

DYNAMIC PRINCIPAL COMPONENT REGRESSION: APPLICATION TO AGE-SPECIFIC MORTALITY FORECASTING

BY

HAN LIN SHANG 

ABSTRACT

In areas of application, including actuarial science and demography, it is increasingly common to consider a time series of curves; an example of this is age-specific mortality rates observed over a period of years. Given that age can be treated as a discrete or continuous variable, a dimension reduction technique, such as principal component analysis (PCA), is often implemented. However, in the presence of moderate-to-strong temporal dependence, *static* PCA commonly used for analyzing independent and identically distributed data may not be adequate. As an alternative, we consider a *dynamic* principal component approach to model temporal dependence in a time series of curves. Inspired by Brillinger's (1974, *Time Series: Data Analysis and Theory*. New York: Holt, Rinehart and Winston) theory of dynamic principal components, we introduce a dynamic PCA, which is based on eigen decomposition of estimated long-run covariance. Through a series of empirical applications, we demonstrate the potential improvement of 1-year-ahead point and interval forecast accuracies that the dynamic principal component regression entails when compared with the static counterpart.

KEYWORDS

Dimension reduction, functional time series, Kernel sandwich estimator, long-run covariance, multivariate time series.

1. INTRODUCTION

In many developed countries, increases in longevity and an aging population have led to concerns regarding the sustainability of pensions, health care, and aged-care systems (e.g., Organization for Economic Co-operation and Development [OECD], 2013). These concerns have resulted in a surge of

interest among government policymakers and planners to engage in accurate modeling and forecasting of age-specific mortality rates. In addition, forecasted mortality rates are an important input for determining annuity prices and thus are very important to pension and insurance industries (see, e.g., Shang and Haberman, 2017). Many statistical methods have been proposed for forecasting age-specific mortality rates (for reviews, see Booth and Tickle, 2008). Of these, a significant milestone in demographic forecasting was the work by Lee and Carter (1992). They implemented a principal component method to model age-specific mortality rates and extracted a single time-varying index of the level of mortality rates, from which the forecasts were obtained by a random walk with drift.

The strengths of the Lee–Carter (LC) method are its simplicity and robustness in situations where age-specific log mortality rates have linear trends (Booth et al., 2006). The main weakness of the LC method is that it attempts to capture the patterns of mortality rates using only one principal component and its associated scores. To rectify this deficiency, the LC method has been extended and modified. For example, from a time series of matrix perspective, Renshaw and Haberman (2003) proposed the use of more than one component in the LC method to model age-specific mortality. From a time series of function perspective, Hyndman and Ullah (2007) proposed a functional time-series method that uses nonparametric smoothing and higher-order principal components.

A common feature of the aforementioned works is that a static principal component analysis (PCA) is often used to decompose a time series of data matrix or curves. Under moderate-to-strong temporal dependence, the extracted principal components may not be consistent because of temporal dependence, leading to erroneous estimators. To overcome this issue, we consider a dynamic approach that extracts principal components based on an estimated long-run covariance instead of estimated variance alone. Note that the long-run covariance includes the variance function as a component, yet also measures temporal cross-covariance at different positive and negative lags. Similar to the finite-dimensional time-series framework, the long-run covariance estimation is the sum of empirical autocovariance functions and is often truncated at some finite lag in practice (see Section 4).

While the LC method is commonly used for analyzing mortality rates at discrete ages, the functional time-series method is often used for analyzing mortality curves where age is treated as a continuum. With these two methods, the contribution of this paper is to demonstrate the improvement of point and interval forecast accuracies that the dynamic principal component regression entails when compared with the static PCA for modeling and forecasting age-specific mortality rates at a 1-year-ahead forecast horizon. In the longer forecast horizon, the difference in forecast accuracy becomes marginal, and these results can be obtained upon request from the author.

The rest of this paper is structured as follows. In Section 4, we describe a kernel sandwich estimator for estimating long-run covariance. Based on the estimated long-run covariance, we introduce an eigen decomposition that

TABLE 1

THE 24 COUNTRIES EXAMINED IN THIS STUDY, WITH THE INITIAL YEAR OF 1950 AND THEIR FINAL YEAR LISTED BELOW.

Country	Abbreviation	Final year	Country	Abbreviation	Final year
Australia	AUS	2014	Italy	ITA	2012
Austria	AUT	2014	Japan	JPN	2014
Belgium	BEL	2015	The Netherlands	NLD	2014
Bulgaria	BGR	2010	Norway	NOR	2014
Canada	CAN	2011	New Zealand	NZ	2013
The Czech Republic	CZE	2014	Portugal	PRT	2015
Denmark	DEN	2014	Spain	SPA	2014
Finland	FIN	2015	Slovakia	SVK	2014
France	FRA	2014	Sweden	SWE	2014
Hungary	HUN	2014	Switzerland	SWI	2014
Iceland	ICE	2013	The United Kingdom	UK	2013
Ireland	IRL	2014	The United States	US	2015

extracts dynamic principal components and their associated scores in Section 5. Illustrated by empirical data obtained from the Human Mortality Database (2019) in Section 2, we evaluate and compare the 1-year-ahead point and interval forecast accuracies between the LC and functional time-series methods described in Section 3, using both the static and dynamic principal component regression models in Section 6. Conclusions are given in Section 7.

2. DATA SETS

The data sets used in this study were taken from the Human Mortality Database (2019). For each sex in a given calendar year, the mortality rates obtained by the ratio between “number of deaths” and “exposure to risk” were arranged in a matrix for age and calendar year. Twenty-four countries, mainly developed nations, were selected, and thus 48 sub-populations of age- and sex-specific mortality rates were obtained for all analyses. The 24 countries selected all had reliable data series commencing during or before 1950. As a result of possible structural breaks (i.e., two world wars), we truncated all data series from 1950 onwards. The omission of Germany was because the Human Mortality Database for a reunited Germany only goes back to 1990. The selected countries are shown in Table 1, alongside their final year of available data. To avoid fluctuation in older ages, we considered ages from 0 to 99 in a single year of age, and the last age group was from 100 onwards.

2.1. Functional time-series plot

To present an evolution of age-specific mortality, we present a functional time-series plot for the raw female log mortality rates in the US in Figure 1(a),

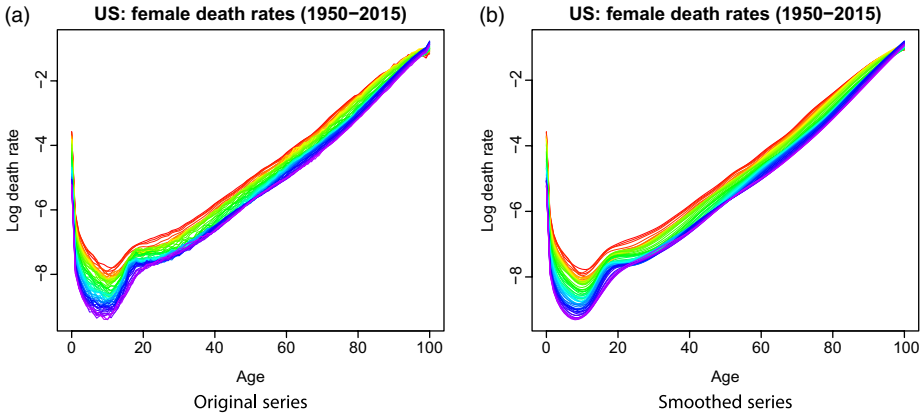


FIGURE 1: Observed and smoothed age-specific female log mortality rates in the US. Data from the distant past are shown in red, and the most recent data are shown in purple.

while the functional time-series plot for the smoothed data is shown in Figure 1(b).

To smooth these functional time series, we assumed there was an underlying L_2 continuous and smooth function $f_t(x)$, such that

$$\mathcal{Y}_t(x_j) = f_t(x_j) + \sigma_t(x_j)\epsilon_{t,j}, \quad j = 1, \dots, p, \quad t = 1, \dots, n,$$

where $\mathcal{Y}_t(x_j)$ denotes the raw log mortality rates, $f_t(x_j)$ denotes the smoothed log mortality rates, $\{\epsilon_{t,j}\}$ represents independent and identically distributed (IID) random variables across t and j with a mean of zero and a unit variance, and $\sigma_t(x_j)$ allows for heteroskedasticity and can be estimated by

$$\hat{\sigma}_t(x_j) = \frac{1}{\exp\{\mathcal{Y}_t(x_j)\}E_t(x_j)},$$

where $E_t(x_j)$ denotes population of age x_j at June 30 in year t (often known as the “exposure-at-risk”).

Given that the log mortality rates increased linearly over age, we used a penalized regression spline with monotonic constraint, where the monotonicity was imposed for ages at and above 65 (for details, see Hyndman and Ullah, 2007). With the weights equal to the inverse variances $w_t(x_j) = 1/\hat{\sigma}_t^2(x_j)$, the smoothed log mortality rate was obtained by

$$f_t(x_j) = \operatorname{argmin}_{\theta_t(x_j)} \sum_{j=1}^M w_t(x_j) |\mathcal{Y}_t(x_j) - \theta_t(x_j)| + \lambda \sum_{j=1}^{M-1} |\theta'_t(x_{j+1}) - \theta'_t(x_j)|,$$

where x_j represents different ages (grid points) in a total of M grid points, λ denotes a smoothing parameter, and θ' denotes the first derivative of smooth function θ , which can both be approximated by a set of B-splines.

Figure 1 is an example of the rainbow plot, where the colors of the curves follow the order of a rainbow with the oldest data shown in red and most recent data shown in violet (see also Hyndman and Shang, 2010). By analyzing the changes in mortality as a function of both age x and year t , it can be seen that mortality rates showed a gradual decline over the years. Mortality rates dipped from their early childhood high, climbed in the teen years, stabilized in the early 20s, and then steadily increased with age. We further noted that, for both males and females, log mortality rates declined over time, especially in the younger and older ages.

2.2. Mortality improvement rate

In demography and actuarial science, a time series of age-specific mortality rates is commonly modeled and forecast at a logarithmic scale. These series are nonstationary, as the mean function changes over time. As an alternative approach, we can model the improvement in mortality rates, rather than the rate itself (see, e.g., Haberman and Renshaw, 2012). The advantage of modeling the mortality improvement is that the data series is stationary. As one way of measuring mortality improvement, the year-on-year mortality improvement rate of Haberman and Renshaw (2012) was considered and expressed as

$$z_{x,t} = 2 \times \frac{1 - m_{x,t}/m_{x,t-1}}{1 + m_{x,t}/m_{x,t-1}} = 2 \times \frac{m_{x,t-1} - m_{x,t}}{m_{x,t-1} + m_{x,t}}, \quad t = 2, \dots, n, \quad (2.1)$$

for age x in year t , where $m_{x,t}$ denotes the raw mortality rate, $z_{x,t}$ denotes the transformed mortality rate, and n symbolizes the number of years.

The expression in Equation (2.1) can be seen as the ratio between the incremental mortality improvement ($m_{x,t-1} - m_{x,t}$) and the average ($m_{x,t} + m_{x,t-1}$)/2 of two adjacent mortality rates. By defining the denominator of the ratio in this way, we avoided the small phase difference between the numerator and denominator that would otherwise be the case. Thus, improving incremental mortality rate changes implied $z_{x,t} > 0$, while deteriorating incremental changes implied $z_{x,t} < 0$.

Via back-transformation of Equation (2.1), we obtained

$$m_{x,t} = \frac{2 + z_{x,t}}{2 - z_{x,t}} \times m_{x,t-1}.$$

In Figure 2(a) and (b), we plot the observed and smoothed curves for the age-specific female mortality rate improvements in the US. The curves are stationary and more volatile in the early ages (i.e., ages between 0 and 40) than the later ages. We obtained smoothed mortality rate improvement by computing the smoothed age-specific mortality curve, and then applying Equation (2.1).

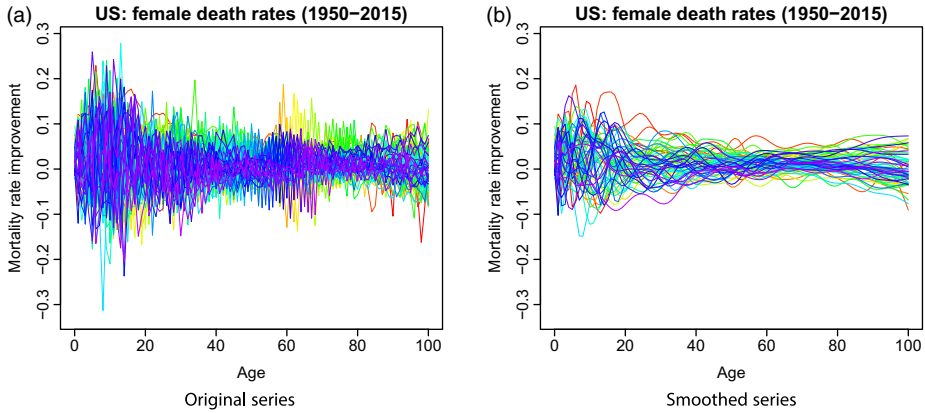


FIGURE 2: Observed and smoothed age-specific female mortality rate improvements in the US.

3. FORECASTING METHODS

Given that the focus of this paper is a comparison of short-term forecast accuracy between the static and dynamic principal component analyses, we revisited the LC and functional time-series methods as two possible methods for forecasting age-specific mortality rates. The LC model considers age a discrete variable, while the functional time-series model treats age as a continuous variable. We denoted with $m_{x,t}$ the observed mortality rate at age x in year t , calculated as the number of deaths aged x in year t , divided by the corresponding mid-year population aged x in year t . With $m_{x,t}$, we first obtained transformed series $z_{x,t}$ from Equation (2.1).

3.1. Adapted LC method

The original LC method was applied to model the log mortality rate (Lee and Carter, 1992). Here, we extended it to model mortality improvement rate. The formulation of our adapted LC model is given by

$$z_{x,t} = a_x + b_x \kappa_t + \epsilon_{x,t}, \tag{3.1}$$

where a_x denotes the age pattern of the mortality rates averaged over years, b_x denotes the first principal component at age x , κ_t denotes the first set of principal component scores at year t and measures the general level of the mortality rates, and $\epsilon_{x,t}$ denotes the residual at age x and year t .

The LC model in Equation (3.1) is overparametrized, in that the model structure is invariant under the following transformations:

$$\begin{aligned} \{a_x, b_x, \kappa_t\} &\mapsto \{a_x, b_x/c, c\kappa_t\}, \\ \{a_x, b_x, \kappa_t\} &\mapsto \{a_x - cb_x, b_x, \kappa_t + c\}, \end{aligned}$$

To ensure the model identifiability, Lee and Carter (1992) imposed two constraints given as

$$\sum_{t=1}^n \kappa_t = 0, \quad \sum_{x=x_1}^{x_p} b_x = 1,$$

where n denotes the number of years and p denotes the number of ages in the observed data set.

Instead of using a random walk with drift, the set of principal component scores, κ_t , can be extrapolated using autoregressive integrated moving average (ARIMA) models. We used the automatic algorithm of Hyndman and Khandakar (2008) to choose the optimal orders of autoregressive p , moving average q , and difference order d . d was selected based on successive Kwiatkowski–Phillips–Schmidt–Shin (KPSS) unit root tests (Kwiatkowski et al., 1992). KPSS tests were used to test the null hypothesis that an observable time series was stationary around a deterministic trend. We tested the original data (i.e., the first set of principal component scores) for a unit root; if the test result was significant, then we tested the differenced data for a unit root. The procedure continued until we obtained our first insignificant result. Having determined d , the orders of p and q were selected based on the optimal Akaike information criterion with a correction for finite sample sizes (Akaike, 1974). After identifying the optimal ARIMA model, the maximum likelihood method could then be used to estimate the parameters. Conditioning on the estimated mean, the estimated first principal component \widehat{b}_x , and the observed mortality rate improvement, the h -step-ahead point forecast of $z_{x,n+h}$ can be expressed as

$$\widehat{z}_{x,n+h|n} = E \left[z_{x,n+h} \mid z_{x,1}, \dots, z_{x,n}, \widehat{a}_x, \widehat{b}_x \right] = \widehat{a}_x + \widehat{b}_x \widehat{\kappa}_{n+h|n},$$

where $\widehat{a}_x = \frac{1}{n} \sum_{t=1}^n z_{x,t}$ denotes the estimated mean, and $\widehat{\kappa}_{n+h|n}$ denotes the h -step-ahead forecast of the principal component scores.

In Haberman and Renshaw (2012), the LC method does not include the mean a_x , since their generalized linear model approach uses a Newton–Raphson iterative fitting algorithm to estimate b_x and κ_t by minimizing a deviance criterion. In contrast, we applied the PCA to mortality improvement $z_{x,t}$. The PCA often requires de-centering the data. In Figure 6, we also compare the one-step-ahead forecast performances of the LC method under a Poisson error structure without centering. We found that the difference regarding whether or not to center the data was marginal in terms of forecast accuracy. In Appendix B, we also compare the 5-step-ahead and 10-step-ahead forecast accuracy using the LC method with and without centering.

3.2. Functional time-series method

Functional time series often consist of random functions observed at regular time intervals. In the context of mortality, functional time series can arise when

observations in a time period can be considered together as finite realizations of an underlying continuous function (e.g., Hyndman and Ullah, 2007). There are several advantages to consider the functional time-series method:

- (1) Data points may be observed sparsely. Via the functional time-series method, the underlying trajectory may be recovered (for details, see Zhang and Wang, 2016).
- (2) With continuity, derivative information can provide new insight into data analysis (see, e.g., Shang, 2019).
- (3) A nonparametric smoothing technique can be incorporated into the modeling procedure to obtain smoothed principal components. Smoothing deals with one criticism of the LC model; namely, that the estimated values, b_x , can be subject to considerable noise and, without smoothing, this would be propagated into forecasts of future mortality rates. Smoothing can reduce measurement error and increase the signal-to-noise ratio, and also deals with estimating missing values for some ages at a given year.
- (4) Given that the functional time-series method can consider more than one component, Shang (2012) indicated that the functional time-series method outperforms the LC method.

Many possible nonparametric smoothing techniques have been proposed, such as basis spline (for details, see de Boor, 2001). We used a penalized regression spline with a monotonic constraint (for details, see Hyndman and Ullah, 2007). The smoothed log mortality rates $f(x) = \{f_1(x), \dots, f_n(x)\}$ were treated as realizations of a stochastic process. In Hyndman and Ullah (2007), they considered modeling the smoothed log mortality rates directly. Given that the log mortality rates are nonstationary, we considered modeling and forecasting mortality rate improvement. From Equation (2.1), we could obtain a set of transformed and smoothed series, denoted by $z(x) = \{z_1(x), \dots, z_n(x)\}$. Using functional PCA, these smoothed mortality improvement curves were decomposed into

$$z_t(x) = a(x) + \sum_{k=1}^K b_k(x)\kappa_{t,k} + e_t(x), \quad t = 1, \dots, n, \quad (3.2)$$

where $a(x)$ denotes the mean function estimated by $\hat{a}(x) = \frac{1}{n} \sum_{t=1}^n z_t(x)$, $\{b_1(x), \dots, b_K(x)\}$ denotes a set of functional principal components, $\{\kappa_{t,1}, \dots, \kappa_{t,K}\}$ denotes a set of principal component scores in year t , $e_t(x)$ denotes the error function with mean zero, and $K < n$ denotes the number of principal components retained.

Decomposition in Equation (3.2) facilitates dimension reduction because the first K terms often provide a reasonable approximation to the infinite sums; thus, the information contained in $z(x)$ can be adequately summarized by the K -dimensional vector, $\Phi = [b_1(x), \dots, b_K(x)]$. In contrast to the LC model, another advantage of the functional time-series model is that more than one component may be used to improve model fitting (see also Renshaw and Haberman, 2003). Here, the number of components is determined as

the minimum that reaches a certain level of the proportion of total variance explained by the leading components, such that

$$K = \operatorname{argmin}_{K:K \geq 1} \left\{ \sum_{k=1}^K \widehat{\lambda}_k / \sum_{k=1}^{\infty} \widehat{\lambda}_k \mathbb{1}_{\{\widehat{\lambda}_k > 0\}} \geq 85\% \right\}, \tag{3.3}$$

where $\widehat{\lambda}_k$ represents the k th estimated eigenvalue, and $\mathbb{1}_{\{\widehat{\lambda}_k > 0\}}$ is to exclude possible zero eigenvalues, and $\mathbb{1}\{\cdot\}$ represents the binary indicator function. The threshold of 85% is advocated in Horvath and Kokoszka (2012, p. 41).

Conditioning on the estimated mean function $\widehat{a}(x)$, the estimated functional principal components Φ , and the observed mortality rate improvement $z(x)$, the h -step-ahead point forecast of $z_{n+h}(x)$ can be expressed as

$$\begin{aligned} \widehat{z}_{n+h|n}(x) &= E [z_{n+h}(x) | z(x), \widehat{a}(x), \Phi] \\ &= \widehat{a}(x) + \sum_{k=1}^K \widehat{b}_k(x) \widehat{\kappa}_{n+h|n,k}, \end{aligned}$$

where $\widehat{a}(x)$ denotes the estimated mean function, $\widehat{b}_k(x)$ denotes the k th estimated functional principal component, and $\widehat{\kappa}_{n+h|n,k}$ denotes the k th estimated principal component scores obtained via a univariate or multivariate time-series forecasting method. Given that it can handle nonstationarity, we considered a univariate forecasting method, such as the ARIMA model with orders selected automatically.

The critical component of the aforementioned forecasting methods is the static PCA which was designed for IID data. In the presence of moderate-to-strong dependent data, the static PCA is not optimal because it does not incorporate autocovariance at different lags in a functional time series. As an alternative, we introduced a dynamic PCA (DPCA) constructed from an eigen decomposition of an estimated long-run covariance. The long-run covariance included the variance and autocovariance at lags greater than zero.

4. LONG-RUN COVARIANCE AND ITS ESTIMATION

In statistics, long-run covariance enjoys vast literature in the case of finite-dimensional time series, beginning with the seminal work of Brillinger (1974), and is still the most commonly used technique for smoothing the periodogram by employing a smoothing weight function and a bandwidth parameter. In the functional time series, long-run covariance plays an important role in modeling temporal dependence (see, e.g., Rice and Shang, 2017).

To provide a formal definition of the long-run covariance, let $\{z_t(x)\}_{t \in Z}$ be a stationary and ergodic functional time series. For example, $z_t(x)$ could be used to denote the density of pollutants in a given city on day t at intraday time x or the mortality rate in year t at age x . If $z_t(x)$ is nonstationary, it

could be suitably transformed, so that the stationarity assumption holds. For a stationary functional time series, the long-run covariance is defined as

$$\begin{aligned}
 C(x, u) &= \sum_{\ell=-\infty}^{\infty} \gamma_{\ell}(x, u) \\
 &= \sum_{\ell=-\infty}^{\infty} \text{cov} [z_0(x), z_{\ell}(u)].
 \end{aligned}$$

Given that $\gamma_{\ell}(x, u)$ is symmetric and nonnegative definite for any ℓ , $C(x, u)$ is also symmetric and nonnegative definite. By applying eigen decomposition to the long-run covariance, $C(x, u)$, we obtained a set of eigenvalues and eigenfunctions.

In practice, we needed to estimate C from a finite sample $\mathbf{z}(x) = [z_1(x), \dots, z_n(x)]$. Given its definition as a bi-infinite sum, a natural estimator of C is

$$\widehat{C}_{h,q}(x, u) = \sum_{\ell=-\infty}^{\infty} W_q \left(\frac{\ell}{h} \right) \widehat{\gamma}_{\ell}(x, u), \tag{4.1}$$

where h is called the bandwidth parameter, and

$$\widehat{\gamma}_{\ell}(x, u) = \begin{cases} \frac{1}{n} \sum_{j=1}^{n-\ell} [z_j(x) - \bar{z}(x)] [z_{j+\ell}(u) - \bar{z}(u)], & \ell \geq 0 \\ \frac{1}{n} \sum_{j=1-\ell}^n [z_j(x) - \bar{z}(x)] [z_{j+\ell}(u) - \bar{z}(u)], & \ell < 0, \end{cases}$$

is an estimator of $\gamma_{\ell}(x, u)$, and W_q is a symmetric weight function with bounded support of order q . The estimator in Equation (4.1) was introduced in Horváth and Kokoszka (2012) and Rice and Shang (2017), among others. As with the kernel estimator, the crucial part is the estimation of bandwidth parameter h . It can be selected through a data-driven approach, such as the plug-in algorithm proposed in Rice and Shang (2017). In Appendix A, we briefly describe the plug-in algorithm. With the estimated long-run covariance, we could obtain dynamic functional principal components and their scores, as described in Section 5.

4.1. Application to US age-specific mortality rates

Figure 3 presents the estimated long-run covariance and variance for the female raw mortality rates in the US. With the input data as age-specific mortality improvement over years, we computed the sample long-run covariance and sample variance for a data matrix. The sample variance was computed by

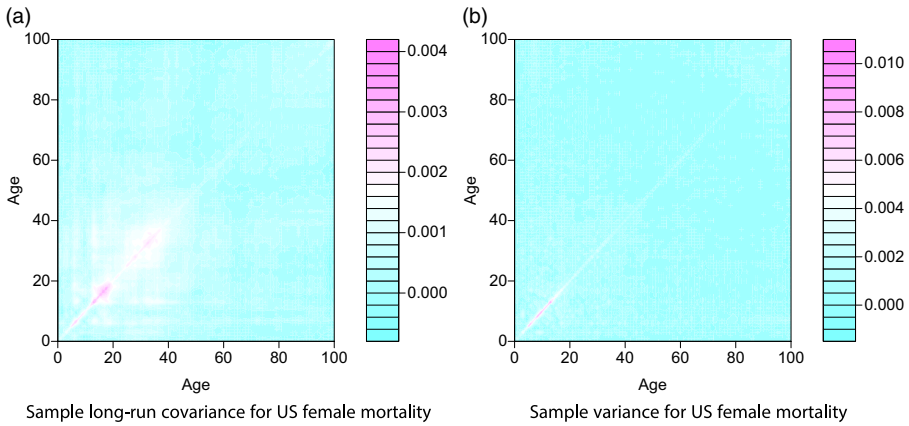


FIGURE 3: Comparison of sample long-run covariance and sample variance for US original female mortality rates.

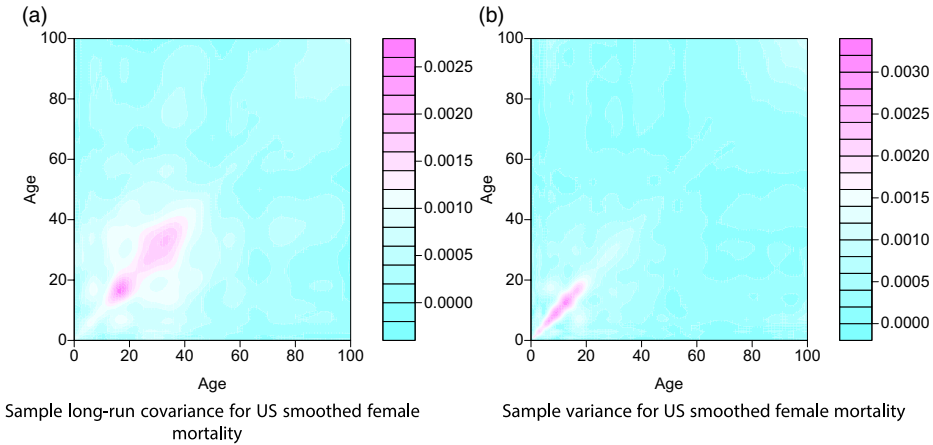


FIGURE 4: Comparison of sample long-run covariance and sample variance for US smoothed female mortality rates.

multiplying the data matrix by its transpose, while the sample long-run covariance was computed by the kernel sandwich estimator. Compared with the sample variance, the sample long-run covariance based on the kernel sandwich estimator with plugged-in bandwidth could incorporate an autocovariance structure, particularly for ages between 0 and 40. The mortality rate at the young ages exhibited higher variance than the mortality rate at the other ages.

Compared with the long-run covariance based on the raw data series, the estimated long-run covariance based on the smoothed data series was smoother and showed a more explicit data structure in Figure 4. For estimating the long-run covariance, the estimated optimal bandwidth was 4.07 for the raw female data and 4.10 for the smooth female data.

The computation of long-run covariance relies heavily on the fact that our mortality improvement data were stationary time series. When the temporal dependence was weak, variance could be sufficient to estimate the long-run covariance. When the temporal dependence was moderate or high, one should include more terms in the long-run covariance estimation. The inclusion of autocovariance could improve forecast accuracy. From Figures 3 and 4, it is clear that the long-run covariance also included the autocovariance at various lags.

5. DYNAMIC FUNCTIONAL PCA

From the long-run covariance, we applied functional principal component decomposition to extract the functional principal components and their associated scores. Via the Karhunen–Loève expansion, a stochastic process, z , can be expressed as

$$z(x) = a(x) + \sum_{k=1}^{\infty} \beta_k \phi_k(x),$$

where $z^c(x) = z(x) - a(x)$. The principal component scores, β_k , are given by the projection of $z^c(x)$ in the direction of the k th eigenfunction ϕ_k —that is, $\beta_k = \langle z^c(x), \phi_k(x) \rangle$. The scores constitute an uncorrelated sequence of random variables with zero mean and variance λ_k . They can be interpreted as the weights of the contribution of the functional principal components $\phi_k(x)$ to $z^c(x)$.

Given that the long-run covariance, $C(x, u)$, is unknown, the population eigenvalues and eigenfunctions can only be approximated through realizations of $z(x)$. A realization of the stochastic process, z , can be written as

$$z_t(x) = \hat{a}(x) + \sum_{k=1}^K \hat{\beta}_{t,k} \hat{\phi}_k(x) + e_t(x), \quad t = 1, 2, \dots, n,$$

where $\hat{\beta}_{t,k}$ is the k th estimated score for the t th year and $e_t(x)$ denotes residual function.

In Figure 5, we present the first eigenfunction extracted from the sample variance and sample long-run covariance, respectively, for the US female mortality. Visually, the first dynamic principal component appears differently to the first static principal component. Conditioning on the estimated mean function $\hat{a}(x)$, the estimated dynamic functional principal components $\Psi = \{\phi_1, \dots, \phi_K\}$, and observed mortality rate improvement, $z(x)$, the h -step-ahead point forecast of $z_{n+h}(x)$ is

$$\begin{aligned} \hat{z}_{n+h|n}(x) &= E [z_{n+h}(x) | z(x), \hat{a}(x), \Psi] \\ &= \hat{a}(x) + \sum_{k=1}^K \hat{\beta}_{n+h|n,k} \hat{\phi}_k(x), \end{aligned}$$

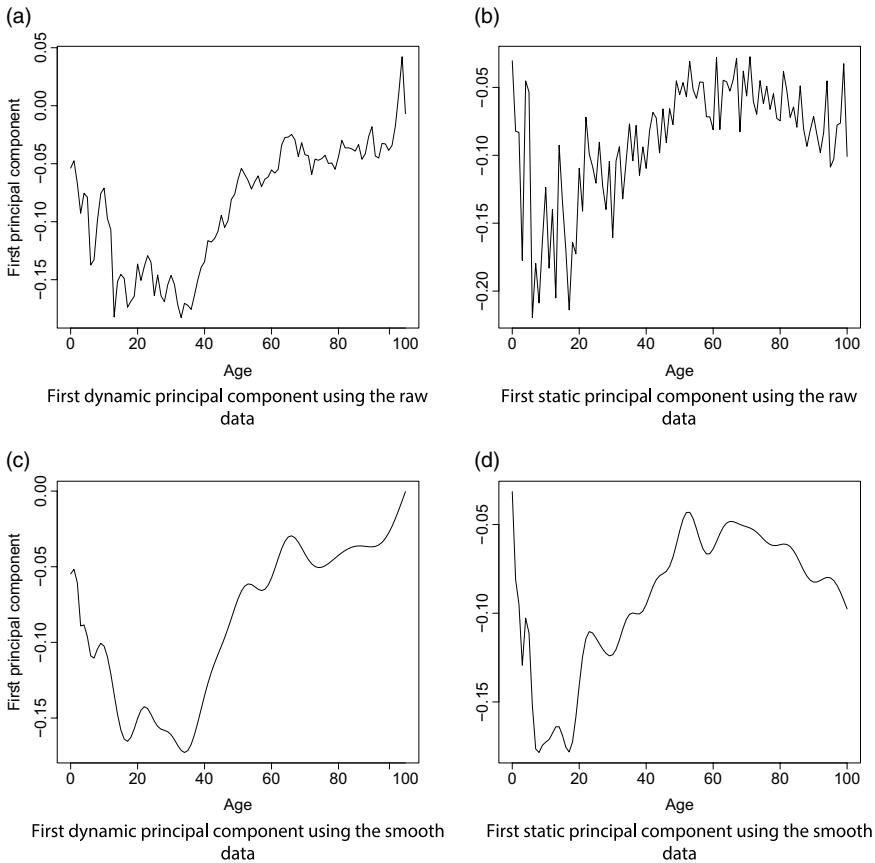


FIGURE 5: First eigenfunction extracted from the static principal component decomposition based on the sample variance, and first eigenfunction extracted from the dynamic principal component decomposition based on the sample long-run covariance. We considered the raw US mortality data in the top row, as well as the smooth US mortality data in the bottom row.

where $\hat{\phi}_k(x)$ denotes the k th estimated functional principal component, $\hat{\beta}_{n+h|n,k}$ denotes the k th estimated principal component scores obtained via a univariate forecasting method, and $K < n$ denotes the number of principal components retained. In practice, the optimal value of K can be selected by explaining at least 85% of total variation in the data; refer to Equation (3.3).

5.1. Constructing prediction intervals

We considered a nonparametric bootstrap method for constructing prediction intervals (for details, see Hyndman and Shang, 2009). The source of uncertainty stemmed from the estimation error in the principal component scores and model residual errors.

Using a univariate time series forecasting method, we could obtain multi-step-ahead forecasts for the principal component scores, $\{\widehat{\beta}_{1,k}, \dots, \widehat{\beta}_{n,k}\}$. Let the h -step-ahead forecast errors denote

$$\widehat{\xi}_{t,h,k} = \widehat{\beta}_{t,k} - \widehat{\beta}_{|t-h,k}, \quad t = h + 1, \dots, n.$$

The estimation errors were then sampled with replacement to give a bootstrap sample of $\beta_{n+h,k}$:

$$\widehat{\beta}_{n+h|n,k}^{(b)} = \widehat{\beta}_{n+h|n,k} + \widehat{\xi}_{*,h,k}^{(b)}, \quad b = 1, \dots, B,$$

where $\widehat{\xi}_{*,h,k}^{(b)}$ are sampled with replacement from $\widehat{\xi}_{t,h,k}$, and B denotes the number of bootstrap replications. As long as the first K principal components approximate the data relatively well, the model residual should be random noise. We could bootstrap the model error by sampling with replacement from the residual term $\{\widehat{e}_1(x), \dots, \widehat{e}_n(x)\}$.

Through combining the two sources of uncertainty, we obtained B variants for $\widehat{z}_{n+h|n}^{(b)}(x)$:

$$\widehat{z}_{n+h|n}^{(b)}(x) = \widehat{a}(x) + \sum_{k=1}^K \widehat{\beta}_{n+h|n,k}^{(b)} \widehat{\phi}_k(x) + \widehat{e}_{n+h|n}^{(b)}(x).$$

Pointwise prediction intervals were produced from the bootstrap variants using quantiles.

6. RESULTS

6.1. Forecast evaluation

We presented 24 countries with data that began in 1950 and ended in the final year listed in Table 1. We retained the final 30 observations for forecasting evaluation, while the remaining observations were treated as initial fitting observations, from which we produced the *one-step-ahead* forecast (i.e., 1-year-ahead forecast). Via an expanding window approach, we re-estimated the parameters in the time-series forecasting models by increasing the fitted observations by 1 year and producing the one-step-ahead forecast. We iterated this process by increasing the sample size by 1 year until reaching the end of the data period. This process produced 30 one-step-ahead forecasts. We compared these forecasts with the holdout samples to determine the out-of-sample forecast accuracy.

6.2. Forecast error criteria

To evaluate the point forecast accuracy, we considered the mean absolute forecast error (MAFE) and root mean squared forecast error (RMSFE). These

criteria measured the closeness of the forecasts in comparison with the actual values of the variable being forecast, regardless of the direction of forecast errors. The MAFE and RMSFE are defined as

$$\begin{aligned} \text{MAFE}_h &= \frac{1}{p \times q} \sum_{j=1}^q \sum_{i=1}^p |m_j(x_i) - \widehat{m}_{j|j-h}(x_i)|, \\ \text{RMSFE}_h &= \frac{1}{p \times q} \sum_{j=1}^q \sum_{i=1}^p \sqrt{[m_j(x_i) - \widehat{m}_{j|j-h}(x_i)]^2}, \end{aligned}$$

where q represents the number of years in the forecasting period, $p \times q$ counts the total number of data points in the forecasting period, $m_j(x_i)$ represents the actual holdout sample for age x_i in year j , and $\widehat{m}_j(x_i)$ represents the forecasts for the holdout sample.

To evaluate the pointwise interval forecast accuracy, we considered the coverage probability deviance (CPD) of Shang (2012) and interval score criterion of Gneiting and Raftery (2007). We considered the common case of the symmetric $100(1 - \alpha)\%$ prediction intervals, with lower and upper bounds that were predictive quantiles at $\alpha/2$ and $1 - \alpha/2$, denoted by $\widehat{m}_j^l(x_i)$ and $\widehat{m}_j^u(x_i)$. The CPD allows comparison of interval forecast accuracy for each method by measuring the differences between the empirical coverage and nominal coverage probabilities. The CPD is defined as

$$\left| \frac{\mathbb{1} \{m_j(x_i) < \widehat{m}_j^l(x_i)\} + \mathbb{1} \{m_j(x_i) > \widehat{m}_j^u(x_i)\}}{p \times q} - \alpha \right|,$$

where $\mathbb{1}\{\cdot\}$ denotes binary indicator, and α denotes the level of significance, customarily $\alpha = 0.2$.

As defined by Gneiting and Raftery (2007), a scoring rule for evaluating the pointwise interval forecast accuracy at time point x_i is

$$\begin{aligned} S_\alpha[\widehat{m}_j^l(x_i), \widehat{m}_j^u(x_i); m_j(x_i)] &= [\widehat{m}_j^u(x_i) - \widehat{m}_j^l(x_i)] \\ &+ \frac{2}{\alpha} [\widehat{m}_j^l(x_i) - m_j(x_i)] \mathbb{1} \{m_j(x_i) < \widehat{m}_j^l(x_i)\} \\ &+ \frac{2}{\alpha} [m_j(x_i) - \widehat{m}_j^u(x_i)] \mathbb{1} \{m_j(x_i) > \widehat{m}_j^u(x_i)\}. \end{aligned}$$

The optimal interval score is achieved when $m_j(x_i)$ lies between $\widehat{m}_j^l(x_i)$ and $\widehat{m}_j^u(x_i)$, with the distance between the upper bound and lower bound being minimal. To obtain summary statistics of the interval score, we take the mean interval score across different ages and forecasting years. The mean interval score can be expressed as

$$\bar{S}_{\alpha,h} = \frac{1}{p \times q} \sum_{j=1}^q \sum_{i=1}^p S_\alpha[\widehat{m}_j^l(x_i), \widehat{m}_j^u(x_i); m_j(x_i)].$$

6.3. Comparisons of forecast errors

In Figure 6, we compare the one-step-ahead point forecast errors between the dynamic and static principal component regression models using the LC method with and without centering and functional time-series method.

In Figure 7, we compare the one-step-ahead interval forecast errors between the dynamic and static principal component regression models using the LC method with centering and functional time-series method.

In Table 2, we compare the summary statistics of the one-step-ahead point forecast errors between the dynamic and static principal component regression models using the LC method with and without centering and functional time-series method. Given that the point forecast errors between the LC method with and without centering are marginal, in Table 3, we compare the summary statistics of the one-step-ahead interval forecast accuracy between the dynamic and static principal component regression models using the LC method with centering and functional time-series method.

We observed the following evidence:

(1) Averaged over the 24 mainly developed countries in Table 1, the dynamic principal component regression outperformed the static principal component regression in terms of point and interval forecast accuracies, as measured by summary statistics of the MAFE and RMSFE, CPD and mean interval score criteria.

(2) In contrast to the results of the LC method, it was advantageous to smooth the data before computing the long-run covariance and variance, because this approach generally produced smaller point and interval forecast errors.

(3) Among the female, male, and total series, it was generally easier to forecast the total series as evident from smaller forecast errors, while it was harder to forecast the male series as evident from larger forecast errors. This could be because the total variation of the male series was larger than that of the total series.

The superiority of the dynamic principal component regression could be because it captures temporal dependence better than the static principal component regression.

We also considered the 5-step-ahead and 10-step-ahead point forecast accuracy, and found a marginal difference between the dynamic and static approaches. The results are provided in Appendix C. In time-series forecasting, many time-series extrapolation models were no longer optimal when we projected long term. As the forecast horizon increased, the proposed DPCA reduced back to the static PCA used in the LC model. The difference between the two derived crucially from the criterion used to extract latent components. In the DPCA, the criterion was long-run covariance, which was a sum of the variance and autocovariance. In the static PCA, the criterion was variance alone. In the long-term forecast, the distant future value was almost IID to the most recent value. In turn, the autocovariance was small, if it existed at all, at a

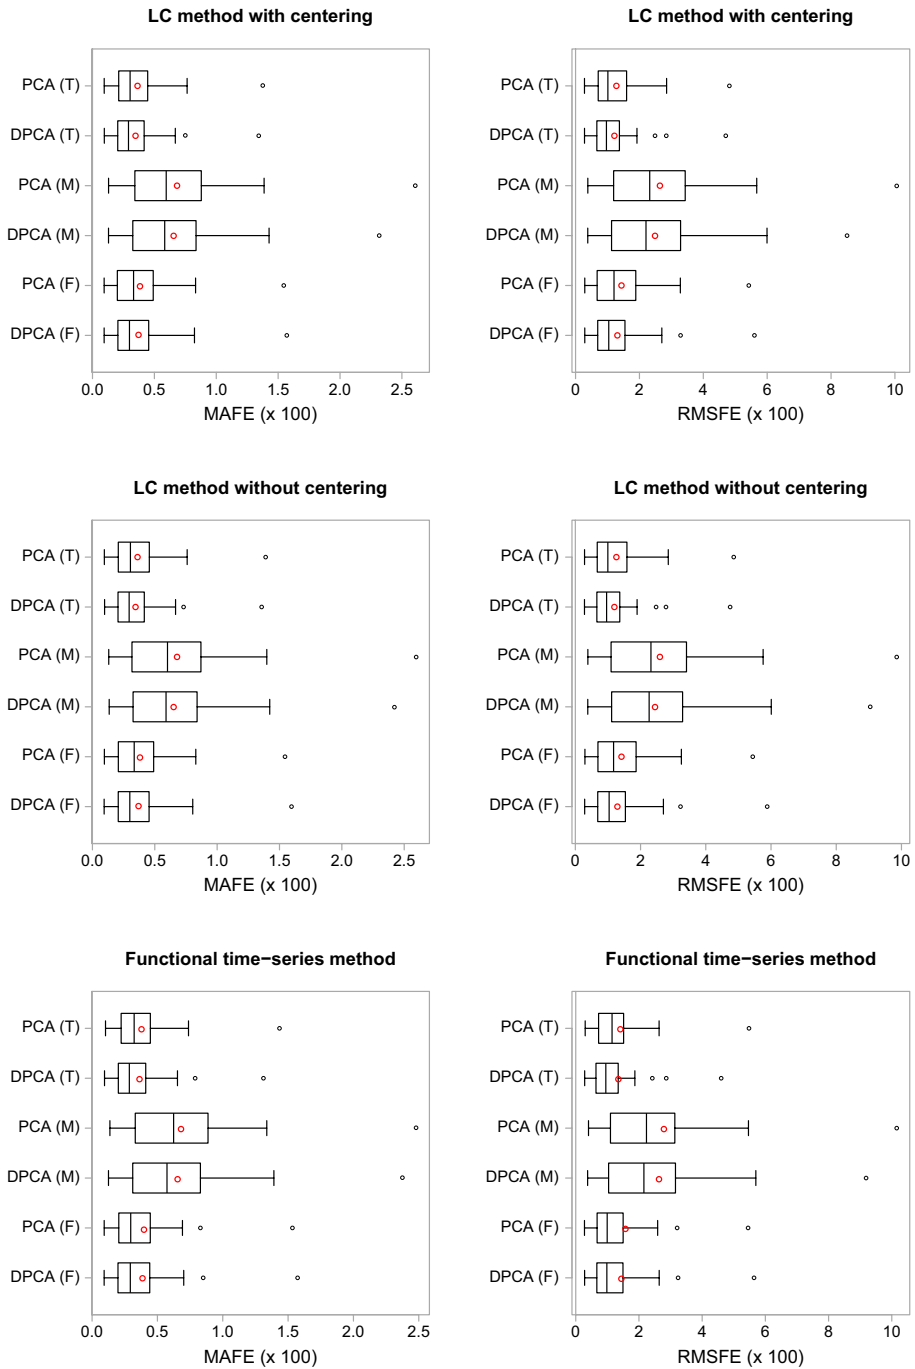


FIGURE 6: Boxplots of one-step-ahead point forecast errors between the DPCA and static PCA using the LC method with and without centering and functional time-series method. The red circle represents the mean.

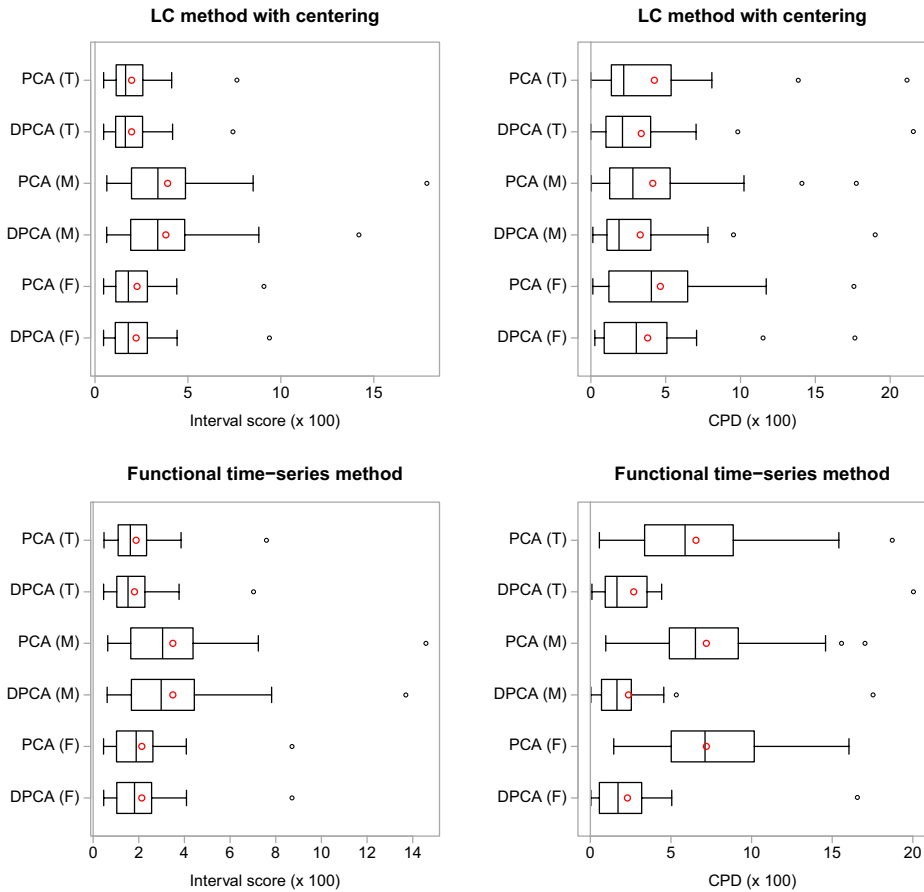


FIGURE 7: Boxplots of one-step-ahead interval forecast errors between the dynamic and static principal component regression using the LC method with centering and functional time-series method. The red circle represents the mean.

long-term forecast horizon. Therefore, the proposed DPCA (almost) reduced back to the static PCA, and the DPCA did not display advantages over the static PCA at a long-term forecast horizon. In addition, as the forecast horizon increased, the forecast of principal component scores was likely to be centered around zero. When that occurred, the forecasts of mortality rates were not so informative and would likely center around the mean function.

7. DISCUSSION

PCA performs dimension reduction (also known as data coarsening) with a minimal loss of information, and is a workhorse in time-series modeling of

TABLE 2

ONE-STEP-AHEAD POINT FORECAST ERRORS BETWEEN THE DPCA AND STATIC PCA USING THE LC METHOD WITH AND WITHOUT CENTERING AND FUNCTIONAL TIME-SERIES METHOD. DUE TO LIMITATIONS OF SPACE, WE REPORT THE SUMMARY STATISTICS OF THE FORECAST ERRORS OBTAINED FROM THE 24 COUNTRIES.

Statistic	MAFE × 100												RMSFE × 100											
	Female				Male				Total				Female				Male				Total			
	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA	DPCA	PCA		
LC method with centering																								
Min.	0.094	0.094	0.129	0.130	0.095	0.095	0.095	0.095	0.288	0.288	0.288	0.288	0.288	0.382	0.384	0.280	0.279	0.720	0.699	1.154	1.247	0.687	0.724	
1st Qu.	0.214	0.208	0.339	0.357	0.211	0.217	0.211	0.217	0.720	0.699	0.720	0.699	1.154	1.247	0.687	0.724	1.011	1.039	1.204	2.322	2.322	0.967	1.011	
Median	0.299	0.332	0.583	0.596	0.291	0.304	0.291	0.304	1.039	1.204	1.039	1.204	2.322	2.661	1.252	1.330	1.330	1.366	1.479	2.532	2.661	1.252	1.330	
Mean	0.381	0.396	0.666	0.693	0.359	0.374	0.359	0.374	1.366	1.479	1.366	1.479	2.532	2.661	1.252	1.330	1.330	1.535	1.867	3.218	3.385	1.360	1.585	
3rd Qu.	0.452	0.488	0.837	0.867	0.416	0.440	0.416	0.440	1.535	1.867	1.535	1.867	3.218	3.385	1.360	1.585	1.585	5.625	5.433	8.516	10.072	4.714	4.819	
Max.	1.569	1.544	2.330	2.621	1.349	1.380	1.349	1.380	5.625	5.433	5.625	5.433	8.516	10.072	4.714	4.819	4.819							
LC method without centering																								
Min.	0.094	0.096	0.134	0.131	0.097	0.096	0.097	0.096	0.286	0.294	0.286	0.294	0.381	0.384	0.281	0.283	0.283	0.712	0.714	1.141	1.103	0.681	0.695	
1st Qu.	0.215	0.214	0.338	0.326	0.211	0.215	0.211	0.215	0.712	0.714	0.712	0.714	1.141	1.103	0.681	0.695	0.695	1.036	1.173	2.263	2.318	0.959	0.997	
Median	0.300	0.335	0.590	0.602	0.294	0.305	0.294	0.305	1.036	1.173	1.036	1.173	2.263	2.318	0.959	0.997	0.997	1.370	1.462	2.555	2.639	1.249	1.314	
Mean	0.383	0.396	0.671	0.689	0.361	0.372	0.361	0.372	1.370	1.462	1.370	1.462	2.555	2.639	1.249	1.314	1.314	1.524	1.855	3.217	3.375	1.350	1.526	
3rd Qu.	0.453	0.488	0.837	0.857	0.415	0.449	0.415	0.449	1.524	1.855	1.524	1.855	3.217	3.375	1.350	1.526	1.526	5.916	5.462	9.053	9.870	4.774	4.875	
Max.	1.602	1.547	2.428	2.602	1.361	1.394	1.361	1.394	5.916	5.462	5.916	5.462	9.053	9.870	4.774	4.875	4.875							
Functional time-series method																								
Min.	0.092	0.093	0.126	0.136	0.095	0.103	0.095	0.103	0.281	0.276	0.281	0.276	0.376	0.404	0.283	0.299	0.299	0.702	0.687	1.042	1.108	0.667	0.748	
1st Qu.	0.210	0.210	0.319	0.341	0.208	0.233	0.208	0.233	0.702	0.687	0.702	0.687	1.042	1.108	0.667	0.748	0.748	0.984	0.995	2.156	2.238	0.952	1.149	
Median	0.292	0.296	0.573	0.624	0.285	0.322	0.285	0.322	0.984	0.995	0.984	0.995	2.156	2.238	0.952	1.149	1.149	1.339	1.338	2.492	2.571	1.226	1.357	
Mean	0.378	0.379	0.662	0.687	0.356	0.386	0.356	0.386	1.339	1.338	1.339	1.338	2.492	2.571	1.226	1.357	1.357	1.491	1.495	3.092	3.123	1.339	1.510	
3rd Qu.	0.441	0.444	0.823	0.887	0.406	0.435	0.406	0.435	1.491	1.495	1.491	1.495	3.092	3.123	1.339	1.510	1.510	5.679	5.470	9.207	10.172	4.630	5.504	
Max.	1.571	1.545	2.381	2.486	1.319	1.441	1.319	1.441	5.679	5.470	5.679	5.470	9.207	10.172	4.630	5.504	5.504							

TABLE 3

ONE-STEP-AHEAD INTERVAL FORECAST ERRORS BETWEEN THE DYNAMIC AND STATIC PRINCIPAL COMPONENT REGRESSION USING THE LC METHOD WITH CENTERING AND FUNCTIONAL TIME-SERIES METHOD. DUE TO LIMITATIONS OF SPACE, WE REPORT THE SUMMARY STATISTICS OF THE FORECAST ERRORS OBTAINED FROM THE 24 COUNTRIES.

Statistic	Score × 100						CPD × 100								
	Female			Male			Female			Male			Total		
	DPCA	PCA	Total	DPCA	PCA	Total	DPCA	PCA	Total	DPCA	PCA	Total	DPCA	PCA	Total
LC method with centering															
Min.	0.461	0.463	0.629	0.637	0.466	0.466	0.264	0.132	0.132	0.132	0.033	0.000	0.000	0.000	0.000
1st Qu.	1.148	1.150	1.931	1.967	1.137	1.152	0.891	1.312	1.312	1.114	1.271	1.048	1.048	1.411	1.411
Median	1.790	1.795	3.376	3.385	1.633	1.646	3.036	4.043	4.043	1.881	2.805	2.112	2.112	2.195	2.195
Mean	2.291	2.307	3.873	3.994	2.048	2.071	3.842	4.642	4.642	3.288	4.212	3.500	3.500	4.130	4.130
3rd Qu.	2.806	2.772	4.779	4.809	2.506	2.506	5.000	6.353	6.353	3.870	5.272	3.911	3.911	5.256	5.256
Max.	9.424	9.144	14.185	17.842	7.446	7.666	17.690	17.591	17.591	19.043	17.789	21.584	21.584	21.155	21.155
Functional time-series method															
Min.	0.463	0.458	0.615	0.640	0.459	0.477	0.033	1.452	1.452	0.033	0.957	0.099	0.099	0.561	0.561
1st Qu.	1.093	1.057	1.706	1.678	1.065	1.138	0.578	5.066	5.066	0.726	5.058	0.924	0.924	3.614	3.614
Median	1.813	1.887	2.982	3.043	1.529	1.631	1.716	7.112	7.112	1.650	6.518	1.650	1.650	5.875	5.875
Mean	2.163	2.182	3.545	3.542	1.871	1.940	2.389	7.241	7.241	2.452	7.218	2.749	2.749	6.624	6.624
3rd Qu.	2.547	2.592	4.352	4.294	2.258	2.328	3.127	10.132	10.132	2.442	9.125	3.507	3.507	8.804	8.804
Max.	8.738	8.745	13.720	14.607	7.046	7.620	16.601	16.040	16.040	17.558	17.063	20.066	20.066	18.779	18.779

age-specific mortality data and application to annuity pricing in actuarial science. The core techniques in the existing demographic literature use static PCA and may not incorporate temporal covariance into the eigen decomposition. When the temporal dependence is moderate or strong, the static principal components extracted from variance are no longer optimal. As an alternative, we proposed a DPCA where the principal components could be extracted from an estimated long-run covariance.

The long-run covariance encompasses the autocovariance at lag 0 (i.e., variance), as well as the autocovariance at other lags. To estimate the long-run covariance, we considered a kernel sandwich estimator used in Andrews (1991). A crucial parameter in the kernel sandwich estimator is the selection of optimal bandwidth. To determine the optimal lags, we presented a plug-in algorithm of Rice and Shang (2017) to determine the optimal bandwidth parameter in the kernel sandwich estimator.

Given that the estimation of the long-run covariance requires stationarity, we chose to work with mortality rate improvement via the forward transformation. Through using static PCA or DPCA, we modeled and forecast mortality rate improvement. Through backward transformation, we obtained the forecast mortality rate in the original scale. Using 24 mainly developed countries, we demonstrated improvement of point and interval forecast accuracies that the dynamic principal component regression entails when compared with the static analysis using the LC and functional time-series methods.

It is noteworthy that, given a sufficient number of static principal components, they can capture a similar amount of information that the dynamic approach entails. For some countries, the temporal dependency is very weak, and then the long-run covariance almost reduces to the variance alone. In that case, the static and dynamic approaches lead to the same or similar point and interval forecast accuracies. We observed that there are small differences between the static and dynamic principal component regression models for the 5-step-ahead and 10-step-ahead forecasts. For the one-step-ahead forecasts, the differences in point and interval forecast accuracies were rather apparent. Thus, it should be recommended as a valuable technique in statistical modeling and short-term forecasting.

There are a number of ways in which the current paper could be further extended, and we briefly mention three. First, a future extension could be to take a cohort perspective on the mortality improvement rate, as taken by Haberman and Renshaw (2013). Second, given that most, if not all, extrapolation methods do not perform well in long-term forecasts, it may be useful to propose a Bayesian version of dynamic functional PCA, where the prior knowledge can be incorporated. Finally, we have demonstrated the usefulness of the DPCA using the LC and functional time-series models; however, it can be applied to other mortality models that use PCA in whole or in part (see, e.g., Shang et al., 2011).

ACKNOWLEDGEMENTS

The author acknowledges a faculty research grant from the College of Business and Economics at the Australian National University. The author is grateful for comments and suggestions from the seminar participants at the United Nations Population Division.

REFERENCES

- AKAIKE, H. (1974) A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, **19**(6), 716–723.
- ANDREWS, D.W.K. (1991) Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, **59**(3), 817–858.
- BOOTH, H., HYNDMAN, R.J., TICKLE, L. and DE JONG, P. (2006) Lee-Carter mortality forecasting: A multi-country comparison of variants and extensions. *Demographic Research*, **15**, 289–310.
- BOOTH, H. and TICKLE, L. (2008) Mortality modelling and forecasting: A review of methods. *Annals of Actuarial Science*, **3**(1–2), 3–43.
- BRILLINGER, D.R. (1974) *Time Series: Data Analysis and Theory*. New York: Holt, Rinehart and Winston.
- DE BOOR, C. (2001), *A Practical Guide to Splines*, Vol. 27, Applied Mathematical Sciences, New York: Springer.
- GNEITING, T. and RAFTERY, A.E. (2007) Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association*, **102**(477), 359–378.
- HABERMAN, S. and RENSHAW, A. (2012) Parametric mortality improvement rate modelling and projecting. *Insurance: Mathematics and Economics*, **50**(3), 309–333.
- HABERMAN, S. and RENSHAW, A. (2013) Modelling and projecting mortality improvement rates using a cohort perspective. *Insurance: Mathematics and Economics*, **53**(1), 150–168.
- HORVÁTH, L. and KOKOSZKA, P. (2012) *Inference for Functional Data with Applications*. New York: Springer.
- HUMAN MORTALITY DATABASE (2019) *University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany)*. Available at www.mortality.org or www.humanmortality.de (data downloaded on 12/June/2017).
- HYNDMAN, R.J. and KHANDAKAR, Y. (2008) Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software*, **27**(3), 1–22. <http://dx.doi.org/10.18637/jss.v027.i03>.
- HYNDMAN, R.J. and SHANG, H.L. (2009) Forecasting functional time series (with discussions). *Journal of the Korean Statistical Society*, **38**(3), 199–221.
- HYNDMAN, R.J. and SHANG, H.L. (2010) Rainbow plots, bagplots, and boxplots for functional data. *Journal of Computational and Graphical Statistics*, **19**(1), 29–45.
- HYNDMAN, R.J. and ULLAH, M.S. (2007) Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics & Data Analysis*, **51**(10), 4942–4956.
- KWIATKOWSKI, D., PHILLIPS, P.C.B., SCHMIDT, P. and SHIN, Y. (1992) Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, **54**(1–3), 159–178.
- LEE, R.D. and CARTER, L.R. (1992) Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, **87**(419), 659–671.
- ORGANIZATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT [OECD] (2013) Pensions at a Glance 2013: OECD and G20 indicators, Technical report, OECD Publishing. Retrieved from http://dx.doi.org/10.1787/pension_glance-2013-en.
- RENSHAW, A.E. and HABERMAN, S. (2003) Lee-Carter mortality forecasting: A parallel generalized linear modelling approach for England and Wales mortality projections. *Journal of the Royal Statistical Society: Series C*, **52**(1), 119–137.

- RICE, G. and SHANG, H.L. (2017) A plug-in bandwidth selection procedure for long run covariance estimation with stationary functional time series. *Journal of Time Series Analysis*, **38**(4), 591–609.
- SHANG, H.L. (2012) Point and interval forecasts of age-specific life expectancies: A model averaging approach. *Demographic Research*, **27**, 593–644.
- SHANG, H.L. (2019) Visualizing rate of change: An application to age-specific fertility rates. *Journal of the Royal Statistical Society: Series A*, **182**(1), 249–262.
- SHANG, H.L., BOOTH, H. and HYNDMAN, R.J. (2011) Point and interval forecasts of mortality rates and life expectancy: A comparison of ten principal component methods. *Demographic Research*, **25**, 173–214.
- SHANG, H.L. and HABERMAN, S. (2017) Grouped multivariate and functional time series forecasting: An application to annuity pricing. *Insurance: Mathematics and Economics*, **75**, 166–179.
- ZHANG, X. and WANG, J.-L. (2016) From sparse to dense functional data and beyond. *The Annals of Statistics*, **44**(5), 2281–2321.

HAN LIN SHANG (Corresponding author)
*Research School of Finance, Actuarial Studies and Statistics, Level 4,
 Building 26C
 Australian National University
 Kingsley Street, Acton, Canberra
 ACT 2601, Australia
 E-Mail: hanlin.shang@anu.edu.au.*

APPENDIX A

A.1. Estimation of long-run covariance

Under the asymptotic mean squared normed error, Rice and Shang (2017) show that the optimal bandwidth parameter h_{opt} has the following forms:

$$h_{\text{opt}} = c_0 n^{\frac{1}{1+2q}},$$

$$c_0 = \left(2q \|C^{(q)}\|^2\right)^{\frac{1}{1+2q}} \left\{ \left[\|C\|^2 + \left(\int_0^1 C(u, u) du \right)^2 \right] \int_{-\infty}^{\infty} W_q^2(x) dx \right\}^{-\frac{1}{1+2q}}, \quad (\text{A.1})$$

where q denotes the order of derivative, and $W_q(x)$ denotes a kernel (weight) function of order q . The crux of the problem is that the quantities involving $C^{(q)}$ and C in Equation (A.1) are unknown, and we use a plug-in algorithm to estimate them, from which we obtain \hat{c}_0 and \hat{h}_{opt} .

The plug-in bandwidth selection method is given as follows:

- (1) Compute pilot estimates of $C^{(p)}$, for $p = 0, q$:

$$\widehat{C}_{h_1, q_1}^{(p)}(x, u) = \sum_{\ell=-\infty}^{\infty} W_{q_1} \left(\frac{\ell}{h_1} \right) |\ell|^p \widehat{\gamma}_\ell(x, u),$$

which utilizes an initial bandwidth choice $h_1 = h_1(n)$, and an initial kernel function W_{q_1} of order q_1 .

- (2) Estimate c_0 by

$$\widehat{c}_0(h_1, q_1) = \left(2q \left\| \widehat{C}_{h_1, q_1}^{(q)} \right\|^2 \right)^{\frac{1}{1+2q}} \left\{ \left[\|C\|^2 + \left(\int_0^1 C(u, u) du \right)^2 \right] \int_{-\infty}^{\infty} W_q^2(x) dx \right\}^{-\frac{1}{1+2q}},$$

where W_q denotes a final kernel function of order q .

- (3) Use the bandwidth

$$\widehat{h}_{\text{opt}}(h_1, q_1) = \widehat{c}_0(h_1, q_1) n^{\frac{1}{1+2q}}$$

in the definition of $\widehat{C}_{h, q}$ in Equation (5) in the paper.

In terms of the initial and final weight functions, Rice and Shang (2017) advocated the use of a flat-top weight function W_∞ for the initial kernel function, and a Bartlett kernel function as the final kernel function. A flat-top weight function W_∞ is of the form

$$W_\infty(t) = \begin{cases} 1 & 0 \leq |t| < k_1; \\ \frac{k_2 - |t|}{k_2 - k_1} & k_1 \leq |t| \leq k_2; \\ 0 & |t| \geq k_2; \end{cases}$$

where $k_2 > k_1$. Let us take $k_2 = 1$ and $k_1 = 0.5$. The Bartlett weight function W_1 is of the form

$$W_1 = \begin{cases} 1 - |x| & \text{for } |x| \leq 1; \\ 0 & \text{otherwise.} \end{cases}$$

APPENDIX B

In Haberman and Renshaw (2012), the LC method does not include the mean term, since their Generalized Linear Model approach uses an iterative fitting algorithm to estimate parameters by minimizing a deviance criterion. In Figure B1, we present the point forecast results for the 1-step-ahead, 5-step-ahead, and 10-step-ahead forecasts. The results are similar to the LC method with the mean term given in Table 2 of the paper and results given in Appendix C.

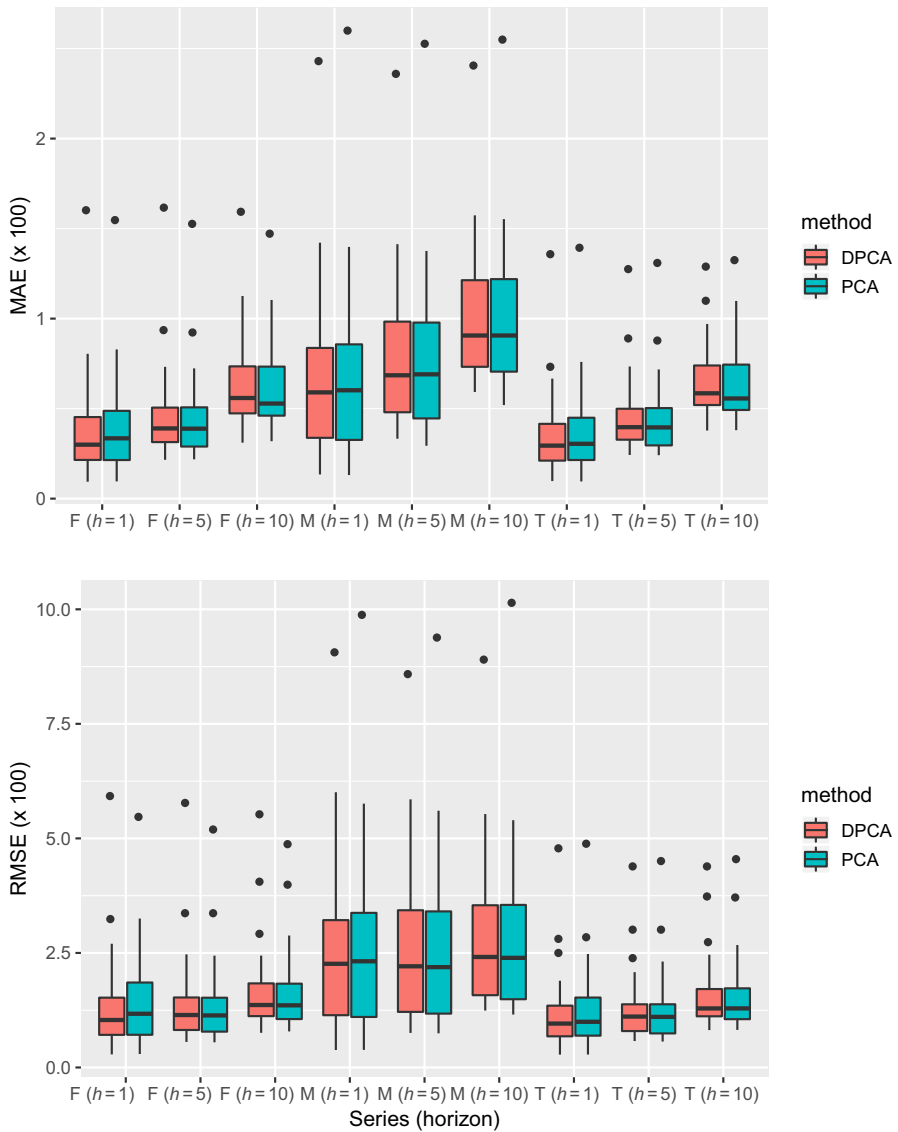


FIGURE B1: Comparison of 1-step-ahead, 5-step-ahead, and 10-step-ahead point forecast errors between the dynamic and static principal component regression using the LC method without centering (i.e., Haberman and Renshaw's (2012) method). h denotes forecast horizon.

APPENDIX C

Using the LC and functional time-series methods, these longer-horizon forecast results are reported in Figure C1 for $h = 5$. For $h = 10$, these results are reported in Figure C2. As measured by MAE, RMSE, and interval score, there are marginal difference between the DPCA and static PCA, although the maximum forecast error of the DPCA is often smaller than that of the static PCA. As measured by CPD, DPCA often outperforms the static PCA.

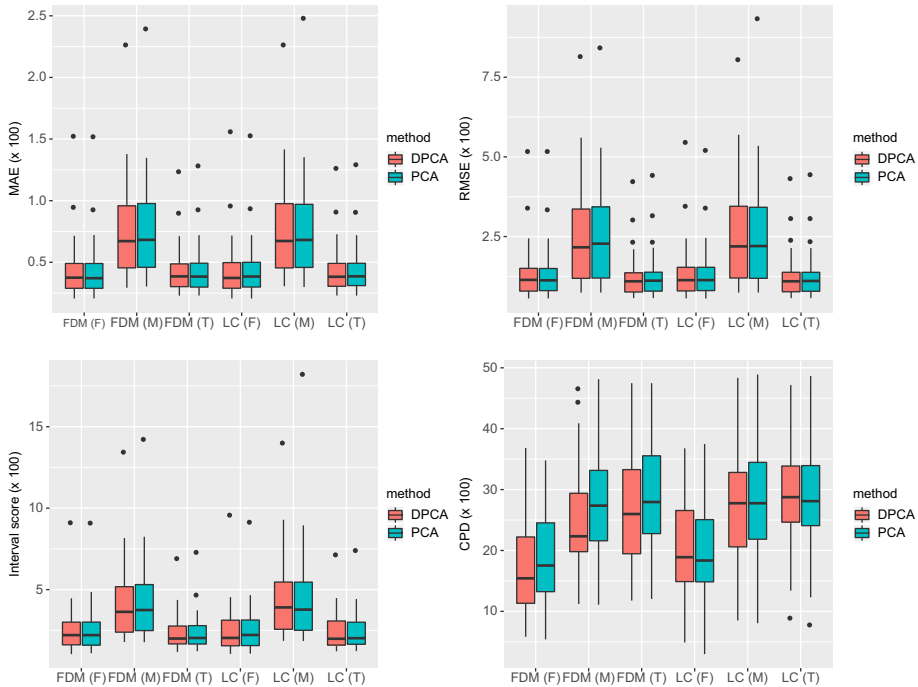


FIGURE C1: Comparison of five-step-ahead point and interval forecast errors between the dynamic and static principal component regression using the functional time-series method and LC method with centering.

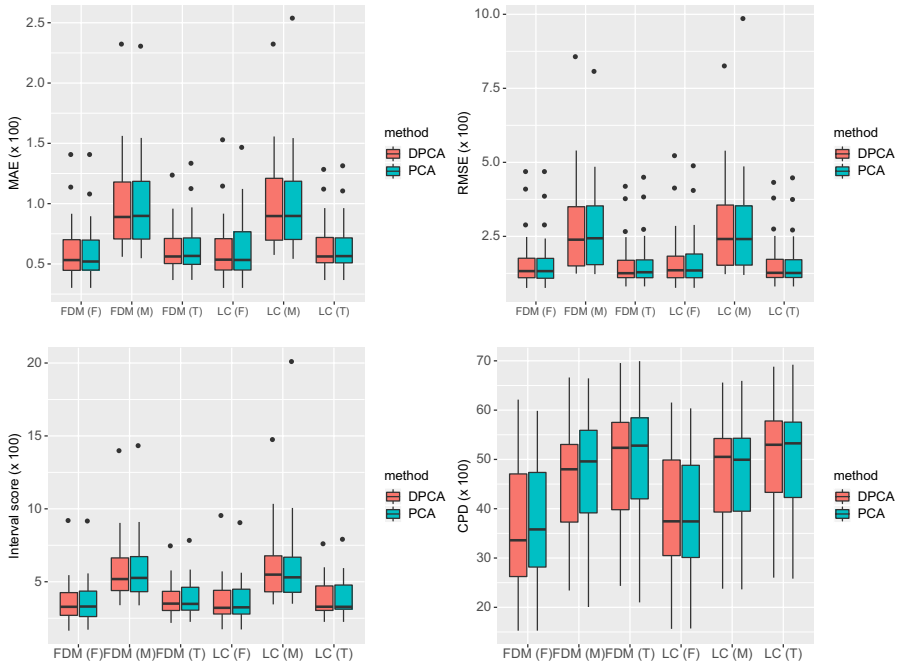


FIGURE C2: Comparison of 10-step-ahead point and interval forecast errors between the dynamic and static principal component regression using the functional time-series method and LC method with centering.

