



Outlier detection in network revenue management

Nicola Rennie¹ · Catherine Cleophas² · Adam M. Sykulski³ · Florian Dost⁴

Received: 29 June 2022 / Accepted: 8 March 2023 / Published online: 23 March 2023
© The Author(s) 2023

Abstract

This paper presents an automated approach for providing ranked lists of outliers in observed demand to support analysts in network revenue management. Such network revenue management, e.g. for railway itineraries, needs accurate demand forecasts. However, demand outliers across or in parts of a network complicate accurate demand forecasting, and the network structure makes such demand outliers hard to detect. We propose a two-step approach combining clustering with functional outlier detection to identify outlying demand from network bookings observed on the leg level. The first step clusters legs to appropriately partition and pools booking patterns. The second step identifies outliers within each cluster and uses a novel aggregation method across legs to create a ranked alert list of affected instances. Our method outperforms analyses that consider leg data without regard for network implications and offers a computationally efficient alternative to storing and analysing all data on the itinerary level, especially in highly-connected networks where most customers book multi-leg products. A simulation study demonstrates the robustness of the approach and quantifies the potential revenue benefits from adjusting demand forecasts for offer optimisation. Finally, we illustrate the applicability based on empirical data obtained from Deutsche Bahn.

Keywords Analytics · Forecasting · Outlier detection · Clustering · Network revenue management

1 Introduction and state-of-the-art

Revenue management (RM) is a challenging task for network service providers. The concept entails controlling the set of product offers over a fixed sales horizon such that, given the predicted demand for the offers, the expected revenue from selling a limited capacity is maximal. To thus maximise revenue, the firm has to forecast the expected demand for all products that require capacity on mutual resources.

✉ Catherine Cleophas
cleophas@bwl.uni-kiel.de

Extended author information available on the last page of the article

Examples include transport itineraries that cross several network legs, and hospitality offers that combine room availability for multiple nights.

Several existing contributions, e.g. Weatherford and Belobaba (2002) and Rennie et al. (2021), demonstrate the negative effects of inaccurate demand forecasts on revenue performance but neglect network effects. Motivated by this, we propose a new approach to detect outliers in network bookings, thereby supporting forecast corrections for improved network revenue management.

1.1 Terminology

For simplicity, we employ a transport-based terminology throughout this paper: a *leg* describes a direct non-stop connection between two stations in a network, and an *itinerary* is any combination of legs that can be jointly booked as one product. A *departure* describes a journey along a connected series of legs that leave the origin station at a unique time and date.

We denote the accumulation of *bookings* across the sales horizon as a *booking pattern*. We define an *outlier* as a booking pattern resulting from short-term systematic demand changes for one or several related network itineraries. These outliers occur when demand deviates from the baseline due to unforeseen events. For example demand increases for a specific destination affect the entire itinerary. In consequence, deviations in the booking patterns are observable on the legs arriving at that destination and in the feeder legs.

Capacity-based revenue management differentiates offers through *fare classes*. Fare classes describe combinations of fares and tariffs at which the firm offers a product. Customers booking a ticket for a specific departure may choose from several offered fare classes. For instance, the cheapest offer could be fare class ‘M’, costing 20 Euros and entailing a no-refund tariff.

1.2 Existing work

RM is a well-studied problem for many different products and services (Talluri and Van Ryzin 2004). Still, only recently the specific issues around network services and demand forecasting have come into focus. E.g. Klein et al. (2020) review how single-leg practices to RM generalise to the network setting. Weatherford (2016) surveys RM forecasting methods and focuses on airline itinerary-level forecasting.

So far, few authors have examined demand outliers in RM data. For historical hotel booking data, Weatherford and Kimes (2003) discuss a simple method of removing observations that are more than $\pm 3\sigma$ away from the mean. Rennie et al. (2021) apply functional analysis to detect outliers on individual legs. Neither, however, consider outliers affecting multiple legs of a network. In Azadeh et al. (2013), the authors identify outliers in network railway bookings via a simple rule to remove them before forecasting future demand. For a slightly different perspective, Kumar and Khani (2020) analyse transit demand for outliers to detect special events. Notably, existing research on outlier detection frequently focuses on binary outlier detection without regard for quantifying how critical an outlier is.

Practical network RM relies on manual forecast adjustments (Quante et al. 2009; Schütze et al. 2020). The previous research has shown that the resulting judgemental forecasts can be biased and even superfluous (Lawrence et al. 2006; De Baets and Harvey 2020). Perera et al. (2019) note that forecasting support tools can improve user judgement by reducing complexity for the analyst. Analysts' time is limited, so they cannot investigate every departure flagged as an outlier. For example Deutsche Bahn experts estimate that they can reasonably adjust less than 1% of forecasts. Therefore, ranking outliers by criticality is crucial.

Beyond RM, Barrow and Kourentzes (2018) propose a functional approach for outlier detection in call arrival forecasting without regard for network effects. General outlier detection in networks often focuses on identifying outlying parts of the network. Fawzy et al. (2013) use this approach in wireless sensor networks to find faulty nodes. Ranshous et al. (2015) consider the extension to identify outlying nodes when the network changes over time. Most research on dynamic networks concentrates on analysing a single time series connected to each node rather than a set of time series, as required when booking patterns are reported for multiple departures. Hyndman et al. (2016) note that the problem of identifying unusual time series within a collection is not as extensively studied as other outlier detection problems. In this paper, we benchmark the approach suggested by Hyndman et al. (2016), which employs principal component analysis (PCA), against our newly proposed approach.

1.3 Contribution

We shall study booking patterns that result when customers book not just a single resource (leg) but network products that require multiple resources (itineraries). Such booking patterns may be reported on the leg or the itinerary level; in this paper, we assume that they are reported per leg and departure. This applies in the case of Deutsche Bahn, which serves as a motivation and empirical demonstration for the work presented here.

Network effects challenge outlier detection in two ways: On the one hand, demand outliers on the itinerary level affect bookings on all legs included in the itinerary. On the other hand, such outliers may not be recognisable when only considering leg bookings independently, given the noise from other itineraries overlapping those legs. As a result, directly extracting outliers from booking data collected in an entire, realistically sized network is likely an intractable problem. To circumvent this problem, in this paper, we aggregate and analyse booking patterns from *legs* instead of *itineraries*, as this allows for computationally and statistically tractable network-wide outlier detection. In Sect. 6, we further discuss the choice of leg-level- vs. itinerary-level-based analysis and point out how our procedures could be adapted to itinerary-level outlier detection.

Our network outlier detection procedure: (i) clusters legs with similar booking patterns and (ii) detects joint outliers within each cluster to compile ranked alert lists of outlying departures and affected legs. Our methodology significantly improves outlier detection performance in a network setting versus alternative methods.

In more detail, our proposed approach first clusters legs by measuring the similarity of booking patterns via functional dynamical correlation (Dubin and Müller 2005). We suggest this measure for its freedom from restrictive assumptions. As the proposed approach is modular, other correlation measures could be used for the same end. In the second step, the proposed approach detects outliers from booking patterns within each cluster by combining the functional data analysis methods of Febrero et al. (2008); Hubert et al. (2012); Rennie et al. (2021) with a novel within-cluster aggregation, which generates a ranked alert list of outliers using extreme value theory. This alert list can help analysts to identify the need for further analysis and adjustments. We consider an outlier as more critical if it indicates a larger demand shift and if it is identified across multiple legs. Factors such as the average fare on legs where outliers are detected, or the revenue at risk from faulty forecasts could also be incorporated into the definition of an outlier's criticality.


Finally, analysts have several choices when tasked with forecast adjustment for network services. The best choice is not obvious, and we further quantify the impact of different potential adjustments on revenue in a simulation study, following concepts outlined in Kimms and Müller-Bungart (2007).

In summary, this paper contributes (i) a method for identifying network legs that will benefit from joint outlier detection and (ii) a method to aggregate outlier detection across any number of legs to create a ranked alert list. To thoroughly evaluate the proposed approach, we offer (iii) wide-ranging simulation studies to benchmark the method's outlier detection performance against and to quantify the potential revenue improvements from forecast adjustments and (iv) a demonstration of applicability on empirical railway booking data from Deutsche Bahn.


2 Method

Several network products may rely on common resources when demand concerns multiple legs at once. In the transport example, even passengers that booked different itineraries often have to traverse the same legs. Therefore, specific legs share common outliers, as, for example a sudden increase in demand from passengers travelling from one end of the network to attend an event at the other end would increase demand for each of the in-between legs. Neither considering each leg independently nor jointly considering the whole network will create the best results when the network spans multiple regions that differ strongly in demand—see Sect. 3.4 and Appendix 6. This raises the question of which legs to consider jointly for outlier detection.

To find an answer, in Sect. 2.1, we adapt a method by Zahn (1971) to *cluster* legs such that (i) legs in the same cluster share demand and can be considered jointly for outlier detection, and (ii) legs in different clusters experience distinct demand and should be considered separately. Subsequently, in Sect. 2.2, we suggest a method for analysing bookings within one such cluster. Based on this, we propose a method to rank departures by the severity of identified outliers.



(a) Original graph where nodes represent stations (b) Inverted graph where nodes represent legs



(c) Minimum spanning tree with edge weights (d) Clusters obtained in inverted graph

Fig. 1 Correlation-based minimum spanning tree clustering

2.1 Clustering legs using correlation-based minimum spanning trees

To cluster legs based on correlations in observed bookings, we first consider the network as a graph where nodes represent the stations and edges represent the legs of a journey. Figure 1a illustrates this on a simple network. To illustrate the concept, we rely on an example from the transport domain: In this example, two train lines (red and blue) intersect at two stations (B and C). The red train arrives at stations B and C before the blue train, which creates two possible transfer connections for passengers: (i) switch from red to blue at B and (ii) switch from red to blue at C. Transfers from the blue to red train are not feasible.

Standard graph clustering algorithms, as exemplified in Schaeffer (2007), seek to cluster the nodes of the graph. In contrast, we wish to cluster similar edges, which correspond to legs in the railway example (Fig. 1a). Hence, we invert the graph to make existing clustering algorithms applicable. In this inversion (Fig. 1b), the directed edges become nodes, e.g. the edge from A to B becomes

node AB. The inverted graph features an undirected edge between two nodes when:

- both legs are in the same train line and share a common station, e.g. legs CD and DE are connected through station D, or
- the legs are in different train lines but share a common transfer station where a connection is possible, e.g. leg FB (red line) and BC (blue line) are connected through station B. However, AB (blue line) and BC (red line) would not be connected by an edge as no connection can be made between them (as we have assumed the red train arrives at B and C before the blue train).

In theory, this transformation could also create edges between legs that share a common entry or exit node, e.g. FB (red line) and AB (blue line), or CG (red line) and CD (blue line). Given that such pairs of legs would never occur in the same itinerary, we would not expect demand outliers to affect both legs. Therefore, inserting an edge between them, potentially allowing them to be in the same cluster, is counter-intuitive. In addition, exploratory analyses of the empirical data found that correlations between these types of legs were very low across the entire network section.

The algorithm aims to assign those legs that experience *similar bookings* to the same cluster and those that experience *dissimilar bookings* to distinct clusters. A corresponding metric only needs to consider the similarity between adjacent legs that share a connecting station since edges do not otherwise exist in the inverted graph. We propose to quantify this similarity via the correlation between booking patterns.

To calculate correlations between booking patterns, we compute the functional dynamical correlation (Dubin and Müller 2005). Functional dynamical correlation is based on calculating scalar products between pairs of smoothed booking patterns; the appendix provides further details. We use the average of these paired correlations over time as the similarity measure between any two legs. Unlike more common statistical correlation measures, such as Pearson correlation, functional dynamical correlation does not assume a specific type of relationship between variables (e.g. linearity). It also accounts for the time dependency between observations within the booking horizon when the intervals between observations vary. For example in the empirical RM data analysed in Sect. 5, the time between observations decreases as the departure date approaches. Further, alternative measures for calculating correlations from functional data (such as functional canonical correlation) often make restrictive assumptions, which real data do not fulfil (He et al. 2003).

We benchmark the clustering algorithm under alternative correlation measures in Appendix 8.

To represent the relationship between legs in the network, i.e. the nodes in the inverted graph, we attach weights to the edges in the inverted graph. These weights are interpreted as distances: A higher edge weight indicates that the connected nodes are more dissimilar. Therefore, an applicable weight function should be non-negative. Further, the weight function needs to ensure that any negatively correlated legs are marked as more dissimilar. Even though a negative correlation may imply that outlier demand jointly affects both legs, we expect it to affect negatively correlated

legs differently. Therefore, these require different adjustments from an analyst and should be in different clusters. To satisfy these requirements, we define the edge weights as:

$$w_{(ij,jk)} = 1 - \rho(ij,jk), \quad (1)$$

where $\rho(ij,jk)$ is the correlation between bookings on legs ij and jk . Though the use of functional dynamical correlation as a measure of similarity between time series is not new, its application as an edge weight in a network setting, to our knowledge, is novel.

To allow for irregular cluster shapes, we recommend a minimum spanning tree (MST) algorithm (Prim 1957). For example in Fig. 1b, a cluster may include AB and DE because they are in the same line, rather than clustering AB and FB. Minimum spanning tree approaches work well for clusters with irregular boundaries (Zahn 1971). Alternative clustering approaches (such as k -means) often assume a specific shape of clusters (spherical, for k -means). MST-based clustering approaches also do not assume that clusters are of similar sizes (Peter and Victor 2010). This makes them particularly suitable for transportation networks constructed as a series of interlocking lines, where the points of intersection are often not equally spaced. For example MST-based approaches have previously been used in optimising layouts of railway networks (Liang et al. 2020).

A *spanning tree* of a graph is a subgraph that includes all vertices in the original graph and a minimum number of edges, such that the spanning tree is connected. Then, the MST is the spanning tree with the minimum summed edge weights—see Fig. 1c. Since the inverted graph is weighted, we use Prim's algorithm (Prim 1957) to calculate the MST—Appendix 2 provides a detailed introduction. Any one-to-one transformation of the weight function, $w_{(ij,jk)}$, will produce an identical minimum spanning tree.

There are two approaches to obtaining clusters from an MST: (i) pre-defining the number of clusters as k and removing the $k - 1$ edges with the highest weight; or (ii) setting a threshold for the edge weights and removing all edges with weights above some threshold, creating an emergent number of clusters. Here, we implement the threshold-based approach, ensuring that each cluster has the same minimum level of correlation. In contrast, setting the number of clusters in advance could result in very heterogeneous levels of correlation across clusters. Further, setting k too low may result in legs with dissimilar features being grouped together. We apply a threshold correlation of 0.5—the level at which legs are more correlated than they are not. This corresponds to a transformed edge weight of 0.5. In the example given in Fig. 1c, this means removing all legs with a weight above 0.5, resulting in the three clusters shown in Fig. 1d. The choice of this clustering threshold will impact the number of alert lists produced. Therefore, we recommend considering factors such as staffing resources and any current (informal) network clustering when choosing this threshold.

While the outlier detection procedure described next applies to individual clusters, it does not require a particular clustering approach. Hence, other implementations may employ alternative approaches, as reviewed in Schaeffer (2007). In

particular, depending on the business context of the network service, alternative clustering algorithms may be more appropriate. The network topology should drive the choice of which clustering algorithm is most appropriate: That topology may differ, e.g. when considering airlines versus bike rentals versus railways, as discussed in Rennie et al. (2022). The choice of an MST-based approach, which often returns linear clusters, is appropriate for the railway application motivating the paper at hand, given the linear nature of the underlying network structure. We further evaluate the performance of MST clustering in this application in Appendix 8.

Furthermore, edge-based clustering could replace the graph inversion and node-based clustering presented here. However, the literature on edge-based clustering is far more limited, and such approaches tend to improve the visualisation of networks with a very high number of edges by reducing the number of edge crossings rather than grouping together the most similar edges (Qu et al. 2007). In contrast, inversion and node-based clustering aim to group network legs that exhibit the highest degree of similarity. However, alongside these advantages, there may be some drawbacks. The node-based approach requires deciding on criteria to select edges to include in the inverted graph.

2.2 Detecting outliers in clusters of legs


Given established clusters, we propose identifying demand outliers within each cluster and quantifying their severity to provide a *ranked alert list* of departures. The previously described clustering allows for processing the outlier detection in parallel for separate clusters, enabling efficient computing.

To identify which departures to include in the alert list, we consider the *functional depth* of the booking patterns, as in Rennie et al. (2021). This step could also rely on other measures of exceedance, including univariate 'threshold' approaches, which look at aggregated bookings and ignore the distribution of bookings over time. We propose to rely on functional depth, as the previous work has found this to be the most effective as an outlier detection mechanism (Rennie et al. 2021).

To compute the functional depth, consider N departures observed over L legs. Let $\mathbf{y}_{nl} = (y_{nl}(t_1), \dots, y_{nl}(t_T))$ be the booking pattern for the n^{th} departure on leg l , observed over T booking intervals $t_1, \dots, t_T$. Let $\mathcal{Y}_l$ be the set of N booking patterns for leg l . For each leg and departure, calculate the functional depth (d_{nl}) given the related booking patterns following the approach given in Hubert et al. (2012) and detailed in Appendix 3. The functional depths take on positive values, with smaller values of the depths relating to more outlying booking patterns.

For each leg l , we calculate a threshold for the functional depth using the approach of Febrero et al. (2008). This method (i) resamples the booking patterns with probability proportional to their functional depths (such that any outlying patterns are less likely to be resampled), (ii) smooths the resampled patterns, and (iii) sets the threshold C_l as the median of the 1^{st} percentiles of the functional depths of the resampled patterns. Here, we use the 1^{st} percentile of the depths as the default threshold, as this has been found to work well in practice (Febrero et al. 2008; Rennie et al.

Fig. 2 z_n as defined in Eq. (3) for a four-leg section of the Deutsche Bahn network



2021). Booking patterns with a functional depth below the threshold C_l are classed as outliers. We explore alternative threshold choices in Appendix 2.

To create ranked alert lists, we first define z_{nl} to be the normalised difference between the functional depth and the threshold:

$$z_{nl} = \frac{C_l - d_{nl}}{C_l}. \quad (2)$$

This transforms the depth measure d_{nl} into a measure of *threshold exceedance*. Values of z_{nl} greater than zero relate to booking patterns classified as outliers. Normalising by the threshold, C_l , ensures that the values of z_{nl} are comparable between different legs.

Next, we define the sums of threshold exceedances across legs:

$$z_n = \sum_{l=1}^L z_{nl} \mathbb{1}_{\{z_{nl} > 0\}}. \quad (3)$$

We sum only those values of z_{nl} that are greater than zero to avoid outliers being masked when they occur only in a subset of legs. This sum implicitly accounts for both the size of an outlier—larger outliers further exceeding the threshold, resulting in larger values of z_{nl} —and for the number of legs where a departure is classified as an outlier (by summing a larger number of nonzero values). To provide an example, Fig. 2 shows those values of z_n that exceed zero for a four-leg section of the Deutsche Bahn network as discussed further in Sect. 5.2. These values of z_n correspond to departures where the booking pattern for *at least* one leg is identified as an outlier. In contrast, all other departures have no detected outliers in any leg such that $z_n = 0$.

To create a ranked list of outlier departures, i.e. those with a nonzero-sum of threshold exceedances, we assign a severity θ_n . A higher value of θ_n indicates that the departure is more likely to be affected by extreme outlier demand and hence should be targeted first by RM analysts.

To model threshold exceedances, we turn to extreme value theory (EVT)—a branch of statistics that deals with modelling rare events occurring in the tails of a distribution. Given that outliers are unusual events, which occur in the tails of distributions, EVT is a clear direction to turn to for modelling outliers—see Talagala et al. (2019). There are two common approaches to EVT: (i) block maxima, which




Fig. 3 Distribution of z_n values from Fig. 2

examine the maximum value in evenly-spaced blocks of time, e.g. annual maxima, and (ii) peaks over the threshold, which examines all observations that exceed some threshold (Leadbetter 1991). The generalised Pareto distribution (GPD) is commonly used to model the tails of distributions in the peaks over threshold approach (Pickands 1975). Motivated thus, we fit a generalised Pareto distribution (GPD) to the sum of threshold exceedances given in equation (3). The GPD has three parameters with probability density function:

$$f(x|\mu, \sigma, \xi) = \frac{1}{\sigma} \left(1 + \frac{\xi(x - \mu)}{\sigma} \right)^{-\frac{1}{\xi} - 1}, \quad (4)$$

for

$$x \in \begin{cases} [\mu, \infty) & \xi \geq 0 \\ [\mu, \mu - \frac{\sigma}{\xi}] & \xi < 0. \end{cases} \quad (5)$$

Here, μ specifies the location, σ the scale, and ξ the shape of the distribution. We fit the parameters using maximum likelihood estimation (Grimshaw 1993) via the R package POT (Ribatet and Dutang 2019). A kernel density estimate of the empirical distribution of $z_n > 0$ from Fig. 2 is shown in Fig. 3a. The resulting fitted GPD is shown in Fig. 3b. As the further analysis in Appendix 17 shows, the GPD fit appears reasonable compared to the empirical distribution.

Two common issues arise in fitting GPDs: (i) the choice of threshold and (ii) the independence of the data points. When the threshold is too low, the assumption of a GPD no longer holds; when it is too high, there are too few data points to fit. We select a threshold of 0, i.e. we fit the GPD to values of $z_n > 0$. Rather than change the threshold at the GPD level, we control the number of observations the GPD is fitted to by varying the percentile used for the individual leg thresholds, C_l . We choose C_l as suggested by Febrero et al. (2008) and find that this choice works well

and provides sufficient outlying points to fit a GPD in both simulated and empirical data.

To account for the second issue, applications of extreme value theory frequently first *decluster* the peaks over the threshold to ensure independence between observations (Fawcett and Walshaw 2007). To that end, the analysis may only consider the maximum of two peaks within a small time window. For transport departures, it is theoretically possible that observed outliers may be dependent; e.g. increased demand caused by Easter affects not only Easter Sunday but also the surrounding days. However, similar outliers may also result from independent events. As we aim to identify outlying departures rather than the underlying events, this argument causes us not to decluster here.

We define θ_n as the non-exceedance probability given by the CDF of the GPD:

$$\theta_n = F_{(\mu, \sigma, \xi)}(z_n) = \begin{cases} 1 - \left(1 + \frac{\xi(z_n - \mu)}{\sigma}\right)^{-\frac{1}{\xi}} & \xi \neq 0 \\ 1 - \exp\left(-\frac{(z_n - \mu)}{\sigma}\right) & \xi = 0 \end{cases} \quad (6)$$

Formally, θ_n is the probability that, given an outlier occurs, the sum of threshold exceedances is at least as large as z_n . Thus, it is *not* the probability that a departure is an outlier. However, we use this non-exceedance probability as a measure of outlier *severity* on a scale of 0–1.

Departures with functional depths that do not fall below the threshold on any legs carry a severity of zero, i.e. they are classified as regular departures. It is conceivable to estimate the uncertainty of θ_n (Smith 1985) to determine further levels of criticality, e.g. if there are several departures with the same outlier severity, the one with the smallest uncertainty would be ranked first. However, given the continuous nature of the data, it is unlikely that multiple departures carry an identical severity. Hence, we leave uncertainty estimation to future research.


From the severity defined in equation (6), we construct a *ranked alert list* containing all departures with a nonzero outlier severity. Although functional depth could be directly used to construct the ranked alert list, computing the severity provides a measure of the difference between ranks and is more easily interpreted by analysts. The top 8 ranked outliers relating to Fig. 2, are shown in Table 1.

In practice, RM analysts' time and resources allow them to examine and adjust controls or forecasts only for a limited number of suspicious booking patterns. Those departures that (i) exceed the functional depth threshold in only one leg or (ii) exceed the threshold only to a small degree have lower but strictly nonzero severity. These outliers are most likely false positives and potentially waste analysts' time. Hence, we suggest limiting the length of the list in practice.

To limit the length of the alert list, we might (i) only include departures if their severity is above some threshold or (ii) set a maximum length. Since we wish to control the number of alerts an analyst will receive, we analyse outlier detection performance as dependent on the maximum length of the alert list. Recall that we classify departures as outliers if and only if their outlier severity exceeds zero. Therefore, if the required length of the alert list exceeds the number of identified outliers,

Table 1 Ranked alert list for cluster = {AB, BC, CD, DE}

Ranking	Departure	Severity	Legs with $z_{nl} > 0$
1	11 May 2019	0.985	AB, BC, CD, DE
2	26 October 2019	0.960	AB, BC, CD, DE
3	09 June 2019	0.942	AB, BC, CD, DE
4	01 June 2019	0.922	AB, BC, CD, DE
5	13 July 2019	0.874	AB, BC, CD, DE
6	13 April 2019	0.865	CD, DE
7	02 February 2019	0.864	CD, DE
8	05 October 2019	0.857	AB, BC, CD, DE
$\vdots$	$\vdots$	$\vdots$	$\vdots$

Fig. 4 Four-leg cluster, dotted lines indicate 10 possible itineraries

we do not include further departures. Appendix 7 features further result on the outlier detection performance when varying the outlier severity threshold.

3 Outlier detection performance

We first implement a simulation study to evaluate the outlier detection performance given known outliers. By varying the demand for itineraries in one cluster, we create outliers that are observable on both the leg and network levels.

The simulation models a network consisting of five stations and four legs, as shown in Fig. 4, mirroring the structure of an empirical railway network cutout. The network includes 10 possible itineraries represented by $\mathcal{O} = \{AB, AC, AD, AE, BC, BD, BE, CD, CE, DE\}$. On each itinerary, the firm offers seven fare classes. In this model, a fare class describes a particular price or fare associated when booking a ticket to travel the itinerary in that class. There are no additional restrictions differentiating classes.

3.1 Demand settings

Extending the demand model described in Rennie et al. (2021) to the network setting, the simulation generates booking requests per customer type i according to

a non-homogeneous Poisson process, where the arrival rate per itinerary o , $\lambda_{i,o}(t)$, at time t , is given by:

$$\lambda_{i,o}(t)|(D_o = d_o) = d_o \times \phi_{io} \frac{t^{a_{io}-1}(1-t)^{b_{io}-1}}{B(a_{io}, b_{io})}. \quad (7)$$

here ϕ_{io} is the fraction of customers of type i and $D_o \sim \text{Gamma}(\alpha_o, \beta_o)$ with probability density function:

$$f(d_o | \alpha_o, \beta_o) = \frac{\beta_o^{\alpha_o}}{\Gamma(\alpha_o) d_o^{\alpha_o-1} e^{\beta_o d_o}}, \quad (8)$$

where a_{io} and b_{io} define are the parameters of a Beta distribution which defines how customers arrive over time. We generate demand over a horizon of 3600 time slices to ensure $\lambda_{i,o}(t) < 1$. This level of detail is required to accurately parameterise the dynamic program for bid price control. The resulting bookings are aggregated into 18 booking intervals.

As in Rennie et al. (2021), we consider differentiated demand from two customer types represented by the set $\mathcal{I} = \{1, 2\}$. We assume that customers book the cheapest available fare class and differ in price sensitivity. We define p_{ijo} as the probability that a customer of type i pays up to fare class j on itinerary o . By combining demand from two customer types that differ in price sensitivity with offers that depend on the current set of offered classes, we mimic a realistic price effect: Offer prices result from the cheapest class currently offered by the firm, as customers will buy the cheapest available class. When a customer's willingness to pay does not equal or exceed the price of the cheapest available class, they do not buy; hence, their price sensitivity translates to decreased demand. Note that the price–demand response depends on the itinerary and time in the booking horizon.

Combining this demand model with the given network creates 210 demand parameters. Table 2 provides a full list of parameter values and interpretations of each parameter. We set the parameters to mirror common RM assumptions (Weatherford and Bodily 1992): (i) valuable customers from type 1 book later than customers from type 2, (ii) customers book earlier for longer journeys, and (iii) customers are willing to pay a higher fare class if they are travelling further. Most passengers book tickets boarding at A and leaving at E; this ensures the correlation between the legs exceeds 0.5 and guarantees that the legs are correctly modelled in the same cluster as detailed in Appendix 19.

We validate that the functional dynamical correlation between the four legs for simulated data is comparable to empirical railway data as detailed in Appendix 19. We generate all regular demand based on these parameters.

The simulation excludes trend and seasonality to evaluate outlier detection approaches in a best-case scenario. In other words, if an algorithm fails on observations from stationary demand, it will likely not perform better given more demand variability. However, additional results based on simulation data that do feature seasonality can be found in Appendix 10.

3.2 Outlier generation and evaluation

We generate demand-volume outliers by changing the gamma distribution parameters that govern the total demand level according to equations (7) and (8). The previous work found that the proportion of outliers had little effect on outlier detection performance in the single-leg case (Rennie et al. 2021). Therefore, we generate booking patterns for 500 departures per demand setting, with 1% of departures experiencing outlier demand. That is, we generate 495 departures from the regular demand distribution and five outliers from a set of twelve outlier distributions where the mean has shifted by $\pm 10\%$, $\pm 20\%$, $\pm 30\%$, $\pm 40\%$, $\pm 50\%$, and $\pm 60\%$. For every shift in the mean, we reduce the variance of the outlier demand distribution by 80%. This still results in an overall increase in the variance of total demand in the presence of outliers but also ensures that we sample sufficiently outlying demand values. Outliers may also occur due to factors such as changes in arrival times or changes in customers' willingness to pay. Rennie et al. (2021) provide results on how the performance of functional depth varies under these different types of outliers. Here, we focus on the different types of outliers caused by varying network effects. In all cases, we consider the application of the outlier detection procedure to the constrained demand—applying the approach directly to the booking patterns without applying any unconstraining approaches first. The problem of unconstraining is one of the major challenges of demand forecasting for revenue management and is beyond the scope of this paper.

We differentiate outlier scenarios in terms of the affected network components. Firstly, we evaluate a scenario where outlier demand affects all network itineraries. We consider the case where each outlier is randomly drawn from one of the twelve outlier distributions, resulting in outliers from a mixture of different distributions. This lets us test whether the ranking of the alert list mirrors the outliers' underlying degree of demand deviation. Then, we consider each of the twelve outlier distributions in isolation to assess the detection sensitivity. Secondly, we evaluate a scenario where outliers only affect a single itinerary. This evaluates the benefits of clustering multiple legs.

In Appendix 6, we consider the practically relevant case of outliers affecting a subset of itineraries and provide further details on all simulation experiments.

Each combination of outcomes can be classified into one of four categories: (i) assigning a nonzero outlier severity to a genuine outlier creates a true positive (TP); (ii) assigning a zero outlier severity to a regular observation creates a true negative (TN); (iii) assigning a nonzero outlier severity to a regular observation creates a false positive (FP); and (iv) assigning a zero outlier severity to a genuine outlier creates a false negative (FN). This classification enables us to compute the true-positive rate (TPR) for the top R ranked departures in the alert list:

$$\text{TPR}_R = \frac{\text{TP}_R}{\text{TP} + \text{FN}}, \quad (9)$$

where TP_R is the number of true positives in the top R departures. The true-positive rate lies between 0 and 1, where 1 means all genuine outliers were identified. We evaluate performance across 1000 stochastic simulations.

In an ideal setting, the alert list should feature, from top to bottom, large outliers and, subsequently, smaller outliers. Therefore, we also use the distribution of outliers within the ranked alert list to evaluate how well the method ranks the most critical outliers.

3.3 Benchmarked outlier detection approaches

For benchmarking, we term the newly proposed approach FD+Agg and compare it to two alternatives from the literature: Principal component analysis combined with high-density regions (PCA+HDR) as inspired by Hyndman et al. (2016), and the leg-based functional depth analysis as proposed in Rennie et al. (2021).

3.3.1 Comparison with PCA+HDR

This benchmark (i) computes features (e.g. mean, variance, and curvature) of the booking patterns for the total demand in a cluster; (ii) uses PCA (Yang and Shahabi 2004) to identify the first two principle components from the features; and (iii) uses HDR, a density-based approach (Hyndman 1996), to find the ν points with the lowest density in the first two principal components. These points are classified as outliers. Extended details of the method, including the list of features, can be found in Appendix 6. This method provides an ordering of the outliers but not a severity measure, as illustrated by Fig. 5.

3.3.2 Comparison with non-ranked, single-leg approaches

To highlight to critical features of FD+Agg, we benchmark (i) the use of severity measures to rank outliers and (ii) the inclusion of network effects. To isolate the effects of each of these features, we perform two separate benchmark tests:

We evaluate the effect of ranking outliers by measuring the increase in precision when ranking outliers. For example we consider the precision in the top 5 ranked departures versus 5 randomly chosen departures with nonzero outlier probabilities (i.e. as in Rennie et al. (2021)). The change in precision when considering the top R departures, $\Delta(\text{Precision})_R$, is given by:

$$\Delta(\text{Precision})_R = \frac{\text{TP}_R}{\text{TP}_R + \text{FP}_R} - \frac{\text{TP}_{R(\text{random})}}{\text{TP}_{R(\text{random})} + \text{FP}_{R(\text{random})}}, \quad (10)$$

where $\text{TP}_{R(\text{random})}$ is the number of true positives in a random selection of R departures with nonzero severity, and $\text{FP}_{R(\text{random})}$ is defined analogously for false positives.

We quantify the value of accounting for network effects by computing ranked alert lists for each leg in isolation. We then compare the true-positive rates to the aggregated, network-driven approach presented in this paper.




Fig. 5 Performance and benchmark comparison with PCA+HDR for demand-volume outliers in all itineraries, showing improved performance

3.4 Detecting outliers in multiple legs

As a first experiment, we consider the scenario where outlier demand equally affects all itineraries and legs within the cluster. For this scenario, Fig. 5a illustrates how the true-positive rate (TPR) increases when ranking outliers for different lengths of the alert list. The red line indicates the number of genuine outliers. The true-positive rates for our method (denoted as FD+Agg) are promising, with a TPR of around 0.2 for a list length of 1. Since there are five genuine outliers, this indicates that a genuine outlier is almost always ranked top. Results under different functional depth thresholds are given in Appendix 2.




Fig. 6 Change in precision from ranking detected outliers in FD+Agg as dependent on the length of the alert list

3.4.1 PCA+HDR benchmark results

The PCA+HDR approach requires a given number of outliers to detect, v , as input. Therefore, we compare the performance of the benchmark method under different choices of v to FD+Agg.

Figure 5b shows that the true-positive rate achieved by FD+Agg consistently exceeds that achieved by PCA+HDR. To achieve the same level of the true-positive rate, PCA+HDR would need to classify around 250 departures (i.e. 50%) as outliers. In comparison, FD+Agg achieves this rate starting at about 30 classified outliers. We consider this a successful validation of the effect of ranking outliers in FD+Agg. Appendix 1 lists these results in tabular format.

Figure 5c shows the distribution of each outlier magnitude in the alert lists. Under FD+Agg, the modes of the distributions generally fall where they should, as larger outliers are ranked higher. The smaller variance in the ranking of the larger magnitude outliers indicates that they are easier to detect. The higher variance of the medium-sized outliers can be explained as the ranking of a medium-sized outlier is dependent on which other types of outliers occur: If there is a large and a medium outlier, the medium outlier is ranked lower; if there is a small and a medium outlier, the medium outlier is ranked higher. The distribution of outliers detected by PCA+HDR, shown in Fig. 5d, also has the modes in the correct order. However, there is much more overlap between the distributions, showing its inability to correctly rank the outliers.

3.4.2 Comparison with non-ranked approach

Figure 6a highlights how the precision improves when ranking outliers instead of listing them in random order. Ranking particularly improves precision when the alert list




Fig. 7 True-positive rate for single itinerary outliers when applying FD+Agg versus detection on isolated legs

covers only a small number of departures. As domain experts indicate that analysts cannot target more than 1% of departures, ranking focuses resources and thereby provides large benefits in practice. Nevertheless, Fig. 6a (when contrasted with Fig. 5a) also highlights the trade-off between reducing the number of false alerts and identifying all outliers. A shorter length of alert list increases precision but reduces the true-positive rate.

The increase in precision from applying our method compared to PCA+HDR is similar to the increase in precision from the inclusion of the ranking (see Fig. 6b). This suggests that PCA+HDR performs reasonably well in terms of outlier detection, but poorly in terms of ranking the outliers.

3.4.3 Comparison with single-leg approach

Figure 7 shows the true-positive rate when a ranked alert list is computed for each leg in isolation versus in the proposed aggregated manner. Here, we consider outlier demand generated by a 50% increase in the affected legs as an illustrative example. We analyse detection performance by breaking down results in terms of which itinerary the outlier demand is generated in. We show only the results relating to itineraries AB, AC, AD, and AE. Figure 26 in Appendix 5 details results for the further itineraries yielding similar conclusions.

For results, when outlier demand is generated across combinations of itineraries, refer to Appendix 6.

In all cases, the true-positive rate for clusters is higher than in any of the individual legs. This is because when considering the leg's bookings in isolation under outlier demand that affects multiple legs, the noise from other itineraries prevents detecting the outlier in every leg. However, clustering increases the number of detected genuine outliers.

Aggregation is most beneficial when the outlier demand affects the most legs. In our example, this applies when itinerary AE experiences outlier demand, as shown in Fig. 7a. The lower true-positive rates in legs AB and DE result because different combinations of itineraries also utilise these legs. The aggregation is less beneficial when outlier demand affects an itinerary consisting of only one or two legs since we aggregate the analysis across legs that are actually not affected by outlier demand. However, there is a modest gain in true-positive rate even in this case—compare Fig. 7(c). This is due to the knock-on effects of decreased capacity on the affected legs, impacting the bid prices for any itineraries which include these legs. For some lengths of the alert list, the leg-level true-positive rates are higher than the aggregated approach, due to false positives from unaffected legs being included in the list. However, even for itinerary AB (Fig. 7d), where false positives from unaffected legs are most likely, the difference is small and cancelled out by the overall increase in true-positive rate.

3.4.4 Sensitivity to different magnitudes of outliers

To better understand outlier detection performance, we break down the results by the magnitude of outliers in Fig. 8.

When outliers result from minor changes in demand levels, they are difficult to detect, resulting in low true-positive rates. Given the significant overlap between the distribution of outlier demand with a 10% change in magnitude and that of regular demand, this is to be expected. Therefore, 10% demand changes effectively provide a lower bound on how big an outlier needs to be in order to be detected.

As the magnitude of the outliers increases, they become easier to detect and true-positive rates are higher, with peak rates reached with shorter alert lists. Thus, genuine outliers are more likely to be ranked higher when they are caused by larger demand changes. For demand decreases of at least 50%, the true-positive rate is very close to the optimal detection rate. Negative demand outliers are slightly easier to detect than positive demand outliers, meaning shorter alert




Fig. 8 Sensitivity of true-positive rate from FD+Agg under different magnitudes of homogeneous demand-volume outliers




Fig. 9 Sensitivity of precision from FD+Agg under different magnitudes of homogeneous demand-volume outliers

lists are required. This is due to the demand censoring imposed by the booking controls and capacity restrictions.

Figure 9 shows the precision gap over randomly ordered lists. Once more, larger magnitude outliers result in larger precision improvements from ranking, while detecting minor outliers gains little over random selection. Similarly, we observe that detecting negative demand outliers gains slightly more precision in comparison with detecting positive outliers of the same magnitude. Additional results regarding false discovery rates are available in the appendix.

4 Simulation study: forecast adjustments

To evaluate the implications of adjusting the demand forecast for further planning steps, we simulate network demand and the optimisation of offered fare classes over the booking horizon. We list and explain all parameters determining the settings in the simulation study in Appendix 7. In this section, we first detail how the simulated RM system uses the demand forecast to compute revenue-optimal offers based on bid prices. In that, it follows a widely implemented industry standard. Subsequently, we describe alternative strategies that analysts may apply to adjust demand forecasts based on identified outliers. Finally, by comparing revenue gained from offers based on different adjusted demand forecasts under the same simulated outlier demand, we highlight the effects of adjustments as dependent on outlier scenarios.

4.1 Network revenue management system

The simulated RM system controls the offered set of fare classes per itinerary to optimise expected revenue. To that end, it implements a dynamic program to compute bid prices per leg and sums them up per itinerary following the methodology described in Strauss et al. (2018) and detailed in the appendix. To test for the sensitivity of results with regard to the revenue optimisation, we compared two industry standards, the leg-based EMSR heuristic as introduced in Belobaba (1987) and dynamic programming in initial simulations studies not further documented here. The results showed that, for the given demand model, the choice of optimisation approach had little effect on the quality of the outlier detection.

The bid price indicates the marginal difference between the value of selling a seat in the current time period and that of reserving it to sell in a future time period. The RM system only offers fare classes where the revenue from a booking exceeds the bid price. Thus, as an RM term, bid prices do not denote the customer's bid but indicate the minimum price a fare class must carry to be included in the offer set. From those classes in the offer set, customers only consider the cheapest offer. Bid prices depend on the time until departure, unsold capacity, and expected demand. Note that in the examples given here, we consider a single capacity per leg, not differentiating, for example 1st or 2nd class compartments with separate capacities.

Booking patterns result as customers arrive and decide to book one of the offered fare classes. The firm does not report booking patterns for each individual itinerary, but only records them on the leg level.

The dynamic programme relies on a given set of expected demand arrival rates per leg l , fare class j , and time slice t of the booking horizon. In the simulation, we derive expected demand arrival rates from our knowledge of the underlying demand model. Arrival rates for each leg l and fare class j are given as

$$\hat{\Lambda}_{j,l}(t) = \sum_{o \in \mathcal{O}_l} \sum_{i \in \mathcal{I}} p_{i,j,o} \lambda_{i,o}(t), \quad (11)$$

where $\lambda_{i,o}(t)$ is the arrival rate of customers of type i requesting itinerary o , and $\mathcal{O}_l$ is the set of itineraries which include leg l . This creates an artificially accurate


demand forecast. Deriving the demand forecast from the actual demand parameter values ensures that the estimation of revenue loss caused by undetected outliers is not affected by flawed forecasts (see Sect. 4.3). In practice, demand parameter values are not known but are estimated based on previously observed demand and time-series forecasting. A recent survey of related research contributions can be found in Banerjee et al. (2020), while Fiig et al. (2019) represent an example of the ongoing discussion on the link between forecast accuracy and RM performance.

4.2 Forecast adjustments for outlier demand

One aim of identifying outlier demand in booking patterns is to support analyst adjustments in RM systems. Without such adjustments, offers would be optimised for a regular demand forecast and thereby not be fit for maximising revenue under outlier demand. This raises the difficulty of predicting the consequences of analyst adjustments throughout the network. As a step in this direction, we analyse a best-case scenario, assuming that the adjustment is made with foresight before the start of the booking horizon. We compare the revenue under three different adjustments:

- *Adjustment 1 (conservative)* Adjust only forecasts of affected single-leg itineraries. E.g. for an outlier creating additional demand for itinerary AC, increase the forecasts of itineraries AB and BC.
- *Adjustment 2 (aggressive)* Adjust forecasts of all itineraries that include at least one of the affected legs. E.g. for additional demand for itinerary AC, adjust all itineraries, including either leg AB or leg BC—i.e. itineraries AB, AC, AD, AE, BC, BD, and BE.
- *Adjustment 3 (balanced)* Adjust forecasts of affected single-leg itineraries and the *cluster-spanning* itinerary; in this case, AE. E.g. for additional demand for itinerary AC, adjust itineraries AB, BC, and AE. The motivation for adjusting AE (ahead of other itineraries) is that, in general, this will be the most popular itinerary in the cluster.

These three adjustments are not the only choices available to analysts. However, they represent options that stretch across the spectrum of how fully network effects should be considered. Adjustments 1 (conservative, leg-based adjustments only) and 2 (aggressive, all potential network effects) are the two extremes. Adjustment 3 (balanced) is a compromise, which is more conservative than Adjustment 2 but still identifies the itinerary most likely to be the source of outlier demand. Further options would be to include more than just the cluster-spanning itinerary in an alternative to Adjustment 3, but this leaves another choice of which itineraries to prioritise. As a lower bound, we compute the revenue when **no adjustment** is made. As an upper bound, we implement an **oracle adjustment**, i.e. only adjusting the forecasts of affected itineraries. We compare the revenue as the level of outlier demand ranges from -60% to +60% of the average leg demand.



(a) Four leg itinerary (AE): mean revenue (left), revenue by demand change (right)




Fig. 10 Revenue under under different forecast adjustments; the subtitle indicates the actual outlier source

4.3 Experimental results: revenue benefits

Figure 10 shows the revenue generated by outlier demand for each of the three adjustments. We show the results for four of ten itineraries contained within these four legs in Fig. 4. The results for the other six itineraries are similar. Appendix 12 further details these results as well as results on adjustments after outlier detection.

When outlier demand affects all four legs in the cluster (Fig. 10a), any type of adjustment is always better than no adjustment. Besides the oracle, the best choice is Adjustment 3, i.e. the balanced approach, which adjusts the forecasts of the cluster-spanning itinerary and the individual leg. Adjustment 3 is able to obtain, on average, 87% of the additional revenue gained under the oracle adjustment. Similar results are obtained when the outlier demand affects three legs (Fig. 10b).

When outlier demand affects only a single-leg itinerary (Fig. 10d), the conservative Adjustment 1 and the oracle adjustment coincide. The aggressive Adjustment




Fig. 11 Distribution of the number of legs per booked itinerary from Deutsche Bahn data

2 yields less revenue than no adjustment. For example although leg AB is correctly adjusted, the erroneous adjustment to itineraries AC, AD, and AE results in incorrect forecasts for legs BC, CD, and DE. The asymmetry between adjustment to positive and negative outlier demand is due to the level of demand being bounded from below by 0. Similar results emerge when the outlier affects only two of the affected legs (Fig. 10c), though the negative consequences of over-adjusting all potentially affected itineraries are less severe, as this causes fewer superfluous adjustments.

The negative impact of adjusting unaffected itineraries highlights the importance of correctly clustering legs ahead of outlier detection. The closer the outlier demand itinerary is to the cluster-spanning itinerary, the less risky it is to adjust all affected itineraries within a cluster, and the more benefit can be gained from doing so. From a managerial perspective, the *best* adjustment (other than the oracle) depends on the firm's objective. To maximise revenue when the most common outlier (e.g. itinerary AE) occurs, the balanced Adjustment 3 is preferable. Conversely, if the objective is to minimise risk to revenue even in the more unlikely scenarios (e.g. an outlier in itinerary AB), conservative Adjustment 1 is preferable. Overall, however, there are clear benefits from forecast adjustment.

5 Empirical study

To demonstrate the practical applicability of the proposed clustering and outlier detection, we apply it to a set of empirical data obtained from Deutsche Bahn. This data set features only bookings of the 2nd class compartment, such that all bookings on one leg require capacity from the same compartment. The Deutsche Bahn long-distance network consists of over 1000 train stations, letting the provider offer more than 110,000 direct origin—destination combinations. The numbers grow further when accounting for alternative transfer itineraries and multiple daily departures. Figure 11 shows the empirical distribution of the number of legs included

in itineraries that passengers booked in November 2019. Only 7% of passengers booked single-leg itineraries, whereas almost half of all booked itineraries span five or more legs.

5.1 Clustering legs in the Deutsche Bahn network

5.1.1 Small network subsection

First, we consider a section of the Deutsche Bahn railway network that consists of two intersecting train lines over a total of 27 stations and 28 legs—see Fig. 12. The red train arrives at the connecting stations before the blue train. Hence, the network offers three transfer connections: changing from red to blue at either Fulda, Kassel-Wilhelmshöhe, or Göttingen. This creates 240 potential travel itineraries. For each leg in this network section, Deutsche Bahn records 359 booking patterns for departures between December 2018 and December 2019. Each booking pattern ranges over 19 booking intervals; the first observation occurs 91 days before departure.



We first apply the correlation-based clustering approach of Sect. 2.1, using a threshold of 0.5, such that only legs with a minimum correlation of 0.5 can be in the same cluster. In Fig. 12a, coloured bubbles indicate the four resulting clusters: Each train line splits into one large and one small cluster.

To evaluate clustering on empirical data, where the true underlying demand for each itinerary is unknown, we use the network topology to check whether the resulting clusters are plausible. To that end, we propose the following set of rules:

- Different train lines must belong to different clusters. Even when passengers can transfer between lines, we expect relatively few passengers to make the same connection. Further, it makes sense to consider train lines separately for forecasting and analyst interventions.
- Train lines are further split into separate clusters on either side of a major station. As many passengers leave the train at a major station and many *different* passengers board, we shall assume a relatively small proportion of passengers book itineraries that pass a major station. Similarly, given that itinerary demand share is driven by which journeys are most common, and passengers often either board or alight at a major station, it is intuitive to have a cluster that contains the legs between major stations.

Deutsche Bahn assigns an ordinal indicator of importance to each station, ranging from 1 to 7. We define a *major station* to be in *Category 1*. The entire Deutsche Bahn network includes 21 major stations, whereas the considered network section includes nine major stations. Figure 12b highlights major stations in grey and shows the clusters resulting from the above rules.

The correlation-based clustering returns four clusters, whereas the rule-based clustering returns nine. Nevertheless, the resulting clusters share similar features. Firstly, the two distinct train lines end up in different clusters in either approach. For legs in distinct train lines, correlation tends to be higher between legs that share a

(a) Correlation-based clustering, $\rho \geq 0.5$ 

(b) Rule-based cluster

Fig. 12 Comparison of correlation-based and rule-based clustering of Deutsche Bahn network

transfer station, but not to a convincing extent—the correlation is at most 0.22. A correlation threshold of 0.27 creates two clusters (one for each train line). Secondly, the breakpoints for the correlation-based approach are a subset of the breakpoints, i.e. major stations, in the rule-based approach. We conclude that the correlation-based approach achieves similar results as the rule-based approach without requiring expert input.

We can formally compare clustering results using the **Normalised Mutual Information (NMI)** (Amelio and Pizzuti 2015). The NMI is 1 if two clusterings are identical, and 0 if they are completely different.




Fig. 13 Comparison of rule-based and correlation-based clustering in a two-line railway network


Figure 13a shows the NMI between the correlation- and rule-based approaches while varying the threshold in the correlation-based approach from 0 to 1. This shows that both approaches achieve similar results, with an NMI reaching 0.899. The approaches are generally more similar at higher correlation thresholds (around 0.7) since the rule-based approach generally creates more clusters. Figure 13b compares the number of clusters of the two approaches—as the correlation threshold changes, the number of clusters ranges from 1 (everything in a single cluster) to 28 (each leg in its own cluster), demonstrating the flexibility of the correlation-based approach.

5.1.2 Large network subsection


We extend the empirical study to five train lines to further demonstrate the complexity that considering the network structure brings to clustering and outlier detection, and show the scalability of the approach. The five-line network consists of 40 stations with 63 legs. As shown in Fig. 14, there are often multiple train lines which cover the same leg or may travel in the opposite direction. As the larger size of the network makes visualisation more difficult, in Fig. 14, stations are represented by circles, with major stations highlighted in black.

Figure 14(a) shows the results of the correlation-based clustering with a default threshold $\rho = 0.5$. This results in nine clusters, with two train lines each forming their own cluster containing all legs. The breakpoints of the clusters occur at major stations, as also previously seen for two train lines. The pattern of breaking clusters at major stations persists as the correlation threshold is varied. In comparison, the output of the rule-based clustering shown in Fig. 14(b) results in 24 clusters, with many being of size 1.

In these empirical studies, we applied rule-based clustering only to evaluate the plausibility of the results from correlation-based clustering. We do not advocate



(a) Correlation-based clustering, resulting in 9 clusters with breakpoints occurring at major stations



(b) Rule-based clustering for a five line network, resulting in 24 clusters

Fig. 14 Comparison of rule-based and correlation-based clustering in a five-line railway network




Fig. 15 Booking patterns for each leg

for it as a method in itself. A rule-based approach, where the clusters are based on domain experts' categorisations, would not be able to respond to the evolving importance of stations across different train lines and departure times. Notably, the correlation-based method not only uncovers major stations but rather identifies legs where multi-leg itineraries cause similar booking patterns and thus could change and adapt over time. We further evaluate clustering performance in a simulation study, where the itinerary-level demand is known, in Appendix 8. The results in the remainder of the paper rely on correlation-based clustering.

5.2 Detecting outliers in the Deutsche Bahn data

Having established clusters, we apply outlier detection independently to each cluster. To exemplify this on empirical data, we apply the outlier detection procedure to a representative four-leg cluster from the Deutsche Bahn network. Applying the proposed outlier detection approach to empirical data cannot precisely judge detection accuracy, given there are no labelled data on genuine outliers. However, this analysis demonstrates the full process of outlier detection on empirical data including, e.g. seasonality and underlines practical implications.

For this analysis, we consider a cluster of four legs from the Deutsche Bahn network with stations anonymised and denoted by A, B, C, D, and E. This cluster results from applying the correlation-based clustering to a new section of the Deutsche Bahn network to Fig. 12.

Figure 15 shows the booking patterns for each of the four legs; bookings are scaled to be between 0 and 1. From initial visual inspection, the structure of the booking patterns appears similar, with some obvious outliers appearing across multiple legs.

To pre-process the data for outlier detection, we transform the booking patterns by applying a functional regression model (Ramsay and Silverman 1997). We then




Fig. 16 Threshold exceedances per leg, z_{nl}

apply the outlier detection to the residual booking patterns. In this pre-processing, we correct for three factors: (i) the departure day of the week; (ii) the departure month of the year; and (iii) the length of the booking horizon.¹

The functional regression fits a mean function to the booking patterns for each different factor in the model. Table 9 in Appendix 13 compares models including different factors. Let $y_{nl}(t)$ be the n^{th} booking pattern for leg l . Then:

$$\begin{aligned}
 y_{nl}(t) = & \beta_{0l}(t) + \beta_{1l}(t)\mathbb{1}_{Mon_{nl}} + \beta_{2l}(t)\mathbb{1}_{Tue_{nl}} + \beta_{3l}(t)\mathbb{1}_{Wed_{nl}} + \\
 & \underbrace{\beta_{4l}(t)\mathbb{1}_{Thu_{nl}} + \beta_{5l}(t)\mathbb{1}_{Fri_{nl}} + \beta_{6l}(t)\mathbb{1}_{Sat_{nl}}}_{\text{Departure Day of the Week}} + \\
 & \beta_{7l}(t)\mathbb{1}_{Jan_{nl}} + \beta_{8l}(t)\mathbb{1}_{Feb_{nl}} + \beta_{9l}(t)\mathbb{1}_{Mar_{nl}} + \\
 & \beta_{10l}(t)\mathbb{1}_{Apr_{nl}} + \beta_{11l}(t)\mathbb{1}_{May_{nl}} + \beta_{12l}(t)\mathbb{1}_{Jun_{nl}} + \beta_{13l}(t)\mathbb{1}_{Jul_{nl}} + \\
 & \underbrace{\beta_{14l}(t)\mathbb{1}_{Aug_{nl}} + \beta_{15l}(t)\mathbb{1}_{Sep_{nl}} + \beta_{16l}(t)\mathbb{1}_{Oct_{nl}} + \beta_{17l}(t)\mathbb{1}_{Nov_{nl}}}_{\text{Departure Month of the Year}} + \\
 & \underbrace{\beta_{18l}(t)\mathbb{1}_{Shorter Horizon_{nl}}}_{\text{Length of Booking Horizon}} + e_{nl}(t).
 \end{aligned} \tag{12}$$

where, e.g. $\mathbb{1}_{Mon_{nl}} = 1$ if departure n relates to a Monday, 0 otherwise. In this model, $\beta_{0l}(t)$ represents the average bookings for Sunday departures in December, with a regular length of booking horizon, and $\beta_{pl}(t)$ for $p > 0$ represent deviations from this mean pattern. The $\beta_{pl}(t)$ are functions of time, which allows for relationships between factors to evolve over the booking horizon. Given that functional depths are calculated independently for each leg, we apply the regression model independently for each leg. The resulting residuals are included in Appendix 14, Fig. 38.

¹ Deutsche Bahn offers a regular booking horizon of 6 months, with the first observation of bookings occurring around 3 months before departure. Due to schedule changes, shorter booking horizons of 3 months apply for departures from mid-December to mid-March.




Fig. 17 Outliers detected in booking patterns

Functional regression preserves the correlation between different legs, as verified in Appendix 19, Table 11b. The clustering approach can consider either the correlations between the booking patterns or the residual booking patterns. Given that the functional depth (the basis for the outlier detection) is calculated on the residuals, we suggest using the correlation between residual patterns to define the clusters. For this data set, the same clusters resulted in either case.

We calculate the functional depth of each booking pattern and compute the threshold as described in Sect. 2.2. We then transform the depths as per equation (2) to obtain z_{nl} , as shown in Fig. 16. The sums of threshold exceedances, z_n , were shown earlier in Fig. 2, with the empirical distribution and fitted generalised Pareto distribution shown in Figs. 3a and 3 b, respectively.

Figure 17 highlights the outliers detected in each leg in pink while depicting outliers detected in other legs but *not* in that leg in blue. Regular patterns are grey.

Of the 40 outliers (11% of departures) detected across all legs, 23 outliers (almost 60%) could be attributed to known events or holidays. When considering only the top 10 outliers, the percentage rose to 70%. A further departure detected as an outlier had been previously flagged by Deutsche Bahn. The firm implemented a booking stop to control sales on that departure for multiple connected legs. Appendix 18 provides further details on the distribution of identified outliers across legs.

6 Conclusion and outlook

In this paper, we proposed a two-step method for (i) clustering legs in a mobility network that could benefit from joint outlier detection and (ii) detecting outlying demand within such clusters. Furthermore, the proposed method, FD+Agg, ranks identified outliers according to their severity, creating an alert list to aid analysts in prioritising demand forecast adjustments.

The simulation study demonstrated the robustness of the method in a range of outlier demand scenarios. It highlighted that aggregating the analysis across

clustered legs improves both detection rate and precision. Further, the ranked alert list often correctly identified the most critical outliers. The advantages of the proposed approach became particularly clear when benchmarking its true-positive rate, distribution of outliers across ranks, and precision, against that from a combination of principal component analysis and high-density regions (PCA+HDR) from Hyndman et al. (2016), and on the non-ranked, leg-based method proposed in Rennie et al. (2021).

Furthermore, we implemented a simulated revenue management system to measure the potential revenue benefits of identifying and adjusting for demand outliers in a network setting by applying forecast adjustments across a cluster of legs. This analysis showed that taking into account, the similarity of the legs can improve revenue in most scenarios. In the less likely scenario where only one or two legs of a cluster are affected by outlier demand, risk-averse firms may prefer individual leg-level adjustments.

Finally, by applying the proposed approach to empirical booking data collected by Deutsche Bahn, we demonstrated its applicability and scalability to the type of data observed in practice. In particular, we used this analysis to showcase the expected cluster results and to demonstrate how to account for additional practical considerations, such as trend and seasonality. Note that once the clustering has been performed, the outlier detection can be performed in parallel within each cluster. Therefore, our methodology is scalable to a much larger data set, such as the entire Deutsche Bahn long-distance train network. Such an analysis is not included in this paper as, beyond giving excessive insight into confidential company data, the research insight to be gained from visualising even more complex network cut-outs is limited.

The remainder of this section discusses design choices taken in the research documented here, related limitations, and open research challenges.

Leg- versus itinerary-level data: Our proposed method aggregates and analyses booking patterns from *legs* instead of *itineraries* based on three considerations. First, when an extensive network features many possible itineraries, most individual itineraries only receive a small share of bookings, challenging any data analysis—the study described in Appendix 11 evaluates such a case. Though the outlier detection may perform well if there are a sufficient number of bookings for a given itinerary, only considering such itineraries risks systematically ignoring outliers from smaller itineraries and feeder legs. Secondly, when offering many potential itineraries, providers rarely store all booking patterns per itinerary. For example capacity-based RM, as described in Strauss et al. (2018), frequently considers leg booking patterns to ensure capacity availability on each leg of a requested itinerary. Accordingly, the methodology proposed here is compatible with capacity-based RM. Finally, even in the idealised case of having large volumes of stored itinerary-level data for every possible itinerary in the network, then running outlier detection algorithms quickly becomes computationally infeasible as the number of possible itineraries grows rapidly with the size of the network. Detecting outlying clusters of legs, rather than individual itineraries, overcomes all three challenges, as we have demonstrated in this paper. We do, however, note that the outlier detection methodology we propose could be applied

directly to itinerary data without performing clustering. However, we only recommend this for densely booked itineraries, as otherwise, zero-inflated data can induce inferior results. We explored this further in Appendix 11.

Constrained versus unconstrained bookings: Observed bookings are constrained by any revenue management controls that were in place at the time of booking, whereas revenue optimisation models rely on unconstrained demand forecasts (Talluri and Van Ryzin 2004, Chapter 9.4). To represent this practice, we analysed constrained bookings in this paper and analysed the effect of adjusting unconstrained forecasts in the computational study. In that vein, further research could also consider the impact of applying the analysis to constrained observations, as showcased here, versus applying it to unconstrained demand estimates, which are frequently used for demand forecasting.

Implications for decision support: Further research is needed to consider the practical aspects of outlier detection from the perspective of decision support. Outliers manifest as changes in arrival rate, price elasticity, or other variables that affect bookings. Outliers can be caused by stochasticity but also by changes in demand patterns as a result of external factors, such as specific events. Complemented by further analysis, successful outlier detection could have three potential uses for RM: 1) detecting outliers early within the booking horizon through online analysis as proposed in Rennie et al. (2021), allowing for rapid interventions; 2) removing any detected outliers from training data for demand forecasting to improve results on predicting reference demand curves; and 3) if outliers can be attributed to specific events, the forecast model could be extended to include such events. Outlier detection can have broader benefits for operational planning in transportation networks, helping service providers to avoid overcrowding and delays. To realise such benefits, future research should particularly focus on effective ways to visualise outliers in networks and to communicate alert lists to planners. To further support analysts in their decision-making, additional measures could be included in the alert list. These might include average fare in the affected cluster, potential revenue loss if the outlier is not accounted for, or the outlier severity resulting from running the outlier detection procedure on *revenue* (instead of booking) patterns. An interesting avenue of further research would be to incorporate a feedback element whereby analysts mark outlier alerts as useful or not useful. A supervised learning approach, e.g. one-class classifiers, could then be combined with our proposed outlier detection routine to filter out false alerts. Analysts could additionally include feedback on the quality of the clustering approach.

Clustering methodology: Investigating the use of alternative clustering approaches is of interest—especially where the clusters are likely to be of different structures compared to the rail industry, e.g. in the airline industry where hub and spoke networks are more common than lines. Whilst this paper relied on clustering to improve outlier detection, we believe that the clustering approach is a useful contribution in and of itself. For example clustering presents additional research avenues such as its application to improving network-level forecasting; supporting the planning for future new stations; and evaluating how the transport network structure is changing over time or defining different travel zones. Finally, further research opportunities lie in considering how the success of network outlier detection

depends on the network structure. This paper featured examples from transport, specifically railway networks. Other application areas of RM, such as hotels, where correlation is induced by bookings for multiple consecutive nights, feature sparser or structurally different service networks.

Appendix A Additional details of method

Appendix A provides additional details on the proposed method described in Sect. 2, including the specifics of the correlation-based minimum spanning tree clustering, and the calculation of the functional depths.

A.1 Functional dynamical correlation

Let $y_{n,ij}(t)$ be the total observed bookings for the n^{th} departure on leg ij up to booking interval t , and similarly for $y_{n,jk}(t)$. The functional dynamical correlation between the booking patterns $y_{n,ij}(t)$ and $y_{n,jk}(t)$ is:

$$\rho_n(ij, jk) = \mathbb{E} \langle y_{n,ij}^*(t), y_{n,jk}^*(t) \rangle. \quad (13)$$

where

$$\langle y_{n,ij}^*(t), y_{n,jk}^*(t) \rangle = \int y_{n,ij}^*(t) y_{n,jk}^*(t) w(t) dt, \quad (14)$$

and $w(t)$ is a weight function that accounts for the time gap between observations. Here, $y_{n,ij}^*(t)$ is a standardised version of $y_{n,ij}(t)$:

$$y_{n,ij}^*(t) = \frac{y_{n,ij}(t) - M_{ij} - \mu_{ij}(t)}{\left[\int \{y_{n,ij}(t) - M_{ij} - \mu_{ij}(t)\}^2 w(t) dt \right]^{1/2}}, \quad (15)$$

where $\mu_{ij}(t)$ is a mean function, and:

$$M_{ij} = \langle y_{n,ij}(t), 1 \rangle. \quad (16)$$

The functional dynamical correlation is then the average across all N departures:

$$\rho(ij, jk) = \frac{1}{N} \sum_{n=1}^N \rho_n(ij, jk). \quad (17)$$

A.2 Prim's algorithm

Prim's algorithm is a greedy algorithm with the following basic steps. Assuming the original graph G has $V(G)$ vertices.

- Initialise the MST, T , with the edge with minimum weight and the two vertices it connects. Let $V(T)$ be the number of edges in T .
- While $V(T) < V(G)$:
 - go through the remaining edges in G in order from smallest to largest weights, until one is found that is connected to T , but does not form a circuit (i.e. the edge does not form a loop such that T is no longer a tree).
 - Add this edge (and the vertices it connects) to T .

More computationally efficient algorithms exist but given the reasonable size of the graphs considered, and more specifically their sparsity (very few stations are adjacent), the computational time is reasonable using Prim's algorithm.

A.3 Functional depth

The functional halfspace depth is given by:

$$d_{nl}(\mathbf{y}_{nl} \in \mathcal{Y}_i; \alpha) = \sum_{j=1}^T w_{\alpha}(t_j) HD_j(\mathbf{y}_{nl}(t_j)), \quad (18)$$

where, using $t_{\tau+1} = t_{\tau} + 0.5(t_{\tau} - t_{\tau-1})$, the weights $w_{\alpha}(t_j)$ are, according to Hubert et al. (2012):

$$w_{\alpha}(t_j) = \frac{(t_{j+1} - t_j) \text{vol} [\{\mathbf{x} \in \mathbb{R}^k : HD_j(\mathbf{x}) \geq \alpha\}]}{\sum_{j=1}^T (t_{j+1} - t_j) \text{vol} [\{\mathbf{x} \in \mathbb{R}^k : HD_j(\mathbf{x}) \geq \alpha\}]}, \quad (19)$$

where $\alpha \in (0, 0.5]$, with a default value of $\alpha = 1/T$. The sample halfspace depth of a K -variate vector x at time t_j is given by (Hubert et al. 2012):

$$HD_j(\mathbf{y}_{nl}(t_j)) = \frac{1}{N} \min_{\mathbf{u}, \|\mathbf{u}\|=1} \#\{\mathbf{y}_{nl}(t_j), n = 1, \dots, N : \mathbf{u}^T \mathbf{y}_{nl}(t_j) \geq \mathbf{u}^T \mathbf{x}\} \quad (20)$$

A.4 Normalised mutual information

For a graph containing M legs, the mutual information between two clusterings $\mathcal{A}$ and $\mathcal{B}$ of the M nodes in the inverted graph is defined as:

$$I(\mathcal{A}, \mathcal{B}) = \sum_{a=1}^{|\mathcal{A}|} \sum_{b=1}^{|\mathcal{B}|} \frac{|\mathcal{A} \cap \mathcal{B}|}{M} \log \left(|\mathcal{A} \cap \mathcal{B}| \frac{M}{M_a M_b} \right), \quad (21)$$

where M_a is the number of nodes in the a^{th} cluster of clustering $\mathcal{A}$, and similarly for M_b . The normalised mutual information (NMI) between two clusterings is defined as (Amelio and Pizzuti 2015):

$$NMI(\mathcal{A}, \mathcal{B}) = \frac{2I(\mathcal{A}, \mathcal{B})}{H(\mathcal{A}) + H(\mathcal{B})}, \quad (22)$$

where $H(\mathcal{A})$ is the entropy (a measure of uncertainty) defined as:

$$H(\mathcal{A}) = - \sum_{a=1}^{|\mathcal{A}|} \frac{M_a}{M} \log \left(\frac{M_a}{M} \right). \quad (23)$$

$NMI(\mathcal{A}, \mathcal{B}) = 1$ if $\mathcal{A}$ and $\mathcal{B}$ are identical, and 0 if they are completely different.

Appendix B Details of computational study

Appendix B contains additional details of the simulation set-up described in Sect. 3, including the computation of the bid prices, and a validation of the chosen parameter values.

B.1 Dynamic programming for bid price control

From Talluri and Van Ryzin (2004), let x be the remaining capacity, and define $V_t(x)$ denote the value function at time t . Define $R(t)$:

$$R(t) = \begin{cases} r_j & \text{if request for fare class } j \text{ arrives in interval } t \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

where r_j denotes the revenue from accepting a request for fare class j . The probability that $R(t) = r_j$ is equal to the arrival rate for fare class j at time t . Note the arrival rates are such that at most one request arrives in each time period. Define:

$$u = \begin{cases} 1 & \text{if request for fare class } j \text{ arrives and is accepted} \\ 0 & \text{otherwise} \end{cases} \quad (25)$$

We wish to maximise the combined revenue in the current time period, and the revenue to come in future time periods:

$$\max_{u \in \{0,1\}} (R(t)u + V_{t+1}(x - u)) \quad (26)$$

The Bellman equation for $V_t(x)$ is:

$$V_t(x) = \mathbb{E} \left[\max_{u \in \{0,1\}} \{R(t)u + V_{t+1}(x - u)\} \right] \quad (27)$$

$$= V_{t+1}(x) + \mathbb{E} \left[\max_{u \in \{0,1\}} \{(R(t) + \Delta V_{t+1}(x))u\} \right] \quad (28)$$

$$V_t(x) = \sum_{j=1}^{|J|} \lambda_j(t) \max\{r_j - \Delta V_{t+1}(x), 0\} \quad (29)$$

where $\lambda_j(t)$ is the arrival rate of demand for fare class j in interval t :

$$\lambda_j(t) = \sum_{i \in I} \lambda_{i,o}(t) p_{ijo}, \quad (30)$$

and $\Delta V_{t+1}(x) = V_{t+1}(x) - V_{t+1}(x-1)$ is the marginal cost of capacity in the next time period. The problem is solved with backwards recursion, with the following boundary conditions apply:

$$V_{T+1}(x) = 0, x = 0, 1, \dots, C \quad (31)$$

$$V_t(0) = 0, t = 1, \dots, T \quad (32)$$

These ensure that (i) no revenue can be generated beyond the booking horizon, i.e. after departure; and (ii) that no further revenue can be generated if there is no capacity remaining. The bid price at time t with remaining capacity x is given by $\Delta V_t(x)$.

B.2 Details of benchmark method

We use the method proposed by Hyndman et al. (2016) as a benchmark comparison for our proposed method in Sect. 3. The method works as follows:

- Define the total demand booking patterns as the sum of the demand for each leg within the cluster.
- Compute f features of the n total demand booking patterns. Features include mean, variance, first-order autocorrelation, trend, linearity, seasonality, peak, trough, entropy, lumpiness, spikiness, change in variance, and Kullback–Leibler score, among others. See Hyndman et al. (2016) for a full list.
- Apply principal component analysis (PCA) as per Yang and Shahabi (2004) to determine the first two principle components, i.e. those that explain the most variance.
- Use a density-based multidimensional approach (Hyndman 1996) to find points in the first two principal components with the lowest density.
- The nu points with the lowest densities relate to the departures which are classified as outliers.

B.3 Parameter values for simulation study


The parameter valued used to generate the demand in the computational study are outlined below.

Table 2 Regular demand generation parameter values


Parameter	Value	Effect of parameter
$\alpha = \{\alpha_{AB}, \alpha_{AC}, \alpha_{AD}, \alpha_{AE}, \alpha_{BC}, \alpha_{BD}, \alpha_{BE}, \alpha_{CD}, \alpha_{CE}, \alpha_{DE}\}$	$\alpha = \{32, 14, 14, 180, 4, 4, 14, 4, 14, 32\}$	Parameters of the gamma distribution which controls the level of total demand across all fare classes and customer types such that the mean demand for itinerary o is: $\mathbb{E}(D_o) = \frac{\alpha_o}{\beta_o}$
$\beta = \{\beta_{AB}, \beta_{AC}, \beta_{AD}, \beta_{AE}, \beta_{BC}, \beta_{BD}, \beta_{BE}, \beta_{CD}, \beta_{CE}, \beta_{DE}\}$	$\beta = \{1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1\}$	Parameters of beta distribution which controls the arrival times of type 1 customers
$a_1 = \{a_{1AB}, a_{1AC}, a_{1AD}, a_{1AE}, a_{1BC}, a_{1BD}, a_{1BE}, a_{1CD}, a_{1CE}, a_{1DE}\}$	$a_1 = \{5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5\}$	
$b_1 = \{b_{1AB}, b_{1AC}, b_{1AD}, b_{1AE}, b_{1BC}, b_{1BD}, b_{1BE}, b_{1CD}, b_{1CE}, b_{1DE}\}$	$b_1 = \{2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2\}$	
$a_2 = \{a_{2AB}, a_{2AC}, a_{2AD}, a_{2AE}, a_{2BC}, a_{2BD}, a_{2BE}, a_{2CD}, a_{2CE}, a_{2DE}\}$	$a_2 = \{2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2\}$	Parameters of beta distribution which controls the arrival times of type 2 customers
$b_2 = \{b_{2AB}, b_{2AC}, b_{2AD}, b_{2AE}, b_{2BC}, b_{2BD}, b_{2BE}, b_{2CD}, b_{2CE}, b_{2DE}\}$	$b_2 = \{2, 3, 5, 7, 2, 3, 5, 2, 3, 2, 2\}$	
$p_{ijo} = \{p_{1Ao}, p_{1Bo}, p_{1Co}, p_{1Do}, p_{1Eo}, p_{1Ro}, p_{1So}, p_{1Mo}\}$	$p_{ijo} = \{0.30, 0.25, 0.20, 0.15, 0.10, 0.0, 0.0\}$	Probability of purchase for each customer type. It is assumed these are constant across itineraries. The no-purchase probability for customer type i is equal to $1 - \sum_{j \in J} p_{ijo}$
$p_{2jo} = \{p_{2Ao}, p_{2Bo}, p_{2Co}, p_{2Do}, p_{2Eo}, p_{2Ro}, p_{2So}, p_{2Mo}\}$	$p_{2jo} = \{0, 0.05, 0.10, 0.15, 0.20, 0.25, 0.25\}$	Proportion of total demand from each customer type for each itinerary. It is assumed that these are constant across itineraries
$\phi_o = \{\phi_{1,o}, \phi_{2,o}\}$	$\phi_o = \{0.5, 0.5\} \forall o$	
$x = \{x_{AB}, x_{BC}, x_{CD}, x_{DE}\}$ A, O, J, P, R, S, M	$x = \{200, 200, 200, 200\}$ 1 leg: {70, 60, 50, 40, 30, 20, 10} 2 legs: {126, 108, 90, 72, 54, 36, 18} 3 legs: {168, 144, 120, 96, 72, 48, 24} 4 legs: {196, 168, 140, 112, 84, 56, 28}	Total capacity of leg Set of fare classes and associated prices

Table 3 Different types of outliers considered in computational study


Experiment	Outlier type	Itineraries affected	Magnitudes
1	Cluster	All	+10%, +20%, +30%, +40%, +50%, +60%, -10%, -20%, -30%, -40%, -50%, -60%
2	Itinerary	AB	+50%
3		AC	+50%
4		AD	+50%
5		AE	+50%
6		BC	+50%
7		BD	+50%
8		BE	+50%
9		CD	+50%
10		CE	+50%
11		DE	+50%
12	Station	AB	+50%
13		AC, BC	+50%
14		AD, BD, CD	+50%
15		AE, BE, CE, DE	+50%



(a) Case 1



(b) Case 2



(c) Case 3

Fig. 18 Benchmark clustering

B.3.1 Outliers considered in computational study

Table 3 shows the different experiments that were carried out as part of the computational study. We consider *cluster* outliers in which every itinerary within the cluster is equally affected; *itinerary* outliers where only a single itinerary within the cluster is affected; and *station* outliers which affect all itineraries that end at a particular station.

Appendix C Computational results

Appendix C includes the extended results from the computational study described in Sect. 3. Results from additional simulation experiments to test the proposed clustering approach are also presented here.




Fig. 19 Itinerary demand per leg

C.1 Evaluation of network clustering

For the correlation-based clustering to perform well it needs to (i) accurately estimate similarity between adjacent legs and (ii) use information about the pairwise similarity between adjacent legs to detect similarity between (potentially) more than two legs to form clusters. We use the proportion of total demand belonging to each itinerary to determine a clustering benchmark. For example in Fig. 18a, when *all* passengers travel the itinerary from A to E, the resulting bookings in each of the four legs would be identical. In this case, the correlation between legs would be 1—giving a single cluster of four legs.

To evaluate the clustering when the underlying demand is known, we define the **common traffic ratio** between two adjacent legs as the proportion of total demand that relates to itineraries over both legs. That is, for two legs ij and jk , we define the common traffic ratio, $r(ij, jk)$, to be:

$$r(ij, jk) = \frac{D_{ik}}{D_{ij} + D_{jk} + D_{ik}}, \quad (33)$$

where D_{ij} is the demand for itinerary ij , and D_{ik} is the total demand for all itineraries which include both legs ij and jk . If all passengers book itineraries that traverse both legs, then $r(ij, jk) = 1$. Conversely, if no passengers book journeys that traverse both legs, then $r(ij, jk) = 0$.

We vary the level of demand for each itinerary to generate different benchmark clusterings. The output of the correlation-based clustering is then compared with benchmark clustering using the NMI. We consider three cases: The four legs belong in a single cluster (Fig. 18a); they belong in two clusters (Fig. 18b); and they belong in four clusters (Fig. 18c).

Table 4 Normalised mutual information

	Fraction of leg demand resulting from cluster itinerary demand					
	50%	60%	70%	80%	90%	100%
Case 1	0.99	1.00	1.00	1.00	1.00	1.00
Case 2	0.98	0.99	1.00	1.00	1.00	1.00
Case 3	0.94	0.97	0.99	1.00	1.00	1.00

- **Case 1:** When itinerary AE accounts for at least 50% of the network demand, we expect legs AB, BC, CD, and DE to belong to the same cluster, as they experience mostly the same demand. The remaining demand is calibrated across itineraries such that the total demand for each leg is reasonably uniformly distributed. We compare the correlation-based clustering with the benchmark clustering of all four legs in a single cluster, when the average percentage of demand on each leg from itinerary AE is 50%, 60%, 70%, 80%, 90%, or 100%. Figure 19a shows the fraction of total demand on each leg, from each itinerary, in the case where 60% of demand is for itinerary AE.
- **Case 2:** We calibrate the majority of demand on leg AB and BC to be for itinerary AC, and the majority of demand on legs CD and DE to be demand for itinerary CE. For simplicity, the distribution of demand is symmetric across the four legs. We compare the performance when the average percentage of demand on each leg belonging to the clustering benchmark itinerary is 50%, 60%, 70%, 80%, 90%, or 100%. Figure 19b shows the case where 60% of demand on each leg is for the respective cluster itineraries (AC or CE).
- **Case 3:** We calibrate the majority of demand on leg AB for itinerary AB, the majority of demand on leg BC for itinerary BC, and so on. We compare the performance when the average percentage of demand on each leg belonging to the leg itinerary is 50%, 60%, 70%, 80%, 90%, or 100%. Figure 19c shows the case where 60% of demand on each leg is for the itinerary consisting of only that leg.

The results are shown in Table 4.

In almost all cases, the normalised mutual information between the correlation-based clustering and the benchmark equals 1, indicating congruence. We now extend the simulation study by comparing the output of the correlation-based clustering under different correlation measures. In addition to the functional dynamical correlation measure described in Sect. 2.1, we compare *Pearson correlation* (Pearson 1895) and *Kendall rank correlation* (Kendall 1938). Let $y_{n,ij}(t)$ be the observed bookings for the n^{th} departure on leg ij , and $y_{n,pq}(t)$ analogous for leg pq .

- **Pearson correlation:** calculate the Pearson correlation between corresponding booking patterns, then average across all booking patterns. That is, for the n^{th} of N booking patterns observed over T booking intervals, we calculate the Pearson correlation coefficient as:

$$\rho_n(ij, pq) = \frac{\sum_{t=1}^T (y_{n,ij}(t) - \overline{y_{n,ij}})(y_{n,pq}(t) - \overline{y_{n,pq}})}{\sqrt{\sum_{t=1}^T (y_{n,ij}(t) - \overline{y_{n,ij}})^2} \sqrt{\sum_{t=1}^T (y_{n,pq}(t) - \overline{y_{n,pq}})^2}} \quad (34)$$

where $\overline{y_{n,ij}}$ is the mean number of bookings for the n^{th} booking pattern. Then:

$$\rho(ij, pq) = \frac{1}{n} \sum_{n=1}^N \rho_n(ij, pq). \quad (35)$$

- **Kendall rank correlation:** observations $(y_{n,ij}(s), y_{n,pq}(s))$ and $(y_{n,ij}(t), y_{n,pq}(t))$ where $s < t$, are *concordant* if their ordering agrees, and *discordant* otherwise. The Kendall rank correlation is defined between the n^{th} booking patterns in legs ij and pq as:

$$\rho_n(ij, pq) = \frac{t_c - t_d}{\sqrt{(t_0 - t_1)(t_0 - t_2)}} \quad (36)$$

where t_c is the number of concordant pairs, t_d is the number of discordant pairs, and t_0 , t_1 , and t_2 are defined as follows:

$$t_0 = \frac{T(T-1)}{2}, \quad (37)$$

$$t_1 = \sum_s u_s(u_s - 1)/2, \quad (38)$$

$$t_2 = \sum_t v_t(v_t - 1)/2, \quad (39)$$

where u_s is the number of tied values in the s^{th} group of ties for in booking patterns for leg ij , and v_t is analogous for leg pq . Then:


$$\rho(ij, pq) = \frac{1}{n} \sum_{n=1}^N \rho_n(ij, pq). \quad (40)$$

We compare the cases where the correlation measure is (i) applied directly to the booking patterns and (ii) applied to the booking patterns where the within-booking pattern relationships, e.g. trends have been removed. The normalised mutual information between the clustering produced by the correlation-based clustering under each of the different correlation measures and the benchmark clustering is shown in Table 5.

For case 1, all three correlation measures seem to be performing equally well, with the normalised mutual information almost always indicating congruence. For cases 2 and 3, the Pearson and Kendall correlation results in extremely poor performance in terms of NMI, with the benchmark clustering never being

Table 5 Normalised mutual information under different correlation measures

Case	Correlation measure	Fraction of leg demand resulting from cluster itinerary demand					
		50%	60%	70%	80%	90%	100%
Case 1	Booking patterns						
	Functional dynamical correlation	0.99	1.00	1.00	1.00	1.00	1.00
	Pearson correlation	1.00	1.00	1.00	1.00	1.00	1.00
	Kendall rank correlation	1.00	1.00	1.00	1.00	1.00	1.00
	Differenced booking patterns						
	Functional dynamical correlation	0.99	1.00	1.00	1.00	1.00	1.00
	Pearson correlation	0.98	1.00	1.00	1.00	1.00	1.00
	Kendall rank correlation	1.00	1.00	1.00	1.00	1.00	1.00
Case 2	Booking patterns						
	Functional dynamical correlation	0.98	0.99	1.00	1.00	1.00	1.00
	Pearson correlation	0.00	0.00	0.00	0.00	0.00	0.00
	Kendall rank correlation	0.00	0.00	0.00	0.00	0.00	0.00
	Differenced booking patterns						
	Functional dynamical correlation	0.98	0.99	1.00	1.00	1.00	1.00
	Pearson correlation	0.00	0.00	0.00	0.00	0.00	0.00
	Kendall rank correlation	0.00	0.00	0.00	0.00	0.00	0.00
Case 3	Booking patterns						
	Functional dynamical correlation	0.94	0.97	0.99	1.00	1.00	1.00
	Pearson correlation	0.00	0.00	0.00	0.00	0.00	0.00
	Kendall rank correlation	0.00	0.00	0.00	0.00	0.00	0.00
	Differenced booking patterns						
	Functional dynamical correlation	0.93	0.96	0.99	1.00	1.00	1.00
	Pearson correlation	0.00	0.00	0.00	0.00	0.00	0.00
	Kendall rank correlation	0.00	0.00	0.00	0.00	0.00	0.00

Fig. 20 Network with two legs

achieved. Functional dynamical correlation, however, continues to perform well with an NMI close to 1.

In order to determine why the Pearson and Kendall rank correlations initially appear to perform well in the single cluster case but fail in the two cluster case, we also compare the value of the correlation coefficient with the known demand share in a simple two-leg example. Consider the simple two-leg network shown in Fig. 20.

The common traffic ratio of legs AB and BC is:

Table 6 Comparison of correlation measures

$r(AB, BC)$	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
<i>Correlation between booking patterns</i>											
Functional dynamical correlation	0.12	0.22	0.35	0.40	0.46	0.55	0.66	0.82	0.86	0.90	1.00
Pearson correlation	0.99	0.99	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Kendall rank correlation	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	1.00	1.00	1.00
<i>Correlation between differenced booking patterns</i>											
Functional dynamical correlation	0.14	0.18	0.29	0.42	0.50	0.53	0.66	0.83	0.88	0.91	1.00
Pearson correlation	0.70	0.71	0.77	0.82	0.85	0.89	0.92	0.95	0.96	0.98	1.00
Kendall rank correlation	0.88	0.90	0.91	0.91	0.92	0.94	0.94	0.95	0.96	0.97	1.00

Table 7 True-positive rate of FD+Agg in comparison to PCA+HDR benchmark under varying lengths of alert list

Length of alert list	1	5	10	50	100	250	500
FD+Agg	0.16	0.44	0.50	0.68	0.68	0.68	0.68
PCA+HDR ($v = 5$)	0.09	0.18	0.18	0.18	0.18	0.18	0.18
PCA+HDR ($v = 10$)	0.09	0.18	0.23	0.23	0.23	0.23	0.23
PCA+HDR ($v = 50$)	0.09	0.18	0.23	0.35	0.35	0.35	0.35
PCA+HDR ($v = 100$)	0.09	0.18	0.23	0.35	0.43	0.43	0.43
PCA+HDR ($v = 250$)	0.09	0.18	0.23	0.35	0.43	0.66	0.66
PCA+HDR ($v = 500$)	0.09	0.18	0.23	0.35	0.43	0.66	1

$$r(AB, BC) = \frac{D_{AC}}{D_{AB} + D_{BC} + D_{AC}}, \quad (41)$$

If $r(AB, BC) = 1$, then the number of bookings on leg AB and leg BC is identical, and the correlation between them is 1. Conversely, if $r(AB, BC) = 0$, then the bookings on leg AB and leg BC are independent with correlation 0. Table 6 shows the estimates of the correlation, compared to the true ratio, $r(AB, BC)$.

Functional dynamical correlation applied directly to the data performs best in all cases. In case 1, where the benchmark clustering is a single cluster, poor clustering performance can only result from underestimating the demand share. Both Pearson and Kendall rank correlation overestimate the correlation between booking patterns, even when the within-booking pattern effects have been removed. This explains the good performance of Pearson and Kendall rank correlation in case 1, despite extremely poor performance in cases 2 and 3.

C.2 Detecting outliers in multiple legs

C.2.1 Results on detecting outliers in multiple legs

Table 7 provides the results shown in Fig. 5 in tabular format.




Fig. 21 Outlier detection performance under different functional depth thresholds

C.2.2 Outlier detection under alternative thresholds

We recognise that the percentage of departures that analysts are able to adjust strongly depends on the ratio of analysts to departures and that this is likely to be domain-dependent. Therefore, here, we consider outlier detection performance as the functional depth threshold varies.

In terms of true-positive rates, the threshold choices of 0.01, 0.05, or 0.1 produce similar results, at least near the top of the alert list. Our method ranks the departures classified as outliers such that genuine outliers are more likely to be at the top of the ranked list, and false positives at the bottom of the list. Therefore, using a higher threshold tends to add more departures to the bottom of the list, and increase the risk of more false positives. As shown in Section 4 of the manuscript, the outliers that a threshold of 0.01 fails to detect tend to be small changes in magnitude. It is these small magnitude outliers that are added to the bottom of the list as the threshold increases. Notably, a threshold of 0.001 results in reduced performance even at the top of the list, suggesting that this would be too low a threshold. Similar results are seen in the change in precision (compared to the non-ranked method with the same threshold).

A higher threshold does result in higher overall true-positive rates as more departures are classified as outliers. However, the maximum true-positive rate for a threshold of 0.05 results in around 1 in 5 departures being classified as outliers. This is quite a high percentage for them all to be considered *outliers* (Fig. 21).




Fig. 22 Fraction of outliers detected in 1, 2, 3, or 4 legs




Fig. 23 Fraction of outliers detected in each leg

C.2.3 Distribution of outliers across multiple legs

In the scenario where all itineraries are equally affected, a high proportion of outliers should be detected in more than one leg. Figure 22a illustrates the proportion of outliers detected in 1, 2, 3, or 4 legs: More than half were detected in multiple legs. Figure 22b shows the proportion of true positives (genuine outliers which were detected), by the number of legs in which they were detected. In contrast with Fig. 22a, a much higher percentage of genuine outliers are detected in all four legs.

Given the clustering is correct, we expect an approximately equal number of single-leg outliers in each leg, as shown in Fig. 23b. If one leg, say DE, had not belonged in this cluster, we would expect a higher proportion of single-leg outliers to have been detected in leg DE. This could be utilised as a method for checking the clustering, after the outlier detection.

Fig. 24 False discovery rate for nonhomogeneous demand-volume outliers





Fig. 25 False discovery rate for homogeneous demand-volume outliers by magnitude

These results motivate aggregating threshold exceedances across legs in two ways: (i) Since less than 100% of genuine outliers were detected in all legs, if outlier detection was carried out only on the leg level, outliers could be missed on some legs. (ii) Given that a much higher proportion of outliers detected in four legs were genuine outliers, by ranking booking patterns detected in all legs as more likely to be outliers, we focus analysts' attention on those more likely to be genuine outliers.

C.2.4 False discovery rate

The *false discovery rate* (*FDR*) is defined as the proportion of booking patterns classified as outliers which were false positives:

$$FDR = \frac{FP}{TP + FP} \quad (42)$$




Fig. 26 True-positive rate for single itinerary outliers (cont.)

See Sect. 3.4 for definitions of true and false positives. Figure 24 shows the FDR for the case where outlier demand affects all itineraries, and the magnitude is randomly chosen from each of the distributions described in Sect. 3.2.

Table 8 Changes in leg demand resulting from an additional 120 passengers in itinerary demand

Event at station	Itineraries affected	Additional 120 passengers in itineraries			
		Resulting demand increase per leg			
		Leg AB	Leg BC	Leg CD	Leg DE
B	A-B	+120 (+50%)	–	–	–
C	A-C, B-C	+60 (+25%)	+120 (+50%)	–	–
D	A-D, B-D, C-D	+40 (+16.6%)	+80 (+33.3%)	+120 (+50%)	–
E	A-E, B-E, C-E, D-E	+30 (12.5%)	+60 (+25%)	+90 (+37.5%)	+120 (+50%)


**Fig. 27** Performance for demand-volume outliers in a subset of itineraries caused by an absolute increase in demand

Figure 25 shows the FDR for each of the magnitudes of outliers considered in the simulation study. Given that smaller magnitude outliers are more similar to the regular demand, these result in higher false discovery rates.

C.2.5 Outliers affecting a single itinerary


Figure 26 shows the true-positive rate for the remaining itineraries in Fig. 7 of Sect. 3.4.

C.2.6 Outliers affecting a subset of itineraries

We consider a case where demand outliers affect only a subset of itineraries. Practical examples for this phenomenon could include trade fairs or conventions as well as regional crises. In such situations, demand towards (or from) a specific destination is most affected. Here, clustering offers additional benefits in guiding analysts towards those itineraries where they should adjust the forecast or controls.

We differentiate four scenarios based on the four-leg network described in Sect. 3, where events affect demand for itineraries travelling to stations B, C, D, and

Fig. 28 True-positive rate for nonhomogeneous demand-volume outliers as minimum outlier severity varies



E, respectively. We expect analogous results when customers aim to travel home from events that happened at stations A, B, C, or D, respectively, given the symmetry of the demand parameters chosen for the computational study.

For each of the four possible events considered, we investigate the case where this generates 50% increase in average leg demand. For simplicity, we assume that these passengers are equally split between the itineraries which alight at the relevant station. Table 8 shows the resulting demand increases for each leg.

Figure 27a shows the true-positive rate for each of the cases. Although the event at E generates outliers in more legs, it is not the case that it has the highest true-positive rate. This shows that though the approach aggregates across legs, it does not ignore outliers only in a subset of those legs, provided they are sufficiently large. These effects may also be caused by interactions between the booking limits on different legs. For example in the case of an event at C, large increases in demand in legs AB and BC may cause booking limits to be reached earlier for these legs, which also limits bookings in itineraries such as AD and AE. Hence, an increase in demand for some legs may cause a decrease in bookings for different legs. By jointly considering multiple legs for outlier detection, we are able to detect the knock-on effects of outliers even when the change in demand only affects a subset of legs. The change in precision can be interpreted similarly, in Fig. 27b.

Had we considered outlier detection on a leg-by-leg basis, the outliers were more likely to be missed in some of the legs. By combining information across legs, we are better able to determine which itineraries are affecting the volume of demand.

C.2.7 Limit alert list length via outlier severity thresholds

The results in this paper focus on limiting the length of the ranked alert list simply by the number of alerts it contains as this is most relevant to analysts. However, an alternative approach limits the length of the list by the outlier severity assigned to each departure. For example classifying a train as an outlier only if its outlier severity is above 80%.

Detection results when outliers affect all itineraries

Figure 28 shows the true-positive rate as the outlier severity decreases from 100% to 0%. Results are similar to those shown in Fig. 5a. Figure 29 shows the




Fig. 29 True-positive rate for homogeneous demand-volume outliers by magnitude




Fig. 30 True-positive rate for single itinerary demand-volume outliers as minimum outlier severity varies

Fig. 31 Seasonal demand model showing distributions of total demand for 10 seasonal groups





Fig. 32 True-positive rate and false discovery rate under a seasonal demand model compared to non-seasonal demand



true-positive rate as the outlier severity decreases from 100% to 0%, for each magnitude of outlier considered. Results are similar to those shown in Fig. 8. Figure 30 indicates the true-positive rate for outliers that affect demand for a single itinerary.

Detection results when outliers affect a single itinerary

C.3 Detection results when demand is seasonal

In this section, we consider the case where demand is non-stationary across departures. In order to make the results of this simulation study more comparable with the other simulation studies included in the paper, we simulate 10 seasonal groups (somewhat analogous to months) as shown in Fig. 31. We then take 50 'days' from each of the 10 groups to produce 500 booking patterns per simulation, as in the other simulation experiments. All other parameters remain fixed in the simulation.

We then fit a functional regression model (similar to the model described in Sect. 5.2 for the empirical study) to remove the seasonality before performing the outlier detection. The quality of the outlier detection will highly depend on the

Fig. 33 True-positive rate and false discovery rate under itinerary-level outlier detection





Fig. 34 True-positive rate and false discovery rate under itinerary-level outlier detection with aggregation



model used to model the seasonal (and trend) components of the demand. A full study of the most suitable approach for modelling this type of demand is outwith the scope of this paper.

In comparison with the non-seasonal demand model, the true-positive rate is slightly lower. The false discovery rate is also slightly lower—see Fig. 32. Due to the seasonal variation the overall variance of the demand is higher which causes a lower functional depth threshold (for the same choice of percentile), meaning fewer observations are fall below the threshold.

C.4 Detection results using itinerary-level data

The methodology described in this paper focuses on the application to leg-level booking data—since in capacity-based RM only leg-level capacities are required to be stored. Often, the itinerary-level data are never recorded. However, if itinerary-level data are available for analysis, one could apply the same approach (either with or without the aggregation step). The results of doing so are presented in this section.

Figure 33 shows the true-positive and false discovery rates from the alert lists generated for each itinerary. It shows that the outlier detection performs well in the most popular itinerary (the itinerary that defines the cluster—A→E). The results are comparable to running the outlier detection procedure on the leg-level data. In the remaining nine itineraries with far fewer bookings, the outlier detection performs poorly, as there are too few bookings to detect a pattern.

Fig. 35 True-positive rate and false discovery rate under varying levels of regular demand showing limited success when demand is zero-inflated




Figure 34 shows that when the itinerary-level outlier detection results are aggregated across itineraries, the overall performance is poorer than the leg-level aggregation. These results suggest that for itinerary-level outlier detection, the better approach would be to identify important itineraries and run the outlier detection routines on those without aggregation. This, however, means that outliers that occur only in a small part of the network (which can cause knock-on effects) would be systematically overlooked. To consider the full network, the leg-based approach that accounts for itineraries through clustering performs better.

To further illustrate that the issue with itinerary-level analysis is primarily caused by an insufficient number of bookings for less-popular itineraries (making it difficult to detect patterns in the demand), we analyse the performance as the regular itinerary-level demand varies. Figure 35 shows that performance (especially false discovery rate) is poorer when demand is especially low.

C.5 Revenue benefits from forecast adjustments

Figure 36 shows the true-positive rate for the remaining itineraries in Fig. 10 of Sect. 4.3.

The analysis in Sect. 4.2 constitutes a best-case scenario in which we assume that, if outlier demand affects a particular leg, the outlier is detected in that leg. However, as we show in Sect. 3.4, even when demand outliers affect multiple legs, the outlier is not always detected in every leg due to noise. Therefore, we additionally compare different adjustments based on the output of the outlier detection, for an outlier in itinerary AE.

- **Adjustment A:** Adjust only the forecasts of the affected single-leg itineraries for those legs in which the outlier is detected.
- **Adjustment B:** Adjust the forecasts of the affected single-leg itineraries for those legs in which the outlier is detected, **and** the cluster spanning itinerary (AE).




Fig. 36 Revenue generated under different itinerary-level forecast adjustments (cont.)

We compare these both to making no adjustment and to the oracle adjustment. This is still a best-case scenario to some extent, given that we assume the correct magnitude of adjustment is made.

Figure 37 shows the revenue under adjustments A and B (as described in Sect. 4.2) depending on the output of the outlier detection procedure. Combining adjustments on




Fig. 37 Revenue generated under different forecast adjustments resulting from the outlier detection for outlier demand in itinerary AE

Table 9 Model comparison for functional regression

Model	Intercept	Day	Month	Short horizon (I)	Short horizon (C)	CV-SSE			
						Leg AB	Leg BC	Leg CD	Leg DE
Model 1	✓					79974160	75034839	79529280	73824611
Model 2	✓			✓		58617546	52622148	52424683	50009080
Model 3	✓				✓	58620898	52863263	52506946	50014984
Model 4	✓	✓				27227350	35376732	32789181	30037659
Model 5	✓	✓		✓		26551341	33724380	32282900	29989390
Model 6	✓	✓			✓	26704943	34154782	32439972	30019196
Model 7	✓		✓			58620649	57895619	52638923	50015645
Model 8	✓		✓	✓		58608640	57865403	52615801	49996331
Model 9	✓		✓		✓	58878374	57885484	52654330	50033157
Model 10	✓	✓	✓			24574978	25700166	21691111	21880038
Model 11	✓	✓	✓	✓		24519539	25691637	21689686	21878259
Model 12	✓	✓	✓		✓	24546715	25697938	21724073	21896889

the leg level with those on the cluster level provides superior results in contrast to leg-level adjustments alone. Though making adjustments to only the single-leg itineraries may be risk-averse in the rare cases where an outlier affects only a small subset of the legs within a cluster, it may be detrimental to revenue when outliers affect multiple legs.

Appendix D Deutsche Bahn booking data

Appendix D contains an additional analysis of the empirical booking data from Deutsche Bahn.

D.1 Model selection for functional regression

Due to the functional nature of the data, in order to determine which of the factors result in a better fitting model, we use the **Cross-Validated Sum Of Integrated Squared Errors** (CV-SSE).

$$CV-SSE = \sum_{n=1}^N \int (y_{nl}(t) - \hat{y}_{nl}(t))dt, \quad (43)$$

where $\hat{y}_{nl}(t)$ is the prediction for the n^{th} booking pattern on the leg l , under the model fitted to all but the n^{th} booking pattern. The model which produces the lowest CV-SSE is chosen as the best fitting. Note that unlike other model selection criteria (e.g. AIC), CV-SSE does not take into account the number of parameters. Given that we are not interested in out-of-sample prediction, only in obtaining the best fitting model for our data, over-fitting is not of great concern. The values of the CV-SSE for each of the 12 models considered are shown in Table 9.

Across all legs, we find that day, month, and shortened booking horizons are all factors that must be taken into account. The inclusion of the days of the week as factors significantly reduces the CV-SSE. In comparison, the inclusion of the booking horizon variable has a smaller, though still positive, effect. We compare two different approaches to accounting for the shortened booking horizon: (i) an indicator function (I) equal to 1 if the booking horizon is shorter, and (ii) a continuous variable (C) between 0 and 1 which gives the length of the shortened horizon as a proportion of the regular length horizon. Based on the CV-SSE scores, shortened booking horizons are best represented by the indicator function, i.e. it is important to know that it is shorter but not by how much. The smaller effect of the horizon length variable may be related to the inclusion of the month variable, which is unsurprising given the overlap in the definition of these variables. The values of the CV-SSE are similar for models 2 and 7, where we only consider one month or horizon length as a factor.

D.2 Residual booking patterns

Figure 38 shows the residual booking patterns resulting from the functional regression applied in equation (12) of Sect. 5.2. Compare with Fig. 15 of Sect. 5.2—the obvious outliers are preserved.




Fig. 38 Residual booking patterns

D.3 Functional linear regression coefficients

Table 10 shows an example of the coefficients for each explanatory variable over the 19 booking intervals, along with their standard errors. Not all days of the week, or months of the year, are significant throughout the entire booking horizon. Similarly, some days of the week have much stronger significance than others, e.g. Saturdays have a much larger impact than Tuesdays.

D.4 Functional depths

Figure 39 shows the functional depths for the empirical residual booking patterns, before the functional depths are transformed into the z_{nl} , as shown in Fig. 16 of Sect. 5.2.

Table 10 Functional linear regression coefficients (and standard errors)

Coefficient	t = 1	2	3	4	5	6	7	8	9
β_{0l}	1.4 (1.0)	1.8 (1.1)	5.4 (1.2)	11.2 (1.4)	13.8 (1.5)	14.8 (1.5)	15.2 (1.6)	16.4 (1.6)	17.7 (1.6)
β_{1l} [Mon]	-0.4 (0.8)	-1.0 (0.9)	-1.5 (1.0)	-2.1 (1.2)	-1.6 (1.2)	-0.7 (1.3)	-0.1 (1.3)	0.5 (1.4)	0.8 (1.4)
β_{2l} [Tue]	-1.2 (0.8)	-2.1 (0.9)	-3.7 (1.0)	-5.1 (1.2)	-5.8 (1.2)	-6.3 (1.3)	-6.8 (1.2)	-6.4 (1.3)	-6.4 (1.4)
β_{3l} [Wed]	-1.2 (0.8)	-1.9 (0.9)	-3.9 (1.0)	-5.2 (1.2)	-6.5 (1.2)	-7.1 (1.3)	-7.3 (1.3)	-7.4 (1.3)	-8.9 (1.4)
β_{4l} [Thu]	-1.0 (0.8)	-1.3 (0.9)	-3.2 (1.0)	-4.3 (1.2)	-5.6 (1.2)	-6.7 (1.3)	-6.7 (1.3)	-7.1 (1.3)	-7.8 (1.4)
β_{5l} [Fri]	-0.5 (0.8)	-0.4 (0.9)	-0.7 (1.0)	-0.3 (1.2)	0.8 (1.2)	1.5 (1.3)	2.6 (1.3)	2.7 (1.4)	2.9 (1.4)
β_{6l} [Sat]	5.5 (0.8)	10.9 (0.9)	17.1 (1.1)	25.5 (1.2)	30.4 (1.3)	33.5 (1.3)	36.9 (1.3)	38.6 (1.4)	41.2 (1.4)
β_{7l} [Jan]	-0.7 (0.9)	-0.3 (1.0)	-1.3 (1.2)	-4.3 (1.3)	-5.0 (1.4)	-4.8 (1.5)	-4.5 (1.5)	-5.0 (1.5)	-4.4 (1.6)
β_{8l} [Feb]	0.3 (1.0)	1.8 (1.1)	0.0 (1.3)	-4.0 (1.4)	-3.5 (1.5)	-2.3 (1.6)	-0.3 (1.6)	0.3 (1.7)	1.6 (1.7)
β_{9l} [Mar]	0.2 (1.0)	0.9 (1.1)	0.4 (1.2)	-3.2 (1.4)	-4.4 (1.5)	-4.0 (1.6)	-3.2 (1.6)	-2.9 (1.7)	-2.5 (1.7)
β_{10l} [Apr]	0.3 (1.0)	2.1 (1.1)	1.2 (1.3)	0.9 (1.4)	3.1 (1.5)	4.8 (1.6)	5.6 (1.6)	5.6 (1.7)	5.9 (1.7)
β_{11l} [May]	3.5 (1.0)	4.4 (1.1)	4.1 (1.2)	0.7 (1.4)	-0.1 (1.5)	0.0 (1.6)	0.6 (1.6)	0.5 (1.7)	0.9 (1.7)
β_{12l} [Jun]	0.2 (1.0)	1.9 (1.1)	0.7 (1.3)	-1.9 (1.4)	-2.6 (1.5)	-2.2 (1.6)	-1.9 (1.6)	-2.0 (1.7)	-1.5 (1.7)
β_{13l} [Jul]	0.4 (0.8)	1.6 (1.1)	1.4 (1.2)	-1.5 (1.4)	-1.8 (1.5)	-0.7 (1.6)	-0.3 (1.6)	0.1 (1.7)	0.9 (1.7)
β_{14l} [Aug]	1.5 (1.0)	3.3 (1.1)	3.3 (1.3)	3.7 (1.5)	5.0 (1.6)	6.4 (1.6)	7.7 (1.7)	8.3 (1.7)	9.8 (1.7)
β_{15l} [Sep]	1.9 (1.1)	2.5 (1.2)	2.1 (1.3)	-1.2 (1.5)	-0.7 (1.6)	0.6 (1.6)	1.5 (1.7)	2.3 (1.7)	3.5 (1.8)
β_{16l} [Oct]	1.6 (1.0)	3.5 (1.1)	3.8 (1.2)	4.7 (1.4)	7.0 (1.5)	8.8 (1.6)	11.2 (1.6)	12.6 (1.7)	15.5 (1.7)
β_{17l} [Nov]	1.0 (0.9)	2.8 (1.1)	4.2 (1.3)	4.5 (1.4)	6.7 (1.5)	9.3 (1.6)	11.0 (1.6)	11.5 (1.7)	12.5 (1.7)
β_{18l} [Horizon]	-1.2 (0.9)	-0.2 (1.0)	0.7 (1.2)	-0.8 (1.3)	-1.1 (1.4)	-0.6 (1.5)	-0.5 (1.5)	-0.2 (1.5)	-0.4 (1.6)
t = 10	11	12	13	14	15	16	17	18	19
18.5 (1.7)	18.7 (1.7)	19.5 (1.7)	20.7 (1.8)	22.2 (1.8)	24.1 (1.8)	26.0 (1.8)	27.9 (1.9)	29.0 (1.9)	30.4 (2.0)
1.1 (1.4)	1.1 (1.4)	1.1 (1.5)	0.4 (1.5)	-0.6 (1.5)	-1.1 (1.5)	-2.5 (1.6)	-3.7 (1.6)	-3.5 (1.6)	-2.1 (1.7)
-6.8 (1.4)	-5.8 (1.4)	-5.7 (1.5)	-6.4 (1.5)	-7.4 (1.5)	-8.0 (1.5)	-9.3 (1.6)	-10.2 (1.6)	-10.0 (1.6)	-9.1 (1.6)
-9.2 (1.4)	-8.9 (1.4)	-9.1 (1.5)	-10.0 (1.5)	-11.2 (1.5)	-12.7 (1.5)	-14.5 (1.6)	-14.7 (1.6)	-15.1 (1.6)	-15.2 (1.7)
-7.8 (1.4)	-7.6 (1.4)	-7.7 (1.5)	-8.9 (1.5)	-10.7 (1.5)	-12.1 (1.5)	-12.9 (1.6)	-13.6 (1.6)	-13.7 (1.6)	-13.7 (1.7)

Table 10 (continued)

t = 10	11	12	13	14	15	16	17	18	19
4.2 (1.4)	5.0 (1.4)	5.4 (1.5)	4.3 (1.5)	4.0 (1.5)	4.6 (1.5)	5.0 (1.6)	5.8 (1.6)	6.4 (1.6)	7.1 (1.7)
45.0 (1.4)	46.1 (1.5)	48.1 (1.5)	49.4 (1.5)	50.7 (1.5)	53.1 (1.6)	55.1 (1.6)	57.1 (1.6)	58.5 (1.7)	59.4 (1.7)
-3.7 (1.6)	-3.0 (1.7)	-3.8 (1.7)	-3.6 (1.7)	-3.1 (1.7)	-3.5 (1.8)	-3.3 (1.8)	-3.8 (1.8)	-3.1 (1.9)	-3.8 (1.9)
3.3 (1.8)	3.7 (1.8)	3.9 (1.8)	4.9 (1.8)	5.9 (1.9)	6.8 (1.9)	7.8 (1.9)	8.2 (2.0)	9.5 (2.0)	9.2 (2.1)
-1.6 (1.7)	-0.8 (1.8)	-1.1 (1.8)	-0.6 (1.8)	-0.2 (1.9)	-0.6 (1.9)	-0.4 (1.9)	-0.8 (2.0)	-0.3 (2.0)	-0.7 (2.0)
6.1 (1.8)	6.3 (1.8)	6.3 (1.8)	6.8 (1.8)	7.2 (1.9)	7.3 (1.9)	7.2 (1.9)	7.1 (2.0)	8.3 (2.0)	8.8 (2.0)
1.1 (1.7)	1.4 (1.8)	1.3 (1.8)	1.8 (1.8)	2.5 (1.9)	2.3 (1.9)	2.4 (1.9)	2.1 (2.0)	3.1 (2.0)	3.5 (2.0)
-1.0 (1.8)	-0.4 (1.8)	-0.2 (1.8)	0.9 (1.8)	1.5 (1.9)	1.1 (1.9)	1.1 (1.9)	0.7 (2.0)	2.0 (2.0)	2.9 (2.0)
1.8 (1.7)	2.5 (1.8)	2.6 (1.8)	3.1 (1.8)	4.0 (1.9)	4.0 (1.9)	4.4 (1.9)	4.6 (2.0)	6.3 (2.0)	6.9 (2.0)
11.6 (1.8)	12.8 (1.8)	13.2 (1.8)	14.0 (1.9)	15.4 (1.9)	16.4 (1.9)	17.1 (2)	17.4 (2.0)	19.2 (2.1)	21.7 (2.1)
5.1 (1.8)	6.8 (1.8)	7.3 (1.9)	8.7 (1.9)	9.9 (1.9)	10.2 (2.0)	11.6 (2.0)	11.6 (2.0)	12.4 (2.1)	13.3 (2.1)
17.0 (1.7)	18.4 (1.8)	19.2 (1.8)	20.7 (1.8)	22.0 (1.9)	22.9 (1.9)	24.6 (1.9)	25.3 (2.0)	27.6 (2.0)	28.0 (2.0)
13.4 (1.8)	14.0 (1.8)	14.0 (1.8)	15.1 (1.8)	15.8 (1.9)	15.9 (1.9)	16.2 (1.9)	16.1 (2.0)	16.6 (2.0)	16.2 (2.0)
-0.2 (1.6)	0.2 (1.7)	1.2 (1.7)	1.9 (1.7)	2.6 (1.7)	2.9 (1.8)	3.2 (1.8)	4.5 (1.8)	5.4 (1.9)	6.0 (1.9)




Fig. 39 Functional depths







Fig. 40 P-P plots

Fig. 41 Fraction of all outliers detected in 1, 2, 3, or 4 legs



(a) Fraction of outliers detected in each leg, for outliers detected in at least one leg



(b) Fraction of outliers detected in each leg, for outliers detected in only one leg

Fig. 42 Fraction of outliers detected in each leg

D.5 Probability plots for GPD and exponential distributions

Given that, if both $\mu = 0$ and $\xi = 0$, the GPD reduces to an exponential distribution, it is appropriate to compare the fit of the GPD with an exponential distribution to check if the inclusion of additional parameters is beneficial. Figure 40 shows the P-P plots, i.e. the fitted theoretical CDF against the empirical CDF for the GPD (Fig. 40a) and the exponential distribution (Fig. 40b). The GPD provides a closer fit to the empirical data and the additional parameters better account for the shape of the distribution.

The GPD does not provide a perfect fit, with the probabilities in the bottom left of Fig. 40a on consistently being underestimated. However, given that we assume points with very low probability are more likely to be false positives, under-estimating may actually be beneficial. Further, only the highly-ranked outliers, i.e. those with high probability, are likely to be considered by an analyst due to time constraints. The GPD provides a very good fit for those data points. If there are a sufficiently large number of threshold exceedances, an empirical distribution could alternatively be used to compute the probabilities.




Fig. 43 Comparison of standard deviation divided by mean of booking patterns

Table 11 Functional dynamical correlation of empirical booking patterns

	Leg AB	Leg BC	Leg CD	Leg DE
<i>(a) Booking patterns</i>				
Leg AB	—	0.95	0.83	0.70
Leg BC	—	—	0.83	0.66
Leg CD	—	—	—	0.78
Leg DE	—	—	—	—
<i>(b) Residual booking patterns</i>				
Leg AB	—	0.92	0.75	0.58
Leg BC	—	—	0.88	0.74
Leg CD	—	—	—	0.84
Leg DE	—	—	—	—

Table 12 Functional dynamical correlation of simulated booking patterns

	Leg AB	Leg BC	Leg CD	Leg DE
Leg AB	—	0.81	0.72	0.60
Leg BC	—	—	0.86	0.68
Leg CD	—	—	—	0.78
Leg DE	—	—	—	—

D.6 Distribution of outliers across multiple legs

The proportion of outliers found in each number of legs is shown in Fig. 41, with over half of the outliers detected in multiple legs. Compared with Fig. 22, this shows a similar proportion of outliers as found in the simulation study.

Figure 42a shows the proportion of total outlying booking patterns in terms of which legs they were detected as outliers in. Figure 42b shows the proportion in each leg of

outlying booking patterns detected in one leg only. The proportions are fairly evenly split between the different legs. This reassures us that the correct clustering was chosen—if leg DE did in fact belong to a separate second cluster, we would expect a higher proportion of single-leg outliers to have been found in leg DE—compare with Fig. 23.

D.7 Simulation verification

In order to validate the parameter choices used to simulate booking patterns, we compare the resulting simulated booking patterns with the empirical booking patterns. We consider the standard deviation and mean of the bookings across the booking horizon of each in Fig. 43. Both the empirical and simulated booking patterns show a similar shape and magnitude of the relationship between the mean and standard deviation across the booking horizon.

We also compare the correlations between the different legs for both the empirical and simulated data. Table 11 shows the functional dynamical correlation between the empirical booking patterns and empirical residual booking patterns, for each leg. Table 12 shows the corresponding correlations between the simulated booking patterns. The values are similar and the rate of decay between legs as they get further apart follows a similar pattern.

Acknowledgements We gratefully acknowledge the support of the EP/L015692/1 STOR-i Centre for Doctoral Training funded by the Engineering and Physical Sciences Research Council. The authors thank Deutsche Bahn for the provision of data and are particularly grateful to Philipp Bartke and Valentin Wagner for helpful discussions and suggestions.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data availability The simulated data and source code that support the findings of this study are available from the corresponding author upon request. R functions for outlier detection are available at: <https://github.com/nrennie/outlier-detection-in-network-revenue-management>. The code also includes some examples of the simulated data. Functions for simulating demand can be found at <https://github.com/nrennie/simnetdemand>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Amelio A, Pizzuti C (2015) Is normalized mutual information a fair measure for comparing community detection methods? Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2015 pp 1584–1585
- Azadeh SS, Labib R, Savard G (2013) Railway demand forecasting in revenue management using neural networks. *Int J Revenue Manag* 7(1):18–36
- Banerjee N, Morton A, Akartunali K (2020) Passenger demand forecasting in scheduled transportation. *Eur J Op Res* 286(3):797–810
- Barrow D, Kourentzes N (2018) The impact of special days in call arrivals forecasting: a neural network approach to modelling special days. *Eur J Op Res* 264(3):967–977
- Belobaba P (1987) Air travel demand and airline seat inventory management. PhD thesis, Massachusetts Institute of Technology
- De Baets S, Harvey N (2020) Using judgment to select and adjust forecasts from statistical models. *Eur J Op Res* 284(3):882–895
- Dubin JA, Müller HG (2005) Dynamical correlation for multivariate longitudinal data. *J Am Stat Assoc* 100:872–881
- Fawcett L, Walshaw D (2007) Improved estimation for temporally clustered extremes. *Environmetrics* 18(2):173–188
- Fawzy A, Mokhtar HM, Hegazy O (2013) Outliers detection and classification in wireless sensor networks. *Egyptian Inf J* 14(2):157–164
- Febrero M, Galeano P, González-Manteiga W (2008) Outlier detection in functional data by depth measures, with application to identify abnormal NOx levels. *Environmetrics* 19(4):331–345
- Fiig T, Weatherford LR, Wittman MD (2019) Can demand forecast accuracy be linked to airline revenue? *J Revenue Pricing Manag* 18(4):291–305
- Grimshaw SD (1993) Computing maximum likelihood estimates for the generalized pareto distribution. *Technometrics* 35(2):185–191
- He G, Müller HG, Wang JL (2003) Functional canonical analysis for square integrable stochastic processes. *J Multivar Anal* 85(1):54–77
- Hubert M, Claeskens G, De Ketelaere B, et al (2012) A new depth-based approach for detecting outlying curves. In: Colubi A, Fokianos K, Gonzalez-Rodriguez G, et al (eds) Proceedings of COMPSTAT 2012, pp 329–340
- Hyndman RJ (1996) Computing and graphing highest density regions. *Am Stat* 50(2):120–126
- Hyndman RJ, Wang E, Laptev N (2016) Large-Scale Unusual Time Series Detection. Proceedings - 15th IEEE International Conference on Data Mining Workshop, ICDMW 2015 pp 1616–1619. <https://doi.org/10.1109/ICDMW.2015.104>
- Kendall MG (1938) A new measure of rank correlation. *Biometrika* 30:81–93
- Kimms A, Müller-Bungart M (2007) Simulation of stochastic demand data streams for network revenue management problems. *Or Spectrum* 29(1):5–20
- Klein R, Koch S, Steinhardt C et al (2020) A review of revenue management: Recent generalizations and advances in industry applications. *Eur J Op Res* 284(2):397–412
- Kumar P, Khani A (2020) Evaluating special event transit demand: a robust principal component analysis approach. *IEEE Trans Intell Transp Syst* 22(12):7370–7382
- Lawrence M, Goodwin P, O'Connor M et al (2006) Judgmental forecasting: a review of progress over the last 25 years. *Int J Forecast* 22(3):493–518
- Leadbetter M (1991) On a basis for 'Peaks over Threshold' modeling. *Stat Probabil Lett* 12(4):357–362
- Liang T, Liu H, Tan Y (2020) Research on the gravity planning model of prefecture city rail transit network. International Academic Exchange Conference on Science and Technology Innovation 2019 145(02005):493–518
- Pearson K (1895) VII Note on regression and inheritance in the case of two parents. *Proc R Soc Lond* 58:87
- Perera HN, Hurley J, Fahimnia B et al (2019) The human factor in supply chain forecasting: A systematic review. *Eur J Op Res* 274(2):574–600
- Peter SJ, Victor SP (2010) A novel algorithm for informative meta similarity clusters using minimum spanning tree. *Int J Computer Sci Inf Secur* 8:1
- Pickands J (1975) Statistical inference using extreme order statistics. *Annals Stat* 3(1):119–131
- Prim R (1957) Shortest connection networks and some generalizations. *Bell Syst Technol J* 36:1389–1401

- Qu H, Zhou H, Wu Y (2007) Controllable and progressive edge clustering for large networks. *Graph Draw Lect Notes Computer Sci* 45:4372
- Quante R, Meyr H, Fleischmann M (2009) Revenue management and demand fulfillment: matching applications, models, and software. *OR Spectrum* 31(1)
- Ramsay JO, Silverman BW (1997) *Functional data analysis*. Springer, New York
- Ranshous S, Shen S, Koutra D et al (2015) Anomaly detection in dynamic networks: A survey. *WIREs: Comput Stat* 7(3):223–247
- Rennie N, Cleophas C, Sykulski AM et al (2021) Identifying and responding to outlier demand in revenue management. *Eur J Op Res* 293:1015–1030. <https://doi.org/10.1016/j.ejor.2021.01.002>
- Rennie N, Cleophas C, Sykulski AM, et al (2022) Analysing and visualising bike sharing demand with outliers. <https://doi.org/10.48550/ARXIV.2204.06112>, <http://arxiv.org/abs/2204.06112>
- Ribatet M, Dutang C (2019) POT: Generalized Pareto Distribution and Peaks Over Threshold. <https://CRAN.R-project.org/package=POT>, r package version 1.1-7
- Schaeffer SE (2007) Graph clustering. *Computer Sci Rev* 1(1):27–64
- Schütze C, Cleophas C, Tarafdar M (2020) Revenue management systems as symbiotic analytics systems: insights from a field study. *Business Res* 13(3):1007–1031
- Smith RL (1985) Maximum likelihood estimation in a class of nonregular cases. *Biometrika* 72(1):67–90
- Strauss AK, Klein R, Steinhardt C (2018) A review of choice-based revenue management: theory and methods. *Eur J Op Res* 271(2):375–387
- Talagala PD, Hyndman RJ, Smith-Miles K et al (2019) Anomaly detection in streaming nonstationary temporal data. *J Comput Gr Stat* 29:13
- Talluri KT, Van Ryzin GJ (2004) *The theory and practice of revenue management*. Kluwer Academic Publishers, Amsterdam
- Weatherford LR (2016) The history of forecasting models in revenue management. *J Revenue Pricing Manag* 15(3):212–221
- Weatherford LR, Belobaba PP (2002) Revenue impacts of fare input and demand forecast accuracy in airline yield management. *J Op Res Soc* 53(8):811–821
- Weatherford LR, Bodily SE (1992) A taxonomy and research overview of perishable-asset revenue management: yield management, overbooking, and pricing. *Op Res* 40:831–844
- Weatherford LR, Kimes SE (2003) A comparison of forecasting methods for hotel revenue management. *Int J Forecast* 19(3):401–415
- Yang K, Shahabi C (2004) A PCA-based similarity measure for multivariate time series. *MMDB 2004: Proceedings of the Second ACM International Workshop on Multimedia Databases* pp 65–74. <https://doi.org/10.1145/1032604.1032616>
- Zahn CT (1971) Graph-theoretical methods for detecting and describing gestalt clusters. *IEEE Trans Computers C*–20(1):68–86

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Nicola Rennie¹ · Catherine Cleophas²  · Adam M. Sykulski³ · Florian Dost⁴

Nicola Rennie
n.rennie@lancaster.ac.uk

Adam M. Sykulski
adam.sykulski@imperial.ac.uk

Florian Dost
florian.dost@b-tu.de

¹ STOR-i Centre for Doctoral Training, Lancaster University, Lancaster LA1 4YW, UK

² Institute of Business, Kiel University, 24098 Kiel, Germany

³ Department of Mathematics, Imperial College London, London SW7 2AZ, UK

⁴ Institute of Business and Economics, Brandenburg University of Technology, 03046 Cottbus, Germany