Putting the "Econ" into Econometrics

Jeffrey H. Dorfman and Christopher S. McIntosh
Department of Agricultural & Applied Economics
University of Georgia

May 1998

Draft for presentation to the 1998 AAEA Meetings in Salt Lake City
Please do not cite without authors' permission

Abstract

Should econometricians always incorporate economic theory in their models or only when unrestricted estimators are found to violate an inviolable theory? Using Monte Carlo experiments, we find that econometricians should use economic theory to the fullest extent possible. To paraphrase Leamer's classic article, we should put the "econ" into econometrics.

Putting the "Econ" into Econometrics

Econometrics is supposed to be a field that combines economics and statistics; however, econometricians rarely make full use of economic theory in their work. Economic theory is frequently used in the selection of explanatory variables for a regression model, is occasionally used in the choice of functional form, and is rarely used in the actual estimation of the parameters. The paucity of economic theory in specifying functional forms should not be blamed completely on econometricians, as economic theory is often uninformative with respect to functional form; some exceptions are noted by Lopez. The same excuse cannot be used with respect to the estimation process, specifically econometricians' failure to make use of economic theory when estimating model parameters. Following Leamer's call to take the "con" out of econometrics, we now suggest it is time to put the "econ" into econometrics. In support of this position, we present evidence that by using economic theory econometricians can more precisely estimate economic models.

Econometricians often impose certain equality restrictions that derive from economic theory (e.g., homogeneity) in their estimation algorithms. However, these restrictions simplify the estimation process by reducing parameters and are easy to impose; we will not focus on this facet of economic theory as it is not an area of dispute. Economic theory often contains additional information that translates into inequality restrictions (e.g., sign restrictions, curvature conditions). Econometricians typically ignore such information during an initial estimation algorithm, check *ex post* to see if the economic theory at issue is violated by the unrestricted estimates, and if so re-estimate the model with the necessary restrictions imposed (cf. Shumway and

Alexander). That is, the econometrician never meant to allow a violation of those restrictions, but did not incorporate them initially for convenience. This results in a pre-test estimator which is known to be inadmissable from the point of view of statistical efficiency (Judge et al.). An alternative would be for econometricians to impose from the beginning of the estimation process all the restrictions based on economic theory which the econometrician is not truly willing to challenge.

Traditionally, econometricians have not imposed such inequality restrictions as a matter of course because of the difficulty of imposing such restrictions (Shumway and Alexander; Weaver). In the past, imposing such restrictions would have forced econometricians to abandon the comfort of least squares estimates and resort to restricted maximum likelihood or restricted least squares algorithms with iterative search routines required to derive the optimal estimates. However, recent advances in numerical Bayesian methods allow econometricians to incorporate inequality restrictions in an easy and straightforward manner (Geweke, 1986). With the advent of these techniques and their growing use in applied econometrics (Chalfant; Chalfant, Gray, and White; DeJong, and Whiteman; Geweke, 1988a; Hayes, Wahl, and Williams; Poirier), it seems appropriate to examine the potential benefit of relying on such estimation algorithms in the place of common pre-test approaches.

This paper examines the benefit of imposing all relevant economic restrictions in econometric estimation through a series of Monte Carlo experiments in which restricted estimates are compared to unrestricted estimates on the basis of mean squared error for certain economic measures. Econometric estimation will be conducted on a simple single equation model with independent sign restrictions on some coefficients and on systems of equations with curvature restrictions that translate into highly nonlinear restrictions on the structural coefficients. The economic measures of interest will be taken to be elasticities (or flexibilities) in all cases; that is, we will measure the mean squared error of the estimated elasticities, not the structural coefficients of each model. After looking at the relative statistical precision of the different elasticity estimates, we find that it is time to put the "econ" into econometrics and fully utilize economic theory in our estimation algorithms.

The Data Generating Models

Three different data generating models were used in this study. The first is based on Waugh's data set on potato demand. This data set consists of fifteen observations on potato prices, quantities, and income. To generate Monte Carlo data sets, a matrix of explanatory variables X (intercept, quantity, and income) is bootstrapped with 50 observations, then a vector of dependent variables y is created using the model $y_i = X_i\beta + \epsilon_i$ where β is the OLS coefficient vector estimates from the original data set and ϵ_i is a vector of i.i. normally distributed error terms with variance that will be varied through a series of experiments. For any single experiment, 1000 data sets are generated through the bootstrapping procedure.

The second data generating model is more complex and is based on a normalized quadratic profit function. The production process is specified to have three variable inputs and one output that is stochastically produced with an additive Gaussian error to represent uncertainty in production efficiency. The firm is assumed to maximize a normalized quadratic profit function with known normalized input prices and a stochastic normalized output price that is forecast to equal the past period's actual price plus a known deterministic drift. The two normalized input prices evolve as independent AR(1) processes, one stationary and one nonstationary. Data sets of 50 observations on quantities and normalized prices are generated randomly by finding the optimal input demands and planned output for the known input prices and projected output price for a single period. Actual output and output price are then determined by the addition of additive error terms to each. The prices then evolve according to their defined dynamic processes, and the next observation of data is generated. Coefficients in the model are based on those published by Lopez and the model was calibrated to loosely represent the pattern of agricultural production data over the last half century. Each Monte Carlo experiment consisted of 1000 data sets generated with a specific pair of variances for the i.i. normally distributed errors added to output price and quantity to represent the stochastics of the production and pricing processes.

The third data generating process is similar to the second, but is based on a normalized quadratic cost function. The production process is assumed to have five inputs (capital, labor, energy, materials, and services) producing one output. Four conditional input demands are solved for to minimize cost based on planned output and known input prices (the numeraire equation is not included as those elasticities can be recovered through the restrictions of symmetry and homogeneity). Input prices are simulated by AR(1) processes with drift through 50 observations meant to represent annual data; output evolves as an AR(1) with drift and trend. The parameters of the cost function and the standard deviations of the random terms in the price and output equations were calibrated to match the features of actual annual data on the chemical and allied products manufacturing sector (SIC 28). This Monte Carlo experiment again consisted of 1000 data sets. However, variances were not changed in the cost function example, but remained at the values estimated using actual data on the chemical and allied products manufacturing sector.

The Estimation Methods

The benefits of imposing economic theory are evaluated by comparing the results of models estimated by different methods. The first method is unrestricted generalized least squares estimation (equivalent here to maximum likelihood estimation for the elasticity estimators) and represents the method used most commonly by econometricians. The second method, used to impose the restrictions from economic theory, is Bayesian. The prior distribution on the parameters is taken to be a product of independent, large variance, normal distributions on the regression coefficients, an indicator function that equals one when all the restrictions are satisfied and zero otherwise, and a Jeffreys (ignorance) prior on the variance components, $p(\Sigma) \propto |\Sigma|^{-(m+1)/2}$, where m is the dimension of the covariance matrix. Importance sampling with antithetic replication (Geweke, 1988b) is employed to generate a Monte Carlo sample from the posterior distribution of the function of interest (e.g., elasticities) and the posterior mean is taken to be the optimal estimate, implying a quadratic loss function. Thus, the Bayesian restricted estimates are based on prior information that is extremely informative with respect to the economic theory (no support for estimates that violate the theory), but very mildly informative with respect to numerical magnitudes (with priors that are quite diffuse while still proper). Each restricted estimate is based on

1000 random draws (500 antithetic pairs) which was empirically determined to be enough to provide suitable numerical accuracy to the posterior means of the elasticities. For both methods, equality restrictions such as homogeneity and symmetry are imposed in all experiments through construction of the data matrices.

A third estimation method is employed in the final experiment. To provide a small piece of evidence on the relative merits of imposing economic theory using the Bayesian estimation technique, a standard restricted maximum likelihood approach is performed with data generated from the cost function model. For this experiment, elasticities are also estimated by a non-Bayesian, restricted maximum likelihood estimation algorithm.

The economic theories imposed on the three experiments are as follows. The single equation demand models are assumed by economic theory to have negative price flexibilities and positive income flexibilities. The profit function is assumed by economic theory to be convex in output and input prices; homogeneity and symmetry are imposed through construction of the data matrix on both restricted and unrestricted estimates. The cost function is assumed by economic theory to be concave in input prices; again homogeneity and symmetry are imposed in all estimations.

Empirical Results

Demand Experiments

Nine experiments were conducted on the Waugh potato data using different variance settings when generating the bootstrapped data, representing a range from 1/3 to 3 times the estimated variance of the original data. The priors on the three regression coefficients for the Bayesian restricted estimates were N(μ = 0, σ = 250) on the intercept and Half-N(0, 250) on the quantity and income coefficients with truncation to restrict the prior support to negative and positive values, respectively. The mean squared errors of the price and income flexibilities are computed for restricted and unrestricted estimates relative to the “true” flexibilities from each data set (although the true model parameters don’t change, the flexibilities vary slightly by data set with the mean of the price, quantity, and income variables). The results of these nine experiments are shown in Table 1.

These results show a general superiority of the restricted estimates with the relative efficiency of the restricted estimator increasing as the number of violations increases (in conjunction with the variance of the error terms). This result is expected; however, it is worth noting that the restricted estimator appears to enjoy an advantage over the unrestricted estimates even with as few rejections as 3 or 5 out of 1000 data sets. All of the violations of the economic theories imposed in these experiments come on the income coefficient. The restricted estimator of the price flexibility still has a MSE that is progressively smaller than the MSE of the unrestricted price flexibility as the number of rejections of the restriction on the income coefficient increases. While the relative advantage of the restricted estimator of the price flexibility is not as great as that of the restricted estimator of the income flexibility, it is still a significant advantage even though that coefficient never violates its associated restriction and the two prior distributions on these parameters are independent.

Profit Function Experiments

Using the normalized quadratic profit function data generating mechanism, the convexity of the profit function in output and input prices was imposed in the restricted estimation algorithm; no individual coefficient signs were restricted. For each of these experiments, the variances of the error terms in the stochastic output and output price processes were varied around a base value of 0.50. One variance was fixed at 0.5, while the other was set to 0.1, 0.25, 0.5, 0.75, and 1.00. Thus, ten Monte Carlo experiments were conducted, each with 1000 data sets. The (0.5, 0.5) variance pair is included twice to check precision of the results.

Results of the experiments with convexity restrictions are shown in Table 2 where only the ratio of the MSEs is shown for these experiments to save space. Out of 90 elasticities, the restricted estimator is preferred on the basis of MSE in 75 cases. Except for the first elasticity (output supply price elasticity) which always favors the unrestricted estimator, there are only five other cases where the restricted estimator is not superior. The explanation of the contrary results on the output price elasticity of supply appears to be that the restricted estimator adds positive bias to this parameter in order to ensure convexity. Overall, these results show a clear superiority by the restricted estimator.

The advantage of the restricted estimator in terms of MSE is not only pervasive, but quite large in percentage terms. For the 75 elasticities where the restricted estimator is superior the average ratio of MSEs is approximately 1.27, implying that imposing the restrictions from economic theory led to a 27 percent improvement in the MSE of the elasticity estimates. In 22 cases, the restricted estimators have MSEs that are less than half the size of the unrestricted estimators (ratios greater than 2).

The size of the two variances appears to have little effect on the results. In the cases with a very small variance in the stochastic output price process (0.10 and 0.25), the two estimators have MSE ratios that are closer to one than in general for the rest of the experiments. Other than for these two cases, no general trend is noticeable with respect to changes in the variances of the two stochastic process. This is probably because there is little correlation between the variances and the number of generated data sets which lead to unrestricted estimates that violate the economic restrictions. The number of draws in the Bayesian importance sampling algorithm which are assigned zero weights due to violating the restrictions is clearly and positively correlated with the output price variance, but not with the output variance.

Further, in about 3 percent of the simulations the Bayesian method relies on fewer than 100 draws due to large numbers of violations in the random sample. In a nonautomated situation, a Bayesian econometrician would surely add draws to the numerical approximation algorithm in cases with many zero-weight draws to ensure numerical accuracy; such a practice should further increase the efficiency of the restricted Bayesian estimator.

Cost Function Experiments

For the cost function experiments, variances for the stochastic processes used to simulate input prices and output quantities are held fixed at values estimated using actual data on the chemical and allied products manufacturing sector of the U.S. economy, not varied as in the profit function experiment. This is due to the lack of any clear pattern in the profit function experiments with respect to the amount of randomness included in the profit function data generating process. Estimation efficiency is measured on the sixteen input price elasticities of the four conditional input demand functions included in the system estimation; the elasticities for the numeraire input are not included as they are all linear functions of the sixteen elasticities which are presented.

In this final experiment, a set of 1000 cost function data sets were generated and the elasticities estimated three different ways: unrestricted generalized least squares, the Bayesian restricted estimation algorithm, and a restricted maximum likelihood estimation algorithm. This allows us to compare the advantage of fully imposing the economic theory through the Bayesian technique which ignores the part of the likelihood surface which violates the restrictions versus the restricted maximum likelihood approach which simply chooses the maximum point on the admissible part of the likelihood surface. One can think of both of these restricted estimators as Bayesian: the first (the posterior mean) is an optimal Bayesian point estimate under quadratic expected loss while the second (posterior mode) is optimal under a zero-one expected loss function (Zellner).

The results of this experiment are presented in Table 3. Three sets of MSE ratios are needed to evaluate this experiment. Table 3 presents ratios of unrestricted over restricted Bayesian, unrestricted over restricted maximum likelihood, and restricted maximum likelihood over restricted Bayesian. The first set of ratios show a very clear advantage for imposing economic restrictions on the cost function's parameter estimates. All 16 elasticity estimators for the Bayesian restricted method have smaller MSEs than the corresponding unrestricted estimators (meaning ratios greater than one). The second ratio shows that the restricted maximum likelihood estimator did not improve on the unrestricted estimator in terms of efficient estimation of the elasticities; in fact, the MSEs are essentially identical for 14 out of the 16 elasticities. That the restricted maximum likelihood estimator is not an improvement is somewhat surprising. Although 458 of the 1000 unrestricted estimators violated the concavity restriction, this restriction is not directly related to single coefficient estimates. Thus, imposing the joint restriction does not appear to improve the estimation accuracy of the individual elasticities. The third ratio shows that the posterior mean estimates derived by the restricted Bayesian method are superior to the restricted maximum likelihood estimates (equal to the posterior mode). Since MSE is compatible with the quadratic loss function used to derive posterior means, this is not completely surprising. All of the 16 ratios favor the Bayesian restricted estimator (ratios greater than one), generally by a 10-20 percent margin.

This result suggests that not only is there a benefit to imposing economic theory, but that the most benefit is gained by fully imposing the theory; that is, using the mean of the posterior distribution which has excluded support for parameter values which violate the restrictions.

Conclusions

For all three sets of experiments conducted in this investigation the Bayesian restricted elasticity estimators are superior in terms of mean squared error. The results of this study suggest that Bayesian restricted estimates frequently offer an increase in statistical efficiency. Further, we have provided evidence that Bayesian methods of imposing nonlinear and inequality restrictions are superior to standard maximum likelihood estimation approaches. Because Bayesian Monte Carlo methods now allow estimators to be computed for cases with either or both inequality and nonlinear restrictions in a few minutes of computer time for virtually all economic models, econometricians need to consider the practice of imposing such nonlinear or inequality restrictions in all cases, rather than only when they are violated by unrestricted estimates.

While most of the paper compared restricted to unrestricted estimates and did not stress that the restricted estimator was Bayesian, the final experiment presents evidence showing the fully Bayesian restricted estimator to be superior to a standard restricted maximum likelihood estimator and that the restricted maximum likelihood estimator is no improvement over the unrestricted estimator. Thus, one may need to be a Bayesian to put the “econ” into econometrics.

Table 1. Flexibility MSEs from Potato Experiments

k	Unrestricted price	Unrestricted income	Restricted price	Restricted income	Ratio price	Ratio income	# of violations
0.33	0.00643	0.00035	0.00642	0.00035	1.00190	1.00010	0
0.67	0.02346	0.00138	0.02350	0.00138	0.99853	0.99977	0
1.00	0.05583	0.00317	0.05483	0.00305	1.01826	1.03886	0
1.33	0.09706	0.00544	0.09304	0.00483	1.04316	1.12654	3
1.67	0.15910	0.00886	0.14615	0.00716	1.08858	1.23816	5
2.00	0.25664	0.01224	0.23226	0.00912	1.10496	1.34243	31
2.33	0.32095	0.01626	0.27933	0.01063	1.14897	1.53020	53
2.67	0.41905	0.02256	0.36142	0.01520	1.15946	1.48409	70
3.00	0.53118	0.02719	0.44107	0.01703	1.20430	1.59628	107

k is the multiple of the original data variance used to generate the data sets. Ratio is the ratio of the unrestricted estimator's MSE to the restricted estimator's MSE; thus values greater than one indicate superiority of the restricted estimates. The # of violations is the number of unrestricted estimates that violated the restrictions (out of the 1000 total data sets).

Table 2. Ratios of MSEs for Elasticities from Profit Function with Convexity Restriction

Variance		Elasticity										
τ^2	ψ^2	1	2	3	4	5	6	7	8	9	# Violations	p=0
0.10	0.50	0.907	2.164	1.151	2.047	2.365	1.215	1.120	1.207	1.267	90	95.6
0.25	0.50	0.908	3.031	1.211	2.614	3.094	1.261	1.139	1.237	1.267	116	101.4
0.50	0.50	0.908	2.243	1.122	2.111	2.476	1.192	1.093	1.197	1.242	102	96.6
0.75	0.50	0.909	2.036	1.086	1.923	2.217	1.155	1.068	1.151	1.318	108	96.7
1.00	0.50	0.909	2.026	1.181	1.923	2.188	1.236	1.136	1.218	1.330	121	98.7
0.50	0.10	0.997	1.003	1.004	1.002	1.002	1.002	1.003	1.002	1.007	3	0.01
0.50	0.25	0.982	1.131	1.108	1.124	1.167	1.147	1.093	1.140	1.124	99	18.6
0.50	0.50	0.911	2.528	1.146	2.310	2.724	1.240	1.117	1.223	1.227	102	90.7
0.50	0.75	0.827	2.918	1.011	2.664	2.269	0.928	0.947	0.934	1.571	77	159.9
0.50	1.00	0.763	2.712	1.354	3.128	2.969	0.866	1.063	0.878	1.682	93	223.8

τ^2 is the variance of the stochastic part of production, ψ^2 is the variance of the stochastic part of the output price's dynamic process. All values in columns 3-11 are MSE of unrestricted estimates over the MSE of restricted estimates. Column 12 (# Violations) is the number of data sets which had unrestricted estimators that violated at least one of the sign restrictions (out of 1000). The last column (p=0) is the average number of draws (out of 1000) in the Bayesian importance sampling algorithm that violated the restrictions and, therefore, received zero prior weight.

Table 3. Ratios of MSEs for Elasticities from Cost Function with Concavity Restrictions: Three Way Comparison

Elasticity	U/B	U/R	R/B
1	1.150	0.999	1.151
2	1.154	1.006	1.147
3	1.102	1.000	1.101
4	1.117	1.000	1.117
5	1.198	0.981	1.221
6	1.137	0.998	1.139
7	2.023	1.002	2.019
8	1.159	0.965	1.201
9	1.072	1.000	1.072
10	1.349	1.003	1.344
11	1.164	0.996	1.167
12	1.222	0.997	1.226
13	1.108	0.999	1.108
14	1.166	1.002	1.164
15	1.224	0.997	1.227
16	1.131	0.999	1.131

Ratios are: U/B = unrestricted MSE over restricted Bayesian MSE; U/R = unrestricted MSE over restricted maximum likelihood MSE; and R/B = restricted maximum likelihood MSE over Bayesian restricted MSE. 458 out of 1000 data sets had unrestricted estimates that violated the concavity restrictions. The average number of draws from the importance sampling algorithm which received zero weight was 602.

References

Chalfant, James A. (1993). “Estimation of demand systems using informative priors.” *American Journal of Agricultural Economics* 75, 1200-1205.

Chalfant, James A., Richard S. Gray, and Kenneth J. White (1991). "Evaluating prior beliefs in a demand system: the case of meats demand in Canada." *American Journal of Agricultural Economics* 73, 476-490.

DeJong, David N., and Charles H. Whiteman (1991). “Reconsidering ‘Trends and random walks in macroeconomic time series’.” *Journal of Monetary Economics* 28, 221-254

Geweke, John (1988a). “The secular and cyclical behavior of real GDP in 19 OECD countries, 1957-1983.” *Journal of Business & Economic Statistics* 6, 479-486.

Geweke, John (1988b). "Antithetic acceleration of Monte Carlo integration in Bayesian inference." *Journal of Econometrics* 38, 73-89.

Geweke, John (1986). "Exact inference in the inequality constrained normal linear regression model." *Journal of Applied Econometrics* 1, 127-141.

Gill, Philip E., Walter Murray, and Margaret H. Wright (1981). *Practical Optimization.* New York: Academic Press.

Hayes, Dermot J., Thomas I. Wahl, and Gary W. Williams (1990). “Testing restrictions on a model of Japanese meat demand.” *American Journal of Agricultural Economics* 72, 556-566.

Judge, George G., W. E. Griffiths, R. Carter Hill, Helmut Lütkepohl, and Tsoung-Chao Lee (1984). *The Theory and Practice of Econometrics*. New York: John Wiley and Sons.

Leamer, Edward E. (1983). “Let’s take the con out of econometrics.” *American Economic Review* 73, 31-43.

Lopez, Ramon E. (1985). “Structural implications of a class of flexible functional forms for profit functions.” *International Economic Review* 26, 593-602.

Lopez, Ramon E. (1984). “Estimating substitution and expansion effects using a profit function framework.” *American Journal of Agricultural Economics* 66, 358-367.

Poirier, Dale J. (1991). “A Bayesian view of nominal money and real output through a new classical macroeconomic window.” *Journal of Business & Economic Statistics* 9, 125-148.

Shumway, C. Richard, and William P. Alexander (1988). “Agricultural product supplies and input demands: regional comparisons.” *American Journal of Agricultural Economics* 70, 153-161.

Waugh, Frederick. V. (1964). “Demand and price analysis--some examples from agriculture.” USDA ERS Technical Bulletin 1316 (Washington DC: United States Department of Agriculture).

Weaver, Robert D. (1983). “Multiple input, multiple output production choices and technology in the U.S. wheat region.” *American Journal of Agricultural Economics* 65, 45-56.

Zellner, Arnold (1971). *An Introduction to Bayesian Inference in Econometrics*. New York: John Wiley & Sons.