

Exercise on Bayesian Basics

© Christopher Sims, 1997. This material may be reproduced for educational and research purposes so long as this copyright notice is included in the copies, the copies are complete and unaltered, and the copies are not sold.

1. Using the law of iterated expectations,

$$E[E[Y|X]] = E[Y],$$

and the fact that

$$P[Y < a] = 1 \text{ implies } E[Y] < a,$$

Show that if $\hat{\mathbf{b}}(Y)$ is a Bayesian posterior mean for \mathbf{b} under some prior distribution on \mathbf{b} , then the classical bias of $\hat{\mathbf{b}}(Y)$ as an estimator for \mathbf{b} , i.e.

$$E[\hat{\mathbf{b}}(Y)|\mathbf{b}] - \mathbf{b},$$

cannot be of the same sign for all values of \mathbf{b} .

2. Suppose $S(X)$ is a Bayesian 95% probability interval for \mathbf{b} under some prior, i.e.

$$P[\mathbf{b} \text{ in } S(X)|X] = .95, \text{ all } X.$$

Then there is a lower bound on the prior probability of the set of \mathbf{b} 's for which the classical coverage probability for $S(X)$,

$$P[\mathbf{b} \text{ in } S(X)|\mathbf{b}],$$

is less than .90. Find the bound.

3. Suppose x and y are jointly $N(0, I)$, and that $r = \sqrt{x^2 + y^2}$, $\mathbf{q} = \arctan(x/y)$ (i.e., polar coordinates). Find $p(r|\mathbf{q} = 0)$ and contrast this to $p(y|x = 0)$. Explain in words why, despite the fact that the condition $x = 0$ puts us in same set of (x, y) values as the restriction $\mathbf{q} = 0$, the conditional p.d.f.'s are different on that set.
4. Consider the model

$$y(t) = c + \mathbf{r}y(t-1) + \mathbf{e}(t), \tag{1}$$

i.e. a univariate AR model with a constant term. We assume $\mathbf{e}(t)$ independent of $y(s)$, all $s < t$ and i.i.d. $N(0,1)$ across t .

- i) Find the marginal distribution of $y(1)$ when $c = 0$, $\mathbf{r} = .999$.
- ii) Generate five random samples from the joint distribution of $\{y(1), \dots, y(100)\}$ under the $c = 0$, $\mathbf{r} = .999$ assumption. For each random sample, find the maximum likelihood estimates of c and \mathbf{r} both using the unconditional likelihood (using a marginal p.d.f. for

$y(1)$ and recognizing the dependence of the p.d.f. of $y(1)$ on the parameters) and using the likelihood conditional on $y(1)$. Display a table of these 10 ML estimates of α and \mathbf{r} .

- iii) For each of your random samples, and for each of the two ML estimates from each sample, produce a plot that displays the time series of values of y and $E[y(s)|y(1), c, \mathbf{r}]$ as a function of $s, s=1, \dots, 100$, with the ML estimates substituted in for c and \mathbf{r} . Also display on the plot $E[y(s)|c, \mathbf{r}]$ and a two-standard-deviation band about this straight line, again substituting ML estimates for true values of α, \mathbf{r} .
- iv) Discuss whether your results show a pattern of “unreasonably” large and persistent predictable changes at the beginning or end of the sample and whether one of the two likelihoods is better behaved than the other in this respect. [You can’t really answer this without drawing on prior beliefs about what kinds of models and data are likely to occur. So it is implicit here that we are interested in whether ML estimates that behave like these might be problematic for use with economic time series data.]