working paper
department
of economics

AGGREGATE PRODUCTION FUNCTIONS: SOME CES EXPERIMENTS

Franklin M. Fisher, Robert M. Solow and
James R. Kearl

August, 1974

No. 136

massachusetts
institute of
technology
50 memorial drive
cambridge, mass. 02139

AGGREGATE PRODUCTION FUNCTIONS:
SOME CES EXPERIMENTS

Franklin M. Fisher, Robert M. Solow and
James R. Kearl

August, 1974

AGGREGATE PRODUCTION FUNCTIONS:

SOME CES EXPERIMENTS

Franklin M. Fisher,

Robert M. Solow

and

James R. Kearl

Massachusetts Institute of Technology

and

National Bureau of Economic Research*

* This research was supported by National Science Foundation Grant #GS-43185 to the National Bureau of Economic Research. We are indebted to Nancy Greene Burstein for programming assistance, and to Martin Osborne for research assistance.

1. Introduction

As a purely theoretical matter, aggregate production functions exist only under conditions too stringent to be believed satisfied by the diverse technological relationships of actual economies.[1] Yet aggregate production functions estimated from real data do appear to give good results, at least sometimes, and to do so in an apparently non-trivial way. Not only do such estimated relationships give good fits to input and output data, but also the calculated marginal products appear to be related to observed factor payments. Alternatively, production functions with parameters estimated from factor payments turn out to fit input and output data pretty well.[2]

It is not a simple matter to decide why this should be so as a matter of theory. Indeed, the problem is sufficiently complicated that perhaps the most promising mode of attack on it is through the construction and analysis of simulation experiments. By constructing simplified economies in which the conditions for aggregation are known not to be satisfied, we can hope to find out inductively the circumstances under which aggregate production functions appear to give good results in the double sense just discussed. Moreover, such experiments can cast light on other aspects of the estimation of aggregate production functions from underlying non-aggregable data.

---

1. See Fisher [2] for a summary discussion and bibliography.

2. A relatively early effort if Solow [5]; there are many others.

This program of research was begun in Fisher [3]; individual firms (with a single homogeneous output and single homogeneous labor but different capital types) were given different Cobb-Douglas production functions, underlying capital and labor data were generated in various ways, and labor was assigned to firms to maximize output. An aggregate Cobb-Douglas production function was then estimated and its wage predictions examined.

A number of subsidiary results were found in these experiments, and we shall comment on some of them below. The principal conclusion, however, was the following. It is obvious, of course, that an aggregate Cobb-Douglas production function which fits input and output data well cannot give good wage predictions unless labor's share of total output happens to be roughly constant. What is not obvious is that such an aggregate production function will turn out to give good wage predictions even if rough constancy of labor's share happens to hold. Yet this is overwhelmingly the case in the experiments reported in [3].

This observation has occasionally been misinterpreted, so we dwell on it here. To say that an estimated Cobb-Douglas function fits input and output data well is to say that it will predict output accurately, given the inputs. Any Cobb-Douglas function will predict that the real wage is proportional to calculated output per worker (and therefore to actual output per worker if it fits well). Therefore, a Cobb-Douglas that fits input-output data well cannot hope to predict wages accurately unless the observed share of wages is roughly constant. If the observed wage share happens to be roughly constant, the fitted Cobb-Douglas can still go wrong, if and only if the elasticity of output with respect to labor input, as estimated from input-output data, deviates from the actual wage share by a lot. The content of Fisher's observation about the earlier experiments is that, whenever the simulated economy exhibited roughly constant factor shares, the Cobb-Douglas elasticity estimated from input-output data was close to the observed wage share, and thus that predicted wages were close to actual wages.

If the same thing should happen in empirical work with real data, it requires interpretation. One cannot simply conclude that under such-and-such a circumstance, the Cobb-Douglas gives a good representation of the "underlying aggregate technology" because it gives a good representation of wages. This is because the simulation examples were generated in just such a way that there is no "underlying aggregate technology" to be represented. No one seriously supposes the situation to be different in real economies. The next thought is likely to be that the estimated aggregate functions work as good approximations only so long as all variables move roughly together. The results of the Cobb-Douglas experiments of [3] suggest that aggregate production functions will work in wider circumstances than that. In particular, an aggregate Cobb-Douglas gives good results whenever labor's share is relatively constant, even though there is quite a lot of relative movement in the underlying variables.

In any case, it is obviously of interest to know the extent to which the results of [3] are limited to a highly simplified case in which not only the aggregate production function estimated is Cobb-Douglas but so are all the underlying micro-production functions. Indeed, there is some reason to suppose that such a limitation might hold because the heuristic argument given in [3] as a partial explanation of the principal result appears heavily dependent on the Cobb-Douglas form of the underlying production functions being aggregated.

The present paper generalizes the experiments of [3] in two obvious ways. First, the underlying production functions of firms are not restricted to being Cobb-Douglas but are taken to be Constant-Elasticity-of-Substitution (CES), functions with the parameters chosen in various ways described below. Second, not only is an aggregate Cobb-Douglas production function estimated and its wage predictions examined, but an aggregate CES production function is also estimated and analyzed.

In general, it turns out that the principal result of [3] continues to hold in this more general context. The aggregate Cobb-Douglas production function predicts wages well whenever labor's share is roughly constant and the relationship is quite close. We have not found any similar organizing principle with which to explain when the aggregate CES production function does or does not give good wage predictions, but this may possibly be due to the fact that the wage predictions of the aggregate CES are generally very good in these experiments, almost always being better than the predictions of the aggregate Cobb-Douglas, even in those cases in which distributive shares are "roughly constant".[3]

In addition to verifying the principal result of [3], however, the present set of experiments is sufficiently rich as to suggest a number of conclusions concerning the estimation of aggregate CES production functions and a good deal of attention will be paid to these results below.

---

3. In particular, the conjecture of [3] that what matters is how closely wages are a log-linear function of output per man is unverified.

2. The Experiments

We worked with underlying units which we shall call "firms". Each firm is distinguished by having its own CES production function. There were N firms. In general we took N = 3, but a set of runs in which N = 5 did not present any very different features. It is important to recognize, however, that the term "firms" is somewhat misleading. Since each firm's production function exhibits constant returns to scale, two firms with the same production function are equivalent (with efficient allocation of labor) to one big firm. Hence, varying the number of true firms in the economy can take the form of varying initial capital stock rather than simply varying N. In effect, N is the number of different "industries" if we suppose that all firms in an industry have the same production function.

Note also that N cannot be varied with other things equal, because variations in N inevitably involve changes in the distribution of the parameters of the production function. Changes in the number of firms that do not involve changes in that distribution are equivalent to changes in firm size rather than in N.

Each of the N firms had a CES production function:

$$Y_f = \gamma \{ \delta_f K_f^{-\rho_f} + (1 - \delta_f) L_f^{-\rho_f} \}^{-1/\rho_f} \tag{2.1}$$

where f = 1, ..., N is the firm subscript; δ_f and ρ_f are parameters; K_f is the fth firm's capital stock; L_f is the amount of labor assigned to the fth firm; and Y_f is the resulting output. The time subscript is omitted.

In each experiment, the elasticity of substitution $\sigma_f = \dfrac{1}{1 + \rho_f}$

and the distribution parameter δ_f were chosen for each of the N firms, as described below. Then a twenty-year time series for each capital stock and for total labor was generated. The labor market is cleared in each period by finding the real wage (in terms of the common product of firms) that equates the sum of all firms' competitive demands for labor to the exogenously given inelastic supply, L. This process determines not only the real wage, $\bar{w}$, but also each firm's employment, L_f, its output, Y_f, and the marginal product of its own type of capital, r_f. Since the competitive process equalizes the marginal product of labor across firms, clearly aggregate output $Y = \Sigma_f Y_f$ is maximized in each period subject to the constraint that $\Sigma_f L_f = L$. We denote the maximized output by Y^*.

In general, the marginal rates of substitution between different capital types vary in the efficient production of total output over time; otherwise there would be no aggregation problem. For each capital type, we compute the average marginal product as:

$$\bar{r}_f = \sum_{t=1}^{20} r_f(t)\, K_f(t) \Big/ \sum_{t=1}^{20} K_f(t) \tag{2.2}$$

and an index of aggregate capital stock, J (t), as:

$$J(t) = \sum_{f=1}^{N} \bar{r}_f\, K_f(t) \tag{2.3}$$

We then estimated both an aggregate Cobb-Douglas production function:

$$Y^* = AL^{\alpha} J^{\beta} \tag{2.4}$$

and an aggregate CES production function:

$$Y^* = \gamma \left(\delta J^{-\rho} + (1-\delta) L^{-\rho}\right)^{-\lambda/\rho} \tag{2.5}$$

by least squares assuming a multiplicative error term e^u.[4] Note that we have

---

4. Equation (2.4) was estimated by the usual logarithmic regression. The estimation procedure used with (2.5) is described more fully in the Appendix, where other computational details are also reported.

not imposed constant returns to scale in either case despite the fact that the underlying data are generated from a constant-returns economy. It turns out to matter whether or not we impose such a constraint, as we shall see when considering the results.

In both cases, we examined the wage prediction obtained for each period as the marginal product of labor in the estimated aggregate production function at the observed L (t) and J (t), and compared those predictions with the actual wages, w.

In addition, an aggregate constant-returns CES production function implies a log-linear relationship between wages and output per man. Moreover, that relationship implies an estimate of the elasticity of substitution, $\sigma = \frac{1}{1 + \rho}$ since:

$$\text{(2.6)} \qquad \text{Log } (Y^* / L) = H + \sigma \text{ Log } w$$

where H is a combination of parameters. Hence, there is another way to estimate the elasticity of substitution for an aggregate CES in addition to estimating it directly from input and output data and we regularly estimated σ by least squares applied to (2.6), which we shall call the "wage equation". Note that estimation of (2.6) does not require the explicit construction of an aggregate capital stock, although the theory which leads to it implicitly supposes that there is such an aggregate.

We must now describe how the parameters were chosen and the data generated.

For each set of experiments, the elasticities of substitution, σ_f, were chosen as indicated in Table 1 below which summarizes the various choices of parameters made. The distribution parameters, δ_f, were chosen in a more complicated way. Except in case H, where all distribution parameters were set at .25, the distribution parameters were first chosen to divide the range of .15 to .35 evenly. (The intent was to generate a share for labor approximately .75, this being roughly labor's share in U.S. output). For each choice of the elasticities of substitution, the distribution parameters were chosen in two sets, half the runs having distribution parameters and substitution elasticities positively correlated and half of them negatively correlated. All this is reported in detail in Table 2.

Since labor's share is a complicated function of the distribution parameters, we adjusted the initial choices for those parameters if the average wage share for the time series was not within the range .65 - .85. We did this by multiplying all the distribution parameters by the ratio of .75 to labor's average share, repeating the process until the average wage share for the time series was within the indicated range. However, such adjustments were only required in Cases D and G.

The series for total labor was generated in all experiments as:

$$L(t) = \mathrm{Exp}\,(.02t + .014\,\varepsilon_t) \tag{2.7}$$

where ε_t is a standard normal deviate.

The capital stock series were generated in a more complicated way. For each case (other than H) indicated in Table 1, there were 22 experiments. Half of these had distribution parameters positively associated with elasticities of substitution and the other half had a negative association. Within each half case, the capital stocks were chosen in eleven different ways, being generated from the equation:

$$K_f(t) = \mathrm{Exp}\,\{\theta_f^0 + \theta_f^1 t + .03\,\eta_{ft}\} \tag{2.8}$$

where η_{ft} is a standard normal deviate (chosen independently in each time period and each run), and the two sets of parameters, θ_f^1 were chosen as follows:

(2.9) a) $N = 3,\ \theta_f^0 = -1, 0, -1$ b) $N = 5,\ \theta_f^0 = -1, \frac{1}{2}, -1$

(2.10) a) $N = 3,\ \theta_f^1 = .04 + (2-f)\mu$ b) $N = 5,\ \theta_f^1 = .04 + \frac{1}{2}(3-f)\mu$

where

(2.11) $\mu = -.05, -.04, \ldots, 0, \ldots, +.04, +.05$

determining eleven different runs.

In other words, with the firms numbered in increasing order of elasticity of substitution, the middle firm starts off with the largest capital stock, the firms at either end having initial capital stocks of 1/e as much (except for the random disturbance). The center firm grows at a rate of 4% (plus a random term) while the other firms grow (or shrink) relative to that with the rates depending on the value of μ. Thus, for example, with three firms, when $\mu = -.05$, the trend terms are -.01, +.04 and +.09 while when $\mu = +.05$, they are +.09, +.04 and -.01. For N = 5, the trend terms cover the same range as for N = 3, with the additional two firms interpolated, in effect. Of course, the random term is far from negligible.[5]

Notice that the stream of investment for each firm is chosen arbitrarily as a distributed trend. We have incorporated no tendency for profitable firms to add to their capital faster than relatively unprofitable firms. There is no strong theoretical reason to believe that endogenous investment of this sort would make aggregation easier, although it does not seem implausible to conjecture that approximate equalization of marginal products across types of capital might in practice tend to make aggregation conditions more nearly satisfied. That may be a subject for further work. In any case, in these

---

5. There have to be some departures from pure trend in the model if perfect collinearity is to be avoided in the aggregate regressions.

experiments, when one of the estimated aggregate production functions works well that outcome is unlikely to be the result of some hidden long-run efficiency consideration of the type just described.

3. Summary of Results

(a) General Remarks:

It will be remembered that each "case" in Tables 1 and 2 consist of 22 "runs". A case is characterized by the parameters of the individual firms' CES production functions. Within each case, the runs fall into two groups of eleven each. In one group of eleven, the firms with the larger elasticities of substitution have the larger distribution parameters: in the other, a large elasticity of substitution goes along with a small distribution parameter. Within each group of eleven, the runs differ according to the relative growth rates of the capital stocks of the different firms.

We then fit Cobb-Douglas and CES production functions to the aggregate time series of each run. We remind the reader that the elasticity of substitution in these production functions is an "estimate" of nothing; there is no "true" aggregate parameter to which it corresponds. Of course, comparison of the calculated aggregate elasticity of substitution within the range of firms' elasticities of substitution, or with their average is irresistible. We comment on this kind of comparison in due course.

Now each estimated aggregate production function gives us a series of 20 wage predictions for each run. In each run, the predicted wage can be compared with the "true" market equilibrium wage series. As a measure of accuracy, we calculate the relative root mean square error, i.e. the root mean square deviation of predicted wage from "true" wage, expressed as a proportion of the average "true" wage over the 20-year period.

In addition, we estimate the wage equation (2.6) from the data for each run. This gives us another estimate of the (non-existent) aggregate elasticity of substitution. The wage equation, introduced in Arrow, Chenery, Minhas and Solow [1], is sometimes used in empirical work as a vehicle for estimation of the elasticity of substitution in the absence of complete data on inputs and outputs. We shall comment on the relation in our artificial data between the estimate obtained from the wage equation and that obtained from the direct estimation of the production function.

In our data, as apparently in much empirical work on production functions, the elasticity of substitution is poorly determined in the sense that the residual sum of squares has a very flat minimum around the best estimate. On the other hand, it is our experience that a small variation in the estimated elasticity of substitution can make a substantial difference to the predicted wage series, even while it make only a trivial difference to the ability of the estimated production function to track aggregate output. This suggests a hybrid estimate of the aggregate production function: we impose the elasticity of substitution estimated from the wage equation on the production function, and use the input-output data only to estimate the distribution and efficiency parameters γ and δ in (2.1). Naturally, this hybrid production function gives us yet another series of wage predictions for each run. As will appear, these hybrid wage predictions are uniformly the best of the lot.

This outcome is, in a vague sense, expectable. The "true" wage series has played a part in the prediction process, because the elasticity of substitution has been lifted from the wage equation. When you think about it, however, the role of the true wage series is rather peripheral: it gives rise to an elasticity of substitution which is grafted on to a production function whose other parameters are estimated without benefit of the wage series, and which is then differentiated to give wage predictions. We confess to some surprise that this hybrid prediction is so consistently and so substantially superior to the others.

Before giving a systematic summary of our "empirical" results, we make two general statements that we shall not bother to document in detail because they lead nowhere in the present series of experiments.

First, in all of our runs, the fit of the aggregate production function, as measured by R^2, is very good. This goes for both Cobb-Douglas and CES functions, estimated from the same data. An attempt to make the fit deteriorate by introducing enough random variation into the exogenously-generated investment and labor-supply series ran into difficulties with the labor-allocation algorithm used. We have not pursued this line because our main interest is not in goodness of fit *per se*, but rather more in wage predictions. Since all of the production functions track aggregate output extremely well, we have not tried to use third and fourth decimal place differences in R^2 as an explanation of differences in the ability to predict wages.

Second, as already mentioned, we did some experiments in which the degree of homogeneity of the fitted production function is determined freely in the estimation process. (Remember that the firms' underlying production functions all have constant returns to scale). We found that, with very few exceptions, the estimated degree of homogeneity is close to unity, say within the range 0.95 - 1.05, and often closer than that. Nevertheless, free estimation of returns to scale is not innocuous. If constant returns to scale is imposed, the other parameters change only slightly, and the goodness of fit deteriorates only slightly, but the wage predictions usually become considerably better. Presumably such small parameter changes have negligible effect on the ability of the estimated function to track aggregate output, but a larger effect on the slopes of the function, and therefore on its wage predictions. It is, of course, interesting that giving the estimation process correct information about the underlying returns to scale actually improves the ability of the (non-existent) production function to predict market wages. It could, in principle, have gone the other way.[6]

(b) Wage Predictions:

We have (see Table 1) 184 runs altogether. Each run gives us an estimated Cobb-Douglas function and two CES functions (one from input-output data only, the other using the elasticity of substitution from the wage equation). Figures 1, A-C show the frequency distributions of the relative root mean square errors of the wage predictions over all runs in the form of a histogram for each of the production functions. (The hybrid estimate is labelled CESW).

Our first observation is that the errors are generally small. The relative root mean square error in predicting the period-by-period market wage is often less than 1%, usually less than 2%, and rarely gets as high as 5%. This observation is of limited value, naturally, because it refers entirely to our model world inside the computer and tells us nothing about practical experience. It does, however, help to calibrate our findings.

A second observation is that when all runs are taken together, the CES function out-performs the CD function, and is in turn quite substantially out-performed by the hybrid CESW function. This conclusion can be read from the medians plotted in Figure 1. The superiority of the CESW predictions is rather remarkable.

---

6. A similar result, however, was found in Fisher [3].

In this respect, Case I is especially interesting: the three firms have elasticities of substitution of 0.95, 1.005, and 1.06 respectively, so they are all essentially Cobb-Douglas, but with different distribution parameters. There are only 11 runs in this case: the CES has better wage predictions than the CD in 4 of them, 2 or 3 of the runs are essentially ties, and the CD is better in the remaining runs. But the CESW wage predictions are the best of all in every one of the 11 runs (although it is essentially tied with CES in one of them). We remark here -- in anticipation of more extensive discussion below -- that the eleven estimated CES functions in this case have elasticities of substitution ranging from 0.49 to 1.93 and averaging 0.99. The eleven wage equations give estimates of the elasticity of substitution ranging from 0.78 to 1.49, and averaging 1.13. Thus, in this case, where the firms all have elasticities of substitution near unity, both methods of estimation of the (non-existent) aggregate elasticity put it near unity on average, though with considerable range from run to run. It is characteristic, as we shall mention later, that estimation of σ from the wage equation provides a narrower range than does direct estimation of a production function.

A third observation, with respect to the CD estimates, is that Fisher's earlier finding is confirmed in these simulations. The estimated aggregate Cobb-Douglas gives good wage predictions when the "observed" wage share is fairly stable, and poorer predictions (as it must) when it is not. As remarked earlier, this means that when the wage share is approximately constant, the CD estimation yields a labor-exponent near the observed wage share. In fact, the accuracy of the CD wage-predictions diminishes rather smoothly as the wage share becomes more variable. This is demonstrated in Figure 2, which is a scatter diagram of the relative root mean square error of the CD predictions against the standard deviation of the "observed" wage share, with each observation representing a 20-period run.

Our fourth observation is a non-observation. We have tried to find some quantity that would "naturally" be correlated with the accuracy of the CES wage predictions, but we have failed. Accuracy is not correlated with the variability of the wage share, but there is no intuitive reason to expect it to be. We wondered if the accuracy with which the fitted CES function predicts wage rates might be correlated with the goodness of fit of the wage equation in the same run (as conjectured in [3]) because they are both indicators of the extent to which the data behave as if they came from an aggregate CES production function. But in fact there was no such correlation to speak of. We did not try to correlate the accuracy of wage prediction with the goodness of fit of the production function, because the variation in R^2 is so small. And we did not try to correlate the accuracy of the CESW wage predictions with anything, because the predictions are almost all very accurate. Indeed, it is quite possible that our failure to find an organizing principle to "explain" the relative accuracy of CES wage predictions in different runs fails simply because almost all the predictions involved are pretty accurate to begin with.

(c) <u>The Aggregate Elasticity of Substitution:</u>

Each of our cases is characterized by six numbers: three elasticities of substitution and three distribution parameters. Recall that the 22 runs within each case fall into two groups of 11 each: in runs 1-11, the elasticities of substitution and distribution parameters are assigned to firms so that the highest value of one goes with the lowest value of the other, while in runs

12-22 the high elasticity of substitution goes with the high distribution parameter, and low with low. Now fix attention on one of these subgroups of 11 runs. These runs differ among themselves in one respect, namely the pattern of investment (and output) across the three firms. In the low-numbered runs within any subgroup, the firm with the low elasticity of substitution is growing fastest and the one with the high elasticity of substitution is growing slowest. In the middle run (numbered 6 or 17) all three firms are growing at about the same rate. And in the higher-numbered runs, the firm with the highest elasticity of substitution grows the fastest. We emphasize this pattern here, because it turns out to be a very important determinant of the behavior of the estimates of the aggregate elasticity of substitution.

Since the technology has been chosen so that no exact aggregate production function exists, there is no true aggregate elasticity of substitution with which the various estimates can be compared. One would like to know, however, whether the estimated aggregate elasticity at least falls within the range of of the single-firm elasticities that generated the data. Table 1 shows, for each subgroup of 11 runs, the underlying single-firm elasticities of substitution, the lowest, average and highest estimated aggregate elasticity from the direct production function, and the lowest, average and highest aggregate elasticity estimated from the wage equation. (The run number at which the highest and lowest estimates occur is shown in parentheses). We have a variety of inductions to make from these experiments.

First, if one looks only at subgroup averages, the estimated aggregate elasticity of substitution is clearly in the ballpark, whether it comes from the production function or from the wage equation. That is to say, the (subgroup average) aggregate elasticity of substitution is inside the range defined by the lowest and highest single-firm elasticities of substitution. Moreover, these (theoretically illegitimate) aggregative methods do succeed in distinguishing a situation like case B from one like case C, i.e. an industry whose firms have moderately weak possibilities of factor substitution from one whose firms can substitute rather more freely. A fortiori, cases D and E are even more sharply distinguishable.

Second, however, if one looks inside the subgroups, the range of estimated aggregate elasticities of substitution is very wide. Indeed, many of them do lie outside the range defined by the single-firm elasticities. To see more clearly what is happening in the various experiments, we must distinguish between the subgroup of runs 1-11 and other subgroup comprising runs 12-22. In the course of run 1, the firm with the lowest elasticity of substitution and the highest distribution parameter is growing relative to the others: in run 11 it is the firm with the highest elasticity of substitution and the lowest distribution parameter. Within this subgroup it is an excellent generalization that the estimated elasticity of substitution is higher for higher-numbered runs. That is to say, when output is shifting to firms with high (low) elasticities of substitution, the estimated aggregate elasticity tends to be high (low). In the extreme runs (say 1,2,10,11), when the firm with high (low) elasticity of substitution is growing much faster than the others, the estimated aggregate elasticity is likely to fall above (below) the range of single-firm elasticities. If this process were to go on long enough, the most rapidly growing firm would come to dominate the model economy and presumably the aggregate production function and the wage equation would come (properly) to estimate the characteristics of that firm. But it appears that during the long interval before that process is complete, the effect of shifting input and output from firms with high elasticity to firms with low elasticity is to make the aggregates look as if the elasticity of substitution were lower still. This is an insight of unfortunately little practical use: in any real case, if one knew enough to expect such shifts one would probably have access to information which would make it unnecessary to resort to highly aggregative methods. Still, it suggests that aggregate estimates of the elasticity of substitution can be very badly thrown off by systematic compositional effects.

Third, these observations have to be modified when we look at runs 12-22. In some cases (B and C) one sees the phenomenon just described, but considerably attenuated. In other cases (A and D) it is reversed. We think this can be understood in the following way. In all of our runs, "capital" accumulates faster than the supply of labor, and so the real wage rises. An interpretation of events using an aggregate production function will certainly want to say that a small wage increase is evidence of a high elasticity of substitution, given the evolution of labor supply and "capital". Now in run 11, for instance, as time goes on the firm with the highest elasticity of substitution and lowest distribution parameter (measuring labor intensity) bulks larger and larger in the industry. On both counts the increasing demand for labor is tempered and the rise in real wages is gentle. So the data suggest a very large elasticity of substitution indeed. In run 22, on the other hand, the rapidly growing firm has a high elasticity of substitution but also a relatively high labor intensity. These two factors have opposite effects on the demand for labor and thus on the course of the equilibrium real wage. So it perhaps is not surprising that the results are variable from case to case.

In cases F and G (for which only runs 1-11 exist) all firms have the same elasticity of substitution, so the increase in estimated elasticity with run number must be due entirely to the effect of the distribution parameter. In case H, on the other hand, all firms have the same distribution parameter, so the outcome is entirely due to the different elasticities of substitution.

Fourth, the effect we have been describing is present in both methods of estimating the aggregate elasticity of substitution, but it is considerably attenuated in the estimates stemming from the wage equation and considerably more pronounced in those coming from the production function. In fact, generally speaking, the wage equation estimates seem more reliable than the production function estimates (though one must remember that the notion of "reliability" is not terribly well defined in these experiments).

Finally, we repeat here an observation made earlier: the elasticity of substitution -- a curvature parameter -- is not well-determined by the data. The other parameters seem to be able to compensate easily for variations in the elasticity of substitution when it comes to tracking the input-output data. But the wage predictions are more seriously affected. The changes in the distribution and efficiency parameters required to compensate for vagaries in the elastcity of substitution seems to distort the calculated marginal products significantly.

(d) The Distribution Parameter:

Less general interest attaches to the estimates of "the" distribution parameter δ (which also doesn't exist as a true aggregate parameter). Nevertheless, we report some observations here because they may be useful in interpreting the results of econometric work on real data. Table 2 shows actual and estimated distribution parameters. In most of the cases, the three firms have distribution parameters equal to 0.15, 0.25 and 0.35. These were chosen so that aggregate quasi-rents would average about 1/4 of total output and aggregate wages about 3/4 of output over the 20 periods. (Of course, the functional distribution of income changes in the course of each run, depending on the inter-firm pattern of growth and its interaction with the firms' technological parameters.) Cases D and G have higher values of δ because they have low elasticities of substitution; for a fixed capital-labor ratio -- so long as it exceeds one -- a lower elasticity of substitution lowers the capital share, and a higher value of δ is needed to pull it back up to 1/4. Case H was designed to give all firms the same distribution parameter.

It is evident from Table 2 that δ is poorly "estimated" in the sense that the estimated distribution parameter in the aggregate CES function may easily lie outside the range of the firms' distribution parameters. This occurs both with the directly estimated CES functions and with the CESW estimates using the wage equation. For cases B, C, D, E, F and G, even the average over the various cases is far from the center of gravity of the microeconomic δ's.

In the columns devoted to Cobb-Douglas estimates, the distribution parameter is simply the estimated elasticity of aggregate output with respect to aggregate "capital". The various cases illustrate a generalization made earlier. In cases like B and D, in which the three firms have elasticities of substitution well below unity, the aggregate Cobb-Douglas underestimates the share of capital, overestimates the share of labor, and therefore gives poor wage predictions. In cases C, E and F all firms have elasticities of substitution above unity, and the Cobb-Douglas overestimates the share of capital. The robustness of the

Cobb Douglas comes from cases like A, G, H, I, the first half of J and the second half of K, in which the Cobb-Douglas manages to estimate a capital-elasticity near 0.25, although the estimating process does not use factor-price data.

More can be said about the estimates of δ, however. Figure 3A is a scatter plot of estimates of δ against estimates of σ for all 184 CES estimates described in Tables 1 and 2. Figure 3B does the same thing for the 184 CESW estimates. Obviously, there is a strong curvilinear relation between estimates of δ and estimates of σ. A run yielding a high estimate of the elasticity of substitution tends to yield a high estimate of the distribution parameter. (By the way, the scatters for individual cases look just like pieces of these combined scatters; there is no misleading composition effect).

One can understand in a loose way why this should happen. Present a least-squares program with time series labelled output, labor and capital, and instruct it to estimate a production function under constant returns to scale. When will it tend to generate a large elasticity of substitution? A natural answer is: when output per worker rises rapidly as capital per worker increases, i.e. when the effect of diminishing returns appears to be mild. One way of making this intuition more precise is to calculate from (2.1) that the elasticity of output per worker with respect to capital per worker (k) is $\delta^{k-\rho}/ (\delta k^{-\rho} + (1-\delta)) = \phi$, say. Next, one can easily verify that $\delta\phi/\delta\delta > 0$ and that $\delta\phi/\delta\rho < 0$ provided that capital per unit of labor exceeds one, as is the case in our model economy with our artificially constructed "capital". Since $\sigma = \frac{1}{1+\rho}$, $\frac{\delta\phi}{\delta\sigma} > 0$ under the same condition. Thus, if output per worker seems to increase freely with capital per worker, the least-squares program can register this fact either by imputing a high elasticity of substitution, or a high distribution parameter. Not surprisingly, it does both of those things, in proportions which must depend on finer details of the artificial time series,

4. Conclusions

There are situations in which aggregate production functions will work (in the sense of the first paragraph of this paper) without any assistance from formal theorems on aggregation. In particular, we had uniformly good results from the hybrid procedure that estimates the elasticity of substitution from the wage equation and other parameters of the production function from input-output data. The resulting production function tracks output very closely, and its partial derivatives capture the behavior of factor prices quite well. Even the straightforward CES estimates performs adequately most of the time, as does the Cobb-Douglas, if only factor shares are not changing drastically. On the other hand, the estimated parameters themselves are sometimes quite far from anything one could sensibly describe as roughly characterizing the real -- i.e. the model -- world. The aggregative data themselves do not tell you very clearly whether the estimated parameters are likely to have average meaning or not.

For many problems, aggregate production functions are simply too useful to pass up, especially since they can work, as our experiments show. Our parting advice is to handle them the way the old garbage man tells the young garbage man to handle garbage wrapped in plastic bags of unknown provenance: " Gingerly, Hector, gingerly."

TABLE 1.

Actual and Estimated Elasticities of Substitution

CASE	N	FIRM ELASTICITIES OF SUBSTITUTION σ_f	RUN #	$\hat{\sigma}$: CES* High	Low	Average	$\hat{\sigma}$: CESW* High	Low	Average
A	3	.6 1.13 1.66	1-11	17.18 (10)	.25 (1)	3.12	2.43 (11)	.52 (1)	1.3
			12-22	2.18 (14)	.60 (13)	1.12	1.21 (16)	1.00 (22)	1.09
B	3	.6 .75 .9	1-11	2.92 (9)	.22 (1)	1.08	1.28 (11)	.47 (1)	.86
			12-22	1.44 (21)	.49 (15)	.89	.92 (19)	.64 (12)	.79
C	3	1.11 1.385 1.66	1-11	8.43 (10)	.47 (1)	2.43	2.79 (11)	.86 (1)	1.7
			12-22	4.43 (18)	.73 (12)	1.98	2.07 (21)	1.18 (12)	1.62
D	3	.3 .45 .6	1-11	1.59 (11)	.33 (2)	.8	.80 (11)	.23 (1)	.5
			12-22	1.00 (16)	.26 (21)	.56	.48 (12)	.37 (22)	.43
E	3	1.66 2.495 3.33	1-8	5.81 (8)	.64 (2)	2.15	4.84 (8)	1.33 (1)	2.5
F	3	1.66 1.66 1.66	1-11	3.75 (11)	.62 (2)	1.63	3.35 (10)	1.23 (1)	2.17
G	3	.6 .6 .6	1-11	1.00 (9)	.36 (1)	.66	.76 (11)	.48 (2)	.62
H	3	.6 1.13 1.66	1-11	3.65 (10)	.45 (1)	1.52	1.54 (11)	.69 (1)	1.08
I	3	.95 1.005 1.06	1-11	1.93 (6)	.59 (1)	.99	1.49 (10)	.78 (1)	1.13
J	3	.25 .95 1.65	1-11	4.37 (11)	.29 (2)	1.5	2.19 (11)	.30 (1)	1.02
			12-22	1.64 (13)	.27 (22)	.83	1.12 (12)	.47 (22)	.71
K	5	.6 .865 1.13 1.66	1-11	12.11 (6)	.10 (1)	3.3	3.2 (9)	.18 (1)	1.79
			12-22	4.79 (12)	.44 (22)	1.44	1.51 (12)	.67 (22)	1.05

* Corresponding run numbers are in parentheses.

TABLE 2.

Actual and Estimated Distribution Parameters

CASE	N	FIRM DISTRIBUTION PARAMETER δ_f	RUN #	$\hat{\delta}$: CES High	$\hat{\delta}$: CES Low	$\hat{\delta}$: CES Average	$\hat{\delta}$: CESW High	$\hat{\delta}$: CESW Low	$\hat{\delta}$: CESW Average	$\hat{\delta}$: CD High	$\hat{\delta}$: CD Low	$\hat{\delta}$: CD Average
A	3	.15 .25 .35	1-11	.51	.006	.28	.43	.09	.28	.31	.23	.27
			12-22	.39	.13	.25	.29	.25	.28	.26	.24	.25
B	3	.15 .25 .35	1-11	.41	.0008	.18	.29	.04	.16	.23	.18	.21
			12-22	.29	.07	.16	.19	.09	.14	.21	.19	.20
C	3	.15 .25 .35	1-11	.51	.09	.33	.46	.23	.36	.33	.27	.29
			12-22	.47	.20	.35	.42	.31	.37	.31	.27	.28
D	3	.30 .38 .46	1-11	.31	.007	.13	.17	.0006	.08	.22	.15	.19
		.29 .37 .46	12-22	.18	.04	.06	.04	.01	.03	.19	.16	.18
E	3	.15 .25 .35	1-8	.50	.21	.39	.50	.37	.44	.34	.32	.33
F	3	.15 .25 .35	1-11	.48	.18	.34	.48	.34	.41	.34	.29	.31
G	3	.25 .35 .43	1-11	.24	.03	.14	.19	.06	.13	.25	.22	.24
H	3	.25 .25 .25	1-11	.46	.07	.28	.36	.16	.27	.28	.25	.26
I	3	.15 .25 .35	1-11	.37	.08	.23	.34	.17	.27	.27	.23	.25
J	3	.15 .25 .35	1-11	.46	.006	.23	.41	.007	.20	.28	.20	.24
			12-22	.33	.003	.15	.25	.04	.15	.22	.19	.21
K	5	.15 .2 .25 .3 .35	1-11	.46	.0001	.32	.43	.007	.31	.35	.26	.30
			12-22	.43	.08	.25	.34	.18	.26	.28	.24	.26

Frequency Distribution of Relative Root Mean Square Error : CES

Frequency Distribution of Relative Root Mean Square Error : CESW

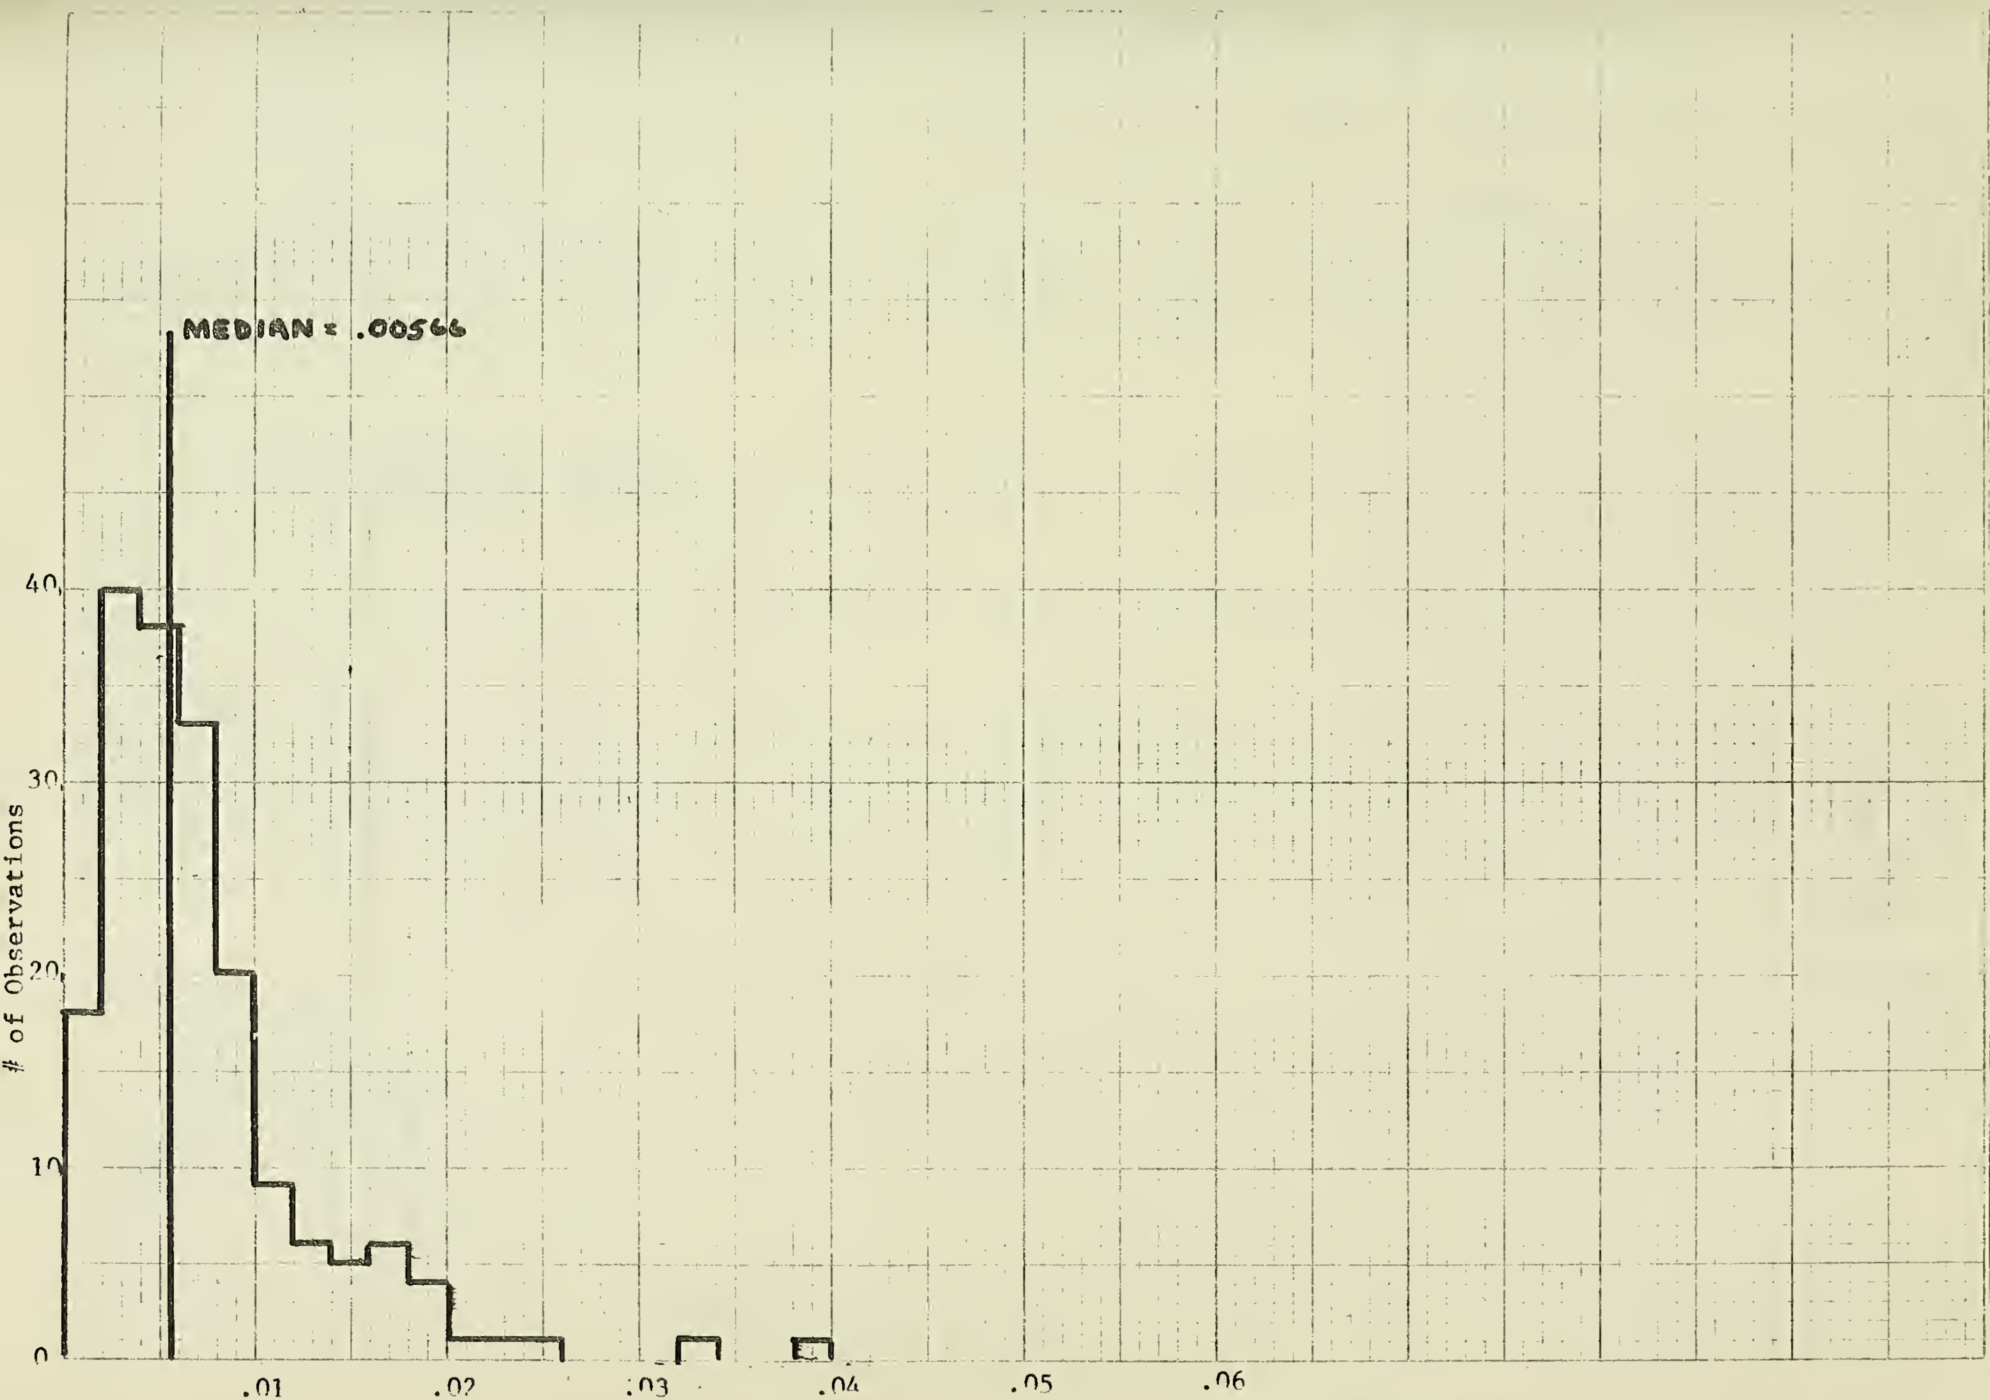

Frequency Distribution of Relative Root Mean Square Errors: CD

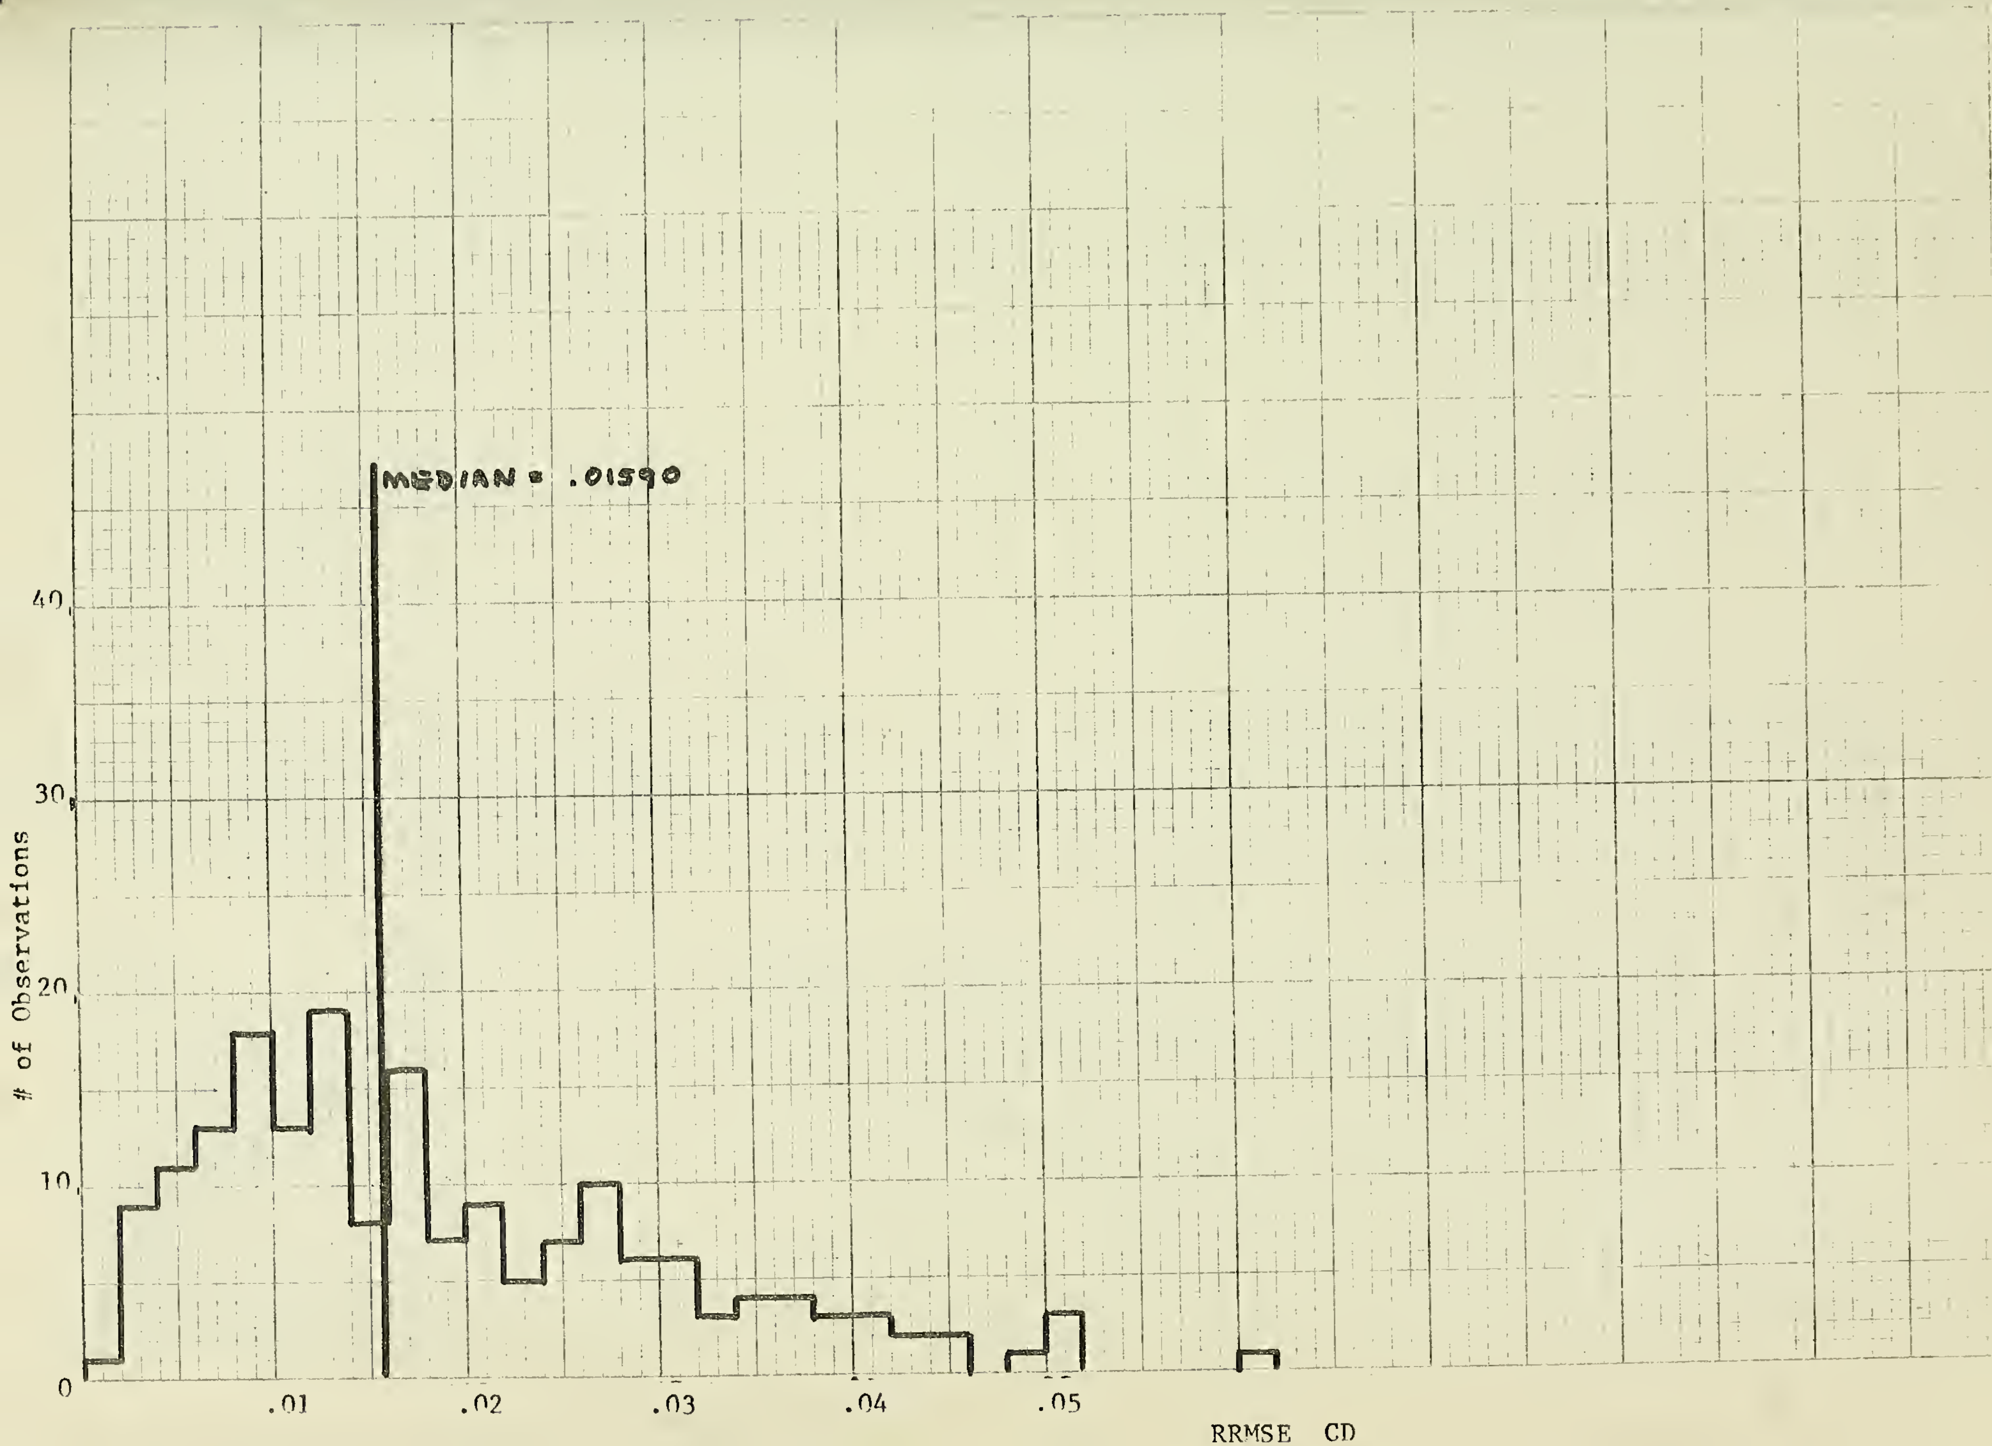

Standard Deviation of Labor's Share and Relative Root Mean Square Error : CD

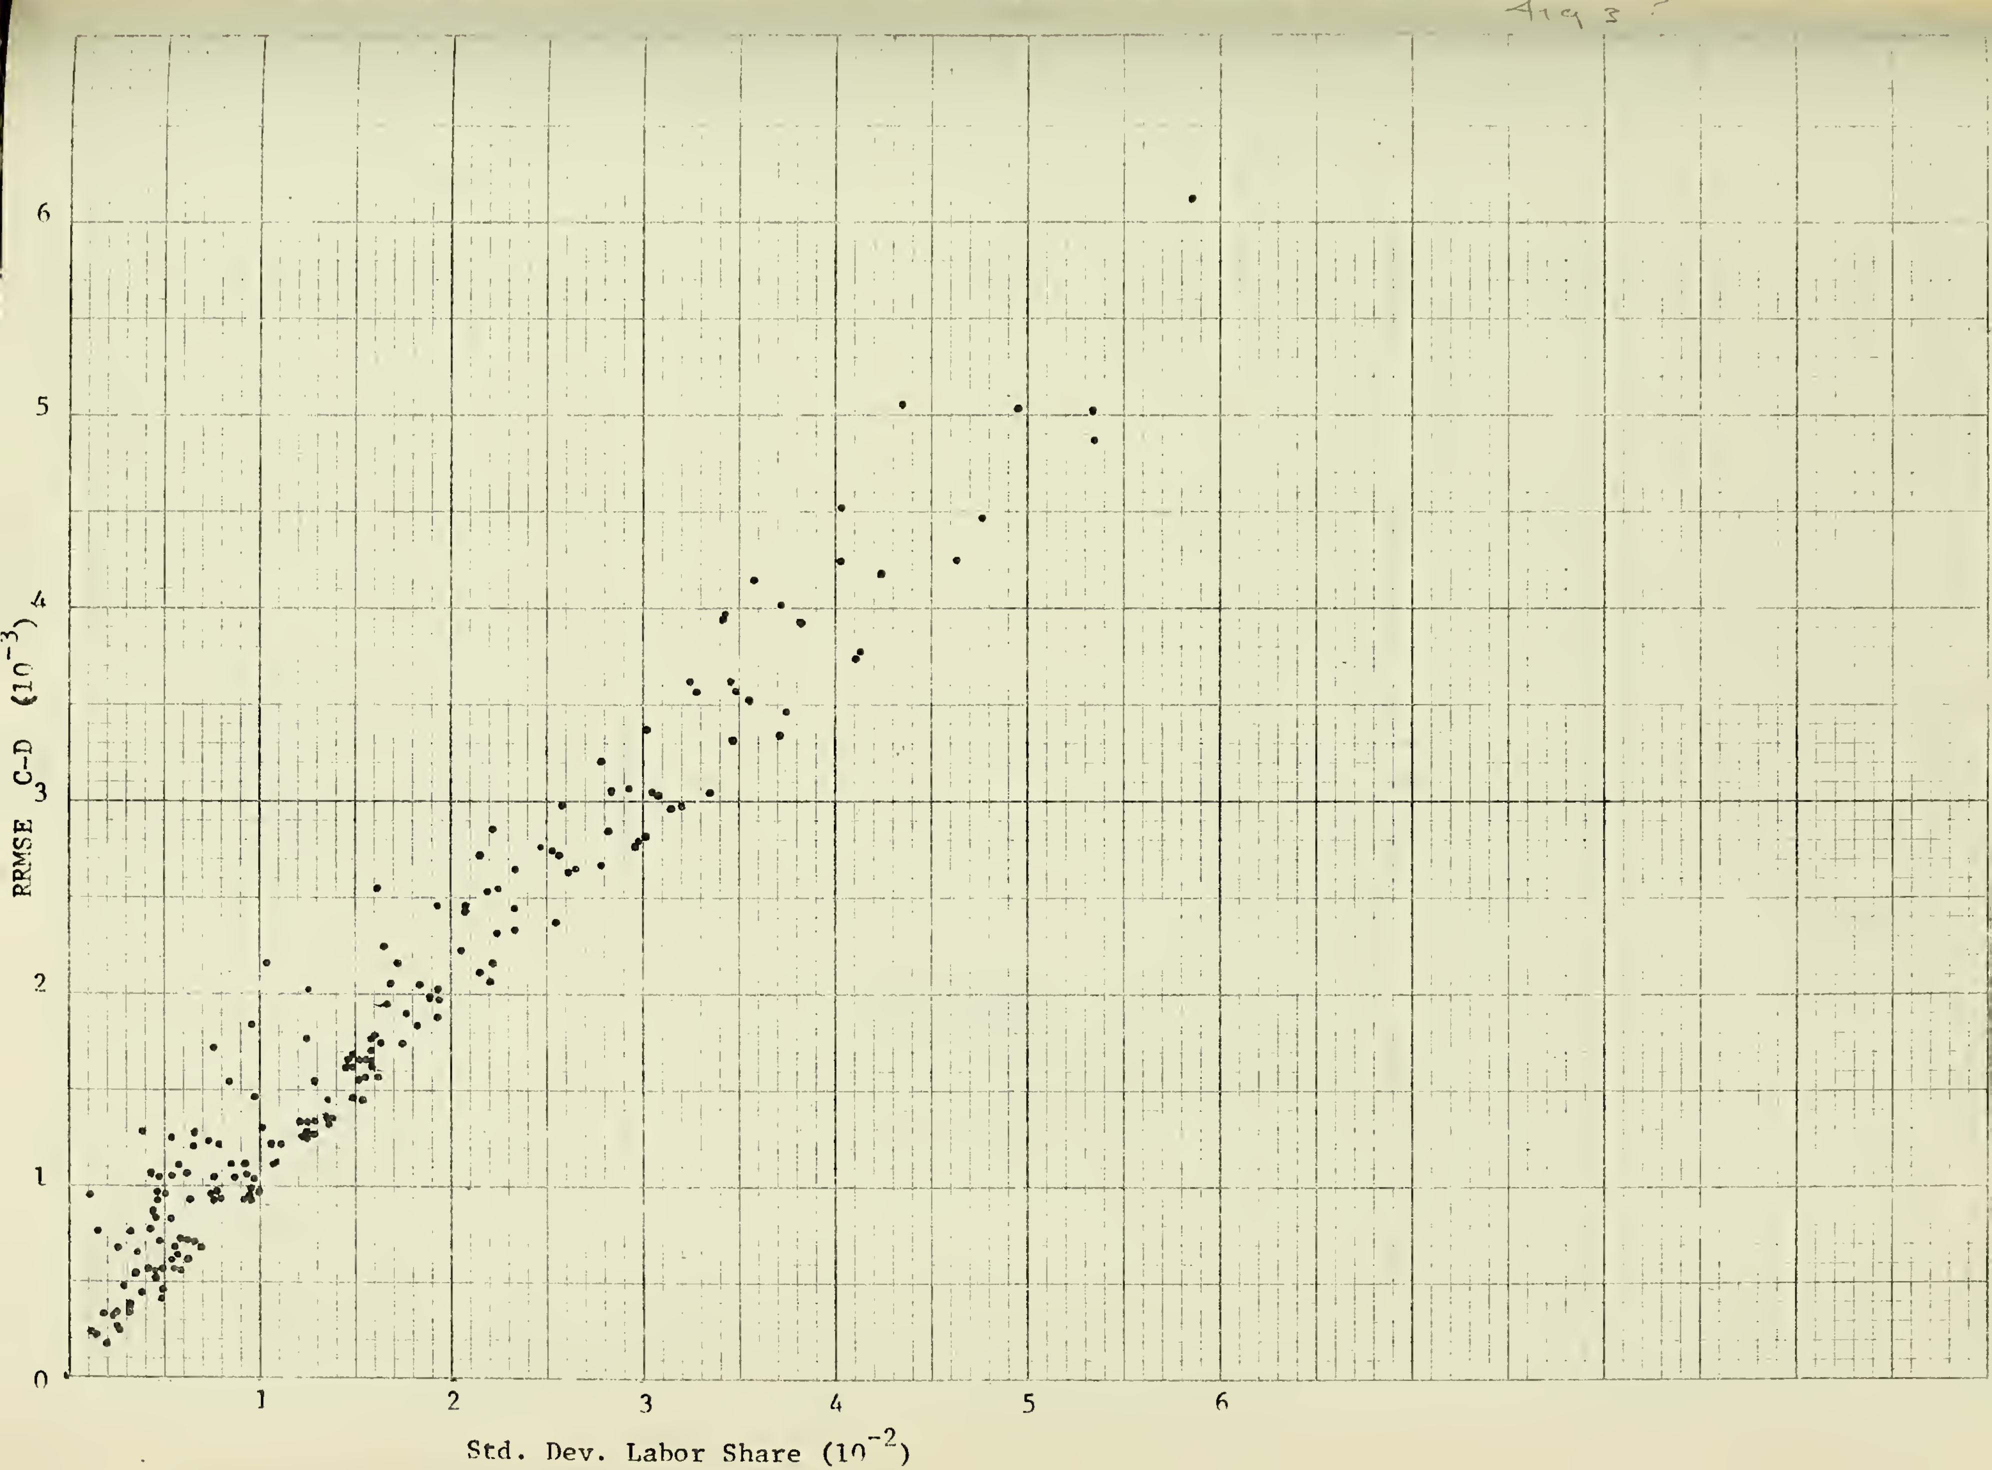

Estimates of δ and of σ :   CES

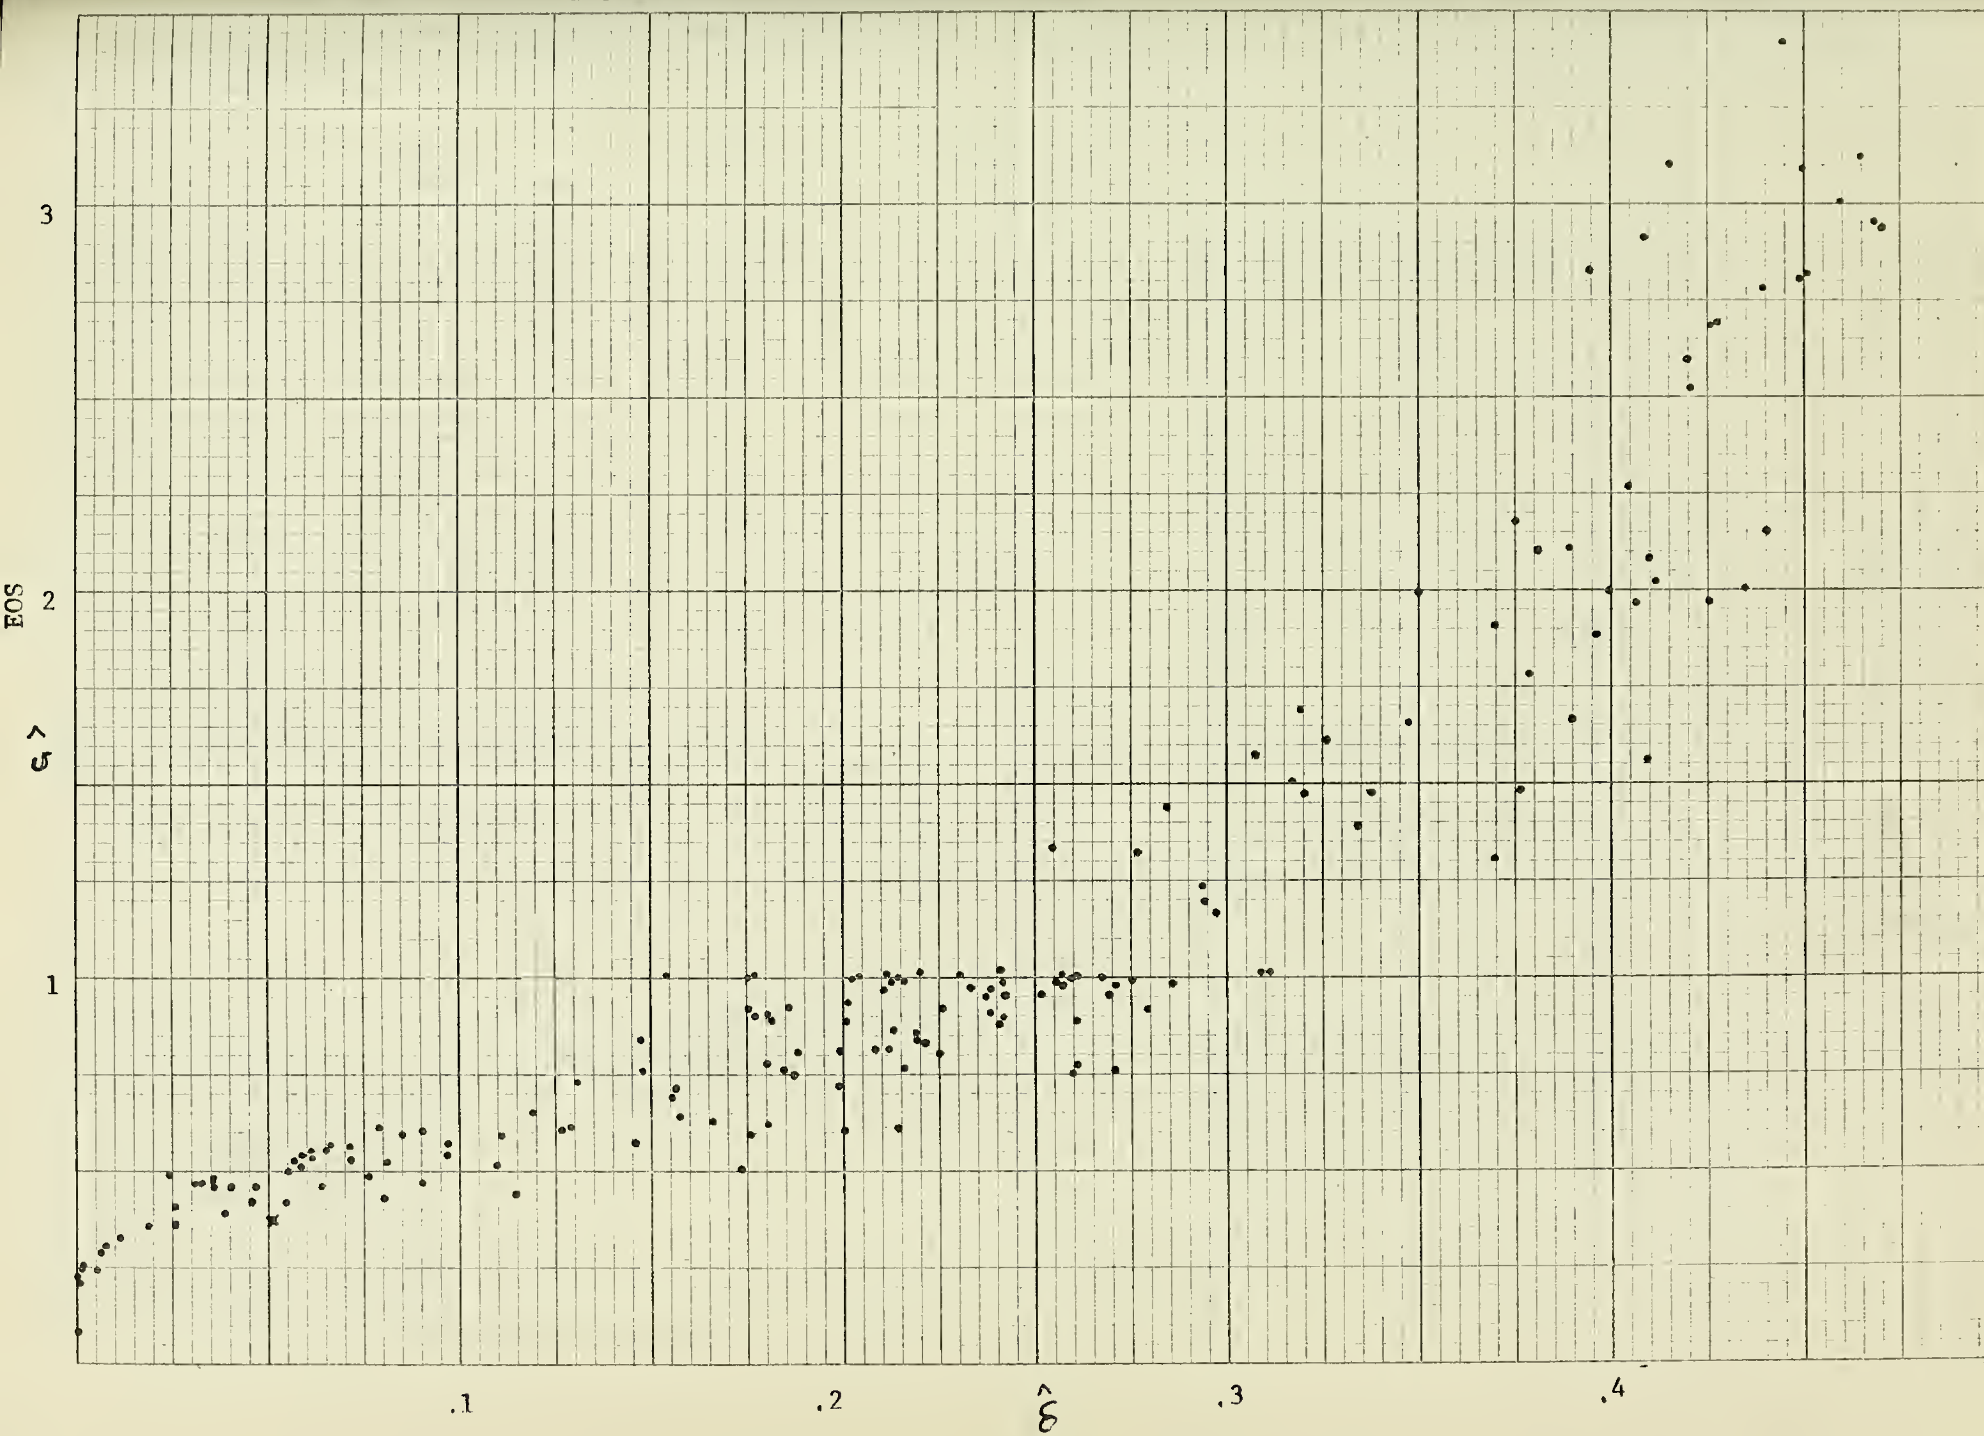

Estimates of δ and of σ : CPSW

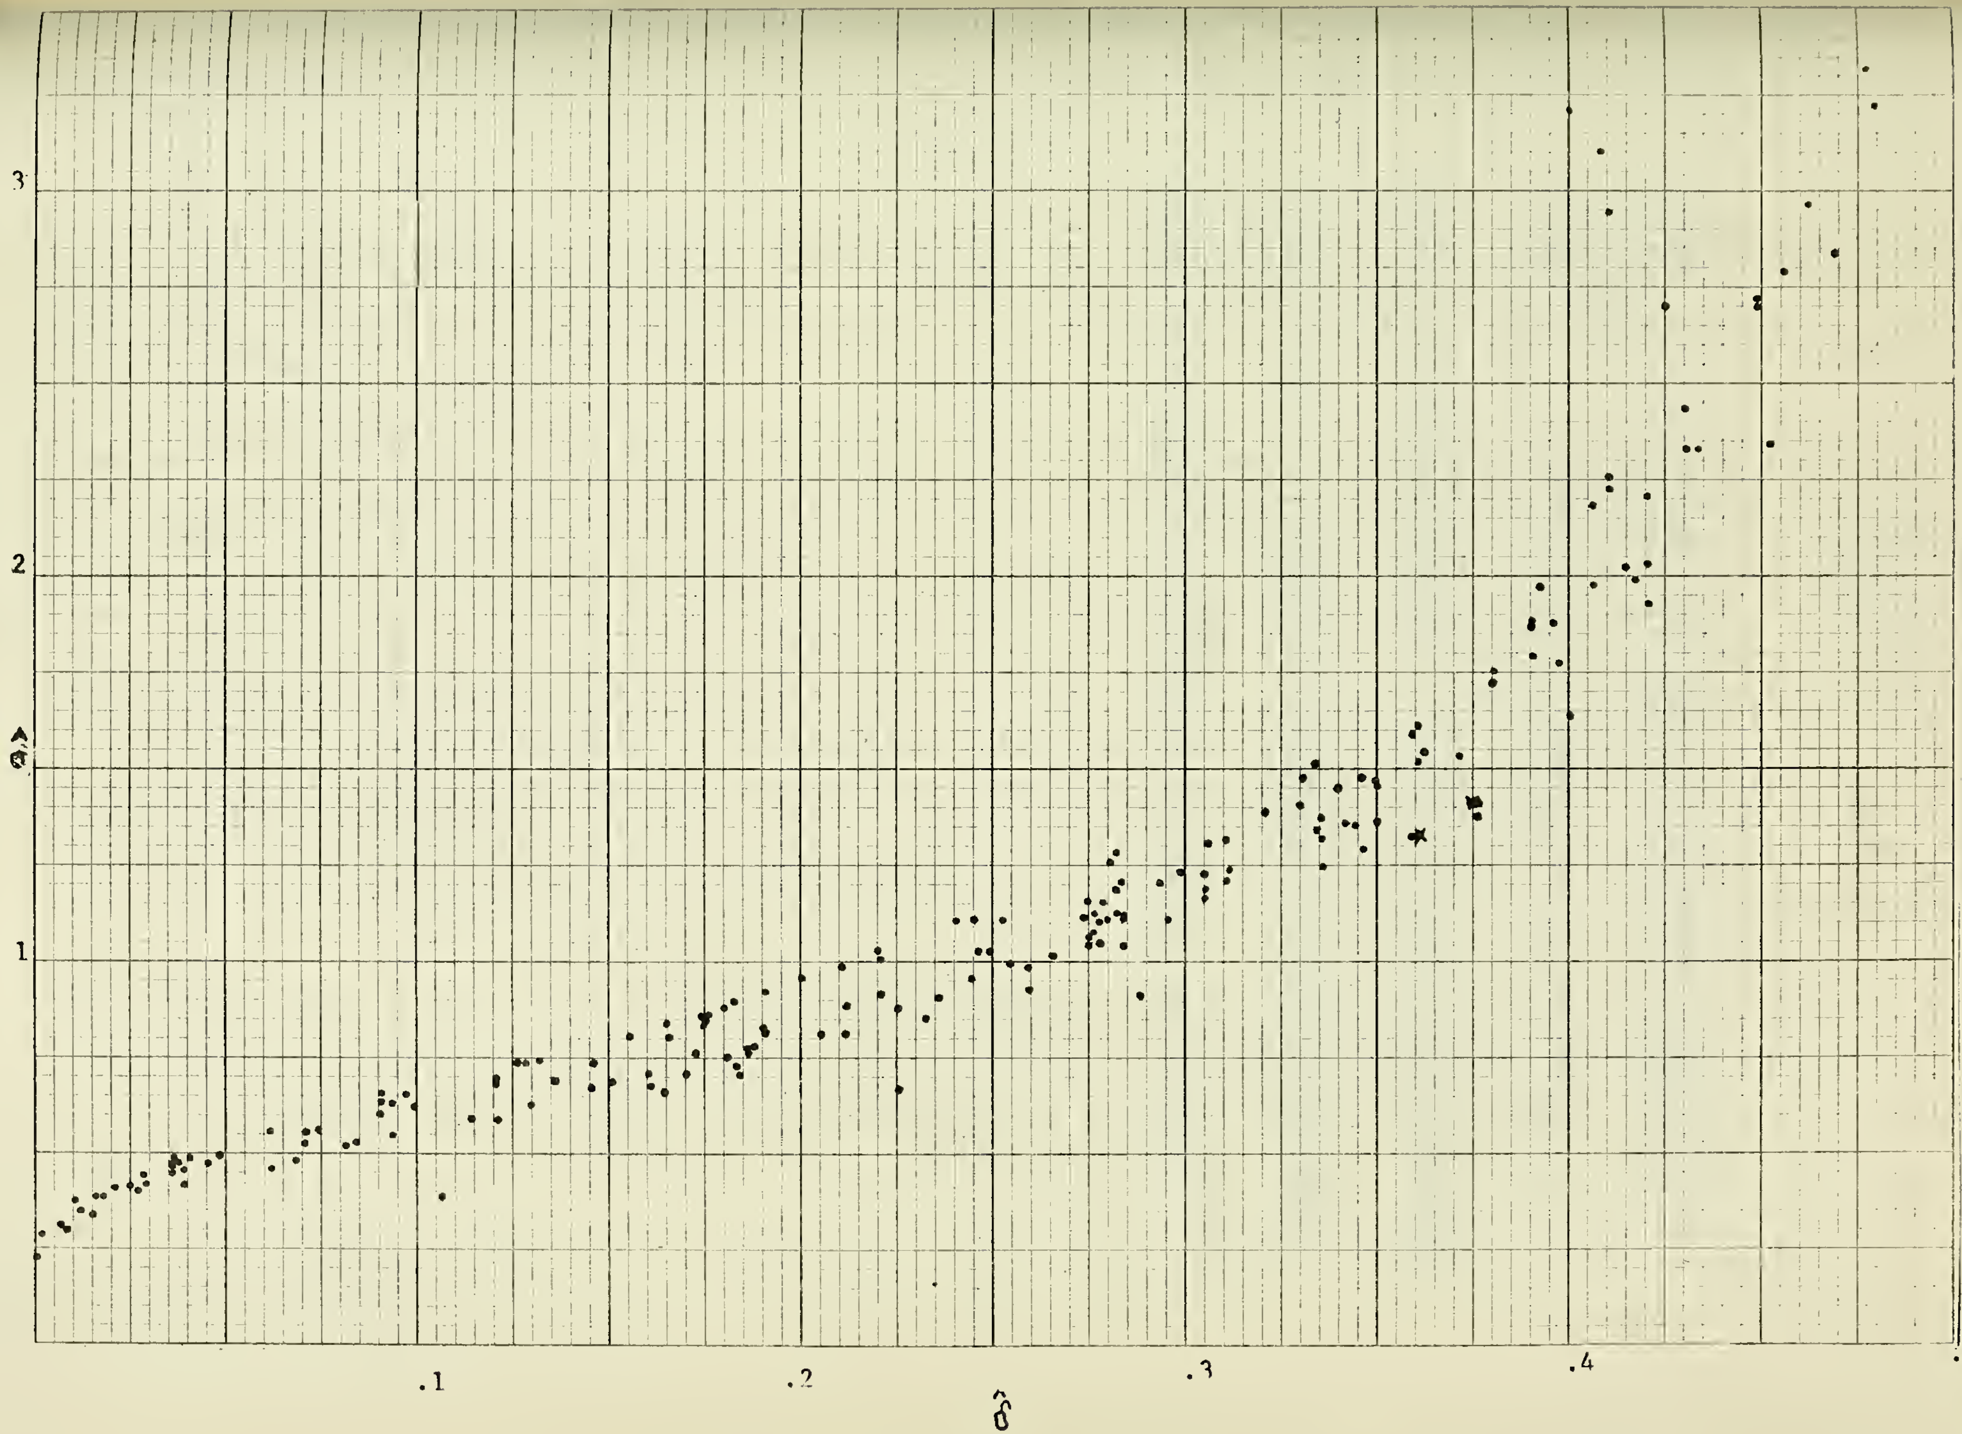

REFERENCES

[1] Arrow, K.J., H. Chenery, B. Minhas and R.M. Solow, "Capital Labor Substitution and Economic Efficiency", Review of Economics and Statistics, 43 (August 1961), 225-50.

[2] Fisher, F.M., "The Existence of Aggregate Production Functions," Econometrica, 37, (October 1969), 553-77.

[3] __________, "Aggregate Production Functions and the Explanation of Wages: A Simulation Experiment," Review of Economics and Statistics, 53, (November 1971), 305-26.

[4] Marquardt,D.W., "An Algorithm for Least-Squares Estimation of Non-linear Parameters," Society for Industrial and Applied Mathematics Journal, (June 1963) 431-41.

[5] Solow, R.M., "Technical Change and the Aggregate Production Function," Review of Economics and Statistics, 39, (August 1957), 312-20.

Appendix

The program developed for this simulation experiment is constructed with three modules:

The first module generates the labor supply for the economy and capital stocks specific to each firm in the economy as described in Section 2 of the paper. The user determines the length of the time series he desires to generate, the number of experiments given firm characteristics, the number of firms, the amount of randomness introduced into the labor supply and capital stocks and provides the appropriate growth and dispersion parameters.

The second module, the "economy" of the experiments, has technologies for Cobb-Douglas, CES and fixed-coefficient firms. The user selects the technology or mix of technologies (all experiments described in this paper used the CES technology) and provides the appropriate firm parameters. The capital stocks generated in Module One are then assigned to individual firms. For each time period, labor is allocated to each firm using an iterative procedure, moving labor from low to high marginal product firms until the marginal product of labor is virtually uniform across firms. This determines the wage rate for the economy for that time period, and all firm marginal products must be within one percent (approximately) of this wage. Simultaneously, the total labor allocated to the firms must be within one half percent of the supply generated in Module One.

Given the allocation of labor to each firm and the capital stock, one can generate a measure of aggregate capital. The firm's technology with the firm's capital and the firm's allocated labor yield output. Thus we get aggregate output and a measure of aggregate capital.

Using the labor supply from Module One, the aggregate measure of capital and the generated output by the economy, Module Three estimates Cobb-Douglas and CES production functions. For the CES estimation we use a nonlinear least-squares algorithm developed by Marquardt [4], modified to impose simultaneous convergence on the parameters and the sum of squared residuals, the percentage change in both being less than 10^{-5} for convergence. The aggregate production function is then used to predict labor's share which is compared with the actual share generated in Modules One and Two.

Date Due

Lib-26-67