Computational Order Statistics

Colin Rose
Murray D. Smith

Using mathStatica [1], Rose and Smith [2] illustrate the automated derivation of the exact density of order statistics obtained from a random sample of size *n* on a *continuous* random variable. This article illustrates the new functionality in mathStatica to allow for *discrete* random variables. Moreover, by building on the new `Piecewise` capability in *Mathematica*, we are able to further generalise to the case of non-identically distributed random variables, thereby providing a completely flexible solution.

■ Introduction

Let X denote a continuous random variable with probability density function (pdf) $f(x)$ and cumulative distribution function (cdf) $F(x)$, and let $(X_1, X_2, \ldots, X_n)$ denote a random sample of size n drawn on X. Let $(X_{(1)}, X_{(2)}, \ldots, X_{(n)})$ denote the random sample ordered such that $X_{(1)} < X_{(2)} < \cdots < X_{(n)}$; then $(X_{(1)}, X_{(2)}, \ldots, X_{(n)})$ are collectively known as the *order statistics derived from the parent* X. For example, $X_{(1)} = \min(X_1, \ldots, X_n)$ is the smallest order statistic and corresponds to the sample minimum, and $X_{(n)}$ is the largest order statistic and corresponds to the sample maximum.

For a detailed discussion of the statistical theory pertaining to order statistics see, for example, [3, 4, 5, 6]. For the case in which X is a discrete random variable, see [7]. On the computational side, Rose and Smith [2, Section 9.4] use mathStatica to obtain the algebraic and numeric properties of order statistics derived from a continuous parent. Whereas Evans et al. [8] appear to be restricted to numeric-only calculations, they also consider settings in which the parent variable is sampled without replacement.

Section 1 illustrates briefly the case of a continuous parent, while Sections 2 and 3 extend to the case of a discrete parent. Section 4 relaxes the independent and identically distributed (iid) assumptions, thereby illustrating some of the new functionality of mathStatica.

We begin by loading mathStatica 1.5 or later.

```
In[1]:=  << mathStatica.m
```

■ 1. Continuous Parent Distribution

Let the continuous random variable X have a U-shaped distribution with pdf $f(x)$:

In[2]:= $\mathbf{f = \frac{1}{\pi\sqrt{A^2 - x^2}};}$

defined on a domain of support $(-A, A)$, where parameter $A > 0$:

```
In[3]:= domain[f] = {x, -A, A} && {A > 0};
```

Figure 1 plots the parent pdf's when parameter $A = 1$, 3 and 5:

```
In[4]:= PlotDensity[f /. A → {1, 3, 5}];
```

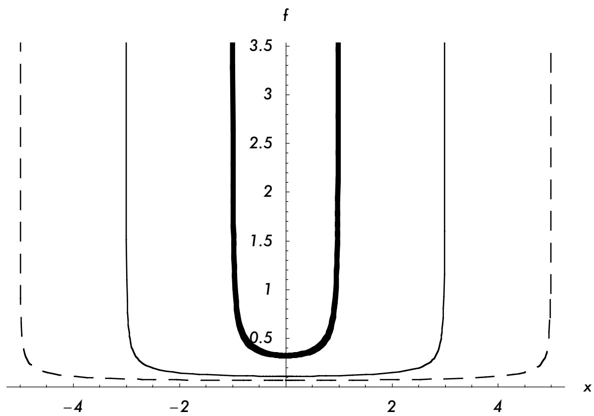

Figure 1.

For a sample of size n on X, the pdf of the r^{th} order statistic $X_{(r)}$ is given by the mathStatica function:

```
In[5]:= g = OrderStat[r, f]
```

Out[5]= $$\frac{2^{1-n}\,\pi^{-n}\left(\pi - 2\,\text{ArcTan}\left[\frac{x}{\sqrt{A^2-x^2}}\right]\right)^{n-r}\left(\pi + 2\,\text{ArcTan}\left[\frac{x}{\sqrt{A^2-x^2}}\right]\right)^{-1+r}\,n!}{\sqrt{(A-x)\,(A+x)}\,(n-r)!\,(-1+r)!}$$

with domain of support:

```
In[6]:= domain[g] = OrderStatDomain[r, f]

Out[6]= {x, -A, A} && {n ∈ Integers, r ∈ Integers, A > 0, 1 ≤ r ≤ n}
```

Figure 2 plots the pdf of $X_{(r)}$ as r increases from 1 to 10, given a sample size $n = 10$, with $A = 3$:

In[7]:= `PlotDensity[g /. {A → 3, r → Range[10], n → 10}];`

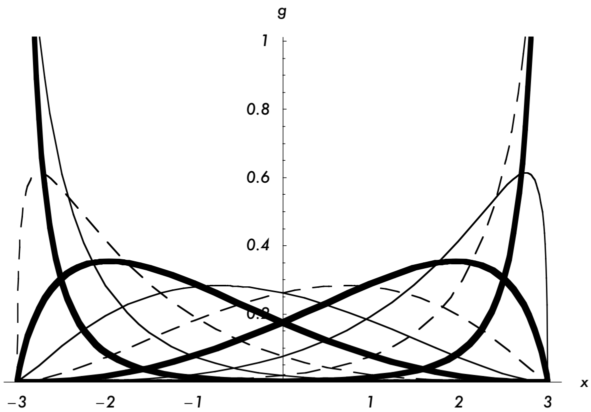

Figure 2.

The bivariate pdf of two order statistics $X_{(r)}$ and $X_{(s)}$, for $r < s$, is given by:

In[8]:= `OrderStat[{r, s}, f]`

Out[8]=
$$\left(\pi^{-1+r-s}\left(\frac{1}{2}+\frac{\mathrm{ArcTan}\left[\frac{x_r}{\sqrt{A^2-x_r^2}}\right]}{\pi}\right)^{-1+r}\left(-\mathrm{ArcTan}\left[\frac{x_r}{\sqrt{A^2-x_r^2}}\right]+\mathrm{ArcTan}\left[\frac{x_s}{\sqrt{A^2-x_s^2}}\right]\right)^{-1-r+s}\left(\frac{1}{2}-\frac{\mathrm{ArcTan}\left[\frac{x_s}{\sqrt{A^2-x_s^2}}\right]}{\pi}\right)^{n-s}\Gamma[1+n]\right)\Big/\left(\Gamma[r]\,\Gamma[1+n-s]\,\Gamma[-r+s]\sqrt{(A^2-x_r^2)(A^2-x_s^2)}\right)$$

■ 2. Discrete Parent Distribution (Function Form)

mathStatica 1.5 expands its functionality to the case of a discrete parent, again for *arbitrary* distributions. To illustrate, let the random variable X have a Poisson two-component-mix distribution with probability mass function (pmf) $f(x)$:

In[9]:=
$$f = \omega\,\frac{e^{-\lambda}\lambda^x}{x!} + (1-\omega)\,\frac{e^{-\theta}\theta^x}{x!};$$

and domain of support:

In[10]:= `domain[f] = {x, 0, ∞} && {λ > 0, θ > 0, 0 < ω < 1} && {Discrete};`

Here, θ and λ are the Poisson parameters, while ω is the mixing weight parameter. Figure 3 illustrates this pmf when $\lambda = 4$, $\theta = 20$ and $\omega = \frac{1}{2}$:

In[11]:=
```
PlotDensity[f /. {λ → 4, θ → 20, ω → 1/2}];
```

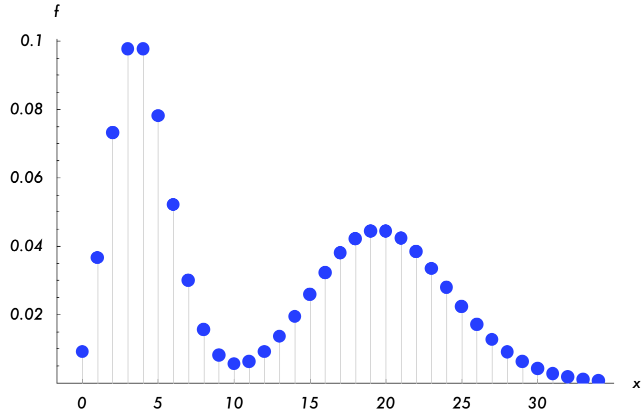

Figure 3.

Given a sample of size n, the pmf of the r^{th} order statistic $X_{(r)}$, denoted $g(x)$, is:

In[12]:=
```
g = OrderStat[r, f]
```

Out[12]=
$$\frac{1}{\text{Beta}[r,\ 1+n-r]}\left(-\text{Beta}\left[\frac{x\,(-(-1+\omega)\,\text{Gamma}[x,\ \theta]+\omega\,\text{Gamma}[x,\ \lambda])}{\Gamma[1+x]},\ r,\ 1+n-r\right]+\text{Beta}\left[-\frac{(1+x)\,((-1+\omega)\,\text{Gamma}[1+x,\ \theta]-\omega\,\text{Gamma}[1+x,\ \lambda])}{\Gamma[2+x]},\ r,\ 1+n-r\right]\right)$$

with domain of support:

In[13]:=
```
domain[g] = OrderStatDomain[r, f]
```

Out[13]=
```
{x, 0, ∞} && {n ∈ Integers, r ∈ Integers,
  θ > 0, λ > 0, 0 < ω < 1, 1 ≤ r ≤ n} && {Discrete}
```

Figure 4 plots the pmf of the minimum order statistic (i.e. $r = 1$) when $\lambda = 4$, $\theta = 20$ and $\omega = \frac{1}{2}$ and the sample size is $n = 10$:

In[14]:=
```
param = {λ → 4, θ → 20, ω → 1/2, r → 1, n → 10};
PlotDensity[g /. param];
```

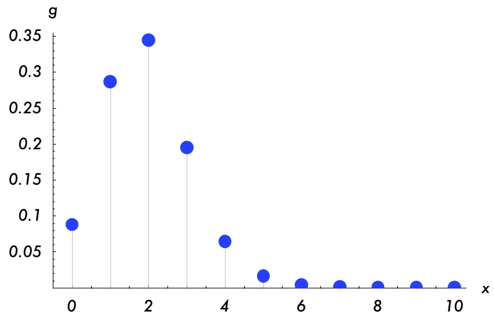

Figure 4.

□ Monte Carlo 'Check' of the Exact Solution We Have Just Plotted

Here is a single pseudorandom sample of size $n = 10$ drawn from the parent two-component-mix Poisson pmf $f(x)$:

In[16]:=
```
DiscreteRNG[10, f /. param]
```

Out[16]=
```
{3, 3, 22, 32, 12, 15, 2, 19, 4, 3}
```

If we want 50,000 such samples (each of size 10), the neatest approach is to generate all $50000 * 10$ drawings in one go:

In[17]:=
```
data = DiscreteRNG[50000 * 10, f /. param]; // Timing
```

Out[17]=
```
{1.96 Second, Null}
```

and then partition this `data` into 50,000 samples (each of size 10). We can then find the minimum of each of the 50,000 samples by mapping the `Min` function across each sample:

In[18]:=
```
samplemin = Map[Min, Partition[data, 10]];
```

If n is *very* large, efficient algorithms specifically designed for pseudorandom generation of order statistics exist; see [9]. Figure 5 plots the *empirical* relative frequency distribution of the sample minimum data (▲) together with the *exact* pmf of the sample minimum (●): a good match should see the former obscure the latter nearly everywhere.

In[19]:=
```
FrequencyPlotDiscrete[samplemin, g /. param];
```

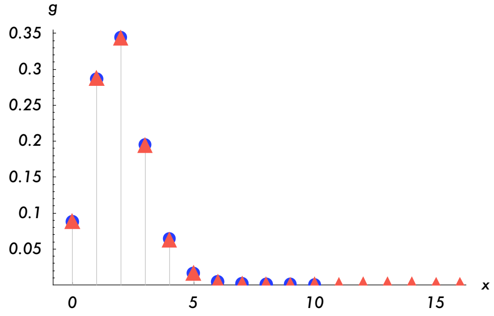

Figure 5.

■ 3. Discrete Parent Distribution (List Form)

Suppose we throw an *unfair* six-sided die onto a flat surface such as a table. Let X denote the upmost face of the die with pmf:

$P(X = x)$:	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{12}$	$\frac{3}{12}$
x :	1	2	3	4	5	6

Table 1. The pmf of X.

Using List Form, we enter this pmf as follows:

In[20]:=
```
f = {1/6, 1/6, 1/6, 1/6, 1/12, 3/12};
domain[f] = {x, {1, 2, 3, 4, 5, 6}} && {Discrete};
```

The pmf of the r^{th} order statistic $X_{(r)}$ is:

In[22]:=
```
g = OrderStat[r, f]
```

Out[22]=

$$\left\{\frac{\text{Beta}[\frac{1}{6}, r, 1+n-r]}{\text{Beta}[r, 1+n-r]}, \frac{-\text{Beta}[\frac{1}{6}, r, 1+n-r]+\text{Beta}[\frac{1}{3}, r, 1+n-r]}{\text{Beta}[r, 1+n-r]}, \frac{-\text{Beta}[\frac{1}{3}, r, 1+n-r]+\text{Beta}[\frac{1}{2}, r, 1+n-r]}{\text{Beta}[r, 1+n-r]}, \frac{-\text{Beta}[\frac{1}{2}, r, 1+n-r]+\text{Beta}[\frac{2}{3}, r, 1+n-r]}{\text{Beta}[r, 1+n-r]}, \frac{-\text{Beta}[\frac{2}{3}, r, 1+n-r]+\text{Beta}[\frac{3}{4}, r, 1+n-r]}{\text{Beta}[r, 1+n-r]}, \frac{-\text{Beta}[\frac{3}{4}, r, 1+n-r]+\text{Beta}[1, r, 1+n-r]}{\text{Beta}[r, 1+n-r]}\right\}$$

with domain of support:

```
In[23]:= domain[g] = OrderStatDomain[r, f]
```

```
Out[23]= {x, {1, 2, 3, 4, 5, 6}} &&
          {n ∈ Integers, r ∈ Integers, 1 ≤ r ≤ n} && {Discrete}
```

Here, for example, is a plot of the pmf of the third order statistic $X_{(3)}$ when the sample size $n = 10$:

```
In[24]:= PlotDensity[g /. {r → 3, n → 10}];
```

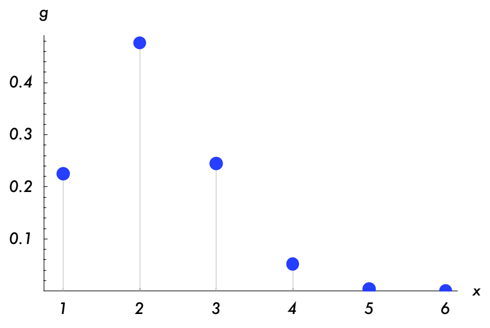

Figure 6.

■ 4. Extensions and Forthcoming Features

Thus far, this article has assumed that we are dealing with samples of iid variables. In this section, we take the major step of relaxing these assumptions. The generalisation to non-identical distributions is an enormously flexible and powerful capability. To do so, we require the new `Piecewise` functionality found in *Mathematica* 5.1 or later. In the examples that follow, we provide an illustration/preview of this new functionality as already implemented in the developmental version of mathStatica and which will be available in its next public release.

□ Non-Identical Parameters

Let X_i denote a continuous random variable with pdf $f(x; \lambda_i)$ and cdf $F(x; \lambda_i)$, such that $(X_1, X_2, \ldots, X_n)$ are independent but *not identical* variables due to differing parameters λ_i, for $i = 1, \ldots, n$. For example, consider an Exponential(λ) parent where identicality is relaxed by replacing λ with λ_i, for $i = 1, \ldots, n$. Thus:

In[25]:= $f = \frac{1}{\lambda_i} e^{-x/\lambda_i}$; domain[f] = {x, 0, ∞} && {λ_i > 0};

Then, the pdf of the minimum order statistic (in, say, a sample of size 4) is:

In[26]:= **OrderStat[1, f$_i$, 4]**

Out[26]= $$\frac{e^{-x\left(\frac{1}{\lambda_1}+\frac{1}{\lambda_2}+\frac{1}{\lambda_3}+\frac{1}{\lambda_4}\right)}\left(\lambda_2\lambda_3\lambda_4+\lambda_1\left(\lambda_3\lambda_4+\lambda_2\left(\lambda_3+\lambda_4\right)\right)\right)}{\lambda_1\lambda_2\lambda_3\lambda_4}$$

The pdf of the next largest order statistic, $X_{(2)}$, is substantially more complicated:

In[27]:= **OrderStat[2, f$_i$, 4]**

Out[27]= $$\frac{1}{\lambda_1\lambda_2\lambda_3\lambda_4}\Big(e^{-x\left(\frac{1}{\lambda_1}+\frac{1}{\lambda_2}+\frac{1}{\lambda_3}+\frac{1}{\lambda_4}\right)}\Big(\Big(-3+e^{\frac{x}{\lambda_2}}+e^{\frac{x}{\lambda_3}}+e^{\frac{x}{\lambda_4}}\Big)\lambda_2\lambda_3\lambda_4+\lambda_1\Big(\Big(-3+e^{\frac{x}{\lambda_1}}+e^{\frac{x}{\lambda_3}}+e^{\frac{x}{\lambda_4}}\Big)\lambda_3\lambda_4+\lambda_2\Big(\Big(-3+e^{\frac{x}{\lambda_1}}+e^{\frac{x}{\lambda_2}}+e^{\frac{x}{\lambda_3}}\Big)\lambda_3+\Big(-3+e^{\frac{x}{\lambda_1}}+e^{\frac{x}{\lambda_2}}+e^{\frac{x}{\lambda_4}}\Big)\lambda_4\Big)\Big)\Big)\Big)$$

□ Non-Identical Distributions

Next, let us suppose we have three completely different distributions defined over three different domains of support. In the following, $f(x)$ is the pdf of an Exponential(λ), $g(x)$ is the pdf of a standard Normal, and $h(x)$ is the pdf of a Uniform(-1, 1) random variable:

In[28]:= **f = $\frac{1}{\lambda}e^{-x/\lambda}$;   domain[f] = {x, 0, ∞} && {λ > 0};**

In[29]:= **g = $\frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}}$;   domain[g] = {x, -∞, ∞};**

In[30]:= **h = $\frac{1}{2}$;   domain[h] = {x, -1, 1};**

We can now solve completely general questions. For example, let us suppose we have a random sample of size $n = 20$. Of this sample, suppose that 10 values are drawn from the Normal, seven from the Exponential, and three from the Uniform. What is the pdf of the second smallest value from the sample, namely the second order statistic? Solving this problem would normally be enormously complicated, but the solution is now given simply by:

In[31]:= **OrderStat[2, {f, g, h}, {10, 7, 3}]**

The output can be viewed in the electronic version of this notebook.

This same technology provides a neat way to solve problems such as finding the pdf of min(X, Y, Z), when X, Y and Z have completely different distributions and different domains of support. For our example, if $X \sim$ Exponential(λ), $Y \sim$ Normal(0, 1) and $Z \sim$ Uniform(-1, 1), then the pdf of min(X, Y, Z) is simply the pdf of the first order statistic:

In[32]:=
```
sol = OrderStat[1, {f, g, h}]
```

Out[32]=

$$\begin{cases} \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}} & x \le -1 \\ -\frac{e^{-\frac{x^2}{2}}(-1+x)}{2\sqrt{2\pi}} + \frac{1}{4}\,\text{Erfc}\left[\frac{x}{\sqrt{2}}\right] & -1 < x \le 0 \\ \frac{e^{-\frac{1}{2}x\left(x+\frac{2}{\lambda}\right)}\left(-\sqrt{\frac{2}{\pi}}\,(-1+x)\,\lambda + e^{\frac{x^2}{2}}\,(1-x+\lambda)\,\text{Erfc}\left[\frac{x}{\sqrt{2}}\right]\right)}{4\lambda} & 0 < x < 1 \end{cases}$$

with domain of support:

In[33]:=
```
domain[sol] = {x, -∞, ∞} && {λ > 0};
```

Here is a plot of the pdf we have just derived:

In[34]:=
```
PlotDensity[sol /. λ → 1, {x, -4, 2}, PlotStyle → Hue[1]];
```

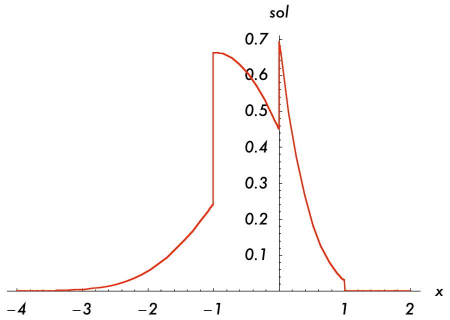

Figure 7.

We can easily 'check' our solution using Monte Carlo methods. Here are 100,000 pseudorandom drawings from each of the three distributions:

In[35]:=
```
<< Statistics`
```

In[36]:=
```
dataf = RandomArray[ExponentialDistribution[1 / λ] /. λ → 1, 100000];
datag = RandomArray[NormalDistribution[0, 1], 100000];
datah = RandomArray[UniformDistribution[-1, 1], 100000];
```

Next, we create 100,000 samples of size 3 containing one drawing from each of the three distributions, and then map the `Min` function across each sample, generating our 100,000 empirical drawings of the sample minimum:

In[39]:=
```
samplemin = Map[Min, Transpose[{dataf, datag, datah}]];
```

Figure 8 compares the empirical pdf (——) of the data we have just generated with the theoretical pdf (---) derived earlier:

```
In[40]:= FrequencyPlot[samplemin, {-4, 2, .05}, sol /. λ → 1];
```

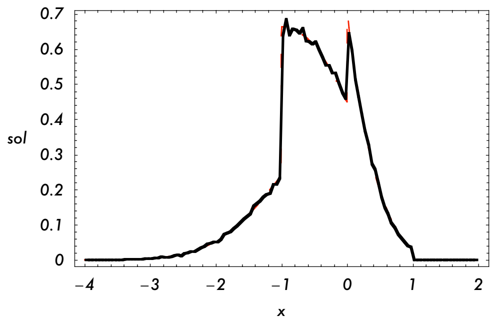

Figure 8.

□ Independence Relaxed, Identicality Maintained

The distribution of an order statistic is derived as a many-to-one transformation from the joint distribution of $(X_1, \ldots, X_n)$. Although there are differing ways in which the distribution can be found, perhaps the simplest method uses equivalence events. To illustrate, the event $X_{(n)} \le x$ is equivalent to $X_i \le x$, for all $i = 1, \ldots, n$, where x is arbitrarily chosen. Accordingly, the cdf of $X_{(n)}$ in terms of x is given by

$$P(X_{(n)} \le x) = P(X_1 \le x, \ldots, X_n \le x). \tag{1}$$

In the case of $X_{(n)}$ (and $X_{(1)}$ too), there is only one equivalent event. The number of equivalent events increases substantially as other inner order statistics are considered; however, in order to keep our discussion as simple as possible we will confine attention to $X_{(n)}$ from here on. This is not necessarily a case without interest, for the extremal $X_{(n)}$ is encountered in many practical contexts. Examples include $X_{(n)}$ as the measure of record high temperatures and record times in sports.

If the standard iid assumptions hold, then $(X_1, \ldots, X_n)$ is a collection of mutually independent random variables, and all are copies of the same parent X (with pdf $f(x)$ and cdf $F(x)$), leading to considerable simplification in the right-hand side of (1):

$$P(X_{(n)} \le x) = P(X \le x)^n. \tag{2}$$

If, for example, X is continuous, then the pdf of $X_{(n)}$ is obtained by differentiation of (2) with respect to x, yielding $g(x) = n\,F(x)^{n-1}\,f(x)$.

Just as identicality can be relaxed in many ways, so too can independence. To introduce a dependence structure, we may begin by rewriting (1) in its copula form

$$\begin{aligned} P(X_{(n)} \le x) &= C(F_1(x), \ldots, F_n(x)) \\ &= C(F(x), \ldots, F(x)), \end{aligned} \tag{3}$$

where the second line recognises that $X_1, \ldots, X_n$ are copies of a common parent X. The n-copula $C : [0, 1]^n \to [0, 1]$ is the function that represents the dependence structure of $X_1, \ldots, X_n$. For example, the special case $C(u_1, \ldots, u_n) = u_1 \ldots u_n$ corresponds to mutual independence amongst $X_1, \ldots, X_n$. The representation (3) is due to Sklar [10] and is unique provided X is continuous.

Tractable results can be obtained by assuming an Archimedean dependence structure for C; for details of these copulas see [11, Chapter 4]. Let $\varphi : [0, 1] \to [0, \infty]$ denote the strict generator function associated with C; it is differentiable such that $\varphi'(t) = \partial\varphi(t) / \partial t < 0$ on $0 < t < 1$, and its inverse φ^{-1} must be completely monotonic if $n \geq 3$. For example, the generator associated with the independence copula $\varphi(t) = -\log t$ is strict, and its inverse $\varphi^{-1}(t) = \exp(-t)$ is completely monotonic. Then, the key property of the generator is that

$$\begin{aligned} \varphi(G(x)) &= \varphi(F(x)) + \cdots + \varphi(F(x)) \\ &= n\,\varphi(F(x)), \end{aligned} \tag{4}$$

where $G(x) = P(X_{(n)} \leq x)$ is the cdf of $X_{(n)}$. Differentiating both sides of (4) with respect to x and rearranging yields the pdf of $X_{(n)}$:

$$g(x) = n \frac{\varphi'(F(x))}{\varphi'(G(x))} f(x), \tag{5}$$

where the denominator would be computed as per $\varphi'(G(x)) = \varphi'(\varphi^{-1}(n\,\varphi(F(x))))$. The resemblance in the structure of the pdf (5) to the pdf in the iid case, $n\,F(x)^{n-1} f(x)$, is striking.

To illustrate, let $X \sim \text{Exponential}(\lambda)$ with pdf $f(x)$:

```
In[41]:=  f = (1/λ) ⅇ^(-x/λ) ; domain[f] = {x, 0, ∞} && {λ > 0};
```

and cdf $F(x) = P(X \leq x)$:

```
In[42]:=  F = Prob[x, f]

Out[42]=  1 - ⅇ^(-x/λ)
```

and summarise our assumptions:

```
In[43]:=  assum = {n ∈ Integers, n > 0, x > 0, λ > 0, θ ≥ 1};
```

Enter the details for a particular case considered by Ballerini [12], namely, that of the Gumbel-Hougaard family of n-copulas with generator $\varphi(t) = (-\log t)^\theta$, with dependence parameter $\theta \geq 1$:

```
In[44]:=  φ[t_] = (-Log[t])^θ;    φi[t_] = Exp[-t^(1/θ)];    φd[t_] = D[φ[t], t];
```

Then, from (5), the pdf of $X_{(n)}$ is given by

```
In[45]:= g = FullSimplify[n φd[F] / φd[φi[n * φ[F]]] f, assum]
```

$$\text{Out[45]=}\quad \frac{\left(1 - e^{-\frac{x}{\lambda}}\right)^{n^{\frac{1}{\theta}}} n^{\frac{1}{\theta}}}{\left(-1 + e^{x/\lambda}\right) \lambda}$$

with domain of support:

```
In[46]:= domain[g] = {x, 0, ∞} && assum;
```

Setting $\theta = 1$ corresponds to the iid case; notice too that replacing n in the iid pdf with $n^{1/\theta}$ yields the pdf of $X_{(n)}$. This, for example, means that many algebraic results on the properties of $X_{(n)}$ can be found simply by replacing n with $n^{1/\theta}$ in the appropriate iid formula. In Figure 9, the solid line denotes the pdf when $\theta = 1$ (the iid case), while the dashed line denotes the pdf when $\theta = 2$.

```
In[47]:= PlotDensity[g /. {θ → {1, 2}, λ → 1, n → 10}];
```

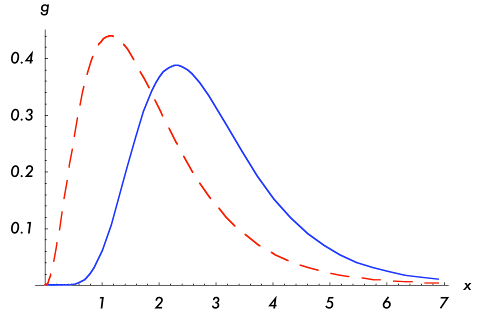

Figure 9.

■ References

[1] mathStatica, Sydney, Australia: MATHSTATICA Pty. Ltd., (Apr 2005) www.mathStatica.com.

[2] C. Rose and M. D. Smith, *Mathematical Statistics with Mathematica*, Springer Texts in Statistics, New York: Springer-Verlag, 2002.

[3] H. A. David and H. N. Nagaraja, *Order Statistics*, 3rd ed., New York: John Wiley & Sons, 2003.

[4] N. Balakrishnan and C. R. Rao, eds., *Order Statistics: Theory and Methods*, Handbook of Statistics, Vol. 16, Amsterdam: Elsevier, 1998.

[5] ———, *Order Statistics: Applications*, Handbook of Statistics, Vol. 17, Amsterdam: Elsevier, 1998.

[6] B. C. Arnold, N. Balakrishnan, and H. N. Nagaraja, *A First Course in Order Statistics*, New York: John Wiley & Sons, 1992.

[7] H. N. Nagaraja, "Order Statistics from Discrete Distributions," *Statistics*, **23**, 1992 pp. 189–216.

[8] D. L. Evans, L. M. Leemis, and J. H. Drew, "The Distribution of Order Statistics for Discrete Random Variables with Applications for Bootstrapping," *INFORMS Journal on Computing* (*JOC*), forthcoming, 2005.

[9] P. R. Tadikamalla and N. Balakrishnan, "Computer Simulation of Order Statistics," in Balakrishnan and Rao, pp. 65–72.

[10] A. Sklar, "Random Variables, Joint Distributions, and Copulas," *Kybernetika*, **9**, 1973 pp. 449–460.

[11] R. B. Nelsen, *An Introduction to Copulas,* Lecture Notes in Statistics, New York: Springer-Verlag, 1999.

[12] R. Ballerini, "A Dependent F^{α}-Scheme," *Statistics & Probability Letters*, **21**(1), 1994 pp. 21–25.

About the Authors

Colin Rose is director of the Theoretical Research Institute, Sydney. In 1998, he was a visiting scholar at Wolfram Research. His recent research interests relate to exact computational methods in mathematical statistics, and in particular, with application to the mathStatica project.

Murray D. Smith is an associate professor in the Discipline of Econometrics and Business Statistics at the University of Sydney. His recent research includes applications of computational algebra in statistics and the uses of copulas in modelling dependence structures in econometrics.

Colin Rose
Theoretical Research Institute
Sydney NSW 2023
Australia
colin@tri.org.au

Murray D. Smith
Econometrics and Business Statistics
The University of Sydney
Sydney NSW 2006
Australia
Murray.Smith@econ.usyd.edu.au