MATCHMAKING AND TESTING FOR EXPONENTIALITY IN THE M/G/∞ QUEUE

RUDOLF GRÜBEL AND HENDRIK WEGENER

Abstract. Customers arrive sequentially at times $x_1 < x_2 < \cdots < x_n$ and stay for independent random times $Z_1, \ldots, Z_n > 0$. The Z-variables all have the same distribution Q. We are interested in situations where the data are incomplete in the sense that only the order statistics associated with the departure times $x_i + Z_i$ are known, or that the only available information is the order in which the customers arrive and depart. In the former case we explore possibilities for the reconstruction of the correct matching of arrival and departure times. In the latter case we propose a test for exponentiality.

1. Introduction

In order to explain the problems discussed in the present paper we consider the years of birth and death of the composers J.S. Bach (1685–1750), Haydn (1732–1809), Mozart (1756–1791), Beethoven (1770–1827), Schubert (1797–1828), Brahms (1833–1897) and Tchaikovsky (1840–1893). The upper part of Figure 1 on the next page shows these on two time axes, for birth and death respectively; the connecting lines indicate the correct matching, i.e. the true year of death associated with any particular year of birth. Suppose now that the correct matching is not known and that we only have the order statistics associated with the birth times and times of death, and that the lifetime distribution Q is known. Can we reconstruct the matching? Or, suppose that only the order of arrivals and departures is known, and that Q is not known. In our running example this would lead to $(1, 2, 1, 3, 4, 3, 5, 2, 4, 5, 6, 7, 7, 6)$, where the first occurrence of a number signifies birth, the second death. Is there still some information about Q in this severely reduced data set?

The first type of data loss typically appears whenever the number of 'customers' in a 'system' is recorded, where the meaning of these words depends on the respective application. For our running example this is illustrated by the lower part of Figure 1, where the number of composers alive at time t is given as a function of t: Knowing the function is obviously equivalent to knowing the ordered arrival and departure data. If arrivals occur at the time instances of a Poisson process with constant rate then we obtain an initial path segment of a queue of the type M/G/∞ in D. G. Kendall's notation.

Formally, we have arrival times $x_1 < x_2 < \cdots < x_n$, which we first regard as fixed. At each of these instances a delay period (or lifetime or service time) begins;

2000 *Mathematics Subject Classification.* Primary 60K25, secondary 62M07, 62M20.

Key words and phrases. Asymptotics, Kendall's τ, log-concave densities, log-convex densities, prediction, permutation, queues.

This work has been supported by grants from the Deutsche Forschungsgemeinschaft and the State of Lower Saxony.

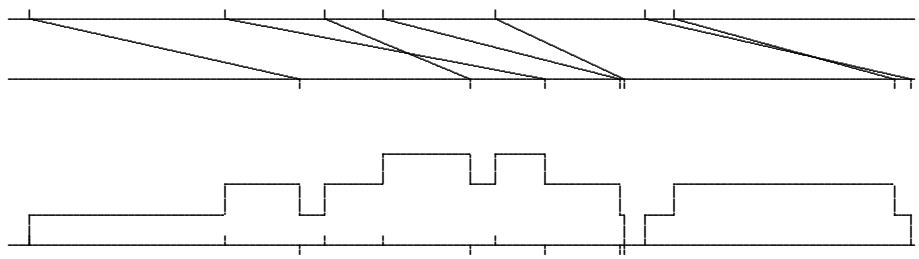

FIGURE 1. Composer data

we assume that these are non-negative, independent, and identically distributed random variables $Z_1, \ldots, Z_n$. We call the distribution Q of the Z-variables the *delay distribution* and generally assume that it has a density f with respect to Lebesgue measure. Let $\tilde{Y}_i = x_i + Z_i$, $1 \le i \le n$, be the departure times (or times of death etc.) and let $Y_1, \ldots, Y_n$ be the increasing order statistics associated with $\tilde{Y}_1, \ldots, \tilde{Y}_n$. With probability 1 there exists a unique random permutation Π in the set $\mathbb{S}_n$ of all permutations of $\{1, \ldots, n\}$ such that $\tilde{Y}_i = Y_{\Pi_i}$, $1 \le i \le n$. Further, $V_i := \#\{1 \le j < i : \tilde{Y}_j > x_i\}$ denotes the number of customers that are in the system at the time of arrival of the ith customer, $1 \le i \le n$.

In Section 2 we consider the problem of predicting or approximating the value of Π on the basis of $x = (x_1, \ldots, x_n)$ and $Y = (Y_1, \ldots, Y_n)$, where we regard Q as known. We deal with maximum a posteriori and minimum distance reconstruction. In Section 3 we introduce and discuss a test for the hypothesis that Q is an exponential distribution, now on the basis of the data (π, v), where $\pi = (\pi_1, \ldots, \pi_n)$ is the value of the matching Π, and $v = (v_1, \ldots, v_n)$ is the value of the random vector $V = (V_1, \ldots, V_n)$. In Section 4 we examine the asymptotic behaviour of our procedures in the case of Poisson process arrivals.

The situation considered in the present paper is similar to the *broken sample problem* introduced in [10], where a sample $(X_1, Y_1), \ldots, (X_n, Y_n)$ from a distribution on $\mathbb{R}^k \times \mathbb{R}^l$ is reduced to the respective multisets of the X- and Y-values; see also [11] and [3]. In [10] it is assumed that the distribution has a density f that can be written as $f(x, y) = \alpha(x)\beta(y)\exp(\gamma(x)\delta(y))$ with suitable functions $\alpha, \beta, \gamma, \delta$; in view of our general assumption $x < y$ this is not possible for delay data. In [11] and [3] related statistical questions are investigated, such as the possibility of consistent estimation of the correlation coefficient for bivariate normal distributions. In the delay situation, estimation of Q on the basis of the ordered arrival and departure data has been considered in the M/G/∞ context in [7] and [13], see also the references given there. The problem that originally motivated the present investigation comes from insurance mathematics: In the context of multivariate risk theory the class of models introduced in [5] incorporates dependence between the components by assuming that original events may trigger later claims in different business lines. The associated ruin functions then depend on the delay times. Finally, from a very general point of view the framework considered here is somewhat similar to the situation with 'decompounding', see [8], in the sense that an inverse problem arises that may be of interest in its own right.

In addition to the model variables defined above the following notation will be used. We write $\mathcal{L}(X)$ for the distribution of a random variable X, occasionally abbreviating $\mathcal{L}(X) = Q$ to $X \sim Q$. We assume throughout the paper that the delay distribution Q has a density f that is strictly positive on the whole of $\mathbb{R}_+$. In our results we further assume that f is log-convex or log-concave. In order to have a compact notation we write $\mathcal{P}_{\rm lcv}$ and $\mathcal{P}_{\rm lcc}$ for the class of distributions with log-convex and log-concave densities respectively, and $\mathcal{P}^\circ_{\rm lcv}$ and $\mathcal{P}^\circ_{\rm lcc}$ if these are strictly log-convex or strictly log-concave. These families are often used in the economics literature ([1], [2]) but also appear in reliability theory ([4, p. 75ff]). Two prominent examples, especially in connection with lifetimes and service times, are the Weibull and the gamma distributions, with densities

$$f_W(z) \;=\; f_W(z|\beta,\eta) \;=\; \beta\eta^\beta z^{\beta-1}\,\exp\bigl(-(\eta z)^\beta\bigr),\;\; z>0, \tag{1}$$

and

$$f_\Gamma(z) \;=\; f_\Gamma(z|\beta,\eta) \;=\; \frac{1}{\Gamma(\beta)}\,\eta^\beta z^{\beta-1}\exp(-\eta z),\;\; z>0, \tag{2}$$

respectively. It is known (and easy to check) that these are strictly log-concave if $\beta > 1$ and strictly log-convex if $\beta < 1$. With $\beta = 1$ we obtain the exponential distribution $\mathrm{Exp}(\eta)$ in the Weibull as well as in the gamma case; indeed, in the above notation, $\mathcal{P}_{\rm lcv} \cap \mathcal{P}_{\rm lcc} = \{\mathrm{Exp}(\eta) : \eta > 0\}$. For further examples of log-convex and log-concave distributions, and of distributions that are neither, we refer the reader to [2].

2. Matchmaking

In this section we regard x as fixed (and known), our data consist of a realization y of Y, and we are interested in the (unknown) value π of the matching Π. As the Z-variables are assumed to be strictly positive π must be an element of the set

$$\mathbb{S}_n^\circ \;=\; \mathbb{S}_n^\circ(x,y) \;:=\; \bigl\{\pi \in \mathbb{S}_n : y_{\pi_i} > x_i \text{ for } i = 1,\dots,n\bigr\}$$

of *permissible permutations*. We use $\pi = (\pi_1,\dots,\pi_n)$ as a notation for the permutation $i \mapsto \pi_i$ and we recall that the elements of $\mathbb{S}_n$ can be coded by the relative ranks of these values,

$$\Phi : \mathbb{S}_n \to \mathop{\times}_{i=1}^{n} \{1,\dots,i\}, \quad \Phi_i(\pi) = \#\{1 \le j \le i : \pi_j \le \pi_i\}. \tag{3}$$

In view of the support assumption on Q all elements of the set $\mathbb{S}_n^\circ$ appear with positive probability. It is easy to see that

$$V_i \;=\; i - 1 \;-\; \#\{1 \le j \le n : Y_j < x_i\}, \tag{4}$$

which means that V is a function of x and Y; in particular, we know its value $v = (v_1,\dots,v_n)$. At the time of the ith arrival $i-1-v_i$ departures have already taken place, so the range of values for the relative rank of the ith departure time within the departures associated with the earlier arrivals is the set $\{i - v_i,\dots,i\}$, and we can define a bijective mapping

$$\Psi \;:\; W_n := \mathop{\times}_{i=1}^{n} \{i - v_i,\dots,i\} \;\to\; \mathbb{S}_n^\circ \tag{5}$$

by the requirement that $\Psi^{-1} = \Phi$ on $\mathbb{S}_n^\circ$ with Φ as in (3). Obviously,

$$\Psi(1,2,\dots,n) = \mathrm{id} := (1,2,\dots,n) \in \mathbb{S}_n^\circ,$$

which is the permissible permutation where customers depart in the order of their arrival. Also, as a consequence of the existence of a bijective mapping, the number of permissible permutations is a simple function of the vector $v = (v_1, \dots, v_n)$,

$$\#\mathbb{S}_n^\circ = \prod_{i=1}^n (1 + v_i).$$

The joint distribution of Y and Π has density

$$g(y_1, \dots, y_n, \pi_1, \dots, \pi_n) \;=\; \prod_{i=1}^n f(y_{\pi_i} - x_i)$$

for $y_1 < y_2 < \cdots < y_n$ and $\pi \in \mathbb{S}_n^\circ$, where the density is with respect to the product of the n-dimensional Lebesgue measure on $\mathbb{R}^n$ and the counting measure on $\mathbb{S}_n$. Regarding the conditional density of Π given $Y = y$ as a function of π we call $\hat{\pi} \in \mathbb{S}_n$ a *maximum a posteriori* (MAP) reconstruction or prediction (or perhaps retrodiction) of the value π of Π if it maximizes this function, or equivalently

$$\pi \;\mapsto\; \sum_{i=1}^n \log f(y_{\pi_i} - x_i) \tag{6}$$

on $\mathbb{S}_n^\circ$. By design, $\hat{\Pi} := \hat{\pi}(Y)$ maximizes the probability that the unknown matching Π is equal to $\hat{\Pi}$ among all functions of the data Y. Our first result gives these maximizers for the classes of distributions introduced at the end of Section 1.

Theorem 2.1. (a) *If $Q = \mathrm{Exp}(\eta)$ then all elements of $\mathbb{S}_n^\circ$ are MAP predictors.*
(b) *If $Q \in \mathcal{P}_{\mathrm{lcc}}$ then $\hat{\pi} = \mathrm{id}$ is an MAP predictor; if $Q \in \mathcal{P}_{\mathrm{lcc}}^\circ$ then the MAP predictor is unique.*
(c) *If $Q \in \mathcal{P}_{\mathrm{lcv}}$ then an MAP predictor is given recursively by*

$$\begin{aligned} \hat{\pi}_n &= \min\{j : y_j > x_n\}, \\ \hat{\pi}_i &= \min\big\{j : y_j > x_i,\, j \notin \{\hat{\pi}_k : k > i\}\big\}, \quad 1 \le i < n. \end{aligned}$$

Moreover, if $Q \in \mathcal{P}_{\mathrm{lcv}}^\circ$, then the MAP predictor is unique.

Proof. For exponential distributions we obtain

$$\begin{aligned} \sum_{i=1}^n \log f(y_{\pi_i} - x_i) \;&=\; n \log \eta - \eta \sum_{i=1}^n (y_{\pi_i} - x_i) \\ &=\; n \log \eta - \eta \Big(\sum_{i=1}^n y_i - \sum_{i=1}^n x_i \Big), \end{aligned}$$

which does not depend on π.

For the proofs of the remaining parts we need some properties of $\mathbb{S}_n^\circ$. The set W_n defined in (5) carries a canonical digraph structure, with a pair $(w, \tilde{w})$ of elements of W_n being a directed edge if and only if there is a $j \in \{1, \dots, n\}$ such that

$$\tilde{w}_i \;=\; \begin{cases} w_j + 1, & \text{if } i = j, \\ w_i, & \text{otherwise.} \end{cases} \tag{7}$$

There are no loops, the vector $\overline{w} = (1, 2, \dots, n)$ is at the top, and $\underline{w} = (1, 2 - v_2, \dots, n - v_n)$ is at the bottom. The digraph is connected in the sense that, for each $w \in W_n$, there is path connecting the bottom to w and a path connecting w to the top. The mapping Ψ transports this structure to $\mathbb{S}_n^\circ$ if we take $(\pi, \tilde{\pi})$ to be

an edge in $\mathbb{S}_n^\circ$ if and only if $(w, \tilde{w})$ is an edge in W_n, with $\pi = \Psi(w)$, $\tilde{\pi} = \Psi(\tilde{w})$. It is easy to check that the transition from w to $\tilde{w}$ in (7) corresponds to the removal of an inversion in $\pi = \Psi(w)$: Position j is swapped with position $i < j$, where customer i is the first to leave after customer j (we have $w_j < j$, so this set is not empty).

The properties that we need can now be formulated as follows: First, if $(\pi, \tilde{\pi})$ is a directed edge then $\tilde{\pi}$ arises from π by a transposition that removes an inversion; second, $\mathbb{S}_n^\circ$ is connected in the sense that any permissible permutation π can be transformed into $\Psi(\underline{w})$ or $\Psi(\overline{w})$ by a finite number of transpositions without leaving $\mathbb{S}_n^\circ$, i.e. with all intermediate results being permissible. The algorithm in part (c) is easily seen to provide the bottom $\Psi(\underline{w})$ of $\mathbb{S}_n^\circ$.

Now suppose that f is log-concave and that $(\pi, \tilde{\pi})$ is an edge in $\mathbb{S}_n^\circ$, so that $\tilde{\pi} = \pi \circ \tau(i,j) \in \mathbb{S}_n^\circ$ for some transposition $\tau(i,j)$. We may assume that $i < j$, and then $\pi_i > \pi_j$. From the fact that $\pi, \tilde{\pi} \in \mathbb{S}_n^\circ$ we obtain that $x_i < x_j < y_{\pi_j} < y_{\pi_i}$; in particular, $a := x_j - x_i$, $b := y_{\pi_j} - x_j$ and $c := y_{\pi_i} - y_{\pi_j}$ are all positive. We need to maximize the function $h : \mathbb{S}_n^\circ \to \mathbb{R}$,

$$h(\pi) := \sum_{k=1}^{n} \log f(y_{\pi_k} - x_k),$$

and we obtain

$$\begin{aligned} h(\pi) - h(\tilde{\pi}) &= \log(f(y_{\pi_i} - x_i)) + \log(f(y_{\pi_j} - x_j)) \\ &\qquad - \log(f(y_{\pi_j} - x_i)) - \log(f(y_{\pi_i} - x_j)) \\ &= \log(f(a+b+c)) + \log(f(b)) \\ &\qquad - \log(f(a+b)) - \log(f(b+c)). \end{aligned} \tag{8}$$

For $\log f$ concave this is easily seen to be ≤ 0, in the case of strict concavity it is even strictly negative (see also Lemma 1 in [1]). This means that by removing an inversion we (strictly) increase the value of h; in view of the structure of $\mathbb{S}_n^\circ$ this implies both statements in (b).

For $\log f$ (strictly) convex the difference is (strictly) positive, hence we obtain the (unique) MAP in the situation (c) by moving to the bottom, which again is possible without leaving $\mathbb{S}_n^\circ$. □

Remark 2.2. (a) The algorithm in part (c) of Theorem 2.1 assigns to each arrival the earliest possible departure, working backwards from the latest arrival, whereas with the solution $\hat{\pi} = \mathrm{id}$ in part (b) departures are delayed as long as possible. Note that the MAP does not depend on the data in the log-concave case; for log-convex densities it depends on the data y only via the set $\mathbb{S}_n^\circ$ of permissible distributions.

(b) For other popular distributions that are neither log-convex nor log-concave, such as the lognormal distribution, we do not have explicit expressions for the MAP predictor(s) $\hat{\pi}$. Algorithmically, the maximization of the function in (6) can be seen as an instance of the *assignment problem*, see e.g. [14], which can be solved with complexity $O(n^3)$, hence we can still find the solutions in a reasonable amount of time if n is not too large. This is an improvement over the naive approach of working through the whole of $\mathbb{S}_n^\circ$ element by element, as this set may be quite large even for moderate values of n. ◁

It should be clear that a reconstruction in the sense of $P(\hat{\Pi} = \Pi)$ being large is somewhat unrealistic in all but extreme cases (such as, in queueing terminology,

cases with a very low traffic intensity). Instead, we could aim for a prediction that is close to the true value in the mean. This requires some loss function, which in turn could be based on a notion of distance. Formally, we have a metric d on $\mathbb{S}_n$ and we call a function ϕ of the data y a minimum mean distance (MMD) predictor with respect to d, if ϕ minimizes $Ed(\Pi, \phi(Y))$ among all measurable functions $\phi : \mathbb{R}^n \to \mathbb{S}_n$. If the transition kernel $K(y, \cdot)$ is a version of $\mathcal{L}(\Pi|Y = y)$ then we obtain an MMD predictor by minimizing $\pi \mapsto \int d(\pi, \sigma)\, K(y, d\sigma)$ for the individual y-values. Of course, MAP prediction for discrete random variables can also be seen as minimizing an expected distance, if we let the distance of π and σ be 1 if $\pi \neq \sigma$, and 0 otherwise (the discrete metric). There are obvious analogies with concepts from frequentist and Bayesian statistics; indeed, what we call MAP prediction occasionally appears as 'maximum likelihood prediction' in the literature.

Our basic reference for distances of permutations and their statistical relevance is [12]. We now consider minimum mean distance prediction of Π if we measure distance by M.G. Kendall's τ. The standard definition of $d_\tau(\pi, \sigma)$ is as the minimum number of transpositions $\tau(i, i+1)$, $1 \leq i < n$, of direct neighbours required to turn π^{-1} into σ^{-1}. There are various other, equivalent, definitions; for us, it will be convenient to use

$$d_\tau(\sigma, \pi) = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} D_{i,j} \tag{9}$$

with

$$D_{i,j} = D_{i,j}(\sigma, \pi) := \begin{cases} 1, & \text{if } (\sigma_i - \sigma_j)(\pi_i - \pi_j) < 0, \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

Here $D_{i,j}(\sigma, \pi)$ indicates whether or not π and σ assign different order to the values i and j. In particular, with the bijection Ψ in (3) and (5), we have

$$d_\tau(\mathrm{id}, \pi) = \sum_{i=1}^{n} (i - w_i) \quad \text{if } \pi = \Psi(w_1, \ldots, w_n). \tag{11}$$

Note that the Cayley distance arises if instead of neighbour transpositions we allow all transpositions $\tau(i, j)$, $1 \leq i < j \leq n$. Apart from being more detailed in a sense, Kendall's τ also seems to us to be more appropriate in the present context as the delay problem makes use of the ordering of the set on which the permutations work via the ordering in time of the events in question.

Stated somewhat informally, the proof of Theorem 2.1 shows that, for log-concave densities, undoing an inversion increases the a posteriori probability, and that we obtain a decrease in the log-convex case. For the first of these it turns out that the MAP predictor is also the MMD predictor; for the corresponding statement in the log-convex case we need an additional assumption. Note that, with $\mathbb{S}_n^\circ = \mathbb{S}_n$, the permutation in part (b) is the same as the permutation defined in Theorem 2.1 (c).

Theorem 2.3. (a) *If $Q \in \mathcal{P}_{\mathrm{lcc}}$ then $\phi(Y) \equiv \mathrm{id}$ is an MMD predictor for Π with respect to Kendall's τ-distance. If $Q \in \mathcal{P}_{\mathrm{lcc}}^\circ$ then the MMD predictor is unique.*
(b) *Suppose that $\mathbb{S}_n^\circ = \mathbb{S}_n$. If $Q \in \mathcal{P}_{\mathrm{lcv}}$, then $\phi(Y) \equiv \hat{\pi}$ with $\hat{\pi}_i = n - i + 1$ for $i = 1, \ldots, n$ is an MMD predictor for Π with respect to Kendall's τ-distance. If $Q \in \mathcal{P}_{\mathrm{lcc}}^\circ$ then the MMD predictor is unique.*

Proof. (a) Let $K(y,\cdot)$ be the conditional distribution of Π given $Y = y$. The computation in (8) shows that for log-concave densities, and with

$$B_{i,j} := \{\pi \in \mathbb{S}_n : \pi_i > \pi_j\}, \quad 1 \le i < j \le n,$$

we have that

$$\{\pi \circ \tau(i,j) : \pi \in B_{i,j} \cap \mathbb{S}_n^\circ\} \subset B_{i,j}^c \cap \mathbb{S}_n^\circ \tag{12}$$

as well as

$$K\big(y, \{\pi \circ \tau(i,j)\}\big) \ge K\big(y, \{\pi\}\big) \quad \text{for all } \pi \in B_{i,j} \cap S_n^\circ. \tag{13}$$

In particular,

$$K(y, \{\pi : \pi_i > \pi_j\}) \le K(y, \{\pi : \pi_i < \pi_j\}) \quad \text{for all } i,j \text{ with } 1 \le i < j \le n, \tag{14}$$

as $\pi \mapsto \pi \circ \tau(i,j)$ maps $B_{i,j} \cap \mathbb{S}_n^\circ$ one-to-one onto a subset of $\{\pi \in \mathbb{S}_n^\circ : \pi_i < \pi_j\}$.

We now use the representation given by (9) and (10). Clearly, for any $\sigma \in \mathbb{S}_n^\circ$, and again with $1 \le i < j \le n$,

$$E[D_{i,j}(\sigma,\Pi)|Y] \;=\; \begin{cases} K(Y, \{\Pi_i > \Pi_j\}), & \text{if } \sigma_i < \sigma_j, \\ K(Y, \{\Pi_i < \Pi_j\}), & \text{if } \sigma_i > \sigma_j, \end{cases}$$

which reduces to $K(Y, \{\Pi_i > \Pi_j\})$ if $\sigma = \mathrm{id}$. With (14) this leads to

$$ED_{i,j}(\sigma,\Pi) = E\big(E[D_{i,j}(\sigma,\Pi)|Y]\big) \ge E\big(E[D_{i,j}(\mathrm{id},\Pi)|Y]\big) = ED_{i,j}(\mathrm{id},\Pi) \tag{15}$$

so that

$$Ed_\tau(\sigma,\Pi) \;=\; \sum_{i=1}^{n-1}\sum_{j=i+1}^{n} ED_{i,j}(\sigma,\Pi) \;\ge\; \sum_{i=1}^{n-1}\sum_{j=i+1}^{n} ED_{i,j}(\mathrm{id},\Pi) \;=\; Ed_\tau(\mathrm{id},\Pi).$$

Again, for strictly log-concave densities, the inequality will be strict.

(b) In the log-convex case, the inequality in (13) is reversed, and the additional assumption $\mathbb{S}_n^\circ = \mathbb{S}_n$ implies that we have equality in (12). Replacing $\sigma = \mathrm{id}$ by $\sigma = \hat{\pi}$ we can now proceed as in the proof of the first part of the theorem. □

Remark 2.4. (a) It follows from the first part of the proof of Theorem 2.1 that for exponential delay distributions the conditional distribution of Π given Y is the discrete uniform distribution on the set of permissible permutations; we abbreviate this to

$$\mathcal{L}(\Pi|Y) = \mathrm{unif}(\mathbb{S}_n^\circ). \tag{16}$$

Indeed, as we regard arrival times as fixed, $\mathbb{S}_n^\circ$ is a function of Y. Moreover, the distribution of the distance between the predictor that is best in the sense of Theorem 2.3 (a) and the true matching can then be written as a convolution of uniform distributions,

$$\mathcal{L}\big(d_\tau(\mathrm{id},\Pi)\big|Y\big) \;=\; \mathop{\star}_{i=1}^{n} \mathrm{unif}(\{0,\ldots,V_i\}), \tag{17}$$

where V_i again denotes the number of customers that are in the system at the time of the ith arrival. To see this, we use (11) and (16), and we note that uniformity on $\mathbb{S}_n^\circ$ implies uniformity on the domain of the bijective function Ψ in (5). In particular,

$$E[d_\tau(\mathrm{id},\Pi)\big|Y] = \frac{1}{2}\sum_{i=1}^{n} V_i, \quad \mathrm{var}[d_\tau(\mathrm{id},\Pi)\big|Y] = \frac{1}{12}\sum_{i=1}^{n} V_i(V_i+2). \tag{18}$$

This will be taken up in Section 4.

(b) In the exponential case we cannot expect uniqueness if $\mathbb{S}_n^\circ = \mathbb{S}_n$ as the distance is right invariant in the sense that $d_\tau(\pi,\sigma) = d_\tau(\pi\circ\nu, \sigma\circ\nu)$ for all $\pi,\sigma,\nu \in \mathbb{S}_n$. Even if $\mathbb{S}_n^\circ \neq \mathbb{S}_n$ the minimizer may not be unique: For example, in the case $n = 3$ and $x_1 < x_2 < y_1 < x_3 < y_2 < y_3$ we have $\mathbb{S}_3^\circ = \{\mathrm{id}, (1,3,2), (2,1,3), (3,1,2)\}$ and some straightforward computations show that for $\sigma = \mathrm{id}$ and $\sigma = (2,1,3)$ the distance $d_\tau(\sigma,\Pi)$ to $\Pi \sim \mathrm{unif}(\mathbb{S}_3^\circ)$ has a binomial distribution with parameters 2 and 1/2; for the other two permutations the distance is uniformly distributed on $\{0,1,2,3\}$. Hence there are two permutations that minimize $\sigma \mapsto E[d_\tau(\sigma,\Pi)|\mathbb{S}_3^\circ]$. Also, for exponential delay distributions, choosing a permutation at random will provide an optimal reconstruction from the MAP point of view. With MMD this is no longer the case: With $\mathbb{S}_3^\circ$ as above, and X and Π independent and uniformly distributed on $\mathbb{S}_3^\circ$, we obtain the expected distances $Ed_\tau(\mathrm{id},\Pi) = 1$ and $Ed_\tau(X,\Pi) = 5/4$.

(c) For the data $x = (1,2,4)$, $y = (3,5,6)$, and with Q the gamma distribution with parameters $\beta = 1/2$ and $\eta = 1$, we obtain the MMD predictor $\pi = (2,1,3)$, whereas the permutation $\hat\pi$ defined in Theorem 2.1 (c) is given by $\hat\pi = (3,1,2)$. Hence, in contrast to the MAP situation, this $\hat\pi$ will not in general be the MMD predictor in the log-convex case — an additional assumption on the structure of $\mathbb{S}_n^\circ$ such as in Theorem 2.3 (b) is needed.

(d) In analogy to the MAP case the question arises how an MMD predictor could be obtained if the lifetime distribution is not of the type considered in Theorem 2.3. We define a matrix A and a vector b, both indexed by the permissible permutations, by

$$A = \big(a(\pi,\sigma)\big)_{\pi,\sigma\in\mathbb{S}_n^\circ}, \text{ with } a(\pi,\sigma) := d_\tau(\pi,\sigma) \text{ for all } \pi,\sigma \in \mathbb{S}_n^\circ,$$

and

$$b = (b_\pi)_{\pi\in\mathbb{S}_n^\circ}, \quad \text{with } b_\pi := \prod_{i=1}^n f_Z(y_{\pi_i} - x_i) \text{ for all } \pi \in \mathbb{S}_n^\circ,$$

respectively. Finding the MMD for a particular y then amounts to finding a component of the vector Ab that has minimal value. We are not aware of a general algorithmic simplification in the style of Remark 2.2(b), but the distance minimization can be carried out separately for the busy periods; see also the proof of Theorem 4.2 below. ◁

3. A test of exponentiality

We now turn to the situation where the original data are reduced to the order of arrivals and departures. Equivalently, we observe the value π of the true matching Π and the value v of the random vector $V = (V_1,\dots,V_n)$ where again V_i denotes the number of customers that are in the system at the time of arrival of the ith customer. If the delay times have a log-concave density then we know from the proof of Theorem 2.1 that the conditional density of Π given Y decreases if the number of inversions increases, it increases in the log-convex case. This motivates the use of $d_\tau(\mathrm{id},\pi)$ as a test statistic for testing exponentiality against log-convex or log-concave alternatives. In the queueing context such procedures test for M/M/∞ within M/G/∞.

For a formal definition let $\alpha \in (0,1)$ be given and let $\phi_+ : \mathbb{S}_n \times \mathbb{N}_0^n \to [0,1]$ be defined by

$$\phi_+(\pi, v) \;=\; \begin{cases} 1, & \text{if } d_\tau(\mathrm{id}, \pi) > c_+(\alpha), \\ \gamma_+(\alpha), & \text{if } d_\tau(\mathrm{id}, \pi) = c_+(\alpha), \\ 0, & \text{if } d_\tau(\mathrm{id}, \pi) < c_+(\alpha), \end{cases} \tag{19}$$

where $c_+(\alpha)$ is the upper α-quantile associated with $Q_0 := \mathcal{L}(d_\tau(\mathrm{id}, \Pi))$, $\Pi \sim \mathrm{unif}(\mathbb{S}_n^\circ)$, and $\gamma_+(\alpha) \in [0,1)$ is such that

$$Q_0\big((c_+(\alpha), \infty)\big) + \gamma_+(\alpha)\, Q_0\big(\{c_+(\alpha)\}\big) \;=\; \alpha. \tag{20}$$

Similarly, we define another test function ϕ_- by

$$\phi_-(\pi, v) \;=\; \begin{cases} 1, & \text{if } d_\tau(\mathrm{id}, \pi) < c_-(\alpha), \\ \gamma_-(\alpha), & \text{if } d_\tau(\mathrm{id}, \pi) = c_-(\alpha), \\ 0, & \text{if } d_\tau(\mathrm{id}, \pi) > c_-(\alpha), \end{cases} \tag{21}$$

where $c_-(\alpha)$ is now the lower α-quantile of Q_0 and $\gamma_-(\alpha)$ is such that

$$Q_0\big([0, c_-(\alpha))\big) + \gamma_-(\alpha)\, Q_0\big(\{c_-(\alpha)\}\big) \;=\; \alpha.$$

In both cases, the variable v enters the test via $\mathbb{S}_n^\circ$; see (5). We point out that the critical values do not depend on η, which is a consequence of the fact that for exponential delay distributions the conditional distribution of Π and hence of $d_\tau(\mathrm{id}, \Pi)$ does not depend on this parameter; see (16). In fact, this distribution is a convolution of discrete uniform distributions, which makes it easy to obtain the critical values. For n small this can be done via Fourier transforms, together with the fast Fourier transform (FFT) algorithm. For moderate or large n and v_i's the normal approximation gives satisfying results; see also Section 4.

Our next theorem collects some properties of these tests. Statements (22) and (24) imply that both are unbiased and that they are similar on the 'boundary' $\mathcal{P}_{\mathrm{lcv}} \cap \mathcal{P}_{\mathrm{lcc}} = \{\mathrm{Exp}(\eta) : \eta > 0\}$, see [15, Chapter 4.3]. Statements (23) and (25) can be regarded as strong unbiasedness 'inside' the hypothesis resp. alternative.

$$\begin{tikzcd}
 & & (\Omega, \mathcal{A}, P_Q) \arrow[dll, "\Pi"'] \arrow[d, "\tilde{Y}"] \arrow[dr, "Y"] \arrow[drr, "V"] & & \\
\mathbb{N}_0 & \mathbb{S}_n \arrow[l, "d_\tau(\mathrm{id},\,\cdot\,)"'] & \mathbb{R}^n \arrow[l, "\text{order}"'] \arrow[r, "\text{sort}"] & \mathbb{R}^n \arrow[r, "\psi"] & \mathbb{N}_0^n
\end{tikzcd}$$

FIGURE 2. General structure of the model variables

The random variables that appear in our model are all defined on some measurable space $(\Omega, \mathcal{A})$. We write P_Q for the probability measure on $(\Omega, \mathcal{A})$ under which the lifetimes are independent and have distribution Q, E_Q denotes expectation with respect to P_Q.

Theorem 3.1. (a) *With* ϕ_+ *as in* (19),

$$E_Q[\phi_+(\Pi,V)|V] \begin{cases} = \alpha, & \text{if } Q = \mathrm{Exp}(\eta) \text{ for some } \eta > 0, \\ \leq \alpha, & \text{if } Q \in \mathcal{P}_{\mathrm{lcc}}, \\ \geq \alpha, & \text{if } Q \in \mathcal{P}_{\mathrm{lcv}}, \end{cases} \tag{22}$$

and

$$E_Q\phi_+(\Pi,V) \begin{cases} < \alpha, & \text{if } Q \in \mathcal{P}^\circ_{\mathrm{lcc}}, \\ > \alpha, & \text{if } Q \in \mathcal{P}^\circ_{\mathrm{lcv}}. \end{cases} \tag{23}$$

(b) *Similarly, with* ϕ_- *as in* (21)

$$E_Q[\phi_-(\Pi,V)|V] \begin{cases} = \alpha, & \text{if } Q = \mathrm{Exp}(\eta) \text{ for some } \eta > 0, \\ \leq \alpha, & \text{if } Q \in \mathcal{P}_{\mathrm{lcv}}, \\ \geq \alpha, & \text{if } Q \in \mathcal{P}_{\mathrm{lcc}}, \end{cases} \tag{24}$$

and

$$E_Q\phi_-(\Pi,V) \begin{cases} < \alpha, & \text{if } Q \in \mathcal{P}^\circ_{\mathrm{lcv}}, \\ > \alpha, & \text{if } Q \in \mathcal{P}^\circ_{\mathrm{lcc}}. \end{cases} \tag{25}$$

Proof. Let $\alpha \in (0,1)$ be given. The various random quantities are summarized in Figure 2, where 'order' and 'sort' are canonically defined (see e.g. their use in the statistical language R), and $\psi = (\psi_1, \ldots, \psi_n)$ is given by

$$\psi_i(y_1, \ldots, y_n) := i - 1 - \#\{1 \leq j < n : y_j < x_i\},$$

see also (4).

We abbreviate P_Q, E_Q to P_η, E_η if $Q = \mathrm{Exp}(\eta)$. Remark 2.4 (a) and (20) lead to

$$E_\eta[\phi_\pm(\Pi,V)|V] = \alpha \quad \text{for all } \eta > 0.$$

Integrating out the variable V we see that the tests have the exact rejection probability α on $\{\mathrm{Exp}(\eta) : \eta > 0\}$.

For the proof of the other assertions we first consider the conditional distribution of $d_\tau(\mathrm{id},\Pi)$ given Y. The structural considerations at the beginning of Section 2 imply that this distribution has support $\{0, 1, \ldots, k_{\max}\}$ with

$$k_{\max} = k_{\max}(Y) := \sum_{i=1}^n \psi_i(Y) = \sum_{i=1}^n V_i.$$

Further, the argument in the proof of Theorem 2.1 can be used to show that

$$k \mapsto \frac{P_Q(d_\tau(\mathrm{id},\Pi) = k|Y)}{P_\eta(d_\tau(\mathrm{id},\Pi) = k|Y)}$$

is strictly decreasing on $\{0, 1, \ldots, k_{\max}\}$ if $Q \in \mathcal{P}^\circ_{\mathrm{lcc}}$, provided that Y is such that $k_{\max}(Y) > 0$. An elementary argument (based, for example, on the proof of a comparison result for likelihood ratio and stochastic order given in [16, p. 12]) shows that this implies that, again for $Q \in \mathcal{P}^\circ_{\mathrm{lcc}}$,

$$P_Q\big(d_\tau(\mathrm{id},\Pi) \geq k\big|Y\big) < P_\eta\big(d_\tau(\mathrm{id},\Pi) \geq k\big|Y\big) \tag{26}$$

for $k = 1, \ldots, k_{\max}$. Further, in (26) we may condition on V instead of Y in view of the fact that V is a function of Y (integrating over $\{y : \psi(y) = v\}$ leads to the conditional probability given $V = v$).

The strict stochastic ordering of the distributions of the test statistics in turn leads to a corresponding strict lexicographic ordering of the pairs $(c_\pm(\alpha), \gamma_\pm(\alpha))$ derived from the distribution of $d_\tau(\mathrm{id}, \Pi)$ under P_Q and P_η, hence

$$\begin{aligned} E_Q[\phi_+(\Pi, V)|V = v] &< \alpha \quad \text{for all } Q \in \mathcal{P}^\circ_{\mathrm{lcc}},\ v \neq 0, \\ E_Q[\phi_-(\Pi, V)|V = v] &> \alpha \quad \text{for all } Q \in \mathcal{P}^\circ_{\mathrm{lcc}},\ v \neq 0. \end{aligned}$$

Here we have written 0 for the n-vector with all components equal to 0. Note that $v \neq 0$ is equivalent to $\#\mathbb{S}^\circ_n > 1$. As the $\tilde{Y}$-variables can take arbitrarily large values, we always have $P_Q(V = 0) < 1$.

Putting this together we see that for $Q \in \mathcal{P}^\circ_{\mathrm{lcc}}$ the rejection probability of the test ϕ_+ is strictly smaller than the preassigned level, and that it is strictly greater for ϕ_-. This proves the upper part of (23) and the lower part of (25). A similar argument, with $\mathcal{P}^\circ_{\mathrm{lcv}}$ instead of $\mathcal{P}^\circ_{\mathrm{lcc}}$, provides the respective other parts.

Removing the qualifier 'strong' in the above arguments, which corresponds to the transition from $\mathcal{P}^\circ_{\mathrm{lcv}}$ to $\mathcal{P}_{\mathrm{lcv}}$ or from $\mathcal{P}^\circ_{\mathrm{lcc}}$ to $\mathcal{P}_{\mathrm{lcc}}$, we obtain inequalities instead of strict inequalities, and this leads to the remaining statements in (22) and (24). □

The conditioning on V means that the tests are invariant under scaling in the sense that the decision will not change if we multiply the arrival and departure data by a fixed positive constant.

The above procedures can easily be adapted to the two-sided case, i.e. for testing $Q = \mathrm{Exp}(\eta)$ for some $\eta > 0$ against $Q \in \mathcal{P}^\circ_{\mathrm{lcv}} \cup \mathcal{P}^\circ_{\mathrm{lcc}}$.

4. Asymptotics

In the previous sections we had a fixed finite sample size n and the arrival times were also regarded as fixed. We now investigate the asymptotic behaviour of the procedures under the assumption that arrivals occur at the points of a Poisson process with constant rate λ. This is, roughly, the M/G/∞ queueing model mentioned in the introduction, but we 'close admittance' after the arrival of the nth customer. Formally, we have a sequence $X = (X_i)_{i\in\mathbb{N}}$ of arrival times, with $X_1 = 0$ and $X_{i+1} - X_i$, $i \in \mathbb{N}$, independent random variables, with distribution $\mathrm{Exp}(\lambda)$. Further, $\tilde{Y}_i = X_i + Z_i$ where Z_i, $i \in \mathbb{N}$, are independent, and all Z_i's have distribution Q. The sequences X and $(Z_i)_{i\in\mathbb{N}}$ are also assumed to be independent. Let $Y = (Y_{n,i})_{n\in\mathbb{N}, i=1,\dots,n}$ be such that $Y_{n,1}, \dots, Y_{n,n}$ are the order statistics for $\tilde{Y}_1, \dots, \tilde{Y}_n$ (we need a second index here as later arrivals may depart earlier than previous ones, see the cases of Mozart and Tchaikovsky in our running example). Again, let

$$V_i := i - 1 - \#\{1 \le j \le i : Y_{i,j} \le X_i\} \tag{27}$$

be the number of customers in the system at the time of the ith arrival (it is easy to see that we do not need a second index for the V-variables; the right hand side of (27) does not change if $Y_{i,j}$ is replaced by $Y_{n,j}$ with an arbitrary $n \ge i$). Finally, let Π_n be the correct matching for the first n arrivals and departures. Our first limit result is conditional on the sequence X and the triangular array Y.

Theorem 4.1. *Suppose that $Q = \mathrm{Exp}(\eta)$ and put $\rho = \lambda/\eta$, Then, with probability 1 as $n \to \infty$,*

$$\frac{1}{n}\, E[d_\tau(\mathrm{id}, \Pi_n)|X, Y] \;\to\; \frac{\rho}{2},$$

$$\frac{1}{n}\, \mathrm{var}[d_\tau(\mathrm{id}, \Pi_n)|X, Y] \;\to\; \frac{\rho(\rho+3)}{12},$$

$$\mathcal{L}\left(\frac{d_\tau(\mathrm{id}, \Pi_n) - E[d_\tau(\mathrm{id}, \Pi_n)|X, Y]}{\mathrm{var}[d_\tau(\mathrm{id}, \Pi_n)|X, Y]^{1/2}} \,\middle|\, X, Y \right) \;\to_{\mathrm{distr}}\; N(0,1).$$

Proof. We need the stochastic dynamics of the sequence $V = (V_i)_{i\in\mathbb{N}}$ as given in [17, Chapter 2.2]: V is an irreducible and aperiodic Markov chain with stationary distribution $\pi = \mathrm{Po}(\rho)$, and the moments EV_i^l, $l \in \mathbb{N}$, converge to the corresponding moments of $\mathrm{Po}(\rho)$, where we have written $\mathrm{Po}(\rho)$ for the Poisson distribution with mean ρ.

By the ergodic theorem for Markov chains, see Theorem 2 in [9, Section I.15], we then have that, with probability 1,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} V_i \;=\; \sum_{j=0}^{\infty} j\pi_j \;=\; \rho,$$

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} V_i(V_i+2) \;=\; \sum_{j=0}^{\infty} j(j+2)\pi_j \;=\; \rho^2 + 3\rho, \tag{28}$$

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} V_i^4 \;=\; \sum_{j=0}^{\infty} j^4\pi_j \;<\; \infty. \tag{29}$$

Using this together with (18) we obtain the first two statements of the theorem. For the distributional limit we start with (17), which in the current notation becomes

$$\mathcal{L}\big(d_\tau(\mathrm{id}, \Pi_n)|X, Y\big) \;=\; \mathop{\star}_{i=1}^{n} \mathrm{unif}(\{0, \ldots, V_i\}).$$

By the Lyapunov condition for the Central Limit Theorem, see e.g. [6, Section 27], it is now enough to show that

$$\lim_{n\to\infty} \frac{\sum_{i=1}^{n} V_i^4}{\big(\frac{1}{12}\sum_{i=1}^{n} V_i(V_i+2)\big)^2} \;=\; 0$$

with probability 1. From (29) we obtain that the numerator is $O(n)$ with probability 1, for the denominator we use (28). □

We now consider the asymptotic behaviour of the tests in Section 3 in the above model. Let $\alpha \in (0,1)$ be fixed. We write $\phi_{n,+}$ and $\phi_{n,-}$ for the level-α tests based on the first n arrivals and departures. The following result shows that these tests are consistent 'inside' their respective hypothesis and alternative.

Theorem 4.2. (a) *If $Q \in \mathcal{P}^\circ_{\mathrm{lcc}}$ then*

$$\lim_{n\to\infty} E_Q\, \phi_{n,+}(\Pi, V) \;=\; 0 \quad \textit{and} \quad \lim_{n\to\infty} E_Q\, \phi_{n,-}(\Pi, V) \;=\; 1.$$

(b) *If $Q \in \mathcal{P}^\circ_{\mathrm{lcv}}$ then*

$$\lim_{n\to\infty} E_Q\, \phi_{n,-}(\Pi, V) \;=\; 0 \quad \textit{and} \quad \lim_{n\to\infty} E_Q\, \phi_{n,+}(\Pi, V) \;=\; 1.$$

Proof. For general Q the sequence $(V_i)_{i\in\mathbb{N}}$ will no longer be a Markov chain, but it still has a regenerative property in the sense that, with $\nu_0 := 1$ and

$$\nu_{k+1} \;:=\; \inf\{i > \nu_k : V_i = 0\} \quad \text{for all } k \in \mathbb{N},$$

the blocks $(W_k)_{k\in\mathbb{N}}$,

$$W_k \;:=\; (V_{\nu_{k-1}}, V_{\nu_{k-1}+1}, \ldots, V_{\nu_k-1})$$

are independent and identically distributed. This is a consequence of the decomposition of M/G/∞ queues into busy and idle periods; in particular,

$$P(\nu_k < \infty \text{ for all } k \in \mathbb{N}) \;=\; 1.$$

The blocks induce a decomposition of the set of permissible permutations (in our running example we have chosen the composers to demonstrate this: Obviously, for π to be permissible, we need $\pi(\{1,2,3,4,5\}) = \{1,2,3,4,5\}$ and $\pi(\{6,7\}) = \{6,7\}$). Consequently, if $n = \nu_k - 1$ for some $k \in \mathbb{N}$ then any $\pi \in \mathbb{S}_n^\circ$ can be written as

$$\pi = \pi^{k-1} \circ \pi^{k-2} \circ \cdots \circ \pi^1,$$

where π^l permutes the integers $\nu_{l-1}, \nu_{l-1}+1, \ldots, \nu_l - 1$ only. This in turn leads to an additive decomposition of the τ-distance,

$$d_\tau(\mathrm{id}, \pi) \;=\; \sum_{l=1}^{k-1} d_\tau(\mathrm{id}, \pi^l).$$

Conditionally on V we thus obtain a representation of the test statistics,

$$d_\tau(\mathrm{id}, \Pi_{\nu_k-1}) \;=\; \sum_{l=1}^{k-1} \xi_l,$$

where the ξ_l's are independent and identically distributed. By the law of large numbers we therefore have, conditionally on V,

$$\lim_{k\to\infty} \frac{1}{\nu_k} d_\tau(\mathrm{id}, \Pi_{\nu_k}) \;=\; E_Q[\xi_1 | V] =: \chi(Q) \quad \text{with probability 1.}$$

The strict stochastic order of the distance distributions that we used in the proof of Theorem 3.1 implies that $\chi(Q) < \chi(\mathrm{Exp}(\eta))$ whenever $Q \in \mathcal{P}_{\mathrm{lcc}}^\circ$. In connection with the conditional asymptotic normality in Theorem 4.1 this means that the conditional probability that $\phi_{\nu_k,+}$ rejects will tend to 0 as $k \to \infty$, and that the conditional probability that $\phi_{\nu_k,-}$ rejects will tend to 1 as $k \to \infty$, both with probability 1.

In order to be able to uncondition, i.e. to integrate out the sequence V, we need to pass from the sequence $(\nu_k)_{k\in\mathbb{N}}$ of random times to the full deterministic sequence $1, 2, \ldots$. This, however, is straightforward, using

$$d_\tau(\mathrm{id}, \Pi_{\nu_k-1}) \;\le\; d_\tau(\mathrm{id}, \Pi_n) \;\le\; d_\tau(\mathrm{id}, \Pi_{\nu_{k+1}-1})$$

on $\nu_k - 1 \le n < \nu_{k+1} - 1$, and the fact that $\nu_{k+1}/\nu_k \to 1$ with probability 1 as $k \to \infty$.

Taken together this proves (a). The arguments for the second part are similar. □

Specialized to the parametric situations in (1) and (2) the theorem gives

$$\lim_{n\to\infty} E_{\beta,\eta}\phi_+(\Pi, V) \;=\; \begin{cases} 0, & \beta > 1, \\ \alpha, & \beta = 1, \\ 1, & \beta < 1, \end{cases}$$

and a similar statement holds for ϕ_-.

Acknowledgements. We wish to thank Ludwig Baringhaus for helpful discussions. We are also grateful to the referee for drawing our attention to [17].

References

[1] An, M.Y. (1998) Logconcavity versus logconvexity: a complete characterization. *J. Econ. Theory* **80**, 350–367.
[2] Bagnoli, M. and Bergstrom, T. (2005) Log-concave probability and its applications. *Econom. Theory*, **26**, 445–469.
[3] Bai, Z. and Hsing, T. 2005) The broken sample problem. *Probab. Theory Related Fields* **131**, 528–552.
[4] Barlow, R.E. and Proschan, F. (1975) *Statistical Theory of Reliability and Life Testing*. Holt Rinehart and Winston, New York.
[5] Bäuerle, N. and Grübel, R. (2005) Multivariate counting processes: copulas and beyond. *ASTIN Bulletin* **35**, 379–408.
[6] Billingsley, P. (1986) *Probability and Measure*. 2nd ed., Wiley, New York.
[7] Bingham, N. H. and Pitts, S. M. (1999) Non-parametric estimation for the $M/G/\infty$ queue. *Ann. Inst. Statist. Math.* **51**, 71–97.
[8] Buchmann, B. and Grübel, R. (2003) Decompounding: an estimation problem for Poisson random sums. *Ann. Stat.* **31**, 1054–1074.
[9] Chung, K.L. (1967) *Markov Chains With Stationary Transition Probabilities*. Springer, Berlin.
[10] DeGroot, M.H., Feder, P.L. and Goel, P.K. (1971) Matchmaking. *Ann. Math. Stat.* **42**, 578–593.
[11] DeGroot, M.H. and Goel, P.K. (1980) Estimation of the correlation coefficient from a broken random sample. *Ann. Stat.* **8**, 264–278.
[12] Diaconis, P. (1988) *Group Representations in Probability and Statistics*. IMS Lecture Notes **11**. Institute of Mathematical Statistics, Hayward, CA.
[13] Hall, P. and Park, J. (2004) Non-parametric inference about service time distribution from indirect measurements. *J. R. Statist. Soc* B **66**, 861–875.
[14] Jungnickel, D. (1999) *Graphs, Networks and Algorithms*. Springer-Verlag, Berlin.
[15] Lehmann, E.L. (1959) *Testing Statistical Hypotheses*. Wiley, New York.
[16] Müller, A. and Stoyan, D. (2002) *Comparison Methods for Stochastic Models and Risk*. Wiley, Chichester.
[17] Takács, L. (1962) *Introduction to the Theory of Queues*. Oxford University Press, New York.

Institut für Mathematische Stochastik, Leibniz Universität Hannover, Postfach 6009, D-30060 Hannover, Germany
E-mail address: rgrubel@stochastik.uni-hannover.de

Institut für Mathematische Stochastik, Leibniz Universität Hannover, Postfach 6009, D-30060 Hannover, Germany
E-mail address: wegener@stochastik.uni-hannover.de