Considering outlier detection to identify extraordinary demand events for quantity-based revenue management

Nicola Rennie[1]

Supervisors: Catherine Cleophas[1,2], Adam Sykulski[1], Florian Dost[1,3]

[1] Lancaster University, [2] Christian-Albrechts-Universität zu Kiel, [3] University of Manchester

28th August 2019

STOR-i excellence with impact

Lancaster University

EPSRC

Erasmus University Rotterdam

Motivation

Motivation: Extraordinary Demand Events in Revenue Management

- Forecasting is a key component of most RM systems.
  - e.g. **passenger demand**, willingness-to-pay, cancellation rates.
- Outliers in demand result in inaccurate forecasts, leading to non-optimal inventory controls, and hence, lost revenue.
- There are two dangers of not detecting outliers:
  - The inability to predict the future in the short-term.
  - Contamination of the data from which forecasts derive in the mid-term.
- Online detection rates are important. If an outlier detection approach only works in hindsight, or close to departure, we can gain little from it.

Motivation: Extraordinary Demand Events in Revenue Management

	Over-forecasting		Under-forecasting	
	% Change in Demand from Forecast			
Forecast Demand Factor	**-25%**	**-12.5%**	**+12.5%**	**+25%**
0.90	-2.6%	-2.0%	+3.5%	+2.2%
1.20	+2.7%	+5.4%	-2.3%	-2.2%
1.50	+10.4%	+2.8%	-7.1%	-7.2%
Avg.	**+3.5%**	**+2.1%**	**-2.0%**	**-2.4%**

Table: % Change in Revenue Resulting from Correcting Inaccurate Demand Forecasts Under EMSRb

- Impacts of unexpected demand are not symmetric.
- Under EMSRb heuristic booking limits, pessimistic forecasting can be beneficial which is in line with previous findings by Weatherford and Belobaba (2002).

Motivation: Extraordinary Demand Events in Revenue Management

	Over-forecasting		Under-forecasting	
	% Change in Demand from Forecast			
Forecast Demand Factor	**-25%**	**-12.5%**	**+12.5%**	**+25%**
0.90	-2.6%	-2.0%	+3.5%	+2.2%
1.20	+2.7%	+5.4%	-2.3%	-2.2%
1.50	+10.4%	+2.8%	-7.1%	-7.2%
Avg.	**+3.5%**	**+2.1%**	**-2.0%**	**-2.4%**

Table: % Change in Revenue Resulting from Correcting Inaccurate Demand Forecasts Under EMSRb

- Impacts of unexpected demand are not symmetric.
- Under EMSRb heuristic booking limits, pessimistic forecasting can be beneficial which is in line with previous findings by Weatherford and Belobaba (2002).

Motivation: Existing Literature

- Weatherford and Belobaba (2002)
  - In terms of loss of potential revenue, ‘the greatest impacts were observed when the fare class demand forecasts proved to be inaccurate.’
- Weatherford and Pölt (2002)
  - Most existing research focuses on accurately forecasting demand, and the ‘better unconstraining of airline demand data in revenue management systems for ... greater revenues’.
- Mukhopadhyay et al (2007)
  - ‘If analysts can reliably improve system-generated forecasts on critical flights at critical times, airlines can generate significantly more revenue.’
- Cleophas et al (2017)
  - ‘Systematically measuring the effect of such interventions and on improving their support is still rare’.
- Aim to improve analyst interventions by identifying critical flights, through incorporating outlier detection methodology from statistics literature into revenue management techniques.

Motivation: Existing Literature

- Weatherford and Belobaba (2002)
  - In terms of loss of potential revenue, 'the greatest impacts were observed when the fare class demand forecasts proved to be inaccurate.'
- Weatherford and Pölt (2002)
  - Most existing research focuses on accurately forecasting demand, and the 'better unconstraining of airline demand data in revenue management systems for ... greater revenues'.
- Mukhopadhyay et al (2007)
  - 'If analysts can reliably improve system-generated forecasts on critical flights at critical times, airlines can generate significantly more revenue.'
- Cleophas et al (2017)
  - 'Systematically measuring the effect of such interventions and on improving their support is still rare'.
- Aim to improve analyst interventions by identifying critical flights, through incorporating outlier detection methodology from statistics literature into revenue management techniques.

Motivation: Existing Literature

- Weatherford and Belobaba (2002)
  - In terms of loss of potential revenue, 'the greatest impacts were observed when the fare class demand forecasts proved to be inaccurate.'
- Weatherford and Pölt (2002)
  - Most existing research focuses on accurately forecasting demand, and the 'better unconstraining of airline demand data in revenue management systems for ... greater revenues'.
- Mukhopadhyay et al (2007)
  - 'If analysts can reliably improve system-generated forecasts on critical flights at critical times, airlines can generate significantly more revenue.'
- Cleophas et al (2017)
  - 'Systematically measuring the effect of such interventions and on improving their support is still rare'.
- Aim to improve analyst interventions by identifying critical flights, through incorporating outlier detection methodology from statistics literature into revenue management techniques.

Motivation: Existing Literature

- Weatherford and Belobaba (2002)
  - In terms of loss of potential revenue, 'the greatest impacts were observed when the fare class demand forecasts proved to be inaccurate.'
- Weatherford and Pölt (2002)
  - Most existing research focuses on accurately forecasting demand, and the 'better unconstraining of airline demand data in revenue management systems for ... greater revenues'.
- Mukhopadhyay et al (2007)
  - 'If analysts can reliably improve system-generated forecasts on critical flights at critical times, airlines can generate significantly more revenue.'
- Cleophas et al (2017)
  - 'Systematically measuring the effect of such interventions and on improving their support is still rare'.
- Aim to improve analyst interventions by identifying critical flights, through incorporating outlier detection methodology from statistics literature into revenue management techniques.

Motivation: Existing Literature

- Weatherford and Belobaba (2002)
  - In terms of loss of potential revenue, 'the greatest impacts were observed when the fare class demand forecasts proved to be inaccurate.'
- Weatherford and Pölt (2002)
  - Most existing research focuses on accurately forecasting demand, and the 'better unconstraining of airline demand data in revenue management systems for ... greater revenues'.
- Mukhopadhyay et al (2007)
  - 'If analysts can reliably improve system-generated forecasts on critical flights at critical times, airlines can generate significantly more revenue.'
- Cleophas et al (2017)
  - 'Systematically measuring the effect of such interventions and on improving their support is still rare'.
- Aim to improve analyst interventions by identifying critical flights, through incorporating outlier detection methodology from statistics literature into revenue management techniques.

Motivation: What Do We Want to Do?

- Examine the use of existing outlier detection methods for the identification of unusual demand, and highlight specific features of their use in the revenue management setting.
- Propose an adaptation to an existing functional outlier detection method which significantly improves performance.

To our knowledge this is the first suggestion of an automated methodology for outlier detection in a revenue management system.

Outlier Detection

Outlier Detection

Outlier

Outlier: ‘an observation which deviates so much from the other observations as to arouse suspicions that it was generated by a different mechanism.’ - Hawkins, 1980

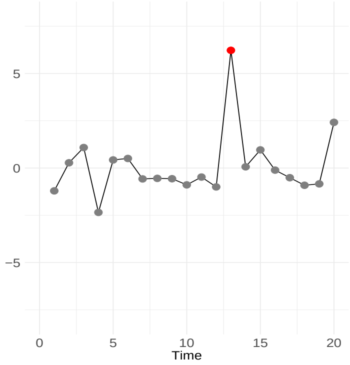

Figure: Outlier Within Time Series

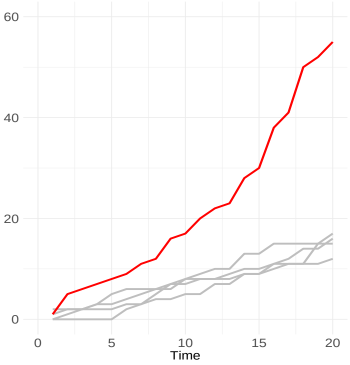

Figure: Time Series as an Outlier

Outlier Detection: Univariate Approaches

- Applied at each time point in independently.
- Ignores dependence between and within booking curves.
- Methods: Nonparametric percentiles, tolerance intervals, and robust Z-score.

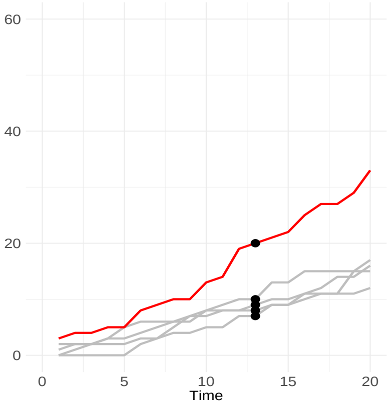

Figure: Booking Curves

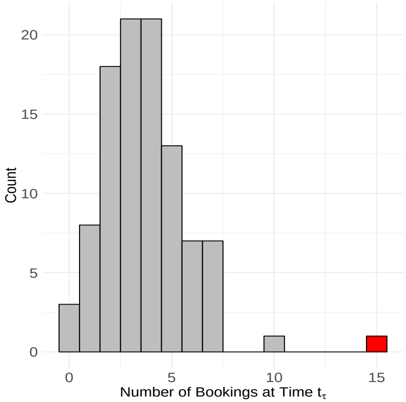

Figure: Histogram of Bookings at Time t_τ

Outlier Detection: Multivariate Approaches

- Applied to a vector of bookings up to time t_τ.
- Ignores dependence between booking curves, and time-dependence within booking curves.
- Methods: Distance-based (k-nearest neighbours), and clustering-based (k-means).

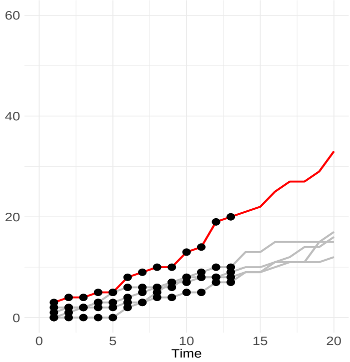

Figure: Booking Curves

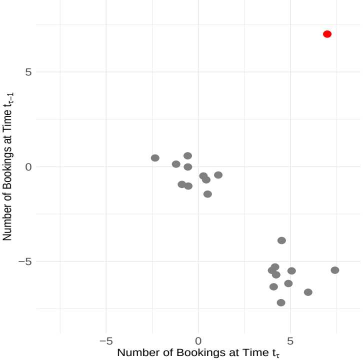

Figure: τ-Dimensional Plot of Bookings at Time t_τ

Outlier Detection: Functional Approaches

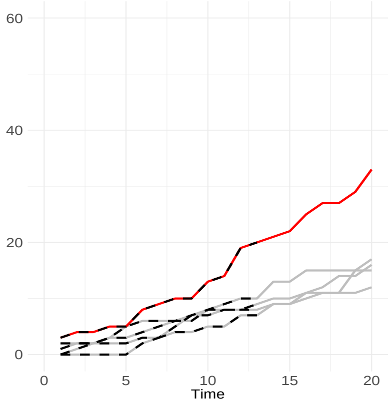

Figure: Booking Curves (2-dimensional)

Functional Data Analysis

- Treat each booking curve as observations of a real function.
- Febrero et al. (2008) define an outlier as a curve generated by a stochastic process with a different distribution than the rest of the curves, which are assumed to be identically distributed.
- Takes into account dependence within booking curves, specifically time dependence.

Outlier Detection: Functional Approaches

Functional Depth

- Functional depth is a measure of the centrality, or 'outlyingness' of an observation with respect to a given dataset (López-Pintado and Romo (2009)).
- In the case of one-dimensional random variables, the **halfspace depth** of a point y_n with respect to a sample $y1, \ldots, y_N$ drawn from distribution F is:

  $$HD(y_n) = min\{F_N(y_n), 1 - F_N(y_n)\}$$

  where F_N is the empirical cumulative distribution of the sample $y_1, \ldots, y_N$.
- This definition has been extended into the multivariate functional data setting.
- We detect outliers by calculating the **multivariate functional halfspace depth** of each booking curve up to time t_τ. Those curves with depths below some threshold are classified as outliers.

Proposed Method: Functional Outlier Detection with Extrapolation

- **Idea**: Combine univariate time series forecasting methods to extrapolate beyond the observed data, then apply function depth outlier detection to the combined observed and extrapolated curve.

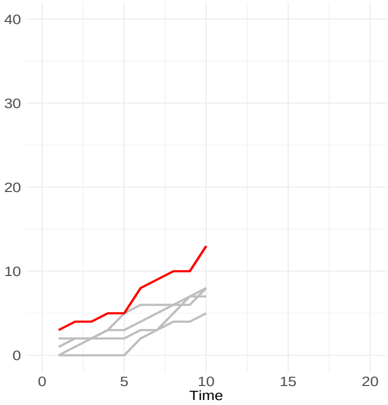

Figure: Booking Curves Until Time t_τ

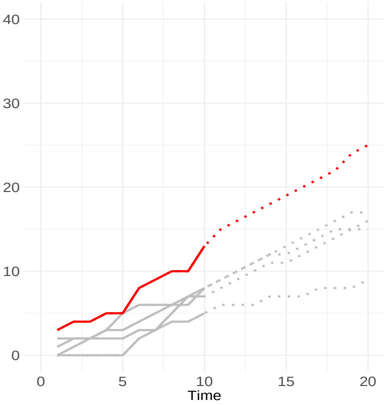

Figure: Booking Curves With Extrapolation

Proposed Method: Functional Outlier Detection with Extrapolation

Algorithm 1: Using Extrapolation to Improve Functional Outlier Detection

At time t_τ forecast the accumulation of bookings at each time $\tau + 1, \dots, T$, $\hat{y}_n(t_{\tau+1}), \dots, \hat{y}_n(t_T)$, for each flight n ;
Calculate $\mathcal{D}_n(\hat{\boldsymbol{y}}_n(t_\tau))$, the functional depth of the observed and extrapolated booking curve $\hat{\boldsymbol{y}}_n(t_\tau) = (y_n(t_1), y_n(t_2), \dots, y_n(t_\tau), \hat{y}_n(t_{\tau+1}), \dots, \hat{y}_n(t_T))$, for each flight n at time t_τ. ;
Calculate a threshold, C, for the functional depth. ;
if $\mathcal{D}_n(\hat{\boldsymbol{y}}_n(t_\tau)) \leq C$ **then**
  Define flight n as an outlier. Delete flight n from the sample of N flights.
end
while $\exists$ n **s.t.** $\mathcal{D}_n(\hat{\boldsymbol{y}}_n(t_\tau)) \leq C$ **do**
  Recalculate functional depths on the new sample, and remove further outliers.
end

Proposed Method: Functional Outlier Detection with Extrapolation

- Why do we need to generate new forecasts for extrapolation?
  - Not all revenue management systems require forecasts of how demand accumulates, only final demand.
  - Not all RM systems store historic forecasts.
  - Forecasts are based on multiple flights which normalises outlying behaviour.

Proposed Method: Functional Outlier Detection with Extrapolation

- Why do we need to generate new forecasts for extrapolation?
  - Not all revenue management systems require forecasts of how demand accumulates, only final demand.
  - Not all RM systems store historic forecasts.
  - Forecasts are based on multiple flights which normalises outlying behaviour.

Proposed Method: Functional Outlier Detection with Extrapolation

- **Simple Exponential Smoothing (SES)**: SES works on the principle of averaging whilst down-weighting older observations. Given a time series $y_n(t_1), y_n(t_2), \ldots, y_n(t_\tau)$, a forecast for time $t_{\tau+1}$, $\hat{y}_n(t_{\tau+1})$ is given by:

$$\hat{y}_n(t_{\tau+1}) = \alpha y_n(t_\tau) + (1-\alpha)\hat{y}_n(t_\tau),$$

for some smoothing constant, α.

- **Autoregressive Integrated Moving Average (ARIMA)**: ARIMA models incorporate a trend component, and assume that future observations are an additive, weighted combination of previous observations and previous errors. Let $\boldsymbol{x}_n(t_\tau)$ be the d^{th} differenced time series relating to $\boldsymbol{y}_n(t_\tau)$. The one-step ahead forecast $\hat{x}_n(t_{\tau+1})$ is given by:

$$\hat{x}_n(t_{\tau+1}) = \mu + \phi_1 x_n(t_\tau) + \ldots + \phi_p x_n(t_{\tau-p+1}) - \theta_1 \epsilon(t_\tau) - \ldots - \theta_q \epsilon(t_{\tau-q+1})$$

for some constant mean μ, parameters $\phi_1, \ldots, \phi_p, \theta_1, \ldots, \theta_q$ and white noise process $\left(\epsilon_{t_j}\right)$.

Proposed Method: Functional Outlier Detection with Extrapolation

- **Simple Exponential Smoothing (SES)**: SES works on the principle of averaging whilst down-weighting older observations. Given a time series $y_n(t_1), y_n(t_2), \dots, y_n(t_\tau)$, a forecast for time $t_{\tau+1}$, $\hat{y}_n(t_{\tau+1})$ is given by:

$$\hat{y}_n(t_{\tau+1}) = \alpha y_n(t_\tau) + (1-\alpha)\hat{y}_n(t_\tau),$$

for some smoothing constant, α.

- **Autoregressive Integrated Moving Average (ARIMA)**: ARIMA models incorporate a trend component, and assume that future observations are an additive, weighted combination of previous observations and previous errors. Let $\boldsymbol{x}_n(t_\tau)$ be the d^{th} differenced time series relating to $\boldsymbol{y}_n(t_\tau)$. The one-step ahead forecast $\hat{x}_n(t_{\tau+1})$ is given by:

$$\hat{x}_n(t_{\tau+1}) = \mu + \phi_1 x_n(t_\tau) + \dots + \phi_p x_n(t_{\tau-p+1}) - \theta_1 \epsilon(t_\tau) - \dots - \theta_q \epsilon(t_{\tau-q+1})$$

for some constant mean μ, parameters $\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q$ and white noise process $\left(\epsilon_{t_j}\right)$.

Proposed Method: Functional Outlier Detection with Extrapolation

- **Integrated Generalised Autoregressive Conditional Heteroskedasticity (IGARCH)**: IGARCH models incorporate a trend component and assume that the variance structure follows an autoregressive moving average model. Again, let $\boldsymbol{x}_n(t_\tau)$ be the d^{th} differenced time series relating to $\boldsymbol{y}_n(t_\tau)$. IGARCH(1,d,1) models assume the following structure:

$$\begin{aligned} x_n(t_{\tau+1}) &= \mu + \epsilon_n(t_{\tau+1}) \\ \epsilon_n(t_{\tau+1}) &= z_n(t_{\tau+1})\sigma_n(t_{\tau+1}) \\ \sigma_n^2(t_{\tau+1}) &= w + \alpha\epsilon_n^2(t_{\tau+1}) + \beta\sigma_n^2(t_\tau) \end{aligned}$$

Simulation

Simulation: Customer Arrivals

- 2 Customer Types:
  - Business and tourist customers, as per Weatherford et al (1993).
  - Business customers arrive later in the booking horizon than tourists.
  - Business customers typically have higher willingness-to-pay, and are less price sensitive.
- 7 Fare Classes:
  - Semi-differentiated fare class structure.

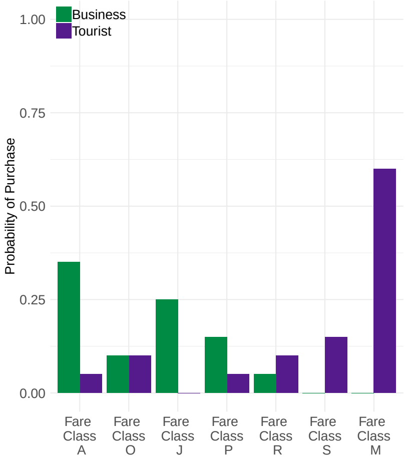

Figure: Probability of Purchase for Each Fare Class

Simulation: Customer Arrivals

Customer Arrivals

Each customer type arrives according to a Poisson-Gamma process with rate $\lambda_i(t)$:

$$\lambda_i(t) = A\phi_i \frac{t^{a_i-1}(1-t)^{b_i-1}}{B(a_i, b_i)}$$

and chooses to purchase a seat in fare class j with probability p_{ij}. where:

- $i \in \mathcal{I} = \{1 = business, 2 = tourist\}$
- $A \sim Gamma(\alpha, \beta)$.
- $\phi_1 + \phi_2 = 1$.
- $\frac{a_1-1}{a_1+b_1-2} > \frac{a_2-1}{a_2+b_2-2}$

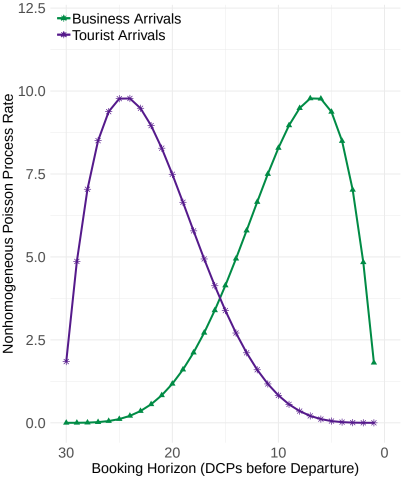

Figure: Arrival Rates for Business and Tourist Passengers

Simulation: Booking Controls

- Maximise revenue by limiting the number of low value tickets sold.
- Forecast demand and allocate capacity to each fare class.

EMSRb

Expected Marginal Seat Revenue-b booking limit for fare class j is given by:

$$PL_j = F_j^{-1}\left(1 - \frac{r_{j+1}}{\tilde{r}_j}\right),$$

- F_j, (Gaussian) distribution of independent demand for fare class j,
- r_j, fare in fare class j,
- $\tilde{r}_j$, weighted-average revenue from classes $1, \ldots, j$.

EMSRb - Marginal Revenue

- Does not assume the distribution of demand is independent across fare classes, and attempts to protect against buy down.
- Fiig et al. (2010) transform the demand and fares into an equivalent independent demand model through a marginal revenue transformation.
- Apply EMSRb to the transformed demand and fares.

Simulation: Booking Controls

- Maximise revenue by limiting the number of low value tickets sold.
- Forecast demand and allocate capacity to each fare class.

EMSRb

Expected Marginal Seat Revenue-b booking limit for fare class j is given by:

$$PL_j = F_j^{-1}\left(1 - \frac{r_{j+1}}{\tilde{r}_j}\right),$$

- F_j, (Gaussian) distribution of independent demand for fare class j,
- r_j, fare in fare class j,
- $\tilde{r}_j$, weighted-average revenue from classes $1, \ldots, j$.

EMSRb - Marginal Revenue

- Does not assume the distribution of demand is independent across fare classes, and attempts to protect against buy down.
- Fiig et al. (2010) transform the demand and fares into an equivalent independent demand model through a marginal revenue transformation.
- Apply EMSRb to the transformed demand and fares.

Simulation: Booking Controls

- Maximise revenue by limiting the number of low value tickets sold.
- Forecast demand and allocate capacity to each fare class.

EMSRb

Expected Marginal Seat Revenue-b booking limit for fare class j is given by:

$$PL_j = F_j^{-1}\left(1 - \frac{r_{j+1}}{\tilde{r}_j}\right),$$

- F_j, (Gaussian) distribution of independent demand for fare class j,
- r_j, fare in fare class j,
- $\tilde{r}_j$, weighted-average revenue from classes $1, \ldots, j$.

EMSRb - Marginal Revenue

- Does not assume the distribution of demand is independent across fare classes, and attempts to protect against buy down.
- Fiig et al. (2010) transform the demand and fares into an equivalent independent demand model through a marginal revenue transformation.
- Apply EMSRb to the transformed demand and fares.

Simulation: Generating Extraordinary Demand Events

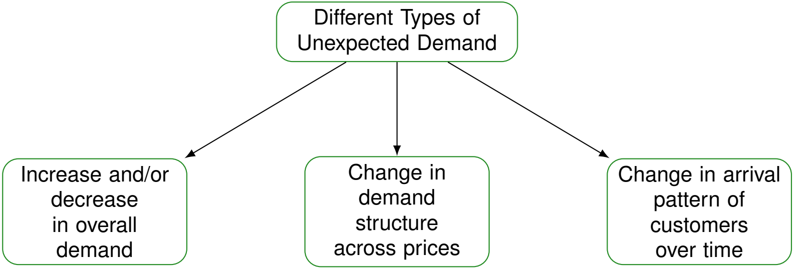

Simulation: Generating Extraordinary Demand Events

- Focus on outliers generated by an increase or decrease in overall demand.
- Consider four types of outliers: $\pm$ 12.5%, $\pm$ 25% change in demand from forecast.

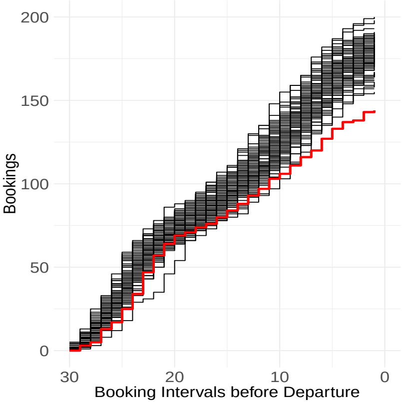

Figure: 25% Decrease

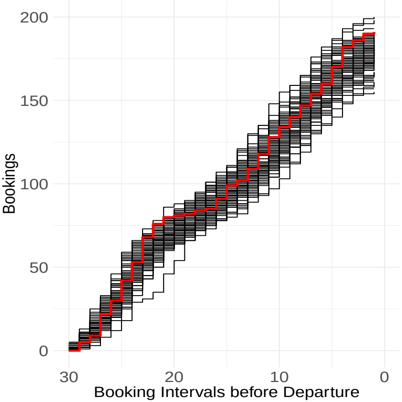

Figure: 25% Increase

Results

Results: Performance Metrics

True Positive Rate (TPR)

$$\frac{TP}{TP + FN}$$

False Positive Rate (FPR)

$$\frac{FP}{FP + TN}$$

Matthews Correlation Coefficient (MCC)

$$\frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

Balanced Classification Rate (BCR)

$$\frac{1}{2}\left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP}\right)$$

Results: Performance Metrics

True Positive Rate (TPR)

$$\frac{TP}{TP + FN}$$

False Positive Rate (FPR)

$$\frac{FP}{FP + TN}$$

Matthews Correlation Coefficient (MCC)

$$\frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

Balanced Classification Rate (BCR)

$$\frac{1}{2}\left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP}\right)$$

Results: Performance Metrics

True Positive Rate (TPR)

$$\frac{TP}{TP + FN}$$

False Positive Rate (FPR)

$$\frac{FP}{FP + TN}$$

Matthews Correlation Coefficient (MCC)

$$\frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

Balanced Classification Rate (BCR)

$$\frac{1}{2}\left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP}\right)$$

Results: Performance Metrics

True Positive Rate (TPR)

$$\frac{TP}{TP + FN}$$

False Positive Rate (FPR)

$$\frac{FP}{FP + TN}$$

Matthews Correlation Coefficient (MCC)

$$\frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

Balanced Classification Rate (BCR)

$$\frac{1}{2}\left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP}\right)$$

Results: EMSRb vs EMSRb-MR

- There is no significant difference between how a method performs under EMSRb vs EMSRb-MR controls.

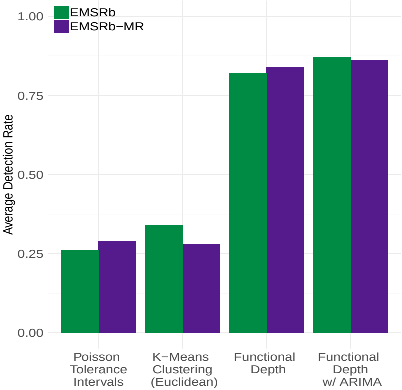

Figure: TPR

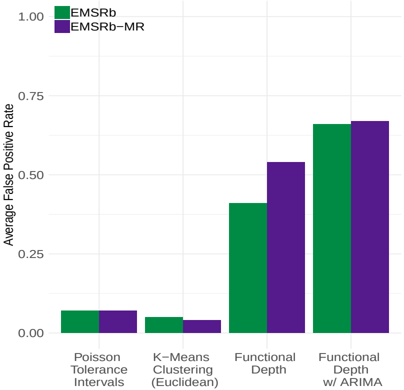

Figure: FPR

Results: EMSRb vs EMSRb-MR

- There is no significant difference between how a method performs under EMSRb vs EMSRb-MR controls.

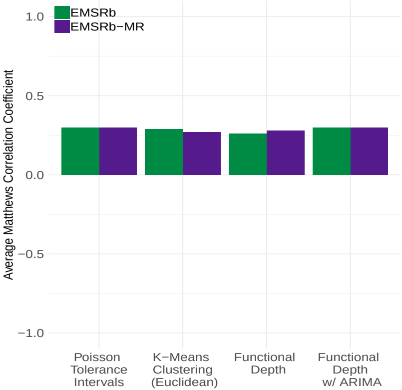

Figure: MCC

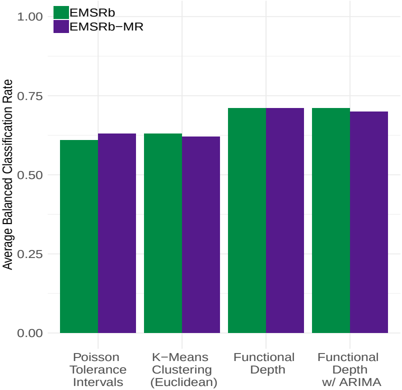

Figure: BCR

Results: Univariate vs Multivariate vs. Functional Approaches

- Univariate methods' performance drops off, and multivariate methods' performance levels off, whereas functional generally increase over time.
- Functional approaches generally outperform other approaches (except in terms of MCC as it penalises high FPR with unbalanced class sizes).

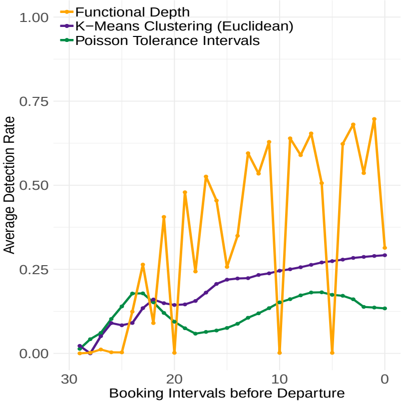

Figure: TPR

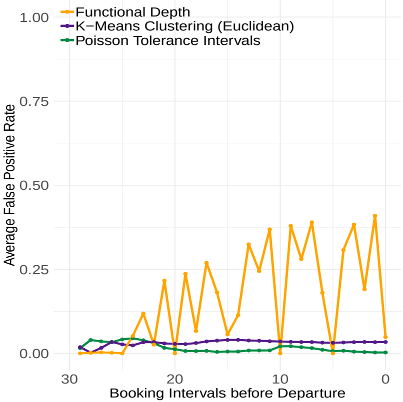

Figure: FPR

Results: Univariate vs Multivariate vs. Functional Approaches

- Univariate methods' performance drops off, and multivariate methods' performance levels off, whereas functional generally increase over time.
- Functional approaches generally outperform other approaches (except in terms of MCC as it penalises high FPR with unbalanced class sizes).

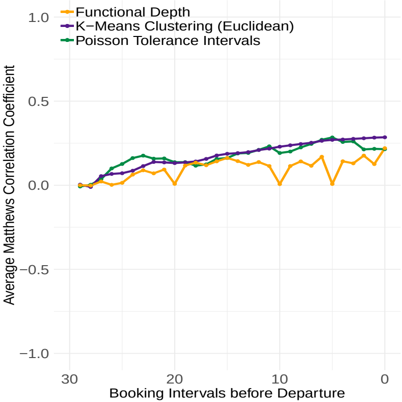

Figure: MCC

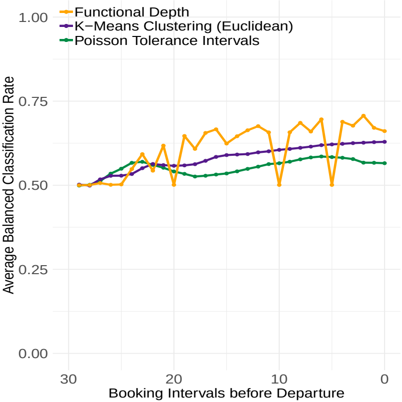

Figure: BCR

Results: Different Types of Outliers

- All methods are better at detecting larger magnitude changes in demand.
- All methods are better at detecting decreases in demand, as opposed to increases.

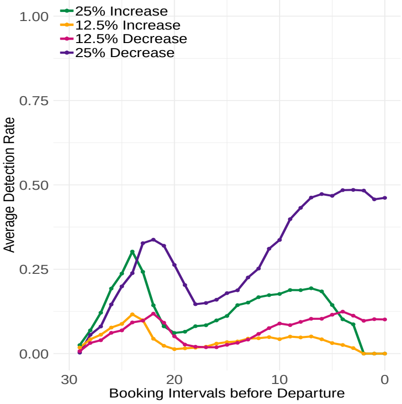

Figure: TPR

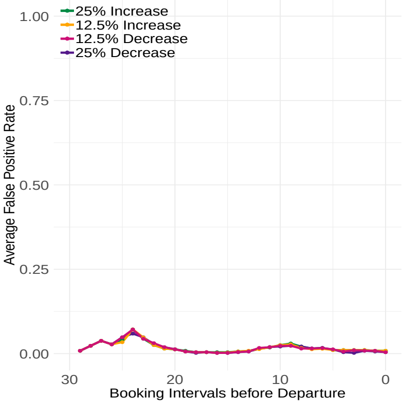

Figure: FPR

Results: Different Types of Outliers

- All methods are better at detecting larger magnitude changes in demand.
- All methods are better at detecting decreases in demand, as opposed to increases.

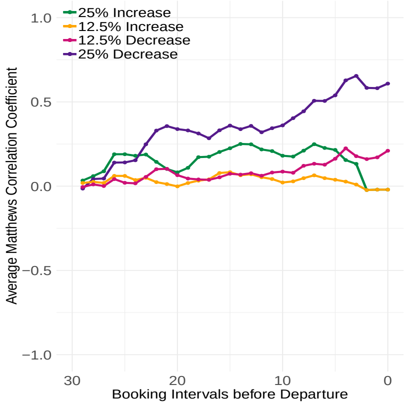

Figure: MCC

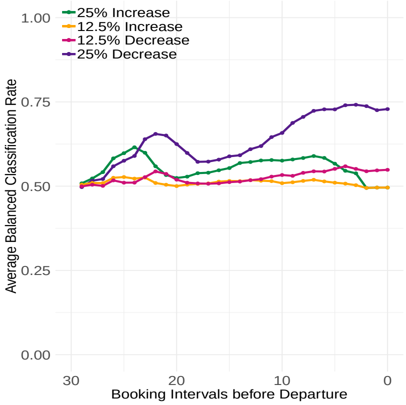

Figure: BCR

Results: Percentage of Outliers

- There is no significant difference between how a method performs under different percentages of outliers.

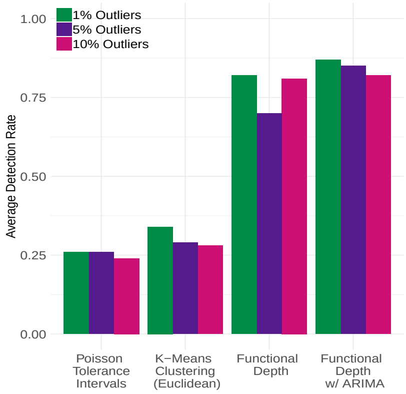

Figure: TPR

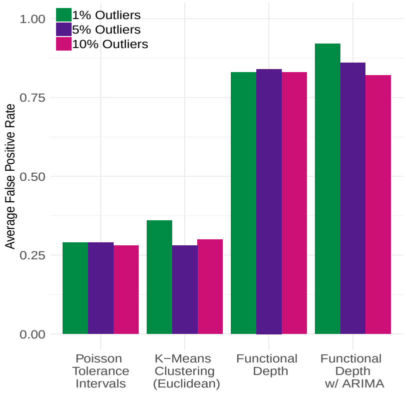

Figure: FPR

Results: Percentage of Outliers

- There is no significant difference between how a method performs under different percentages of outliers.

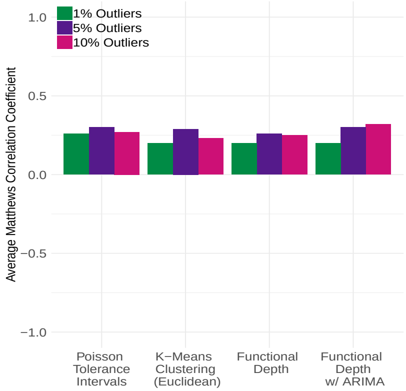

Figure: MCC

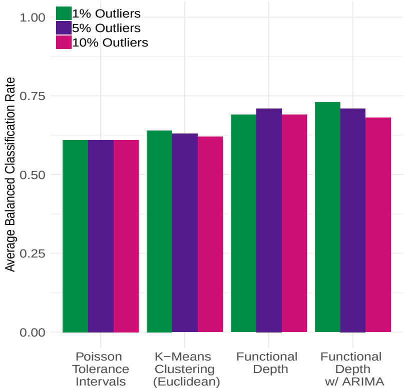

Figure: BCR

Results: Improvement from Extrapolation

- Extrapolation improves outlier detection performance, specifically early in the booking horizon.

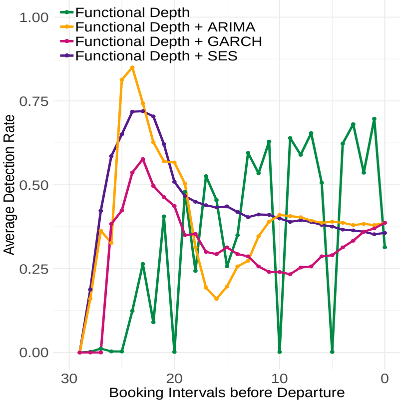

Figure: TPR

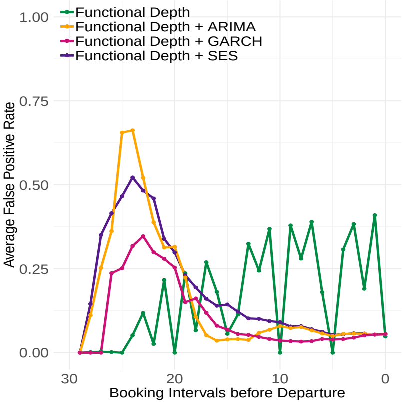

Figure: FPR

Results: Improvement from Extrapolation

- Extrapolation improves outlier detection performance, specifically early in the booking horizon.

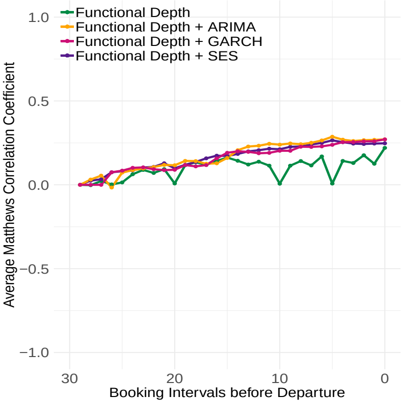

Figure: MCC

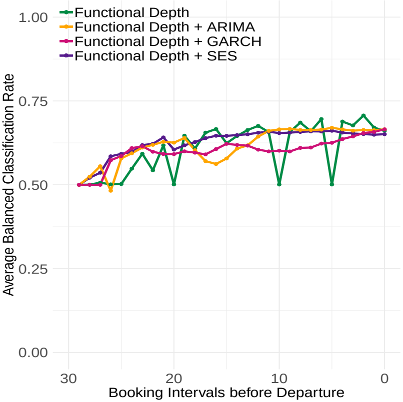

Figure: BCR

Conclusions

Conclusions

- Outlier detection is a viable method for automating the identification of critical flights which require analyst adjustment.
- Our simulation framework demonstrates that functional approaches are more promising than univariate or multivariate approaches to outlier detection.
- Our proposed extrapolation step improves outlier detection performance, particularly early in the booking horizon when it is most valuable.

What are the Implications for Revenue Management Analysts?

- Implementing outlier detection becomes an automated part of the RM system.
- An alert is sent to the relevant analyst if a flight is deemed an outlier.
- The analyst updates the forecast (and therefore booking controls).

Further Work

- Take into account time-dependence between curves, and include seasonality in demand.
- Further develop the method by making recommendations to analysts about which action should be taken, after a critical flight is identified.
- Extend to a multivariate setting to jointly monitor booking curves and revenue curves.
- Investigate the impact of unexpected demand in the dynamic pricing setting.

References

- L. R. Weatherford, S. E. Bodily, and P. E. Pfeifer. *Modeling the Customer Arrival Process and Comparing Decision Rules in Perishable Asset Revenue Management Situations*. Transportation Science, 27(3):239–251, 1993.
- L. R. Weatherford and P. P. Belobaba. *Revenue impacts of fare input and demand forecast accuracy in airline yield management*. The Journal of the Operational Research Society, 53(8):811–821, 2002.
- K. T. Talluri and G. J. Van Ryzin. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, 2004.
- S. Mukhopadhyay, S. Samaddar, and C. Colville. *Improving Revenue Management Decision Making for Airlines by Evaluating Analyst-Adjusted Passenger Demand Forecasts*. Decision Sciences, 38(2):309–327, 2007.
- C. Cleophas, D. Kadatz, and S. Vock. *Resilient Revenue Management: A Literature Survey of Recent Theoretical Advances*. Journal of Revenue and Pricing Management, 16(5): 483–498, 2017.

Questions?

n.rennie@lancaster.ac.uk