Lambda = \{2^k, k \in \mathbb{N}\} \cap \{1, \dots, n\}$. At each step t , we compute $\hat{\theta}_{B,\lambda}^\tau[t]$ for each $\lambda \in \Lambda$ and use for prediction $\hat{\theta}_{B,\lambda(t)}^\tau[t]$ where $\lambda(t)$ is defined by

$$\lambda(t) = \arg \min_{\lambda \in \Lambda} \sum_{j=3}^{t-1} \ell_\tau(\Delta GDP_j, f_{\hat{\theta}_{B,\lambda}^\tau[j]}(X_{j-1}, X_{j-2})),$$

namely, the value that is currently the best for online prediction. This choice leads to good numerical results.

In practice, the choice of B has less importance. As soon as B is large enough, simulation shows that the estimator does not really depend on B , only the theoretical bound does. As a consequence we take $B = 100$ in our experiments.

5.3 Implementation

We use the importance sampling method to compute $\hat{\theta}_{B,\lambda}^\tau[t]$, see, e.g., [36]. We draw an iid sample T_1, \dots, T_N of vectors in \mathbb{R}^4 , from the distribution $\mathcal{N}(\hat{\theta}^\tau, vI)$ where $v > 0$ and $\hat{\theta}^\tau$ is simply the τ -quantile regression estimator of θ in (1), as computed by the “quantile regression package” of the R software [37]. Let $g(\cdot)$ denote the density of this distribution. By the law of large numbers we have

$$\sum_{i=1}^N \frac{T_i \exp[-\lambda r_t(T_i)] \mathbf{1}_{\Theta(B+1)}(T_i)}{g(T_i) \sum_{j=1}^N \frac{\exp[-\lambda r_t(T_j)] \mathbf{1}_{\Theta(B+1)}(T_j)}{g(T_j)}} \xrightarrow[N \rightarrow \infty]{a.s.} \hat{\theta}_{B,\lambda}^\tau[t].$$

Remark that this is particularly convenient as we only simulate the sample T_1, \dots, T_N once and we can use the previous formula to approximate $\hat{\theta}_{B,\lambda}^\tau[t]$ for several values of τ .

5.4 Results

The results are shown in Figure 1 for prediction, $\tau = 0.5$, in Figure 2 for confidence interval of order 50%, i.e. $\tau = 0.25$ and $\tau = 0.75$ (left) and for confidence interval of order 90%, i.e. $\tau = 0.05$ and $\tau = 0.95$ (right). We report only the results for the period 2000-Q1 to 2011-Q3 (using the period 1988-Q1 to 1999-Q4 for learning).




Fig. 1. French GDP online prediction using the quantile loss function with $\tau = 0.5$.

Note that we can compare the ability of our predictor $\hat{\theta}_{B,\lambda}^{0.5}$ with the predictor used in [35] that relies on a least square estimation of (1), that we will denote by $\hat{\theta}^*$. Interestingly, both are quite similar but $\hat{\theta}_{B,\lambda}^{0.5}$ is a bit more precise. We remind that

$$\begin{aligned} \text{mean abs. pred. error} &= \frac{1}{n} \sum_{t=1}^n \left| \Delta GDP_t - f_{\hat{\theta}_{B,\lambda(t)}^{0.5}}[t](X_{t-1}, X_{t-2}) \right| \\ \text{mean quad. pred. error} &= \frac{1}{n} \sum_{t=1}^n \left[\Delta GDP_t - f_{\hat{\theta}_{B,\lambda(t)}^{0.5}}[t](X_{t-1}, X_{t-2}) \right]^2. \end{aligned}$$

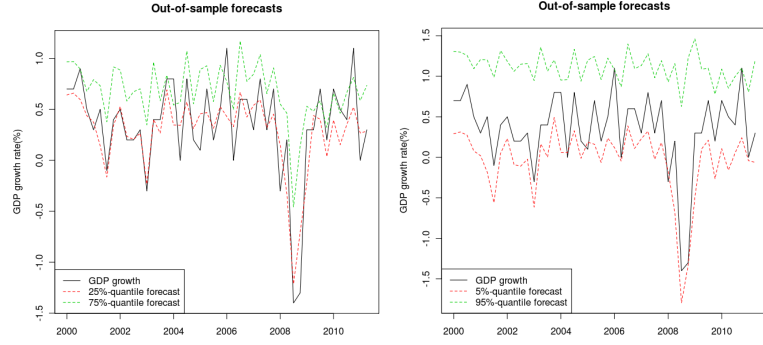


Fig. 2. French GDP online 50%-confidence intervals (left) and 90%-confidence intervals (right).

Predictor	Mean absolute prevision error	Mean quadratic prevision error
$\hat{\theta}_{B,\lambda}^{0.5}$	0.22360	0.08033
$\hat{\theta}_\star$	0.24174	0.08178

We also report the frequency of realizations of the GDP falling above the predicted τ -quantile for each τ . Note that this quantity should be close to τ .

Estimator	Frequency
$\hat{\theta}_{B,\lambda}^{0.05}$	0.065
$\hat{\theta}_{B,\lambda}^{0.25}$	0.434
$\hat{\theta}_{B,\lambda}^{0.5}$	0.608
$\hat{\theta}_{B,\lambda}^{0.75}$	0.848
$\hat{\theta}_{B,\lambda}^{0.95}$	0.978

As the INSEE did, we miss the value of the 2008 crisis. However, it is interesting to note that our confidence interval shows that our prediction at this date is less reliable than the previous ones: so, at this time, the forecaster could have been aware of some problems in their predictions.

6 Conclusion

We proposed some theoretical results to extend learning theory to the context of weakly dependent time series. The method showed good results on an application to GDP forecasting. A next step will be to give theoretical results on the online risk of our method, e.g. using tools from [12, 11]. From both theoretical and practical perspective, an adaptation with respect to the dependence coefficient $\theta_{\infty,n}(1)$ would also be really interesting but is probably a more difficult objective.

Acknowledgments

We deeply thank Matthieu Cornec (INSEE) who provided the data and the model in the application. We also thank Pr. Paul Doukhan (Univ. Cergy-Pontoise) for his helpful comments. Research partially supported by the “Agence Nationale pour la Recherche”, grant ANR-09-BLAN-0128 “PARCIMONIE”.

7 Proofs

7.1 Some preliminary lemmas

Our main tool is Rio’s Hoeffding type inequality [38]. The reference [38] is written in French and unfortunately, up to our knowledge, there is no English version of this result. So for the sake of completeness, we provide this result.

Lemma 1 (Rio [38]). *Let h be a function $(\mathbb{R}^p)^n \rightarrow \mathbb{R}$ such that for all $x_1, \dots, x_n, y_1, \dots, y_n \in \mathbb{R}^p$,*

$$|h(x_1, \dots, x_n) - h(y_1, \dots, y_n)| \leq \sum_{i=1}^n \|x_i - y_i\|. \quad (2)$$

Then for any $t > 0$ we have

$$\mathbb{E} \left(e^{t\{\mathbb{E}[h(X_1, \dots, X_n)] - h(X_1, \dots, X_n)\}} \right) \leq e^{\frac{t^2 n(\mathcal{B} + \theta_{\infty, n(1)})^2}{2}}.$$

Others exponential inequalities can be used to obtain PAC-Bounds in the context of time series: the inequalities in [24] for mixing time series, and [23, 39] under weakest “weak dependence” assumptions, [40] for martingales. However, Lemma 1 is particularly convenient here, and will lead to optimal learning rates.

Lemma 2. *We remind that $\kappa = K(1+L)(\mathcal{B}+\mathcal{C})/\sqrt{2}$. We assume that Assumptions **LipLoss**, **WeakDep** and **Lip** are satisfied. For any $\lambda > 0$ and $\theta \in \Theta$,*

$$\mathbb{E} \left(e^{\lambda[R(\theta) - r_n(\theta)]} \right) \leq e^{\frac{\lambda^2 \kappa^2}{n(1 - \frac{\kappa}{n})^2}} \text{ and } \mathbb{E} \left(e^{\lambda[r_n(\theta) - R(\theta)]} \right) \leq e^{\frac{\lambda^2 \kappa^2}{n(1 - \frac{\kappa}{n})^2}}.$$

Proof of Lemma 2. Let us fix $\lambda > 0$ and $\theta \in \Theta$. Let us define the function h by:

$$h(x_1, \dots, x_n) = \frac{1}{K(1+L)} \sum_{i=k+1}^n \ell(f_\theta(x_{i-1}, \dots, x_{i-k}), x_i).$$

We now check that h satisfies (2), remember that $\ell(x, x') = g(x - x')$ so

$$\begin{aligned} & \left| h(x_1, \dots, x_n) - h(y_1, \dots, y_n) \right| \\ & \leq \frac{1}{K(1+L)} \sum_{i=k+1}^n \left| g(f_\theta(x_{i-1}, \dots, x_{i-k}) - x_i) - g(f_\theta(y_{i-1}, \dots, y_{i-k}) - y_i) \right| \end{aligned}$$

$$\leq \frac{1}{1+L} \sum_{i=k+1}^n \left\| (f_\theta(x_{i-1}, \dots, x_{i-k}) - x_i) - (f_\theta(y_{i-1}, \dots, y_{i-k}) - y_i) \right\|$$

where we used Assumption **LipLoss** for the last inequality. So we have

$$\begin{aligned} & \left| h(x_1, \dots, x_n) - h(y_1, \dots, y_n) \right| \\ & \leq \frac{1}{1+L} \sum_{i=k+1}^n \left(\left\| f_\theta(x_{i-1}, \dots, x_{i-k}) - f_\theta(y_{i-1}, \dots, y_{i-k}) \right\| + \left\| x_i - y_i \right\| \right) \\ & \leq \frac{1}{1+L} \sum_{i=k+1}^n \left(\sum_{j=1}^k a_j(\theta) \|x_{i-j} - y_{i-j}\| + \|x_i - y_i\| \right) \\ & \leq \frac{1}{1+L} \sum_{i=1}^n \left(1 + \sum_{j=1}^k a_j(\theta) \right) \|x_i - y_i\| \leq \sum_{i=1}^n \|x_i - y_i\| \end{aligned}$$

where we used Assumption **Lip**. So we can apply Lemma 1 with $h(X_1, \dots, X_n) = \frac{n-k}{K(1+L)} r_n(\theta)$, $\mathbb{E}(h(X_1, \dots, X_n)) = \frac{n-k}{K(1+L)} R(\theta)$, and $t = K(1+L)\lambda/(n-k)$:

$$\mathbb{E} \left(e^{\lambda[R(\theta) - r_n(\theta)]} \right) \leq e^{\frac{\lambda^2 K^2 (1+L)^2 (\mathcal{B} + \theta_{\infty, n(1)})^2}{2n(1-\frac{k}{n})^2}} \leq e^{\frac{\lambda^2 K^2 (1+L)^2 (\mathcal{B} + \mathcal{C})^2}{2n(1-\frac{k}{n})^2}}$$

by Assumption **WeakDep**. This ends the proof of the first inequality. The reverse inequality is obtained by replacing the function h by $-h$. ■

We also remind the following result concerning the Kullback divergence.

Lemma 3. *For any $\pi \in \mathcal{M}_+^1(E)$, for any measurable upper-bounded function $h : E \rightarrow \mathbb{R}$ we have:*

$$\pi[\exp(h)] = \exp \left(\sup_{\rho \in \mathcal{M}_+^1(E)} \left(\rho[h] - \mathcal{K}(\rho, \pi) \right) \right). \quad (3)$$

Moreover, the supremum with respect to ρ in the right-hand side is reached for the Gibbs measure $\pi\{h\}$ defined by $\pi\{h\}(dx) = e^{h(x)}\pi(dx)/\pi[\exp(h)]$.

Actually, it seems that in the case of discrete probabilities, this result was already known by Kullback (Problem 8.28 of Chapter 2 in [41]). For a complete proof in the general case, we refer the reader for example to [25, 13]. We are now ready to state the following key result.

Lemma 4. *Let us assume that Assumptions **LipLoss**, **WeakDep** and **Lip** are satisfied. Let us fix $\lambda > 0$. Let k be defined as in Lemma 2. Then,*

$$\mathbb{P} \left\{ \begin{array}{l} \forall \rho \in \mathcal{M}_+^1(\Theta), \\ \int R d\rho \leq \int r_n d\rho + \frac{\lambda \kappa^2}{n(1-\frac{k}{n})^2} + \frac{\mathcal{K}(\rho, \pi) + \log(\frac{2}{\varepsilon})}{\lambda} \\ \text{and} \\ \int r_n d\rho \leq \int R d\rho + \frac{\lambda \kappa^2}{n(1-\frac{k}{n})^2} + \frac{\mathcal{K}(\rho, \pi) + \log(\frac{2}{\varepsilon})}{\lambda} \end{array} \right\} \geq 1 - \varepsilon. \quad (4)$$

Proof of Lemma 4. Let us fix $\theta > 0$ and $\lambda > 0$, and apply the first inequality of Lemma 2. We have:

$$\mathbb{E} e^{\lambda \left[R(\theta) - r_n(\theta) - \frac{\lambda \kappa^2}{n \left(1 - \frac{k}{n}\right)^2} \right]} \leq 1,$$

and we multiply this result by $\varepsilon/2$ and integrate it with respect to $\pi(d\theta)$. Fubini's Theorem gives:

$$\mathbb{E} \int e^{\lambda \left[R(\theta) - r_n(\theta) - \frac{\lambda^2 \kappa^2}{n \left(1 - \frac{k}{n}\right)^2} - \log\left(\frac{2}{\varepsilon}\right) \right]} \pi(d\theta) \leq \frac{\varepsilon}{2}.$$

We apply Lemma 3 and we get:

$$\mathbb{E} e^{\sup_{\rho} \left\{ \lambda \int [R(\theta) - r_n(\theta)] \rho(d\theta) - \frac{\lambda^2 \kappa^2}{n \left(1 - \frac{k}{n}\right)^2} - \log\left(\frac{2}{\varepsilon}\right) - \mathcal{K}(\rho, \pi) \right\}} \leq \frac{\varepsilon}{2}.$$

As $e^x \geq \mathbf{1}_{\mathbb{R}_+}(x)$, we have:

$$\mathbb{P} \left\{ \sup_{\rho} \left\{ \lambda \int [R(\theta) - r_n(\theta)] \rho(d\theta) - \frac{\lambda^2 \kappa^2}{n \left(1 - \frac{k}{n}\right)^2} - \log\left(\frac{2}{\varepsilon}\right) - \mathcal{K}(\rho, \pi) \right\} \geq 0 \right\} \leq \frac{\varepsilon}{2}.$$

Now, we follow the same proof again but starting with the second inequality of Lemma 2. We obtain:

$$\mathbb{P} \left\{ \sup_{\rho} \left\{ \lambda \int [r_n(\theta) - R(\theta)] \rho(d\theta) - \frac{\lambda^2 \kappa^2}{n \left(1 - \frac{k}{n}\right)^2} - \log\left(\frac{2}{\varepsilon}\right) - \mathcal{K}(\rho, \pi) \right\} \geq 0 \right\} \leq \frac{\varepsilon}{2}.$$

A union bound ends the proof. \blacksquare

7.2 Proof of Theorems 1 and 2

Proof of Theorem 1. We apply Lemma 4. So, with probability at least $1 - \varepsilon$ we are on the event given by (4). From now, we work on that event. The first inequality of (4), when applied to $\hat{\rho}_\lambda(d\theta)$, gives

$$\int R(\theta) \hat{\rho}_\lambda(d\theta) \leq \int r_n(\theta) \hat{\rho}_\lambda(d\theta) + \frac{\lambda \kappa^2}{n \left(1 - \frac{k}{n}\right)^2} + \frac{1}{\lambda} \log\left(\frac{2}{\varepsilon}\right) + \frac{1}{\lambda} \mathcal{K}(\hat{\rho}_\lambda, \pi).$$

According to Lemma 3 we have:

$$\int r_n(\theta) \hat{\rho}_\lambda(d\theta) + \frac{1}{\lambda} \mathcal{K}(\hat{\rho}_\lambda, \pi) = \inf_{\rho} \left(\int r_n(\theta) \rho(d\theta) + \frac{1}{\lambda} \mathcal{K}(\rho, \pi) \right)$$

so we obtain

$$\int R(\theta) \hat{\rho}_\lambda(d\theta) \leq \inf_{\rho} \left[\int r_n(\theta) \rho(d\theta) + \frac{\lambda \kappa^2}{n \left(1 - \frac{k}{n}\right)^2} + \frac{\mathcal{K}(\rho, \pi) + \log\left(\frac{2}{\varepsilon}\right)}{\lambda} \right]. \quad (5)$$

We now want to bound from above $r(\theta)$ by $R(\theta)$. Applying the second inequality of (4) and plugging it into Inequality 5 gives

$$\int R(\theta) \hat{\rho}_\lambda(d\theta) \leq \inf_{\rho} \left[\int R d\rho + \frac{2}{\lambda} \mathcal{K}(\rho, \pi) + \frac{2\lambda\kappa^2}{n(1 - \frac{\kappa}{n})^2} + \frac{2}{\lambda} \log\left(\frac{2}{\varepsilon}\right) \right].$$

We end the proof by the remark that $\theta \mapsto R(\theta)$ is convex and so by Jensen's inequality $\int R(\theta) \hat{\rho}_\lambda(d\theta) \geq R(\int \theta \hat{\rho}_\lambda(d\theta)) = R(\hat{\theta}_\lambda)$. \blacksquare

Proof of Theorem 2. We can apply Theorem 1 with $R = R_\tau$. We have, with probability at least $1 - \varepsilon$,

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\rho \in \mathcal{M}_+^1(\Theta)} \left[\int R^\tau d\rho + \frac{2\lambda\kappa^2}{n(1 - \frac{2}{n})^2} + \frac{2\mathcal{K}(\rho, \pi) + 2\log(\frac{2}{\varepsilon})}{\lambda} \right].$$

Now, let us fix $\delta \in (0, 1]$ and $\theta \in \Theta(B)$. We define the probability distribution $\rho_{\theta,\delta}$ as the uniform probability measure on the set $\{T \in \mathbb{R}^4, \|\theta - T\|_1 \leq \delta\}$. Note that $\rho_{\theta,\delta} \ll \pi_B$ as π_B is defined as uniform on $\Theta(B+1) \supset \Theta(B+\delta)$. Then:

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\substack{\theta \in \Theta(B) \\ \delta > 0}} \left[\int R^\tau d\rho_{\theta,\delta} + \frac{2\lambda\kappa^2}{n(1 - \frac{2}{n})^2} + \frac{2\mathcal{K}(\rho_{\theta,\delta}, \pi) + 2\log(\frac{2}{\varepsilon})}{\lambda} \right]. \quad (6)$$

Now, we have to compute or to upper-bound all the terms in the right-hand side of this inequality. First, note that:

$$\begin{aligned} \int R^\tau d\rho_{\theta,\delta} &= \int_{\{\|\theta - T\|_1 \leq \delta\}} R^\tau(T) d\rho_{\theta,\delta}(T) \\ &\leq R^\tau(\theta) + 2\mathcal{B}\delta \max(\tau, 1 - \tau) \leq R^\tau(\theta) + 2\mathcal{B}\delta. \end{aligned} \quad (7)$$

We compute $\mathcal{K}(\rho_{\theta,\delta}, \pi_B) = 3\log(\frac{B+1}{\delta})$ and plug this with (7) into (6) to obtain:

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\theta, \delta} \left\{ R^\tau(\theta) + 2 \left[\frac{\lambda\kappa^2}{n(1 - \frac{2}{n})^2} + \mathcal{B}\delta + \frac{3\log(\frac{B+1}{\delta}) + \log(\frac{2}{\varepsilon})}{\lambda} \right] \right\}.$$

It can easily be seen that the minimum of the right-hand side w.r.t. δ is reached for $\delta = 3/(\mathcal{B}\lambda) \leq 1$ as soon as λ is large enough, and so:

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\theta} \left\{ R^\tau(\theta) + \frac{2\lambda\kappa^2}{n(1 - \frac{2}{n})^2} + \frac{6\log\left(\frac{(B+1)\mathcal{B}\lambda\varepsilon}{3}\right) + 2\log(\frac{2}{\varepsilon})}{\lambda} \right\}.$$

We finally propose $\lambda = \sqrt{3n}/\kappa$, this leads to:

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\theta} \left\{ R^\tau(\theta) + \frac{2\sqrt{3}\kappa}{\sqrt{n}} \left[\frac{1}{(1 - \frac{2}{n})^2} + \log\left(\frac{(B+1)\mathcal{B}\varepsilon}{\kappa} \sqrt{\frac{n}{3}}\right) + \frac{\log(\frac{2}{\varepsilon})}{3} \right] \right\}.$$

Remark that the condition $\delta < 1$ is satisfied as soon as $n > \kappa^2/(3\mathcal{B}^2)$. Moreover, $\forall n \geq 10$, $1/(1 - \frac{2}{n})^2 \leq \frac{25}{16}$ and we can re-organize the terms to obtain:

$$R^\tau(\hat{\theta}_{B,\lambda}^\tau) \leq \inf_{\theta} \left\{ R^\tau(\theta) + \frac{2\sqrt{3}\kappa}{\sqrt{n}} \left[2.25 + \log \left(\frac{(B+1)\mathcal{B}\sqrt{n}}{\kappa} \right) + \frac{\log(\frac{1}{\varepsilon})}{3} \right] \right\}.$$

■

References

1. Koenker, R., Bassett, G.J.: Regression quantiles. *Econometrica* **46** (1978) 33–50
2. Hamilton, J.: Time Series Analysis. Princeton University Press (1994)
3. Brockwell, P., Davis, R.: Time Series: Theory and Methods (2nd Edition). Springer (2009)
4. Cesa-Bianchi, N., Lugosi, G.: Prediction, Learning, and Games. Cambridge University Press, New York (2006)
5. Stoltz, G.: Agrégation séquentielle de prédicteurs : méthodologie générale et applications à la prévision de la qualité de l’air et à celle de la consommation électrique. *Journal de la SFDS* **151** (2010) 66–106
6. Modha, D.S., Masry, E.: Memory-universal prediction of stationary random processes. *IEEE transactions on information theory* **44** (1998) 117–133
7. Meir, R.: Nonparametric time series prediction through adaptive model selection. *Machine Learning* **39** (2000) 5–34
8. Alquier, P., Wintenberger, O.: Model selection for weakly dependent time series forecasting. *Bernoulli* **18** (2012) 883–193
9. Biau, G., Patra, B.: Sequential quantile prediction of time series. *IEEE Transactions on Information Theory* **57** (2011) 1664–1674
10. Dalalyan, A., Tsybakov, A.: Aggregation by exponential weighting, sharp PAC-Bayesian bounds and sparsity. *Machine Learning* **72** (2008) 39–61
11. Gerchinovitz, S.: Sparsity regret bounds for individual sequences in online linear regression. In: *Proceedings of COLT’11*. (2011)
12. Catoni, O.: Statistical Learning Theory and Stochastic Optimization. Springer Lecture Notes in Mathematics (2004)
13. Catoni, O.: PAC-Bayesian Supervised Classification (The Thermodynamics of Statistical Learning). Volume 56 of Lecture Notes-Monograph Series. IMS (2007)
14. Littlestone, N., Warmuth, M.: The weighted majority algorithm. *Information and Computation* **108** (1994) 212–261
15. Vovk, V.: Aggregating strategies. In: *Proceedings of the 3rd Annual Workshop on Computational Learning Theory (COLT)*. (1990) 372–283
16. Koenker, R.: Quantile Regression. Cambridge University Press, Cambridge (2005)
17. Belloni, A., Chernozhukov, V.: L1-penalized quantile regression in high-dimensional sparse models. *The Annals of Statistics* **39** (2011) 82–130
18. Shawe-Taylor, J., Williamson, R.: A PAC analysis of a bayes estimator. In: *Proceedings of the Tenth Annual Conference on Computational Learning Theory, COLT’97*, ACM (1997) 2–9
19. McAllester, D.A.: PAC-Bayesian model averaging. In: *Procs. of of the 12th Annual Conf. On Computational Learning Theory, Santa Cruz, California (Electronic)*, ACM, New-York (1999) 164–170

20. Alquier, P.: PAC-Bayesian bounds for randomized empirical risk minimizers. *Mathematical Methods of Statistics* **17** (2008) 279–304
21. Audibert, J.Y.: PAC-Bayesian aggregation and multi-armed bandits. HDR Université Paris Est (2010)
22. Audibert, J.Y., Catoni, O.: Robust linear least squares regression. *The Annals of Statistics* **39** (2011) 2766–2794
23. Dedecker, J., Doukhan, P., Lang, G., León, J.R., Louhichi, S., Prieur, C.: *Weak Dependence, Examples and Applications*. Volume 190 of *Lecture Notes in Statistics*. Springer-Verlag, Berlin (2007)
24. Doukhan, P.: *Mixing*. *Lecture Notes in Statistics*. Springer, New York (1994)
25. Catoni, O.: A PAC-Bayesian approach to adaptative classification. Preprint Laboratoire de Probabilités et Modèles Aléatoires (2003)
26. Devilliers, M.: Les enquêtes de conjoncture. In: *Archives et Documents*. Number 101, INSEE (1984)
27. Clavel, L., Minodier, C.: A monthly indicator of the french business climate. *Documents de Travail de la DESE* (2009)
28. Dubois, E., Michaux, E.: étalonnages à l’aide d’enquêtes de conjoncture: de nouveaux résultats. In: *Économie et Prévision*. Number 172, INSEE (2006)
29. Biau, G., Biau, O., Rouvière, L.: Nonparametric forecasting of the manufacturing output growth with firm-level survey data. *Journal of Business Cycle Measurement and Analysis* **3** (2008) 317–332
30. Diebold, F.X., Tay, A.S., Wallis, K.F.: Evaluating density forecasts of inflation: the survey of professional forecasters. Discussion Paper No.48, ESRC Macroeconomic Modelling Bureau, University of Warwick and Working Paper No.6228, National Bureau of Economic Research, Cambridge, Mass. (1997)
31. Tay, A.S., Wallis, K.F.: Density forecasting: a survey. *Journal of Forecasting* **19** (2000) 235–254
32. Britton, E., Fisher, P., Whitley, J.: The inflation report projections: Understanding the fan chart. *Bank of England Quarterly Bulletin* **38** (1998) 30–37
33. Cornec, M.: Constructing a conditional gdp fan chart with an application to french business survey data. 30th CIRET Conference, New York (2010)
34. Dowd, K.: The inflation fan charts: An evaluation. *Greek Economic Review* **23** (2004) 99–111
35. Li, X.: Agrégation de prédicteurs appliquée à la conjoncture. Rapport de stage de M2 - Université Paris 6 - INSEE sous la direction de Matthieu Cornec (2010)
36. Casella, G., Robert, C.: *Monte Carlo Statistical Methods*. Springer-Verlag (2nd Edition 2004)
37. R Development Core Team: *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna. (2008)
38. Rio, E.: Inégalités de Hoeffding pour les fonctions lipschitziennes de suites dépendantes. *Comptes Rendus de l’Académie des Sciences de Paris, Série I* **330** (2000) 905–908
39. Wintenberger, O.: Deviation inequalities for sums of weakly dependent time series. *Electronic Communications in Probability* **15** (2010) 489–503
40. Seldin, Y., Laviolette, F., Cesa-Bianchi, N., Auer, P., Shawe-Taylor, J.: PAC-Bayesian inequalities for martingales. Preprint arXiv:1110.6886 (2011)
41. Kullback, S.: *Information theory and statistics*. Wiley, New York (1959)