



Simulated Likelihood Estimation of the Normal-Gamma Stochastic Frontier Function

WILLIAM H. GREENE

wgreene@stern.nyu.edu; <http://www.stern.nyu.edu/~wgreene>

Department of Economics, Stern School of Business, New York University, New York University, 44 West 4th St., New York, NY, 10012

Abstract

The normal-gamma stochastic frontier model was proposed in Greene (1990) and Beckers and Hammond (1987) as an extension of the normal-exponential proposed in the original derivations of the stochastic frontier by Aigner, Lovell and Schmidt (1977). The normal-gamma model has the virtue of providing a richer and more flexible parameterization of the inefficiency distribution in the stochastic frontier model than either of the canonical forms, normal-half normal and normal-exponential. However, several attempts to operationalize the normal-gamma model have met with very limited success, as the log likelihood is possessed of a significant degree of complexity. This note will propose an alternative approach to estimation of this model based on the method of maximum simulated likelihood estimation as opposed to the received attempts which have approached the problem by direct maximization.

JEL classification: C1, C2, C5

Keywords: frontier, gamma, simulated likelihood

The stochastic frontier model was proposed by researchers in 1977 by Aigner, Lovell and Schmidt (1977) and Meeusen and van den Broeck (1977). The original form of the model,

$$y = \beta' \mathbf{x} + v - u$$

where $\beta' \mathbf{x} + v$ constitute a conventional regression model and u is a one side disturbance that is distributed either as half normal or exponential, has stood the test of nearly 25 years as the workhorse of the literature on frontier estimation. Notwithstanding the original model's longevity and distinguished service record, researchers have proposed many variants of the model in attempts to generalize the distribution of the inefficiency distribution, $f(u)$.

The normal-gamma stochastic frontier model was proposed in Greene (1990), Beckers and Hammond (1987) and Stevenson (1990) as an extension of the normal-exponential proposed

in the original derivations of the stochastic frontier by Aigner, Lovell and Schmidt (1977). The normal-gamma model provides a richer and more flexible parameterization of the inefficiency distribution in the stochastic frontier model than either of the canonical forms, normal-half normal and normal-exponential. However, several attempts to operationalize the normal-gamma model have met with very limited success, as the log likelihood is possessed of a significant degree of complexity. Greene (1990) attempted a direct, but crude maximization procedure which, as documented by Ritter and Simar (1997) was not sufficiently accurate to produce satisfactory estimates. (Ritter and Simar concluded from their work that even an accurate estimator of this model would suffer from significant identification problems. That appears not to be the case. We will return to this point below.) Stevenson (1980) made note of the difficulties of estimation early on and proposed limiting attention to the Erlang form (noted below), which is a significant restriction of the model. Beckers and Hammond (1987) derived a form of the log likelihood which showed some potential, but, in the end, remained exceedingly complicated and was never operationalized. We note in passing at this point the work of Tsionas (2000), who made considerable progress on the model in a Bayesian framework, but provided estimates of the posterior distribution of u , rather than direct estimation of the parameters. This note will confine attention to classical, parametric analysis of the model.

We will propose an alternative approach to estimation of this model based on the method of maximum simulated likelihood estimation as opposed to the received attempts which have approached the problem by direct maximization. The previous work on this model has approached the estimation problem by direct maximization of what turns out to be a very complicated log likelihood function. As shown below, the log likelihood function for the normal-gamma model is actually the expectation of a random variable which can be simulated. This suggests the method of maximum simulated likelihood as a method of estimation.

In Section 1 below, we will briefly review the modeling framework for the stochastic frontier model leading to the normal-gamma model. Section 2 will analyze the problem of maximizing the log likelihood function. The simulated log likelihood function is derived, then two practical aspects of the estimation process are presented. First, we note a useful method of producing the necessary random draws for the simulation—the model requires simulation of draws from a truncated distribution, which is, itself, a significant complication. Second, we introduce a new tool which has proved very useful in the simulated likelihood maximization of models for discrete choice, the Halton draw. Halton draws provide a method of dramatically speeding up the process of maximization by simulation. A brief application of the technique is presented in Section 3. Some conclusions are drawn in Section 4.

1. The Stochastic Frontier Model

The motivation for and mechanical details of estimation of the stochastic frontier model have appeared at many places in the literature, so it is not necessary to reiterate them here. For further details, the reader is referred, for example, to one of the many surveys on the subject, such as Kumbhakar and Lovell (2000). In this section, we will recall the aspects of the model that appear in the results below.

For convenience, we suppress the observation subscripts that denote an observed sample of N observations, $i = 1, \dots, N$. The stochastic frontier model is

$$y = \beta' \mathbf{x} + \varepsilon$$

$$\varepsilon = v - u$$

$$v \sim N[0, \sigma_v^2]$$

$u \geq 0$, with continuous density, $f(u | \theta)$, where θ is a vector of parameters.

The objective of estimation is $\beta, \theta, \sigma_v^2$ then u . As a practical matter, the firm specific inefficiency, u is the ultimate objective of the estimation. As has been widely documented, however, this is not feasible. Jondrow, Lovell, Materov and Schmidt (1982) suggested the feasible alternative

$$E[u | \varepsilon] = g(\varepsilon, \theta)$$

which can be estimated using the estimated parameters and the observed data. (Properties of this estimator have been explored, e.g., by Horrace and Schmidt, 1996.)

The received literature has relied on two specific formulations of the inefficiency distribution,

$$\text{half normal: } u = |U|, \quad U \sim N[0, \sigma_u^2],$$

and

$$\text{exponential } f(u) = \theta \exp(-\theta u), \quad \theta > 0.$$

As noted, the parameters of the distributions are of secondary importance in the estimation process. What is actually of greatest interest is the inefficiency component of the underlying model and estimation of values that will enable the comparison of individuals in the sample. The Jondrow, et al. formulas for the two models suggested are

$$E[u | \varepsilon] = \sigma \lambda / (1 + \lambda^2) [\phi(\varepsilon \lambda / \sigma) / \{1 - \Phi(\varepsilon \lambda / \sigma)\} - \varepsilon \lambda / \sigma].$$

where

$$\varepsilon = y - \beta' \mathbf{x},$$

$$\lambda = \sigma_u / \sigma_v,$$

$$\sigma = \sqrt{\sigma_u^2 + \sigma_v^2},$$

$\phi(t)$ and $\Phi(t)$ = standard normal density and distribution functions, respectively,
for the normal-half normal model and

$$E[u | \varepsilon] = z + \sigma_v \phi(q / \sigma_v) / \Phi(q / \sigma_v)$$

$$q = \varepsilon - \theta \sigma_v^2$$

for the normal-exponential model.

The appeal of these two distributions is their known and quite straightforward closed form. See, e.g., the original paper by Aigner et al. (1977) or the survey by Kumbhakar and Lovell (2000). Many extensions have been proposed which layer deeper parameterizations

on the mean and variance components of the half normal random variable, such as replacing the zero mean in the half normal distribution with a regression,

$$E[U | \mathbf{z}] = \delta' \mathbf{z}.$$

or adding heteroscedasticity in the variances, as in

$$\sigma_v^2 = \exp(\alpha' \mathbf{h}).$$

(Recent extensive surveys which discuss these are Coelli et al., 1997 or Kumbhakar and Lovell, 2000.) However, even with these extensions, the normal-half normal has remained the workhorse of the literature.

2. The Normal-Gamma Stochastic Frontier Model

The normal-gamma frontier model provides an extension to the normal-exponential model;

$$f(u) = \theta^P / \Gamma(P) \exp(-\theta u) u^{P-1}.$$

This formulation provides a more flexible parameterization of the distribution. Values of P larger than one allow the mass of the distribution to move away from zero-values of P less than one produce densities that resemble the exponential distribution. Thus, the prior assumption of a value of P (e.g., 1) amounts to a substantive assumption about the distribution of inefficiencies in the population.

2.1. Log Likelihood for the Normal-Gamma Stochastic Frontier Model

The log likelihood function for the normal-gamma model is derived in Greene (1990) and in a different form in Beckers and Hammond (1987). We will proceed directly to the end result here. For the normal-exponential (NE) model,

$$\log L_{NE} = N \left\{ \log \theta + \frac{1}{2} \theta^2 \right\} + \sum_{i=1}^N \{ \theta \varepsilon_i + \log \Phi[-(\varepsilon_i / \sigma_v + \theta \sigma_v)] \}$$

where $\varepsilon_i = y_i - \beta' \mathbf{x}_i$. The log likelihood for the normal-gamma (NG) model is that for the normal-exponential model plus the term which has complicated the analysis to date;

$$\log L_{NG} = \log L_{NE} + N[(P-1) \log \theta - \log \Gamma(P)] + \sum_{i=1}^N \log h(P-1, \varepsilon_i)$$

$$h(r, \varepsilon_i) = \frac{\int_0^\infty z^r \left(\frac{1}{\sigma_v}\right) \phi\left(\frac{z-\mu_i}{\sigma_v}\right) dz}{\int_0^\infty \left(\frac{1}{\sigma_v}\right) \phi\left(\frac{z-\mu_i}{\sigma_v}\right) dz}, \quad \mu_i = -\varepsilon_i - \theta \sigma_v^2.$$

Note if $P = 1$, then $h(r, \varepsilon_i) = 1$ and $\log h(r, \varepsilon_i) = 0$ which returns the log likelihood for the exponential model.

2.2. The Simulated Log Likelihood

In principle, the parameters of the model are estimated by maximizing the log likelihood function. The ‘problem’ that has heretofore complicated the matter is computing $h(P-1, \varepsilon_i)$. Stevenson simplified the problem by suggesting the Erlang form, that is the cases with integer P . This does produce a simpler model, but drastically restricts the model. Beckers and Hammond reformulated the function in terms of a series of Pochhammer symbols. (The Pochhammer symbol is $a_x = \Gamma(a+x)/\Gamma(a)$. Accurate computation when x is small relative to a requires special routines.) While the Beckers and Hammond (1987) formulation did provide a complete (if not closed) functional form for the integral, it was not operationalized. Greene’s original application of the model used a very crude approximation to the integral based on Simpson’s rule and the areas of trapezoids. Simar showed in two later papers that this approach was insufficiently accurate. Simar’s results might have implied a full implementation, however he did not do the analysis from that viewpoint, so a method of proceeding for the practitioner remains to be developed. The purpose of this note is to suggest such a procedure.

We base an alternative approach on the result that, as can be seen by inspection, $h(r, \varepsilon_i)$ is the expectation of a random variable;

$$h(r, \varepsilon_i) = E[z^r \mid z \geq 0] \quad \text{where } z \sim N[\mu_i, \sigma_v^2] \quad \text{and} \quad \mu_i = -\varepsilon_i - \theta\sigma_v^2.$$

In words, $h(r, \varepsilon_i)$ is the expected value of z^r where z has a truncated at zero normal distribution with underlying mean μ_i and variance σ_v^2 . We propose to estimate $h(r, \varepsilon_i)$ by using the mean of a sample of draws from this distribution. For given values of (i.e., conditioned on) ε_i and μ_i (i.e., $y_i, \mathbf{x}_i, \beta, \sigma_v, \theta, r$), by the Lindberg-Levy variant of the central limit theorem (see Greene, 2000), $h(r, \varepsilon_i)$ would be consistently estimated by

$$\hat{h}_i = \frac{1}{Q} \sum_{q=1}^Q z_{iq}^r$$

where z_{iq} is a random draw from the truncated normal distribution with mean parameter μ_i and variance parameter σ_v . (The truncated normal distribution has finite moments of all orders, so this is an application of the most narrow version of the central limit theorem.) We propose, then, to maximize the simulated log likelihood function

$$\log L_{NG,S} = \log L(\text{exponential}) + N[(P-1) \log \theta - \log \Gamma(P)] + \sum_{i=1}^N \log \hat{h}(P-1, \varepsilon_i)$$

(Properties of this method of maximum likelihood estimation are discussed elsewhere, such as in the November, 1994 symposium in the *Review of Economics and Statistics*. The techniques has been widely used in estimation of multinomial probit models (see the aforementioned symposium) and in estimation of discrete choice models with random parameters (see Revelt and Train (1999) for example).

2.3. Constructing the Simulated Log Likelihood Function

We now consider two practical issues, generating the random draws from the truncated distribution and the prior problem of producing the primitive draws for the simulation.

Computing the simulated log likelihood function will require Q draws from the truncated normal distribution for each observation. In principle, these can be drawn by the simple rejection method. Draws are taken from an untruncated distribution and rejected until a draw from the desired region is obtained. The problems with this approach (notwithstanding its raw inelegance) are that it can require huge numbers of draws if the desired region is near the tail of the underlying distribution, and, most importantly for maximum simulated estimation likelihood, has the result that different draws will be used for different computations of the log likelihood function. As such, the simulated log likelihood will no longer be a continuous function of the parameters. The iterations will never converge. An alternative method used, e.g., in Geweke, Keane and Runkle (1994) is a superior approach: This method requires only a single draw; the procedure is carried out as follows: Let

T = the truncation point

μ = the mean of untruncated distribution

σ = the standard deviation of untruncated distribution

$P_L = \Phi[(T - \mu)/\sigma]$

F = a draw from the standard continuous uniform distribution $U[0, 1]$

and

$$z = \mu + \sigma \Phi^{-1}[P_L + F \times (1 - P_L)]$$

Then,

z = the draw from the desired truncated distribution.

For implementing Geweke's method to compute $h(r, \varepsilon)$, we have the following:

$$T = 0,$$

$$\varepsilon_i = y_i - \beta' \mathbf{x}_i,$$

$$\mu_i = -\varepsilon_i - \theta \sigma_v^2,$$

$$\sigma = \sigma_v,$$

$$P_L = \Phi[(\varepsilon_i + \theta \sigma_v^2)/\sigma_v] = \Phi(-\mu_i/\sigma).$$

Collecting all terms, we have the simulated log likelihood function:

$$\begin{aligned} \text{Log } L = & N \left\{ \log \theta + \frac{1}{2} \sigma_v^2 \theta^2 \right\} + \sum_{i=1}^N \{ \theta \varepsilon_i + \log \Phi(\mu_i/\sigma_v) \} + N[(P-1) \log \theta - \log \Gamma(P)] \\ & + \sum_{i=1}^N \log \left\{ \frac{1}{Q} \sum_{q=1}^Q \left[\mu_i + \sigma_v \Phi^{-1} \left(F_{iq} + (1 - F_{iq}) \Phi \left(\frac{-\mu_i}{\sigma_v} \right) \right) \right]^{P-1} \right\} \end{aligned}$$

The simulated log likelihood function is a smooth continuous function of the parameters, $[\beta', \sigma_v, \theta, P]$. Derivatives are fairly simple as well. (They are presented in the Appendix.) Conventional maximization methods such as DFP or BFGS can be used for the optimization. We have used the BHHH (outer product of gradients) estimator to compute the asymptotic covariance matrix of the simulated maximum likelihood estimator.

For maximizing the simulated log likelihood we emphasize that the simulation would be over the F_{iq} draws from the standard uniform distribution. Use Q points for the simulation. In order to achieve continuity, it is necessary to use the same set of draws, $[F_{i1}, F_{i2}, \dots, F_{iQ}]$, for every computation. Every sum over $q = 1, \dots, Q$ uses the same set of random draws. (Each observation has its own unchanging vector of draws.) How this is done by researchers who employ this method varies from one application to another. Ruud and McFadden (1994) and Bhat (1999) recommend maintaining a fixed, indexed reservoir of random draws. For our simulations, we have, instead, controlled the draws by associating with each individual observation in the sample a unique seed for the random number generator, and restarting the generator at that value for each observation as needed. In connection to the point noted in the next section, we note that this method is slightly more time consuming than the fixed pool approach. But, it requires no additional computer memory while the fixed pool method will be profligate with memory when analyzing a large sample and using a large number of draws for the simulation.

2.4. *Efficient Computation of the Simulated Log Likelihood Function*

For implementation, there remains the practical consideration of how best to obtain the underlying random draws, F_{iq} , from the $U[0, 1]$ distribution that drive the simulation. We consider two issues, how many draws to obtain and how to create them. On the first point, the literature varies. The glib advice, “the more the better,” is not helpful when time becomes a consideration and, in fact, the marginal benefit of additional draws eventually becomes nil. Again, researchers differ, but received studies seem to suggest that several hundred to over 1,000 are needed (see, e.g., Bhat, 1999). The second consideration concerns how to obtain the draws. For most practitioners, the conventional approach amounts to relying on a random number generator buried within, say, Gauss, under the assumption that the draws it produces are truly random by the standards of received tests of randomness. A recently emerging literature (see, e.g., Bhat, 1999 or Revelt and Train, 1999), based on work in numerical analysis, suggests that this view of the process neglects some potentially large gains in computational efficiency. An alternative approach based on Halton draws (derived below) promises to improve the computations of simulated likelihoods such as ours.

The Halton sequence of draws is based on an ‘intelligent’ set of values for the simulation. The process is motivated by the idea that true randomness is not really the objective in producing the simulation. Coverage of the range of the variation with a sequence of draws that are uncorrelated with the variables in the model is the objective—the simulation is intended, after all, to estimate an integral, that is, an expectation. Numerical analysts have found that a small number of Halton draws is as effective as or more so than a large number of pseudorandom draws using a random number generator.

Halton sequences are generated as follows: Let s be a prime number larger than 2. Expand the sequence of integers $g = 1, \dots$ in terms of the base s as

$$g = \sum_{i=0}^I b_i s^i$$

where by construction, $0 \leq b_i \leq s-1$ and $s^I \leq g < s^{I+1}$. Then, the Halton sequence of values that corresponds to this series is

$$H(g) = \sum_{i=0}^I b_i s^{-i-1}$$

Halton values are contained in the unit interval. They are not random draws, but they are well spaced in the interval. Halton sequences are being used with great success in estimating “mixed logit” models that require huge amounts of simulation (see, e.g., Train, 1999). The computational efficiency compared to pseudorandom draws appears to be as much as 10 to 1 in moderately sized models. The same results are obtained with only 1/10 of the number of draws. Computation time in these models is roughly linear in the number of replications, so time savings are potentially very large. (The advantage diminishes somewhat as the dimensions of the model become larger. In our application, we are integrating over only a single dimension.)

3. An Application

We have applied the preceding to Christensen and Greene’s (1976) electricity data. We used the 1970 sample that contains 158 firms and holding companies. The regression function is a simple Cobb-Douglas cost function with a quadratic term in log output;

$$\log(C/P_f) = \beta_1 + \beta_2 \log(Pl/P_f) + \beta_3 \log(Pk/P_f) + \beta_4 \log Y + \beta_5 \log^2 Y + v + u$$

where C is total cost of generation, Y is output, and P_f , Pk and Pl are the unit prices of fuel, capital and labor, respectively. (Translation of the original model to a cost function requires only a trivial change of sign of some terms in the log likelihood and its derivatives.) Table 1 presents the estimation results. The Halton results are based on $Q = 200$.

The estimated inefficiencies from the two sets of estimates are very similar, as the last two rows of Table 1 would suggest. Also, the estimate of P in the gamma models is not particularly large and, moreover, is close to one. Table 2 lists the descriptive statistics of the Jondrow et al. estimator of $E[u | \varepsilon]$ for the models. The JLMS efficiency measure has the simple form

$$E[u | \varepsilon] = h(P, \varepsilon_i) / h(P - 1, \varepsilon_i).$$

for the normal-gamma model.

Table 1. Estimated stochastic frontier functions.

	Exponential		Gamma-Halton	
Parameter	Est.	Std. Err.	Est.	Std. Err.
β_1	−7.0345	0.207	−7.0336	0.2148
β_2	0.1449	0.0421	0.1451	0.0417
β_3	0.1391	0.0390	0.1379	0.0388
β_4	0.4413	0.0302	0.4445	0.0318
β_5	0.0286	0.00208	0.0284	0.00217
θ	11.012	2.697	9.5770	6.231
σ_v	0.1030	0.0127	0.1070	0.0170
P	1.0000	0.0000	1.0820	0.7374
Log likelihood	95.05542		93.14108	
<i>Estimated standard deviations of the underlying variable</i>				
v	0.10296		0.10182	
u	0.09083		0.08967	

Table 2. Descriptive statistics for JLMS estimated inefficiencies.

	Mean	Minimum	Maximum
Exponential	0.090813	0.022991	0.443508
Gamma	0.076925	0.014862	0.440013

4. Conclusions

The preceding has proposed a method of estimating the normal-gamma frontier model. The operational aspects of the proposed method are fairly straightforward, and implementation should be relatively simple. The estimator used here was built into *LIMDEP* (2002, forthcoming release) but could easily be programmed in *Gauss* or in a low level language if the analyst prefers. Experience with the estimator is limited, but the results do suggest a useful extension of the stochastic frontier model. If it is established that the estimator actually does work well in practice, then familiar extensions might be added to it. For example—and this could be added to the exponential model though we have not seen it—the location parameter can be parameterized to include heterogeneity, in the form $\theta_i = \exp(\delta' \mathbf{z}_i)$ for example.

Whether Ritter and Simar's observations about the (non)identifiability of the normal-gamma model prove general is an empirical question. Their result was a matter of degree, not a definitive result. That is, they found that identification of the model would be "difficult," not impossible. Moreover, the conclusion was drawn from a small simulation in which the "model" contained only a constant term. As such, they essentially fit a four parameter distribution from a sample on a single variable. The results are not surprising. Their main result concerns the number of observations one might need to do effective estimation, which has no generality in the context of a model with covariates and wider variability

in the dependent variable. We find that although the peak in the likelihood function is not sharp, it is well defined and computation of the estimator in the preceding was not especially difficult. As such, as might be expected, further research is called for.

Appendix: Derivatives of the Simulated Log Likelihood

The i th term in the log likelihood for the normal-gamma model is

$$\begin{aligned} \log L_i = & \log \theta + \frac{1}{2} \sigma_v^2 \theta^2 + \theta \varepsilon_i + \log \Phi(\mu_i / \sigma_v) + (P-1) \log \theta - \log \Gamma(P) \\ & + \log \left\{ \frac{1}{Q} \sum_{q=1}^Q \left[\mu_i + \sigma_v \Phi^{-1} \left(F_{iq} + (1 - F_{iq}) \Phi \left(\frac{-\mu_i}{\sigma_v} \right) \right) \right]^{P-1} \right\}. \end{aligned}$$

It is convenient to gather the terms and collect a few to rewrite this as

$$\begin{aligned} \log L_i = & P \log \theta - \log \Gamma(P) - \theta \sigma W_i - \theta^2 \sigma^2 / 2 + \log \Phi(W_i) \\ & + \log \left\{ \frac{1}{Q} \sum_{q=1}^Q \left[\sigma W_i + \sigma \Phi^{-1}(F_{iq} + (1 - F_{iq}) \Phi(-W_i)) \right]^{P-1} \right\} \end{aligned}$$

where $W_i = \mu_i / \sigma$ and we have dropped the subscript on σ_v to simplify the notation. We have now written the function in terms of P , θ , σ and W_i which is a function of θ and σ as well as β . It will also be convenient to define a symbol for the bracketed term in the summation, so we further compress the function to

$$\log L_i = P \log \theta - \log \Gamma(P) - \theta \sigma W_i - \theta^2 \sigma^2 / 2 + \log \Phi(W_i) + \log \left\{ \frac{1}{Q} \sum_{q=1}^Q [C_{iq}]^{P-1} \right\}$$

To avoid some notational clutter, let

$$H_i = \left\{ \frac{1}{Q} \sum_{q=1}^Q [C_{iq}]^{P-1} \right\}, \quad \lambda_i = \frac{\phi(W_i)}{\Phi(W_i)}, \quad \text{and} \quad D_{iq} = (P-1)[C_{iq}]^{P-2}$$

Recall that $W_i = \mu_i / \sigma = -\varepsilon_i / \sigma - \theta \sigma$. With these preliminaries, then,

$$\frac{\partial \log L_i}{\partial \theta} = \frac{P}{\theta} - \sigma W_i - \theta \sigma \frac{\partial W_i}{\partial \theta} - \theta \sigma^2 + \lambda_i \frac{\partial W_i}{\partial \theta} + \frac{1}{H_i} \frac{1}{Q} \sum_{q=1}^Q D_{iq} \frac{\partial C_{iq}}{\partial \theta}, \quad \frac{\partial W_i}{\partial \theta} = -\sigma,$$

$$\frac{\partial \log L_i}{\partial \sigma} = \theta W_i - \theta \sigma \frac{\partial W_i}{\partial \sigma} - \theta^2 \sigma + \lambda_i \frac{\partial W_i}{\partial \sigma} + \frac{1}{H_i} \frac{1}{Q} \sum_{q=1}^Q D_{iq} \frac{\partial C_{iq}}{\partial \sigma}, \quad \frac{\partial W_i}{\partial \sigma} = \varepsilon_i / \sigma^2 - \theta,$$

$$\frac{\partial \log L_i}{\partial \varepsilon_i} = \theta \sigma \frac{\partial W_i}{\partial \varepsilon_i} + \lambda_i \frac{\partial W_i}{\partial \varepsilon_i} + \frac{1}{H_i} \frac{1}{Q} \sum_{q=1}^Q D_{iq} \frac{\partial C_{iq}}{\partial \varepsilon_i}, \quad \frac{\partial W_i}{\partial \varepsilon_i} = -1/\sigma, \quad \frac{\partial \varepsilon_i}{\partial \beta} = -\mathbf{x}_i$$

$$\frac{\partial \log L_i}{\partial P} = \log \theta - \Psi(P) + \frac{1}{H_i} \frac{1}{Q} \sum_{q=1}^Q (\log C_{iq}) [C_{iq}]^{P-1}$$

To complete the derivation, we require the derivatives of

$$C_{iq} = \sigma W_i + \sigma \Phi^{-1}(F_{iq} + (1 - F_{iq})\Phi(-W_i))$$

Let $a_{iq} = F_{iq} + (1 - F_{iq})\Phi(-W_i)$ so $C_{iq} = \sigma[W_i + \Phi^{-1}(a_{iq})]$. Then,

$$\frac{\partial C_{iq}}{\partial W_i} = \sigma + \sigma \frac{F_{iq}\phi(-W_i)}{\phi(a_{iq})} = \sigma[1 + E_{iq}].$$

Then, inserting derivatives of W_i where needed, we have

$$\begin{aligned}\frac{\partial C_{iq}}{\partial \theta} &= -\sigma[1 + E_{iq}](-\sigma) \\ \frac{\partial C_{iq}}{\partial \sigma} &= [W_i + \Phi^{-1}(a_{iq})] + \sigma[1 + E_{iq}](\varepsilon_i/\sigma^2 - \theta) \\ \frac{\partial C_{iq}}{\partial \varepsilon_i} &= \sigma[1 + E_{iq}](-1/\sigma)\end{aligned}$$

Acknowledgments

I would like to thank the participants at the North American Productivity Workshop, held at Union College in Schenectady, New York, June 15–17, 2000, for helpful comments.

References

- Aigner, D., K. Lovell and P. Schmidt. (1977). "Formulation and Estimation of Stochastic Frontier Production Function Models." *Journal of Econometrics* 6, 21–37.
- Battese, G. and G. Corra. (1977). "Estimation of a Production Frontier Model, with Application to the Pastoral Zone off Eastern Australia." *Australian Journal of Agricultural Economics* 21, 169–179.
- Beckers, D. and C. Hammond. (1987). "A Tractable Likelihood Function for the Normal-Gamma Stochastic Frontier Model." *Economics Letters* 24, 33–38.
- Bhat, C. (1999). "Quasi-Random Maximum Simulated Likelihood Estimation of the Mixed Multinomial Logit Model." Working Paper, Department of Civil Engineering, University of Texas, Austin.
- Christensen, L. and W. Greene. (1976). "Economies of Scale in U.S. Electric Power Generation." *Journal of Political Economy* 84, 655–676.
- Coelli, T., D. Rao and G. Battese. (1997). *An Introduction to Efficiency and Productivity Analysis*. Kluwer Academic Publishers, Boston.
- Geweke, J., M. Keand and D. Runkle. (1994). "Statistical Inference in the Multinomial, Multiperiod Probit Model." *Review of Economics and Statistics* 76, 609–632.
- Greene, W. (1990). "A Gamma Distributed Stochastic Frontier Model." *Journal of Econometrics* 46, 141–164.
- Greene, W. (2000). *Econometric Analysis*, 4th ed. Prentice Hall, Upper Saddle River, NJ.
- Horrace, W. and P. Schmidt. (1996). "Confidence Statements for Efficiency Estimates from Stochastic Frontier Models." *Journal of Productivity Analysis* 7, 257–282.
- Jondrow, J., I. Materov, K. Lovell and P. Schmidt. (1982). "On the Estimation of Technical Inefficiency in the Stochastic Frontier Production Function Model." *Journal of Econometrics* 19, 233–238.
- Kumbhakar, S. and K. Lovell. (2000). *Stochastic Frontier Analysis*. Cambridge University Press, Cambridge.
- McFadden, D. and P. Ruud. (1994). "Estimation by Simulation." *Review of Economics and Statistics* 76, 591–608.
- Meeusen, W. and J. van den Broeck. (1977). "Efficiency Estimation from Cobb-Douglas Production Functions with Composed Error." *International Economic Review* 18, 435–444.

- Revelt, D. and K. Train. (1999). "Customer Specific Taste Parameters and Mixed Logit Models." Working Paper, Department of Economics, University of California, Berkeley.
- Ritter, C. and L. Simar. (1997). "Pitfalls of Normal-Gamma Stochastic Frontier Models." *Journal of Productivity Analysis* 8, 167–182.
- Stevenson, R. (1980). "Likelihood Functions for Generalized Stochastic Frontier Estimation." *Journal of Econometrics* 13, 57–66.
- Tsionas, E. (2000). "Full Likelihood Inference in Normal-Gamma Stochastic Frontier Models." *Journal of Productivity Analysis* 13, 183–296.