

# Nonparametric Risk Bounds for Time-Series Forecasting

**Daniel J. McDonald**

*Department of Statistics  
Indiana University  
Bloomington, IN 47404*

DAJMCDON@INDIANA.EDU

**Cosma Rohilla Shalizi**

*Department of Statistics  
Carnegie Mellon University  
Pittsburgh, PA 15208*

CSHALIZI@CMU.EDU

**Mark Schervish**

MARK@CMU.EDU

**Editor:** Mehryar Mohri

## Abstract

We derive generalization error bounds for traditional time-series forecasting models. Our results hold for many standard forecasting tools including autoregressive models, moving average models, and, more generally, linear state-space models. These non-asymptotic bounds need only weak assumptions on the data-generating process, yet allow forecasters to select among competing models and to guarantee, with high probability, that their chosen model will perform well. We motivate our techniques with and apply them to standard economic and financial forecasting tools—a GARCH model for predicting equity volatility and a dynamic stochastic general equilibrium model (DSGE), the standard tool in macroeconomic forecasting. We demonstrate in particular how our techniques can aid forecasters and policy makers in choosing models which behave well under uncertainty and mis-specification.

**Keywords:** generalization error, prediction risk, model selection, VC dimension, state-space models, linear time-invariant systems

## 1. Introduction

Generalization error bounds are probabilistically valid, non-asymptotic tools for characterizing the predictive ability of forecasting models. This methodology is fundamentally about choosing particular prediction functions out of some class of plausible alternatives so that, with high reliability, the resulting predictions will be nearly as accurate as possible (“probably approximately correct”). While many of these results are aimed at classification problems with independent and identically distributed (i.i.d.) data, this paper adapts and extends these methods to time-series models, so that economic and financial forecasting techniques can be evaluated rigorously. In particular, these methods control the expected accuracy of future predictions from mis-specified models based on finite samples. This allows for immediate model comparisons which neither appeal to asymptotics nor make strong assumptions about the data-generating process, in stark contrast to such popular model-selection tools as AIC.

To fix ideas, imagine i.i.d. data  $((Y_1, X_1), \dots, (Y_n, X_n))$  with  $(Y_i, X_i) \in \mathcal{Y} \times \mathcal{X}$ , some prediction function  $f : \mathcal{X} \rightarrow \mathcal{Y}$ , and a loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$  which measures the cost of bad predictions. The *generalization error* or *risk* of  $f$  is

$$R(f) := \mathbb{E}[\ell(Y, f(X))],$$

where the expectation is taken with respect to  $\mathbb{P}$ , the joint distribution of  $(Y, X)$ . The generalization error measures the inaccuracy of our predictions when we use  $f$  on future data, making it a natural criterion for model selection, and a target for performance guarantees. To actually calculate the risk, we would need to know the data-generating distribution  $\mathbb{P}$  and have a single fixed prediction function  $f$ , neither of which is common. Because explicitly calculating the risk is infeasible, forecasters typically try to estimate it, which calls for detailed assumptions on  $\mathbb{P}$ . The alternative we employ here is to find upper bounds on risk which hold uniformly over large classes of models  $\mathcal{F}$  from which some particular  $f$  is chosen, possibly in a data-dependent way, and uniformly over distributions  $\mathbb{P}$ .

Our main results in Section 5 assert that for wide classes of time-series models, the expected cost of poor predictions is bounded by the model’s in-sample performance inflated by a term which balances the amount of observed data with the complexity of the model. The bound holds with high probability under the unknown distribution  $\mathbb{P}$  assuming only mild conditions—existence of some moments, stationarity, and the decay of temporal dependence as data points become widely separated in time. We give applications in Section 6.

Our goal in this paper is to provide general bounds for common time-series models with unbounded loss functions, no explicit regularization, and potential dependence on the entirety of the observed data. The bounds we derive here are relevant for the time-series models typically used in applied settings—finance, economics, engineering, etc.—as well as covering models more common in machine learning. In particular, we derive results for non-linear models which depend only on a fixed quantity of recent data and linear time invariant systems, state-space models, which use the entire past to predict new data. These results however do not cover, e.g. HMMs in the strictest sense, as they require absolutely continuous latent states rather than discrete valued ones.

The remainder of this paper is structured as follows. Section 2 provides motivation and background for our results, giving intuition in the i.i.d. setting by focusing on concentration of measure ideas and characterizations of model complexity. Section 3 gives the explicit assumptions we make and describes how to leverage powerful ideas from time series to generalize the i.i.d. methods. Section 4 introduces linear time-invariant systems and discusses how such forecasters are different from, e.g., autoregressive models. Section 5 states and proves risk bounds for the time-series forecasting setting, while we demonstrate how to use the results in Section 6 and give some properties of those results in Section 7. Finally, Section 8 concludes and illustrates the path toward generalizing our methods to more elaborate model classes.

## 1.1 Related Work

Yu (1994) and Nobel and Dembo (1993) showed that it was possible to transfer some i.i.d. results to  $\beta$ -mixing sequences, albeit without explicit rates of convergence (c.f. Vidyasagar, 2003, §3.4). Karandikar and Vidyasagar (2002); Meir (2000); Vidyasagar (2003) presented

bounds for model classes with finite covering numbers (a sufficient condition for which is finite VC-dimension of a related class) but additionally require that  $\mathcal{Y}$  is compact and  $\ell$  is bounded. Early work in signal processing (Modha and Masry, 1998) proposes predictors based on sequences of parametric models of increasing memory which minimize a complexity regularized least squares criterion and establish that these predictors deliver the same statistical performance as oracle predictors. Steinwart and Christmann (2009) prove an oracle inequality regularized ERM algorithms when observations are  $\alpha$ -mixing which are close to the optimal i.i.d. rates. Mohri and Rostamizadeh (2009) give results using Rademacher complexity which are both tighter than those using VC-dimension or covering numbers as well as being computable from the data in many cases. Mohri and Rostamizadeh (2010) and Agarwal and Duchi (2013) consider another family of bounds for  $\phi$ -mixing and  $\beta$ -mixing sequences when the predictors are algorithmically stable. Many classes of common machine learning algorithms are amenable to either Rademacher or algorithmic-stability bounds: Kernel-regularized methods, support-vector machines, relative-entropy based regularization, and kernel ridge regression among others. However, methods common to time-series such as AR models, ARIMA models, ARCH and GARCH models (Engle, 1982, 2001), state-space models, and other Box-Jenkins type predictors are not because they are not explicitly regularized, the loss functions are not bounded, and the predictions can depend on more than simply a fixed dimensional past. McDonald et al. (2011b) shows that stationarity alone can be used to impose a kernel-type regularization on an AR model, and hence, following the results of Mohri and Rostamizadeh (2009), is amenable to Rademacher complexity for a bounded loss function.

Other dependence conditions apart from stationary and strong mixing are also considered in the literature. Alquier et al. (2012) develop oracle inequalities and model selection procedures for linear models, neural networks, and non-parametric autoregressions when observations come from causal Bernoulli shifts or bounded, weakly-dependent processes. Under the same weak-dependence conditions, Alquier et al. (2014) extends this result to convex Lipschitz loss functions and examines forecasting of the French GDP. Finally, recent work by Kuznetsov and Mohri (2014) examines both average-path generalization and path-dependent generalization for certain types of non-stationary mixing processes and derives Rademacher complexity bounds.

## 2. Statistical Learning Theory for I.I.D. Data

Our goal is to control the risk of predictive models, that is, their expected inaccuracy on new data from the same stochastic source as the data used to fit the model. To orient readers, we present some standard results for i.i.d. data, which are adapted to the dependent setting in Section 5.

Let  $f : \mathcal{X} \rightarrow \mathcal{Y}$  be some function used for making predictions of  $Y$  from  $X$ . We define a loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$  which measures the cost of making poor predictions. Throughout this paper, we will assume that  $\ell(y, y')$  is a function solely of the difference  $y - y'$  where  $\ell(\cdot)$  is nonnegative and  $\ell(0) = 0$ ; we take the liberty of denoting that function  $\ell(y - y')$ . Then the risk of any predictor  $f \in \mathcal{F}$  (where  $f$  is fixed independently of the data)

is given by

$$R(f) = \mathbb{E}[\ell(Y - f(X))],$$

where  $(X, Y) \sim \mathbb{P}$ . The risk or generalization error is the expected cost of using  $f$  to predict  $Y$  from  $X$  on a new observation.

Since the true distribution  $\mathbb{P}$  is unknown, so is  $R(f)$ , but we can try to estimate it based on our observed data. The *training error* or *empirical risk* of  $f$  is

$$\widehat{R}_n(f) := \frac{1}{n} \sum_{i=1}^n \ell(Y_i - f(X_i)).$$

In other words, the in-sample training error,  $\widehat{R}_n(f)$ , is the average loss over the actual training points. For any given  $f$ , we can bound  $R(f)$  in terms of  $\widehat{R}_n(f)$  using deviation inequalities, as illustrated below.

When we use the data to choose an  $\widehat{f}$  from  $\mathcal{F}$ , we would like to bound  $R(\widehat{f})$ . To do so, we must consider not just  $\widehat{R}_n(\widehat{f})$ , but also the size, in some sense, of  $\mathcal{F}$ . There are a number of measures for the size or capacity of a model many of which lead to learning theoretic risk bounds. Algorithmic stability (Bousquet and Elisseeff, 2002; Kearns and Ron, 1999) quantifies the sensitivity of the chosen function to small perturbations to the data. Similarly, maximal discrepancy (Vapnik, 2000) asks how different the predictions could be if two functions are chosen using two separate data sets. A more direct, functional-analytic approach partitions  $\mathcal{F}$  into equivalence classes under some metric, leading to covering numbers (Pollard, 1984, 1990). Rademacher complexity (Bartlett and Mendelson, 2002; Koltchinskii and Panchenko, 2002) directly describes a model’s ability to fit random noise. We focus on a measure which is both intuitive and powerful: Vapnik-Chervonenkis (VC) dimension (Vapnik, 2000; Vapnik and Chervonenkis, 1971).

VC dimension starts as an idea about collections of sets.

**Definition 1** *Let  $\mathbb{U}$  be some (infinite) set and  $S$  a finite-cardinality subset of  $\mathbb{U}$ . Let  $\mathcal{C}$  be a family of subsets of  $\mathbb{U}$ . We say that  $\mathcal{C}$  shatters  $S$  if for every  $S' \subseteq S$ ,  $\exists C \in \mathcal{C}$  such that  $S' = S \cap C$ .*

Essentially,  $\mathcal{C}$  can shatter a set  $S$  if it can pick out every subset of points in  $S$ . This says that the collection  $\mathcal{C}$  is very complicated or flexible. The cardinality of the largest set  $S$  that can be shattered by  $\mathcal{C}$  is the latter’s VC dimension.

**Definition 2 (VC dimension)** *The Vapnik-Chervonenkis (VC) dimension of a collection  $\mathcal{C}$  of subsets of  $\mathbb{U}$  is*

$$\text{VCD}(\mathcal{C}) := \sup\{|S| : S \subseteq \mathbb{U} \text{ and } S \text{ is shattered by } \mathcal{C}\}.$$

To see why this is a “dimension”, we need one more notion.

**Definition 3 (Growth function)** *The growth function  $G(n, \mathcal{C})$  of a collection  $\mathcal{C}$  of subsets of  $\mathbb{U}$  is the maximum number of subsets which can be formed by intersecting a set  $S \subset \mathbb{U}$  of cardinality  $n$  with  $\mathcal{C}$ ,*

$$G(n, \mathcal{C}) := \sup_{S \subset \mathbb{U} : |S|=n} |S \wedge \mathcal{C}|,$$

where  $\mathcal{A} \wedge \mathcal{B}$  is the class of all sets  $A \cap B$ ,  $A \in \mathcal{A}$ ,  $B \in \mathcal{B}$ .

The growth function counts how many *effectively* distinct sets the collection contains, when we can only observe what is going on at  $n$  points, not all of  $\mathbb{U}$ . If  $n \leq \text{vcd}(\mathcal{C})$ , then from the definitions  $G(n, \mathcal{C}) = 2^n$ . If the VC dimension is finite, however, and  $n > \text{vcd}(\mathcal{C})$ , then  $G(n, \mathcal{C}) < 2^n$ , and in fact it can be shown (Vapnik and Chervonenkis, 1971) that

$$G(n, \mathcal{C}) \leq \left( \frac{en}{\text{vcd}(\mathcal{C})} \right)^{\text{vcd}(\mathcal{C})} \leq (n+1)^{\text{vcd}(\mathcal{C})}.$$

This polynomial growth of capacity with  $n$  is why  $\text{vcd}$  is a “dimension”.

Using VC dimension to measure the capacity of function classes is straightforward. Define the indicator function  $\mathbf{1}_A(x)$  to take the value 1 if  $x \in A$  and 0 otherwise. Suppose that  $f \in \mathcal{F}$ ,  $f : \mathbb{U} \rightarrow \mathbb{R}$ . Each  $f$  corresponds to the set

$$C_f = \{(u, u_0) \in \mathbb{U} \times \mathbb{R} : \mathbf{1}_{[0, \infty)}(f(u) - u_0) = 1\},$$

so  $\mathcal{F}$  corresponds to the class  $\mathcal{C}_{\mathcal{F}} := \{C_f : f \in \mathcal{F}\}$ . This extension is sometimes called the pseudo dimension (see e.g. Anthony and Bartlett, 1999; Pollard, 1990).

**Theorem 4 (Vapnik and Chervonenkis, 1971)** *Suppose that  $0 \leq \ell(y, y') \leq M < \infty$ . Then,*

$$\mathbb{P} \left( \sup_{f \in \mathcal{F}} |R(f) - \widehat{R}_n(f)| > \epsilon \right) \leq 4G(2n, \mathcal{C}_{\mathcal{F}}) \exp \left\{ -\frac{n\epsilon^2}{\Upsilon} \right\},$$

where  $\Upsilon$  depends only on  $M$  and not  $n$  or  $\mathcal{F}$ .

When the loss function is unbounded, similar results hold, but we must consider the composition of the loss function with  $f$ . This leads to the set

$$C_{\ell \circ f} = \{(u, u_0, u'_0) \in \mathbb{U} \times \mathbb{R} \times \mathbb{R} : \mathbf{1}_{[0, \infty)}(\ell(u_0 - f(u)) - u'_0) = 1\},$$

and the corresponding class  $\mathcal{C}_{\ell \circ \mathcal{F}} := \{C_{\ell \circ f} : f \in \mathcal{F}\}$ .

The concentration result in Theorem 4 works well for independent data. However, for time series, we must be able to handle dependent data. In particular the length  $n$  of a sample path  $Y_1, \dots, Y_n$  exaggerates how much information it contains relative to independent observations. Knowing the past allows forecasters to predict future data (at least to some degree), so actually observing those future data points gives less information about the underlying process than in the i.i.d. case. Thus, while in Theorem 4 the probability of large discrepancies between empirical means and their expectations decreases exponentially in  $n$ , in the dependent case, the effective sample size may be much less than  $n$ , resulting in looser bounds.

### 3. Time Series

In moving from the i.i.d. setting to time-series forecasting, we need a number of modifications to our initial setup. Rather than observing input/output pairs  $(Y_i, X_i)$ , we observe a single sequence of random variables  $Y_1, \dots, Y_n$  where each  $Y_i$  takes values in  $\mathbb{R}^p$ , though we can generalize to arbitrary metric spaces at some cost in notational clarity. We are interested in using functions which take past observations as inputs and predict future values of the

process. Specifically, given data from time 1 to time  $n$ , we wish to predict time  $n + 1$ . To be clear about notation, we will use the following conventions:  $Y_{i:j} := (Y_i, Y_{i+1}, \dots, Y_j)$ ,  $Y_\infty := Y_{-\infty:\infty}$  is an infinite dimensional sequence; we also have the associated joint distributions  $\mathbb{P}_{i:j}$  and  $\mathbb{P}_\infty$  and  $\sigma$ -fields  $\sigma_{i:j} = \sigma(Y_{i:j})$  and  $\sigma_\infty = \sigma(Y_\infty)$ .

While we no longer presume i.i.d. data, we still need to restrict the sort of dependent process we work with. We first remind the reader of the notion of (strict or strong) stationarity.

**Definition 5 (Stationarity)** *A random sequence  $Y_\infty$  is stationary when all its finite-dimensional distributions are time-invariant: for all  $t$  and all non-negative integers  $i$  and  $j$ , the random vectors  $Y_{t:t+i}$  and  $Y_{t+j:t+i+j}$  have the same distribution.*

Stationarity does not imply that the random variables  $Y_i$  are independent across time  $i$ , only that the marginal distribution of  $Y_i$  is constant in time. (And similarly for  $Y_{i:i+j}$ .) We limit ourselves not just to stationary processes, but also to ones in which widely-separated observations are asymptotically independent. Without this restriction, convergence of the training error to the expected risk could occur arbitrarily slowly, and finite-sample bounds may not exist. In fact, Adams and Nobel (2010) demonstrate that for ergodic processes, finite VC dimension is enough to give consistency, but cannot itself provide rates. The next definition describes the sort of serial dependence which we entertain.

**Definition 6 ( $\beta$ -Mixing)** *Consider a stationary random sequence  $Y_\infty$  defined on a probability space  $(\Omega, \Sigma, \mathbb{P}_\infty)$ . Let  $\mathbb{P}_0$  be the restriction of  $\mathbb{P}_\infty$  to  $\sigma_{-\infty:0}$ ,  $\mathbb{P}_a$  be the restriction of  $\mathbb{P}_\infty$  to  $\sigma_{a:\infty}$ , and  $\mathbb{P}_{0 \otimes a}$  be the restriction of  $\mathbb{P}_\infty$  to  $\sigma(Y_{-\infty:0}, Y_{a:\infty})$ . The coefficient of absolute regularity, or  $\beta$ -mixing coefficient,  $\beta_a$ , is given by*

$$\beta_a := \|\mathbb{P}_0 \times \mathbb{P}_a - \mathbb{P}_{0 \otimes a}\|_{TV},$$

where  $\|\cdot\|_{TV}$  is the total variation norm. A stochastic process is absolutely regular, or  $\beta$ -mixing, if  $\beta_a \rightarrow 0$  as  $a \rightarrow \infty$ .

This is only one of many equivalent characterizations of  $\beta$ -mixing (see Bradley, 2005, for others). This definition makes clear that a process is  $\beta$ -mixing if the joint probability of events which are widely separated in time approaches the product of the individual probabilities, that is that  $Y_\infty$  is asymptotically independent. Many common time-series models are known to be  $\beta$ -mixing, and the rates of decay are known up to constant factors which are functions of the true parameters of the process. Among the processes for which such results are known are ARMA models (Mokkadem, 1988), GARCH models (Carrasco and Chen, 2002), and certain Markov processes (see Doukhan, 1995, for an overview). Additionally, functions of  $\beta$ -mixing processes are  $\beta$ -mixing, so if  $\mathbb{P}_\infty$  could be specified by a linear time-invariant system (see below), state-space model, vector auto regression, or some function of a hidden Markov model, the observed data would satisfy this condition.

Knowing  $\beta_a$  would let us determine the effective sample size of a time series  $Y_{1:n}$ . In effect, having  $n$  dependent-but-mixing data points is like having  $\mu < n$  independent ones. Once we determine the correct  $\mu$ , we can (as we will show) use concentration results for i.i.d. data like Theorem 4 with small corrections.

---

**Algorithm 1:** Kalman filter

---

Recursively generate minimum mean squared error predictions  $\widehat{Y}_i$  using the state space model in (1).

**Input:** Initial guesses for the mean and variance of  $\alpha_1$ :  $\widehat{\alpha}_1$  and  $\widehat{P}_1$

**Set**  $\widehat{Y}_1 = T\widehat{\alpha}_1$ .

**for**  $1 \leq i \leq n$  **do**

Filter

$$\begin{aligned} v_i &= Y_i - \widehat{Y}_i, & F_i &= (Z\widehat{P}_iZ^\top + H)^{-1}, \\ K_i &= T\widehat{P}_iZ^\top F_i, & L_i &= T - K_iZ, \\ \widehat{\alpha}_{i+1} &= T\widehat{\alpha}_i + K_iv_i, & \widehat{P}_{i+1} &= T\widehat{P}_iL_i^\top + GQG^\top. \end{aligned}$$

Predict  $\widehat{Y}_{i+1} = Z\widehat{\alpha}_{i+1}$ .

**end**

**return**  $\widehat{Y}_{1:n+1}$

---

#### 4. Linear Time-Invariant Dynamical Systems

Our goal in this paper will be to derive risk bounds for linear time-invariant dynamical systems (LTIs). Such models presume an underlying latent process, and attempt, given observations, to learn that process, predict future values of the latent process, and forecast future observations. Learning algorithms for these goals are linear functions of all previously observed values. Such models nest many common time-series forecasting techniques—ARIMA models, GARCH models, linear-Gaussian state-space models—but, due to their (in general) dependence on the entire past, are not covered by the work discussed in Section 1.1. We present the general form of such models here, provide a canonical forecasting algorithm, and discuss some properties of such models.

Linear dynamical systems model observations  $Y_i$  as

$$\begin{aligned} Y_i &= Z\alpha_i + \epsilon_i, \\ \alpha_{i+1} &= T\alpha_i + G\eta_{i+1}. \end{aligned} \tag{1}$$

This is essentially a hidden markov model under certain conditions:  $Y_i$  denotes observations,  $\alpha_i$  are hidden “state” variables,  $\epsilon_i$  and  $\eta_i$  are both absolutely continuous random noise with  $\mathbb{E}[\epsilon_i] = \mathbb{E}[\eta_i] = 0$ ,  $\mathbb{E}[\epsilon_i\epsilon_j^\top] = \delta_{ij}H$ , and  $\mathbb{E}[\eta_i\eta_j^\top] = \delta_{ij}Q$  for all  $i, j$ . We further assume that  $\epsilon$  and  $\eta$  are mutually independent even though this is not strictly necessary, because it makes notation simpler. We require stationarity for our results, and so we also require the LTI to be stationary. This amounts to forcing the complex eigenvalues of  $T$  to lie inside the unit circle. We note that the condition that the noise distributions are absolutely continuous means that HMMs in the strictest sense are not members of this family. We also allow the parameter matrices  $Z, T, H, G$ , and  $Q$  to depend on a (possibly unknown) parameter vector  $\theta$ , and assume that  $H$  and  $Q$  are positive definite for all  $\theta$ .

The *filtering* problem uses observations  $Y_{1:i}$  up to time  $i$  to learn information about the distribution of  $\alpha_i$ . Then, conditional on an estimate  $\widehat{\alpha}_i$ , we can forecast  $\widehat{\alpha}_{i+1}$  and hence

derive a prediction  $\widehat{Y}_{i+1}$ . For models of this form, one uses the Kalman filter (Anderson and Moore, 2012; Durbin and Koopman, 2001; Kalman, 1960) both to estimate the latent variables,  $\alpha_{1:n+1}$  and to generate predictions  $\widehat{Y}_{1:n+1}$  (Algorithm 1). This procedure gives the minimum mean-squared error predictions of  $\alpha_{i+1}$  (and hence of  $Y_{i+1}$ ) given  $Y_{1:i}$  in the sense that

$$\widehat{\alpha}_{i+1} = \operatorname{argmin}_a \mathbb{E} \left[ \operatorname{tr} \left( (\alpha_{i+1} - a)(\alpha_{i+1} - a)^\top \right) \mid Y_{1:i} \right].$$

Furthermore, if  $\epsilon_i$  and  $\eta_i$  are Gaussian, then Algorithm 1 also gives the likelihood for the unknown parameter vector  $\theta$ . To estimate the unknown parameters, we either: (1) maximize the likelihood returned by the filter; or (2) use the EM algorithm, alternating between running the Kalman filter (the E-step) and maximizing the conditional likelihood by least squares (the M-step). Bayesian estimation works like EM, replacing the M-step with Bayesian updating.

Predictions based on Algorithm 1 are linear functions of previous observations, but these predictions depend on all previous observations rather than simply a fixed number as would be the case with, say, autoregressive models. More specifically,  $\widehat{Y}_{1:n+1} = \mathbf{B}Y_{1:n}$  where

$$(\mathbf{B})_{ij} = b_{ij} = \begin{cases} Z \prod_{k=j+1}^{i-1} L_k K_k & i - j > 1 \\ Z K_j & i - j = 1 \\ 0 & i - j \leq 0 \end{cases} \quad (2)$$

for all  $1 \leq i \leq n + 1$  (Durbin and Koopman, 2001). Here  $b_{ij}$  is the weight for predicting  $Y_i$  based on  $Y_j$ . Because of this dependence on the entire past, we will require some information about the behavior of the matrices  $b_{ij}$  in terms of  $i$  and  $j$ . Define  $\lambda_k(A)$  to be the  $k^{\text{th}}$  largest absolute eigenvalue of a square matrix  $A$ , and let  $\lambda_{\max}(A) = \max_k |\lambda_k(A)|$  and  $\lambda_{\min}(A) = \min_k |\lambda_k(A)|$ .

Conditions for good behavior or LTIs are governed entirely by stationarity. In particular, if  $r = \lambda_{\max}(T) < 1$ , then the sequence of predictions generated by Algorithm 1 will be stationary, and their dependence on the remote past will decay quickly. The rate of this decay depends on  $r$  and  $\rho = \lambda_{\max}(T - \overline{K}Z)$ , where  $K_i \rightarrow \overline{K}$ . The next three results show that the convergence is exponential in  $\rho$ . Proofs of these results are given in Section B.

**Lemma 7** *If  $H$  is positive definite and  $r = \lambda_{\max}(T) < 1$ , then  $\lambda_{\max}(b_{ij}) = O(r^{i-j-1})$  for any  $i > 1, j < i$ .*

While  $\widehat{P}_i = \operatorname{Var}(\alpha_{i+1} | Y_{1:i})$  changes with  $i$ , for stationary LTIs,  $\widehat{P}_i$  converges to a limiting value as  $i \rightarrow \infty$ . This means that the algorithm converges to a steady state as  $i$  grows. The next result gives the values to which the algorithm converges, and shows that this convergence occurs exponentially fast.

**Lemma 8** *If  $\lambda_{\max}(T) < 1$  then:*

1. *The solution,  $\overline{P}$  to the matrix equation  $P = TPT^\top - TPZ^\top(ZPZ^\top + H)^{-1}ZPT^\top + GQG^\top$  exists and is positive definite.*
2.  *$\widehat{P}_i \rightarrow \overline{P}$  and  $K_i \rightarrow \overline{K} = TP\overline{P}Z^\top(Z\overline{P}Z^\top + H)^{-1}$ .*



---

**Algorithm 2:** Steady-state approximate filter

---

Recursively generate approximate minimum mean squared error predictions  $\widehat{Y}_i$  using the state space model in (1).

**Input:** Initial guess for the mean of  $\alpha_1$ :  $\widehat{\alpha}_1$

**Set**  $\widehat{Y}_1 = Z\widehat{\alpha}_1$ .

**Solve**  $P = TP^T - TPZ^T(ZPZ^T + H)^{-1}ZPT^T + GQG^T$  for  $P$ , denote the solution as  $\overline{P}$

**Set**  $\overline{K} = T\overline{P}Z^T(Z\overline{P}Z^T + H)^{-1}$

**for**  $1 \leq i \leq n$  **do**

Filter

$$v_i = Y_i - \widehat{Y}_i, \quad \widehat{\alpha}_{i+1} = T\widehat{\alpha}_i + \overline{K}v_i,$$

Predict  $\widehat{Y}_{i+1} = Z\widehat{\alpha}_{i+1}$ .

**end**

**return**  $\widehat{Y}_{1:n+1}$

---

3.  $0 < \lambda_{\max}(T - \overline{K}Z) = \rho < 1$ .

4. For any matrix norm  $\|\cdot\|$ ,  $\|\widehat{P}_i - \overline{P}\| = O(\rho^i)$ ,  $\|F_i - \overline{F}\| = O(\rho^i)$ , and  $\|K_i - \overline{K}\| = O(\rho^i)$ .

Since the Kalman filter algorithm converges quickly to a steady state, one could instead use Algorithm 2 which approximates Algorithm 1 but is more computationally efficient.

For Algorithm 2, we can similarly write predictions as linear functions of previous observations. In this case,  $\widehat{Y}_{1:n+1} = \mathbf{S}Y_{1:n}$  where the prediction weights are given by

$$(\mathbf{S})_{ij} = s_{ij} = s_{i-j-1} = \begin{cases} Z(T - \overline{K}Z)^{i-j-1}\overline{K} & i - j > 0 \\ 0 & i - j \leq 0 \end{cases} \quad (3)$$

To arrive at this formulation, simply substitute  $L_k = (T - \overline{K}Z)$  and  $K_j = \overline{K}$  in (2). Notice in particular that the weights depend only on the difference  $i - j$  for this algorithm. The next result shows that the prediction weights for the two algorithms converge rapidly.

**Lemma 9** *If  $\lambda_{\max}(T) < 1$  and  $H$  positive definite. Then,  $\|s_{ij}\| = O(\rho^{i-j-1})$  and for  $j < i$ ,  $\|b_{ij} - s_{ij}\| = O(\rho^j)$ .*

We will refer to the class of predictors given by Algorithm 1 as  $\mathcal{F}_1$  and those given by Algorithm 2 as  $\mathcal{F}_2$ .

## 5. Risk Bounds

With the relevant background in place, we can put the pieces together to derive our results. We use  $\beta$ -mixing to find out how much information is in the data and VC dimension to measure the capacity of the state-space model's prediction functions. The result is a

bound on the generalization error of the chosen function  $\hat{f}$ . After slightly modifying the definition of “risk” to fit the time-series forecasting scenario and stating necessary technical assumptions, we derive risk bounds for traditional time-series forecasting models.

### 5.1 Setup and Assumptions

We observe a finite subsequence of random vectors  $Y_{1:n}$  from a process  $Y_\infty$  defined on a probability space  $(\Omega, \Sigma, \mathbb{P}_\infty)$ , with  $Y_i \in \mathbb{R}^p$ . We make the following assumption on the process.

**Assumption A**  $\mathbb{P}_\infty$  is a stationary,  $\beta$ -mixing process with mixing coefficients  $\beta_a, \forall a > 0$ .<sup>1</sup>

Under stationarity, the marginal distribution of  $Y_t$  is the same for all  $t$ . We deal mainly with the joint distribution of  $Y_{1:n+1}$ , where we observe the first  $n$  observations and try predicting  $Y_{n+1}$ . For the remainder of this paper, we will call this joint distribution  $\mathbb{P}$ . Our results extend to predicting more than one step ahead, but the notation becomes cumbersome.

We must define generalization error and training error slightly differently for time series than in the i.i.d. setting. Using the same notion of loss functions as before, we consider prediction functions  $f : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^p$ . The function  $f$  may use some or all of the past to generate predictions. A function using only the most recent  $d$  observations as inputs will be said to have *fixed memory* of length  $d$ . Other functions, in particular, the linear time-invariant systems we discuss below, have *growing memory* which means that  $f$  may use all the previous data to predict the next data point. These concepts require us to state with some care what we mean by prediction functions, and by time-series training error and risk.

#### Definition 10 (Time-series risk)

$$R_n(f) := \mathbb{E} \left[ \ell(Y_{n+1} - f(Y_{1:n})) \right].$$

The expectation is taken with respect to the joint distribution  $\mathbb{P}$  and therefore depends on  $n$ .

#### Definition 11 (Time-series training error)

$$\hat{R}_n(f) := \frac{1}{n-d} \sum_{i=d+1}^n \ell(Y_i - f(Y_{1:i-1})).$$

In order to make use of this single definition of training error, we let  $d \geq 0$ . In fixed memory cases—say an AR(2)— $d$  has an obvious meaning, and  $f(Y_{1:i}) = f(Y_{i-d+1:i})$  by definition of fixed memory, while with growing memory,  $d = 0$ , and we define  $Y_{1:0} := \emptyset$ .

To control the generalization error for time-series forecasting, we make one final assumption, about the possible magnitude of the losses. Specifically, we weaken the bounded loss assumption we used in Section 2 to allow for unbounded loss as long as we retain some control on moments of the loss.

---

1. In order to apply the results, one must either know  $\beta_a$  for some  $a$  or be able to estimate it with sufficient precision and accuracy. McDonald et al. (2011a, 2015) show how to estimate the mixing coefficients non-parametrically, based on a single sample from the process. However, those results (and those contained in this paper) only apply if the data generating process is  $\beta$ -mixing, an assumption that cannot be verified.

**Assumption B** Assume that for all  $f \in \mathcal{F}$  and all  $d \in \mathbb{N}$

$$Q_d(f) := \sqrt{\mathbb{E}_{\mathbb{P}} \left[ \ell(Y_{d+1} - f(Y_{1:d}))^2 \right]} \leq M < \infty.$$

Assumption B will be satisfied for  $Y_i$  sub-Gaussian, as well as other distributions with bounded second moment. These include, for instance, heavy-tailed Lévy noises where the tails of the pdf decay faster than an inverse cubic.

## 5.2 Fixed Memory

We can now state our results giving finite sample risk bounds for the problem of time-series forecasting. We begin with the fixed memory setting; the next section will allow the memory length to grow.

**Theorem 12** Suppose that Assumption A and Assumption B hold, that the model class  $\mathcal{F}$  has a fixed memory length  $d < n$ , and that we have a sample  $Y_{1:n}$ . Let  $\mu$  and  $a > d$  be integers such that  $2\mu a + d = n$ .<sup>2</sup> Then, for all  $0 < \epsilon < \frac{e^{3/2}}{\sqrt{2}}$ ,

$$\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_n(f)}{Q_d(f)} > \epsilon \right) \leq 8G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) \exp \left\{ - \frac{\mu \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\} + 2\mu\beta_{a-d},$$

where  $W(\cdot)$  is the Lambert W function.

The implications of this theorem are considerable. Given a finite sample of length  $n$ , we can say that with high probability, future prediction errors will not be much larger than our observed training errors. It makes no difference whether the model is correctly specified. This stands in stark contrast to model selection tools like AIC or BIC which appeal to asymptotics. Moreover, given a model class  $\mathcal{F}$ , we can say exactly how much data we need to have good control of the prediction risk. As the effective data size increases, the training error is a better and better estimate of the generalization error, uniformly over all of  $\mathcal{F}$ , provided  $\mu, a \rightarrow \infty$  such that  $\beta_{a-d} = o(1/\mu)$ .

The Lambert W function in the exponential term deserves some explanation. The Lambert W function is defined as the inverse of  $f(w) = w \exp w$  (c.f. Corless et al., 1996). A strictly, but only slightly, worse bound can be achieved by noting that

$$\exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right) \leq \frac{\epsilon^{8/3}}{4^{2/3}}$$

for all  $\epsilon \in [0, 1]$  (see Cortes et al., 2010, for the derivation).

The difference between expected and empirical risk is only interesting when  $R_n(f)$  exceeds  $\widehat{R}_n(f)$ . Due to the supremum preceding  $\frac{R_n(f) - \widehat{R}_n(f)}{Q_d(f)}$ , events where the training error exceeds the expected risk are irrelevant (as this term will be negative). Therefore, we are

---

2. By making appropriate modifications to the definition of the training error and some of the proof elements, one could allow  $2\mu a + d < n$ , but we obviate this issue for the sake of clarity.




Figure 1: The blocking procedure divides  $Y_{d+1:n}$  into  $2\mu$  alternating blocks  $U_j$  (green) and  $V_j$  (orange) each of length  $a$ . It ignores the first  $d$  observations  $Y_{1:d}$  (blue).

only concerned with  $0 \leq \widehat{R}_n(f) \leq R_n(f)$ . Of course, as discussed in Section 2, for most estimation procedures,  $f$  is chosen to make  $\widehat{R}_n(f)$  as small as possible.

Before we prove Theorem 12, we will state a corollary which puts the same result in a form that is sometimes easier to use.

**Corollary 13** *Under the conditions of Theorem 12, for any  $f \in \mathcal{F}$ , the following bound holds with probability at least  $1 - \eta$ , for all  $\eta > 2\mu\beta_{a-d}$ ,*

$$R_n(f) \leq \widehat{R}_n(f) + Me^2 \sqrt{\frac{\mathcal{E}(4 - \log \mathcal{E})}{2}},$$

with

$$\mathcal{E} = \frac{4 \log G(n, \mathcal{C}_{\ell_0 \mathcal{F}}) + 4 \log 8/\eta'}{\mu},$$

and  $\eta' = \eta - 2\mu\beta_{a-d}$ .

We now prove both Theorem 12 and Corollary 13 to provide the reader with some intuition for the types of arguments necessary. We defer proof of the remainder of our results to Section B.

**Proof** [of Theorem 12 and Corollary 13] The first step is to move from the actual sample size  $n$  to the effective sample size  $\mu$  which depends on the  $\beta$ -mixing behavior. Let  $a$  and  $\mu$  be non-negative integers such that  $2a\mu + d = n$ . Now divide  $Y_{d+1:n}$  into  $2\mu$  blocks, each of length  $a$ , ignoring the first  $d$  observations. Identify the blocks as follows:

$$\begin{aligned} U_j &= \{Y_i : 2(j-1)a + d + 1 \leq i \leq (2j-1)a + d\}, \\ V_j &= \{Y_i : (2j-1)a + d + 1 \leq i \leq 2ja + d\}. \end{aligned}$$

Let  $\mathbf{U}$  be the sequence of odd blocks  $U_j$ , and let  $\mathbf{V}$  be the sequence of even blocks  $V_j$ . A graphical depiction of the blocking procedure is shown in Figure 1. Finally, let  $\mathbf{U}'$  be a sequence of blocks which are mutually independent but such that each block has the same distribution as a block from the original sequence. That is construct  $U'_j$  such that

$$\mathcal{L}(U'_j) = \mathcal{L}(U_j) = \mathcal{L}(U_1),$$

where  $\mathcal{L}(\cdot)$  means the probability law of the argument.

Let  $\widehat{R}_{\mathbf{U}}(f)$ ,  $\widehat{R}_{\mathbf{U}'}$ , and  $\widehat{R}_{\mathbf{V}}(f)$  be the empirical risk of  $f$  based on the block sequences  $\mathbf{U}$ ,  $\mathbf{U}'$ , and  $\mathbf{V}$  respectively. We have

$$\begin{aligned}
 \widehat{R}_n(f) &= \frac{1}{n-d} \sum_{i=d+1}^n \ell(Y_i - f(Y_{1:i-1})) \\
 &= \frac{1}{n-d} \left[ \sum_{i:Y_i \in \mathbf{U}} \ell(Y_i - f(Y_{1:i-1})) + \sum_{j:Y_j \in \mathbf{V}} \ell(Y_j - f(Y_{1:j-1})) \right] \\
 &= \frac{1}{2} \left[ \frac{2}{n-d} \sum_{i:Y_i \in \mathbf{U}} \ell(Y_i - f(Y_{1:i-1})) + \frac{2}{n-d} \sum_{j:Y_j \in \mathbf{V}} \ell(Y_j - f(Y_{1:j-1})) \right] \\
 &= \frac{1}{2} \left[ \widehat{R}_{\mathbf{U}}(f) + \widehat{R}_{\mathbf{V}}(f) \right].
 \end{aligned}$$

Then,

$$\begin{aligned}
 &\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_n(f)}{Q_d(f)} > \epsilon \right) \\
 &= \mathbb{P} \left( \sup_{f \in \mathcal{F}} \left[ \frac{R_n(f) - \widehat{R}_{\mathbf{U}}(f)}{2Q_d(f)} + \frac{R_n(f) - \widehat{R}_{\mathbf{V}}(f)}{2Q_d(f)} \right] > \epsilon \right) \\
 &\leq \mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_{\mathbf{U}}(f)}{Q_d(f)} + \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_{\mathbf{V}}(f)}{Q_d(f)} > 2\epsilon \right) \tag{4}
 \end{aligned}$$

$$\leq \mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_{\mathbf{U}}(f)}{Q_d(f)} > \epsilon \right) + \mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_{\mathbf{V}}(f)}{Q_d(f)} > \epsilon \right) \tag{5}$$

$$= 2\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_{\mathbf{U}}(f)}{Q_d(f)} > \epsilon \right)$$

$$= 2\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_d(f) - \widehat{R}_{\mathbf{U}}(f)}{Q_d(f)} > \epsilon \right). \tag{6}$$

Here, (4) follows by the convexity of the supremum and (5) by a union bound. Now, for (6), as  $f$  has fixed memory  $d$ , we have that  $f(Y_{1:j}) = f(Y_{j-d+1:j})$ . Thus

$$\begin{aligned}
 R_n(f) &= \mathbb{E}[\ell(Y_{n+1} - f(Y_{1:n}))] = \mathbb{E}[\ell(Y_{n+1} - f(Y_{n-d+1:n}))] \\
 &= \mathbb{E}[\ell(Y_{d+1} - f(Y_{1:d}))] = R_d(f)
 \end{aligned}$$

by stationarity and, similarly,  $Q_d(f) = \sqrt{\mathbb{E}[\ell(Y_{d+1} - f(Y_{1:d}))^2]}$ . Thus  $R_d(f)$  and  $Q_d(f)$  depend on only  $d+1$  data values. Likewise, a ‘‘point’’ in the training error  $\widehat{R}_{\mathbf{U}}(f)$  depends on  $d+1$  data values. Therefore, a prediction at any  $Y_i$  in some block  $U_j$  is separated by at least  $a-d$  observations from any  $Y_{i'}$  in different block  $U_{j'}$ . Furthermore, we are estimating an expectation  $R_d(f)$  which depends on  $d+1$  values with an empirical expectation  $\widehat{R}_{\mathbf{U}}(f)$  which is a dependent sum of components each of which depends on  $d+1$  observed data values.

Therefore, we can apply Lemma 4.1 in Yu (1994) (reproduced as Lemma 21 in Section A) to the event  $\left\{ \sup_{f \in \mathcal{F}} \frac{R_d(f) - \widehat{R}_{\mathbf{U}}(f)}{Q_d(f)} > \epsilon \right\}$ . This allows us to move from statements about the dependent blocks in  $\widehat{R}_{\mathbf{U}}(f)$  to statements about the independent blocks in  $\widehat{R}_{\mathbf{U}'}(f)$  with a slight correction which accounts for the worst-case dependence between adjacent blocks:  $\beta_{a-d}$ . Therefore,

$$2\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_d(f) - \widehat{R}_{\mathbf{U}}(f)}{Q_d(f)} > \epsilon \right) \leq 2\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_d(f) - \widehat{R}_{\mathbf{U}'}(f)}{Q_d(f)} > \epsilon \right) + 2\mu\beta_{a-d}, \quad (7)$$

where the probability on the right of (7) is for the  $\sigma$ -field generated by the independent block sequence  $\mathbf{U}'$ . Let us now introduce the growth function for the class of “blocked” predictor losses  $\mathcal{C}_{\ell_{\circ}\mathcal{F}}^{\Sigma} = \{C_{\ell_{\circ}\mathcal{F}}^{\Sigma} : f \in \mathcal{F}\}$ , where<sup>3</sup>

$$\mathcal{C}_{\ell_{\circ}\mathcal{F}}^{\Sigma} = \left\{ (\mathbf{u}, \mathbf{d}, b) : \mathbf{1}_{[0, \infty)} \left( \sum_{j=1}^a \ell(d_j - f(u_j)) - b \right) = 1, \mathbf{u} \in \mathbb{U}^a, \mathbf{d} \in \mathbb{R}^{k \times a}, b \in \mathbb{R} \right\}.$$

This allows us to state that

$$2\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_d(f) - \widehat{R}_{\mathbf{U}'}(f)}{Q_d(f)} > \epsilon \right) \leq 8G(2\mu, \mathcal{C}_{\ell_{\circ}\mathcal{F}}^{\Sigma}) \exp \left\{ -\frac{\mu \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\}, \quad (8)$$

where we have applied Corollary 23 to bound the independent blocks  $\mathbf{U}'$ . Since there are  $\mu$  independent blocks, this upper bound is in terms of  $\mu$  effectively independent data “points” penalized by the correction  $\beta_{a-d}$  which adjusts for the worst case dependence between “points” rather than  $n$  dependent data points. However, the growth function is in terms of the class of blocked predictors and depends on  $a$ . To remove this dependence, we present a novel result, Lemma 24. This lemma shows that  $G(2\mu, \mathcal{C}_{\ell_{\circ}\mathcal{F}}^{\Sigma}) \leq G(2\mu a, \mathcal{C}_{\ell_{\circ}\mathcal{F}}) \leq G(n, \mathcal{C}_{\ell_{\circ}\mathcal{F}})$ , giving the desired result.

To prove the corollary, set the right hand side of (8) to  $\eta$ , take  $\eta' = \eta - 2\mu\beta_{a-d}$ , and solve for  $\epsilon$ . We get that with probability at least  $1 - \eta$ , for all  $f \in \mathcal{F}$ ,

$$\frac{R_n(f) - \widehat{R}_n(f)}{Q_d(f)} \leq \epsilon.$$

Solving the equation

$$\eta' = 8G(n, \mathcal{C}_{\ell_{\circ}\mathcal{F}}) \exp \left\{ -\frac{\mu \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\}$$

implies

$$\epsilon = e^2 \sqrt{\frac{\mathcal{E}(4 - \log \mathcal{E})}{2}}$$

---

3. The  $\Sigma$  indicates a sum over the coordinates of the block.

with

$$\mathcal{E} = \frac{4 \log G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) + 4 \log 8/\eta'}{\mu}.$$

■

The only obstacle to the use of Theorem 12 is knowledge of the  $G(n, \mathcal{C}_{\ell \circ \mathcal{F}})$  or  $\text{VCD}(\mathcal{C}_{\ell \circ \mathcal{F}})$ . For some models, these can be calculated explicitly.

**Theorem 14** *If  $\ell(y - y')$  is monotone increasing in  $|y - y'|$ , then for the class of  $AR(d)$  models,*

$$G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) \leq \left( \frac{2en}{d+1} \right)^{d+1}.$$

*If  $\ell(y - y')$  is monotone increasing in each coordinate of  $|y_j - y'_j|$ , then for the class of vector autoregressive models of order  $d$  with  $p$  time-series,*

$$G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) \leq \left( \frac{2en}{d+1} \right)^{p(d+1)}.$$

The proof of Theorem 14 is given in Section B. We note that this result applies equally to Bayesian VARs, however, this is likely conservative as the prior tends to restrict the effective complexity of the function class.<sup>4</sup>

### 5.3 Bounds for LTIs

As discussed in Section 4, our goal is to derive similar bounds for linear time-invariant dynamical systems which produce forecasts via Algorithm 1 or Algorithm 2. We begin with a result for the simple case with predictions generated by Algorithm 2.

**Theorem 15** *Suppose that Assumption A and Assumption B hold, and that the model class  $\mathcal{F}_2$  is generated by Algorithm 2 with  $\lambda_{\max}(T) < 1$ . Further assume that the loss function  $\ell$  is a norm and let  $\ell^*(A) = \sup_{z \neq 0} \ell(Az)/\ell(z)$  be the matrix norm induced by  $\ell$ . Given a time-series of length  $n$ , fix some  $1 \leq d < n$ , and let  $\mu$  and  $a$  be integers such that  $2\mu a + d = n$ .*

---

4. Here we should mention that these risk bounds are frequentist in nature. We mean that if one treats Bayesian methods as a regularization technique and predicts with the posterior mean or mode, then our results hold. However, from a subjective Bayesian perspective, our results add nothing since all inference can be derived from the posterior. For further discussion of the frequentist risk properties of Bayesian methods under mis-specification, see for example Kleijn and van der Vaart (2006); Müller (2013) or Shalizi (2009).

Then

$$\begin{aligned} & \mathbb{P} \left( \sup_{f_2 \in \mathcal{F}_2} \frac{R_n(f_2) - \widehat{R}_n(f_2) - \delta_d(f_2)}{Q_d(f_2)} > \epsilon \right) \\ & \leq 8G(n, \mathcal{C}_{\ell \circ \mathcal{F}'_2}) \exp \left\{ -\frac{\mu \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\} + 2\mu\beta_{a-d}, \\ & \leq 8 \left( \frac{2en}{d+1} \right)^{p(d+1)} \exp \left\{ -\frac{\mu \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\} + 2\mu\beta_{a-d}, \end{aligned}$$

where

$$\delta_d(f_2) = \frac{1}{n-d} \sum_{i=d+1}^n \left[ \ell \left( Y_i - \sum_{j=i-d}^{i-1} s_{ij} Y_j \right) - \ell \left( Y_i - \sum_{j=1}^{i-1} s_{ij} Y_j \right) \right] + \mathbb{E}[\ell(Y_1)] \sum_{j=1}^{n-d} \ell^*(s_{nj}).$$

The  $\delta_d(f_2)$  term arises from taking a fixed-memory approximation, of length  $d$ , to predictors with growing memory<sup>5</sup>. As becomes clear in the proof (see Section B), we make this approximation to apply the previous theorem, but it involves a trade-off. As  $d \nearrow n$ ,  $\delta_d(f_2) \searrow 0$ , but this drives  $\mu \searrow 0$ , resulting in fewer effective training points, whereas smaller  $d$  has the opposite effect. The summation inside square braces on the left is the difference between empirical risks for  $f_2$  and that of the truncated predictor  $f'_2$  which uses only the most recent  $d$  data values. That is

$$\frac{1}{n-d-1} \sum_{i=d}^{n-1} \left[ \ell \left( Y_{i+1} - \sum_{j=i-d+1}^i s_{ij} Y_j \right) - \ell \left( Y_{i+1} - \sum_{j=1}^i s_{ij} Y_j \right) \right] = \widehat{R}_n(f'_2) - \widehat{R}_n(f_2).$$

This term is easily calculated from the data. Also,  $\delta_d(f_2)$  depends on  $\mathbb{E}[\ell(Y_1)]$  which is not necessarily desirable. However, Assumption B has the consequence that  $\mathbb{E}[\ell(Y_1)] \leq M < \infty$  as long as  $\mathcal{F}_2$  allows  $s_{ij} = 0$ . Finally, we reiterate that  $s_{ij} = s_{i-j-1}$  is a function only of the difference between  $i$  and  $j$ . Finally, note that the upper bound depends on the growth function of the truncated class  $\mathcal{F}'_2$ , which can be bounded using Theorem 14.

**Corollary 16** *Under the conditions of either Theorem 15, for any  $f_2 \in \mathcal{F}_2$ , with probability at least  $1 - \eta$ ,*

$$R_n(f_2) \leq \widehat{R}_n(f_2) + \delta_d(f_2) + Me^2 \sqrt{\frac{\mathcal{E}(4 - \log \mathcal{E})}{2}},$$

where

$$\mathcal{E} = \frac{4 \log G(n, \mathcal{C}_{\ell \circ \mathcal{F}'_2}) + 4 \log 8/\eta'}{\mu},$$

and  $\eta' = \eta - 2\mu\beta_{a-d}$ . Furthermore,

$$\delta_d(f_2) = O_{\mathbb{P}}(\rho^d).$$

---

5. There are several ways one could make such an approximation. To simplify the proof, we have set all coefficients at lags beyond  $d$  to zero rather than, e.g., asking for the  $d$ -memory linear predictor coming closest in  $L_2$  to the infinite-memory predictor.



For the case of Algorithm 1, Assumption B no longer makes sense as it fails to reflect the fact each  $f_1 \in \mathcal{F}_1$  is a time varying function of previous observations as illustrated in (2). We therefore strengthen the assumption slightly. Notice in particular that  $f'_1$  depends on  $n$  as illustrated in (2).

**Assumption C** For all  $f'_1$  generated by truncating Algorithm 1 to depend on the most recent  $d$  observations, and all  $n > d$ ,

1.

$$Q_d(f'_1) := \sqrt{\mathbb{E}_{\mathbb{P}} \left[ \ell(Y_{n+1} - f'_1(Y_{n-d+1:n}))^2 \right]} \leq M < \infty.$$

2. For every  $\epsilon > 0$  there exists  $\delta > 0$  such that, for all measurable sets  $A$ , with  $\mathbb{P}(A) \leq \delta$ , and every  $f'_1$ ,

$$\int_A d\mathbb{P} \left[ \ell(Y_{n+1} - f'_1(Y_{n-d+1:n}))^2 \right] \leq \epsilon.$$

The first part of this assumption is analogous to Assumption B but for predictors  $f_1 \in \mathcal{F}_1$ . However, we need only impose the assumption on a fixed-memory version of these growing memory predictors. Note also that for the case of Algorithm 2, as in Theorem 15, for any  $n$ ,  $f_2 \in \mathcal{F}_2$ , we have  $Q_d(f'_2) = \sqrt{\mathbb{E}_{\mathbb{P}} \left[ \ell(Y_n - f'_2(Y_{n-d:n-1}))^2 \right]}$ . For Algorithm 1, this is no longer the case (as  $f'_1$  depends on  $n$ ). However, the dependence on  $n$  will not be necessary. The second part of this assumption ensures that the loss of the truncated versions of these predictors is uniformly integrable. Under this assumption, since predictions  $f'_1(Y_{n-d+1:n}) - f'_2(Y_{n-d+1:n}) = O_{\mathbb{P}}(1)$ , we have that  $\lim_{n \rightarrow \infty} Q_d(f'_1) = Q_d(f_2)$ . This statement is made rigorous in the proof.

**Theorem 17** Suppose that Assumption A and Assumption C hold, and that the model class  $\mathcal{F}_1$  is generated by Algorithm 1 with  $\lambda_{\max}(T) < 1$ . Further assume that the loss function  $\ell$  is a norm and let  $\ell^*(A) = \sup_{z \neq 0} \frac{\ell(Az)}{\ell(z)}$  be the matrix norm induced by  $\ell$ . Given a time-series of length  $n$ , fix some  $1 \leq d < n$ , and let  $\mu$  and  $a$  be integers such that  $2\mu a + d = n$ . Then

$$\begin{aligned} & \mathbb{P} \left( \sup_{f_1 \in \mathcal{F}} \frac{R_n(f_1) - \widehat{R}_n(f_1) - \delta_d(f_1)}{\lim_{n \rightarrow \infty} Q_d(f'_1)} > \epsilon \right) \\ & \leq 8G(n, \mathcal{C}_{\ell \circ \mathcal{F}'_2}) \exp \left\{ - \frac{\mu \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\} + 2\mu\beta_{a-d}, \\ & \leq 8 \left( \frac{2en}{d+1} \right)^{p(d+1)} \exp \left\{ - \frac{\mu \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\} + 2\mu\beta_{a-d}, \end{aligned}$$

where

$$\begin{aligned} \delta_d(f_1) &= \frac{1}{n-d} \sum_{i=d+1}^n \left[ \ell \left( Y_i - \sum_{j=i-d}^{i-1} b_{ij} Y_j \right) - \ell \left( Y_i - \sum_{j=1}^{i-1} s_{i-j-1} Y_j \right) \right] \\ &+ \mathbb{E}[\ell(Y_1)] \sum_{j=1}^{n-d} \ell^*(b_{nj}). \end{aligned}$$

For this result, the summation inside square braces on the left is the difference between empirical risk for  $f_1$  and that of the truncated steady-state predictor  $f'_2$  which uses only the most recent  $d$  data values. That is

$$\frac{1}{n-d} \sum_{i=d+1}^n \left[ \ell \left( Y_i - \sum_{j=i-d}^{i-1} b_{ij} Y_j \right) - \ell \left( Y_{i+1} - \sum_{j=1}^i s_{ij} Y_j \right) \right] = \widehat{R}_n(f_1) - \widehat{R}_n(f'_2)$$

where  $f'_2$  is given by Algorithm 2. Again, this term is easily calculated from the data. Furthermore, the upper bound depends on the growth function of the level sets of  $\mathcal{F}'_2$  which is easily bounded using Theorem 14.

**Corollary 18** *Under the conditions of Theorem 17, for any  $f_1 \in \mathcal{F}_1$ , with probability at least  $1 - \eta$ ,*

$$R_n(f_1) \leq \widehat{R}_n(f_1) + \delta_d(f_1) + M e^2 \sqrt{\frac{\mathcal{E}(4 - \log \mathcal{E})}{2}},$$

where

$$\mathcal{E} = \frac{4 \log G(n, \mathcal{C}_{\ell \circ \mathcal{F}'_2}) + 4 \log 8/\eta'}{\mu},$$

and  $\eta' = \eta - 2\mu\beta_{a-d}$ . Furthermore,

$$\delta_d(f_1) = O_{\mathbb{P}}(r^d) + O(\rho^{n-d+1}) + O_{\mathbb{P}}((n-d)^{-1}).$$

This corollary shows that  $\delta_d(f_1)$  decays rapidly as long as  $d \rightarrow \infty$  and  $d/n \rightarrow 0$ . But for LTIs, it is simple to compute  $\delta_d(f_1)$  or  $\delta_d(f_2)$  using results from Algorithm 1 or Algorithm 2, so this and Corollary 16 let us compute risk bounds for common time-series forecasting models without appealing to the asymptotic form of  $\delta_d(f)$ .

As this bound is rather complicated, in that  $d$ ,  $a$ ,  $\mu$  and  $\eta$  are expected to go to either 0 or  $\infty$  as functions of  $n$ , we provide the following simple example demonstrating convergence.

**Example** *Suppose  $Y_\infty$  is a 1-dimensional geometrically  $\beta$ -mixing time series with  $\beta_a = O(k^a)$  for some constant  $0 < k < 1$ . We will choose as our predictor the state-space model as in (1)*

$$\begin{aligned} Y_i &= Z\alpha_i + H\epsilon_i, \\ \alpha_{i+1} &= T\alpha_i + G\eta_{i+1}, \end{aligned}$$

with a univariate state  $\alpha$ . We will maximize the likelihood to estimate  $\theta = (Z, H, T, G)$  and apply Algorithm 1 to generate forecasts from  $n$  samples. Thus

$$\mathcal{F}_1 = \{f_1 : Z \in R, H, G > 0, -1 < T < 1\}.$$

To achieve reasonable bounds, we fix  $0 < \gamma < 1$  and choose  $\mu = (n-d)^{1-\gamma}$ ,  $a = 0.5(n-d)^\gamma$ , and  $\eta = 3\mu\beta_{a-d}$ . With these choices, and using  $C$  as a mutable constant independent of  $n$  and  $d$ ,

$$\mathcal{E} = C \frac{4 \log(d+1) \log(2en/(d+1)) - (1-\gamma) \log(n-d) + [(n-d)^\gamma - d] \log(1/k)}{(n-d)^{1-\gamma}}.$$

Here, we have used Theorem 14 to bound the logarithm of the growth function of the growing memory predictor with  $\log(d+1) \log(2en/(d+1))$  for truncation length  $d$ . Therefore, for  $n$  and  $d$  such that  $(n-d)^\gamma > C \log(n) \log(d)$ ,  $\mathcal{E} < C(n-d)^{-(1-2\gamma)}$ . To complete the bound, we need to choose  $d$  such that  $\eta \rightarrow 0$ ,  $r^d \rightarrow 0$ , and  $n-d \rightarrow \infty$ . To achieve the first, we have

$$\eta = \mu\beta_{a-d} = g(n-d)^{1-\gamma} k^{0.5(n-d)^\gamma - d}.$$

Therefore, we require  $d = o(n^\gamma)$ . Thus, as long as  $d = o(n^\gamma)$  and  $d \rightarrow \infty$ , for  $n$  large enough (depending on many constants), Corollary 18 gives

$$R_n(f_1) \leq \widehat{R}_n(f_1) + C \sqrt{\frac{1}{n^{1-2\gamma}}}$$

for some  $C$  with high probability.

## 6. Bounds in Practice

We now show how the theorems of the previous section can be used both to quantify prediction risk and to select models. We first estimate a simple stochastic volatility model using IBM return data and calculate the bound for the predicted volatility using Corollary 18. Then we show how the same methods can be used for typical macroeconomic forecasting models.

### 6.1 Stochastic Volatility Model

We estimate a standard stochastic volatility model using daily log returns for IBM from January 1962 until October 2011 ( $n = 12541$  observations). Figure 2 shows the squared log-return series.

The model we investigate is

$$\begin{aligned} Y_i &= \sigma z_i \exp(\alpha_i/2), & z_i &\sim \text{N}(0, 1), \\ \alpha_{i+1} &= \phi \alpha_i + \eta_i, & \eta_i &\sim \text{N}(0, \sigma_\eta^2), \end{aligned}$$

where the disturbances  $z_i$  and  $w_i$  are mutually and serially independent. Following Harvey et al. (1994), we linearize this by taking the natural log of the observation equation:

$$\begin{aligned} \log Y_i^2 &= \varsigma + \frac{1}{2} \alpha_i + \epsilon_i, & \epsilon_i &= \log z_i^2 - \xi, \\ \varsigma &= \log \sigma^2 + \xi, & \xi &= \mathbb{E}[\log z_1^2]. \end{aligned}$$

The noise term  $\epsilon_i$  is no longer Gaussian (it has a shifted log-gamma distribution), but the Kalman filter will still give the minimum-mean-squared-error linear estimate of the variance sequence  $\alpha_{1:n+1}$ . The observation variance is now  $\pi^2/2$ .

To match the data to the model, let  $y_i$  be the log returns and remove 688 observations where the return was 0 (the price did not change from one day to the next). Using the Kalman filter, the negative log likelihood is given by

$$\mathcal{L}(Y_{1:n} | \varsigma, \phi, \sigma_\eta^2) \propto \sum_{i=1}^n \log F_i + v_i^2 F_i^{-1}.$$




Figure 2: Daily realized volatility (squared log returns) for IBM from 1962–2011.

Model	Training error	AIC-Baseline	Risk bound ( $1 - \eta > 0.95$ )	Effective VC
SV	1.44	-2816	15.92	3
AR(2)	1.49	-348	15.00	3
Mean	1.51	0	12.63	1

Table 1: This table shows the training error and risk bounds for 3 models. AIC is given as the difference from predicting with the global mean (the smaller the value, the more support for that model). The “Effective VC” dimension column reports the exponent in the bound on the growth function. This is slightly different from the VC dimension of the model.

Minimizing this gives estimates  $\varsigma = -9.62$ ,  $\phi = 0.996$ , and  $\sigma_\eta^2 = 0.003$ . Using absolute error loss gives training error  $\widehat{R}_n(f) = 1.44$ .

To actually calculate the bound, we need a few more values. First, using the methods in McDonald et al. (2011a, 2015), we can estimate  $\beta_8 = 0.017$ . For  $a > 8$ , the optimal point estimate of  $\beta_a$  is 0. While this is presumably an underestimate, we will take  $\beta_a = 0$  for  $a > 8$ . For the upper bound in Assumption B, we use  $M = \sqrt{\pi^2/2}$  which corresponds to the SV model being true.

Combining these values with the VC dimension for the stochastic volatility model, we can bound the prediction risk. Finally, we take  $\mu = 538$ ,  $a = 11$ ,  $d = 2$ , and  $\mathbb{E}[|Y_1|] = M$ . The result is the bound  $R_n(f) \leq 15.92$  with probability at least 0.95. In other words, the bound is much larger than the training error, but this is to be expected: the data are highly dependent, so the large  $n$  translates into a relatively small effective sample size  $\mu$ .

For comparison, we also computed the bound for forecasts produced with an AR(2) model (with intercept) and with the global mean alone. In the case of the mean, we take  $\mu = 658$  and  $a = 9$  since in this case,  $d = 0$ . The results are shown in Table 1. The stochastic volatility model reduces the training error by 5% relative to predicting with the mean, an improvement which is marginal at best. But the resulting risk bound clearly demonstrates that given the small effective sample size, even this gain may be spurious: it is likely that the stochastic volatility model is simply over-fitting.




Figure 3: Time series used to estimate the RBC model. These are quarterly data from 1948:I until 2010:I. The blue line is GDP (output), the red line is consumption, the green line is investment, and the orange line is hours worked. These data are plotted as percentage deviations from trend as discussed in the Online Appendix.

## 6.2 Real Business Cycle Model

In this section, we will discuss the methodology for applying risk bounds to the forecasts generated by the real business cycle (RBC) model. This is a standard tool in macroeconomic forecasting. For a discussion of the RBC model and the standard methods used to bring such models to the data, see, for example DeJong and Dave (2007); DeJong et al. (2000); Fernández-Villaverde (2010); Kydland and Prescott (1982); Romer (2011); Sims (2002); Smets and Wouters (2007).

To estimate the parameters of this model, we use four data series. These are GDP  $y_t$ , consumption  $c_t$ , investment  $i_t$ , and hours worked  $n_t$  which are from the Federal Reserve Economic Database. The series we use are shown in Figure 3.

The basic idea of the estimation is to transform the model from an inter-temporal optimization form into a state space model. This leads to a linear, Gaussian state-space model with four observed variables (listed above), and two unobserved state variables. The mapping from parameters of the optimization problem to parameters of the state-space model is nonlinear, but, for each parameter setting, the Kalman filter returns the likelihood, so that likelihood methods are possible. As the data are fairly uninformative about many of the parameters, we estimate by maximizing a penalized likelihood, rather than a simple likelihood. Then the Kalman filter produces in-sample forecasts which are linear in past values of the data, so that we could potentially apply the growing memory bound.

For macroeconomic time series, there is not enough data to give nontrivial bounds, regardless of the mixing coefficients or the size of the finite memory approximation. Figure 3

shows  $n = 245$  observations. The minimal possible finite approximation model is a VAR with one lag and four time series. In this case, since we are dealing with vector valued forecasts, we take  $\ell(y - y') = \|y - y'\|_2$  so that the induced matrix norm of  $A$  is  $\ell^*(A) = \sigma_{\max}(A)$ , the largest singular value of  $A$ . We assume that Assumption C is satisfied with  $M = 0.1$  and demand confidence 0.95 ( $\eta = 0.05$ ). By Theorem 14, this class has growth function bounded by  $(en)^8$ .

Again, using the methods of McDonald et al. (2011a), we can estimate the  $\beta$ -mixing coefficients of the macroeconomic data set. The result is a point estimate  $\beta_4 = 0$ . Assuming that this is approximately accurate (0 is of course an underestimate), this suggests that the effective size of the macroeconomic data set is no more than about  $\mu = 25$ , much smaller than  $n = 245$ . To calculate the bound, we assume that  $\mathbb{E}[\|Y\|_2] < 0.1$ . The training error of the fitted RBC model is  $\widehat{R}_n(f) = 0.046$ . The bound is  $R_n(f) < 2.29$ .

The bound here is three orders of magnitude larger than the training error. If the bound is tight, then this suggests that the training error severely underestimates the true prediction risk. Of course, this should not be too surprising since the RBC model has 11 parameters and we are trying to get confidence intervals using only 25 effective data points.

In some sense, the empirical results in this section may seem slightly unreasonable. Since the results are only upper bounds, it is important to get an idea as to how tight they may be. We address this issue in the next section.

## 7. Properties of Our Results

In the previous section, we showed that the upper bound for the risk of standard macroeconomic forecasting models may be large. This of course raises the question “How tight are these bounds?” We address this question next and then discuss how to use the bounds for model selection.

### 7.1 How Tight Are the Bounds?

Here we give some idea of how tight the bounds presented in Section 5 are. Denote the function that minimizes the training error (or penalized training error) over  $\mathcal{F}$  by  $\widehat{f}_{erm}$ , and let

$$f^* = \operatorname{argmin}_{f \in \mathcal{F}} R_n(f)$$

be the oracle predictor with respect to  $\mathcal{F}$ . We call

$$L_n(\Pi) := \sup_{\mathbb{P} \in \Pi} \mathbb{E}_{\mathbb{P}}[R_n(\widehat{f}_{erm}) - R_n(f^*)] = \sup_{\mathbb{P} \in \Pi} \mathbb{E}_{\mathbb{P}}[R_n(\widehat{f}_{erm})] - R_n(f^*)$$

the “oracle loss”; it describes how well empirical risk minimization works relative to the best possible predictor  $f^*$  over the worst distribution  $\mathbb{P}$ . Vapnik (1998) shows that for classification and i.i.d. data, for sufficiently large  $n$ , there exist constants  $\Upsilon_*$  and  $\Upsilon^*$  such that

$$\Upsilon_* \sqrt{\frac{\operatorname{VCD}(\mathcal{C}_{\mathcal{F}})}{n}} \leq L_n(\Pi) \leq \Upsilon^* \sqrt{\frac{\operatorname{VCD}(\mathcal{C}_{\mathcal{F}}) \log n}{n}},$$

where  $\Pi$  is the class of all distributions on  $\mathbb{U} \times \{0, 1\}$  and  $\mathcal{C}_{\mathcal{F}}$  has finite VC-dimension  $\text{VCD}(\mathcal{C}_{\mathcal{F}})$ . In other words, for i.i.d. data, the best we can hope to do is a rate of

$$O\left(\sqrt{\frac{\text{VCD}(\mathcal{C}_{\mathcal{F}})}{n}}\right)$$

, and prediction methods which perform worse than  $O\left(\sqrt{\text{VCD}(\mathcal{C}_{\mathcal{F}}) \log n/n}\right)$  are inefficient. We will derive similar bounds for the  $\beta$ -mixing setting. First, we need a slightly different version of Theorem 12.

**Theorem 19** *Suppose that  $\ell(y - y') < M$ , that Assumption A holds, and that  $\mathcal{F}$  has a fixed memory length  $d < n$ . Let  $\mu$  and  $a$  be integers such that  $2\mu a + d \leq n$ . Then, for all  $\epsilon > 0$ ,*

$$\mathbb{P}\left(\sup_{f \in \mathcal{F}} (R_n(f) - \widehat{R}_n(f)) > \epsilon\right) \leq 8G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) \exp\left\{-\frac{\mu \epsilon^2}{\Upsilon}\right\} + 2\mu \beta_{a-d}.$$

where  $\Upsilon$  depends only on  $M$ .

The proof of Theorem 19 is exactly like that for Theorem 12 using the result for bounded loss in Theorem 4 rather than that for relative loss in Corollary 23 to control the probability that empirical risk over blocks deviates from its expectation (see (8)).

**Assumption D** *The time series  $Y_\infty$  is geometrically  $\beta$ -mixing, that is  $\beta_a = v_1 \exp(-v_2 a^\kappa)$  for some constants  $v_1, v_2, \kappa$ .*

**Theorem 20** *Suppose  $\ell(y - y') < M$  and that Assumption D holds. Further assume that  $\mathcal{C}_{\ell \circ \mathcal{F}}$  has finite VC-dimension  $\text{VCD}(\mathcal{C}_{\ell \circ \mathcal{F}})$ . Then, for sufficiently large  $n$ , there exist constants  $\Upsilon_*$  and  $\Upsilon^*$ , independent of  $n$  and  $\text{VCD}(\mathcal{C}_{\ell \circ \mathcal{F}})$ , such that*

$$\Upsilon_* \sqrt{\frac{\text{VCD}(\mathcal{C}_{\ell \circ \mathcal{F}})}{n}} \leq L_n(\Pi) \leq \Upsilon^* \sqrt{\frac{\text{VCD}(\mathcal{C}_{\ell \circ \mathcal{F}}) \log n}{n^{\kappa/(1+\kappa)}}}.$$

If we instead assume *algebraic mixing*, that is  $\beta_a = va^{-\rho}$ , then we can retrieve the same rate where  $0 < \kappa < (\rho - 1)/2$  (see Meir, 2000). Theorem 20 says that in dependent data settings, using the blocking approach developed here, we may pay a penalty: the upper bound on  $L_n(\Pi)$  goes to zero more slowly than in the i.i.d. case. But, the lower bound cannot be made any tighter since i.i.d. processes are still allowed under Assumption D (and of course under the more general Assumption A). In other words, we may have  $\kappa \rightarrow \infty$  so we can not rule out the faster learning rate of  $O\left(\sqrt{\frac{\text{VCD}(\mathcal{C}_{\ell \circ \mathcal{F}}) \log n}{n}}\right)$ .

## 7.2 Structural Risk Minimization

Our presentation so far has focused on choosing one function  $\widehat{f}$  from a model  $\mathcal{F}$  and demonstrating that the prediction risk  $R_n(\widehat{f})$  is well characterized by the training error inflated by a complexity term. The procedure for actually choosing  $\widehat{f}$  has been ignored. Common ways of choosing  $\widehat{f}$  are frequently referred to as *empirical risk minimization* or ERM: approximate the expected risk  $R_n(f)$  with the empirical risk  $\widehat{R}_n(f)$ , and choose  $\widehat{f}$  to minimize

the empirical risk. Many likelihood based methods have exactly this flavor. But more frequently, forecasters have many different models in mind, each with a different empirical risk minimizer. Regularized model classes (ridge regression, lasso, Bayesian methods) implicitly have this structure — altering the amount of regularization leads to different models  $\mathcal{F}$ . Or one may have many different forecasting models from which the forecaster would like to choose the best. This scenario leads to a generalization of ERM called *structural risk minimization* or SRM.

Given a collection of models  $\mathcal{F}_1, \mathcal{F}_2, \dots$  each with associated empirical risk minimizers  $\hat{f}_1, \hat{f}_2, \dots$ , we wish to use the function which has the smallest risk. Of course different models have different complexities, and those with larger complexities will tend to have smaller empirical risk. To choose the best function, we therefore penalize the empirical risk and select that function which minimizes the penalized version. Model selection tools like AIC or BIC have exactly this form, but they rely on specific knowledge of the data likelihood and use asymptotics to derive approximate penalties. In contrast, we have finite-sample bounds for the expected risk. This leads to a natural model selection rule: choose the predictor which has the smallest bound on the expected risk.

The generalization error bounds in Section 5 allow one to perform model selection via the SRM principle without knowledge of the likelihood or appeals to asymptotic results. The penalty accounts for the complexity of the model through the VC dimension. Most useful however is that by using generalization error bounds for model selection, we are minimizing the prediction risk. So in the volatility forecasting exercise above, we would choose the mean.

If we want to make the prediction risk as small as possible, we can minimize the generalization error bound simultaneously over models  $\mathcal{F}$  and functions within those models. This amounts to treating VC dimension as a control variable. Therefore, by minimizing both the empirical risk and the VC dimension, we can choose that model and function which has the smallest prediction risk, a claim which other model selection procedures cannot make (Massart, 2007; Vapnik, 2000).

## 8. Conclusion

This paper demonstrates how to control the generalization error of common time-series forecasting models, especially those used in economics and engineering—ARMA models, vector autoregressions (Bayesian or otherwise), linearized dynamic stochastic general equilibrium models, and linear state-space models. We derive upper bounds on the risk, which hold with high probability while requiring only weak assumptions on the data-generating process. These bounds are finite sample in nature, unlike standard model selection penalties such as AIC or BIC. Furthermore, they do not suffer the biases inherent in other risk estimation techniques such as the pseudo-cross validation approach often used in the economic forecasting literature.

While we have stated these results in terms of standard economic forecasting models, they have very wide applicability. Theorem 12 applies to any forecasting procedure with fixed memory length, linear or non-linear. Theorem 17 applies only to methods whose forecasts are linear in the observations, but a similar result for nonlinear methods would



just need to ensure that the dependence of the forecast on the past decays in some suitable way.

Rather than deriving bounds theoretically, one could attempt to estimate bounds on the risk. While cross-validation is tricky (Racine, 2000), nonparametric bootstrap procedures may do better. A fully nonparametric version is possible, using the circular bootstrap (reviewed in Lahiri, 1999). Bootstrapping lengthy out-of-sample sequences for testing fitted model predictions yields intuitively sensible estimates of  $R_n(f)$ , but there is currently no theory about the coverage level. Also, while models like VARs can be fit quickly to simulated data, general state-space models, let alone DSGEs, require large amounts of computational power, which is an obstacle to any resampling method.

While our results are a crucial first step for the learning-theoretic analysis of time-series forecasts, many avenues remain for future exploration. To gain a more complete picture of the performance of forecasting algorithms, we would want minimax lower bounds (c.f. Tsybakov, 2008). These would tell us the smallest risk we could hope to achieve using any forecaster in some larger model class, letting us ask whether any of the models in common use actually approach this minimum. Another possibility is to target not the *ex ante* risk of the forecast, but the *ex post* regret: how much better might our forecasts have been, in retrospect and on the actually-realized data, had we used a different prediction function from the model  $\mathcal{F}$  (Cesa-Bianchi and Lugosi, 2006; Rakhlin et al., 2010)? Remarkably, we can find forecasters which have low *ex post* regret, even if the data came from an adversary trying to make us perform badly. If we target regret rather than risk, we can actually ignore mixing, and even stationarity (Shalizi et al., 2011).

An increased recognition of the abilities and benefits of statistical learning theory can be of tremendous aid to financial and economic forecasters. The results presented here represent an initial yet productive foray in this direction. They allow for principled model comparisons as well as high probability performance guarantees. Future work in this direction will only serve to sharpen our ability to measure predictive power.

## Acknowledgments

This work is partially supported by grants from the Institute for New Economic Thinking. DJM is partially supported by NSF grant DMS1407439. CRS is partially supported by NIH Grant R01 NS047493 and by NSF grants DMS1207759 and DMS 1418124. The authors wish to thank David N. DeJong, Larry Wasserman, Alessandro Rinaldo and Darren Homrighausen for valuable suggestions on early drafts. Finally, we wish to thank two anonymous referees for their thorough and invaluable comments. Their suggestions have dramatically improved this paper.

## Appendix A. Auxiliary Results

**Lemma 21 (Yu, 1994 Lemma 4.1)** *Let  $Z$  be an event with respect to the block sequence  $\mathbf{U}$ . Then,*

$$|\mathbb{P}(Z) - \tilde{\mathbb{P}}(Z)| \leq \beta_a(\mu - 1),$$

where the first probability is with respect to the dependent block sequence,  $\mathbf{U}$ , and  $\tilde{\mathbb{P}}$  is with respect to the independent sequence,  $\mathbf{U}'$ .

This lemma essentially gives a way to apply i.i.d. results to  $\beta$ -mixing data. Because the dependence decays as we increase the separation between blocks, widely spaced blocks are nearly independent of each other. In particular, the difference between expectations over these nearly independent blocks and expectations over blocks which are actually independent can be controlled by the  $\beta$ -mixing coefficient. Very similar results are also given in Nobel and Dembo (1993) and Vidyasagar (2003).

**Lemma 22 (Cortes et al., 2010 Theorem 7)** *Under Assumption B,*

$$\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \hat{R}_n(f)}{Q_d(f)} > \epsilon \sqrt{2 + \log \frac{1}{\epsilon}} \right) \leq 4G(2n, \mathcal{C}_{\ell_0 \mathcal{F}}) \exp \left\{ -\frac{n\epsilon^2}{4} \right\}.$$

**Corollary 23** *Under Assumption B, for  $0 < \epsilon \leq e^{3/2}/\sqrt{2}$ ,*

$$\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \hat{R}_n(f)}{Q_d(f)} > \epsilon \right) \leq 4G(2n, \mathcal{C}_{\ell_0 \mathcal{F}}) \exp \left\{ -\frac{n \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\}.$$

**Proof** Letting  $\epsilon' = \epsilon \sqrt{2 + \log \frac{1}{\epsilon}}$  and solving for  $\epsilon$  gives

$$\epsilon = \exp \left\{ \frac{1}{2} \left( W \left( -\frac{2\epsilon'^2}{e^4} \right) + 4 \right) \right\}$$

if  $0 < \epsilon' \leq e^{3/2}/\sqrt{2}$ , so

$$\begin{aligned} \mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \hat{R}_n(f)}{Q_d(f)} > \epsilon \sqrt{2 + \log \frac{1}{\epsilon}} \right) &\leq 4G(2n, \mathcal{C}_{\ell_0 \mathcal{F}}) \exp \left\{ -\frac{n\epsilon^2}{4} \right\} \\ \Rightarrow \mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \hat{R}_n(f)}{Q_d(f)} > \epsilon \right) &\leq 4G(2n, \mathcal{C}_{\ell_0 \mathcal{F}}) \exp \left\{ -\frac{n \exp \left( W \left( -\frac{2\epsilon^2}{e^4} \right) + 4 \right)}{4} \right\}. \end{aligned}$$

■

**Lemma 24** *Suppose we have a function class  $\mathcal{F}$  with growth function  $G(n, \mathcal{F})$ . Let  $2\mu = n$ . Consider the class  $\mathcal{F}^\Sigma = \{\sum_{j=1}^a f(z_j) : f \in \mathcal{F}\}$ .*

$$G(\mu, \mathcal{F}^\Sigma) \leq G(n/2, \mathcal{F}).$$

**Proof** Let  $z_1, \dots, z_\mu \in \mathbb{U}^a$ ,  $b_1, \dots, b_\mu \in \mathbb{R}$  be a set such that  $\exists f_1^\Sigma, \dots, f_K^\Sigma \in \mathcal{F}^\Sigma$  with  $K = G(\mu, \mathcal{F}^\Sigma)$  which shatters  $\{z_i\}$ . That is, for  $v_{ik} \in \{0, 1\}$ ,  $i = 1, \dots, \mu$ ,  $k = 1, \dots, K$ , we have

$$I(f_k^\Sigma(z_i) > b_i) \text{ iff } v_{ik} = 1.$$

But this means,  $I(a^{-1} \sum_{j=1}^a f_k(z_{ij}) > b_i/a) \text{ iff } v_{ik} = 1$ . Since this is a convex combination of the  $f_k(z_{ij})$ , there is a subset of the  $z_{ij}$  such that

$$I(f_k(z_{ij}) > b_i/a) \text{ iff } v_{ik} = 1.$$

Thus  $G(\mu, \mathcal{F}^\Sigma) \leq G(\mu a, \mathcal{F}) = G(n/2, \mathcal{F})$ . ■

**Lemma 25** *Let  $\mathcal{F}$  be a class of predictor functions  $\mathbb{U} \mapsto \mathbb{R}$ , and let  $h : \mathbb{R} \mapsto \mathbb{R}$ , be a monotone-increasing function. Define  $h \circ \mathcal{F} = \{h \circ f : f \in \mathcal{F}\}$ . Then  $\forall n \geq 1$ ,  $G(n, h \circ \mathcal{F}) \leq G(n, \mathcal{F})$ .*

**Proof** Let  $z_1, \dots, z_n \in \mathbb{U}$ ,  $b_1, \dots, b_n \in \mathbb{R}$  be a set such that  $\exists f_1, \dots, f_K \in \mathcal{F}$  and  $K = G(\mu, h \circ \mathcal{F})$  which shatters  $\{z_i\}$ . That is, for  $v_{ik} \in \{0, 1\}$ ,  $i = 1, \dots, n$ ,  $k = 1, \dots, K$ , we have

$$I(h(f_k(z_i)) > b_i) \text{ iff } v_{ik} = 1.$$

Set

$$c_i = \max_{k \in 1:K} \{f_k(z_i) : h(f_k(z_i)) \leq b_i\}$$

Then, because  $h$  is monotone,

$$b_i \geq h(f_k(z_i)) \Leftrightarrow c_i \geq f_k(z_i)$$

Hence this set of points is shattered by the same collection of functions, and  $G(n, h \circ \mathcal{F}) \leq G(n, \mathcal{F})$ . ■

## Appendix B. Proofs of Selected Results

**Proof** [Lemma 7] We have that

$$\begin{aligned} b_{ij} &= Z \prod_{k=j+1}^{i-1} L_k K_j \\ &= Z(I - P_{j+1} Z^\top F_{j+1} Z) T(I - P_{j+2} Z^\top F_{j+2} Z) \cdots T(I - P_{i-1} Z^\top F_{i-1} Z) K_j. \end{aligned}$$

Now, by assumption  $\lambda_{\max}(T) < 1$ . Furthermore,  $\lambda_{\max}((I - P_k Z^\top F_k Z)) \leq 1$  for all  $k$ . To see this, write

$$\begin{aligned} I - P_k Z^\top F_k Z &= I - P_k Z^\top (Z P_k Z^\top + H)^{-1} Z \\ &= I - Z^\dagger [Z P_k Z^\top] [(Z P_k Z^\top + H)^{-1}] Z, \end{aligned}$$

where  $Z^\dagger$  is a generalized inverse for  $Z$ . As,  $P_k$  and  $H$  are positive definite,  $[ZP_kZ'][(ZP_kZ' + H)^{-1}]$  is positive definite with  $\lambda_{\max}([ZP_kZ'][(ZP_kZ' + H)^{-1}]) < 1$  and therefore  $I - Z^\dagger[ZP_kZ'][(ZP_kZ' + H)^{-1}]Z$  is positive semi-definite if  $\dim(H) > \dim(P_k)$  and positive definite if  $\dim(P_k) \leq \dim(H)$ . Therefore,  $0 < \lambda_{\max}((I - P_kZ'F_kZ)) \leq 1$ . Finally,  $\lambda_{\max}(b_{i,j}) \leq \|Z\|_2 \|T\|_2^{i-j-1} \prod_{k=j+1}^{i-1} \|I - P_kZ'F_kZ\|_2 \leq \lambda_{\max}(T)^{i-j-1} \|Z\|_2 = O(r^{i-j-1})$ . ■

**Proof** [Lemma 8]

1. This result is given in Anderson and Moore (2012, 4.4).
2. The convergence of  $\widehat{P}_i$  is in Anderson and Moore (2012, 4.4). As

$$K_i = T\widehat{P}_iZ^\top(Z\widehat{P}_iZ^\top + H)^{-1}$$

is continuous in  $\widehat{P}_i$ , it converges as well.

3. This result is given in Anderson and Moore (2012, 4.4).
4.  $\|\widehat{P}_i - \bar{P}\| = O(\lambda_{\max}(T - \bar{K}Z)^i)$  is given Anderson and Moore (2012, 4.4). For  $F_i$ ,

$$\|F_i - \bar{F}\| = \|F_i(F_i^{-1} - \bar{F}^{-1})\bar{F}\| = O(\|F_i^{-1} - \bar{F}^{-1}\|)$$

since  $\lambda_{\min}(F_i^{-1}) > \lambda_{\min}(H) > 0 \Rightarrow \lambda_{\max}(F_i) < 1/\lambda_{\min}(H) = O(1)$ . Proceeding,

$$O(\|F_i^{-1} - \bar{F}^{-1}\|) = O(\|Z\widehat{P}_iZ^\top - Z\bar{P}Z^\top\|) = O(\|\widehat{P}_i - \bar{P}\|) = O(\rho^i).$$

For  $K_i$ , note that

$$K_i - K = TP_iZ^\top F_i - T\bar{P}Z^\top \bar{F} = TP_iZ^\top F_i - T\bar{P}Z^\top F_i + T\bar{P}Z^\top F_i - T\bar{P}Z^\top \bar{F}.$$

We have  $\|TP_iZ^\top F_i - T\bar{P}Z^\top F_i\| = O(\rho^i)$  by the result for  $\widehat{P}_i$  and  $\|T\bar{P}Z^\top F_i - T\bar{P}Z^\top \bar{F}\| = O(\rho^i)$  by the result for  $F_i$ . ■

**Proof** [Lemma 9] By (3),  $s_{ij} = Z(T - \bar{K}Z)^{i-j-1}\bar{K}$  with  $\lambda_{\max}(T - \bar{K}Z) = \rho < 1$ . Thus  $\|s_{ij}\| = \|s_{i-j-1}\| = O(\rho^{i-j-1})$ . For  $j < i - 1$ ,

$$b_{ij} = Z \prod_{k=j+1}^{i-1} (T - K_kZ)K_j \qquad s_{ij} = Z(T - \bar{K}Z)^{i-j}\bar{K}.$$

Therefore,

$$\begin{aligned}
 \|b_{ij} - s_{ij}\| &\leq \|Z\| \left\| \prod_{k=j+1}^{i-1} (T - K_k Z) K_j - (T - \bar{K} Z)^{i-j-1} \bar{K} \right\| \\
 &\leq \|Z\| \left\| \prod_{k=j+1}^{i-1} (T - K_k Z) K_j - \prod_{k=j+1}^{i-1} (T - K_k Z) \bar{K} \right\| \\
 &\quad + \|Z\| \left\| \prod_{k=j+1}^{i-1} (T - K_k Z) \bar{K} - (T - \bar{K} Z)^{i-j-1} \bar{K} \right\| \\
 &\leq \|Z\| \left\| \prod_{k=j+1}^{i-1} (T - K_k Z) \right\| \|K_j - \bar{K}\| \\
 &\quad + \|Z\| \|\bar{K}\| \left\| \prod_{k=j+1}^{i-1} (T - K_k Z) - (T - \bar{K} Z)^{i-j-1} \right\| \\
 &= O(\lambda_{\max}(T - \bar{K} Z)^j) + O\left(\left\| \prod_{k=j+1}^{i-1} (T - K_k Z) - (T - \bar{K} Z)^{i-j-1} \right\|\right) \quad (9) \\
 &= O(\lambda_{\max}(T - \bar{K} Z)^j) + O\left(\|(T - K_{j+1} Z)^{i-j-1} - (T - \bar{K} Z)^{i-j-1}\|\right) \quad (10) \\
 &= O(\lambda_{\max}(T - \bar{K} Z)^j).
 \end{aligned}$$

Here, (9) follows from Lemma 8 and (10) follows because  $\|(T - K_k Z) - (T - \bar{K} Z)\|$  is increasing in  $k$ , so

$$\begin{aligned}
 \left\| \prod_{k=j+1}^{i-1} (T - K_k Z) - (T - \bar{K} Z) \right\| &\leq \left\| \prod_{k=j+1}^{i-1} (T - K_{j+1} Z) - (T - \bar{K} Z) \right\| \\
 &= \|(T - K_{j+1} Z)^{i-j-1} - (T - \bar{K} Z)\|.
 \end{aligned}$$

For  $j = i - 1$ ,

$$\|b_{i,i-1} - s_{i,i-1}\| = \|ZK_j - Z\bar{K}\| = O(\lambda_{\max}(T - \bar{K} Z)^j).$$

■

**Proof** [Theorem 14] Consider first an AR( $d$ ) model and the class

$$\mathcal{C}_{|\mathcal{F}|} = \{(u_0, u_1, u) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^d : \mathbf{1}(|u_0 - a^\top u| > u_1) = 1\}.$$

This is the class contains sets of points  $u$  for which linear predictions of  $u_0$  are made via  $a$  with threshold  $b$ . Rewrite  $\mathcal{C}_{|\mathcal{F}|}$  as

$$\mathcal{C}_{|\mathcal{F}|} = \left\{ (u_0, u_1, u) : \max\{\text{sgn}((1 - a - u_1)^\top (u_0 \ u \ 1)), \text{sgn}((-1 \ a - u_1)^\top (u_0 \ u \ 1))\} = 1 \right\}.$$

Then  $\mathcal{C}_{|\mathcal{F}|}$  is a 2-combination of  $\mathcal{F}$  as in Anthony and Bartlett (1999, Theorem 7.3). As  $\mathcal{F}$  is a linear function, it has solution set components bound  $B = 1$  (Anthony and Bartlett, 1999, p. 91). We add an intercept to  $\mathcal{F}$  so that it will be closed under addition. Then by (Anthony and Bartlett, 1999, Theorem 7.6),

$$G(n, \mathcal{C}_{|\mathcal{F}|}) \leq \left( \frac{2en}{d+1} \right)^{d+1}.$$

Then apply Lemma 25. The result for VARs is similar. Consider the class

$$\begin{aligned} \mathcal{C}_{\mathcal{G}} = & \left\{ [\text{sgn}((1 \ -a_1 \ 0 \cdots 0 \ -u_1/k)^\top (u_0 \ u \ 1))] \vee [\text{sgn}((-1 \ a_1 \ 0 \cdots 0 \ -u_1/k)^\top (u_0 \ u \ 1))] \right. \\ & \vee [\text{sgn}((1 \ 0 \ -a_2 \ 0 \cdots 0 \ -u_1/k)^\top (u_0 \ u \ 1))] \vee [\text{sgn}((-1 \ 0 \ a_2 \ 0 \cdots 0 \ -u_1/k)^\top (u_0 \ u \ 1))] \\ & \vee \cdots \vee [\text{sgn}((1 \ 0 \cdots 0 \ -a_k \ -u_1/k)^\top (u_0 \ u \ 1))] \\ & \left. \vee [\text{sgn}((-1 \ 0 \cdots 0 \ a_k \ -u_1/k)^\top (u_0 \ u \ 1))] \right\}. \end{aligned}$$

This is the class which correctly classifies at least one of the  $k$  coordinates. By (Anthony and Bartlett, 1999, Theorem 7.6),

$$G(n, \mathcal{C}_{\mathcal{G}}) \leq \left( \frac{2ken}{k(d+1)} \right)^{k(d+1)} = \left( \frac{2en}{d+1} \right)^{k(d+1)}.$$

Then apply Lemma 25. As  $|\mathcal{F}| \subseteq \mathcal{G}$ , we have the result.  $\blacksquare$

**Proof** [Theorem 15 and Corollary 16] Let  $\mathcal{F}_2$  be the class of predictors given by Algorithm 2. Let  $\mathcal{F}'_2$  be the same class, but predictions are made based on the truncated memory length  $d$ . That is for  $f_2 \in \mathcal{F}_2$ ,

$$\widehat{Y}_{d+1:n+1} = \mathbf{S}Y_{1:n}$$

where

$$\mathbf{S} = \begin{bmatrix} s_{d+1,1} & \cdots & s_{d+1,d} & & & \mathbf{0} \\ s_{d+2,1} & \cdots & s_{d+2,d} & s_{d+2,d+1} & & \\ \vdots & & \vdots & & \ddots & \\ s_{n+1,1} & \cdots & s_{n+1,d} & s_{n+1,d+1} & \cdots & s_{n+1,n} \end{bmatrix} = \begin{bmatrix} s_{d-1} & \cdots & s_0 & & & \mathbf{0} \\ s_d & \cdots & s_1 & s_0 & & \\ \vdots & & \vdots & & \ddots & \\ s_{n-1} & \cdots & s_{n-d} & s_{n-d+1} & \cdots & s_0 \end{bmatrix}$$

as in (3) (redefining  $s_{i,j} \rightarrow s_{i-j-1}$ ), but for  $f'_2 \in \mathcal{F}'_2$ ,

$$\widehat{Y}'_{d+1:n+1} = \mathbf{S}'Y_{1:n}$$

where

$$\mathbf{S}' = \begin{bmatrix} s_{d-1} & s_{d-2} & \cdots & s_0 & 0 & \mathbf{0} \\ 0 & s_{d-1} & \cdots & s_1 & s_0 & \\ & & & & \ddots & \\ & & \mathbf{0} & & & \\ & & & & s_{d-1} & \cdots & s_0 \end{bmatrix}. \quad (11)$$

Note that by using Algorithm 2, predictions  $\widehat{Y}_i'$  are made using the same linear combination of previous observations at all times  $i$ . Thus, truncating the memory, makes these predictions simple autoregressive models of order  $d$ .

Define  $\widetilde{R}_n(f'_2)$  to be the training error of this truncated predictor  $f'_2$ . Then, for any  $f_2 \in \mathcal{F}_2$ ,

$$\begin{aligned} R_n(f_2) - \widehat{R}_n(f_2) &= (R_n(f_2) - R_n(f'_2)) + (R_n(f'_2) - \widetilde{R}_n(f'_2)) + (\widetilde{R}_n(f'_2) - \widehat{R}_n(f_2)) \\ &= \Delta_1(d) + (R_n(f'_2) - \widetilde{R}_n(f'_2)) + \Delta_2(d) \\ &= (R_n(f'_2) - \widetilde{R}_n(f'_2)) + \delta_d(f_2), \end{aligned}$$

where we have defined  $\Delta_1(d)$ ,  $\Delta_2(d)$ , and  $\delta_d(f_2)$  in the obvious way. For now, we will simply proceed with the analysis incorporating the approximation term  $\delta_d(f_2)$ , before showing that  $\delta_d(f_2)$  decays rapidly for this class of models.

Now,

$$R_n(f_2) - \widehat{R}_n(f_2) - \delta_d(f_2) = R_n(f'_2) - \widetilde{R}_n(f'_2) = R_d(f'_2) - \widetilde{R}_n(f'_2).$$

Dividing through by  $Q_d(f'_2) = Q_d(f_2)$  and taking the supremum over  $\mathcal{F}_2$  (and therefore over  $\mathcal{F}'_2$ ) gives

$$\sup_{f_2 \in \mathcal{F}_2} \frac{R_n(f_2) - \widehat{R}_n(f_2) - \delta_d(f_2)}{Q_d(f_2)} \leq \sup_{f'_2 \in \mathcal{F}'_2} \frac{R_n(f'_2) - \widetilde{R}_n(f'_2)}{Q_d(f'_2)}$$

Finally,

$$\mathbb{P} \left( \sup_{f \in \mathcal{F}} \frac{R_n(f) - \widehat{R}_n(f) - \delta_d(f)}{Q_d(f)} > \epsilon \right) \leq \mathbb{P} \left( \sup_{f' \in \mathcal{F}'} \frac{R_n(f') - \widetilde{R}_n(f')}{Q_d(f')} > \epsilon \right)$$

Since  $\mathcal{F}'$  is a class with finite memory, we can apply Theorem 12 and Corollary 13 to get the results.

To show that the finite approximation  $\delta_d(f)$  decays rapidly for Algorithm 2, consider both components separately. For the case of the difference in expected risks, we need only consider the last rows of  $\mathbf{S}$  and  $\mathbf{S}'$ . As

$$\Delta_1(d) = R_n(f) - R_n(f') = \mathbb{E}[\ell(Y_{n+1} - \mathbf{s}_{n+1}Y_{1:n})] - \mathbb{E}[\ell(Y_{n+1} - \mathbf{s}'_{n+1}Y_{1:n})],$$

where  $\mathbf{s}_{n+1}$  indicates the  $(n+1)^{st}$  row of  $\mathbf{S}$  and similarly for  $\mathbf{s}'_{n+1}$ ,

$$\begin{aligned}
 \Delta_1(d) &\leq \mathbb{E}[\ell((\mathbf{s}_{n+1} - \mathbf{s}'_{n+1})Y_{1:n})] && \text{(by the Triangle inequality for } \ell) \\
 &= \mathbb{E}\left[\sum_{j=1}^{n-d} \ell(s_{n-j}Y_j)\right] \\
 &\leq \sum_{j=1}^{n-d} \ell^*(s_{n-j})\mathbb{E}[\ell(Y_j)] && \left(\text{as } \ell \text{ is a norm, with } \ell^*(s) = \sup_{x \neq 0} \frac{\ell(sx)}{\ell(x)}\right) \\
 &= \mathbb{E}[\ell(Y_1)] \left(\sum_{j=1}^{n-d} \ell^*(s_{n-j})\right) && \text{(by stationarity)} \\
 &= \mathbb{E}[\ell(Y_1)] \left(\sum_{j=1}^{n-d} O(\rho^{n-j})\right) && \text{(Lemma 9 and equivalence of norms)} \\
 &= O\left(\frac{\rho^d - \rho^n}{1 - \rho}\right) = O(\rho^d).
 \end{aligned}$$

Similarly,

$$\begin{aligned}
 \Delta_2(d) &= \tilde{R}_n(f') - \hat{R}_n(f) \\
 &\leq \frac{1}{n-d} \sum_{i=d+1}^n \ell((\mathbf{s}_i - \mathbf{s}'_i)Y_{1:i-1}) \\
 &= \frac{1}{n-d} \sum_{i=d+2}^n \sum_{j=1}^{i-d-1} \ell(s_{i-j-1}Y_j) \\
 &\leq \frac{1}{n-d} \sum_{i=d+2}^n O_p(1) \sum_{j=1}^{i-d-1} \ell^*(s_{i-j-1}) \\
 &\leq O_{\mathbb{P}}((n-d)^{-1}) \sum_{i=d+2}^n \sum_{j=1}^{i-d-1} O(\rho^{i-j-1}) \\
 &= O_{\mathbb{P}}((n-d)^{-1}) O\left[\frac{(n-d-1)\rho^d - (n-d)\rho^{d+1} + \rho^n}{(1-\rho)^2}\right] \\
 &= O_{\mathbb{P}}(\rho^d)
 \end{aligned}$$

■

**Proof** [Theorem 17 and Corollary 18] Let  $\mathcal{F}_1$  be the class of predictors given by Algorithm 1. Let  $\mathcal{F}'_2$  be the class of predictors given by Algorithm 2 based on the truncated memory length  $d$ . That is for  $f_1 \in \mathcal{F}_1$ ,

$$\hat{Y}_{d+1:n+1} = \mathbf{B}Y_{1:n}$$



where

$$\mathbf{B} = \begin{bmatrix} b_{d,1} & \cdots & b_{d,d} & & & & 0 \\ b_{d+1,1} & \cdots & b_{d+1,d} & b_{d+1,d+1} & & & \\ \vdots & & \vdots & & \ddots & & \\ b_{n+1,1} & \cdots & b_{n+1,d} & b_{n+1,d+1} & \cdots & b_{n+1,n} \end{bmatrix}$$

as in (2), but for  $f'_2$

$$\widehat{Y}'_{d+1:n+1} = \mathbf{S}'Y_{1:n}$$

as in (11).

Then, proceeding as in the previous proof, define  $\widetilde{R}_n(f'_2)$  to be the training error of the associated truncated limiting predictor  $f'_2$  and write analogously

$$\begin{aligned} R_n(f_1) - \widehat{R}_n(f_1) &= (R_n(f_1) - R_n(f'_2)) + (R_n(f'_2) - \widetilde{R}_n(f'_2)) + (\widetilde{R}_n(f'_2) - \widehat{R}_n(f_1)) \\ &= \Delta_1(d) + (R_n(f'_2) - \widetilde{R}_n(f'_2)) + \Delta_2(d) \\ &= (R_n(f'_2) - \widetilde{R}_n(f'_2)) + \delta_d(f_1), \end{aligned}$$

where

$$\delta_d(f_1) = (R_n(f_1) - R_n(f'_2)) + (\widetilde{R}_n(f'_2) - \widehat{R}_n(f_1)) = \Delta_1(d) + \Delta_2(d).$$

Now, by Lemma 9, the truncated version of  $f_1, f'_1$  converges to  $f'_2$  so that  $f'_1(Y_{n-d:n-1})$  converges in distribution to  $f'_2(Y_{n-d:n-1})$ . And because  $\ell(y-x)^2$  is continuous in  $x$ ,  $\ell(Y_n - f'_1(Y_{n-d:n-1}))^2$  converges in distribution to  $\ell(Y_n - f'_2(Y_{n-d:n-1}))^2$  by the continuous mapping theorem. By Assumption C,  $\ell(Y_n - f'_1(Y_{n-d:n-1}))^2$  is uniformly integrable, so

$$\lim_{n \rightarrow \infty} \mathbb{E} [\ell(Y_n - f'_1(Y_{n-d:n-1}))^2] = \mathbb{E} [\ell(Y_n - f'_2(Y_{n-d:n-1}))^2].$$

Therefore,

$$\frac{R_n(f_1) - \widehat{R}_n(f_1) - \delta_d(f_1)}{\lim_{n \rightarrow \infty} Q_d(f_1)} = \frac{R_n(f'_2) - \widetilde{R}_n(f'_2)}{Q_d(f'_2)}$$

For each  $f_1 \in \mathcal{F}_1$  (parameterized by  $\theta$ ), there exists one function  $f_2 \in \mathcal{F}_2$  which results from the same parameter vector, so that

$$\sup_{f_1 \in \mathcal{F}_1} \frac{R_n(f_1) - \widehat{R}_n(f_1) - \delta_d(f_1)}{\lim_{n \rightarrow \infty} Q_d(f_1)} \leq \sup_{f'_2 \in \mathcal{F}'_2} \frac{R_n(f'_2) - \widetilde{R}_n(f'_2)}{Q_d(f'_2)}.$$

Therefore,

$$\mathbb{P} \left( \sup_{f_1 \in \mathcal{F}_1} \frac{R_n(f_1) - \widehat{R}_n(f_1) - \delta_d(f_1)}{\lim_{n \rightarrow \infty} Q_d(f_1)} > \epsilon \right) \leq \mathbb{P} \left( \sup_{f'_2 \in \mathcal{F}'_2} \frac{R_n(f'_2) - \widetilde{R}_n(f'_2)}{Q_d(f'_2)} > \epsilon \right)$$

Since  $\mathcal{F}'_2$  is a class with finite memory, we can apply Theorem 12 and Corollary 13 to get the results.

To show that the finite approximation  $\delta_d(f_1)$  decays rapidly, consider both components separately. For the case of the difference in expected risks, we need only consider the last rows of  $\mathbf{B}$  and  $\mathbf{S}'$ . As

$$\Delta_1(d) = R_n(f_1) - R_n(f'_2) = \mathbb{E}[\ell(Y_{n+1} - \mathbf{b}_{n+1}Y_{1:n})] - \mathbb{E}[\ell(Y_{n+1} - \mathbf{s}'_{n+1}Y_{1:n})],$$

where  $\mathbf{b}_{n+1}$  indicates the  $(n+1)^{st}$  row of  $\mathbf{B}$  and similarly for  $\mathbf{s}'_{n+1}$ . Then, as in the previous proof,

$$\Delta_1(d) \leq \mathbb{E}[\ell((\mathbf{b}_{n+1} - \mathbf{s}'_{n+1})Y_{1:n})] \tag{12}$$

$$\begin{aligned} &\leq \mathbb{E} \left[ \left( \sum_{j=1}^{n-d} \ell(b_{n+1,j}Y_j) + \sum_{j=n-d+1}^n \ell((b_{n+1,j} - s_{n-j})Y_j) \right) \right] \\ &\leq \sum_{j=1}^{n-d} \ell^*(b_{n+1,j})\mathbb{E}[\ell(Y_j)] + \sum_{j=n-d+1}^n \ell^*(b_{n+1,j} - s_{n-j})\mathbb{E}[\ell(Y_j)] \end{aligned} \tag{13}$$

$$= \mathbb{E}[\ell(Y_1)] \left( \sum_{j=1}^{n-d} \ell^*(b_{n+1,j}) + \sum_{j=n-d+1}^n \ell^*(b_{n+1,j} - s_{n-j}) \right)$$

$$= \mathbb{E}[\ell(Y_1)] \left( \sum_{j=1}^{n-d} O(r^{n-j}) + \sum_{j=n-d+1}^n O(\rho^j) \right)$$

$$= O\left(\frac{r^d - r^n}{1-r}\right) + O\left(\frac{\rho^{n-d+1}(1-\rho^d)}{1-\rho}\right)$$

$$= O(r^d) + O(\rho^{n-d+1}).$$

(12) and B are via the triangle inequality for  $\ell$ , while (13) uses the fact that  $\ell$  is a norm, with  $\ell^*(b) = \sup_{x \neq 0} \frac{\ell(bx)}{\ell(x)}$ . Similarly,

$$\begin{aligned}
 \Delta_2(d) &= \tilde{R}_n(f'_2) - \hat{R}_n(f_1) \\
 &\leq \frac{1}{n-d} \sum_{i=d+1}^n \ell((\mathbf{b}_i - \mathbf{s}'_i)Y_{1:n-1}) \\
 &\leq \frac{1}{n-d} \sum_{i=d+1}^n \sum_{j=1}^{i-d-1} \ell(b_{ij}Y_j) + \frac{1}{n-d} \sum_{i=d+1}^n \sum_{j=i-d}^{i-1} \ell((b_{ij} - s_{n-j})Y_j) \\
 &\leq \frac{1}{n-d} \sum_{i=d+2}^n O_{\mathbb{P}}(1) \sum_{j=1}^{i-d-1} \ell^*(b_{ij}) + \frac{1}{n-d} \sum_{i=d+1}^n O_{\mathbb{P}}(1) \sum_{j=i-d}^{i-1} \ell^*(b_{ij} - s_{ij}) \\
 &\leq O_{\mathbb{P}}((n-d)^{-1}) \sum_{i=d+2}^n \sum_{j=1}^{i-d-1} r^{i-j-1} + O_{\mathbb{P}}((n-d)^{-1}) \sum_{i=d+1}^n \sum_{j=i-d}^{i-1} \rho^j \\
 &= O_{\mathbb{P}}((n-d)^{-1}) O \left[ \frac{(n-d-1)r^d - (n-d)r^{d+1} + r^n}{(1-r)^2} \right] \\
 &\quad + O_{\mathbb{P}}((n-d)^{-1}) \left[ \frac{(1-r^d)(r^d - r^n)}{r^{d-1}(1-r)^2} \right] \\
 &= O_{\mathbb{P}}(r^d) + O_{\mathbb{P}}((n-d)^{-1})
 \end{aligned}$$

■

**Proof** [Theorem 20] Theorem 19 implies that simultaneously

$$\mathbb{P} \left( R_n(\hat{f}_{erm}) - \hat{R}_n(\hat{f}_{erm}) > \epsilon \right) \leq 8G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) \exp \left\{ -\frac{\mu\epsilon^2}{\Upsilon} \right\} + 2\mu\beta_{a-d}$$

and

$$\mathbb{P} \left( R_n(f^*) - \hat{R}_n(f^*) > \epsilon \right) \leq 8G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) \exp \left\{ -\frac{\mu\epsilon^2}{\Upsilon} \right\} + 2\mu\beta_{a-d}.$$

Since  $\hat{R}_n(\hat{f}_{erm}) - \hat{R}_n(f^*) \leq 0$ , then

$$\mathbb{P} \left( R_n(\hat{f}_{erm}) - R_n(f^*) > \epsilon \right) \leq 16G(n, \mathcal{C}_{\ell \circ \mathcal{F}}) \exp \left\{ -\frac{\mu\epsilon^2}{\Upsilon} \right\} + 4\mu\beta_{a-d}.$$

For any nonnegative random variable  $Z$ ,  $\mathbb{E}[Z] = \int_0^\infty \mathbb{P}(Z > \epsilon) d\epsilon$ , so

$$\begin{aligned}
 L_n^2(\Pi) &= \int_0^\infty d\epsilon \mathbb{P} \left( (R_n(\widehat{f}_{erm}) - R_n(f^*))^2 > \epsilon \right) \\
 &= \int_0^\xi d\epsilon \mathbb{P} \left( (R_n(\widehat{f}_{erm}) - R_n(f^*))^2 > \epsilon \right) + \int_\xi^\infty d\epsilon \mathbb{P} \left( R_n(\widehat{f}_{erm}) - R_n(f^*) > \sqrt{\epsilon} \right) \\
 &= \int_0^\xi d\epsilon \mathbb{P} \left( (R_n(\widehat{f}_{erm}) - R_n(f^*))^2 > \epsilon \right) + \int_\xi^\infty d\epsilon \mathbb{P} \left( (R_n(\widehat{f}_{erm}) - R_n(f^*))^2 > \sqrt{\epsilon} \right) \\
 &= \int_0^\xi d\epsilon \mathbb{P} \left( (R_n(\widehat{f}_{erm}) - R_n(f^*))^2 > \epsilon \right) + \int_\xi^{M^2} d\epsilon \mathbb{P} \left( R_n(\widehat{f}_{erm}) - R_n(f^*) > \sqrt{\epsilon} \right) \tag{14} \\
 &\leq \xi + \int_\xi^{M^2} d\epsilon \left[ 16G(n, \mathcal{C}_{\ell_o\mathcal{F}}) \exp \left\{ -\frac{\mu\epsilon}{\Upsilon} \right\} + 4\mu\beta_{a-d} \right] \\
 &\leq \xi + \int_\xi^\infty d\epsilon 16G(n, \mathcal{C}_{\ell_o\mathcal{F}}) \exp \left\{ -\frac{\mu\epsilon}{\Upsilon} \right\} + \int_\xi^{M^2} d\epsilon 4\mu\beta_{a-d} \\
 &= \xi + \frac{16\Upsilon G(n, \mathcal{C}_{\ell_o\mathcal{F}}) \exp \left\{ -\frac{\mu\xi}{\Upsilon} \right\}}{\mu} + 4(M^2 - \xi)\mu\beta_{a-d}
 \end{aligned}$$

for all  $0 < \xi < \Upsilon$ . Here, (14) follows because  $\mathbb{P} \left( R_n(\widehat{f}_{erm}) - R_n(f^*) > M \right) = 0$ . Using Assumption D, take

$$a = (n-d)^{1/(1+\kappa)} = O(n^{1/(1+\kappa)}), \quad \mu = .5(n-d)^{\kappa/(1+\kappa)} = O(n^{\kappa/(1+\kappa)}),$$

and  $\xi = \frac{\log G(n, \mathcal{C}_{\ell_o\mathcal{F}})}{n^{\kappa/(1+\kappa)}}$  to balance the exponential and linear terms. Then

$$L_n(\Pi) = O \left( \sqrt{\frac{\log G(n, \mathcal{C}_{\ell_o\mathcal{F}})}{n^{\kappa/(1+\kappa)}}} \right) = O \left( \sqrt{\frac{\text{VCD}(\mathcal{C}_{\ell_o\mathcal{F}}) \log n}{n^{\kappa/(1+\kappa)}}} \right).$$

as  $\mathcal{C}_{\ell_o\mathcal{F}}$  has finite VC-dimension.

For the lower bound, apply the i.i.d. version, as classification is a special case of bounded regression. The result follows.  $\blacksquare$

## References

- Terrence M. Adams and Andrew B. Nobel. Uniform convergence of Vapnik-Chervonenkis classes under ergodic sampling. *Annals of Probability*, 38:1345–1367, 2010.
- Alekh Agarwal and John C. Duchi. The generalization ability of online algorithms for dependent data. *IEEE Transactions on Information Theory*, 59(1):573–587, 2013.
- Pierre Alquier, Olivier Wintenberger, et al. Model selection for weakly dependent time series forecasting. *Bernoulli*, 18(3):883–913, 2012.

- Pierre Alquier, Xiaoyin Li, and Olivier Wintenberger. Prediction of time series by statistical learning: general losses and fast rates. *Dependence Modelling*, 1:65–93, 2014.
- Brian D.O. Anderson and John B. Moore. *Optimal filtering*. Prentice-Hall, Englewood Cliffs, NJ, 2012.
- Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge, England, 1999.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, 2:499–526, 2002.
- Richard C. Bradley. Basic properties of strong mixing conditions: A survey and some open questions. *Probability Surveys*, 2:107–144, 2005.
- Marine Carrasco and Xiaohong Chen. Mixing and moment properties of various GARCH and stochastic volatility models. *Econometric Theory*, 18:17–39, 2002.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, England, 2006.
- Robert M. Corless, Gaston H. Gonnet, D. E. G. Hare, David J. Jeffrey, and Donald E. Knuth. On the Lambert  $W$  function. *Advances in Computational Mathematics*, 5:329–359, 1996.
- Corinna Cortes, Yishay Mansour, and Mehryar Mohri. Learning bounds for importance weights. In John Lafferty, C. K. I. Williams, John Shawe-Taylor, Richard S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing (NIPS)*, pages 442–450. MIT Press, 2010.
- David N. DeJong and Chetan Dave. *Structural Macroeconometrics*. Princeton University Press, Princeton, New Jersey, 2007.
- David N. DeJong, Beth F. Ingram, and Charles H. Whiteman. A Bayesian approach to dynamic macroeconomics. *Journal of Econometrics*, 98:203–223, 2000.
- Paul Doukhan. *Mixing: Properties and Examples*. Springer-Verlag, New York, 1995.
- James Durbin and Siem Jam Koopman. *Time Series Analysis by State Space Methods*. Oxford University Press, Oxford, 2001.
- Robert F. Engle. Autoregressive conditional heteroscedasticity with estimates of variance of United Kingdom inflation. *Econometrica*, 50:987–1008, 1982.
- Robert F. Engle. Garch 101: The use of arch/garch models in applied econometrics. *Journal of Economic Perspectives*, 15:157–168, 2001.
- Jesús Fernández-Villaverde. The econometrics of DSGE models. *SERIEs*, 1:3–49, 2010.

- Andrew Harvey, Esther Ruiz, and Neil Shephard. Multivariate stochastic variance models. *The Review of Economic Studies*, 61:247–264, 1994.
- Rudolf E. Kalman. A new approach to linear filtering and prediction problems. *ASME Transactions, Journal of Basic Engineering*, 82D:35–50, 1960.
- Rajeeva L. Karandikar and Mathukumalli Vidyasagar. Rates of uniform convergence of empirical means with mixing processes. *Statistics and Probability Letters*, 58:297–307, 2002.
- Michael J. Kearns and Dana Ron. Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. *Neural Computation*, 11:1427–1453, 1999.
- Bas J. K. Kleijn and Aad W. van der Vaart. Misspecification in infinite-dimensional Bayesian statistics. *Annals of Statistics*, 34:837–877, 2006.
- Vladimir Koltchinskii and Dmitriy Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, 30(1):1–50, 02 2002.
- Vitaly Kuznetsov and Mehryar Mohri. Generalization bounds for time series prediction with non-stationary processes. In Peter Auer, Alexander Clark, Thomas Zeugmann, and Sandra Zilles, editors, *Algorithmic Learning Theory (ALT)*, pages 260–274, 2014.
- Finn E. Kydland and Edward C. Prescott. Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370, November 1982.
- Soumendra N. Lahiri. Theoretical comparisons of block bootstrap methods. *Annals of Statistics*, 27:386–404, 1999.
- Pascal Massart. *Concentration Inequalities and Model Selection*. Springer-Verlag, Berlin, 2007.
- Daniel J. McDonald, Cosma Rohilla Shalizi, and Mark Schervish. Estimating  $\beta$ -mixing coefficients. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 516–524, 2011a.
- Daniel J. McDonald, Cosma Rohilla Shalizi, and Mark Schervish. Risk bounds for autoregressive models. Technical report, Statistics Department, CMU, 2011b.
- Daniel J. McDonald, Cosma Rohilla Shalizi, and Mark Schervish. Estimating beta-mixing coefficients via histograms. *Electronic Journal of Statistics*, 9:2855–2883, 2015.
- Ron Meir. Nonparametric time series prediction through adaptive model selection. *Machine Learning*, 39:5–34, 2000.
- Dharmendra S. Modha and Elias Masry. Memory-universal prediction of stationary random processes. *IEEE Transactions on Information Theory*, 44(1):117–133, 1998.

- Mehryar Mohri and Afshin Rostamizadeh. Rademacher complexity bounds for non-I.I.D. processes. In Daphne Koller, D. Schuurmans, Y. Bengio, and Léon Bottou, editors, *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1104, 2009.
- Mehryar Mohri and Afshin Rostamizadeh. Stability bounds for stationary  $\phi$ -mixing and  $\beta$ -mixing processes. *Journal of Machine Learning Research*, 11, 2010.
- Abdelkader Mokkadem. Mixing properties of ARMA processes. *Stochastic Processes and their Applications*, 29(2):309–315, 1988.
- Ulrich K. Müller. Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix. *Econometrica*, 81:1805–1849, 2013.
- Andrew Nobel and Amir Dembo. A note on uniform laws of averages for dependent processes. *Statistics and Probability Letters*, 17:169–172, 1993.
- David Pollard. *Convergence of Stochastic Processes*. Springer-Verlag, Berlin, 1984.
- David Pollard. *Empirical Processes: Theory and Applications*, volume 2 of *NSF-CBMS Regional Conference Series in Probability and Statistics*. Institute of Mathematical Statistics, Hayward, California, 1990.
- Jeff Racine. Consistent cross-validated model-selection for dependent data: hv-block cross-validation. *Journal of Econometrics*, 99:39–61, 2000.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Random averages, combinatorial parameters, and learnability. In John Lafferty, C. K. I. Williams, John Shawe-Taylor, Richard S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing (NIPS)*, pages 1984–1992. MIT Press, 2010.
- David Romer. *Advanced Macroeconomics*. McGraw-Hill, New York, 4 edition, 2011.
- Cosma Rohilla Shalizi. Dynamics of Bayesian updating with dependent data and misspecified models. *Electronic Journal of Statistics*, 3:1039–1074, 2009.
- Cosma Rohilla Shalizi, Abigail Z. Jacobs, Kristina Lisa Klinkner, and Aaron Clauset. Adapting to non-stationarity with growing expert ensembles. Technical report, Statistics Department, CMU, 2011.
- Christopher A. Sims. Solving linear rational expectations models. *Computational Economics*, 20:1–20, 2002.
- Frank Smets and Rafael Wouters. Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606, June 2007.
- Ingo Steinwart and Andreas Christmann. Fast learning from non-i.i.d. observations. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22, pages 1768–1776. MIT Press, Cambridge, MA, 2009.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Verlag, New York, 2008.

Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley, New York, 1998.

Vladimir N. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, Berlin, 2nd edition, 2000.

Vladimir N. Vapnik and Alexey Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16: 264–280, 1971.

Mathukumalli Vidyasagar. *Learning and Generalization: With Applications to Neural Networks*. Springer-Verlag, Berlin, second edition, 2003.

Bin Yu. Rates of convergence for empirical processes of stationary mixing sequences. *Annals of Probability*, 22:94–116, 1994.