distributions.

Consider a non-negative random variable w representing the precision. Let δ be a non-negative parameter representing an arbitrarily imposed lower-bound support for the Bayesian prior distribution of the precision w. Assume that the Bayesian prior distribution of the precision w is a truncated-gamma probability density function (with truncation parameter δ) of the form

$$\varphi_0^\delta(w) \;=\; w^{a_0-1}\, e^{-b_0 w} \Big/ \int_\delta^\infty w^{a_0-1}\, e^{-b_0 w}\, dw \qquad \textbf{(27)}$$

for $w \geq \delta$, while $\varphi_0^\delta(w)=0$ for $w<\delta$. When choosing δ to be positive, the model is effectively eliminating a priori all variances above $1/\delta$. The technical reason for declaring impermissible worlds of unboundedly high variance is to make the integral defining the moment generating function of x converge to a finite value. An economic rationale presumably has to do with the difficulty of envisioning the unbounded loss function arising from unlimited variability in growth rates. The implicit message is that the appropriate value of δ is far from being known *a priori.*

The three non-negative parameters $\delta,\ a_0,\ b_0$ of the truncated gamma distribution (27) represent prior beliefs about the precision. In the limit as $\delta \to 0^+$, the mean of the gamma prior approaches a_0/b_0, while the variance of the gamma prior approaches a_0/b_0^2. Thus, at least for small δ, the prior mean and prior variance of the precision can be assigned any values just by judiciously selecting a_0 and b_0. Classical statistical analysis is exactly isomorphic to the limiting case of a diffuse prior: $\delta \to 0^+$, $a_0 \to 0^+$, $b_0 \to 0^+$. Therefore, the analysis presented here can be viewed as paralleling the classical specification very closely, except that it is slightly more general by allowing positive parameter values other than the limiting value 0^+.

Let $\varphi_n^\delta(w)$ be the *posterior* distribution of the precision w at a time just after observing the n independent realizations $g_1,\ \ldots\ ,\ g_n$. When $\delta=0$, it is well known (see any of the references cited in footnote 4) that the normal-gamma distribution constitutes a conjugate family of priors. When $\delta>0$, we have the same conjugate family of priors, except that w is subject to a lower-bound constraint. Therefore, the posterior is in the same form as the prior, and subject to the same bounding constraint. The modification of a basic conjugate-prior result in the Bayesian statistical