Cranking through Bayesian Calculus, Part III

Ed Herbst

March 14, 2025

Last time we talked about the regression model,

$$y_t = x'_t \beta + u_t, \quad u_t \sim iidN(0, \sigma^2). \tag{1}$$

We focused on the "new" parameter, σ^2 , and talked about how to construct a prior for it. We then described a few different parameterizations of the prior. Finally, we derived the posterior for σ^2 , under the likelihood defined in (1) with the restriction that $\beta = 0$. Today we're going to focus on jointly estimating the two parameters of our regression model: (β, σ^2) . We'll refer to this vector of parameters as θ . Also, let's make it explicit that β is a $k \times 1$ vector; that is, there are k explanatory variables in our regression.

First, we'll introduce a new distribution defined jointly over (β, σ^2) , which we'll use as our prior distribution. To do so, we'll factorize the prior as follows:

$$p(\beta, \sigma^2) = p(\beta|\sigma^2)p(\sigma^2).$$

 (β, σ^2) follows a normal inverse gamma distribution with parameters $(\nu_0, s_0^2, \mu_0, V_0)$ if σ^2 follows an inverse gamma distribution with parameters (ν_0, s_0^2) and β conditional on σ^2 follows a normal distribution with mean μ_0 and variance $\sigma^2 V_0$. Use the above factorization, the joint density of the distribution can be written as:

$$p(\beta, \sigma^2) = (2\pi)^{-k/2} [\det(\sigma^2 V_0)]^{-1/2} \exp\left\{-\frac{1}{2}(\beta - \mu_0)'[\sigma^2 V_0]^{-1}(\beta - \mu_0)\right\} \\ \times \frac{\nu_0/2}{\Gamma(\nu_0/2)} s_0^{\nu_0} (1/\sigma^2)^{\nu_0/2+1} \exp\left\{-\nu_0 s_0^2/(2\sigma^2)\right\}$$

Problem 1: Write two scripts is **python** or **R** that: (1) generate draws from this distribution and (2) evaluate the log pdf of the distribution, given some values.

#test

Note that in our formulation, we've constructed β conditional on σ^2 . It's also interesting to examine the *marginal* distribution of β ,

$$p(\beta) = \int p(\beta, \sigma^2) d\sigma^2.$$

Problem 2: Derive the marginal distribution of β by integrating out σ^2 . Validate your derivation by comparing a density estimated from the simulations in Problem 1 to the analytic formulation. What is the name of this distribution?

#test

The normal inverse gamma prior is convenient because it's *conjugate* for the normal regression model in (1). This means that the posterior distribution of the parameters is also a normal inverse gamma distribution.

Problem 3: Derive the posterior distribution for the model in (1). Use a normal inverse gamma prior with parameters $(\nu_0, s_0^2, \mu_0, V_0)$. For notation, let $X = [x_1, \ldots, x_T]'$ and $Y = [y_1, \ldots, y_T]'$.

#test

Let's run a Bayesian regression! The data in the table below come from T. Haavelmo, "Methods of Measuring the Marginal Propensity to Consume," J. Am. Statist. Assoc, 42, p. 88 (1947). Using (1) to relate income, y_t , to a constant and "autonomous" investment, the independent variable. The coefficient associated with investment is termed the investment multiplier.

Problem 4: Pick a parameterization of the normal inverse gamma distribution that is not very informative; that is, it doesn't impose strong beliefs about the plausible values one the coefficients. Let's center the prior for β at $\mu_0 = 0$, and for the inverse gamma portion set $s_0^2 = 600$. What should you do with ν_0 and V_0 ? Construct the posterior distribution for β and σ^2 . What is the posterior mean of β_2 , the coefficient associated with investment? What happens when you increase the "strength" of the prior, by increasing ν_0 or decreasing V_0 ?

Year	Income	Investment
1922	433	39
1923	483	60
1924	479	42
1925	486	52
1926	494	47
1927	498	51
1928	511	45
1929	534	60
1930	478	39
1931	440	41
1932	372	22
1933	381	17
1934	419	27
1935	449	33
1936	511	48
1937	520	51
1938	477	33
1939	517	46
1940	548	54
1941	629	100

Table 1: Haavelmo's Data on Income and Investment