University of Groningen

Forecast covariances in the linear multiregression dynamic model

Queen, Catriona M.; Wright, Ben J.; Albers, Casper J.

Published in:
Journal of Forecasting

DOI:
[10.1002/for.1050](https://doi.org/10.1002/for.1050)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2008

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Queen, C. M., Wright, B. J., & Albers, C. J. (2008). Forecast covariances in the linear multiregression dynamic model. *Journal of Forecasting*, 27(2), 175-191. <https://doi.org/10.1002/for.1050>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Forecast Covariances in the Linear Multiregression Dynamic Model

CATRIONA M. QUEEN,* BEN J. WRIGHT AND
CASPER J. ALBERS

The Open University, Milton Keynes, UK

ABSTRACT

The linear multiregression dynamic model (LMDM) is a Bayesian dynamic model which preserves any conditional independence and causal structure across a multivariate time series. The conditional independence structure is used to model the multivariate series by separate (conditional) univariate dynamic linear models, where each series has contemporaneous variables as regressors in its model. Calculating the forecast covariance matrix (which is required for calculating forecast variances in the LMDM) is not always straightforward in its current formulation. In this paper we introduce a simple algebraic form for calculating LMDM forecast covariances. Calculation of the covariance between model regression components can also be useful and we shall present a simple algebraic method for calculating these component covariances. In the LMDM formulation, certain pairs of series are constrained to have zero forecast covariance. We shall also introduce a possible method to relax this restriction. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS multivariate time series; dynamic linear model; conditional independence; forecast covariance matrix; component covariances

INTRODUCTION

A multiregression dynamic model (MDM) (Queen and Smith, 1993) is a non-Gaussian multivariate state-space time series model. The MDM has been designed for forecasting time series for which, at any one time period, a conditional independence structure across the time series and causal drive through the system can be hypothesized. Such series can be represented by a directed acyclic graph (DAG) which, in addition to giving a useful pictorial representation of the structure of the multivariate series, is used to decompose the multivariate time series model into simpler (conditional) univariate dynamic models (West and Harrison, 1997). As such, the MDM is a graphical model (see, for example, Cowell *et al.*, 1999) for multivariate time series, which simplifies computation in the model through local computation. Like simultaneous equation models, in the MDM each univariate series has contemporaneous variables as regressors in its model. When these regressions are linear, we have a linear MDM (LMDM) and in this case we have a set of univariate regression dynamic linear

* Correspondence to: Catriona Queen, Department of Statistics, The Open University, Milton Keynes, MK7 6AA, UK.
E-mail: C.Queen@open.ac.uk

models (DLMs) (West and Harrison, 1997, Section 9.2). Like a non time series graphical model, the conditional univariate models in an MDM are computationally simple to work with (in this case univariate DLMs), while the joint distribution can be arbitrarily complex. Also, because the model is a collection of univariate dynamic models, the MDM avoids the difficult problem of eliciting an observation covariance matrix for the multivariate problem. It also avoids the alternative of learning about the covariance matrix online via the matrix normal DLM (West and Harrison, 1997, Section 16.4) whose use is restricted to problems in which all the individual series are similar.

Many multivariate time series may be suitable for use with an LMDM. For example, Queen (1994) uses the LMDM to model monthly brand sales in a competitive market. In this application, the competition in the market is the causal drive within the system and is used to define a conditional independence structure across the time series. Queen (1997) and Queen *et al.* (1994) focus on how DAGs may be elicited for market models. DAGs are also used to represent brand relationships when forecasting time series of brand sales in Goldstein *et al.* (1993), Farrow *et al.* (1997) and Farrow (2003). As another example, Whitlock and Queen (2000) and Queen *et al.* (2007) use the LMDM to model hourly vehicle counts at various points in a traffic network. Here, as in Sun *et al.* (2006), the direction of traffic flow produces the causal drive in the system and the possible routes through the network are used to define a conditional independence structure across the time series. The LMDM then produces a set of regression DLMs where contemporaneous traffic flows at upstream links in the network are used as linear regressors. Tebaldi *et al.* (2002) also use regression DLMs when modelling traffic flows. Like Queen *et al.* (2007), they use traffic flows at upstream links in the network as linear regressors. However, whereas the vehicle counts in Queen *et al.* (2007) are for 1-hour intervals, those in Tebaldi *et al.* (2002) are for 1-minute intervals, and so regression on lagged flows (rather than contemporaneous flows) is required. Fosen *et al.* (2006) use similar ideas to the LMDM when proposing a dynamical graphical model combining the additive hazard model and classical path analysis to analyse a trial of cancer patients with liver cirrhosis. Like the LMDM, their model uses linear regression with parents from the DAG as regressors. The focus, however, is to investigate the treatment effects on the variates over time, rather than forecasts of the variables directly.

There are many other potential application areas for the LMDM, including problems in economics (modelling various economic indicators such as energy consumption and GDP), environmental problems (such as water and other resource management problems), industrial problems (such as product distribution flow problems) and medical problems (such as patient physiological monitoring).

When using the LMDM it is important to be able to calculate the forecast covariance matrix for the series. Not only is this of interest in its own right, but it is also required for calculating the forecast variances for the individual series. Queen and Smith (1993) presented a recursive form for the covariance matrix, but this is not in an algebraic form which is always simple to use in practice. In this paper we introduce a simple algebraic method for calculating forecast covariances in the LMDM.

Following the superposition principle (see West and Harrison, 1997, p. 98), each conditional univariate DLM in an LMDM can be thought of as the sum of individual model components. For example, a DLM may have a trend component, a regression component, and so on. In an LMDM it can be useful to calculate the covariance between individual DLM regression components for different series and here we introduce a simple algebraic method for calculating such component covariances.

The paper is structured as follows. In the next section the LMDM is described and its use is illustrated using a bivariate time series of brand sales. The third section presents a simple algebraic method for calculating the LMDM one-step- and k -step-ahead forecast covariance matrices.

Component covariances are introduced in the fourth section, along with a simple method to calculate them. In the LMDM formulation, certain pairs of series are constrained to have zero forecast covariance and the fifth section introduces a possible method to relax this restriction. Finally, the sixth section gives some concluding remarks.

THE LINEAR MULTIREGRESSION DYNAMIC MODEL

We have a multivariate time series $\mathbf{Y}_t = (Y_t(1), \dots, Y_t(n))^T$. Suppose that the series is ordered and that the same conditional independence and causal structure are defined across the series through time, so that at each time $t = 1, 2, \dots$, we have

$$Y_t(r) \perp\!\!\!\perp \{\{Y_t(1), \dots, Y_t(r-1)\} \setminus \text{pa}(Y_t(r))\} \mid \text{pa}(Y_t(r)) \quad \text{for } r = 2, \dots, n$$

which reads ‘ $Y_t(r)$ is independent of $\{Y_t(1), \dots, Y_t(r-1)\} \setminus \text{pa}(Y_t(r))$ given $\text{pa}(Y_t(r))$ ’ (Dawid, 1979), where the notation ‘ $\setminus$ ’ reads ‘excluding’ and $\text{pa}(Y_t(r)) \subseteq \{Y_t(1), \dots, Y_t(r-1)\}$. Each variable in the set $\text{pa}(Y_t(r))$ is called a *parent* of $Y_t(r)$ and $Y_t(r)$ is known as a *child* of each variable in the set $\text{pa}(Y_t(r))$. We shall call any series for which $\text{pa}(Y_t(\cdot)) = \emptyset$ a *root node* and list all root nodes before any children in the ordered series $\mathbf{Y}_t$.

The conditional independence relationships at each time point t can be represented by a DAG, where there is a directed arc to $Y_t(r)$ from each of its parents in $\text{pa}(Y_t(r))$. To illustrate, Figure 1 shows a DAG for five time series at time t , where $\text{pa}(Y_t(2)) = \emptyset$, $\text{pa}(Y_t(3)) = \{Y_t(1), Y_t(2)\}$, $\text{pa}(Y_t(4)) = \{Y_t(3)\}$ and $\text{pa}(Y_t(5)) = \{Y_t(3), Y_t(4)\}$. Note that both $Y_t(1)$ and $Y_t(2)$ are root nodes.

Suppose further that there is a conditional independence and causal structure defined for the processes so that, if $\mathbf{Y}^t(r) = (Y_1(r), \dots, Y_t(r))^T$:

$$Y_t(r) \perp\!\!\!\perp \{\{\mathbf{Y}^t(1), \dots, \mathbf{Y}^t(r-1)\} \setminus \text{pa}(\mathbf{Y}^t(r))\} \mid \text{pa}(\mathbf{Y}^t(r)), \mathbf{Y}^{t-1}(r) \quad \text{for } r = 2, \dots, n$$

Denote the information available at time t by D_t . An LMDM has the following system equation and n observation equations for all times $t = 1, 2, \dots$:

Observation equations:

$$Y_t(r) = \mathbf{F}_t(r)^T \boldsymbol{\theta}_t(r) + v_t(r), \quad v_t(r) \sim N(0, V_t(r)) \quad 1 \leq r \leq n$$

System equation:

$$\boldsymbol{\theta}_t = \mathbf{G}_t \boldsymbol{\theta}_{t-1} + \mathbf{w}_t, \quad \mathbf{w}_t \sim N(\mathbf{0}, \mathbf{W}_t)$$

Initial information:

$$(\boldsymbol{\theta}_0 | D_0) \sim N(\mathbf{m}_0, \mathbf{C}_0)$$

Figure 1. DAG representing five time series at time t

The vector $F_i(r)$ contains the parents $\text{pa}(Y_i(r))$ and possibly other known variables (which may include $Y^{r-1}(r)$ and $\text{pa}(Y^{r-1}(r))$); $\theta_i(r)$ is the parameter vector for $Y_i(r)$; $V_i(1), \dots, V_i(n)$ are the scalar observation variances; $\theta_i^T = (\theta_i(1)^T, \dots, \theta_i(n)^T)$; and the matrices G_r , W_i and C_0 are all block diagonal. The error vectors, $v_i^T = \{v_i(1), \dots, v_i(n)\}$ and $w_i^T = \{w_i(1)^T, \dots, w_i(n)^T\}$, are such that $v_i(1), \dots, v_i(n)$ and $w_i(1), \dots, w_i(n)$ are mutually independent and $\{v_i, w_i\}_{i \geq 1}$ are mutually independent with time.

The LMDM therefore uses the conditional independence structure to model the multivariate time series by n separate univariate models—for $Y_i(1)$ and $Y_i(r)|\text{pa}(Y_i(r))$, $r = 2, \dots, n$. For those series with parents, each univariate model is simply a regression DLM with the parents as linear regressors. For root nodes, any suitable univariate DLM may be used. For example, consider the DAG in Figure 1. As $Y_i(1)$ and $Y_i(2)$ are both root nodes, each of these series can be modelled separately in an LMDM using any suitable univariate DLMs. $Y_i(3)$, $Y_i(4)$ and $Y_i(5)$ all have parents and so these would be modelled by (separate) univariate regression DLMs with the two regressors $Y_i(1)$ and $Y_i(2)$ for $Y_i(3)$'s model, the single regressor $Y_i(3)$ for $Y_i(4)$'s model and the two regressors $Y_i(3)$ and $Y_i(4)$ for $Y_i(5)$'s model.

As long as $\theta_i(1), \theta_i(2), \dots, \theta_i(n)$ are mutually independent a priori, each $\theta_i(r)$ can be updated separately in closed form from $Y_i(r)$'s (conditional) univariate model. A forecast for $Y_i(1)$ and the conditional forecasts for $Y_i(r)|\text{pa}(Y_i(r))$, $r = 2, \dots, n$, can also be found separately using established DLM results (see West and Harrison, 1997, for details). For example, in the context of the DAG in Figure 1, forecasts can be found separately (using established DLM results) for

$$Y_i(1), Y_i(2), Y_i(3)|(Y_i(1), Y_i(2)), Y_i(4)|Y_i(3) \text{ and } Y_i(5)|(Y_i(3), Y_i(4))$$

However, $Y_i(r)$ and $\text{pa}(Y_i(r))$ are observed simultaneously. So the marginal forecast for each $Y_i(r)$, without conditioning on the values of its parents, is required. Unfortunately, the marginal forecast distributions for $Y_i(r)$, $r = 2, \dots, n$, will not generally be of a simple form. However (under quadratic loss), the mean and variance of the marginal forecast distributions for $Y_i(r)$, $r = 2, \dots, n$, are adequate for forecasting purposes, and these can be calculated. So, returning to the context of the DAG in Figure 1, this means that the marginal forecast means and variances for $Y_i(3)$, $Y_i(4)$ and $Y_i(5)$ need to be calculated. Note that as $Y_i(1)$ and $Y_i(2)$ do not have any parents, their marginal forecasts have already been calculated from DLM theory.

For calculating the marginal forecast variance for $Y_i(r)$, the marginal forecast covariance matrix for $\text{pa}(Y_i(r))$ is required. Returning to the example DAG in Figure 1 to illustrate, this means that the marginal forecast variance of $Y_i(5)$, for example, requires the forecast covariance for $Y_i(3)$ and $Y_i(4)$, without conditioning on either of their parents $Y_i(1)$ and $Y_i(2)$ —that is, the marginal covariance of $Y_i(3)$ and $Y_i(4)$. The marginal forecast covariance matrix also provides information about the structure of the multivariate series and, as such, is of interest in its own right. Queen and Smith (1993) gave a recursive form for calculating the marginal forecast covariance matrix. However, this is not always easy to use in practice. The third section presents a simple algebraic form for calculating marginal forecast covariances. First, we shall illustrate how the LMDM works in practice.

Example: forecasting a bivariate time series of brand sales

We shall illustrate using a simple LMDM for a bivariate series of monthly brand sales. The series consist of 34 months of data (supplied by Unilever Research) for two brands, which we shall call B1 and B2, who compete with each other in a product market.

Denote the sales at each time t for brands B1 and B2 by $X_i(1)$ and $X_i(2)$, respectively. Let Y_i be the total number of sales of B1 and B2 in month t (so that $Y_i = X_i(1) + X_i(2)$). Then Y_i can be

modelled by a Poisson distribution with some mean μ_t , denoted $Po(\mu_t)$. Suppose that for each individual purchase of this product in month t

$$P(\text{B1 purchased} | \text{B1 or B2 purchased}) = \theta_t$$

Then, $X_t(1)$, the number of B1 purchased in month t , can be modelled by a binomial distribution with sample size Y_t , the total number of B1 and B2 purchased in month t , and parameter θ_t —that is, $X_t(1) | Y_t \sim Bi(Y_t, \theta_t)$.

These distributions for Y_t and $X_t(1)$ can be represented by the DAG in Figure 2. $X_t(2)$ is a logical function of its parents and is known once its parents are known. Following the terminology of WinBUGS software (<http://www.mrc-bsu.cam.ac.uk/bugs/>) we shall call this a logical variable and denote it on the DAG by a double oval. Note that if we had defined a conditional binomial model for $X_t(2) | Y_t$ instead, $X_t(1)$ would have been the logical variable.

Approximating the Poisson distribution for Y_t and the conditional binomial distribution for $X_t(1) | Y_t$ to normality, the observation equations in an LMDM for the DAG in Figure 2 are of the following form:

$$\begin{aligned} Y_t &= \mu_t + v_t(y), \quad v_t(y) \sim N(0, V_t(y)) \\ X_t(1) &= Y_t \theta_t + v_t(1), \quad v_t(1) \sim N(0, V_t(1)) \end{aligned}$$

and

$$X_t(2) = Y_t - X_t(1)$$

The exact form of the observation equation for Y_t can be far more complicated to account for trend, seasonality and so on—whatever its form, its *mean* at time t is μ_t . The time series plot of Y_t suggests that in fact a linear growth model is appropriate, and this has been implemented here. A plot of $X_t(1)$ against Y_t is given in Figure 3. For these data, using Y_t as a linear regressor for $X_t(1)$ clearly seems sensible. If there was a non-linear relationship between Y_t and $X_t(1)$, then Y_t and/or $X_t(1)$ could be transformed so that the linear regression model implied by the LMDM is appropriate, or the more general MDM could be used.

For simplicity, in this illustration the observation variances, $V_t(y)$ and $V_t(1)$, were fixed throughout and were simply estimated by fitting simple linear regression models for the 34 observations of Y_t and $X_t(1)$ (with regressor Y_t), respectively. Discount factors of 0.8 and 0.71 were used for Y_t and $X_t(1)$'s models, respectively, and were chosen so as to minimize the mean squared error (MSE) and the mean absolute deviation (MAD). The first two observations for each series were used to

Figure 2. DAG representing monthly sales of brands B1 and B2 ($X_t(1)$ and $X_t(2)$, respectively), and their sum (Y_t)

Figure 3. Plot of $X_t(1)$ against Y_t

calculate initial values for the prior means for the parameters, while their initial prior variances were each set to be large (10,800 for the two parameters for Y_t and 1 for the parameter for $X_t(1)$) to allow forecasts to adapt quickly.

To get an idea of how well the LMDM is performing for these series, a standard multivariate DLM (MV DLM) was also used, using linear growth models for both $X_t(1)$ and $X_t(2)$. In order to try to get a fair comparison, the observation covariance matrix was assumed fixed and was calculated in exactly the same way as the observation variances were for the LMDM. As with the LMDM, initial estimates of prior means for the parameters were calculated using the first two observations of the series, the associated prior variances were set to be large to allow quick adaptation in the model and discount factors were chosen (0.85 for $X_t(1)$, 0.70 for $X_t(2)$) so as to minimize the MSE/MAD.

Figure 4 shows plots of the time series for $X_t(1)$ and $X_t(2)$, together with the (marginal) one-step forecasts and ± 1.96 (marginal) forecast standard deviation error bars calculated using both the MV DLM and the LMDM. From these plots, it can be seen that the LMDM is performing better than the MV DLM. This is also clearly reflected in the MSE/MAD values for the one-step-ahead forecasts for the two models, shown in Table I. Also shown in Table I are the MSE/MAD values for the two- and three-step-ahead forecasts using both models. Clearly, the LMDM is also performing better with respect to these k -step-ahead forecasts.

In addition to giving a better forecast performance, the LMDM also has the advantage that intervention can be simpler to implement. For example, suppose that total sales of the brands were expected to increase suddenly. Then to accommodate this information into the model, only interven-

Figure 4. Plots of the one-step forecasts (joined dots) and ± 1.96 forecast standard deviation error bars (dotted lines) for $X_t(1)$ and $X_t(2)$ calculated using the LMDM, together with the data (filled dots)

tion for Y_t would be required in the LMDM, whereas the MV DLM would require intervention for both $X_t(1)$ and $X_t(2)$. The LMDM also has the advantage that only observation variances need to be elicited, whereas the MV DLM has the additional difficult task of eliciting the observation covariance between $X_t(1)$ and $X_t(2)$ as well.

Table I. MSE and MAD values for the MV DLM and LMDM

		MSE		MAD	
		MV DLM	LMDM	MV DLM	LMDM
One-step forecasts	$X_t(1)$	3541.72	1601.39	26.89	23.26
	$X_t(2)$	4622.96	3988.42	49.31	40.77
Two-step forecasts	$X_t(1)$	4579.25	1513.81	27.71	23.56
	$X_t(2)$	5486.01	3998.98	51.96	37.86
Three-step forecasts	$X_t(1)$	3854.69	1509.43	39.54	22.51
	$X_t(2)$	7198.47	4145.12	68.43	39.56

SIMPLE CALCULATION OF MARGINAL FORECAST COVARIANCES

From Queen and Smith (1993), the marginal forecast covariance between $Y_t(i)$ and $Y_t(r)$, $i < r$, $r = 2, \dots, n$, can be calculated recursively using

$$\text{cov}(Y_t(i), Y_t(r)|D_{t-1}) = E(Y_t(i) \cdot E(Y_t(r)|Y_t(1), \dots, Y_t(r-1), D_{t-1})|D_{t-1}) - E(Y_t(i)|D_{t-1})E(Y_t(r)|D_{t-1})) \quad (1)$$

In this paper we shall use this to derive a simple algebraic form for calculating these forecast covariances. In what follows let $\mathbf{a}_t(r)$ be the prior mean for $\boldsymbol{\theta}_t(r)$.

Theorem 1 In an LMDM, let $Y_t(j_1), \dots, Y_t(j_{m_r})$ be the m_r parents of $Y_t(r)$. Then for $i < r$, the one-step-ahead forecast covariance between $Y_t(i)$ and $Y_t(r)$ can be calculated using

$$\text{cov}(Y_t(i), Y_t(r)|D_{t-1}) = \sum_{j=1}^{m_r} \text{cov}(Y_t(i), Y_t(j_t)|D_{t-1}) a_t(r)^{(j_t)}$$

where $a_t(r)^{(j_t)}$ is the element of $\mathbf{a}_t(r)$ associated with the parent regressor $Y_t(j_t)$ —i.e., $a_t(r)^{(j_t)}$ is the prior mean for the parameter for regressor $Y_t(j_t)$.

Proof Consider equation (1). From the observation equations for the LMDM

$$E(Y_t(r)|Y_t(1), \dots, Y_t(r-1), D_{t-1}) = \mathbf{F}_t(r)^T \mathbf{a}_t(r)$$

Also, using the result that for two random variables X and Y , $E(Y) = E(E(Y|X))$

$$E(Y_t(r)|D_{t-1}) = E(\mathbf{F}_t(r)^T \mathbf{a}_t(r)|D_{t-1}) = E(\mathbf{F}_t(r)^T|D_{t-1}) \mathbf{a}_t(r)$$

So

$$\begin{aligned} \text{cov}(Y_t(i), Y_t(r)|D_{t-1}) &= E(Y_t(i) \cdot \mathbf{F}_t(r)^T|D_{t-1}) \mathbf{a}_t(r) - E(Y_t(i)|D_{t-1}) E(\mathbf{F}_t(r)^T|D_{t-1}) \mathbf{a}_t(r) \\ &= \text{cov}(Y_t(i), \mathbf{F}_t(r)^T|D_{t-1}) \mathbf{a}_t(r) \end{aligned}$$

Now $Y_t(r)$ has the m_r parents $Y_t(j_1), \dots, Y_t(j_{m_r})$, so

$$\mathbf{F}_t(r)^T = (Y_t(j_1) \dots Y_t(j_{m_r}) \mathbf{x}_t(r)^T)$$

where $\mathbf{x}_i(r)^T$ is a vector of known variables. Then, $\text{cov}(Y_i(i), \mathbf{x}_i(r)^T | D_{t-1})$ is simply a vector of zeros and so

$$\text{cov}(Y_i(i), Y_i(r) | D_{t-1}) = \sum_{l=1}^{m_r} \text{cov}(Y_i(i), Y_i(j_l) | D_{t-1}) a_i(r)^{(j_l)}$$

as required. $\square$

The marginal forecast covariance between $Y_i(i)$ and $Y_i(r)$ is therefore simply the sum of the covariances between $Y_i(i)$ and each of the parents of $Y_i(r)$. Consequently it is simple to calculate the one-step forecast covariance matrix recursively.

Corollary 1 *In an LMDM, let $Y_i(j_1), \dots, Y_i(j_{m_r})$ be the m_r parents of $Y_i(r)$. Suppose that the series have been observed up to, and including, time t . Define*

$$\mathbf{a}_i(r, k) = E(\boldsymbol{\theta}_{i+k}(r) | D_t)$$

so that

$$\mathbf{a}_i(r, k) = G_{i+k}(r) \mathbf{a}_i(r, k-1)$$

with $\mathbf{a}_i(r, 0) = \mathbf{m}_i(r)$, the posterior mean for $\boldsymbol{\theta}_i(r)$ at time t . Denote the parameter associated with parent regressor $Y_i(j_l)$ by $\theta_i(r)^{(j_l)}$, and $E(\theta_i(r)^{(j_l)} | D_t)$ by $a_i(r, k)^{(j_l)}$, the associated element of $\mathbf{a}_i(r, k)$.

Then, for $i < r$ and $k \geq 1$, the k -step-ahead forecast covariance between $Y_{i+k}(i)$ and $Y_{i+k}(r)$ can be calculating recursively using

$$\text{cov}(Y_{i+k}(i), Y_{i+k}(r) | D_t) = \sum_{l=1}^{m_r} \text{cov}(Y_{i+k}(i), Y_{i+k}(j_l) | D_t) a_i(r, k)^{(j_l)}$$

Proof Let $\mathbf{X}_{i+k}(r)^T = (Y_{i+k}(1) \dots Y_{i+k}(r-1))$. Using the result that $E(XY) = E[X \cdot E(Y|X)]$

$$\begin{aligned} \text{cov}(\mathbf{X}_{i+k}(r), Y_{i+k}(r) | D_t) &= E(\mathbf{X}_{i+k}(r) \cdot E(Y_{i+k}(r) | Y_{i+k}(1), \dots, Y_{i+k}(r-1), D_t) | D_t) \\ &\quad - E(\mathbf{X}_{i+k}(r) | D_t) E(Y_{i+k}(r) | D_t) \end{aligned} \quad (2)$$

From DLM theory (see, for example, West and Harrison, 1997, pp. 106–107)

$$E(Y_{i+k}(r) | Y_{i+k}(1), \dots, Y_{i+k}(r-1), D_t) = \mathbf{F}_{i+k}(r)^T \mathbf{a}_i(r, k) \quad (3)$$

where

$$\mathbf{a}_i(r, k) = G_{i+k}(r) \mathbf{a}_i(r, k-1)$$

with $\mathbf{a}_i(r, 0) = \mathbf{m}_i(r)$, the posterior mean for $\boldsymbol{\theta}_i(r)$ at time t . Also

$$\begin{aligned} E(Y_{i+k}(r) | D_t) &= E[E(Y_{i+k}(r) | Y_{i+k}(1), \dots, Y_{i+k}(r-1), D_t) | D_t] \\ &= E[\mathbf{F}_{i+k}(r)^T \mathbf{a}_i(r, k) | D_t] \end{aligned}$$

by equation (3).

Thus Equation (2) becomes

$$\text{cov}(X_{t+k}(r), Y_{t+k}(r)|D_t) = \text{cov}(X_{t+k}(r), F_{t+k}(r)^T|D_t)a_t(r, k)$$

and so the i th row of this gives us

$$\text{cov}(Y_{t+k}(i), Y_{t+k}(r)|D_t) = \sum_{l=1}^{m_r} \text{cov}(Y_{t+k}(i), Y_{t+k}(jl)|D_t)a_t(r, k)^{(jl)}$$

as required. $\square$

Corollary 2 For two root nodes $Y_i(i)$ and $Y_i(r)$, under the LMDM

$$\text{cov}(Y_i(i), Y_i(r)|D_{t-1}) = 0$$

Proof From the proof of Theorem 1

$$\text{cov}(Y_i(i), Y_i(r)|D_{t-1}) = \text{cov}(Y_i(i), F_i(r)^T|D_{t-1})a_i(r)$$

Since $Y_i(r)$ is a root node, $F_i(r)$ only contains known variables, so that

$$\text{cov}(Y_i(i), F_i(r)^T|D_{t-1}) = 0$$

The result then follows directly. $\square$

To illustrate calculating a one-step forecast covariance using Theorem 1 and Corollary 2, consider the following example.

Example 1

Consider the DAG in Figure 1. For $r = 2, \dots, 5$ and $j = 1, \dots, 4$, let $a_t(r)^{(j)}$ be the prior mean for parent regressor $Y_i(j)$ in $Y_i(r)$'s model. The forecast covariance between $Y_i(1)$ and $Y_i(5)$, for example, is then calculated as follows:

$$\text{cov}(Y_i(1), Y_i(5)|D_{t-1}) = \text{cov}(Y_i(1), Y_i(4)|D_{t-1})a_t(5)^{(4)} + \text{cov}(Y_i(1), Y_i(3)|D_{t-1})a_t(5)^{(3)}$$

$$\text{cov}(Y_i(1), Y_i(4)|D_{t-1}) = \text{cov}(Y_i(1), Y_i(3)|D_{t-1})a_t(4)^{(3)}$$

$$\text{cov}(Y_i(1), Y_i(3)|D_{t-1}) = \text{cov}(Y_i(1), Y_i(1)|D_{t-1})a_t(3)^{(1)} + \text{cov}(Y_i(1), Y_i(2)|D_{t-1})a_t(3)^{(2)}$$

$Y_i(1)$ and $Y_i(2)$ are both root nodes and so their covariance is 0. Thus

$$\text{cov}(Y_i(1), Y_i(5)|D_{t-1}) = \text{var}(Y_i(1)|D_{t-1})a_t(3)^{(1)}(a_t(4)^{(3)}a_t(5)^{(4)} + a_t(5)^{(3)})$$

where $\text{var}(Y_i(1)|D_{t-1})$ is the marginal forecast variance for $Y_i(1)$. This is both simple and fast to calculate.

When a DAG has a logical node, the logical node will be some (logical) function of its parents. If the function is linear, then the covariance between $Y_i(i)$ and a logical node $Y_i(r)$ will simply be a linear function of the covariance between $Y_i(i)$ and the parents of $Y_i(r)$. For example, for the DAG in Figure 5, suppose that $Y_i(4) = Y_i(3) - Y_i(2)$. Then

Figure 5. DAG with a logical node $Y_t(4)$

$$\text{cov}(Y_t(1), Y_t(4)|D_{t-1}) = \text{cov}(Y_t(1), Y_t(3)|D_{t-1}) - \text{cov}(Y_t(1), Y_t(2)|D_{t-1})$$

The covariance can then be found simply by applying Theorem 1.

COMPONENT COVARIANCES

Suppose that $Y_t(r)$ has the m_r parents $Y_t(r'_1), \dots, Y_t(r'_m)$ for $r = 2, \dots, n$. Write the observation equation for each $Y_t(r)$ as the sum of regression components as follows:

$$Y_t(r) = \sum_{i=1}^{m_r} Y_t(r, r'_i) + Y_t(r, \mathbf{x}_t(r)) + v_t(r), \quad v_t(r) \sim N(0, V_t(r)) \quad (4)$$

with

$$Y_t(r, r'_i) = Y_t(r'_i) \theta_t(r)^{(r'_i)}$$

$$Y_t(r, \mathbf{x}_t(r)) = \mathbf{x}_t(r)^T \boldsymbol{\theta}_t(r)^{(\mathbf{x}_t(r))}$$

where $\theta_t(r)^{(r'_i)}$ is the parameter associated with the parent regressor $Y_t(r'_i)$ and $\boldsymbol{\theta}_t(r)^{(\mathbf{x}_t(r))}$ is the vector of parameters for known variables $\mathbf{x}_t(r)$. It can sometimes be helpful to find the covariance between two components $Y_t(i, i')$ and $Y_t(r, r')$ from the models for $Y_t(i)$ and $Y_t(r)$ respectively. Call $\text{cov}(Y_t(i, i'), Y_t(r, r'))$ the *component covariance* for $Y_t(i, i')$ and $Y_t(r, r')$. We shall illustrate why component covariances may be useful by considering two examples.

Example 2: component covariances in traffic networks

Queen *et al.* (2007) consider the problem of forecasting hourly vehicle counts at various points in a traffic network. The possible routes through the network are used to elicit a DAG for use with an LMDM. It may be useful to learn about driver route choice probabilities in such a network, i.e., the probability that a vehicle starting at a certain point A will travel to destination B. Unfortunately, these are not always easy to estimate from vehicle count data. However, component covariances could be useful in this respect, as illustrated by the following hypothetical example.

Consider the simple traffic network illustrated in Figure 6. There are five data collection sites, each of which records the hourly count of vehicles passing that site. There are four possible routes through the system: A to C, A to D, B to C and B to D. Because all traffic from A and B to C and D is counted at site 3, it can be difficult to learn about driver route choices using the time series of vehicle counts alone.

Figure 6. Hypothetical traffic network for Example 2 with starting points A and B and destinations C and D. The grey diamonds are the data collection sites, each of which is numbered. White arrows indicate the direction of traffic flow

Figure 7. DAG representing the conditional independence relationships between $Y_t(1), \dots, Y_t(5)$ in the traffic network of Figure 6

Let $Y_t(r)$ be the vehicle count for hour t at site r . A suitable DAG representing the conditional independence relationships between $Y_t(1), \dots, Y_t(5)$ is given in Figure 7. (For details on how this DAG can be elicited, see Queen *et al.*, 2007.) Note that all vehicles at site 3 flow to site 4 or 5 so that, conditional on $Y_t(3)$ and $Y_t(4)$, $Y_t(5)$ is a logical node with $Y_t(5) = Y_t(3) - Y_t(4)$.

From Figure 6, $Y_t(3)$ receives all its traffic from sites 1 and 2, while $Y_t(4)$ receives all its traffic from site 3. Thus possible LMDM observation equations for $Y_t(3)$ and $Y_t(4)$ are given by

$$Y_t(3) = Y_t(1)\theta_t(3)^{(1)} + Y_t(2)\theta_t(3)^{(2)} + v_t(3), \quad v_t(3) \sim N(0, V_t(3)),$$

$$Y_t(4) = Y_t(3)\theta_t(4)^{(3)} + v_t(4), \quad v_t(4) \sim N(0, V_t(4))$$

where $0 \leq \theta_t(3)^{(1)}, \theta_t(3)^{(2)}, \theta_t(4)^{(3)} \leq 1$. Then

$$Y_t(1)\theta_t(3)^{(1)} = Y_t(3, 1) = \text{number of vehicles travelling from site 1 to 3 in hour } t;$$

$$Y_t(3)\theta_t(4)^{(3)} = Y_t(4, 3) = \text{number of vehicles travelling from site 3 to 4 in hour } t.$$

Thus $\text{cov}(Y_t(3, 1), Y_t(4, 3))$ is informative about the use of route A to C. A high correlation between $Y_t(3, 1)$ and $Y_t(4, 3)$ indicates a high probability that a vehicle at A travels to C and a small correlation indicates a low probability. Of course, the actual driver route choice probabilities still cannot be estimated from these data. However having some idea of the relative magnitude of the choice probabilities can still be very useful.

Example 3: accommodating changes in the DAG

In many application areas the DAG representing the multivariate time series may change over time, either temporarily or permanently. For example, in a traffic network a temporary diversion may change the DAG temporarily, or a change in the road layout may change the DAG permanently. Because of the structure of the LMDM, much of the posterior information from the original DAG can be used to help form priors in the new DAG (see Queen *et al.*, 2007, for an example illustrating this). In this respect, component covariances may be informative about covariances in the new DAG. We shall illustrate this using a simple example.

Figure 8 shows part of a DAG representing four time series $Y_t(1), \dots, Y_t(4)$ (the DAG continues with children of $Y_t(3)$ and $Y_t(4)$, but we are only interested in $Y_t(1), \dots, Y_t(4)$ here). Using an LMDM and equation (4), suppose we have the following observation equations for $Y_t(3)$ and $Y_t(4)$:

$$Y_t(3) = Y_t(3, 1) + v_t(3), \quad v_t(3) \sim N(0, V_t(3)) \quad (5)$$

$$Y_t(4) = Y_t(4, 1) + Y_t(4, 2) + v_t(4), \quad v_t(4) \sim N(0, V_t(4)) \quad (6)$$

Now suppose that the DAG is changed so that a new series, X_t , is introduced which lies between $Y_t(1)$ and $Y_t(4)$ in Figure 8. Suppose further that $Y_t(1)$ and $Y_t(4)$ are no longer observed. The new DAG is given in Figure 9 (again the DAG continues, this time with children of $Y_t(3)$, X_t and $Y_t(2)$). The

Figure 8. Part of a DAG for Example 3 for four time series $Y_t(1), \dots, Y_t(4)$

Figure 9. New DAG for Example 3 (adapting the DAG in Figure 8) where X_t is introduced into the DAG and $Y_t(1)$ and $Y_t(4)$ are no longer observed

component $Y_t(3, 1)$ from equation (5) is informative about $Y_t(3)$ in the new model, and the component $Y_t(4, 1)$ from equation (6) is informative about X_t in the new model. Thus the component covariance $\text{cov}(Y_t(3, 1), Y_t(4, 1))$ is informative about $\text{cov}(Y_t(3), X_t)$ in the new model.

The following theorem presents a simple method for calculating component covariances.

Theorem 2 Suppose that $Y_t(i')$ is the parent of $Y_t(i)$ and $Y_t(r')$ is a parent of $Y_t(r)$, with $i < r$. Then

$$\text{cov}(Y_t(i, i'), Y_t(r, r') | D_{t-1}) = \text{cov}(Y_t(i'), Y_t(r') | D_{t-1}) a_t(i)^{(i')} a_t(r)^{(r')}$$

where $a_t(i)^{(i')}$ and $a_t(r)^{(r')}$ are, respectively, the prior means for the regressor $Y_t(i')$ in $Y_t(i)$'s model and the regressor $Y_t(r')$ in $Y_t(r)$'s model.

Proof For each r , let

$$\mathbf{Z}_t^{(r)T} = (Y_t(r, r'_1) \ Y_t(r, r'_2) \ \dots \ Y_t(r, r'_{m_r}) \ Y_t(r, \mathbf{x}_t(r)) \ v_t(r))$$

and

$$\mathbf{Z}_t(r)^T = (\mathbf{Z}_t^{(1)T} \ \mathbf{Z}_t^{(2)T} \ \dots \ \mathbf{Z}_t^{(r-1)T})$$

Using the result that for two random variables X and Y , $E(XY) = E(X \cdot E(Y|X))$, we have, for a specific $r' \in \{r'_1, \dots, r'_{m_r}\}$

$$\text{cov}(\mathbf{Z}_t(r), Y_t(r, r') | D_{t-1}) = E(\mathbf{Z}_t(r) \cdot E(Y_t(r, r') | \mathbf{Z}_t(r), D_{t-1}) | D_{t-1}) - E(\mathbf{Z}_t(r) | D_{t-1}) E(Y_t(r, r') | D_{t-1}) \quad (7)$$

Now, from Equation (4)

$$E(Y_t(r, r') | \mathbf{Z}_t(r), D_{t-1}) = E(Y_t(r, r') | Y_t(1), \dots, Y_t(r-1), D_{t-1}) = Y_t(r') a_t(r)^{(r')}$$

Also

$$E(Y_t(r, r') | D_{t-1}) = E[E(Y_t(r, r') | \mathbf{Z}_t(r), D_{t-1})] = E(Y_t(r') a_t(r)^{(r')} | D_{t-1})$$

Therefore, equation (7) becomes

$$\begin{aligned} \text{cov}(\mathbf{Z}_t(r), Y_t(r, r') | D_{t-1}) &= E(\mathbf{Z}_t(r) Y_t(r') a_t(r)^{(r')} | D_{t-1}) - E(\mathbf{Z}_t(r) | D_{t-1}) E(Y_t(r') a_t(r)^{(r')} | D_{t-1}) \\ &= \text{cov}(\mathbf{Z}_t(r), Y_t(r') | D_{t-1}) a_t(r)^{(r')} \end{aligned}$$

The single row from $\mathbf{Z}_t(r)$ corresponding to $Y_t(i, i')$ then gives us

$$\text{cov}(Y_t(i, i'), Y_t(r, r') | D_{t-1}) = \text{cov}(Y_t(i, i'), Y_t(r') | D_{t-1}) a_t(r)^{(r')} \quad (8)$$

Let

$$\mathbf{X}_t(r, i)^T = (Y_t(1) \ Y_t(2) \ \dots \ Y_t(i-1) \ Y_t(i+1) \ \dots \ Y_t(r-1))$$

Then

$$\text{cov}(X_t(r, i), Y_t(i, i') | D_{t-1}) = E(X_t(r, i) \cdot E(Y_t(i, i') | X_t(r, i), D_{t-1}) | D_{t-1}) - E(X_t(r, i) | D_{t-1}) E(Y_t(i, i') | D_{t-1}) \quad (9)$$

Now

$$E(Y_t(i, i') | X_t(r, i), D_{t-1}) = Y_t(i') a_t(i)^{(i')}$$

and equation (9) becomes

$$\text{cov}(X_t(r, i), Y_t(i, i') | D_{t-1}) = \text{cov}(X_t(r, i), Y_t(i') | D_{t-1}) a_t(i)^{(i')}$$

So taking the individual row of $X_t(r, i)$ corresponding to $Y_t(r')$ we get

$$\text{cov}(Y_t(r'), Y_t(i, i') | D_{t-1}) = \text{cov}(Y_t(r'), Y_t(i') | D_{t-1}) a_t(i)^{(i')}$$

Substituting this into equation (8) gives us

$$\text{cov}(Y_t(i, i'), Y_t(r, r') | D_{t-1}) = \text{cov}(Y_t(i'), Y_t(r') | D_{t-1}) a_t(i)^{(i')} a_t(r)^{(r')}$$

as required. $\square$

Theorem 2 allows the simple calculation of the component correlations. For example, in Example 2:

$$\text{cov}(Y_t(3, 1), Y_t(4, 3) | D_{t-1}) = \text{cov}(Y_t(1), Y_t(3) | D_{t-1}) a_t(3)^{(1)} a_t(4)^{(3)}$$

and $\text{cov}(Y_t(1), Y_t(3) | D_{t-1})$ is simple to calculate using Theorem 1.

Corollary 3 Suppose that $Y_t(i')$ is the parent of $Y_t(i)$ and $Y_t(r')$ is a parent of $Y_t(r)$, with $i < r$. Then, using the notation presented in Corollary 1, for $i < r$ and $k \geq 1$, the k -step-ahead forecast covariance between components $Y_{t+k}(i, i')$ and $Y_{t+k}(r, r')$ can be calculated recursively using

$$\text{cov}(Y_{t+k}(i, i'), Y_{t+k}(r, r') | D_t) = \sum_{l=1}^{m_r} \text{cov}(Y_{t+k}(i'), Y_{t+k}(r') | D_t) a_t(i, k)^{(i')} a_t(r, k)^{(r')}$$

Proof This is proved using the same argument to that used in the proof of Theorem 2, where t is replaced by $t + k$ and noting the $E(\theta_{t+k}(i)^{(i')} | D_t) = a_t(i, k)^{(i')}$ and $E(\theta_{t+k}(r)^{(r')} | D_t) = a_t(r, k)^{(r')}$. $\square$

COVARIANCE BETWEEN ROOT NODES

Recall from Corollary 2 that the covariance between two root nodes is zero in the LMDM. However, this is not always appropriate. For example, consider the traffic network in Example 2. Both $Y_t(1)$ and $Y_t(2)$ are root nodes which may in fact be highly correlated—they may have similar daily patterns with the same peak times, etc., and they may be affected in a similar way by external events such as weather conditions.

Figure 10. DAG representing time series of vehicle counts in Example 4, with an extra node X_t to account for any correlation between $Y_t(1)$ and $Y_t(2)$

One possible way to introduce non-zero covariances between root nodes is to add an extra node as a parent of all root nodes in the DAG. This extra node represents any variables which may account for the correlation between the root nodes.

Example 4

In the traffic network in Example 2, suppose that X_t is a vector of variables which can account for the correlation between $Y_t(1)$ and $Y_t(2)$. So, X_t might include such variables as total traffic volume entering the system, hourly rainfall, temperature, and so on. Then X_t can be introduced into the DAG as a parent of $Y_t(1)$ and $Y_t(2)$ as in Figure 10.

The observation equations for $Y_t(1)$ and $Y_t(2)$ now both have X_t as regressors and applying Theorem 1 we get

$$\text{cov}(Y_t(1), Y_t(2)|D_{t-1}) = \text{var}(X_t|D_{t-1})\mathbf{a}_t(1)^{(X_t)}\mathbf{a}_t(2)^{(X_t)}$$

where $\mathbf{a}_t(2)^{(X_t)}$ and $\mathbf{a}_t(2)^{(X_t)}$ are the prior mean vectors for regressors X_t in $Y_t(1)$ and $Y_t(2)$'s model.

CONCLUDING REMARKS

In this paper we have presented a simple algebraic form for calculating the one-step-ahead covariance matrix in LMDMs. We have also introduced a simple method for calculating covariances between regression components of different DLMS within the LMDM. Component covariances may be of interest in their own right, and may also prove to be useful for forming informative priors following any changes in the DAG for the LMDM. Their use in practice now needs to be investigated in further research.

One of the problems with the LMDM is the imposition of zero covariance between root nodes. To allow non-zero covariance between root nodes we have proposed introducing X_t into the model as a parent of all the root nodes, where X_t is a set of variables which may help to explain the correlation between parents. Further research is now required to investigate how well this might work in practice.

ACKNOWLEDGEMENTS

We would like to thank the reviewer for their very useful comments on an earlier version of this paper and also Unilever Research for supplying the data.

REFERENCES

- Cowell RG, Dawid AP, Lauritzen SL, Spiegelhalter DJ. 1999. *Probabilistic Networks and Expert Systems*. Springer: New York.
- Dawid AP. 1979. Conditional independence in statistical theory (with discussion). *Journal of the Royal Statistical Society, Series B* **41**: 1–31.
- Farrow M. 2003. Practical building of subjective covariance structures for large complicated systems. *Statistician* **52**(4): 553–573.
- Farrow M, Goldstein M, Spiropoulos T. 1997. Developing a Bayes linear decision support system for a brewery. In *The Practice of Bayesian Analysis*, French S, Smith JQ (eds). Arnold: London; 71–106.
- Fosen J, Ferkingstad E, Borgan Ø, Aalen OO. 2006. Dynamic path analysis: a new approach to analyzing time-dependent covariates. *Lifetime Data Analysis* **12**(2): 143–167.
- Goldstein M, Farrow M, Spiropoulos T. 1993. Prediction under the influence: Bayes linear influence diagrams for prediction in a large brewery. *Statistician* **42**(4): 445–459.
- Queen CM. 1994. Using the multiregression dynamic model to forecast brand sales in a competitive product market. *Statistician* **43**(1): 87–98.
- Queen CM. 1997. Model elicitation in competitive markets. In *The Practice of Bayesian Analysis*, French S, Smith JQ (eds). Arnold: London; 229–243.
- Queen CM, Smith JQ. 1993. Multiregression dynamic models. *Journal of the Royal Statistical Society, Series B* **55**: 849–870.
- Queen CM, Smith JQ, James DM. 1994. Bayesian forecasts in markets with overlapping structures. *International Journal of Forecasting* **10**: 209–233.
- Queen CM, Wright BJ, Albers CJ. 2007. Eliciting a directed acyclic graph for a multivariate time series of vehicle counts in a traffic network. *Australian and New Zealand Journal of Statistics* **49**: 221–239.
- Sun SL, Zhang CS, Yu GQ. 2006. A Bayesian network approach to traffic flows forecasting. *IEEE Transactions on Intelligent Transportation Systems* **7**(1): 124–132.
- Tebaldi C, West M, Karr AF. 2002. Statistical analyses of freeway traffic flows. *Journal of Forecasting* **21**: 39–68.
- West M, Harrison PJ. 1997. *Bayesian Forecasting and Dynamic Models* (2nd edn). Springer: New York.
- Whitlock ME, Queen CM. 2000. Modelling a traffic network with missing data. *Journal of Forecasting* **19**(7): 561–574.

Authors' biographies:

Catriona Queen is a Lecturer in Statistics in the Statistics Department at the Open University, UK. Her research interests include dynamic models, multivariate forecasting, graphical models and Bayesian statistics.

Casper Albers is a Research Fellow in the Statistics Department at the Open University, UK. His research interests include multivariate forecasting, Bayesian statistics and statistical geometry.

Ben Wright is a Research Associate in the Department of Health Sciences at the University of Leicester. His research interests include Bayesian Statistics, dynamic models, multivariate time series, Markov chain Monte Carlo methods and genome-wide association studies.

Authors' addresses:

Catriona Queen and **Casper Albers**, Department of Statistics, The Open University, Milton Keynes, MK7 6AA, UK.

Ben Wright, Department of Health Sciences, University of Leicester, Adrian Building, Leicester, LE1 6TP, UK.