since dx/dz is positive and given by the term $(b/a^b)z^{b-1}$.[14]

The mean of the Weibull distribution is $\mu_W = a\Gamma(1+(1/b))$, the variance is $\sigma_W^2 = a^2[\Gamma(1+(2/b)) - (\Gamma(1+(1/b)))^2]$, whereas the mode exists and is given by $\tilde{\mu}_W = a((b-1)/b)^{(1/b)}$ when $b > 1$.

4.2.4. *Inverted Gamma and Inverted Wishart Distributions*

A random variable $z > 0$ has an *inverted gamma distribution* with shape parameter $a > 0$ and scale parameter $b > 0$, denoted by $z \sim IG(a,b)$, if and only if its pdf is given by

$$p_{IG}(z|a,b) = \frac{2}{\Gamma(a)b^a} z^{-(2a+1)} \exp\left(\frac{-1}{bz^2}\right). \tag{4.10}$$

This pdf has a unique mode at $\tilde{\mu}_{IG} = (2/(b(2a+1)))^{1/2}$; cf. Zellner (1971).[15] Moreover, the statement $z \sim IG(a,b)$ is equivalent to $z = 1/\sqrt{x}$ where $x \sim G(a,b)$.

The inverted gamma distribution is an often used prior for a standard deviation parameter. Letting $\sigma = z$, $a = q/2$, and $b = 2/qs^2$, we get

$$p_{IG}(\sigma|s,q) = \frac{2}{\Gamma(q/2)}\left(\frac{qs^2}{2}\right)^{q/2} \sigma^{-(q+1)} \exp\left(\frac{-qs^2}{2\sigma^2}\right), \tag{4.11}$$

where $s, q > 0$. The parameter q is an integer (degrees of freedom) while s is a location parameter. This pdf has a unique mode at $\tilde{\mu}_{IG} = s(q/(q+1))^{1/2}$. Hence, the mode is below s for finite q and converges to s when $q \to \infty$.

The moments of this distribution exists when q is sufficiently large. For example, if $q \geq 2$, then the mean is

$$\mu_{IG} = \frac{\Gamma((q-1)/2)}{\Gamma(q/2))}\left(\frac{q}{2}\right)^{1/2} s,$$

while if $q \geq 3$ then the variance is given by

$$\sigma_{IG}^2 = \frac{q}{q-2}s^2 - \mu_{IG}^2.$$

Hence, both the mean and the variance are decreasing functions of q; see Zellner (1971) for details.

Moreover, if $q \geq 4$ then the third moment also exists. The exact expression can be found in Zellner (1971, eq. (A.48)), but since that expression is very messy an alternative skewness measure may be of interest. One such simpler alternative is the *Pearson measure of skewness*, defined as the mean minus the mode and divided by the standard deviation. For the inverted gamma we here find that

$$S_{P,IG} = \frac{\mu_{IG} - \tilde{\mu}_{IG}}{\sigma_{IG}} = \frac{\dfrac{\Gamma\big((q-1)/2\big)}{\Gamma(q/2)}\left(\dfrac{q}{2}\right)^{1/2} - \left(\dfrac{q}{q+1}\right)^{1/2}}{\left[\dfrac{q}{q-2} - \left(\dfrac{\Gamma\big((q-1)/2\big)}{\Gamma(q/2)}\right)^2 \dfrac{q}{2}\right]^{1/2}}, \qquad q \geq 3.$$

This expression is positive for finite q and the inverted gamma distribution is therefore right-skewed. As q gets large, the skewness measure $S_{P,IG} \to 0$. Both the numerator and the denominator are decreasing in q and for $q > 5$ the ratio is decreasing.[16]

---

[14] With $p_X(x) = \exp(-x)$ and $z = ax^{1/b}$ we find from equation (4.3) that $f^{-1}(z) = (z/a)^b$. Moreover, $f'(x) = (a/b)x^{(1-b)/b}$ so that $f'(f^{-1}(z)) = (a^b/b)z^{1-b}$. By multiplying terms we obtain the density function in (4.9). Notice that $dx/dz = 1/f'(f^{-1}(z))$, the Jacobian of the transformation z into x.

[15] Bauwens, Lubrano, and Richard (1999) refer to the inverted gamma distribution as the inverted gamma-1 distribution. The inverted gamma-2 distribution is then defined for a variable $x = z^2$, where z follows an inverted gamma-1 distribution.

[16] Skewness is defined as the third standardized central moment, i.e., the third central moment divided by the standard deviation to the power of 3. There is no guarantee that the sign of this measure always corresponds to the sign of the Pearson measure.

A few examples of the inverted gamma distribution have been plotted in the upper right panel of Figure 1. The location parameter is for simplicity kept fixed at 0.1, while the number of degrees of freedom are given by $q = (1, 2, 5, 10)$. It can be seen that the height of the density increases as q becomes larger.[17] Moreover, the variance is smaller while skewness appears to be lower for $q = 10$ than for $q = 5$. The latter is consistent with the results for the Pearson measure of skewness, $S_{P,IG}$.

Another parameterization of the inverted gamma distribution is used in the software developed by Adolfson et al. (2007b). Letting $a = d/2$ and $b = 2/c$, the pdf in (4.10) can be written as:

$$p_{IG}(z|c,d) = \frac{2}{\Gamma(d/2)}\left(\frac{c}{2}\right)^{d/2} z^{-(d+1)} \exp\left(\frac{-c}{2z^2}\right).$$

The mode of this parameterization is found by setting $\tilde{\mu}_{IG} = (c/(d+1))^{1/2}$. With $c = qs^2$ and $d = q$ this parameterization is equal to that in equation (4.11) with $z = \sigma$.

A multivariate extension of the inverted gamma distribution is given by the *inverted Wishart distribution*. Specifically, when a $p \times p$ positive definite matrix Ω is inverted Wishart, denoted by $\Omega \sim IW_p(A, \nu)$, its density is given by

$$p(\Omega) = \frac{|A|^{\nu/2}}{2^{\nu p/2}\pi^{p(p-1)/4}\Gamma_p(\nu)}|\Omega|^{-(\nu+p+1)/2} \exp\left(-\frac{1}{2}\mathrm{tr}\left[\Omega^{-1}A\right]\right), \tag{4.12}$$

where $\Gamma_b(a) = \prod_{i=1}^{b}\Gamma([a-i+1]/2)$ for positive integers a and b, with $a \geq b$, and $\Gamma(\cdot)$ being the gamma function in (4.4). The parameters of this distribution are given by the positive definite location matrix A and the degrees of freedom parameter $\nu \geq p$. The mode of the inverted Wishart is given by $(1/(p+\nu+1))A$, while the mean exists if $\nu \geq p+2$ and is then given by $E[\Omega] = (1/(\nu-p-1))A$; see, e.g., Zellner (1971, Appendix B.4) and Bauwens et al. (1999, Appendix A) for details.

Suppose for simplicity that $p = 2$ and let us partition Ω and A conformably

$$\Omega = \begin{bmatrix}\Omega_{11} & \Omega_{12}\\ \Omega_{12} & \Omega_{22}\end{bmatrix}, \qquad A = \begin{bmatrix}A_{11} & A_{12}\\ A_{12} & A_{22}\end{bmatrix}.$$

It now follows from, e.g., Bauwens et al. (1999, Theorem A.17) that:

(1) Ω_{11} is independent of Ω_{12}/Ω_{11} and of $\Omega_{22\cdot1} = \Omega_{22} - \Omega_{12}^2/\Omega_{11}$;
(2) $\Omega_{11} \sim IW_1(A_{11}, \nu-1)$;
(3) $\Omega_{12}/\Omega_{11}|\Omega_{22\cdot1} \sim N(A_{12}/A_{11}, \Omega_{22\cdot1}/A_{11})$, where $N(\mu, \sigma^2)$ denotes the univariate normal distribution with mean μ and variance σ^2 (see, e.g., Section 4.2.6 for details); and
(4) $\Omega_{22\cdot1} \sim IW_1(A_{22\cdot1}, \nu)$, where $A_{22\cdot1} = A_{22} - A_{12}^2/A_{11}$.

From these results it is straightforward to deduce that the multivariate random matrix Ω may be represented by three independent univariate random variables. Specifically, let

$$\sigma_1 \sim IG(s_1, \nu-1), \quad \sigma_2 \sim IG(s_2, \nu), \quad \text{and } \rho \sim N(0,1). \tag{4.13}$$

With $\Omega_{11} = \sigma_1^2$ and $A_{11} = (\nu-1)s_1^2$, it can be shown that $\Omega_{11} \sim IW_1(A_{11}, \nu-1)$ by evaluating the inverted gamma density at Ω_{11}, A_{11}, and multiplying this density with the inverse of the derivative of Ω_{11} with respect to σ_1, i.e., by $(1/2)\Omega_{11}^{-1/2}$; recall equation (4.3) in Section 4.2.1. Furthermore, letting $\Omega_{22\cdot1} = \sigma_2^2$ and $A_{22\cdot1} = \nu s_2^2$ we likewise find that $\Omega_{22\cdot1} \sim IW_1(A_{22\cdot1}, \nu)$. Trivially, we also know that $\Omega_{12}/\Omega_{11} = A_{12}/A_{11} + \sqrt{\Omega_{22\cdot1}/A_{11}}\rho$ implies that $\Omega_{12}/\Omega_{11}|\Omega_{22\cdot1} \sim N(A_{12}/A_{11}, \Omega_{22\cdot1}/A_{11})$.

[17] This can also be seen if we let $\sigma = \tilde{\mu}_{IG}$ in equation (4.11).

Together, these results therefore ensure that $\Omega \sim IW_2(A, \nu)$, where

$$\Omega_{11} = \sigma_1^2,$$

$$\Omega_{12} = \Omega_{11} \left[\frac{A_{12}}{A_{11}} + \sqrt{\frac{\sigma_2^2}{A_{11}}} \rho \right],$$

$$\Omega_{22} = \sigma_2^2 + \frac{\Omega_{12}^2}{\Omega_{11}}.$$

It may also be noted that one can derive an inverted Wishart distribution for the general $p \times p$ case based on p univariate inverted gamma random variables and $p(p-1)/2$ univariate standard normal random variables, and where all univariate variables are independent. The precise transformations needed to obtain Ω from these univariate variables can be determined by using Theorem A.17 from Bauwens et al. (1999) in a sequential manner.

4.2.5. *Beta, Snedecor (F), and Dirichlet Distributions*

A random variable $c < x < d$ has a *beta distribution* with parameters $a > 0$, $b > 0$, $c \in \mathbb{R}$ and $d > c$, denoted by $x \sim B(a, b, c, d)$ if and only if its pdf is given by

$$p_B(x|a, b, c, d) = \frac{1}{(d-c)\beta(a, b)} \left(\frac{x-c}{d-c} \right)^{a-1} \left(\frac{d-x}{d-c} \right)^{b-1}. \tag{4.14}$$

The standardized beta distribution can directly be determined from (4.14) by defining the random variable $z = (x-c)/(d-c)$. Hence, $0 < z < 1$ has a beta distribution with parameters $a > 0$ and $b > 0$, denoted by $z \sim B(a, b)$ if and only if its pdf is given by

$$p_{SB}(z|a, b) = \frac{1}{\beta(a, b)} z^{a-1}(1-z)^{b-1}. \tag{4.15}$$

For $a, b > 1$, the mode of (4.15) is given by $\tilde{\mu}_{SB} = (a-1)/(a+b-2)$. Zellner (1971) provides general expressions for the moments of the beta pdf in (4.15). For example, the mean of the standardized beta is $\mu_{SB} = a/(a+b)$, while the variance is $\sigma_{SB}^2 = ab/((a+b)^2(a+b+1))$.

The a and b parameters of the beta distribution can be expressed as functions of the mean and the variance. Some algebra later we find that

$$\begin{aligned} a &= \frac{\mu_{SB}}{\sigma_{SB}^2} \left[\mu_{SB} (1 - \mu_{SB}) - \sigma_{SB}^2 \right], \\ b &= \frac{(1 - \mu_{SB})}{\mu_{SB}} a. \end{aligned} \tag{4.16}$$

From these expressions we see that a and b are defined from μ_{SB} and σ_{SB}^2 when $\mu_{SB}(1-\mu_{SB}) > \sigma_{SB}^2 > 0$ with $0 < \mu_{SB} < 1$.

Letting μ_B and σ_B^2 be the mean and the variance of $x \sim B(a, b, c, d)$, it is straightforward to show that:

$$\begin{aligned} \mu_B &= c + (d-c)\mu_{SB}, \\ \sigma_B^2 &= (d-c)^2 \sigma_{SB}^2. \end{aligned} \tag{4.17}$$

This means that we can express a and b as functions of μ_B, σ_B, c, and d:

$$\begin{aligned} a &= \frac{(\mu_B - c)}{(d-c)\sigma_B^2} \left[(\mu_B - c)(d - \mu_B) - \sigma_B^2 \right], \\ b &= \frac{(d - \mu_B)}{(\mu_B - c)} a. \end{aligned} \tag{4.18}$$

The conditions that $a > 0$ and $b > 0$ means that $c < \mu_B < d$, while $(\mu_B - c)(d - \mu_B) > \sigma_B^2$. The mode still exists when $a, b > 1$ and is in that case given by $\tilde{\mu}_B = c + (d-c)\tilde{\mu}_{SB} = c + (d-c)(a-1)/(a+b-2)$.