Journal of Econometrics 30 (1985) 445–458. North-Holland

NON-PARAMETRIC ANALYSIS OF OPTIMIZING BEHAVIOR WITH MEASUREMENT ERROR*

Hal R. VARIAN

University of Michigan, Ann Arbor, MI 48109, USA

We consider how one might test observed choice data for consistency with optimizing models in the presence of measurement error. We derive an appropriate test statistic and conduct case study involving cost minimization behavior by electric utility plants.

1. Introduction

In several earlier papers listed in the references I have described methods for testing observed economic behavior for consistency with optimizing models. These tests have built on the work of Afriat (1967, 1972, 1976), Diewert (1973), Diewert and Parkan (1978, 1980) and Hanoch and Rothschild (1972). A defect of the methods proposed in these works is that there it seems difficult to incorporate measurement error into the analysis. The data are assumed to be observed without error, so that the tests are 'all or nothing': either the data satisfy the optimization hypothesis or they don't.

Despite this stringent nature of the tests, they may well be worth doing. Indeed, if some data pass such a test *without* resorting to any specification of measurement error one might feel more confident than usual about the veracity of the null hypothesis. (Or perhaps feel more dubious than usual about the power of the data to reveal violations of the null hypothesis.)

However, it seems that if some data fail the tests, but only by a small amount, we might well be tempted to attribute this failure to measurement error, left out variables, or other sorts of stochastic influences rather than to reject the hypothesis outright. The problem here of course is to give formal content to the phrase 'only a small amount'.

That is the goal of this paper. In the following sections I offer a general method that is, in principle, capable of measuring the magnitude of departure from the underlying model of optimizing behavior. I am able to interpret this procedure in terms of the classical statistical framework of hypothesis testing, and I provide a case study to illustrate the feasibility of the method.

*This research was supported by the National Science Foundation. I received helpful comments on earlier drafts from Bruce Hill, Peter Schmidt, and Adonis Yatchew.

0304-4076/85/$3.30

2. Stochastic considerations in non-parametric analysis

The non-parametric tests mentioned above usually take the form of asking whether there exists a solution to a certain set of linear inequalities. For example, suppose that we have n observations on the output (y_i), the factor prices (w_i) and the factor demands (x_i) for a particular firm. It is shown in Varian (1984) that the following 'Weak Axiom of Cost Minimization' (WACM) is a necessary and sufficient condition for the observed behavior of the firm to be compatible with cost minimizing behavior:

$$w_i x_i \leq w_i x_j \quad \text{for all} \quad y_i \leq y_j.$$

The interpretation of this condition is that the cost at the observed operating positions should be no greater than the cost of using any other factors capable of producing at least as much output. If there are n observations, this gives rise to n^2 inequalities that must be satisfied by the observed choices if they are to be consistent with cost minimizing behavior.

This sort of condition is easy to test, but it may be a bit more difficult to interpret. If the data do not satisfy WACM, what are we to do? The answer seems to depend on the *magnitude* of the violation. If the data fail to satisfy WACM by only a small amount, then we might well be tempted to attribute this violation to measurement error and accept the hypothesis of cost minimization.

The problem here is to give specific content to the words 'magnitude of the violation' and if possible to phrase the discussion in the formal language of statistical hypothesis testing. I will describe my progress towards achieving these goals.

Let us suppose that the observed demand for factor k in observation i, x_{ik}, is related to the 'true' factor demand z_{ik} in the following way:

$$x_{ik} = z_{ik} + \varepsilon_{ik}, \qquad i = 1, \ldots, n, \quad k = 1, \ldots, m,$$

where ε_{ik} is a random error term. We will suppose that this error term ε_{ik} is iid $N(0, \sigma^2)$. Of course other stochastic specifications are possible; but this choice is a convenient one for discussion.[1]

The null hypothesis that I wish to consider can be stated as

$$H_0\colon \quad \text{the data } (w_i, z_i, y_i) \text{ satisfy WACM.}$$

It is convenient to think of the matrix of observations (x_{ik}) as a vector with mn components which we will denote by X. Thus the non-negative orthant of

[1] In most applications, an assumption of proportional measurement error is often more appropriate. This is in fact the specification that we use in the empirical work presented below.

R^{mn} is the set of all possible data. The set of data consistent with the null hypothesis is then that subset H of R^{mn} that satisfies WACM. The observed choices, X, is not an element of this set, but under the null hypothesis, the true choices, Z, is an element of H. The situation is depicted in fig. 1.

How can we conveniently test this null hypothesis? There are several approaches that one might consider. The first approach is somewhat Bayesian in flavor. Since $z_{ik} = x_{ik} - \varepsilon_{ik}$ we might regard z_{ik} as a Normal random variable with mean x_{ik} and variance σ^2. The vector Z can then be thought of as being multivariate Normal with a probability density $f(Z)$. We can integrate the density over the region H to compute the probability that the null hypothesis actually holds.

There are three sorts of problems with this approach that I can see. First, this procedure is not in the spirit of classical statistical hypothesis testing, since we generally want to consider the distribution of the observed data given that the null hypothesis is true, not the distribution of the true data given the observations. Secondly, the test may be computationally quite demanding, especially since it may need to be performed for several different values of σ^2. Thirdly, it does not generalize in a convenient way to other sorts of non-parametric tests.

For example, if we are given observations on a consumer's choices (x_i) when facing prices (p_i), a necessary and sufficient condition for the data to be consistent with maximizing behavior is that there exist positive numbers (U_i, λ_i) that satisfy the following system of inequalities:

$$U_i \leq U_j + \lambda_j p_j (x_i - x_j), \qquad i, j = 1, \ldots, n.$$

Thus in the case of consumer maximization, the region H will be all data sets for which these inequalities have a positive solution. This may be rather difficult to calculate, much less to integrate over.

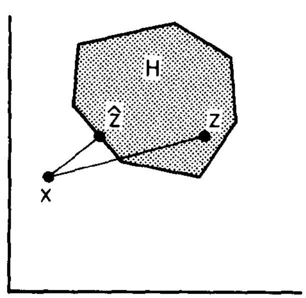

Fig. 1. Non-parametric test.

For these reasons I have adopted a different approach. There are two arguments that lead to the same test procedure, which we will examine in the following two sections.

3. A chi-squared test

Suppose that we could somehow observe the true data (z_{ik}). Then since $\varepsilon_{ik} = z_{ik} - x_{ik}$, we could compute the 'test statistic':

$$T = \sum_{i=1}^{n} \sum_{k=1}^{m} (z_{ik} - x_{ik})^2/\sigma^2.$$

Under the null hypothesis H_0, this 'statistic' has a chi-squared distribution. Thus we can find a critical value C_α for any desired level of significance, α. If $T > C_\alpha$ we would reject the null hypothesis.

The problem is of course that the 'statistic' T is not observable. However, it turns out that we can calculate an observable *lower bound* on T that will still allow us to apply the above testing method. Consider the following quadratic programming problem:

$$S = \min \sum_{i=1}^{n} \sum_{k=1}^{m} (\zeta_{ik} - x_{ik})^2/\sigma^2,$$

subject to

$$w_i \zeta_i \leq w_i \zeta_j \quad \text{for} \quad y_i \leq y_i.$$

Under the null hypothesis, the 'true data' (w_i, z_i, y_i) satisfy the constraint. Hence the minimum of the sum of squares S must be no larger than the 'test statistic' T.

Thus if we reject H_0 whenever $S > C_\alpha$, we are certain that in fact $T > C_\alpha$, and thus we have *at least* the desired level of significance. That is, the probability of rejecting H_0 when it is true will be less than α. In this sense, the test is very conservative.

The basic trick in the above method is using the mathematical programming problem to derive a bound on the unobserved random variable. The exact stochastic specification and the distributional assumptions were chosen in order to present a specific example and are not critical to the structure of the test. The choice of specification in these non-parametric methods rests on the trade-off between generality and computability here just as it does in all statistical work.[2]

[2] This idea has been used before in the statistics literature. The classic upper and lower bounds on the Durbin–Watson statistic were calculated by using a similar bounding argument.

Although I believe the hypothesis test described above is in the spirit of classical hypothesis testing, it is comforting to note that it is also quite sensible. We are simply asking for the minimal perturbation of the data that satisfies WACM. If the minimal perturbation is small relative to the amount of noise thought to be present in the data then it seems reasonable to accept the null hypothesis.

One particularly nice feature of the test outlined above is that it can handle 'nuisance parameters' quite easily. For example, in the case of the utility maximization problem described above the programming problem becomes

$$S = \min \sum_{i=1}^{n} \sum_{k=1}^{m} (\zeta_{ik} - x_{ik})^2/\sigma^2,$$

subject to

$$A_i \leq A_j + b_j p_j(\zeta_i - \zeta_j), \qquad i, j = 1, \dots, n.$$

This time the minimization takes place over the variables (ξ_{ik}, A_i, b_i). Under the null hypothesis, there is a set of variables (z_i, U_i, λ_i) that satisfies the constraints, so the answer to this minimization problem will be an appropriate bound to the desired test statistic. It appears that it would be quite difficult, if not impossible, to analyze this problem by the integration technique mentioned earlier.

A further advantage of the method is that it actually constructs the perturbation of the data that satisfies cost minimization. If we want to go on to calculate bounds on the underlying technology or to forecast demand behavior as described in Varian (1984), we can use this constructed technology. The integration approach described above lacks this feature.

There is a nice geometric interpretation of the proposed test that is depicted in fig. 1. Here the Euclidean distance between X and Z is a sum of squared residuals, and this distance in obviously bounded by the minimal distance from X to H.

However, despite these observations, there are some unpleasant features of the proposed test. The major difficulty is the fact that one needs to specify a known variance.[3] However, the fact that one must postulate a value for this parameter does not make the undertaking entirely arbitrary. For example, one could use estimates of the error variance derived from parametric fits or from knowledge about how the variables were actually measured. In any event, it seems that postulating one parameter, an indication of how noisy the investigator believes the data to be, is much less objectionable than the common practice of postulating an entire functional form. It must be remembered that

[3] If there are several observed choices at each price it may be possible to actually estimate the variance. However, this type of data is rather rare.

the usual estimates of error variances are correct only under the maintained hypothesis of the specified functional form; and this maintained hypothesis is often arbitrary.

Even if we are not able to estimate the error variance as in parametric models, we can still derive bounds on the error variance that is necessary in order to reject the maintained hypothesis of maximizing behavior. Let us consider this point in more detail.

As above, let C_α be the critical value for our proposed test, and let S be the value of our objective function. Then by inspection of the objective function, $S = R/\sigma^2$, where σ^2 is the 'true' variance of the error term and R is the sum of squared residuals. We are proposing to reject the null hypothesis when $S > C_\alpha$ which means when $\sigma^2 < R/C_\alpha$. Let us refer to the $\bar{\sigma}^2 = R/C_\alpha$ as the *critical value* of σ^2, and let $\bar{\sigma}$ be its square root. Note that $\bar{\sigma}$ is easily computable once we have solved the quadratic programming problem.

The critical value, $\bar{\sigma}$, measures what the standard error of the data would have to be for us to consider the rejection of the maximization hypothesis to be a statistically significant rejection. If $\bar{\sigma}$ is much smaller than our prior opinions concerning the precision with which these data have been measured, we may well want to accept the maximization hypothesis.

4. A constrained maximum likelihood approach

Another approach to the sort of test described above is through the method of constrained maximum likelihood. Given some observations (x_{ik}) and the specification of a normal error term we can write the log-likelihood function

$$\log L = mn \log(2\pi)/2 - mn \log \sigma - \sum_{i=1}^{n} \sum_{k=1}^{m} (z_{ik} - x_{ik})^2 / 2\sigma^2.$$

We think of (z_{ik}) and σ^2 as unknown parameters to be estimated. Under the null hypothesis, $Z = (z_{ik})$ is an element of the set H, so we can consider the *constrained* maximum likelihood estimates derived by maximizing the likelihood over the set H. It is easy to see that this gives us as our estimator for (z_{ik}) the values that solve the quadratic programming problem described above. Let these values be denoted by $(\hat{z}_{ik})$. The associated estimator for σ^2 is

$$\hat{\sigma}^2 = \sum_{i=1}^{n} \sum_{k=1}^{m} (\hat{z}_{ik} - x_{ik})^2 \Big/ mn.$$

Thus the 'fitted values' $(\hat{z}_{ik})$ are a (constrained) maximum likelihood estimate of the true unknown values. However, I am unsure whether it is possible to establish any useful statistical properties for these estimates. After all, we are estimating one more parameter that we have observations. The only thing that allows our estimates to be identified at all is the constraint.

However, the statistic $\hat{\sigma}^2$ can be used for hypothesis testing purposes. As before we simply note that

$$mn\hat{\sigma}^2/\sigma^2 = R$$

is no larger than a chi-squared variable with mn degrees of freedom. Hence if we reject the null hypothesis that Z is in H whenever $R > C_\alpha$, we are guaranteed a test of *at least* the desired level of significance.

Rearranging the test condition, we have

$$\sigma^2 < mn\hat{\sigma}^2/C_\alpha = R/C_\alpha,$$

which is exactly the condition given earlier. If our prior beliefs suggest that σ^2 is less than R/C_α, we should reject the optimization hypothesis. Otherwise, it should not be rejected.

5. A comparison with parametric methods

The diagram in fig. 1 can be used to establish a nice link with standard parametric estimation techniques in models with optimizing behavior. Suppose that we have some parametric form for the underlying technology which we can use to derive functional forms for the factor demand for factor k at observation i as a function of the factor prices, output levels, and an unknown vector of parameters β. We denote this factor demand by $g_k(w_i, y_i, \beta)$.

Then under the null hypothesis of optimizing behavior and known parametric form, the true data, (z_{ik}), will satisfy the parametric relationship:

$$z_{ik} = g_k(w_i, y_k, \beta), \qquad i = 1, \ldots, n, \quad k = 1, \ldots, m.$$

If there are b unknown parameters, the set of all Z that satisfy this relationship will generically be a b-dimensional manifold in R^{mn} which we denote by M. Since Z satisfies the optimization conditions by construction, the manifold M must be a subset of the set H. This relationship is depicted in fig 2.

The usual approach to constrained parametric estimation involves maximizing the likelihood function over the manifold M. This simply means that we find some fitted values (z_{ik}) in M that maximize the likelihood function. The associated values of β are the maximum likelihood estimates of the unknown parameters. The maximum likelihood estimate for σ^2 is just the sum of squared residuals divided by mn.

These estimates are entirely analogous to the non-parametric estimates given above. The only difference is that we are maximizing the likelihood over the mn-dimensional set H rather than the b-dimensional manifold M. Despite this

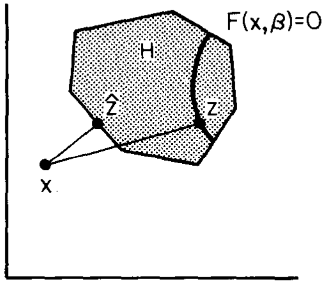

Fig. 2. Parametric test.

resemblance in method, the statistical properties of the estimators may be quite different. I suspect that the major differences arise because the dimension of the manifold M stays fixed as the number of observations increases while the dimension of H keeps growing.

Let us use R_p to denote the minimal sum of squared residuals for the parametric estimate and R_n the minimal sum of squared residuals for the non-parametric estimate. The ratio R_n/R_p is, in some sense, a measure of the 'goodness of fit' of the parametric model, conditional on the optimization hypothesis. If the perturbation of the data necessary to satisfy the optimization hypothesis in the presence of a specific parametric form is very large relative to the perturbation necessary to satisfy optimization alone, we might not find the parametric hypothesis very convincing.

Of course again we need to ask 'large relative to what?' I have not yet come up with a satisfactory answer to that question, and until someone does we will simply have to make do with an intuitive notion of what magnitudes are plausible.

6. A case study

In order to examine the feasibility of the methods described above I undertook a case study involving data on California electric power generation. The data in question were obtained from Woo (1982) and are discussed in detail in his Chapter 4. The data consist of eighteen time series observations from 1960–1979 on the factor inputs and output of two electric power generation plants.[4] The factors used in this study are labor, fuel and capital.

We first checked the plants for consistency with the cost minimization hypothesis. We found that the observed behavior *violated* the WACM inequalities. The natural question is: by how much were the inequalities violated?

[4] Woo's original data set contained two observations in which the cost of capital variable was negative. These observations were discarded.

To answer this question we solved up the quadratic programming problem described above. Each program involved fifty-four variables (three factors, eighteen observations) and around two hundred constraints. The actual program solved was different from that described above in that we postulated a proportional measurement error rather than an additive one. Since the factor demands were measured in very different units this seemed like a much more plausible specification.

Specifically, we assumed that the true demand was related to the observed demand in the following way:

$$z_{ik} = x_{ik}(1 + \varepsilon_{ik}),$$

where ε_{ik} is an iid Normal disturbance with mean zero and constant variance. This leads to an objective function of the form

$$\sum_{i=1}^{m} \sum_{k=1}^{n} (z_{ik}/x_{ik} - 1)^2.$$

The quadratic programming problem was solved by a quadratic programming package called MINOS by Murtagh and Saunders (1977). We found the cost of each quadratic programming problem was around \$2.00 during normal priority on the University of Michigan computer system, an Amdahl 5860 running MTS. The result of these calculations are presented in tables 1 and 2.

We will describe the data in table 1 since the story in table 2 is much the same. Columns 2–4 give the observed output and input levels, while columns 5–7 give the percent 'residuals' that would make the observed firm choices consistent with WACM and minimize the sum of squared deviations. The value *SSR*, the sum of squared residuals, is given for each factor and overall, as well as the square root of the overall *SSR* divided by the number of observations. The latter variable may be thought of as something like a standard error so we have denoted it by *SE*.

There are several interesting things about the 'fit' described in table 1. Note for example the size of the perturbations for the different factors. One would imagine that the labor cost component of these electric power plants would be the easiest to measure, with the fuel costs a close second. The most difficult factor to measure is certainly the capital stock. Indeed, when we look at the sum of squared residuals for each factor we find that they conform to this expected pattern.[5]

[5]Of course if the analyst thought that there were extreme differences in the degree of measurement error across the factors of production he would presumably modify the null hypothesis of equal percentage error variance. This seems to me to be analogous to the problem of heteroskedasticity in ordinary regression analysis. Just as in the case of ordinary least squares, we can minimize a *weighted* sum of squares if this problem is felt to be a serious one. In most cases, I doubt that it is.

Table 1

Alamitos.[a]

Year	Actual fuel	Actual labor	Actual capital	Residual ×100	Residual ×100	Residual ×100
1	24150	55.00	33.27	0.01	0.00	0.00
2	26624	67.00	63.24	0.00	0.00	0.00
3	50010	81.00	95.12	0.00	0.00	0.00
4	48608	67.00	90.39	0.00	0.00	0.00
5	56676	98.00	86.48	4.37	−0.15	1.94
6	52526	126.00	119.28	−4.00	−0.20	−2.68
7	88855	132.00	141.67	−1.03	−0.03	−0.53
8	97145	139.00	137.03	−6.32	−0.55	−4.42
9	95378	141.00	130.55	−3.23	−0.10	−2.74
10	98783	161.00	135.12	3.29	−0.12	−1.65
11	91613	158.00	128.72	0.49	0.24	−1.82
12	92427	156.00	122.53	−0.36	0.13	3.20
13	83392	154.00	116.54	0.00	0.00	0.00
14	90072	149.00	111.56	0.00	0.00	0.00
17	77048	140.00	99.49	0.00	0.00	0.00
18	101445	143.00	94.59	2.63	0.30	3.81
19	88150	147.00	90.22	0.00	0.00	0.00
20	111575	148.00	86.13	5.58	0.18	1.47
SSR				0.0136	0.0001	0.0710
SE				0.0275	0.0018	0.0199

[a] Overall $SSR = 0.0208$, overall $SE = 0.00196$, critical value $\bar{\sigma} = 0.0005$.

Secondly note how small the perturbations are. The largest perturbation is in the period 8 fuel usage, and this is only about 6.3 per cent. These perturbations seem to be quite small relative to my beliefs about the likely magnitude of the measurement error associated with these data.

We can be more precise about this statement. Following the discussion in section 2, I have computed the critical value of σ outlined there. The 95% critical value of a chi-squared variable with 54 degrees of freedom is about 44. This implies that $\bar{\sigma}$ is 0.0005. This means that one would have to believe the data were measured with a standard error of *less than* 0.05 percent in order to reject the null hypothesis of cost minimization. This seems a substantially smaller measurement error than anyone would be likely to attribute to these data. On these grounds I am willing to accept the hypothesis of cost minimization.

How do these data fare when confronted with standard sorts of parametric methods? To answer this question I found the minimal perturbation of the data to satisfy factor demands derived from Cobb–Douglas and CES production functions. The fit is described in table 3. Note that in the Cobb–Douglas case the sum of squared residuals is over 500 times as large as the perturbation needed to satisfy minimization alone! The CES case fares somewhat better,

Table 2

Pittsburg.[a]

Year	Actual fuel	Actual labor	Actual capital	Residual ×100	Residual ×100	Residual ×100
1	52336	100.00	110.22	0.57	0.01	0.40
2	76601	135.00	138.15	0.00	0.00	0.00
3	70065	135.00	131.62	−4.80	−0.12	−2.92
4	63000	129.00	125.07	0.00	0.00	0.00
5	69602	129.00	116.44	0.96	0.03	0.52
6	55398	134.00	110.25	2.32	−1.62	−4.75
7	70906	127.00	104.78	3.86	0.09	1.85
8	52797	132.00	99.67	3.41	1.62	3.77
9	61286	133.00	94.15	0.00	0.00	0.00
10	37792	139.00	89.42	0.00	0.00	0.00
11	44233	152.00	85.16	0.00	0.00	0.00
12	47155	157.00	81.08	0.00	0.00	0.00
13	108952	163.00	146.41	−1.01	−0.03	−0.21
14	84772	170.00	145.69	0.00	0.00	0.00
17	105566	185.00	150.17	0.67	0.47	2.51
18	105935	199.00	155.81	0.37	−0.47	−2.39
19	82745	204.00	150.76	0.00	0.00	0.00
20	86842	209.00	142.48	0.00	−0.00	0.00
SSR				0.0058	0.0006	0.0061
SE				0.0179	0.0056	0.0184

[a] Overall $SSR=0.0125$, overall $SE=0.0152$, critical value $\bar{\sigma}=0.0002$,

Table 3

Comparison with parametric method.

Plant	*SSR* Cobb–Douglas	CES	Non-parametric
Alamitos	12.359	1.0697	0.0208
Pittsburg	6.444	0.8704	0.0125

with a *SSR* only 50 times as large as that needed in the non-parametric case. Even so this perturbation seems quite large. The conclusion seems to be that the data are consistent with cost minimization – but not Cobb–Douglas or CES cost minimization.

7. Objections and replies

After describing the above method, I have often encountered various objections to it. In this section I will describe several of the most common objections and my replies.

(1) *You have specified error terms only on the quantity terms; the price terms may also be measured with error.*

I agree. It would be desirable to incorporate error terms on the prices as well. However, note that the resulting programming problem would then have non-linear constraints and thus be considerably more difficult to solve. Furthermore, note that standard regression methods typically specify that regressors are non-stochastic. If one is estimating conditional factor demand equations, this means that price and output variables are hypothesized to be measured without error, and only factor demands themselves are assumed to be measured with error – exactly as specified here.

(2) *Simply because you fail to reject the null hypothesis doesn't mean that the data were generated by optimization.*

Of course not. If we fail to reject the null hypothesis then we have simply stated that the observed departures from the model are not extremely unlikely given the null hypothesis. That is exactly what is done in ordinary statistical hypothesis testing.

(3) *It would be useful to have a measure of the power of this test.*

Absolutely. But of course the power of a test *depends* on the specific alternative hypothesis. Given some reasonable alternative and the distribution of the errors it would be possible to compute the probability that the proposed test would be satisfied. This would almost certainly have to be done by Monte Carlo methods, since I see little hope of an analytic solution. However, by way of comparison, let me note that it is quite rare that one sees power reported in parametric econometric studies.

(4) *The fact that one needs to specify the error variance effectively renders this approach worthless.*

I think not. Rather than speculate idly on this point, it is worth considering the particular example presented above. Does anyone really believe that the factor demand data described in table 1 were measured with a standard error of less than 0.05 percent? If not, the procedures outlined above indicate that the departures from cost minimizing behavior depicted in that table are not statistically significant. This seems to me to be a perfectly satisfactory statement. Furthermore, specifying the likely magnitude of the measurement error seems to me to be much *less difficult* a task than specifying a plausible functional form for a production function, as is required in the conventional approach.

And of course, one may be able to construct plausible estimates or bounds on the error variance by other means. One obvious choice is to use the standard errors that are generated by parametric methods. Moreover, Epstein and Yatchew (1984) have suggested a method of estimating σ that is consistent under certain hypotheses about the family of functional forms that could have generated the data. If the Epstein–Yatchew method works well in small samples, it may eliminate this objection entirely.

(5) *It is difficult to argue that we know enough to specify the way disturbances enter the demand equations and can specify a parametric form of the distribution while at the same time arguing that we do not know enough to specify the form for the demand system.*

I find the specification of a Normal error term much less difficult than the specification of a particular parametric form for technology or demand. In any event parametric studies usually require a specification of *both* the functional form of the demand relationship and the parametric form of the distribution.

(6) *What is the relationship of this approach to the literature on frontier estimation?*[6]

In this paper I have specified the error term as a measurement error associated with the factor demands, since in my opinion these are the variables that are the most poorly measured in this sort of study. However, one could consider alternative approaches in a non-parametric context as well.

For example, suppose that we thought that the output levels (y_i) were underestimates of the 'true' output frontier levels (y_i^*). That is,

$$y_i = y_i^* + \varepsilon_i, \qquad i = 1, \dots, n,$$

where $\varepsilon_i \leq 0$. In this case we might find a minimal perturbation of the observations that would satisfy WACM, while respecting the sign restriction on the perturbations implied by the non-positivity of the error term. Under the null hypothesis this would be a lower bound on the actual perturbation and all the analysis of this paper applies. The only difficulty with this approach is that the minimization problem does not take a standard form.

8. Summary

I have shown how one can extend the non-parametric methods described in the introduction to cases involving measurement error. The logic of the approach involves asking for the minimal perturbation of the data that satisfies the inequality relations implied by the underlying theory. This sort of test can be given an interpretation consistent with the classical theory of statistical hypothesis testing. Furthermore, the methods are practical from the computational perspective and can be applied in a wide variety of stochastic specifications.

[6]For a survey of frontier estimation, see Aigner and Schmidt (1980).

References

Afriat, S.N., 1967, The construction of a utility function from expenditure data, International Economic Review 8, 67–77.

Afriat, S.N., 1972, Efficiency estimation of production functions, International Economic Review 13, 568–598.

Afriat, S.N., 1976, The combinatorial theory of demand (Input–Output Publ. Co., London).

Aigner, Dennis and Peter Schmidt, 1980, Specification and estimation of frontier production functions, profit and cost functions, Journal of Econometrics 13.

Diewert, Erwin, 1973, Afriat and revealed preference theory, Review of Economic Studies 40, 419–426.

Diewert, Erwin and Celik Parkan, 1978, Tests for consistency of consumer data and nonparametric index numbers, Working paper 78-27 (University of British Columbia, Vancouver).

Diewert, Erwin and Celik Parkan, 1980, Linear programming tests of regularity conditions for production functions, in: W. Eichorn, K. Henn, K. Neumann, R. Shepherd, eds. Quantitative studies on production and prices (Physica-Verlag) forthcoming.

Epstein, Lawrence and Adonnis Yatchew, 1984, Nonparametric hypothesis testing procedures and applications to demand analysis, this issue.

Hanoch, Giore and Michael Rothschild, 1972, Testing the assumptions of production theory, Journal of Political Economy 80, 256–275.

Murtagh, Bruce and Michael Saunders, 1967, MINOS: large scale nonlinear programming system (for problems with linear constraints), Department of Operations Research technical report SOL 77-9 (Stanford University, Stanford, CA).

Varian, Hal, 1982a, Nonparametric methods in demand analysis, Economics Letters 9, 23–29.

Varian, Hal, 1982b, The nonparametric approach to demand analysis, Econometrica 50, 945–973.

Varian, Hal, 1982c, Trois evaluations de l'impact 'social' d'un changement de prix, Cahiers du Seminar d'Econométrie 24, 13–30.

Varian, Hal, 1983a, Nonparametric tests of consumer behavior, Review of Economic Studies 50, 99–110.

Varian, Hal, 1983b, Nonparametric tests of models of investor behavior, Journal of Financial and Quantitative Analysis 18, 269–278.

Varian, Hal, 1984, The nonparametric approach to production analysis, Econometrica 52, 579–597.

Woo, Chi-Keung, 1982, The nonparametric approach to production analysis: A case study on a regulated electric utility, Ph.D. thesis (University of California, Davis, CA).