

Forecasting Emergency Medical Service Call Arrival Rates

David S. Matteson* Mathew W. McLean Dawn B. Woodard
Shane G. Henderson

School of Operations Research and Information Engineering
Cornell University

November 19, 2010

Abstract

We introduce a new method for forecasting emergency call arrival rates that combines integer-valued time series models with a dynamic latent factor structure. Covariate information is captured via simple constraints on the factor loadings. We directly model the count-valued arrivals per hour, rather than using an artificial assumption of normality. This is crucial for the emergency medical service context, in which the volume of calls may be very low. Smoothing splines are used in estimating the factor levels and loadings to improve long-term forecasts. We impose time series structure at the hourly level, rather than at the daily level, capturing the fine-scale dependence in addition to the long-term structure.

Our analysis considers all emergency priority calls received by Toronto EMS between January 2007 and December 2008 for which an ambulance was dispatched. Empirical results demonstrate significantly reduced error in forecasting call arrival volume. To quantify the impact of reduced forecast errors we design a queueing model simulation that approximates the dynamics of an ambulance system. The results show better performance as the forecasting method improves. This notion of quantifying the operational impact of improved statistical procedures may be of independent interest.

KEY WORDS: Ambulance planning; Dynamic factor model; Non-homogeneous Poisson process; Integer-valued time series; Smoothing splines.

Short title: Forecasting EMS Call Arrival Rates

**Address for Correspondence:* School of Operations Research and Information Engineering, Cornell University, 282 Rhodes Hall, Ithaca, NY 14853, dm484@cornell.edu

1. INTRODUCTION

Considerable attention has been paid to the problem of how to best deploy ambulances within a municipality to minimize their response times to emergency calls while keeping costs low. Sophisticated operations research models have been developed to address issues such as the optimal number of ambulances, where to place bases, and how to move ambulances in real time via system-status management (Swersey 1994; Goldberg 2004; Henderson 2009). However, methods for estimating the inputs to these models, such as travel times on road networks and call arrival rates, are *ad hoc*. Use of inaccurate parameter estimates in these models can result in poor deployment decisions, leading to low performance and diminished user confidence in the software. We introduce methods for estimating the demand for ambulances, i.e., the total number of emergency calls per period, that are highly accurate, straightforward to implement, and have the potential to simultaneously lower operating costs while improving response times.

Current practice for forecasting call arrivals is often rudimentary. For instance, to estimate the call arrival rate in a small region over a specific time period, e.g. next Monday from 8-9am, simple estimators have been constructed by averaging the number of calls received in the corresponding period in four previous weeks: the immediately previous two weeks and the current and previous weeks of the previous year. Averages of so few data points can produce highly noisy estimates, with resultant cost and efficiency implications. Excessively large estimates lead to over-staffing and unnecessarily high costs, while low estimates lead to under-staffing and slow response times. Setzler et al. (2009) document an emergency medical service (EMS) agency which extends this simple moving average to twenty previous observations: the previous four weeks from the previous five years. A more formal time series approach is able to account for possible differences from week to week and allows inclusion of neighboring hours in the estimate.

We generate improved forecasts of the call-arrival volume by introducing an integer-valued time series model with a dynamic latent factor structure for the hourly call arrival rate. Day-of-week and week-of-year effects are included via simple constraints on the factor loadings. The factor structure allows for a significant reduction in the number of model parameters. Further, it provides a systematic approach to modeling the diurnal pattern observed in intraday

counts. Smoothing splines are used in estimating the factor levels and loadings. This may introduce a small bias in some periods, but it offers a significant reduction in long-horizon out-of-sample forecast-error variance. This is combined with integer-valued time series models to capture residual dependence and to provide adaptive short-term forecasts. Our empirical results demonstrate significantly reduced error in forecasting hourly call-arrival volume.

Few studies have focused specifically on EMS call arrival rates, and of those that have proposed methods for time series modeling, most have been based on Gaussian linear models. Even with a continuity correction, this method is highly inaccurate when the call arrival rate is low, which is typical of EMS calls at the hourly level. Further, it conflicts with the Poisson distribution assumption used in operations research methods for optimizing staffing levels. For example, Channouf et al. (2007) forecast EMS demand by modeling the daily call arrival rate as Gaussian, with fixed day-of-week, month-of-year, special day effects and fixed day-month interactions. They also consider a Gaussian autoregressive moving-average (ARMA) model with seasonality and special day effects. Hourly rates are later estimated either by adding hour-of-day effects or assigning a multinomial distribution to the hourly volumes, conditional on the daily call volume estimates.

Setzler et al. (2009) provide a comparative study of EMS call volume predictions using an artificial neural network (ANN). They forecast at various temporal and spatial granularities with mixed results. Their approach offered a significant improvement at low spatial granularity, even at the hourly level. At a high spatial granularity, the mean square forecast error (MSFE) of their approach did not improve over simple binning methods at a temporal granularity of three hours or less.

Methods for the closely related problem of forecasting call center demand have received much more study. Bianchi et al. (1998), Andrews and Cunningham (1995) use ARMA models to improve forecasts for daily call volumes in a retail company call center and a telemarketing center, respectively. A dynamic harmonic regression model for hourly call center demand is shown in Tych et al. (2002) to outperform seasonal ARMA models. Their approach accounts for possible non-stationary periodicity in a time series. The major drawback common to these studies is that the integer-valued observations are assumed to have a continuous distribution,

which is problematic during periods with low arrival rates.

The standard industry assumption is that hourly call-arrival volume has a Poisson distribution. The Palm-Khintchine theorem - stating that the superposition of a number of independent point processes is approximately Poisson - provides a theoretical basis for this assumption (see, e.g., Whitt 2002). Brown et al. (2005) provide a comprehensive analysis of operational data from a bank call center and thoroughly discuss validating classical queueing theory, including this theorem. Henderson (2005) states that we can expect the theorem to hold for typical EMS data because there are a large number of callers who can call at any time and each has a very low probability of doing so.

Weinberg et al. (2007) use Bayesian techniques to fit a non-homogenous Poisson process model for call arrivals to a commercial bank's call center. This approach has the advantage that forecast distributions for the rates and counts may be easily obtained. They incorporate smoothness in the within-day pattern. They implement a variance stabilizing transformation to obtain approximate normality. This approximation is most appropriate for a Poisson process with high arrival rates, and would not be appropriate for our application in which very low counts are observed in many time periods.

Shen and Huang (2008b) apply the same variance stabilizing transformation and achieve better performance than Weinberg et al. (2007). They use a singular value decomposition (SVD) to reduce the number of parameters in modeling arrival rates. Their approach is used for intraday updating and forecasts up to one day ahead.

Shen and Huang (2008a) propose a dynamic factor model for 15-minute call arrivals to a bank call center. They assume that call arrivals are a Cox process. A Cox process (cf. Cox and Isham 1980) is a Poisson process with a stochastic intensity, i.e., a doubly stochastic Poisson process. The factor structure reduces the number of parameters by explaining the variability in the call arrival rate with a small number of unobserved variables. Estimation proceeds by iterating between an SVD and fitting Poisson generalized linear models to successively estimate the factors and their respective loadings. The intensity functions are assumed to be serially dependent. Forecasts are made by fitting a simple autoregressive time series model to the factor score series.

We assume that the hourly EMS call-arrival volume has a Poisson distribution. This allows parsimonious modeling of periods with small counts, conforms with the standard industry assumption, and avoids use of variance stabilizing transformations. We assume the intensity function is a random process and that it can be forecast using previous observations. This has an interpretation very similar to a Cox process, but is not equivalent since the random intensity is allowed to depend on not only its own history, but also on previous observations. We partition the random intensity function into stationary and non-stationary components.

Section 2 describes the general problem and our dataset. Section 3 presents the proposed methodology. We consider a dynamic latent factor structure to model the non-stationary pattern in intraday call arrivals and greatly reduce the number of parameters. We include day-of-week and week-of-year covariates via simple constraints on the factor loadings of the non-stationary pattern. Smoothing splines are easily incorporated into estimation of the proposed model to impose a smooth evolution in the factor levels and loadings, leading to improved long-horizon forecast performance. We combine the factor model with stationary integer-valued time series models to capture the remaining serial dependence in the intensity process. This is shown to further improve short-term forecast performance of our approach. A simple iterative algorithm for estimating the proposed model is presented. It can be implemented largely through existing software. Section 4 assesses the performance of our approach using statistical metrics and a queueing model simulation. Section 5 gives our concluding remarks.

2. NOTATION AND DATA DESCRIPTION

We assume that over short, equal-length time intervals, e.g. one hour periods, the latent call arrival intensity function can be well approximated as being constant, and that all data have been aggregated in time accordingly. We suppose aggregated call arrivals follow a non-homogeneous counting process $\{Y_t : t \in \mathbb{Z}\}$, with discrete time index t . Underlying this is a latent, real-valued, non-negative intensity process $\{\lambda_t : t \in \mathbb{Z}\}$. We further assume that conditional on λ_t , Y_t has a Poisson distribution with mean λ_t .

As shown in Figure 1, the pattern of call arrivals over the course of a typical day has a distinct shape. After quickly increasing in the late morning it peaks in the early afternoon,

then slowly falls until it troughs between 5 and 6 a.m. See Section 4 for more discussion. In our analysis, we consider an arrival process that has been repeatedly observed over a particular time span, specifically a 24 hour day. Let

$$\{y_t : t = 1, \dots, n\} \equiv \{y_{ij} : i = 1, \dots, d; j = 1, \dots, m\}$$

denote the sequence of call arrival counts, observed over time period t , which corresponds one-to-one with the j^{th} sub-period of the i^{th} day, so that $n = dm$. Our baseline approach is to model the arrival intensity λ_t for the distinct shape of intraday call arrivals using a small number of smooth curves.

We consider two disjoint information sets for predictive conditioning. Let $\mathcal{F}_t = \sigma(Y_1, \dots, Y_t)$ denote the σ -field generated by Y_1, \dots, Y_t , and let $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ denote any available deterministic covariate information about each observation. We incorporate calendar information such as day-of-week and week-of-year in our analysis. We define λ_t as the conditional expectation of Y_t given \mathcal{F}_{t-1} and \mathbf{X} . We defined this above as the mean of Y_t . In our *model* these coincide; however, this mean may not be the same as the conditional expectation since λ_t may depend on other unobserved random variables. Let $\mu_t = E(Y_t|\mathbf{X}) > 0$ denote the conditional mean of Y_t given only the covariates \mathbf{X} . Let

$$\lambda_t = E(Y_t|\mathcal{F}_{t-1}, \mathbf{X}) = \mu_t E(Y_t/\mu_t|\mathcal{F}_{t-1}, \mathbf{X}) = \mu_t \eta_t, \quad (1)$$

in which $\eta_t > 0$ is referred to as the conditional intensity inflation rate (CIIR). By construction

$$E(\eta_t|\mathbf{X}) = E(E(Y_t|\mathcal{F}_{t-1}, \mathbf{X})|\mathbf{X})/\mu_t = E(Y_t|\mathbf{X})/\mu_t = 1.$$

The CIIR process is intended to model any remaining serial dependence in call arrival counts after accounting for available covariates. In the EMS context we hypothesize that this dependence is due to sporadic events such as inclement weather or unusual traffic patterns. Since information regarding these events may not be available or predictable in general, we argue that an approach such as ours which explicitly models the remaining serial dependence will lead to improved short-term forecast accuracy. In Section 3 we consider a dynamic latent factor model estimated with smoothing splines for modeling μ_t , various time series models for

modeling η_t , and finally a conditional likelihood algorithm for estimating the latent intensity process λ_t via estimation of η_t given μ_t .

The call arrival data used consists of all emergency priority calls received by Toronto EMS between January 1, 2007 and December 31, 2008 for which an ambulance was dispatched. This includes some calls not requiring lights-and-sirens response, but does not include scheduled patient transfers. We include only the first call arrival time in our analysis when multiple calls are received for the same event. The data were processed to exclude calls with no reported location. These removals totaled less than 1% of the data.

Many calls resulted in multiple ambulances being dispatched. Exploratory analysis revealed that the number of ambulances deployed for a single emergency did not depend on the day of the week, the week of the year, or exhibit any serial dependence. However, such instances were slightly more prevalent in the morning hours. Our analysis of hourly ambulance demand defines an event as a call arrival if *one or more* ambulances are deployed.

We removed seven days from the analysis because there were large gaps, over at least two consecutive hours, in which no emergency calls were received. These days most likely resulted from malfunctions in the computer-aided dispatch system which led to failures in recording calls for extended periods. Strictly speaking, it is not necessary to remove the entire days; however, we did so since it had a negligible impact on our results and it greatly simplified out-of-sample forecast comparisons and implementation of the simulation studies in Section 4.

Finally, we gave special consideration to holidays. We found that the intraday pattern on New Year's Eve and Day was fundamentally different from the rest of the year and removed these days from our analysis. This finding is similar to the conclusions of Channouf et al. (2007) who found that New Year's Day and the dates of the Calgary Stampede were the only days requiring special consideration in their methodology when applied to the city of Calgary. In practice, staffing decisions for holidays require special planning and consideration of many additional variables.

3. MODELING

Factor models provide a parsimonious representation of high dimensional data in many applied sciences, e.g., econometrics (cf. Stock and Watson 2002). We combine a dynamic latent factor

model with integer-valued time series models. We include covariates via simple constraints on the factor loadings. We estimate the model using smoothing splines to impose smooth evolution in the factor levels and loadings. The factor model provides a parsimonious representation of the non-stationary pattern in intraday call arrivals while the time series models capture the remaining serial dependence in the arrival rate process.

3.1 Dynamic Latent Factor Model

For notational simplicity, assume m consecutive observations per day are available for d consecutive days with no omissions in the record. Let $\mathbf{Y} = (y_{ij})$ denote the $d \times m$ matrix of observed counts for each day i over each sub-period j . Let $\mu_{ij} = E(Y_{ij}|\mathbf{X})$, and let $\mathbf{M} = (\mu_{ij})$ denote the corresponding $d \times m$ latent non-stationary intensity matrix. To reduce the dimension of the intensity matrix \mathbf{M} , we introduce a K -factor model.

We assume that the intraday pattern of expected hourly call arrivals on the log scale can be well approximated by a linear combination of (a small number) K factors or functions, denoted by \mathbf{f}_k for $k = 1, \dots, K$. The factors are orthogonal length- m vectors. The intraday arrival rate model $\boldsymbol{\mu}_i$ over a particular day i is given by

$$\log \boldsymbol{\mu}_i = L_{i1}\mathbf{f}_1 + \dots + L_{iK}\mathbf{f}_K. \quad (2)$$

Each of the factors \mathbf{f}_k varies as a function over the periods within a day, but they are constant from one day to the next. Day-to-day changes are modeled by allowing the various factor loadings L_{ik} to vary across days. When K is much smaller than either m or d , the dimensionality of the general problem is greatly reduced. In practice, K must be chosen by the practitioner; we provide some discussion on choosing K in Section 4.

In matrix form we have

$$\log \mathbf{M} = \mathbf{L}\mathbf{F}^T, \quad (3)$$

in which $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_K)$ denotes the $m \times K$ matrix of underlying factors and \mathbf{L} denotes the corresponding $d \times K$ matrix of factor loadings, both of which are assumed to have full column rank. Although other link functions are available, the component-wise log transformation implies a multiplicative structure among the K common factors and ensures a positive estimate of each hourly intensity μ_{ij} . Since neither \mathbf{F} nor \mathbf{L} are observable, the expression (3) is not

identifiable. We further require $\mathbf{F}^T \mathbf{F} = \mathbf{I}$ to alleviate this ambiguity and we iteratively estimate \mathbf{F} and \mathbf{L} .

3.2 Factor Modeling with Covariates via Constraints

To further reduce the dimensionality we impose a set of constraints on the factor loading matrix \mathbf{L} . Let \mathbf{H} denote a $d \times r$ full rank matrix ($r < d$) of given constraints (we discuss later what these should be for EMS). Let \mathbf{B} denote an $r \times K$ matrix of unconstrained factor loadings. These unconstrained loadings \mathbf{B} linearly combine to constitute the constrained factor loadings \mathbf{L} , such that $\mathbf{L} = \mathbf{H}\mathbf{B}$. Our factor model may now be written as

$$\log \mathbf{M} = \mathbf{L}\mathbf{F}^T = \mathbf{H}\mathbf{B}\mathbf{F}^T.$$

A considerable reduction in dimensionality occurs when r is much smaller than d .

Constraints to assure identifiability are standard in factor analysis. The constraints we now consider incorporate auxiliary information about the rows and columns of the observation matrix \mathbf{Y} to simplify estimation and to improve out-of-sample predictions. Similar constraints have been used in Takane and Hunter (2001), Tsai and Tsay (2009), Matteson and Tsay (2010).

For example, the rows of \mathbf{H} might consist of incidence vectors for particular days of the week, or special days which might require unique loadings on the common factors. We may choose to constrain all weekdays to have identical factor loadings and similarly constrain weekend days. However, this approach is much more general than simple equality constraints, as demonstrated below.

The intraday pattern of hourly call arrivals varies from one day to the next, although the same general shape is maintained. As seen in Figure 1, different days of the week exhibit distinct patterns. We do not observe large changes from one week to the next, but there are significant changes over the course of the year. We allow loadings to slowly vary from week to week. Both of these features are incorporated into the factor loadings \mathbf{L} by specifying appropriate constraints \mathbf{H} . Let

$$\log \mathbf{M} = \mathbf{L}\mathbf{F}^T = \mathbf{H}\mathbf{B}\mathbf{F}^T = (\mathbf{H}^{(1)} \quad \mathbf{H}^{(2)}) \begin{pmatrix} \mathbf{B}^{(1)} \\ \mathbf{B}^{(2)} \end{pmatrix} \mathbf{F}^T = (\mathbf{H}^{(1)}\mathbf{B}^{(1)} + \mathbf{H}^{(2)}\mathbf{B}^{(2)})\mathbf{F}^T, \quad (4)$$

in which the first term corresponds to day-of-week effects and the second to smoothly varying week-of-year effects. $\mathbf{H}^{(1)}$ is a $d \times 7$ matrix in which each row $\mathbf{H}_i^{(1)}$ is an incidence vector for

the day-of-week. Similarly, $\mathbf{H}^{(2)}$ is a $d \times 53$ matrix in which each row $\mathbf{H}_i^{(2)}$ is an incidence vector for the week-of-year. (We use a 53 week year since the first and last weeks may have fewer than 7 days.) The $7 \times K$ matrix $\mathbf{B}^{(1)} = (\mathbf{b}_1^{(1)}, \dots, \mathbf{b}_K^{(1)})$ contains unconstrained factor loadings for the day-of-week and $\mathbf{B}^{(2)} = (\mathbf{b}_1^{(2)}, \dots, \mathbf{b}_K^{(2)})$ is a $53 \times K$ matrix of factor loadings for the week-of-year.

3.3 Factor Model Estimation via Smoothing Splines

We assume that as the non-stationary intensity process μ_{ij} varies over the hours j of each day i , it does so smoothly. If each of the common factors $\mathbf{f}_k \in \mathbb{R}^m$ varies smoothly over sub-periods j , then the smoothness of μ_{ij} is guaranteed for each day. Increasing the number of factors reduces possible discontinuities between the end of one day and the beginning of the next. To incorporate smoothness into the model (2), we use Generalized Additive Models (GAMs) in the estimation of the common factors \mathbf{f}_k . GAMs extend generalized linear models, allowing for more complicated relationships between the response and predictors by modeling some predictors nonparameterically, (see, e.g., Hastie and Tibshirani 1990; Wood 2006). GAMs have been successfully used for count-valued data in the study of fish populations (cf. Borchers et al. 1997; Daskalov 1999). The factors $\mathbf{f}_k = f_k(j)$ are a smooth function of the intraday time index covariate j . The loadings \mathbf{L} are defined as before. If the loadings \mathbf{L} were known covariates, Equation (2) would be a varying coefficient model (cf. Hastie and Tibshirani 1993).

There are several excellent libraries available in the statistical package *R* (R Development Core Team 2009) for fitting GAMs, thus making them quite easy to implement. We used the *gam* function from the *mgcv* library (Wood 2008) extensively. Other popular libraries include the *gam* package (Hastie 2009) and the *gss* package (Gu 2010). See Wood (2006, Section 5.6) for an introduction to GAM estimation using *R*.

In estimation of the model (2) via the *gam* function, we have used thin plate regression splines with a ten-dimensional basis, the Poisson family, and the log-link function. Thin plate regression splines are a low rank, isotropic smoother with many desirable properties. For example, no decisions on the placement of knots is needed. They are an optimal approximation to thin plate splines and, with the use of Lanczos iteration, they can be fit quickly even for large datasets (cf. Wood 2003).

When the factors \mathbf{F} are treated as a fixed covariate, the factor model can again be interpreted as a varying coefficient model. Given the calendar covariates \mathbf{X} , let

$$\begin{aligned}
\log \mu_{ij} &= F_{j1} L_{1i}^T + \cdots + F_{jK} L_{Ki}^T \\
&= \sum_{k=1}^K F_{jk} \left\{ \mathbf{H}_i^{(1)T} \mathbf{b}_k^{(1)} + \mathbf{H}_i^{(2)T} \mathbf{b}_k^{(2)} \right\} \\
&= \sum_{k=1}^K F_{jk} \left\{ b_k^{(1)}(\mathbf{x}_i) + b_k^{(2)}(\mathbf{x}_i) \right\}, \tag{5}
\end{aligned}$$

in which $b_k^{(1)}(\mathbf{x}_i)$ is a piece-wise constant function of the day-of-week, and $b_k^{(2)}(\mathbf{x}_i)$ is a smoothly varying function over the week-of-year. We may again proceed with estimation via the *gam* function in *R*. Day-of-week covariates are simply added to the linear predictor as indicator variables. These represent a level shift in the daily loadings on each of the factors \mathbf{f}_k . In our application it is appropriate to assume a smooth transition between the last week of one year and the first week of the next. To ensure this in estimation of $b_k^{(2)}(\mathbf{x}_i)$, we use a cyclic cubic regression spline for the basis (cf. Wood 2006, Section 5.1). Iterative estimation of \mathbf{F} , and \mathbf{L} via \mathbf{B} , for a given number of factors K is discussed in Section 3.5.

We allow the degree of smoothness for the factors f_k and the loadings function $b_k^{(2)}(\mathbf{x}_i)$ to be automatically estimated by generalized cross validation (GVC). We expect short term serial dependence in the residuals for our application. For smoothing methods in general, if autocorrelation between the residuals is ignored, automatic smoothing parameter selection may break down (see, e.g., Opsomer et al. 2001). The proposed factor model may be susceptible to this if the number of days included is not sufficiently large compared to the number of smooth factors and loadings, or if the residuals are long-range dependent. We use what is referred to as a *performance* iteration (cf. Gu 1992) versus an *outer* iteration strategy which requires repeated estimation for many trial sets of the smoothing parameters. The performance iteration strategy is much more computationally efficient for use in the proposed algorithm, but convergence is not guaranteed, in general. In particular, *cycling* between pairs of smoothing parameters and coefficient estimates may occur (cf. Wood 2006, Section 4.5), especially when the number of factors K is large.

3.4 Adaptive Forecasting with Time Series Models

Let $\hat{e}_t = Y_t/\hat{\mu}_t$ denote the multiplicative residual in period t implied by the fitted values $\hat{\mu}_t$ from a factor model estimated as described in the previous sections. Time series plots of this residual process appear stationary, but exhibit some serial dependence. In this section, we consider time series models for the latent CIIR process $\eta_t = E(Y_t/\mu_t|\mathcal{F}_{t-1}, \mathbf{X})$ to account for this dependence.

To investigate the nature of the serial dependence, we study the bivariate relationship between the \hat{e}_t process versus several lagged values of the process $\hat{e}_{t-\ell}$. Scatterplots reveal a roughly linear relationship. Residual autocorrelation and partial autocorrelation plots for one of the factor models fit in Section 4 are given in Figures 5(b,c). These quantify the strength of the linear relationship as the lag ℓ increases. It appears to persist for many periods, with an approximately geometric rate of decay as the lag increases.

To explain this serial dependence we first consider a generalized autoregressive linear model, defined by the recursion

$$\eta_t = \omega + \alpha\hat{e}_{t-1} + \beta\eta_{t-1}. \quad (6)$$

To ensure positivity, we restrict $\omega > 0$ and $\alpha, \beta \geq 0$. When μ_t is constant, the resulting model for Y_t is an Integer-GARCH(1,1) (IntGARCH) model (e.g., Ferland et al. 2006). It is worth noting some properties of this model for the constant μ_t case. To ensure the stationarity of η_t , we further require that $\alpha + \beta < 1$. This sum determines the persistence of the process, with larger values of α leading to more adaptability. When this stationarity condition is satisfied, and η_t has reached its stationary distribution, the expectation of η_t given \mathbf{X} is

$$E(\eta_t|\mathbf{X}) = \omega/(1 - \alpha - \beta).$$

To ensure $E(\eta_t|\mathbf{X}) = 1$ for the fitted model, we may parameterize $\omega = 1 - \alpha - \beta$. This constraint is simple enough to enforce for the model (6) and we do so. However, additional parameter constraints such as this may make numerical estimation intractable in more complicated models and they are not enforced by us in the models outlined below.

When μ_t is a non-stationary process, the conditional intensity

$$\lambda_t = \mu_t\eta_t$$

is also non-stationary. Since $E(\eta_t|\mathbf{X}) = 1$, we interpret η_t as the stationary multiplicative deviation, or inflation rate, between λ_t and μ_t . The λ_t process is mean reverting to the μ_t process. Let

$$\widehat{\varepsilon}_t = Y_t/\widehat{\lambda}_t$$

denote the multiplicative *standardized* residual process given an estimated CIIR process $\widehat{\eta}_t$. If a fitted model defined by (6) sufficiently explains the observed linear dependence in $\widehat{\varepsilon}_t$, then an autocorrelation plot of $\widehat{\varepsilon}_t$ should be statistically insignificant for all lags ℓ . As a preview, the standardized residual autocorrelation plot for one such model fit in Section 4 is given in Figure 5(d). The serial correlation appears to have been adequately removed.

Next, we formulate three different nonlinear generalizations of (6) that may better characterize the serial dependence, and possibly lead to improved forecasts. The first is an exponential autoregressive model defined as

$$\eta_t = \alpha\widehat{\varepsilon}_{t-1} + [\beta + \delta \exp(-\gamma\eta_{t-1}^2)]\eta_{t-1}, \quad (7)$$

in which $\alpha, \beta, \delta, \gamma > 0$. Exponential autoregressive models are attractive in application because of their threshold-like behavior. For large η_{t-1} , the functional coefficient for η_{t-1} is approximately β , and for small η_{t-1} it is approximately $\beta + \delta$. Additionally, the transition between these regimes remains smooth. As in Fokianos et al. (2009), for $\alpha + \beta < 1$ one can verify the η_t process has a stationarity version when μ_t is constant.

We also consider a piecewise linear threshold model

$$\eta_t = \omega + \alpha\widehat{\varepsilon}_{t-1} + \beta\eta_{t-1} + (\gamma\widehat{\varepsilon}_{t-1} + \delta\eta_{t-1}) I_{\{\widehat{\varepsilon}_{t-1} \notin (c_1, c_2)\}}, \quad (8)$$

in which I is an indicator variable and the threshold boundaries satisfy $0 < c_1 < 1 < c_2 < \infty$. To ensure positivity of η_t we assume $\omega, \alpha, \beta > 0$, $(\alpha + \gamma) > 0$, and $(\beta + \delta) > 0$. Additionally, we take $\delta \leq 0$ and $\gamma \geq 0$, such that when $\widehat{\varepsilon}_{t-1}$ is outside the range (c_1, c_2) the CIIR process η_t is more adaptive, i.e., puts more weight on $\widehat{\varepsilon}_{t-1}$ and less on η_{t-1} . When μ_t is constant, the η_t process has a stationary version under the restriction $\alpha + \beta + \gamma + \delta < 1$, see Woodard et al. (2010). In practice, the threshold boundaries c_1 and c_2 are fixed during estimation, and may be adjusted as necessary after further exploratory analysis. We chose $c_1 = 1/1.15$ and $c_2 = 1.15$, i.e. thresholds at 15% above and below 1.

Finally, we consider a model with regime switching at deterministic times, let

$$\eta_t = (\omega_1 + \alpha_1 \widehat{e}_{t-1} + \beta_1 \eta_{t-1}) I_{\{t \in (t_1, t_2)\}} + (\omega_2 + \alpha_2 \widehat{e}_{t-1} + \beta_2 \eta_{t-1}) I_{\{t \notin (t_1, t_2)\}}. \quad (9)$$

This model is appropriate assuming the residual process has two distinct regimes for different periods of the day. For example, one regime could be for normal workday hours with the other regime being for the evening and early morning hours. No stationarity is possible for this model. A drawback of this model is that the process has jumps at t_1 and t_2 . As was the case for c_1 and c_2 in (8), t_1 and t_2 are fixed during estimation. After exploratory analysis, we chose $t_1 = 10$ am and $t_2 = 4$ pm.

3.5 Estimation Algorithm

The estimation procedure below begins with an iterative algorithm for estimating the factor model from Sections 3.1 - 3.3 through repeated use of the *gam* function from the *mgcv* library in *R*. Any serial dependence is ignored during estimation of μ_t for simplicity. Given estimates for the factor model $\widehat{\mu}_t$, conditional maximum likelihood is used to estimate the conditional intensity λ_t via one of the time series models given in (6-9) for the CIIR process η_t .

1. Initialization:

- (a) Fix K and \mathbf{H} .
- (b) Choose some $c \in (0, 1)$ and define $\mathbf{Y}_c = (y_{ij} \vee c)$.
- (c) Apply a singular value decomposition (SVD) to find $\log(\mathbf{Y}_c) = \mathbf{U}_0 \mathbf{D}_0 \mathbf{V}_0^T$.
 - i. Let $\mathbf{U}_0^{(1:K)}$ denote the first K columns of the left singular matrix \mathbf{U}_0 .
 - ii. Let $\mathbf{V}_0^{(1:K)}$ denote the first K columns of the right singular matrix \mathbf{V}_0 .
 - iii. Let $\mathbf{D}_0^{(1:K)}$ denote the upper-left $K \times K$ sub-matrix of \mathbf{D}_0 , the diagonal matrix of singular values.
- (d) Assign $\mathbf{L}_0 = \mathbf{U}_0^{(1:K)} \mathbf{D}_0^{(1:K)}$ and $\mathbf{F}_0 = \mathbf{V}_0^{(1:K)}$.
No smoothing is performed and the constraints \mathbf{H} are omitted in initialization.

2. Update:

- (a) Fit the Poisson GAM model described in Section 3.3 with $\mathbf{F} = \mathbf{F}_n$ and \mathbf{H} as fixed covariates.
 - Assign \mathbf{B}_{n^*} as the estimated parameter values from this fit and let $\mathbf{L}_{n^*} = \mathbf{H} \mathbf{B}_{n^*}$.
- (b) Fit the Poisson GAM model described in Section 3.3 with $\mathbf{L} = \mathbf{L}_{n^*}$ as a fixed covariate.
 - Assign \mathbf{F}_{n^*} as the estimated parameter values from this fit.
- (c) Apply an SVD to find $\mathbf{B}_{n^*} \mathbf{F}_{n^*}^T = \mathbf{U}_{n+1} \mathbf{D}_{n+1} \mathbf{V}_{n+1}^T$.
 - i. Assign $\mathbf{B}_{n+1} = \mathbf{U}_{n+1}^{(1:K)} \mathbf{D}_{n+1}^{(1:K)}$.
 - ii. Assign $\mathbf{F}_{n+1} = \mathbf{V}_{n+1}^{(1:K)}$.

- iii. Assign $\mathbf{L}_{n+1} = \mathbf{H}\mathbf{B}_{n+1}$.
 - (d) Let $\log \mathbf{M}_{n+1} = \mathbf{L}_{n+1}\mathbf{F}_{n+1}^T$.
3. Repeat the *update* steps recursively until convergence.

Convergence is reached when the relative change in \mathbf{M} is sufficiently small. After convergence we can recover $\log \hat{\mu}_t$ from the rows of the final estimate of $\log \mathbf{M}$. These values are then treated as fixed constants during estimation of η_t . We use conditional maximum likelihood to estimate the parameters $(\omega, \alpha, \beta, \dots)$ associated with a time series model for η_t . The recursion defined by (6-9) require initialization by choosing a value for η_1 ; the estimates are conditional on the chosen initialization.

We may always specify the joint distribution $P_{\mathbf{Y}}$ of the observations \mathbf{Y} as an iterated product of successive conditional distributions P_{Y_t} for Y_t given (Y_{t-1}, \dots, Y_1) as

$$P_{\mathbf{Y}}(y_T, y_{T-1}, \dots, y_2, y_1) = P_{Y_1}(y_1) \prod_{t=2}^T P_{Y_t}(y_t | y_{t-1}, \dots, y_1).$$

We follow the standard convention of fixing $P_{Y_1}(y_1) = 1$ in estimation. For large sample sizes the practical impact of this decision is negligible. We may therefore write the log likelihood function as the sum of iterated conditional log likelihood functions. The conditional distribution for the observations is assumed to be Poisson with mean $\lambda_t = \mu_t \eta_t$.

For uninterrupted observations over periods $1, \dots, T$, we define the log likelihood function as

$$\begin{aligned} \ell(\omega, \alpha, \beta, \dots | \widehat{\mathbf{M}}, \mathbf{Y}, \eta_1) &= \sum_{t=2}^T \ell_t(\omega, \alpha, \beta, \dots | y_t, y_{t-1}, \hat{\mu}_t, \hat{\mu}_{t-1}, \eta_{t-1}) \\ &= \sum_{t=2}^T (y_t \log \lambda_t - \lambda_t - \log y_t!) \\ &= \sum_{t=2}^T (y_t \log(\hat{\mu}_t \eta_t) - \hat{\mu}_t \eta_t - \log y_t!). \end{aligned} \tag{10}$$

This recursion requires an initial value for η_1 . For simplicity, we use its expected value, $\eta_1 = 1$. When there are gaps in the observation record, Equation (10) is calculated over every contiguous block of observations. This requires reinitialization of $\eta_t = 1$ at the beginning of each block. The log likelihood for the blocks are then added together to form the entire log likelihood. The maximum likelihood estimate is the argmax of this quantity, subject to the

constraints given in Section 3.4. Finally, η_t is estimated by the respective recursion given by Equations (6-9) with parameters replaced by their estimates, again with reinitialization of $\eta_t = 1$ at the beginning of each contiguous block of observations. Blocks were large enough in our application that the effect of reinitialization was negligible.

4. EMPIRICAL ANALYSIS

Using the data described in Section 2 we perform the following analysis: (a) we define various statistical goodness of fit metrics suitable for the proposed models; based on in-sample performance, these metrics are used to determine the number of factors K for use in the dynamic factor models. (b) We compare the out-of-sample forecast performance for the factor model in (3), the factor model with constraints in (4), and the factor model with constraints and smoothing splines in (5). These comparisons help ascertain the improvement from each refinement and validate the proposed selection methods for K . (c) For the latter factor model, we compare the out-of-sample forecast performance with the addition of the CIIR process, via use of the various time series models defined in Section 3.4. (d) We quantify the practical impact of these successive statistical improvements with a queuing application constructed to approximate ambulance operations.

4.1 Interpreting the Fitted Model

The mean number of calls was approximately 24 per hour for 2007 and 2008, and no increasing or decreasing linear trend in time was detected during this period. We partition the observations by year into two datasets referred to as *2007* and *2008*, respectively. Each year is first regarded as a *training set*, and each model is fit individually to each year. The opposite year is subsequently used as a *test set* to evaluate the out-of-sample performance of each fitted model. To account for missing days, we reinitialize the CIIR process η_t in the first period following each day of missing data. This was necessary at most five times per year including the first day of the year.

We found the factor model fit with constraints, smoothing splines, and $K = 4$ factors to be the most appropriate of the factor models considered. The estimated factors \mathbf{f}_k for 2008 are shown in Figure 2(a). Each of the four factors varies smoothly over the hours of the day

via use of smoothing splines. The first factor \mathbf{f}_1 is strictly positive and the least variable. It appears to capture the mean diurnal pattern. The factor \mathbf{f}_2 appears to isolate the dominant relative differences between weekdays and weekend days. The defining feature of \mathbf{f}_3 and \mathbf{f}_4 is the large increase late in the day, corresponding closely to the relative increase observed on Friday and Saturday evenings. However, \mathbf{f}_3 decreases in the morning, while \mathbf{f}_4 increases in the morning and decreases in the late afternoon. As K increases, additional factors become increasingly more variable over the hours of the day. Too many factors results in over-fitting the model as the extra factors capture noise.

The corresponding daily factor loadings \mathbf{L} for the first four weeks of 2008 are shown in Figure 2(b). The loadings ($\mathbf{L}_1 - 14.5$) are shown to simplify comparisons. The much higher loadings on \mathbf{f}_1 confirm its interpretation as capturing the mean. The peaks on Fridays coincide with Friday having the highest average number of calls, as seen in Figure 1. Weekdays get a positive loading on \mathbf{f}_2 , while weekend days get negative loading. Loadings on \mathbf{f}_3 are lowest on Sundays and Mondays and loadings on \mathbf{f}_4 are largest on Fridays and Saturdays. As K increases, the loadings on additional factors become increasingly close to zero. This partially mitigates the over-fitting described above. Factors with loadings close to zero have less impact on the fitted values $\hat{\mu}_t$. Nevertheless, they can still reduce out-of-sample forecast performance.

The daily factor loadings for all of 2008 are shown in Figure 2(c). The relative magnitude of each loading vector with respect to day-of-week is constant. This results from use of the constraint matrix $\mathbf{H}^{(1)}$ in (4). As the loadings vary over the days of the week, they also vary smoothly over the course of the year, via use of the constraint matrix $\mathbf{H}^{(2)}$ and the use of cyclic smoothing splines in estimation of $\mathbf{B}^{(2)}$ in (4). The loadings on \mathbf{f}_1 show how the expected number of calls per day varies over the year. The week to week variability in the other loadings influences how the days of the week change relative to each other over the year. Figure 3 shows the estimated intensity process $\hat{\mu}_i$ for every day in 2008, shaded by day-of-week. The curves vary smoothly over the hours of the day. The fit for each day of the week keeps the same relative shape, but it varies smoothly over the weeks of the year.

Section 3.4 described incorporating time series models to improve the short-term forecasts of a factor model. The models capture the observed serial dependence in the multiplicative

residuals from a fitted factor model; see Figure 5. Parameter estimates and approximate standard errors for the IntGARCH model are given in Table ???. A fitted factor model $\hat{\mu}_t$ using constraints, smoothing splines and $K = 4$, as well as the factor model including a fitted IntGARCH(1,1) model $\hat{\lambda}_t$ are also shown in Figure 6(a), with the observed call arrivals per hour for Weeks 8 and 9 of 2007. The $\hat{\lambda}_t$ process is mean reverting about the $\hat{\mu}_t$ process. They are typically close to each other, but when they differ by a larger amount, they tend to differ for several hours at a time. The corresponding fitted CIIR process $\hat{\eta}_t$ is shown in Figure 6(b). This clearly illustrates the dependence and persistence exhibited in Figure 6(a). The CIIR process ranges between $\pm 6\%$ during this period. With a mean of 24 calls per hour, this range corresponds to $\hat{\lambda}_t$ varying about $\hat{\mu}_t$ by about ± 1.5 expected calls per hour.

4.2 Goodness of Fit and Model Selection

To evaluate the fitted values and forecasts of the proposed models, three types of residuals are computed: multiplicative, Pearson and Anscombe. Their respective formulas for the Poisson distribution are given by

$$\hat{r}_{M,t} = \frac{y_t}{\hat{\lambda}_t} - 1, \quad \hat{r}_{P,t} = \frac{y_t - \hat{\lambda}_t}{\sqrt{\hat{\lambda}_t}}, \quad \hat{r}_{A,t} = \frac{\frac{3}{2}(y_t^{2/3} - \hat{\lambda}_t^{2/3})}{\hat{\lambda}_t^{1/6}}.$$

We refer to the root mean square error (RMSE) of each metric as RMSME, RMSPE, and RMSAE, respectively. The multiplicative residual is defined as before and is a natural choice given the definition for the CIIR. Since the variance of a Poisson random variable is equal to its mean, the Pearson residual is quite standard. However, the Pearson residual can be quite skewed for the Poisson distribution (cf. McCullagh and Nelder 1989, Section 2.4). The Anscombe residual is derived as a transformation that makes the distribution of the residuals as close to Gaussian as possible while suitably scaling to stabilize the variance. See Pierce and Schafer (1986) for further discussion of residuals for generalized linear models. While the three methods always yielded the same conclusion, we found use of the Anscombe residuals gave a more robust assessment of model accuracy and simplified paired comparisons between the residuals of competing models.

The three RMSE metrics were used for both in- and out-of-sample model comparisons. For in-sample comparisons of the factor models, we also computed the *deviance* of each fitted

model $\hat{\mu}_t$. As a goodness of fit metric, deviance is derived from the logarithm of a ratio of likelihoods. For a log likelihood function $\ell(\boldsymbol{\mu}|\mathbf{Y})$, it is defined as

$$-2\{\ell(\boldsymbol{\mu} = \hat{\boldsymbol{\mu}}|\mathbf{Y}) - \ell(\boldsymbol{\mu} = \mathbf{Y}|\mathbf{Y})\},$$

in general. For a fitted factor model, ignoring serial dependence, the deviance corresponding to a Poisson distribution is

$$2 \sum_{t=1}^n \{y_t \log(y_t/\hat{\mu}_t) - (y_t - \hat{\mu}_t)\},$$

in which the first term is zero if $y_t = 0$.

We compare the fitted models' relative reduction in deviance and RMSE as we increase the number of factors K . Figure 4 shows these results for factor models fit to 2007 data with constraints and smoothing splines. The results for other models and for 2008 were very similar. This plot may be interpreted similarly to a *scree* plot in PCA by identifying the point at which performance tapers off and the marginal improvement from additional factors is negligible. Under each scenario we consistently selected $K = 4$ factors through this graphical criterion. To further justify this as a factor selection strategy, we also consider the impact the number of factors K has on out-of-sample performance for each of the proposed models below. This approach is straightforward, but it does not fully account for the uncertainty on the number of factors. Bayesian estimation would require specialized computation, but it may improve model assessment (see, e.g., Lopes and West 2004).

4.3 Out-of-sample Forecast Performance

Out-of-sample comparisons were made by fitting models to the 2007 training set and forecasting on the 2008 test set, and vice versa. To make predictions comparable from one year to the next, we align corresponding calendar weeks of the year, not days of the year. This ensures that estimates for Sundays are appropriately compared to Sundays, etc.

The first model considered was the *simple prediction* (SP) method. This simple moving average involving four observations was defined in the Introduction. Next, the forecasts of various factor models (FM) were considered. For $K = 1, \dots, 6$, we evaluated the forecasts from the FM in (3), the FM with constraints in (4), and the FM with constraints and smoothing

splines in (5). Finally, for the latter FM, with $K = 4$, we calculate the implied fit from the training set with the inclusion of the CIIR process via the various time series models defined in Section 3.4. We compute the forecast RMSE of each model for the three residual types, for both years.

The forecast results are shown in Table 1. The basic FMs did slightly worse than the SP both years. With only one year of observations, these FMs tend to over-fit the training set data, even with a small number of factors. The FMs with constraints give a very significant improvement over the previous models. The forecast RMSE is lowest at $K = 4$ for the 2007 test set, and at $K = 3$ for the 2008 test set. There was also a very large decrease between $K = 1$ and $K = 2$. The FMs with constraints and smoothing splines offered an additional improvement. The forecast RMSE is lowest at $K = 4$ for both test sets. With the addition of the IntGARCH model for the CIIR process to this model, the RMSE improved again. Application of the nonlinear time series models instead offered only a slight improvement over the IntGARCH model.

With only one year of training data, each FM begins to over-fit with $K = 5$ factors. Results were largely consistent regardless of the residual used, but the Anscombe residuals were the least skewed and allowed the simplest pairwise comparisons. Although the FMs with constraints had superior in-sample performance, the use of smoothing splines reduced the tendency to over-fit and resulted in improved forecast performance. The CIIR process offered improvements in fit over FMs alone.

We also fit each of the nonlinear time series models discussed in Section 3.4 using a FM with $K = 4$. The regime switching model had the best performance. It had the lowest RMSE for both test sets. The exponential autoregressive and the piecewise linear threshold models performed similarly to the IntGARCH model for both test sets. Although the nonlinear models consistently performed better in-sample, their out-of-sample performance was similar to the IntGARCH model.

4.4 Queueing Model Simulation to Approximate Ambulance Operations

To comprehensively improve ambulance operations it would be advantageous to simultaneously model the service duration of dispatched ambulances in addition to the demand for ambulance

service. Unfortunately, such information was not available. We are currently working with Toronto EMS to use our improved estimates of call arrival rates to improve staffing in their dispatch call center. Extending our approach to a spatial-temporal forecasting model will likely be used to help determine when *and* where to deploy ambulances.

We present a simulation study that uses a simple queueing system to quantify the impact that improved forecasts have on staffing decisions and relative operating costs, for the Toronto data. The queueing model is a simplification of ambulance operations that ignores the spatial component. Similar queueing models have been used frequently in EMS modeling (see Swersey 1994, p. 173). This goodness of fit measure facilitates model comparisons and a similar approach may be useful in other contexts.

We use the terminology employed in the call center and queueing theory literature throughout the section; for our application, servers are a proxy for ambulances, callers or customers are those requiring EMS, a server completing service is equated to an ambulance completing transport of a person to a hospital, etc. As before, let y_t denote the observed number of call arrivals during hour t . Our experiment examines the behavior of a simple $M/M/s$ queueing system. The arrival rate in time period t is λ_t . During this period, let s_t denote the number of servers at hand. For simplicity, we assume that the service rate ν for each server is the same, and constant over time. Furthermore, intra-hour arrivals occur according to a Poisson process with rate λ_t , and service times of callers are independent and exponentially distributed with rate ν .

As in Section 4.3, models are calibrated on one year of observations and forecasts for λ_t are made for the other year. Each model's forecasts $\hat{\lambda}_t$ are then used to determine corresponding staffing levels \hat{s}_t for the system.

To facilitate comparisons of short-term forecasts we assume that the number of servers can be changed instantaneously at the beginning of each period. In practice it is possible to adjust the number of ambulances in real time, but not to the degree that we assume here.

Each call has an associated arrival time and service time. When a call arrives, the caller goes immediately into service if a server is available, otherwise it is added to the end of the queue. A common goal in EMS is to ensure that a certain proportion of calls are reached by

an ambulance within a pre-specified amount of time. We approximate this goal by instead aiming to answer a proportion, θ , of calls immediately; this is a standard approximation in queueing applications in many areas including EMS (Kolesar and Green 1998). For each call arrival, we note whether or not the caller was served immediately. As servers complete service, they immediately begin serving the first caller waiting in the queue, otherwise they await new arrivals if the queue is currently empty. One simulation replication of the queueing system simulates all calls in the test year.

To implement the queueing system simulation, it is first necessary to simulate arrival and service times for each caller in the forecast period. We use the observed number of calls for each hour y_t as the number of arrivals to the system in period t . Since arrivals to the system are assumed to follow a Poisson process, we determine the y_t call arrival times using the well known result that, conditional on the total number of arrivals in the period $[t, t + 1]$, the arrival times have the same distribution as the order statistics of y_t independent $\text{Uniform}(t, t + 1)$ random variables. We exploit this relationship to generate the intra-hour arrival times given the observed arrival volume y_t . The service times for each call are generated independently with an $\text{Exponential}(\nu)$ distribution.

The final input is the initial state of the queue within the system. We generate an initial number of callers in the queue as $\text{Poisson}(y_1)$, then independently generate corresponding $\text{Exponential}(\nu)$ residual service times for each of these callers. This initialization is motivated through an infinite-server model; see, e.g., Kolesar and Green (1998). Whenever there is a missing day, in either the test set or corresponding training set period, we similarly reinitialize the state of the queue but with y_1 replaced by the number of calls observed in the first period following the missing period. These initializations are common across the different forecasting methods to allow direct comparisons.

To evaluate forecast performance, we define a cost function and an appropriate method for determining server levels from arrival rate estimates. Let n_t denote the number of callers served immediately in period t . The hourly cost function is given by

$$C(n_t, y_t, s_t) = \text{Pen}(n_t, y_t) + s_t,$$

in which

$$\text{Pen}(n_t, y_t) = \begin{cases} 0 & \text{if } n_t \geq \theta y_t, \\ q(y_t - n_t) & \text{otherwise,} \end{cases}$$

$\theta \in (0, 1)$ is the targeted proportion of calls served immediately, and $q \geq 0$ is the cost of not immediately serving a customer, *relative* to the cost of staffing one server for one hour. The total cost, with respect to the hourly server cost, for the entire forecast period is

$$C = \sum_t C(n_t, y_t, s_t) = \sum_t \text{Pen}(n_t, y_t) + \sum_t s_t.$$

This approach, where penalties for poor service are balanced against staffing costs, is frequently used; see, e.g., Andrews and Parsons (1993), Harrison et al. (2005).

At time $t - 1$, the number of call arrivals and the number served immediately are random variables, denoted as Y_t and N_t , respectively. A natural objective is to choose staffing levels that minimize the hourly expected cost as

$$\hat{s}_t = \underset{s_t \in \mathbb{N}}{\text{argmin}} E\{C(N_t, Y_t, s_t) | \mathcal{F}_{t-1}, \mathbf{X}\}, \quad (11)$$

in which Y_t is assumed to have a Poisson distribution with mean equal to the arrival rate forecast $\hat{\lambda}_t$. The staffing levels are then a function of arrival rate forecasts, $\hat{s}_t(\hat{\lambda}_t)$. We approximate this expectation numerically by randomly generating J independent realizations as $Y_{t,j} \sim \text{Poisson}(\hat{\lambda}_t)$. Then, for each $Y_{t,j}$ we simulate one independent realizations of N_t . For a fixed value of s_t the expectation is approximated by $J^{-1} \sum_{t=1}^J \{\text{Pen}(N_{t,j}, Y_{t,j}) + s_t\}$. We found that $J = 25,000$ provided adequate accuracy.

Independent realizations of $N_t | Y_t$ require running the queueing system forward one hour, but this is very computationally intensive. To approximate $N_t | Y_t$ we use a Binomial distribution. Let $N_{t,j} | Y_{t,j} \sim \text{Binomial}\{Y_{t,j}, g(\hat{\lambda}_t, s_t, \nu)\}$. The function g gives the *steady state* probability that a customer is served immediately for a queueing system with a *constant* arrival rate, server level, and service rate, $\hat{\lambda}_t, s_t$, and ν , respectively. Derivation of this function is available in any standard text on queueing theory (e.g., Gross and Harris 1998, Chapter 2).

Let p_i denote the long run proportion of time such a system contains i customers and let $\rho = \lambda/(\nu s)$. Then

$$g(\lambda, s, \nu) = \begin{cases} 1 - \frac{\lambda^s p_0}{s! \nu^s (1-\rho)} & \text{if } \rho < 1, \\ 0 & \text{if } \rho \geq 1. \end{cases}, \quad \text{in which } p_0^{-1} = \sum_{u=0}^{c-1} \frac{r^u}{u!} + \frac{r^c}{c!(1-\rho)} \quad \text{for } \rho < 1.$$

When $\rho \geq 1$, the arrival rate is faster than the net service rate, and the system is unstable; the long run probability that a customer is served immediately is zero. The binomial approximation greatly reduces the computational costs and provides reasonable results, though it tends to underestimate the true variability of $N_t|Y_t$ due to the positive correlation in successive caller delays.

A final deliberation is needed on the removal of servers when \hat{s}_t decreases. In our implementation, idle servers were removed first, and if necessary, busy servers were dropped in ascending order with respect to remaining service time. We also considered random selection of servers to be dropped. Doing so produced highly variable results, and is under further study. To further simplify the implementation, if it was necessary to drop a busy server, it was simply discarded, along with any remaining service time for that caller. The effect of this simplification depends on the service rate ν ; our results did not appear to be sensitive to this simplification.

Simulation of the queueing system is now rather straightforward. On each iteration i , we note whether each caller was served immediately or not. Forecast performance is assessed by examining the total cost $C^{(i)} = \sum_t C(n_t^{(i)}, y_t, \hat{s}_t)$ over the test period. For both years, we performed 100 simulations over the test year for each forecast method. To demonstrate the robustness of this methodology, we performed the experiment for several different values of the queueing system's parameters. Specifically, all combinations of $q \in \{2, 5, 10\}$, $\nu \in \{1, \frac{2}{3}\}$, and $\theta \in \{.8, .9\}$ were considered, after consultation with EMS experts.

Results for the mean hourly cost over the 100 simulations for each forecasting method, for each test year, are summarized in Figure 7. We see that the mean hourly cost is lowest for the FM w/ IntGARCH, followed by the FM only, and finally by SP. All pairwise differences in mean were highly significant; the smallest t-ratio was 80. In fact, this ordering in performance held for almost every iteration of the queueing system, not just on average.

The mean percentage of callers served immediately can be found in Figure 8. The total number of server hours $\sum \hat{s}_t$ used was also recorded for each model for each set of parameter values. A table containing the values of all these quantities can be found in the online supplemental material. Both mean percentage served immediately and mean hourly cost increase

with q . For each test year, for each level of (q, ν, θ) , $\sum_t \hat{s}_t$ differed by between one and three thousand server-hours, for the different models.

5. CONCLUSIONS

Our analysis was motivated by a dataset provided by Toronto EMS. The proposed forecasting method allows parsimonious modeling of the dependent and nonstationary count-valued EMS call arrival process. Our method is straightforward to implement and demonstrates substantial improvements in forecast performance relative to simpler forecasting methods. We measured the impact of our successive refinements to the model, showing the merit of factor model estimation with covariates and smoothing splines. The factor model was able to capture the nonstationary behavior exhibited in call arrivals. Introduction of the CIIR process allowed adaptive forecasts of deviations from this diurnal pattern.

Assessing the impact that different arrival rate forecasts can have on call centers and related applications has received very little attention in the literature. Our data-based simulation approach is straightforward to implement, and was able to clearly distinguish the effectiveness of each forecasting method. The simulation results coincide with the out-of-sample RMSE analysis in Section 4.3 and provide a practical measure of forecast performance. Relative operating cost is a natural metric for measuring call arrival rate forecasts, and our implementation may easily be extended to many customized cost functions and a wide variety of applications.

Ultimately, we seek to strengthen emergency medical service by improving upon relevant statistical methodology. Future work will consider inclusion of additional covariates and study of other nonlinear time series models. Bayesian methods which directly model count-valued observations have desirable properties for inference and many applications, and are under study. Spatial and spatial-temporal analysis of call arrivals will also offer new benefits to EMS.

ACKNOWLEDGMENTS

This work was partially supported by National Science Foundation Grant Number CMMI-0926814. The authors sincerely thank Toronto EMS for sharing their data; in particular, Mr. Dave Lyons for his comments and support.

SUPPLEMENTAL MATERIALS

Additional Tables: Tables 2 & 3 (Tables23.pdf).

Estimation Algorithms: *R* code for estimating the models in Section 3 and for calculating the RMSE metrics in Section 4 (ModelFitting.R).

Simulation Algorithms: *R* code for implementing the queueing model simulation in Section 4.4 (QueueSimulations.R).

REFERENCES

- Andrews, B. and Cunningham, S. (1995), “LL Bean Improves Call-center Forecasting,” *Interfaces*, 25, 1–13.
- Andrews, B. and Parsons, H. (1993), “Establishing Telephone-agent Staffing Levels Through Economic Optimization,” *Interfaces*, 23, 14–20.
- Bianchi, L., Jarrett, J., and Choudary Hanumara, R. (1998), “Improving Forecasting for Telemarketing Centers by ARIMA Modeling with Intervention,” *International Journal of Forecasting*, 14, 497–504.
- Borchers, D., Buckland, S., Priede, I., and Ahmadi, S. (1997), “Improving the Precision of the Daily Egg Production Method Using Generalized Additive Models,” *Canadian Journal of Fisheries and Aquatic Sciences*, 54, 2727–2742.
- Brown, L., Gans, N., Mandelbaum, A., Sakov, A., Shen, H., Zeltyn, S., and Zhao, L. (2005), “Statistical Analysis of a Telephone Call Center,” *Journal of the American Statistical Association*, 100, 36–50.
- Channouf, N., L’Ecuyer, P., Ingolfsson, A., and Avramidis, A. (2007), “The Application of Forecasting Techniques to Modeling Emergency Medical System Calls in Calgary, Alberta,” *Health Care Management Science*, 10, 25–45.
- Cox, D. and Isham, V. (1980), *Point Processes*, New York: Chapman & Hall/CRC.
- Daskalov, G. (1999), “Relating Fish Recruitment to Stock Biomass and Physical Environment in the Black Sea Using Generalized Additive Models,” *Fisheries Research*, 41, 1–23.
- Ferland, R., Latour, A., and Oraichi, D. (2006), “Integer-valued GARCH Process,” *Journal of Time Series Analysis*, 27, 923–942.
- Fokianos, K., Rahbek, A., and Tjøstheim, D. (2009), “Poisson Autoregression,” *Journal of the American Statistical Association*, 104, 1430–1439.
- Goldberg, J. B. (2004), “Operations Research Models for the Deployment of Emergency Services Vehicles,” *EMS Management Journal*, 1, 20–39.
- Gross, D. and Harris, C. (1998), *Fundamentals of Queueing Theory*, New York: Wiley, 3rd ed.
- Gu, C. (1992), “Cross-Validating Non-Gaussian Data,” *Journal of Computational and Graphical Statistics*, 1, 169–179.
- (2010), *gss: General Smoothing Splines*, R Package Version 1.1-3.
- Harrison, J., Zeevi, A., and Shum, S. (2005), “A Method for Staffing Large Call Centers Based on Stochastic Fluid Models,” *Manufacturing & Service Operations Management*, 7, 20–36.
- Hastie, T. (2009), *gam: Generalized Additive Models*, R Package Version 1.01.
- Hastie, T. and Tibshirani, R. (1990), *Generalized Additive Models*, London, UK: Chapman and Hall.
- (1993), “Varying-coefficient Models,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 757–796.

- Henderson, S. G. (2005), “Should We Model Dependence and Nonstationarity, and if so How?” in *Proceedings of the 37th Conference on Winter Simulation*, eds. Kuhl, M., Steiger, N., Armstrong, F., and Joines, J., Winter Simulation Conference, Piscataway, NJ: IEEE, pp. 120–129.
- (2009), *Operations Research Tools for Addressing Current Challenges in Emergency Medical Services*, New York: Wiley, in J. J. Cochran, editor, Wiley Encyclopedia of Operations Research and Management Science.
- Kolesar, P. and Green, L. (1998), “Insights on Service System Design from a Normal Approximation to Erlang’s Delay Formula,” *Production and Operations Management*, 7, 282–293.
- Lopes, H. F. and West, M. (2004), “Bayesian Model Assessment in Factor Analysis,” *Statistica Sinica*, 14, 41–68.
- Matteson, D. S. and Tsay, R. S. (2010), “Constrained Independent Component Analysis,” Under Revision.
- McCullagh, P. and Nelder, J. A. (1989), *Generalized Linear Models*, London: Chapman and Hall, 2nd ed.
- Opsomer, J., Yang, Y., and Yang, Y. (2001), “Nonparametric Regression with Correlated Errors,” *Statistical Science*, 16, 134–153.
- Pierce, D. and Schafer, D. (1986), “Residuals in Generalized Linear Models,” *Journal of the American Statistical Association*, 81, 977–986.
- R Development Core Team (2009), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.
- Setzler, H., Saydam, C., and Park, S. (2009), “EMS Call Volume Predictions: A Comparative Study,” *Computers & Operations Research*, 36, 1843–1851.
- Shen, H. and Huang, J. (2008a), “Forecasting Time Series of Inhomogeneous Poisson Processes with Application to Call Center Workforce Management,” *The Annals of Applied Statistics*, 2, 601–623.
- (2008b), “Interday Forecasting and Intraday Updating of Call Center Arrivals,” *Manufacturing and Service Operations Management*, 10, 391–410.
- Stock, J. and Watson, M. (2002), “Macroeconomic Forecasting Using Diffusion Indexes,” *Journal of Business & Economic Statistics*, 20, 147–162.
- Swersey, A. (1994), “The Deployment of Police, Fire, and Emergency Medical Units,” *Handbooks in Operations Research and Management Science*, 151–200.
- Takane, Y. and Hunter, M. (2001), “Constrained Principal Component Analysis: A Comprehensive Theory,” *Applicable Algebra in Engineering, Communication and Computing*, 12, 391–419.
- Tsai, H. and Tsay, R. S. (2009), “Constrained Factor Models,” Under Revision.
- Tych, W., Pedregal, D., Young, P., and Davies, J. (2002), “An Unobserved Component Model for Multi-rate Forecasting of Telephone Call Demand: The Design of a Forecasting Support System,” *International Journal of Forecasting*, 18, 673–695.
- Weinberg, J., Brown, L., and Stroud, J. (2007), “Bayesian Forecasting of an Inhomogeneous Poisson Process with Applications to Call Center Data,” *Journal of the American Statistical Association*, 102, 1185–1198.

- Whitt, W. (2002), *Stochastic-Process Limits*, Springer Series in Operations Research, New York: Springer.
- Wood, S. (2003), “Thin Plate Regression Splines,” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 65, 95–114.
- (2006), *Generalized Additive Models: An Introduction with R*, New York: CRC Press.
- (2008), *gss: General Smoothing Splines*, R Package Version 1.6-1.
- Woodard, D. B., Matteson, D. S., and Henderson, S. G. (2010), “Stationarity of Count-valued and Nonlinear Time Series Models,” Submitted.

Model	Constraint	Smoothing	2007 model, 2008 residuals			2008 model, 2007 residuals		
			RMSME	RMSPE	RMSAE	RMSME	RMSPE	RMSAE
Simple Prediction	NA	NA	0.2696	1.1955	1.1849	0.2661	1.1902	1.1925
Factor Model, $K = 1$	No	No	0.2722	1.2369	1.2237	0.2657	1.2183	1.2263
Factor Model, $K = 2$	No	No	0.2721	1.2357	1.2225	0.2661	1.2197	1.2277
Factor Model, $K = 3$	No	No	0.2727	1.2374	1.2239	0.2659	1.2182	1.2262
Factor Model, $K = 4$	No	No	0.2729	1.2383	1.2249	0.2666	1.2206	1.2283
Factor Model, $K = 5$	No	No	0.2732	1.2395	1.2260	0.2670	1.2220	1.2294
Factor Model, $K = 6$	No	No	0.2733	1.2401	1.2270	0.2668	1.2217	1.2294
Factor Model, $K = 1$	Yes	No	0.2638	1.1863	1.1756	0.2575	1.1633	1.1721
Factor Model, $K = 2$	Yes	No	0.2402	1.0938	1.0888	0.2333	1.0722	1.0875
Factor Model, $K = 3$	Yes	No	0.2392	1.0877	1.0829	0.2324	1.0688	1.0848
Factor Model, $K = 4$	Yes	No	0.2413	1.0945	1.0889	0.2347	1.0761	1.0912
Factor Model, $K = 5$	Yes	No	0.2425	1.0994	1.0933	0.2363	1.0817	1.0961
Factor Model, $K = 6$	Yes	No	0.2436	1.1051	1.0988	0.2377	1.0858	1.0999
Factor Model, $K = 1$	Yes	Yes	0.2633	1.1837	1.1731	0.2573	1.1615	1.1703
Factor Model, $K = 2$	Yes	Yes	0.2371	1.0844	1.0805	0.2310	1.0643	1.0803
Factor Model, $K = 3$	Yes	Yes	0.2347	1.0744	1.0710	0.2289	1.0561	1.0728
Factor Model, $K = 4$	Yes	Yes	0.2344	1.0730	1.0696	0.2288	1.0549	1.0715
Factor Model, $K = 5$	Yes	Yes	0.2347	1.0740	1.0706	0.2289	1.0549	1.0714
Factor Model, $K = 6$	Yes	Yes	0.2347	1.0739	1.0705	0.2289	1.0551	1.0716
Time Series & FM, $K = 4$	Yes	Yes	-	-	-	-	-	-
IntGARCH	-	-	0.2308	1.0571	1.0570	0.2274	1.0442	1.0580
IntExpGARCH	-	-	0.2308	1.0570	1.0569	0.2274	1.0441	1.0579
IntThreshGARCH	-	-	0.2308	1.0571	1.0570	0.2275	1.0443	1.0580
IntRsGARCH	-	-	0.2299	1.0540	1.0554	0.2274	1.0433	1.0565

Table 1: Root mean square multiplicative, Pearson, and Anscombe errors for fitting model to 2007 and forecasting 2008, and vice versa. A Yes in the constraints column implies that the factor model was fit using the constraints outlined in Section 3.2. A Yes in the smoothing column indicates that the model was fit using smoothing splines as described in Section 3.3.

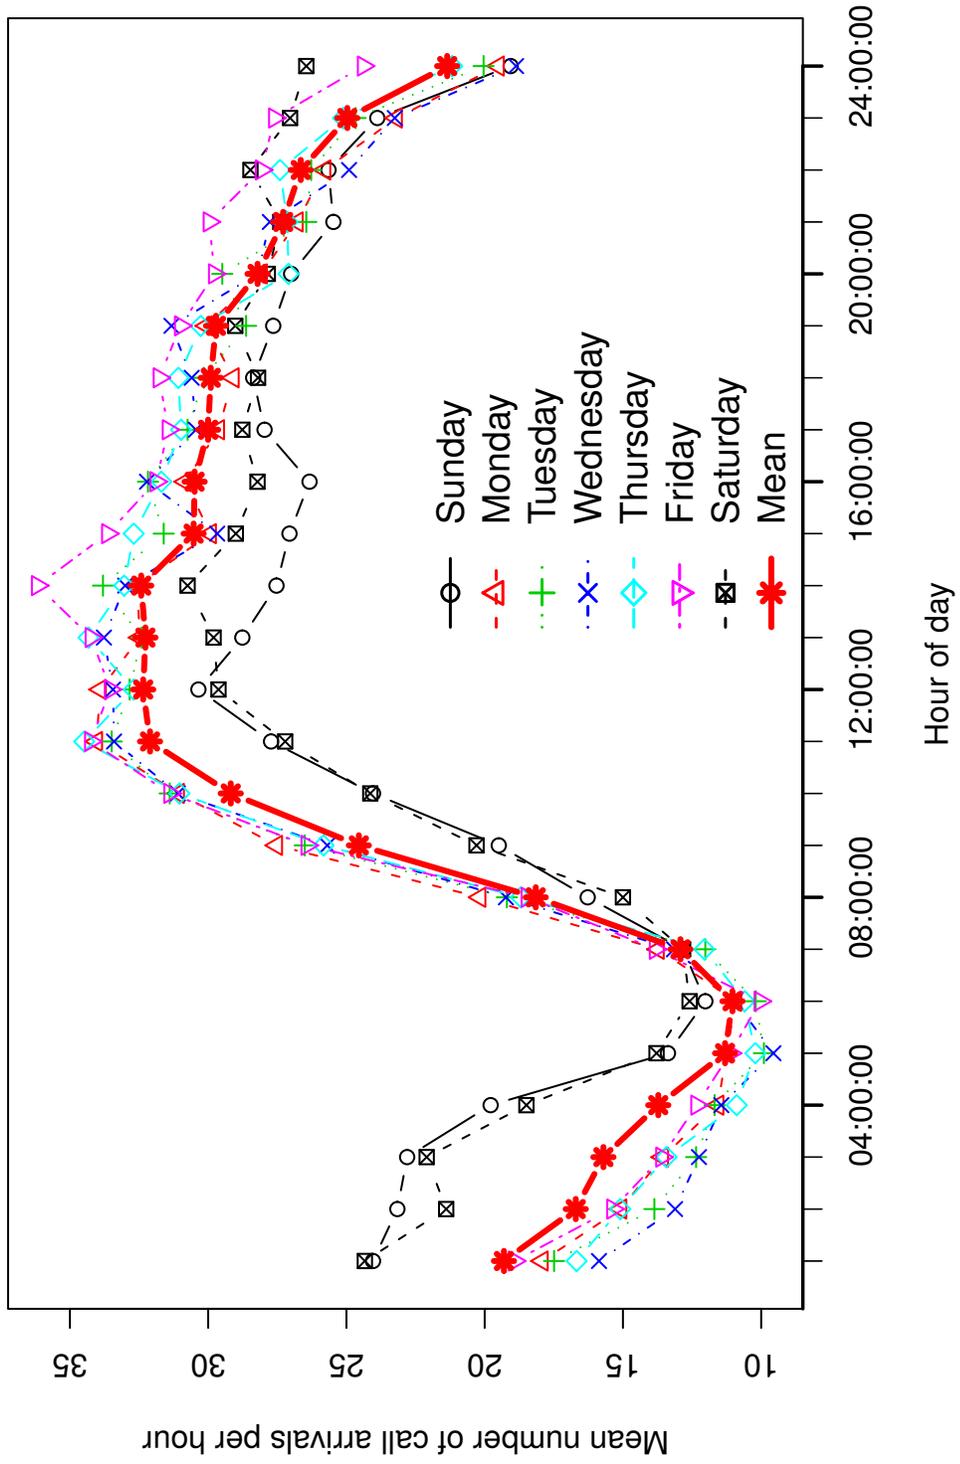


Figure 1: Mean number of calls per hour by day of the week.

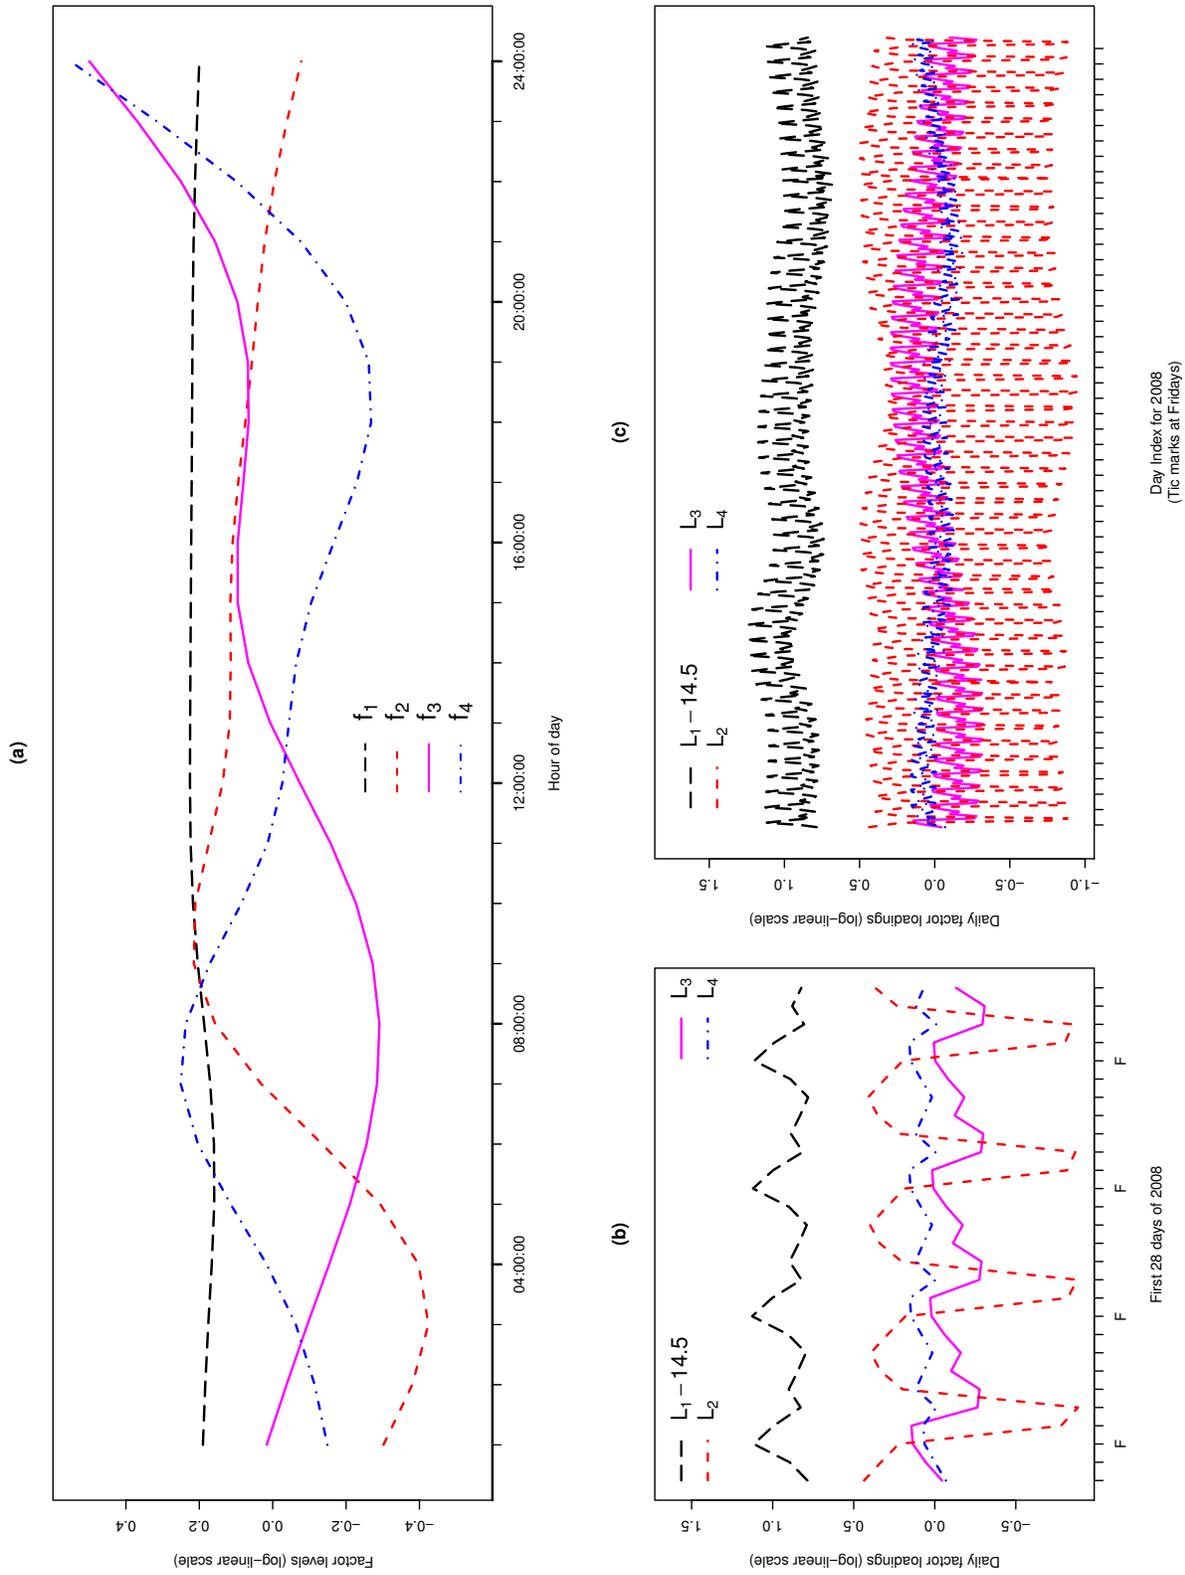


Figure 2: 2008 fitted (a) factor levels f_k (log-linear scale) and (b) corresponding factor loadings L_k . (log-linear scale) for a factor model fit with constraints, smoothing splines, and $K = 4$ factors. ($L_1, - 14.5$) is shown for easier comparison.

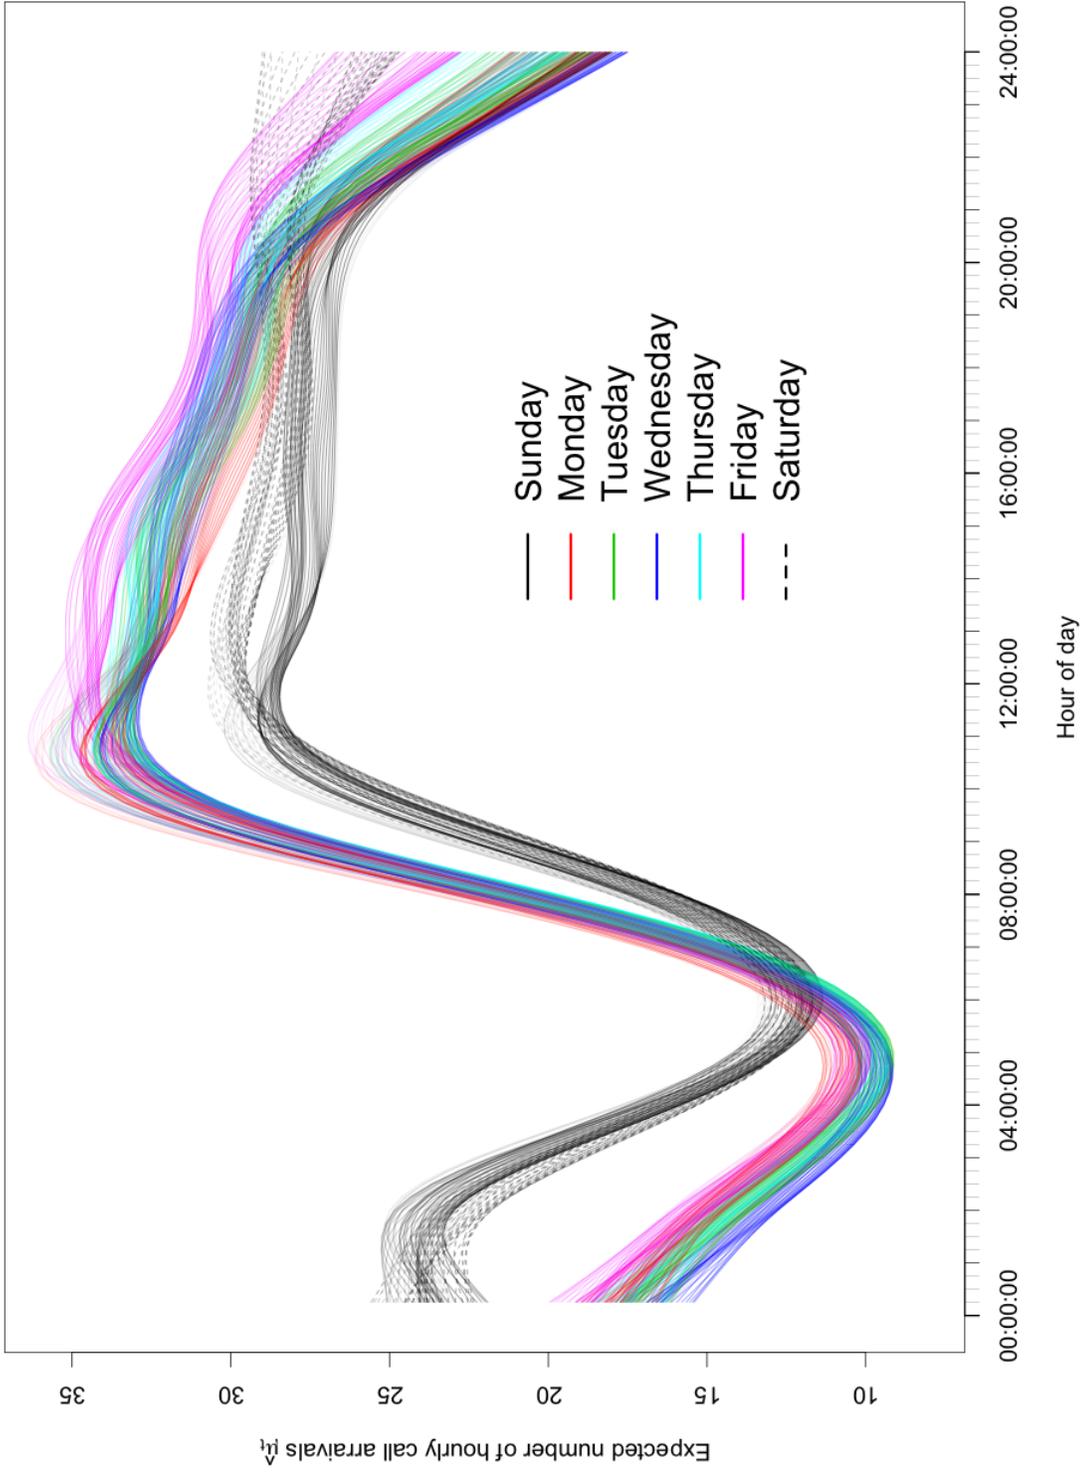


Figure 3: The estimated intensity process $\hat{\mu}_t$, for every day in 2008, for a factor model fit with constraints, smoothing splines, and $K = 4$ factors; colored by day-of-week; shaded light to dark by week-of-year.

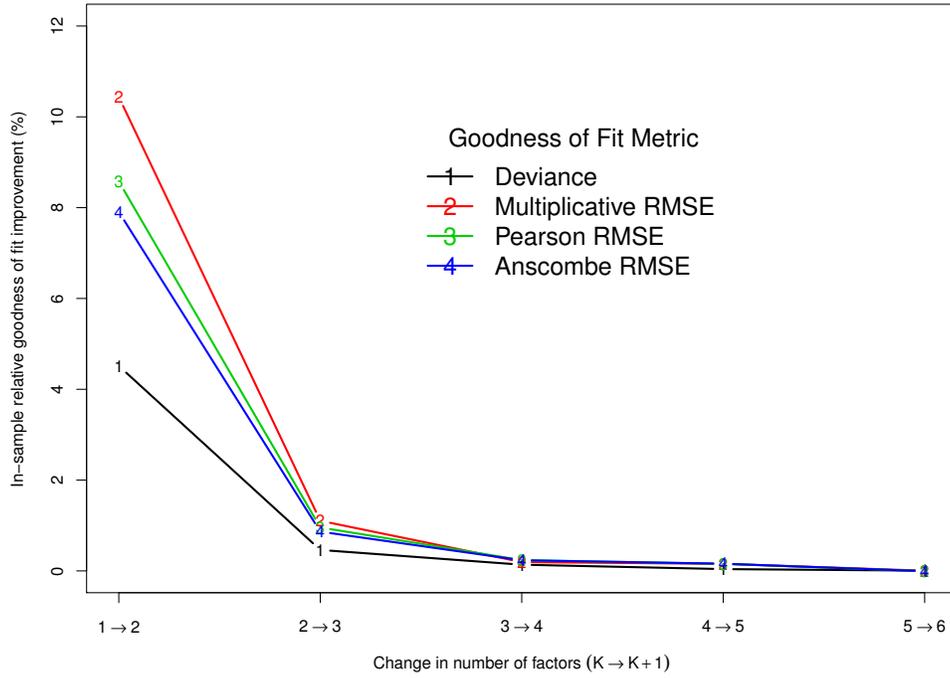


Figure 4: 2007 percentage in-sample relative goodness of fit improvement by addition of one factor ($K \rightarrow K + 1$) for a factor model fit with constraints and smoothing splines.

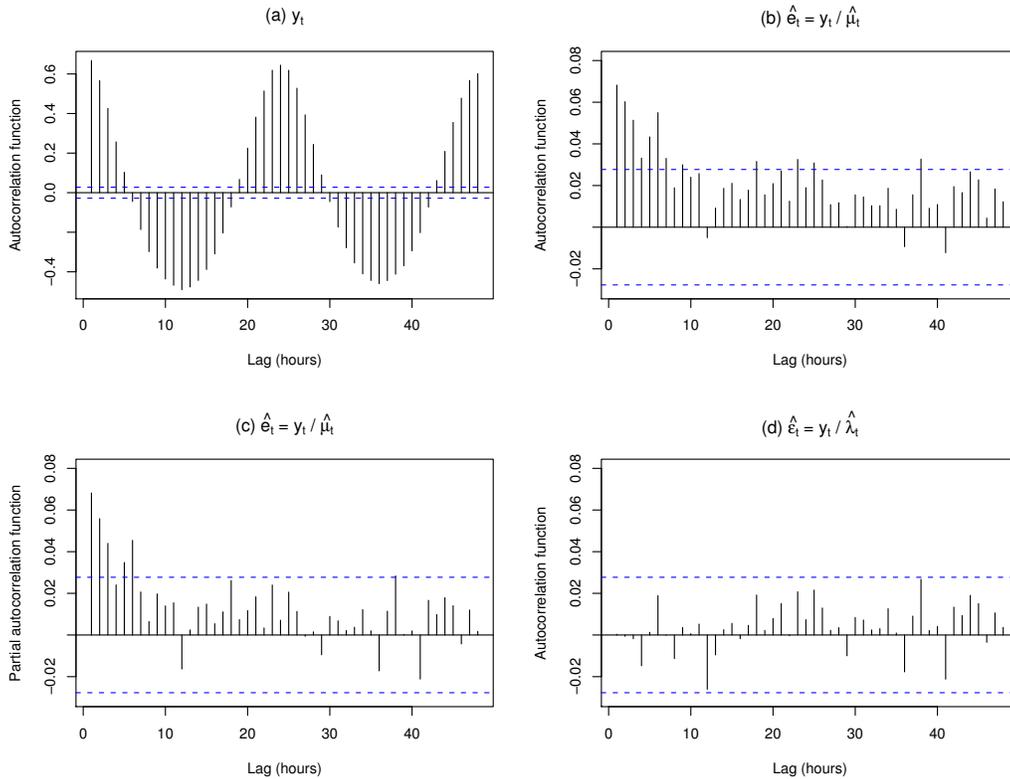


Figure 5: (a) Sample autocorrelation function for hourly call arrival counts y_t . Residual $\hat{\epsilon}_t = y_t / \hat{\mu}_t$ (b) autocorrelation and (c) partial autocorrelation functions for fitted factor model $\hat{\mu}_t$ with $k = 4$ factors using factor and loading constraints and smoothing splines. (d) Standardized residual $\hat{\epsilon}_t = y_t / \hat{\lambda}_t = y_t / (\hat{\mu}_t \hat{\eta}_t)$ autocorrelation function for fitted factor model with fitted IntGARCH(1,1) model for η_t . Dashed lines give approximate 95% confidence levels.

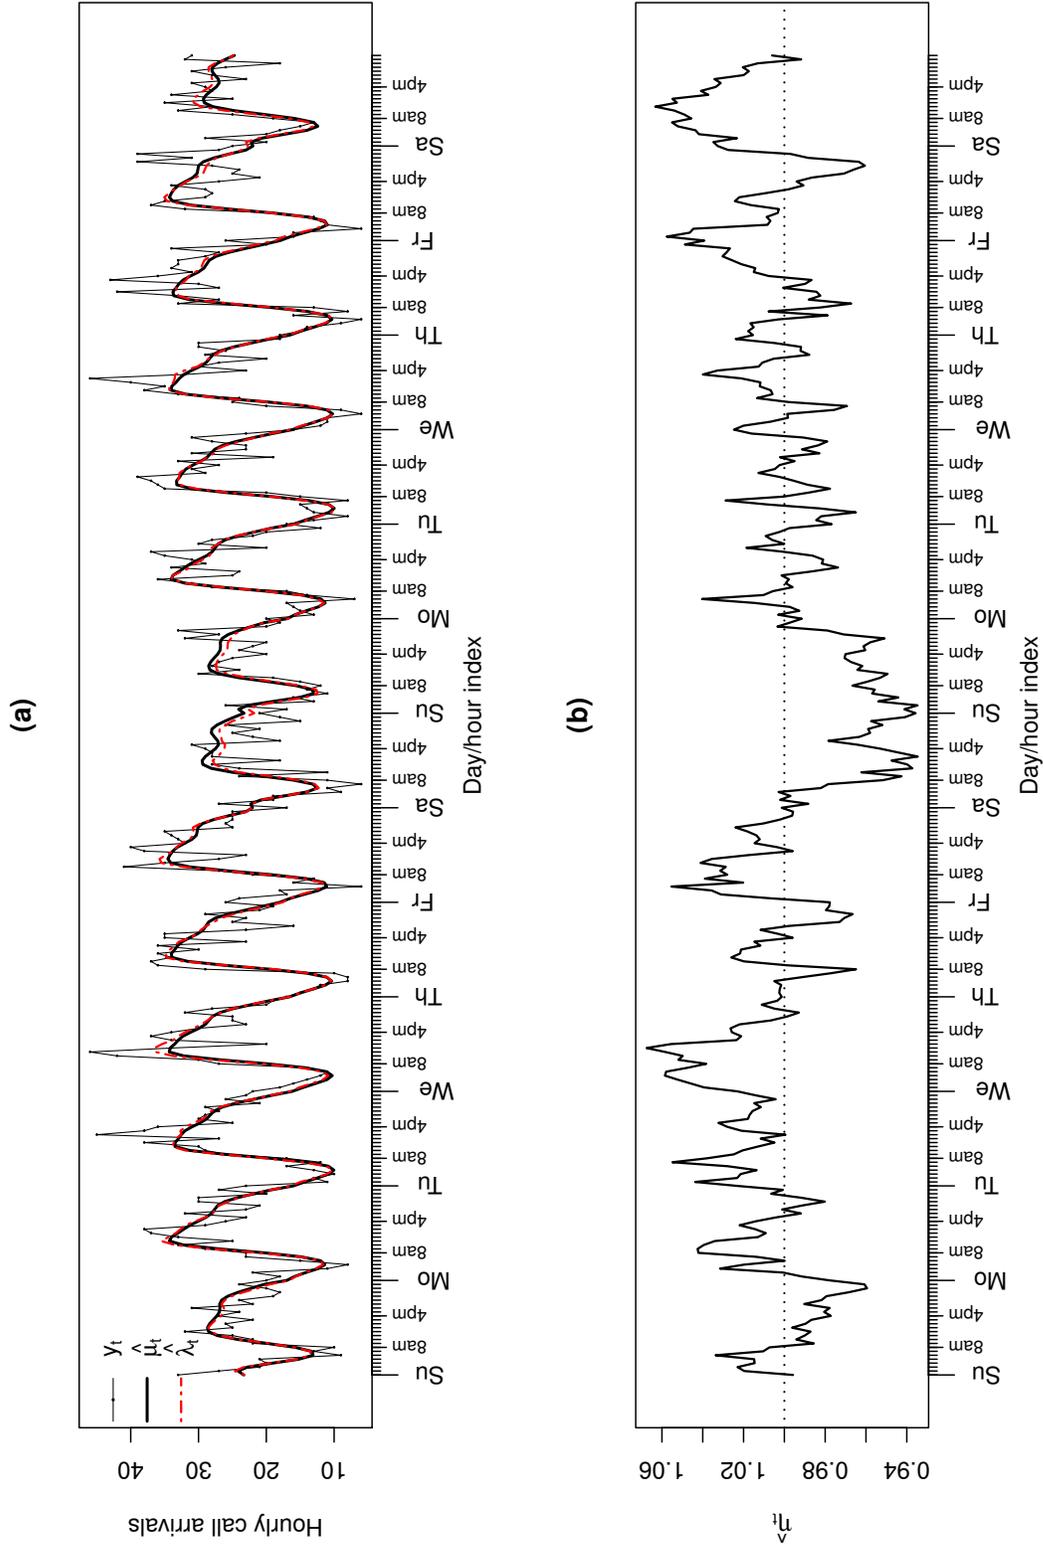


Figure 6: Weeks 8 & 9 of 2007: (a) observed call arrivals per hour y_t , fitted $K = 4$ dynamic factor model $\hat{\mu}_t$ using constraints and smoothing splines, and factor model $\hat{\lambda}_t$ including fitted IntGARCH(1,1); (b) the fitted conditional intensity inflation process $\hat{\eta}_t$ from the IntGARCH(1,1) model.

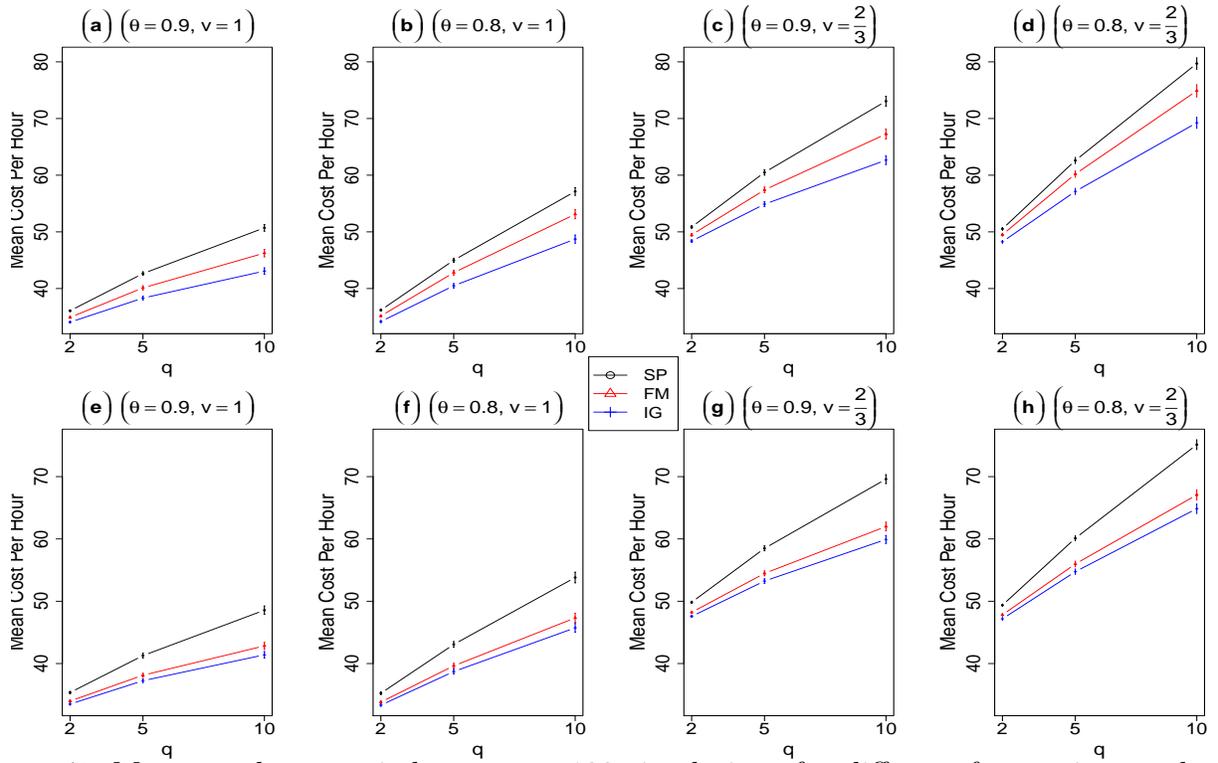


Figure 7: Mean total per period cost over 100 simulations for different forecasting methods and different values of q , ν , and θ . Plots (a)-(d) use the 2008 test set and plots (e)-(h) use 2007 as the test set. The vertical lines represent ± 1 standard deviation.

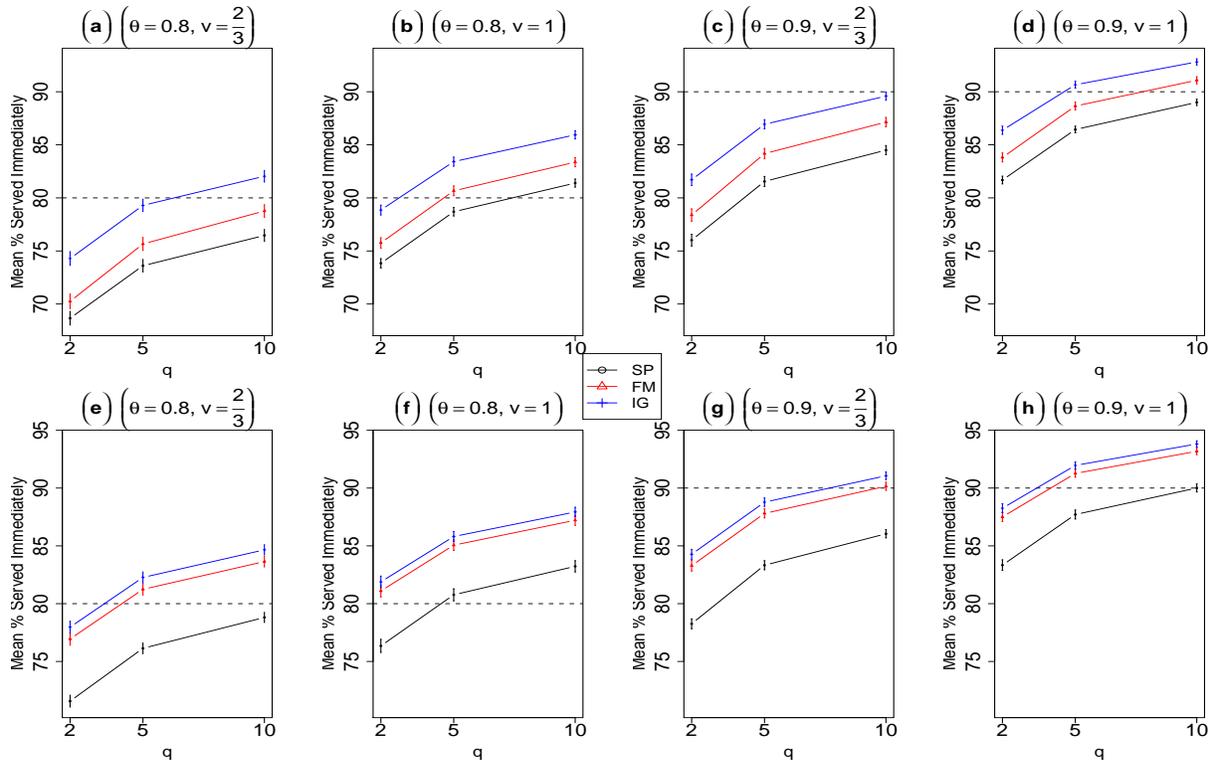


Figure 8: Mean percentage served immediately for the entire test set over 100 simulations for different forecasting methods and different values of q , ν , and θ . Plots (a)-(d) use the 2008 test set and plots (e)-(h) use 2007 as the test set. The vertical lines represent ± 1 standard deviation.