International Journal of Forecasting 32 (2016) 1340–1351

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Constrained functional time series: Applications to the Italian gas market

Antonio Canale[a,*], Simone Vantini[b]

[a] Department of Economics and Statistics, University of Turin and Collegio Carlo Alberto, Italy

[b] MOX, Department of Mathematics, Politecnico di Milano, Italy

ARTICLE INFO

Keywords:
Autoregressive model
Demand and offer model
Energy forecasting
Functional data analysis
Functional ridge regression

ABSTRACT

Motivated by market dynamic modelling in the Italian Natural Gas Balancing Platform, we propose a model for analyzing time series of functions, subject to equality and inequality constraints at the two edges of the domain, respectively, such as daily demand and offer curves. Specifically, we provide the constrained functions with suitable pre-Hilbert structures, and introduce a useful isometric bijective map that associates each possible bounded and monotonic function to an unconstrained one. We introduce a functional-to-functional autoregressive model that is used to forecast future demand/offer functions, and estimate the model via the minimization of a penalized mean squared error of prediction, with a penalty term based on the Hilbert–Schmidt squared norm of autoregressive lagged operators. The approach is of general interest and could be generalized to any situation in which one has to deal with functions that are subject to the above constraints which evolve over time.

1. Introduction

Energy markets in general, and natural gas markets in particular, are emerging fields that pose a great variety of forecasting problems, including load forecasting (Hong, 2014), price forecasting (Weron, 2014), daily price curve profile forecasting (Chen & Li, 2015), consumption forecasting (Brabec, Konár, Pelikán, & Malı, 2008), and many others. Motivated by price prediction in the Italian natural gas balancing market, this paper proposes a model for forecasting the day-to-day evolution of supply and demand curves. The proposed model is innovative from both the methodological and applied perspectives.

The supply and demand curves model is indeed a well-known microeconomic model of price determination, but its application is typically descriptive and static rather than strategic and predictive, which clearly does not help gas traders with either the forecasting of future prices or decision making and bidding. At the same time, while the usual forecasting methods, such as classical time series analysis, produce useful predictions of scalar quantities of interest (e.g., prices), they do not provide the insights into the market that are given by the supply and demand model. Furthermore, in markets with a moderate number of traders, the effect of a single offer or demand cannot be incorporated directly into either the inferential procedure or what-if simulations. For all of these reasons, the prediction of the entire supply and demand curves, and hence of their intersection, can be of strong interest.

We deal with this problem using a functional data analysis (FDA) approach. FDA is an extremely useful set of tools for dealing with data that can be modeled as functions, such as our demand and supply curves; for a quick introduction, refer to Ferraty and Vieu (2006), Ramsay and

* Corresponding author.
E-mail address: antonio.canale@unito.it (A. Canale).

http://dx.doi.org/10.1016/j.ijforecast.2016.05.002
0169-2070/

Silverman (2002, 2005), or Sørensen, Goldsmith, and Sangalli (2013). However, our approach differs from the most common FDA framework in two ways. First, we focus on functions that are constrained (i.e., monotonic and with an equality constraint on one edge of the domain and an inequality constraint on the other edge), and second, we embed such constraints for curves that are temporally dependent. The statistical literature has focused separately on (a) the problem of obtaining a constrained estimation of the underlying function given some point-wise evaluations of it, and (b) the problem of modeling functional data with temporal dependence (i.e., functional time series). To the best of our knowledge, the present work is the first to tackle the temporal dependence jointly with constraints pertaining to monotonicity, boundedness, and values of the function at the boundary of the domain. We refer to this joint framework henceforth as *constrained functional time series*.

Before going into detail about the mathematical modeling and the estimation method that we propose herein, we provide a brief overview of the state of the art pertaining to both monotonic estimation and functional time series estimation.

The problem of having monotonic estimates of unknown functions that are observed only at a few sparse points in the domain, possibly with some measurement error, has been being tackled in the literature for many decades, even before the recent outbreak of FDA. Isotonic regression was the first approach presented in the literature, and has been the most common approach to this issue for years (see for example Mammen, 1991; Mammen & Thomas-Agnan, 1999; Mukerjee, 1988; Passow & Roulier, 1977; Ramsay, 1988; Winsberg & Ramsay, 1980, 1981). The basic idea is to introduce a flexible functional basis (e.g., splines) for representing the function and estimating the coefficients of the basis expansion by minimizing the residual sum of squares under the constraint of monotonicity of the estimated functions. Typical choices rely on the use of either an I-spline basis with a positive constraint on the coefficients or a B-spline basis with equally spaced knots and a monotonicity constraint on the coefficients. A similar approach has been proposed in the framework of kernel regression (Hall & Huang, 2001; Henderson, List, Millimet, Parmeter, & Price, 2008), where the kernels are modified locally in order to achieve monotonicity. Another approach is the projection method (Bloch & Silverman, 1997; Friedman & Tibshirani, 1984; Mammen, Marron, Turlach, & Wand, 2001), where the unknown monotonic function is estimated in an unconstrained fashion and then projected onto the convex subspace of the monotonic functions.

The approach that we are going to use here instead of these is in line with the so-called transform/back-transform method. This method is in common use in the FDA literature, and was initially proposed by Ramsay and Silverman (2002, 2005). Basically, the idea is to transform the functions so as to perform an unconstrained estimation, and then back-transform the estimated function to the convex subspace of the monotonic functions. Some very recent work (Boogaart, Egozcue, & Pawlowsky-Glahn, 2014; Egozcue, Díaz-Barrero, & Pawlowsky-Glahn, 2006; Menafoglio, Guadagnini, & Secchi, 2014) focusing on modeling the cumulative distribution functions of absolutely continuous random variables (inspired by the pioneering work on compositional data by Aitchison, 1982) formalized this approach by imposing a suitable Hilbert structure on the set of probability density functions and an isometric bijective map on L^2 for transforming and back-transforming functional data and conveniently mapping the entire statistical analysis in a linear subspace of L^2 (i.e., the zero-mean L^2 functions). The present paper, on the other hand, imposes a suitable pre-Hilbert structure, i.e., a non-complete vector space provided with an inner product, on the set of monotonic, lower and upper bounded functions that satisfy an equality constraint on one edge of the domain and an inequality constraint on the other, with an associated isometric bijective map to L^2 that allows us to model the temporal dependence in an unconstrained framework. In the remainder of this paper, we will refer to this as the $\mathcal{M}^2$ space. To the best of our knowledge, this is the first time that a geometry in a functional space has been introduced and formalized in order to obtain a sound theoretical framework for modeling the temporal dependencies among constrained functional data.

The literature dealing with the temporal dependencies among functional data is more recent, dating to the end of last century. The pioneering contribution to the topic is that of Bosq (1991), who derived a functional Yule–Walker estimator for time-dependent functional data. Functional autoregressive models (FAR) are the most commonly used approach for modeling temporal dependencies among functional data, due to both their ease of interpretation and their good performances in applications (Elezović, 2009). In FAR models, autoregressive parameters are replaced by Hilbert–Schmidt operators, and thus, model estimation is defined by the estimation of the autoregressive operators. Various different methods have been presented in the literature; for recent surveys, see Hormann and Kokoszka (2012) and Horváth and Kokoszka (2012). Autoregressive operators are linked directly with lagged autocovariance operators (e.g., Kargin & Onatski, 2008), and thus one possible approach would be to estimate the lagged autocovariance operators from the functional time series, and then to estimate the autoregressive operators accordingly. However, sample autocovariance operators are typically replaced by reduced rank approximations because of the infinite dimensionality of functional data, and in order to obtain more stable estimates. A spread approach relies on functional principal component decomposition (e.g., Aue, Norinho, & Hörmann, 2015; Hyndman & Shang, 2009; Hyndman & Ullah, 2007; Shang, 2013) and the use of reduced numbers of principal components. Other alternative reduced rank approximations that have been presented in the literature are based on wavelet expansions of the original data (Antoniadis & Sapatinas, 2003) and on predictive factors (Kargin & Onatski, 2008).

Another approach to the estimation of the autoregressive operators is the direct minimization of the mean squared error of prediction. However, the minimization problem has to be approached with some care in order to avoid over-fitting, due to the infinite dimensionality of functional data. For instance, Fan and Zhang (2000) and