ARTICLE IN PRESS
JID: EOR [m5G;August 26, 2016;9:35]

European Journal of Operational Research 000 (2016) 1–12

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Development and implementation of a real time statistical control method to identify the start and end of the winter surge in demand for paediatric intensive care

Christina Pagel[a,*], Padmanabhan Ramnarayan[b], Samiran Ray[b], Mark J. Peters[b,c]

[a] UCL Clinical Operational Research Unit, University College London, 4 Taviton Street, London, WC1H 0BT, UK
[b] Children's Acute Transport Service (CATS), Great Ormond Street Hospital NHS Foundation Trust, 26-27, Boswell Street, London WC1N 3JZ, UK
[c] Respiratory, Anaesthesia and Critical Care Group, Institute of Child Health, 30 Guilford St, London WC1N 1EH, United Kingdom

A R T I C L E I N F O

Article history:
Received 30 December 2015
Accepted 9 August 2016
Available online xxx

Keywords:
OR in health services
Analytics
Statistical process control
Implementation
Emergency medical services

A B S T R A C T

Winter surge management in intensive care is hampered by the annual variability in the winter surge. We aimed to develop a real-time monitoring system that could promptly identify the start, and accurately predict the end, of the winter surge in a paediatric intensive care (PIC) setting. We adapted a statistical process control method from the stock market called "Bollinger bands" that compares current levels of demand for PIC services to thresholds based on the medium term average demand. Algorithms to identify the start and end of the surge were developed for a specific PIC service: the North Thames Children's Acute Transport Service (CATS) using eight winters of data (2005–12) to tune the algorithms and one winter to test the final method (2013/14). The optimal Bollinger band thresholds were 1.2 and 1 standard deviations above and below a 41-day moving average of demand respectively. A simple linear model was found to predict the end of the surge and overall demand volume as soon as the start had been identified. Applying the method to the validation winter of 2013/14 showed excellent performance, with the surge identified from 18th November 2013 to 4th January 2014.

An Excel tool running the algorithms has been in use within CATS since September 2014. There were three factors which facilitated the successful implementation of this tool: the perceived problem was pressing and identified by the clinical team; there was close clinical engagement throughout and substantial effort was made to develop an easy-to-use Excel tool for sustainable use.

1. Introduction

Clinicians from Great Ormond Street Hospital (GOSH) in the UK run the Children's Acute Transport Service (CATS) which is responsible for transporting very sick children from non-specialist hospitals to paediatric intensive care units (PICUs) in north London. They are staffed by intensive care doctors and nurses and have two emergency paediatric ambulances for the North Thames area. While the numbers of children transported using such specialist retrieval teams are relatively small, all such retrievals are potentially life-saving and represent an important and expensive NHS resource (Ramnarayan, 2009). If the two teams are out on a call and cannot meet demand, that child must be transferred to another transport service (if possible) or else wait for the next CATS team to be free, with a risk of further clinical deterioration while waiting. If there is no local specialist PICU bed available, teams will need to transport the child further afield to the nearest hospital that has capacity (which can be as far away as the very north of England), which further impacts the service as that team is then unavailable for a considerable amount of time. Thus, when demand stretches capacity there is a risk of a worse clinical outcome for the children waiting longer for transport to a specialist unit and significant experienced stress for the CATS team.

In the UK, every winter brings with it an increase in the number of emergency admissions (particularly patients with respiratory disease). The British Medical Association ("BMA – Winter pressures") and most hospitals (e.g. "Winter preparedness 2013–14") have preparedness plans in place to try to cope with this increase. However, these plans tend to focus on what do when pressures arise and less on whether we can forecast the start of the annual surge (Hanratty and Robinson, 1999). The pressure on children's services is often particularly acute since PICU beds are not an abundant resource and young children, particularly the most vulnerable ones (for instance those with cystic fibrosis), are

* Corresponding author.
E-mail address: c.pagel@ucl.ac.uk (C. Pagel).

http://dx.doi.org/10.1016/j.ejor.2016.08.023
0377-2217/

Please cite this article as: C. Pagel et al., Development and implementation of a real time statistical control method to identify the start and end of the winter surge in demand for paediatric intensive care, European Journal of Operational Research (2016), http://dx.doi.org/10.1016/j.ejor.2016.08.023

proportionately more susceptible to respiratory illnesses than adults (O'Donnell, Parslow & Draper, 2010). Every winter, the CATS service has been severely stretched by increasing demand with no accompanying increase in capacity. However, there is now the possibility for CATS to temporarily increase its capacity during the winter surge through staffing additional retrieval teams. There is also potential for triggering different procedures for elective admissions to paediatric intensive care units when there is indication that the annual winter peak is beginning (for instance preemptively re-scheduling some operations). Currently, the UK Paediatric Intensive Care Society (PICS) identifies the winter surge as starting in mid-November and ending the first week of January and discusses various operational ways services can try to cope with the extra demand, including regular regional and national conference calls sharing information on experienced demand.

Previous approaches to the winter surge in emergency demand in adult intensive care have involved using disease surveillance (Hiller, Stoneking, Min, & Rhodes, 2013; Moriña, Puig, Ríos, Vilella, & Trilla, 2011; Nguyen et al., 2016), and/or weather and seasonal information (Batal, Tench, McMillan, Adams, & Mehler, 2001; Boyle et al., 2012; Diehl, Morris, & Mannis, 1981; Jones, Joy, & Pearson, 2002; Marcilio, Hajat, & Gouveia, 2013; Shiue, Perkins, & Bearman, 2016), and/or previous demand (Abraham, Byrnes, & Bain, 2009; Jones, 2007; Jones et al., 2008; Proudlove, Gordon, & Boaden, 2003), using a range of techniques including regression, stochastic Markov models and time series analysis methods such as Autoregressive Integrated Moving Average (ARIMA) models. In general, while emergency demand was universally found to be strongly seasonal and autoregressive it was also extremely variable, with its stochastic nature making accurate forecasts beyond seasonal or monthly means difficult. Additional problems for paediatric emergency retrievals at CATS is that vulnerable children tend to be the first population cohort to fall ill every winter. This means that the peak in paediatric intensive care is often a month or two earlier than in adults meaning that using sentinel disease methods are less useful. At the same time case volume and resources are much lower in this context, so that the stochastic nature of the demand tends to have even greater influence on experienced daily demand.

The primary question we address in this paper is thus: can we find and implement a simple, feasible, method to improve on the PICS identification of the winter surge and so help the CATS team with their winter planning? Real time identification of the surge could also be of benefit beyond CATS as their data feed into other regional and national services throughout the winter. We describe the novel adaptation of a statistical process control method adopted from the stock market to build a signalling algorithm to identify the start and end of the winter surge. We also discuss the process of developing this solution with the clinical team and how this method was successfully implemented within CATS.

2. Data

Data on all CATS referrals are checked and entered onto a dedicated database daily so that the database is up to date and of high quality in almost real time. In practice, at 9 a.m. on any given day, data exist up until the day before. We used anonymised reports generated from routinely collected data in our analysis. The study was discussed with the local Independent Review Board Chair who confirmed that ethical approval was not required.

To develop our signalling algorithm, we used data on all calls to CATS from April 2005 to July 2013. We quarantined data from July 2013 to February 2014 to validate the final algorithm. All analysis was carried out using a combination of Stata IC12 (StataCorp) and Microsoft Excel 2010.

3. Time scales and variable of interest

Each call is logged in the database with the exact time of call. We first had to decide what time scale to use (for instance, daily or weekly) and what the main variable of interest should be.

All calls that come into CATS need to be answered and dealt with by a member of the clinical team. However, only a subset of these calls result in the retrieval team leaving to pick up and transport a critically ill child. Discussing this with the clinical team at CATS, it was decided that "busy-ness" was experienced as demand for retrieval rather than just volume of calls; an increase in call volume without a corresponding increase in retrievals (for instance, many more calls for advice) could be absorbed relatively easily by the team. Thus it was decided to consider "demand for retrievals" as the variable of interest. This was defined as the sum of calls that actually ended in a retrieval and calls that would have resulted in a retrieval if capacity had been available (i.e. refusals due to no available team). The latter were identified by calls with an outcome of "refused" where the reason for refusal was either no available retrieval team or no available paediatric intensive care bed.

A surge in demand occurs whenever current demand is significantly higher than the recent average. However, from an operational point of view, surges only matter to a service if demand is high enough to strain capacity. For instance, in a service that has experienced consistently low demand (e.g. during August), a surge to medium demand is unlikely to present any problems in meeting that demand. To be useful for winter planning purposes, the identification of the start of the winter surge should occur only when demand has reached a level that strains available capacity. Reviewing the time series data with the clinical team and discussing "busy-ness", revealed that 4 or more retrievals (or demand for retrievals) a day was experienced as "busy". It was agreed to consider the start of the CATS winter surge as when demand first consistently breached "4 a day" or "28 a week" and the end of the surge when "4 a day on average" period was over.

To explore appropriate timescales for the data, we first transformed the raw data into three time series at daily, weekly and monthly timescales. For the weekly data, we worked with consecutive seven day periods rather than "week of year" to avoid disproportionately low numbers in the last week of the year which is less than 7 days. For monthly data, we used calendar month to define each time period. Fig. 1 shows the full monthly and weekly time series from 2005–2013 while Fig. 2 shows example daily data comparing the summer (Panel A) and winter (Panel B) of 2009. The monthly time series shows the annual winter surge in demand very clearly (Panel A, Fig. 1).

The monthly time series (Fig. 1, Panel A) highlights that we can predict roughly when the winter surge will occur – December will be the busiest month of the year. However, the week the winter surge starts is variable – certainly more variable than the current identification of 14th November–1st January allows for.

The time series of demand for emergency retrieval by the CATS team is very variable, particularly at timescales of days and weeks which are most useful for short term service planning (see Figs. 1 and 2). Fig. 2 illustrates that even in summer, when demand is lowest, demand can still be high on any given day and conversely, in winter when demand is highest, demand can still be very low on any given day. For winter surge identification to be operationally useful, it is important to be able to monitor demand levels every day. However, using raw daily demand data, given we are looking for a move from an average daily demand of about "3 a day" to one of about "4 a day (or more)", seems futile given the variability of the daily demand (Fig. 2).

Instead, we used the rolling 7-day total demand as a daily time series. That is, if $\{y_t\}$ is the daily demand for retrievals at time t,