Data Driven Alerts in Airline Revenue Management

The Identification of Inaccurate Demand Estimates

Nicola Rennie[1]

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

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

24th June 2019

Motivation

Motivation: Revenue Management

- Relates to making **demand-management decisions**, with the objective being to increase revenue.
- Combines forecasting with optimisation.
- Three types of revenue management decisions:
  - Structural decisions
  - Price decisions
  - Quantity decisions

Motivation: Revenue Management

- Relates to making **demand-management decisions**, with the objective being to increase revenue.
- Combines forecasting with optimisation.
- Three types of revenue management decisions:
  - Structural decisions
  - Price decisions
  - Quantity decisions

Motivation: Revenue Management

- Relates to making **demand-management decisions**, with the objective being to increase revenue.
- Combines forecasting with optimisation.
- Three types of revenue management decisions:
  - Structural decisions
  - Price decisions
  - Quantity decisions

Motivation: Revenue Management

- Relates to making **demand-management decisions**, with the objective being to increase revenue.
- Combines forecasting with optimisation.
- Three types of revenue management decisions:
  - Structural decisions
  - Price decisions
  - Quantity decisions

Motivation: Revenue Management

- Relates to making **demand-management decisions**, with the objective being to increase revenue.
- Combines forecasting with optimisation.
- Three types of revenue management decisions:
  - Structural decisions
  - Price decisions
  - Quantity decisions

Motivation: Revenue Management

- Relates to making **demand-management decisions**, with the objective being to increase revenue.
- Combines forecasting with optimisation.
- Three types of revenue management decisions:
  - Structural decisions
  - Price decisions
  - Quantity decisions

Motivation: Revenue Management

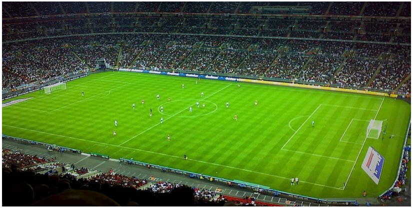

Unexpected Demand

↓

Inaccurate Forecasts

↓

Non-optimal Capacity Allocation

↓

Lost Revenue

Airline Revenue Management System

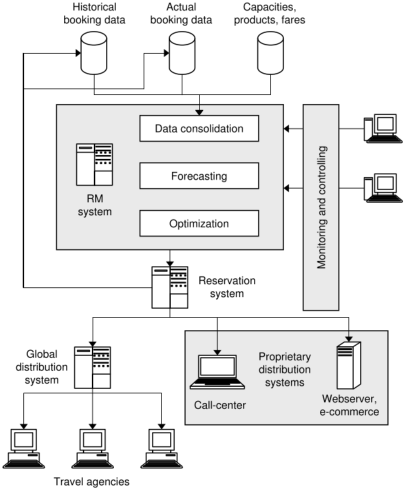

Figure: *The Theory and Practice of Revenue Management*, Talluri & van Ryzin, 2004

Airline Revenue Management System

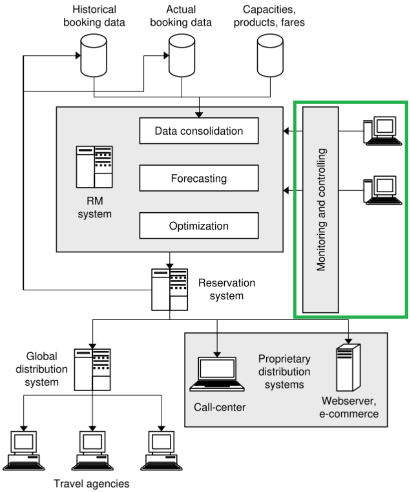

Figure: *The Theory and Practice of Revenue Management*, Talluri & van Ryzin, 2004

Airline Revenue Management System

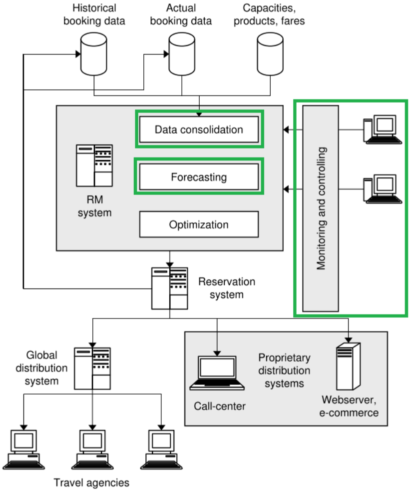

Figure: *The Theory and Practice of Revenue Management*, Talluri & van Ryzin, 2004

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.'
- 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.’
- 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.'
- 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.'
- 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: Impact on Potential Revenue

	% Change in Demand from Forecast			
Demand Factor	**-25%**	**-12.5%**	**+12.5%**	**+25%**
0.90	-13.6%	-7.7%	+10.0%	+13.8%
1.20	-11.3%	-3.4%	+1.4%	+2.8%
1.50	+2.1%	-0.7%	+7.7%	+18.0%
Avg.	-7.6%	-3.9%	+6.4%	+11.5%

Table: % Change in Revenue from Identifying Inaccurate Demand Forecasts Under EMSR-b Controls

- In line with previous findings by Weatherford and Belobaba (2002).
  - Impacts of unexpected demand are not symmetric.
  - Under EMSR-b heuristic booking limits, optimistic forecasting can be beneficial.
- Potential impact of detecting outliers depends on the optimisation routine used to set booking limits.

Motivation: Impact on Potential Revenue

	% Change in Demand from Forecast			
Demand Factor	**-25%**	**-12.5%**	**+12.5%**	**+25%**
0.90	-13.6%	-7.7%	+10.0%	+13.8%
1.20	-11.3%	-3.4%	+1.4%	+2.8%
1.50	+2.1%	-0.7%	+7.7%	+18.0%
Avg.	-7.6%	-3.9%	+6.4%	+11.5%

Table: % Change in Revenue from Identifying Inaccurate Demand Forecasts Under EMSR-b Controls

- In line with previous findings by Weatherford and Belobaba (2002).
  - Impacts of unexpected demand are not symmetric.
  - Under EMSR-b heuristic booking limits, optimistic forecasting can be beneficial.
- Potential impact of detecting outliers depends on the optimisation routine used to set booking limits.

Simulation

Simulation: Customer Arrivals

- Two customer types: business and tourist, 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.

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:

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

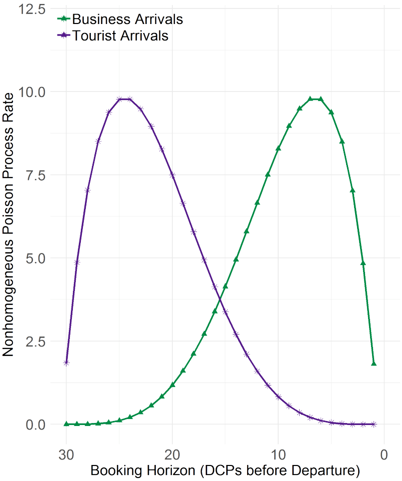

Figure: Arrival Rates for Business and Tourist Passenger

Simulation: Customer Arrivals

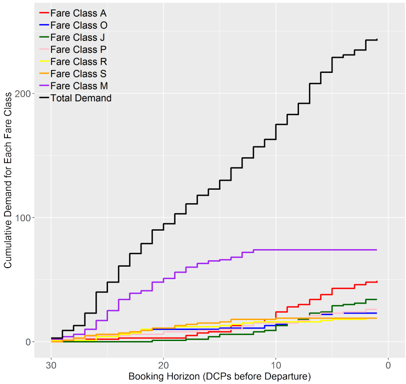

Figure: Fare Class Demand from all Passengers

Fare Class	$\hat{\mu}$	$\hat{\sigma}^2$
A	46.2	25.3
O	24.2	18.8
J	28.6	25.5
P	22.9	26.6
R	18.5	16.5
S	16.9	11.2
M	69.8	28.2

Table: Mean and Variance Forecasts $\alpha = 240$, $\beta = 1$, $\phi_1 = \phi_2 = 0.5$, $a_1 = 5$, $b_1 = 2$, $a_2 = 2$, $b_2 = 5$

Simulation: Booking Limits

- Maximise revenue by limiting the number of low value tickets sold.
- Allocate capacity to each fare class.
- 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 demand for fare class j,
  - r_j, fare in fare class j,
  - $\tilde{r}_j$, weighted-average revenue from classes $1, \ldots, j$.

Fare Class	$\hat{\mu}$	**BL**
A	46.2	43
O	24.2	23
J	28.6	29
P	22.9	24
R	18.5	19
S	16.9	17
M	69.8	45

Table: Booking Limits

Simulation: Booking Limits

- Maximise revenue by limiting the number of low value tickets sold.
- Allocate capacity to each fare class.
- 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 demand for fare class j,
  - r_j, fare in fare class j,
  - $\tilde{r}_j$, weighted-average revenue from classes $1, \ldots, j$.

Fare Class	$\hat{\mu}$	**BL**
A	46.2	43
O	24.2	23
J	28.6	29
P	22.9	24
R	18.5	19
S	16.9	17
M	69.8	45

Table: Booking Limits

Simulation: Booking Limits

- Maximise revenue by limiting the number of low value tickets sold.
- Allocate capacity to each fare class.
- 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 demand for fare class j,
  - r_j, fare in fare class j,
  - $\tilde{r}_j$, weighted-average revenue from classes $1, \ldots, j$.

Fare Class	$\hat{\mu}$	BL
A	46.2	43
O	24.2	23
J	28.6	29
P	22.9	24
R	18.5	19
S	16.9	17
M	69.8	45

Table: Booking Limits

Simulation: Booking Limits

- Maximise revenue by limiting the number of low value tickets sold.
- Allocate capacity to each fare class.
- 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 demand for fare class j,
  - r_j, fare in fare class j,
  - $\tilde{r}_j$, weighted-average revenue from classes $1, \ldots, j$.

Fare Class	$\hat{\mu}$	**BL**
A	46.2	43
O	24.2	23
J	28.6	29
P	22.9	24
R	18.5	19
S	16.9	17
M	69.8	45

Table: Booking Limits

Simulation: Booking Data

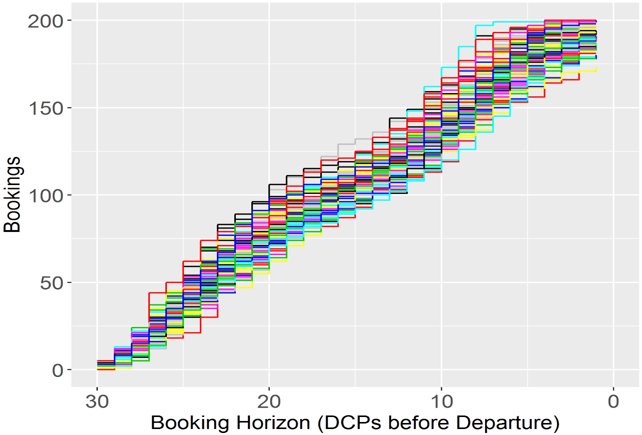

Figure: Simulated Booking Data (Aggregated by Departure Date)

Simulation: Generating Unexpected Demand

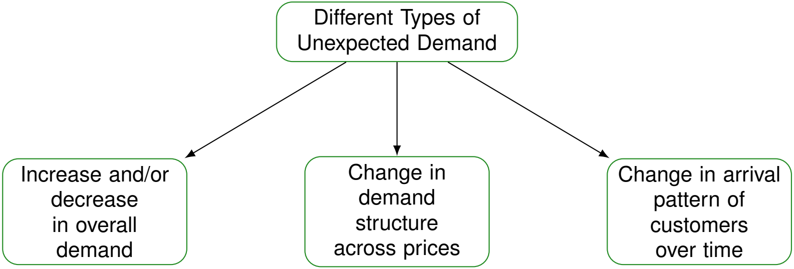

Simulation: Generating Unexpected Demand

- 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.
  - e.g. 25% increase in demand: generate 475 (normal) flights which have expected demand 240, and generate 25 (outlier) flights which have expected demand 300.
  - Aim to detect those 25 outlier flights as early in the booking horizon as possible.

Simulation: Generating Unexpected Demand

- 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.
  - e.g. 25% increase in demand: generate 475 (normal) flights which have expected demand 240, and generate 25 (outlier) flights which have expected demand 300.
  - Aim to detect those 25 outlier flights as early in the booking horizon as possible.

Methodology

Outlier Detection

- Univariate outlier detection
  - Applied at each time point independently.
  - Ignores time dependence within and between booking curves.
- Multivariate outlier detection
  - Treats each booking curve at time t, as a point in t-dimensional space.
  - Ignores time dependence within and between booking curves.
  - Issues with high-dimensionality.
- Functional outlier detection
  - Treat booking curves as observations of a real function.
  - 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.

Outlier Detection

- Univariate outlier detection
  - Applied at each time point independently.
  - Ignores time dependence within and between booking curves.
- Multivariate outlier detection
  - Treats each booking curve at time t, as a point in t-dimensional space.
  - Ignores time dependence within and between booking curves.
  - Issues with high-dimensionality.
- Functional outlier detection
  - Treat booking curves as observations of a real function.
  - 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.

Outlier Detection

- Univariate outlier detection
  - Applied at each time point independently.
  - Ignores time dependence within and between booking curves.
- Multivariate outlier detection
  - Treats each booking curve at time t, as a point in t-dimensional space.
  - Ignores time dependence within and between booking curves.
  - Issues with high-dimensionality.
- Functional outlier detection
  - Treat booking curves as observations of a real function.
  - 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.

Univariate Outlier Detection

- Percentile Bootstrapping
  - At each time point, calculate an lower and upper limit for the number of bookings which are classified as normal.
  - Bootstrap number of bookings at each time point, and find 2.5^{th}, and 97.5^{th} percentile of each bootstrap sample. Take median across bootstrap samples as lower/upper limit.

- Tolerance Intervals
  - Nonparametric or parametric approaches.
  - Require two parameters: the coverage proportion, β, and confidence level, $1-\alpha$.
  - For $X_1, X_2, \dots, X_n$ a random sample from a population with distribution $F(X)$, if

$$P(F(U) - F(L) \geq \beta) = 1 - \alpha, \quad (1)$$

then the interval (L, U) is called a $(\beta, 1-\alpha)$ two-sided tolerance interval.

- Robust Z-Score
  - Let $y_i(t)$ be the cumulative number of bookings for flight i at time t. The robust Z-score can be calculated as:

$$Z = \frac{0.6745\,(y_i(t) - \tilde{y}(t))}{MAD(t)} \quad (2)$$

where $\tilde{y}(t)$ is the median number of bookings at time t across all flights. Flights with a robust Z-score above 3.5 are classified as outliers. (Iglewicz and Hoaglin (1993))

Univariate Outlier Detection

- Percentile Bootstrapping
  - At each time point, calculate an lower and upper limit for the number of bookings which are classified as normal.
  - Bootstrap number of bookings at each time point, and find 2.5^{th}, and 97.5^{th} percentile of each bootstrap sample. Take median across bootstrap samples as lower/upper limit.

- Tolerance Intervals
  - Nonparametric or parametric approaches.
  - Require two parameters: the coverage proportion, β, and confidence level, $1 - \alpha$.
  - For $X_1, X_2, \ldots, X_n$, a random sample from a population with distribution $F(X)$, if:

$$\mathbb{P}\left(F(U) - F(L) > \beta\right) = 1 - \alpha, \tag{1}$$

then the interval (L, U) is called a $(\beta, 1 - \alpha)$ two-sided tolerance interval.

- Robust Z-Score

Univariate Outlier Detection

- Percentile Bootstrapping
  - At each time point, calculate an lower and upper limit for the number of bookings which are classified as normal.
  - Bootstrap number of bookings at each time point, and find 2.5^{th}, and 97.5^{th} percentile of each bootstrap sample. Take median across bootstrap samples as lower/upper limit.

- Tolerance Intervals
  - Nonparametric or parametric approaches.
  - Require two parameters: the coverage proportion, β, and confidence level, $1-\alpha$.
  - For $X_1, X_2, \ldots, X_n$, a random sample from a population with distribution $F(X)$, if:

    $$\mathbb{P}\left(F(U) - F(L) > \beta\right) = 1 - \alpha, \tag{1}$$

    then the interval (L, U) is called a $(\beta, 1-\alpha)$ two-sided tolerance interval.

- Robust Z-Score
  - Let $y_i(t)$ be the cumulative number of bookings for flight i at time t. The robust Z-score can be calculated as:

    $$\tilde{Z}_i = \frac{0.6745\left(y_i(t) - \tilde{y}(t)\right)}{MAD(t)}, \tag{2}$$

    where $\tilde{y}(t)$ is the median number of bookings at time t across all flights. Flights with a robust Z-score above 3.5 are classified as outliers, (Iglewicz and Hoaglin (1993)).

Univariate Outlier Detection: Results

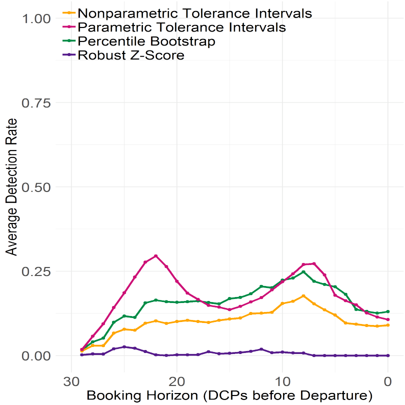

Figure: TDR

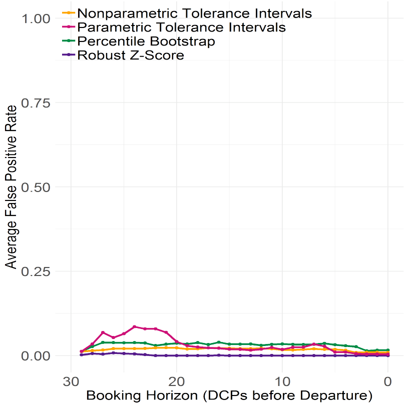

Figure: FPR

Univariate Outlier Detection: Results

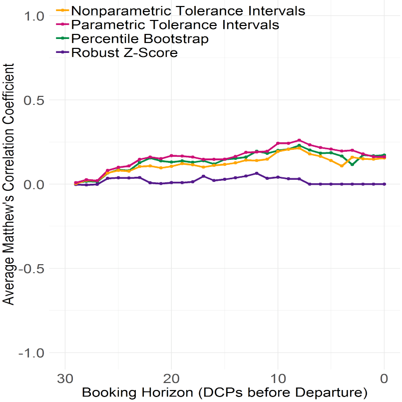

Figure: MCC

Matthew's Correlation Coefficient

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

MCC lies between ± 1:

- 1 = perfect classification
- 0 = equivalent to random classification
- -1 = perfectly incorrect classification

Univariate Outlier Detection: Results

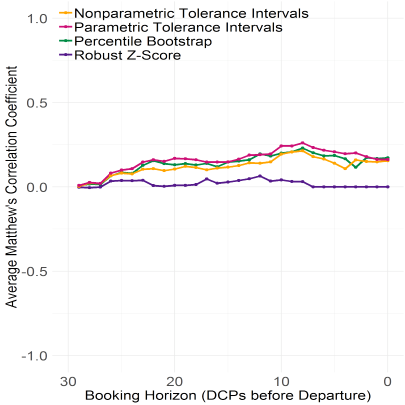

Figure: MCC

- Easier to detect outliers which have a large magnitude increase/decrease in demand.
- Easier to detect decreases in demand, as opposed to increases.
- Drop in ability to detect outliers shortly before departure, due to demand censoring from booking controls.

Multivariate Outlier Detection

- Distance Metrics
  - For each booking curve, calculate the mean (Euclidean or Manhattan) distance between it and booking curves for all other flights.
  - The distance between two vectors $\boldsymbol{x} = (x_1, x_2, \ldots, x_N)$ and $\boldsymbol{y} = (y_1, y_2, \ldots, y_N)$ is given by:
    - Euclidean: $D(\boldsymbol{x}, \boldsymbol{y}) = \left(\sum_{n=1}^{N} (x_n - y_n)^2\right)^{\frac{1}{2}}$
    - Manhattan: $D(\boldsymbol{x}, \boldsymbol{y}) = \sum_{n=1}^{N} |(x_n - y_n)|$

- K-Means Clustering
  - Split the booking curves into groups (clusters).
  - Iteratively minimise (Euclidean or Manhattan) distance between observations and cluster centres.
  - Those curves with a distance from their cluster centre above some threshold, are classified as outliers.

Multivariate Outlier Detection

- Distance Metrics
  - For each booking curve, calculate the mean (Euclidean or Manhattan) distance between it and booking curves for all other flights.
  - The distance between two vectors $\boldsymbol{x} = (x_1, x_2, \ldots, x_N)$ and $\boldsymbol{y} = (y_1, y_2, \ldots, y_N)$ is given by:

$$\text{Euclidean: } D(x, y) = \left(\sum_{n=1}^{N} (x_n - y_n)^2\right)^{\frac{1}{2}}$$

$$\text{Manhattan: } D(x, y) = \sum_{n=1}^{N} |(x_n - y_n)|$$

- K-Means Clustering
  - Split the booking curves into groups (clusters).
  - Iteratively minimise (Euclidean or Manhattan) distance between observations and cluster centres.
  - Those curves with a distance from their cluster centre above some threshold, are classified as outliers.

Multivariate Outlier Detection: Results

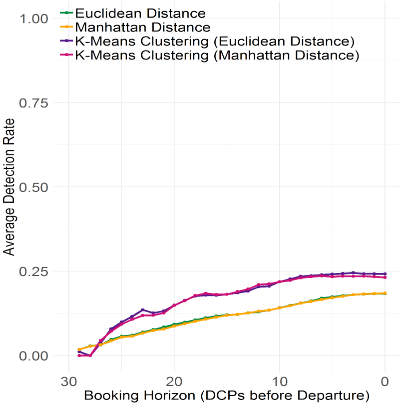

Figure: TDR

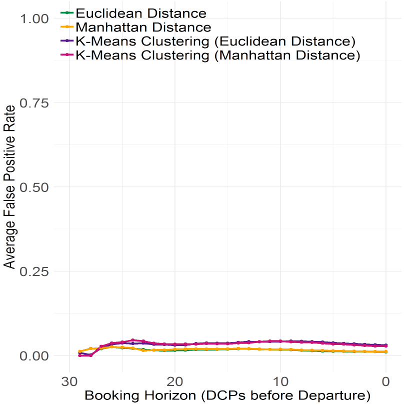

Figure: FPR

Multivariate Outlier Detection: Results

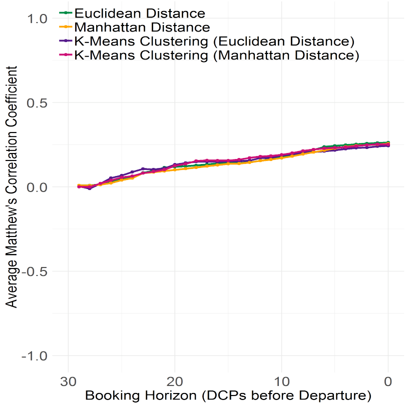

Figure: MCC

- Easier to detect outliers which have a large magnitude increase/decrease in demand.
- Easier to detect decreases in demand, as opposed to increases.
- Less of a drop in performance later in the booking horizon, as more previous information is used.

Multivariate Outlier Detection: Results

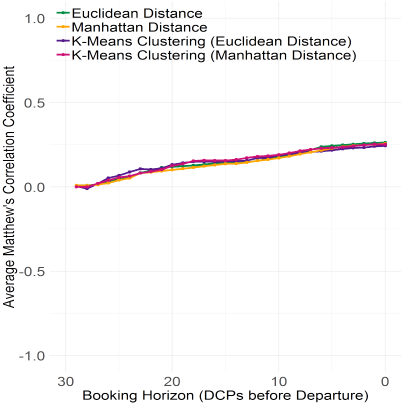

Figure: MCC

- Easier to detect outliers which have a large magnitude increase/decrease in demand.
- Easier to detect decreases in demand, as opposed to increases.
- Less of a drop in performance later in the booking horizon, as more previous information is used.

Functional Outlier Detection

- Functional Depth
  - A measure of the centrality, or ‘outlyingness’ of an observation with respect to a given dataset (López-Pintado and Romo (2009)).
    - For $\left\{x_i(t_j); i = 1, \ldots, n; j = 1, \ldots, m\right\}$, define the sample Fraiman-Muniz depth as:

$$SFMD_n(x_i) = \sum_{j=2}^{m} \Delta_j \left(1 - \left|\frac{1}{2} - F_{n,t_j}(x_i(t_j))\right|\right)., \quad i = 1, \ldots, n$$

  - Those curves with depths below some threshold are classified as outliers.

Functional Outlier Detection: Results

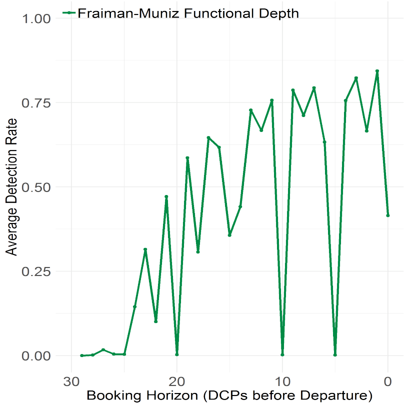

Figure: TDR

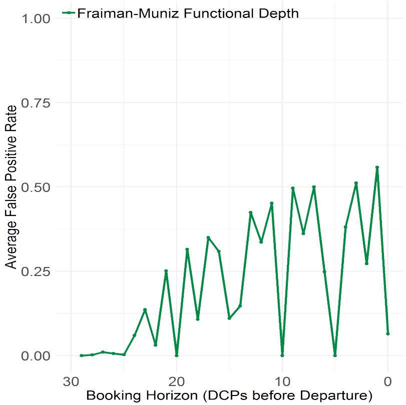

Figure: FPR

Functional Outlier Detection: Results

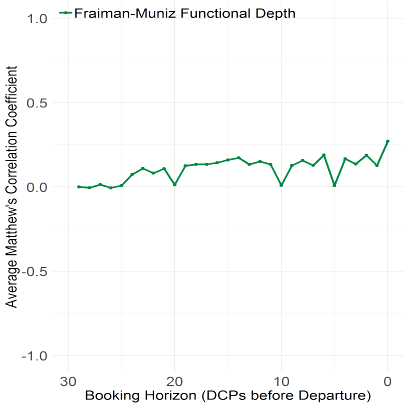

Figure: MCC

- Easier to detect outliers which have a large magnitude increase/decrease in demand.
- Difference between ability to detect positive and negative outliers is smaller.
- Unusual spikes in outlier detection performance.

Functional Outlier Detection: Results

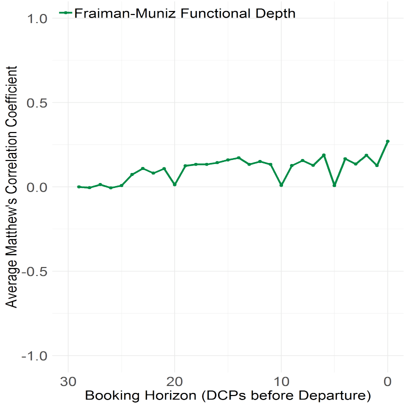

Figure: MCC

- Easier to detect outliers which have a large magnitude increase/decrease in demand.
- Difference between ability to detect positive and negative outliers is smaller.
- Unusual spikes in outlier detection performance.

Comparison of Methods

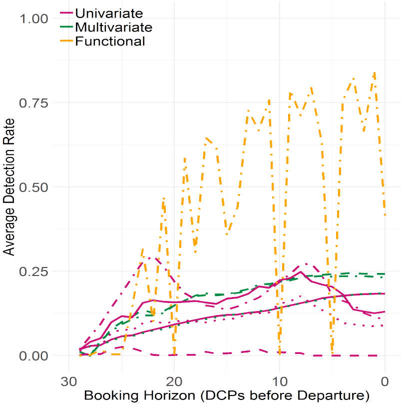

Figure: TDR

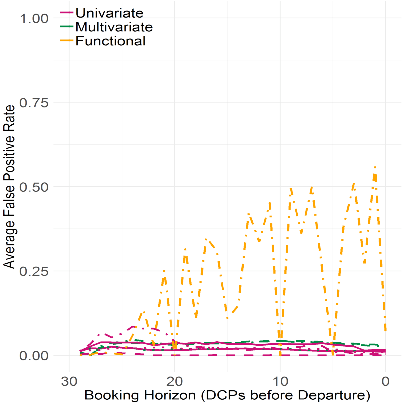

Figure: FPR

Comparison of Methods

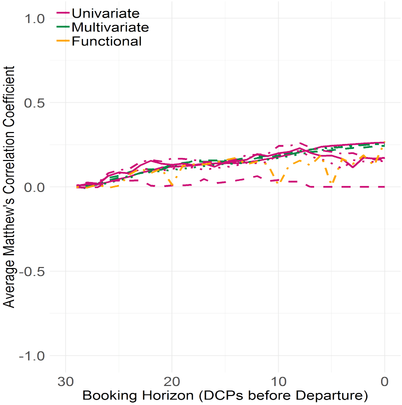

Figure: MCC

- Multivariate (clustering) approaches appear to perform best.
- Disadvantages include issues with high-dimensionality as number of DCPs increases, and specifying number of clusters in advance.
- Functional approaches have more scope for extension.

Comparison of Methods

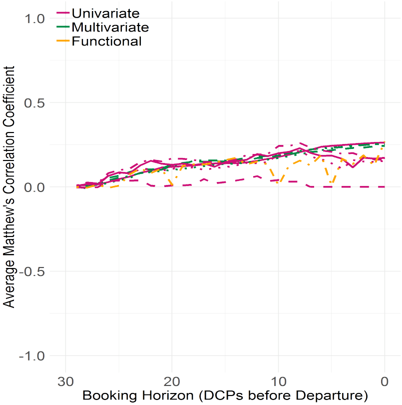

Figure: MCC

- Multivariate (clustering) approaches appear to perform best.
- Disadvantages include issues with high-dimensionality as number of DCPs increases, and specifying number of clusters in advance.
- Functional approaches have more scope for extension.

Conclusions

- Although the impact of outlier detection depends on the optimisation method, identifying situations where demand is not as expected is beneficial.
- Multivariate and functional approaches are more promising than univariate approaches to outlier detection.
- Demand censoring from booking controls creates issues in outlier detection.
- Outlier detection can be beneficial in identifying critical flights.

Conclusions

- Although the impact of outlier detection depends on the optimisation method, identifying situations where demand is not as expected is beneficial.
- Multivariate and functional approaches are more promising than univariate approaches to outlier detection.
- Demand censoring from booking controls creates issues in outlier detection.
- Outlier detection can be beneficial in identifying critical flights.

Conclusions

- Although the impact of outlier detection depends on the optimisation method, identifying situations where demand is not as expected is beneficial.
- Multivariate and functional approaches are more promising than univariate approaches to outlier detection.
- Demand censoring from booking controls creates issues in outlier detection.
- Outlier detection can be beneficial in identifying critical flights.

Conclusions

- Although the impact of outlier detection depends on the optimisation method, identifying situations where demand is not as expected is beneficial.
- Multivariate and functional approaches are more promising than univariate approaches to outlier detection.
- Demand censoring from booking controls creates issues in outlier detection.
- Outlier detection can be beneficial in identifying critical flights.

Further Work

- Extend the functional outlier detection approaches to incorporate forecasts, to improve difficulties with censoring.
- Take into account time dependence between curves, and include seasonality.
- Extend to a multivariate setting to jointly monitor booking curves and revenue curves.
- Investigate the impact of unexpected demand in the dynamic pricing setting.

Further Work

- Extend the functional outlier detection approaches to incorporate forecasts, to improve difficulties with censoring.
- Take into account time dependence between curves, and include seasonality.
- Extend to a multivariate setting to jointly monitor booking curves and revenue curves.
- Investigate the impact of unexpected demand in the dynamic pricing setting.

Further Work

- Extend the functional outlier detection approaches to incorporate forecasts, to improve difficulties with censoring.
- Take into account time dependence between curves, and include seasonality.
- Extend to a multivariate setting to jointly monitor booking curves and revenue curves.
- Investigate the impact of unexpected demand in the dynamic pricing setting.

Further Work

- Extend the functional outlier detection approaches to incorporate forecasts, to improve difficulties with censoring.
- Take into account time dependence between curves, and include seasonality.
- 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