

Nonparametric Applications of Bayesian Inference

Gary CHAMBERLAIN

Department of Economics, Harvard University, Cambridge, MA 02138 (gary_chamberlain@harvard.edu)

Guido W. IMBENS

Department of Economics and Department of Agricultural and Resource Economics, University of California, Berkeley, CA 94720 (imbens@econ.berkeley.edu)

This article evaluates the usefulness of a nonparametric approach to Bayesian inference by presenting two applications. Our first application considers an educational choice problem. We focus on obtaining a predictive distribution for earnings corresponding to various levels of schooling. This predictive distribution incorporates the parameter uncertainty, so that it is relevant for decision making under uncertainty in the expected utility framework of microeconomics. The second application is to quantile regression. Our point here is to examine the potential of the nonparametric framework to provide inferences without relying on asymptotic approximations. Unlike in the first application, the standard asymptotic normal approximation turns out not to be a good guide.

KEY WORDS: Bayesian inference; Dirichlet distributions; Nonparametric models; Semiparametric models.

This article evaluates in the context of two applications the usefulness of a nonparametric approach to Bayesian inference. The basic approach is due to Ferguson (1973, 1974) and Rubin (1981). It has three key features. First, it has the basic benefits of Bayesian inference in providing a well-defined posterior distribution that is an important ingredient in many decision problems. Second, it has some of the advantages of semiparametric models used in frequentist analyses by not relying on a tightly parameterized likelihood function, based, for example, on a normal distribution. Third, it avoids pitfalls arising in Bayesian analyses from using high-dimensional parameter spaces with flat or other conventional prior distributions by using a prior distribution that arguably reflects lack of prior knowledge. These three features are illustrated in the two applications.

Our first application considers an educational choice problem. Specifically, we look at an individual's decision on the level of schooling when the individual is uncertain about the return to schooling. Following Angrist and Krueger (1991) we allow for endogeneity of the schooling measure by using a quarter of birth dummy as an instrumental variable. A standard parametric model would require distributional assumptions on the joint distribution of earnings and schooling given the instrument. On the other hand, standard instrumental variables methods that do not require such assumptions do not lead to the predictive earnings distributions required for the educational choice problem. The Bayesian nonparametric approach discussed in this article allows us to obtain a predictive distribution for earnings corresponding to various levels of schooling that incorporates the parameter uncertainty, so that it is relevant for decision making under uncertainty in the expected utility framework of microeconomics. At the same time in this application this approach avoids strong distributional assumptions without introducing strong sensitivity to the prior distribution.

The second application is to quantile regression. Our point here is to examine the potential of the nonparametric Bayesian

framework to provide inferences without making asymptotic approximations. Unlike in the first application, in this application the standard asymptotic normal distribution turns out to be a poor approximation to the sampling distribution of the estimator in some cases. If the standard normal distribution provides a good approximation to the finite sample distribution, posterior intervals obtained through the Bayesian nonparametric approach discussed in this article are close to confidence intervals. When the large sample normal approximation fails to provide a good approximation to the finite sample distribution, the interpretation of our posterior distribution is not affected.

1. DIRICHLET PRIOR DISTRIBUTIONS

Here we present a concise review of the basic theory, extended to allow for parameters defined by moment restrictions, that is sufficient to follow the applications. For more details, see the work of Ferguson (1973, 1974), Rubin (1981), Chamberlain and Imbens (1995), and Hirano (2002). There is a family of probability distributions $\{P_\theta; \theta \in \Theta\}$, and we observe $\{Z_i\}_{i=1}^n$, where the random variables Z_i are independently and identically distributed according to P_θ for some unknown value of θ in the parameter space Θ . To simplify notation, let Z denote a random variable that is distributed according to P_θ . We assume that the distributions P_θ have a common, finite support, $P_\theta(Z = a_j) = \theta_j$ ($j = 1, \dots, J$), where θ_j denotes the j th component of θ , and we take Θ to be the unit simplex in \mathcal{R}^J . Because J can be arbitrarily large and our data are measured with finite precision, the finite support assumption is arguably not restrictive. In fact,

Ferguson's (1973) discussion does not rely on discreteness. See also Hirano (2002).

Typically we are interested in some function of θ rather than elements of θ itself: $\beta = g(\theta)$, where the function $g(\cdot)$ may depend on the points of support $\{a_j\}_{j=1}^J$. For example, we consider cases where $g(\cdot)$ is defined implicitly through moment restrictions,

$$E_\theta \psi(Z, \beta) = \sum_{j=1}^J \psi(a_j, \beta) \cdot \theta_j = 0, \quad (1)$$

where ψ is a given function with dimension equal to that of β , and there is a unique solution for all $\theta \in \Theta$. Although it may appear to be restrictive to limit this discussion to the case with the dimension of β equal to that of ψ , one can apply the same approach to overidentified gmm models where the dimension of ψ is higher than the dimension of β by augmenting the parameter vector and the moment functions. Specifically, let $\gamma = (\beta_0, \beta_1, \Gamma_0, \Gamma_1, \Delta)$, and let

$$\tilde{\psi}(Z, \gamma) = \begin{pmatrix} \Gamma_0' \psi(Z, \beta_0) \\ \text{vec}(\partial \psi(Z, \beta_0) / \partial \beta' - \Gamma_0) \\ \text{vec}(\psi(Z, \beta_0) \psi(Z, \beta_0)' - \Delta) \\ \Gamma_1' \Delta^{-1} \psi(Z, \beta_1) \\ \text{vec}(\partial \psi(Z, \beta_1) / \partial \beta' - \Gamma_1) \end{pmatrix}.$$

Then the solution to $\sum_{i=1}^n \tilde{\psi}(z_i, \gamma) = 0$ gives the standard optimal two-step generalized method of moments estimator for β , motivating our interest in the posterior distribution for the parameter defined as the solution to $E[\tilde{\psi}(Z, \gamma)] = 0$. Our proposed procedure will give a posterior distribution for this parameter given the data.

A second example concerns cases where $g(\cdot)$ is defined as the solution to an optimization problem,

$$\beta = \arg \min_t E_\theta[\rho(Z, t)] = \arg \min_t \sum_{j=1}^J \rho(a_j, t) \cdot \theta_j, \quad (2)$$

where ρ is a given scalar-valued function and there is a unique solution for all $\theta \in \Theta$. In both cases we obtain draws from the posterior distribution of β by first drawing from the posterior distribution of θ and then solving (1) or (2).

We limit ourselves to prior distributions in the Dirichlet family with density

$$p(\theta) \propto \prod_{j=1}^J \theta_j^{b_j-1} \quad \text{for } \theta \in \Theta \quad (b_j > 0), \quad (3)$$

which, with J free parameters b_j , is fairly flexible. Similar to the way the Beta distribution is the conjugate prior distribution for the parameter of a binomial distribution, the Dirichlet distribution is the conjugate prior distribution for the parameters of a multinomial distribution. Let $d = \{z_i\}_{i=1}^n$ denote the data, that is, the observed values of the Z_i , and let $n_j = \sum_{i=1}^n 1(z_i = a_j)$ be the number of sample observations equal to a_j . The posterior density is proportional to the product of the prior density and the likelihood function,

$$p_n(\theta | d) \propto \prod_{j=1}^J \theta_j^{n_j+b_j-1},$$

and thus also Dirichlet with parameters $n_j + b_j$, $j = 1, \dots, J$. Within this family of Dirichlet prior distributions we focus on the improper prior distribution with all the $b_j \rightarrow 0$. There are three important features of this improper prior distribution.

First, the improper prior distribution avoids the potential pitfall in using the Dirichlet prior with large J and all of the b_j bounded away from zero. Because we rely on J being large to make the model flexible, this potentially would be an important drawback of the method. To see the problem, let ϕ denote the probability that Z is in some set B : $\phi = \sum_{j:a_j \in B} \theta_j$. Then the posterior distribution for ϕ is a beta distribution with

$$E(\phi | d) = \sum_{j:a_j \in B} (n_j + b_j) / \sum_{j=1}^J (n_j + b_j),$$

$$\text{Var}(\phi | d) = E(\phi | d)[1 - E(\phi | d)] / \left(1 + \sum_{j=1}^J (n_j + b_j)\right).$$

Suppose $b_j = \epsilon > 0$ for all j , and consider increasing the number of support points while keeping the data d fixed. Let the fraction of support points in B approach a limit r : $\frac{1}{J} \sum_{j=1}^J 1(a_j \in B) \rightarrow r$ as $J \rightarrow \infty$. Then $E(\phi | d) \rightarrow r$, $\text{Var}(\phi | d) \rightarrow 0$, and both prior and posterior distribution of ϕ become concentrated at r , regardless of the data. In particular, this argument covers a flat prior for θ ($b_j \equiv 1$), suggesting that a flat prior distribution does not capture a lack of prior information very well when J is large.

The second point is computational. The algorithm for evaluation of $\beta = g(\theta)$ defined through moment functions takes a particularly simple form for the limiting posterior distribution that results from letting all the $b_j \rightarrow 0$ in (3). Then the θ_j corresponding to the support points a_j not observed in the sample are all zero with posterior probability one. Let $\{V_i\}_{i=1}^n$ be independently distributed according to a standard exponential distribution [i.e., the gamma distribution $\mathcal{G}(1, 1)$]. Then, for a given function $\lambda(\cdot)$,

$$\sum_{i=1}^n \lambda(z_i) V_i / \sum_{i=1}^n V_i = \sum_{j:n_j > 0} \lambda(a_j) U_j / \sum_{j:n_j > 0} U_j,$$

where $U_j = \sum_{i:z_i=a_j} V_i \sim \mathcal{G}(n_j, 1)$, using the fact that a sum of independent exponential random variables has a gamma distribution. Thus to simulate the posterior distribution of β based on (1), instead of drawing from the posterior distribution of θ and then solving

$$\sum_{j=1}^J \psi(a_j, \beta) \cdot \theta_j = 0,$$

we draw sets of iid exponential random variables $\{V_i^{(l)}\}_{i=1}^n$ and solve

$$\sum_{i=1}^n \psi(z_i, \beta^{(l)}) V_i^{(l)} = 0, \quad (4)$$

and similarly for β based on (2) we solve

$$\beta^{(l)} = \arg \min_t \sum_{i=1}^n \rho(z_i, t) V_i^{(l)}. \quad (5)$$

Repeating this for $l = 1, \dots, L$ gives us L independent draws from the posterior distribution of β . Rubin (1981) developed this simulation algorithm (using a representation for the ratio of exponentials to the sum of exponentials as gaps in order statistics from a uniform distribution), and it was applied by Lancaster (1994) in the analysis of choice-based samples.

The third issue is that the improper prior distribution for θ does not imply a unique prior distribution for the parameter of interest. Although for proper prior distributions for θ the prior distribution for β is well defined, the limiting prior distribution for β as the $b_j \rightarrow 0$ depends on the limits of the ratios b_j/b_l . To see this, consider the example discussed in which we are interested in ϕ , the probability that Z is in some set B : $\phi = \sum_{j:a_j \in B} \theta_j$. For fixed b_j the prior mean of ϕ is $E(\phi) = \sum_{j:a_j \in B} b_j / \sum_{j=1}^J b_j$. As we let the $b_j \rightarrow 0$, the limiting mean depends on the limit of the ratios of b_j/b_l . The posterior mean is $E(\phi | d) = \sum_{j:a_j \in B} (n_j + b_j) / \sum_{j=1}^J (n_j + b_j)$, which, after taking the limit $b_j \rightarrow 0$, equals $\sum_{j:a_j \in B} n_j / \sum_{j=1}^J n_j$, which does not depend on the limit of the ratios b_j/b_l . As this example illustrates, it is important to understand the implications of the choice of the limiting Dirichlet distribution. To measure the informativeness of the prior distribution for β , we propose calculating the expected posterior distribution given a small number m of observations, where we take the expectation over the empirical distribution. Let F_n denote the empirical distribution of our sample: $F_n(B) = \frac{1}{n} \sum_{i=1}^n 1(z_i \in B)$. Let $\pi_m^\beta(\cdot | \{t_i\}_{i=1}^m)$ denote the posterior distribution for β based on the m observations $Z_i = t_i$ (and assume for a moment that this posterior distribution is proper). The expected posterior distribution for β based on a random sample (with replacement) of size m from F_n is given by $\bar{\pi}_m^\beta(\cdot) = \int \pi_m^\beta(\cdot | \{t_i\}_{i=1}^m) \prod_{i=1}^m dF_n(t_i)$. To allow for the possibility of an improper posterior distribution, we modify this formula as

$$\bar{\pi}_m^\beta(\cdot) = \int \pi_m^\beta(\cdot | \{t_i\}_{i=1}^m) 1(\{t_i\}_{i=1}^m \in C_m) \prod_{i=1}^m dF_n(t_i) / \int 1(\{t_i\}_{i=1}^m \in C_m) \prod_{i=1}^m dF_n(t_i), \quad (6)$$

where the set C_m consists of the points $\{t_i\}_{i=1}^m$ such that $\pi_m^\beta(\cdot | \{t_i\}_{i=1}^m)$ is a proper distribution. If the prior distribution is not very informative for β , different small samples $\{t_i\}_{i=1}^m$ could lead to very different posterior distributions, and thus the average posterior distribution should be relatively dispersed. If we find, therefore, that this average small sample posterior distribution is dispersed compared to the full posterior distribution, we interpret that as evidence that our prior distribution does not dominate the data.

2. INSTRUMENTAL VARIABLES

The first application illustrates how the described general method can generate posterior distributions without tightly parameterized models. Such a posterior distribution is called for to include parameter uncertainty in the decision making formulation; see, for example, the work of Rossi, McCulloch, and Allenby (1995), Kandel and Stambaugh (1996), and

Barberis (2000). In this first example the large sample normal approximation to the sampling distribution can be used to approximate this posterior distribution fairly accurately. If, however, the objective is a posterior distribution for the parameter of interest, then our procedure is more direct than having to first approximate a sampling distribution by a normal distribution and then to argue that this normal distribution can be used to approximate a posterior distribution.

We use a very simple model relating earnings and schooling with a constant, additive treatment effect, linear in years of schooling. An individual may choose schooling levels by maximizing expected lifetime discounted utility, with utility depending on earnings at various schooling levels as well as costs associated with schooling. Such a decision requires the posterior distribution of earnings at the relevant schooling levels as one of the inputs. The potential outcome with treatment level s is $Y_s = Y_0 + \gamma s$, where Y_0 is the potential outcome with treatment level 0 and γ is the unknown return to schooling, common to all individuals and common to all schooling levels. The actual treatment level is X , which gives an actual outcome Y of $Y = Y_0 + \gamma X$. Let α be the population mean of Y_0 , and define the disturbance $U = Y_0 - \alpha$ so that $E_\theta(U) = 0$. The instrumental variable W satisfies $E_\theta(WU) = 0$ and $\text{Cov}_\theta(W, X) \neq 0$. We are abstracting from the presence of exogenous covariates—they could be incorporated into the presented analyses without any problems.

Let $Z = (Y, X, W)$ and $\beta' = (\alpha, \gamma)$. Then β satisfies the moment condition $E_\theta \psi(Z, \beta) = 0$ with

$$\psi(Z, \beta) = (Y - \alpha - \gamma X) \begin{pmatrix} 1 \\ W \end{pmatrix}.$$

Assuming finite support for the distribution of Z , we use the improper Dirichlet prior [with all the $b_j \rightarrow 0$ in (3)] for the parameters of this, and the posterior distribution of β can be simulated as in (4).

Our data is a subset of the data used by Angrist and Krueger (1991) containing males born in either the first or the fourth quarters between 1930 and 1939. The sample size is $n = 162,515$. The outcome variable Y is the log of weekly earnings in 1979. The treatment X is years of schooling completed, and the instrumental variable W is an indicator equal to one if the individual was born in the fourth quarter and equal to zero otherwise.

First we evaluate the information content of the prior distribution for the parameter of interest γ . To do so, we calculate the expected posterior distribution $\bar{\pi}_m^\gamma$ as in (6), with $m = 10$ observations. We compare these expected posteriors with the actual posterior distribution based on the full sample with $n = 162,515$ observations. Here are some of the quantiles for the γ distributions:

quantile:	.025	.05	.25	.50	.75	.95	.975,
$\bar{\pi}_{10}^\gamma$:	-2.43	-1.02	-.09	.07	.23	1.22	2.51,
$\pi_n^\gamma(\cdot d)$:	.047	.054	.075	.089	.104	.124	.132,
$\mathcal{N}(.089, .021^2)$:	.048	.055	.075	.089	.103	.124	.130.

It appears that the prior distribution is reasonably uninformative for γ , so that the posterior distribution mainly reflects the sample information.

The instrumental-variables estimate $\hat{\gamma}$ [i.e., the solution to $\sum_{i=1}^n \psi(z_i, \hat{\beta}) = 0$, where $\hat{\beta}' = (\hat{\alpha}, \hat{\gamma})$] is .089. An asymptotic approximation to its sampling distribution (allowing for heteroscedasticity of unknown form) gives a normal distribution with mean γ and standard deviation .021. A normal distribution with mean .089 and standard deviation .021 would provide a good approximation to our posterior distribution.

3. QUANTILE REGRESSION

The second application illustrates how the posterior distribution can be well defined when standard approximations to the sampling distribution are not appropriate. Let $Z = (X, Y)$, where Y is scalar and X is $K \times 1$. We can define a linear predictor corresponding to the τ th quantile as follows: $E_{\theta}^*(Y | X = x) = \beta'x$, where

$$\beta = \arg \min_{\beta} E_{\theta}[c_{\tau}(Y - t'X)],$$

$$c_{\tau}(t) = |t| \cdot [(1 - \tau) \cdot 1(t < 0) + \tau \cdot 1(t \geq 0)].$$

(β in general depends on τ , but this should be clear from the context.) If $\tau = .5$, then this reduces to minimizing the mean absolute error: $\min_{\beta} E_{\theta}(|Y - t'X|)$. By weighting the absolute error differently for positive and negative values, the check function $c_{\tau}(\cdot)$ extends this notion of linear predictor to other quantiles. The role of the check function in quantile regression was developed by Koenker and Bassett (1978, 1982).

Our simulation procedure produces independent draws $\{\beta^{(l)}\}_{l=1}^L$ from the posterior distribution of β . To obtain $\beta^{(l)}$, first take iid draws $\{V_i^{(l)}\}_{i=1}^n$ from a standard exponential distribution. Then solve

$$\beta^{(l)} = \arg \min_{\beta} \sum_{i=1}^n V_i^{(l)} c_{\tau}(y_i - t'x_i),$$

where the observed value of Z_i is $z_i = (x_i, y_i)$. The computations are simplified by exploiting the fact that $rc_{\tau}(t) = c_{\tau}(rt)$ if $r \geq 0$. Thus define $Y_i^{(l)} = V_i^{(l)} y_i$ and $X_i^{(l)} = V_i^{(l)} x_i$. Then

$$\beta^{(l)} = \arg \min_{\beta} \sum_{i=1}^n c_{\tau}(Y_i^{(l)} - t'X_i^{(l)}).$$

This is a linear programming problem, and we use the Barrodale–Roberts (1973) modification of the standard simplex algorithm.

Our application is based on the work of Meyer, Viscusi, and Durbin (1995), who obtained data for two states, Kentucky and Michigan, on a random sample of indemnity claims. We focus on Kentucky. The claims were filed by workers seeking compensation for work-related injury or illness. Meyer et al. concentrate on temporary total disability claims. Such a claim is filed when the person is unable to work but is expected to recover fully and return to work. The data include date injured, duration of temporary total benefits, total medical costs, previous wage, weekly benefit amount, type of injury (body part affected and the type of damage), age, sex, marital status, and an industry code.

Table 1. Quantile Regression Coefficients for Log of Duration, Kentucky High and Low Earnings Groups Pooled

Variables	Quantile					OLS
	.10	.25	.50	.75	.90	
Intercept	-5.555 (.817)	-3.067 (.497)	-1.749 (.403)	-.811 (.490)	-1.239 (.692)	-1.994 (.410)
After increase *High earnings group	.136 (.102)	.141 (.057)	.164 (.053)	.170 (.060)	.137 (.088)	.145 (.051)
After increase	-.008 (.073)	-.039 (.042)	-.029 (.034)	.013 (.040)	.074 (.057)	.000 (.033)
High earnings group	1.755 (1.352)	.525 (.931)	.024 (.771)	-.792 (1.014)	-3.191 (1.692)	-.696 (.806)

NOTE: The dependent variable is $\ln(.5 + \text{duration})$. The sample size is 5,349. The additional regressors are $\ln(\text{previous wage})$, $\ln(\text{previous wage}) * \text{High earnings group}$, Male, Married, $\ln(\text{age})$, $\ln(\text{total medical costs})$, Hospital stay indicator; Industry indicators: Manufacturing, Construction; Injury type indicators: Head, Neck, Upper extremities, Trunk, Low back, Lower extremities, Occupational diseases. The omitted industry is other industries, and the omitted injury is other injuries.

The amount of the weekly benefit is based on a schedule that determines the benefit as a function of previous earnings. The schedule has a ceiling, with earnings levels above a threshold corresponding to the same weekly benefit. Kentucky raised the maximum benefit from \$131 to \$217 per week on July 15, 1980.

Meyer et al. worked with claims with injury dates during the year before or the year after the change in the benefit schedule. They also limited the sample to a high earnings group and a low earnings group. The weekly benefit amount for the high earnings group was affected by the increase in the benefit ceiling, whereas the benefit amount for the low earnings group was not affected. Thus the low earnings group can provide a control for period effects. The basic specification in the work of Meyer et al. is

$$E_{\theta}(Y | X = x) = \beta_1 + \beta_2 \cdot x_2 \cdot x_3 + \beta_3 \cdot x_2 + \beta_4 \cdot x_3 \quad (7)$$

($x_1 \equiv 1$ denotes a constant). Here $Y = \log$ of duration, with duration measured by weeks of temporary total benefits paid; $x_2 = 1$ if injured after the benefit increase, $x_2 = 0$ otherwise;

Table 2. Quantile Regression Coefficients for Duration, Kentucky High and Low Earnings Groups Pooled

Variables	Quantile					OLS
	.10	.25	.50	.75	.90	
Intercept	-6.199 (1.157)	-7.258 (1.441)	-8.972 (1.779)	-11.566 (3.310)	-19.848 (7.254)	-25.886 (8.412)
After increase *High earnings group	.229 (.143)	.302 (.165)	.873 (.230)	1.351 (.554)	2.661 (1.339)	1.665 (1.043)
After increase	-.052 (.085)	-.032 (.097)	-.116 (.138)	.122 (.289)	.498 (.629)	.457 (.674)
High earnings group	.051 (2.546)	-.356 (2.848)	-1.655 (3.528)	-11.541 (9.299)	-56.802 (27.400)	-41.783 (16.539)

NOTE: The dependent variable is duration (in weeks). The sample size is 5,349. The additional regressors are the same as those in Table 1. The omitted industry is other industries, and the omitted injury is other injuries.




Figure 1. Posterior Histogram. $q = .5$ (long list for x).

$x_3 = 1$ if high earnings group, $x_3 = 0$ otherwise. The key coefficient is β_2 , measuring the effect of the benefit increase on time out of work, with controls for period and for the earnings group:

$$\beta_2 = [E_\theta(Y | x_2 = 1, x_3 = 1) - E_\theta(Y | x_2 = 0, x_3 = 1)] - [E_\theta(Y | x_2 = 1, x_3 = 0) - E_\theta(Y | x_2 = 0, x_3 = 0)].$$

An appealing aspect of the Meyer et al. analysis is that it is plausible to regard the injury date, and hence the applicable benefit schedule, as if it were randomly assigned.

To account for possible changes in the composition of the sample after the benefit increase, Meyer et al. included regression controls for attributes of the individual, the job, and the injury—16 regressors in addition to the 4 in (7). The last column of Table 1 presents least squares estimates (and conventional standard errors) corresponding to Table 6 in the Meyer et al. work. The first five columns of Table 1 present estimates of the linear predictor coefficients corresponding to the .10, .25, .50, .75, and .90 quantiles. These estimates are based on the simulation procedure described earlier. The point estimates are posterior medians, and the standard errors in parentheses are constructed so that the point estimate plus or minus 1.96 standard errors gives an interval with a .95 posterior probability. The key coefficients [corresponding to β_2 in (7)] are in the second row. The effect of the benefit increase is fairly constant across the quantiles, suggesting a location model in which the distribution of log duration shifts rigidly in response to the benefit increase.

Table 2 presents results using duration out of work (in weeks) instead of its logarithm. Now the estimates show a substantial increase as we go from low to high quantiles, suggesting that the effect of the benefit increase is concentrated on the upper half of the duration distribution. The estimated effect on the median of the distribution is .87 weeks, with a standard error of .23. In contrast, the least squares estimate of the effect on the mean of the distribution is quite imprecise, with a point estimate of 1.66 and a standard error of 1.04.

The histogram of the draws from the posterior distribution of β_2 is shown in Figure 1 for $\tau = .5$, using duration in weeks. The posterior mean is .87, and the posterior standard deviation is .23. Thus assuming the posterior distribution is normal and using $.87 \pm 1.96 \times .23$ gives a probability interval close to the one we constructed without assuming normality.

We examine the influence of the prior distribution by calculating the expected posterior distribution $\tilde{\pi}_m^\beta$ as in (6), for $m = 21$ observations, and comparing this distribution with the posterior distribution $\pi_n^\beta(\cdot | d)$ based on the full sample with $n = 5,349$ observations. Here are some of the quantiles of the β_2 distributions for $\tau = .5$, using duration in weeks:

quantile:	.025	.05	.25	.50	.75	.95	.975,
$\tilde{\pi}_{21}^{\beta_2}$:	−290	−157	−20.4	1.01	24.3	184	323,
$\pi_n^{\beta_2}(\cdot d)$:	.41	.49	.71	.87	1.03	1.25	1.32.

The prior distribution is dominated by the sample information.

Now consider dropping all the predictor variables except for the four that appear in (7): $1, x_2 \cdot x_3, x_2, x_3$. We compare




Figure 2. Posterior Histogram. $q = .5$ (short list for x).




Figure 3. Posterior Histogram. $q = .9$ (short list for x).

the expected posterior distribution for $m = 5$ observations with the posterior distribution based on the full sample. Here are quantiles of these distributions for β_2 with $\tau = .5$, using duration in weeks:

quantile:	.025	.05	.25	.50	.75	.95	.975,
$\bar{\pi}_5^{\beta_2}$:	-121	-36	-6	1	9	59	110,
$\pi_n^{\beta_2}(\cdot d)$:	0	0	1	1	2	2	2.

The posterior histogram for β_2 is in Figure 2. It is concentrated on just four points: $-1, 0, 1,$ and 2 weeks, with posterior probabilities of $.01, .14, .55,$ and $.30$. This reflects the discreteness of the benefit duration distribution. The upper tail of that distribution is somewhat continuous, but 56% of the distribution is concentrated on the integers from 0 to 4 weeks. The $(.5, .75, .9, .95, .975)$ quantiles are $(4, 8, 15, 25, 49)$ weeks. Including the long list of predictor variables smoothes out this discreteness in the outcome variable, in the sense of producing a residual distribution (for $Y - \beta'X$) that is much closer to being continuous.

Here are the quantiles of the β_2 distributions for $\tau = .9$, using just the four regressors in (7) and duration in weeks:

quantile:	.025	.05	.25	.50	.75	.95	.975,
$\bar{\pi}_5^{\beta_2}$:	-145	-41	-7	1	10	72	124,
$\pi_n^{\beta_2}(\cdot d)$:	2	3	5	7	8	11	12.

The posterior histogram for β_2 is in Figure 3. This is closer to a normal distribution, corresponding to the continuity in the upper tail of the duration distribution.

The standard asymptotic distribution theory for quantile regression requires that the distribution of the residual $Y - \beta'X$ (conditional on θ) be absolutely continuous with a positive density in a neighborhood of zero. This requirement may be satisfied because the distribution of Y conditional on X is continuous. Alternatively, even if Y is discrete, it may be satisfied because $X'\beta$ is continuous. For example, with Y binary and X uniform on $[0, 1]$, and $E[Y | X] = X$, the limiting distribution of the coefficient in a quantile regression is normal despite the binary nature of Y . In our example Y is discrete with most mass concentrated on a few values. With only three binary regressors, the resulting distribution of the residual is still highly discrete. With the long list of regressors, although many of them are discrete, the continuity requirement for the

residual is much closer to being satisfied, and the standard large sample approximation to the sampling distribution is more accurate. In contrast, our posterior distributions provide straightforward inferences that do not rely on the approximate normality of a sampling distribution.

4. CONCLUSION

The Bayesian approach to inference provides an attractive conceptual framework because of its connection with optimization concepts in decision theory and its lack of reliance on large-sample approximations. In practice, its use has been limited by the requirement of a fully specified parametric model because many econometric models are only partly specified. In this article we presented two applications of a less parametric Bayes approach that are due to Ferguson (1973, 1974) and Rubin (1981). In the first application, the decision-theoretic nature of the underlying question forces the use of posterior distributions rather than sampling distributions. In the second application, the assumptions underlying the asymptotic normality of the sampling distributions are clearly violated, but inference based on posterior distributions is straightforward.

ACKNOWLEDGMENTS

The authors thank David Cox, Jinyong Hahn, and Neil Shephard for helpful comments and Alan Krueger and Bruce Meyer for making their data available to us. The National Science Foundation provided financial support.

[Received December 2000. Revised June 2001.]

REFERENCES

Angrist, J., and Krueger, A. (1991), "Does Compulsory School Attendance Affect Schooling and Earnings?," *Quarterly Journal of Economics*, 106, 979-1014.
 Barberis, N. (2000), "Investing for the Long Run when Returns are Predictable," *Journal of Finance*, 55, 225-264.
 Barrodale, I., and Roberts, F. (1973), "An Improved Algorithm for Discrete l_1 Linear Approximation," *SIAM Journal of Numerical Analysis*, 10, 839-848.
 Chamberlain, G., and Imbens, G. (1995), "Semiparametric Applications of Bayesian Inference," Discussion Paper 1716, Harvard Institute of Economic Research Cambridge, MA.
 Ferguson, T. (1973), "A Bayesian Analysis of Some Nonparametric Problems," *The Annals of Statistics*, 1, 209-230.
 Ferguson, T. (1974), "Prior Distributions on Spaces of Probability Measures," *The Annals of Statistics*, 2, 615-629.

- Hirano, K. (2002), "Semiparametric Bayesian Inference in Autoregressive Panel Data Models," *Econometrica*, 70, 781–800.
- Kandel, S., and Stambaugh, R. (1996), "On the Predictability of Stock Returns: An Asset-Allocation Perspective," *Journal of Finance*, 51, 385–424.
- Koenker, R., and Bassett, G. (1978), "Regression Quantiles," *Econometrica*, 46, 33–50.
- Koenker, R., and Bassett, G. (1982), "Robust Tests for Heteroscedasticity Based on Regression Quantiles," *Econometrica*, 50, 43–61.
- Lancaster, T. (1994), "Bayes WESML: Posterior Inference From Choice-Based Samples," unpublished manuscript, Brown University, Providence, RI.
- Meyer, B., Viscusi, W. K., and Durbin, D. (1995), "Workers' Compensation and Injury Duration: Evidence From a Natural Experiment," *American Economic Review*, 85, 322–340.
- Rossi, P., McCulloch, R., and Allenby, G. (1995), "Hierarchical Modelling of Consumer Heterogeneity: An Application to Target Marketing," in *Case Studies in Bayesian Statistics* (Vol. II), *Lecture Notes in Statistics*, 105, eds. C. Gatsonis, J. Hodges, R. Kass, and N. Singpurwalla, New York: Springer-Verlag, 323–349.
- Rubin, D. (1981), "The Bayesian Bootstrap," *The Annals of Statistics*, 9, 130–134.