

The Calibration of CES Production Functions

Jonathan R. W. Temple

Discussion Paper No. 08/606

September 2008

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN

The Calibration of CES Production Functions*

Jonathan R. W. Temple
Department of Economics, University of Bristol
8 Woodland Road, Bristol, BS8 1TN, UK

September 25, 2008

Abstract

This note addresses some issues that arise when using “normalized” CES production functions, an approach that has become popular in the literature. The results of Klump and de La Grandville (2000) provide a simple way to calibrate the parameters of the CES production function when the necessary data are available. But some of the other applications of normalized CES production functions appear problematic, especially when used to argue that productivity is increasing in the elasticity of substitution.

Keywords: CES production functions, elasticity of substitution, normalization

JEL classifications: D24, O40

*I am grateful to Patrick Carter and Rosaida Mohd Rasep for helpful discussions, but the usual disclaimer applies.

1 Introduction

In recent years, the CES production technology has returned to centre-stage in growth economics and short-run macroeconomics. There are many reasons for this, including the observed variation in factor shares in advanced economies (Blanchard 1997, Bentolila and Saint-Paul 2006); the tendency for empirical studies using single-country data or microeconomic data to estimate the elasticity of substitution between capital and labor as well below unity (for example, Antràs 2004; for surveys see Chirinko 2008 and Klump et al. 2008); and cross-country studies that reject a unitary elasticity of substitution (Duffy and Papageorgiou 2000). In growth models, the consideration of varying factor shares, and factor-biased technical change and questions of appropriate technology, all require technologies more flexible than Cobb-Douglas (for example, Acemoglu 2003, and Caselli and Coleman 2006). The wider use of CES technologies means that it is increasingly common for studies of growth models to consider the effects of variation in the elasticity of substitution (for example, Turnovsky 2002).

Although the CES production technology seems relatively straightforward, its mathematical simplicity can be deceptive. de La Grandville (1989), Klump and de La Grandville (2000) and Klump et al. (2008) have emphasized that the economic interpretation of the CES production technology requires care. In particular, they recommend “normalizing” CES technologies when analyzing the consequences of variation in the elasticity of substitution.

It may seem odd that there is any normalization issue to raise at all. Perhaps the easiest way to demonstrate the underlying problem is to imagine a productivity comparison between two firms, with production functions $AF(K, L)$ and $BG(K, L)$ respectively. At first glance, the parameters A and B enter the production functions symmetrically and have the same interpretation, as TFP parameters. But since the production technologies differ, a comparison of the relative magnitudes of A and B has no economic meaning. The two are not on the same scale,

and the mathematical symmetry is misleading about the economic content of the comparison.

This is a simple illustration of a far more general problem, which emerges especially clearly in the CES case. If the elasticity of substitution is allowed to vary, this is rather like moving from one function $F(K, L)$ to another, $G(K, L)$. This raises the issue of whether other technology parameters will retain the same economic interpretation as before, and what it means, in economic terms, to vary the elasticity of substitution while holding other parameters “constant”. Different proposals for normalizing the CES technology are different proposals about what, exactly, should be held constant as the elasticity of substitution is varied.

Most researchers have worked with the standard (ACMS) form, due to Arrow et al. (1961):

$$Y = A(bK^\rho + (1 - b)L^\rho)^{\frac{1}{\rho}}$$

where Y , K and L are output, capital and labor respectively, and where the elasticity of substitution $\sigma = 1/(1 - \rho)$. Much of the discussion that follows centres on assumptions about the TFP parameter A , and the distribution parameter b , for which the admissible range is $0 < b < 1$. de La Grandville (1989) and some later authors argue that, when varying the elasticity of substitution, both A and b should be considered functions of the elasticity of substitution. They derive explicit relationships that can be used to normalize the function as the elasticity of substitution varies.

The easiest interpretation of normalization is to view the inputs of capital and labor as index numbers, so that each could be measured relative to arbitrarily-chosen benchmark values. We can then write the CES production function in the “calibrated share form” of Rutherford (1995):

$$Y = Y_0 \left(\pi_0 \left(\frac{K}{K_0} \right)^\rho + (1 - \pi_0) \left(\frac{L}{L_0} \right)^\rho \right)^{\frac{1}{\rho}} \quad (1)$$

where the parameter π_0 can be interpreted as the capital share that arises at a benchmark capital-labor ratio K_0/L_0 and output per worker level Y_0/L_0 , under perfect competition and marginal productivity factor pricing. The ACMS form can then be seen as normalizing the function so that the distribution parameter b is the capital share that arises when the capital-labor ratio K_0/L_0 is unity.

It is well known that the distribution parameter cannot be defined independently of the units of measurement of capital and labor, but the problem is deeper than this. If we want to study the effect of varying the substitution parameter ρ , the problem is that the function can be normalized using any benchmark capital-labor ratio, and this arbitrary choice will influence how the production surface is reshaped by changing the elasticity of substitution. For example, with the ACMS form, an increase in the elasticity of substitution reduces the curvature of the isoquants, but the old and new isoquants are tangential at a capital-labor ratio of unity (Kamien and Schwartz 1968, p.12). Not least because we can interpret the measures of capital and labor as dimensionless index numbers, there is no reason to privilege this capital-labor ratio over others, and therefore no inherent reason to prefer one normalization of the CES production function to another. The calibrated share form (1) implies that the tangency of the old and new isoquants will occur at the benchmark capital-labor ratio, rather than a capital-labor ratio of unity.

There is considerable room for misunderstanding here. This note aims to distinguish between instances where it is useful to normalize a CES production function in the way that Klump and de La Grandville (2000) recommend, and instances where the idea could be misused. For simplicity, the discussion will assume there are just two inputs, capital and labor, and constant returns to scale.

2 The uses of normalization

Consider the problem faced by a researcher studying the transitional dynamics of a growth model which includes a CES production function, written in the ACMS form. How should the researcher choose the distribution parameter b ? Conventionally, it is interpreted as the capital share that would arise when the elasticity of substitution is unity, but this is not much use if the researcher is primarily interested in other cases. More generally, we have seen that the ACMS distribution parameter can be interpreted as the capital share that will arise when the capital-labor ratio is unity. But this relationship is not much help, since the researcher will rarely be able to gauge a sensible magnitude for the capital share at that point.

If the researcher has multiple observations on factor shares and factor ratios, the distribution and substitution parameters can be estimated from the data using standard methods. Since the substitution parameter is treated as a fixed constant to be estimated from the data, no issue of normalization arises. Alternatively, if the researcher has just one observation on the capital share, corresponding to particular (known) values of the capital-output ratio and output per worker, then the following expressions can be used to calibrate the TFP and distribution parameters in the ACMS form:

$$A = ((1 - \pi_0)(Y_0/L_0)^\rho + \pi_0(Y_0/K_0)^\rho)^{\frac{1}{\rho}} \quad (2)$$

$$b = \pi_0(Y_0/K_0)^\rho A^{-\rho} \quad (3)$$

These expressions are equivalent to alternative forms in Klump and de La Grandville (2000) and Klump and Saam (2008), and also imply the calibrated share form (1). In this light, normalizing the function is no more than using data to pin down key parameters, and is closer to calibration than normalization.

The conceptual problems are more significant when there are no data available to pin down the parameters. For example, a researcher might want to study the transition dynamics of the Solow model in the CES case. Here, one natu-

ral approach would be to think about the capital share believed to apply in the steady-state. The capital-output ratio in the Solow model is a simple function of the saving rate, and the rates of population growth, technical progress and the depreciation rate. Hence the capital-output ratio can be calculated independently of the production function parameters, and used in equations (2) and (3) together with the researcher's best guess of the capital share that will obtain asymptotically. Alternatively, the researcher could guess the marginal product of capital that will obtain in the steady-state.

But things get especially complicated when the researcher wants to simulate the growth model for distinct values of the elasticity of substitution, and compare the outcomes, in the absence of data that could pin down parameter values. The traditional approach would be to hold the distribution parameter fixed and vary the elasticity of substitution. This will imply different factor shares apply at any given capital-output ratio, and different levels of output per worker apply at any given capital-labor ratio. More fundamentally, however, we have already seen that the economic interpretation of holding the distribution parameter fixed is not straightforward. As noted above, there are many ways of normalizing the CES production function, and the precise choice affects how the production surface is reshaped by variation in the elasticity of substitution.

With this in mind, de La Grandville (1989), Klump and de La Grandville (2000) and Klump and Preissler (2000) argue that the choices of the elasticity of substitution, TFP parameter and the distribution parameter are best seen as interdependent. If the researcher simulating a growth model varies the elasticity of substitution, they should also vary the TFP and distribution parameters. Their recommendation is to express the TFP parameter and the distribution parameter as functions of the elasticity of substitution so that, as the elasticity is varied, the production function always yields the same output per worker and marginal rate of technical substitution at a specific capital-labor ratio. Put differently, the procedure forces production surfaces that differ in the elasticity of substitution to

be tangent to one another along a particular ray $K = k_0L$ where k_0 is a baseline capital-labor ratio. The resulting “normalized” production function can then be written in a number of ways, with equation (8) in Klump and de La Grandville (2000) as one of the simplest:

$$\frac{Y}{L} = \frac{Y_0}{K_0} \left(\frac{\pi_0}{\pi} \right)^{\frac{1}{\rho}} \frac{K}{L} \quad (4)$$

using the same notation as before. Hence, if the researcher knows the capital share which applies at a particular capital-output ratio, the above equation can be used to eliminate any role for the TFP and distribution parameters. This ensures that experiments which vary the elasticity of the substitution are always consistent with the known data.

But if there are not enough data to pin down Y_0/K_0 and π_0 in equation (4), the TFP and distribution parameters remain “free” parameters, and the usefulness of (4) is less clear. Klump and de La Grandville (2000) clearly intend that the normalized production function should allow comparison between economies that differ “only” in the elasticity of substitution. This exercise does not seem conceptually straightforward: since the original production function has three parameters, it would be surprising if the variation across technologies could be reduced to a one-dimensional parameter space. It is clear that, instead of reducing the number of dimensions, normalization replaces the variation in the distribution parameter with variation in the benchmark capital-labor ratio, a quantity that is even harder to interpret.

To give a specific example, Klump and Saam (2008) study the dependence of the convergence rate on the elasticity of substitution. In their Table 1, they report a range of values for the convergence rate, which vary not only with the elasticity of substitution but also with the “baseline capital intensity” k_0 . They write (p. 258) that “The effect a given rise in the elasticity of substitution has on the speed of convergence depends on the relative magnitude of baseline and steady state capital intensity”. The last column of their Table 1 makes clear that

a higher elasticity of substitution either lowers or raises the rate of convergence, depending on the baseline capital-labor ratio. Since the latter can be arbitrarily chosen by a researcher, the economic meaning of this variation is unclear.

3 The misuses of normalization

The discussion above regards the normalization of CES production functions as a useful way of calibrating TFP and distribution parameters when the factor shares (or alternatively, the marginal product of capital, or even the marginal rate of technical substitution) are observed at a particular, known capital-output ratio and level of output per worker. Alternatively, a researcher may have a strong prior about the factor shares that are expected to hold in a long-run steady-state. But for researchers who lack data or priors with which to calibrate CES production functions, perhaps when simulating a dynamic model or carrying out a theoretical analysis, the use of normalization does not eliminate the multi-dimensional nature of the parameter space, and may risk obscuring it.

The associated problems are most serious when normalization is used to derive theoretical results. It has been argued, especially in Klump and de La Grandville (2000), that normalization increases the usefulness of CES production functions for growth theorists, and this has led to its use in subsequent work such as Miyagiwa and Papageorgiou (2007) and Papageorgiou and Saam (2008). In their paper, Klump and de La Grandville prove two theorems which essentially state that, for two economies with constant-returns CES production functions “differing only in the elasticity of substitution”, and sharing common values for the initial capital-labor ratio, population growth and investment rate, the economy with a higher elasticity of substitution will have a higher level of per capita income at other capital-labor ratios, and will have a higher capital intensity and income per capita in the long-run steady-state.

Some have taken this to mean that the elasticity of substitution can be reinter-

preted as an index of technology: as economies become more advanced, perhaps their elasticity of substitution increases, and with it their productivity. If true, this would be a surprising and important claim. It is not guaranteed that labor productivity should be monotonically increasing (or decreasing) in the elasticity of substitution. Arrow et al. (1961) pointed out that the elasticity of substitution may vary with the level of development, but there is no reason to expect productivity to respond in a straightforward way to such a change, or to expect that a rising elasticity of substitution can be the “engine of growth” implied by Klump and de La Grandville (2000, p. 287).

To see this, consider the much simpler case of two countries that differ in their Cobb-Douglas production functions. In country A, output is equal to $Y = AK^\theta L^{1-\theta}$ and in country B, output is equal to $Y = BK^\mu L^{1-\mu}$ where $\mu \neq \theta$. A quick calculation will show that the two production surfaces intersect: for some capital-labor ratios, output per worker will be higher in country A, and for others, it will be higher in country B. It is easy to show that there will be a threshold value for the capital-labor ratio, with the identity of the more productive country changing as this threshold is crossed. This brief example shows that productivity is not uniformly increasing in the output-capital elasticity, and accordingly, increases in the latter would rarely be considered an engine of growth.¹

With the above discussion in mind, it is clear that we may not be able to derive a monotonic relationship between productivity and the elasticity of substitution, contrary to Theorem 1 in Klump and de La Grandville (2000). And examining that theorem more closely raises significant problems of interpretation. A higher elasticity of substitution is said to ensure that a country will have a higher level of productivity “at any stage in its development” (at any capital-labor ratio). But

¹This is not to say that other comparative static results are ruled out. In the Cobb-Douglas version of the Solow model, for example, it is possible to derive the effect of the output-capital elasticity on steady-state output per capita, and on the convergence rate in the vicinity of the steady-state. What the example shows, however, is that simple relationships between productivity levels and specific technology parameters do not always exist.

the normalization procedure implies, by construction, that the productivity level is unchanged at the benchmark capital-labor ratio used to normalize the production function as the elasticity of substitution varies. Moreover, precisely because the different production surfaces are tangent at the benchmark capital-labor ratio, the posited increase in productivity will be very small in the vicinity of the benchmark capital-labor ratio. This is clear, for example, from figure 1 in Klump and de La Grandville (2000).

These points immediately call into question the ambitious interpretation of the elasticity of substitution as a meaningful index of productivity or the level of technology. It is an unusual index of technology which, when it increases, does not raise productivity at all capital-labor ratios; and, more seriously, which raises productivity by varying amounts depending on a benchmark capital-labor ratio that is arbitrarily chosen. It would be hard to use this index to quantify the associated effect on productivity. To clarify just how odd the interpretation becomes, we can make the discussion more concrete by translating it into a slightly different context. Here is an imaginary conversation between a company engineer and the company's manager:

Engineer: I have changed our production process so that it now has a higher elasticity of substitution.

Manager: Excellent. I understand that this will lower our unit costs, at any given ratio of inputs?

Engineer: Not exactly. It will lower them at all input ratios except one, and in the vicinity of that exception, the reduction in unit costs will be modest.

Manager: And what is the input ratio at which unit costs remain the same as before?

Engineer: That depends, and in principle, it could be any of them.

Translating the problem into such concrete terms makes clear that normalizing the production function, by ensuring that different production surfaces are tangent at a particular input ratio, raises significant problems of interpretation. The nor-

malization ultimately fails to resolve the underlying difficulty, namely that when a production technology changes form, the economic interpretation of all the technology parameters is likely to be altered. As a result, it is hard to interpret comparative static results for a parameter such as the elasticity of substitution.

These problems are especially serious when the normalization is used to derive theoretical results that contrast outcomes associated with the different surfaces. To be clear, the issue raised here does not concern the logical or mathematical consistency of the proofs given in Klump and de La Grandville (2000) but rather their economic interpretation. In particular, it is not clear what it means to say that these alternative economies are “differing only by their elasticity of substitution”, and to assume that such a statement is unambiguous carries significant risks.

In interpreting such theorems, precision in language is everything. Given the discussion above, an appropriate restatement of their Theorem 1 would be “If the parameters of the CES production functions of two different economies are calibrated so that both functions yield the same capital share at a specific level of output per worker and capital per worker, then the economy with the higher elasticity of substitution will be predicted to have higher output per worker at all other values of the capital-labor ratio”. This claim seems internally consistent, but it is a claim about what happens when CES production functions are calibrated to match specific quantities. It does not imply that the elasticity of substitution can be interpreted as an index of productivity, or that economies with a higher elasticity of substitution are more productive, or that we should expect higher levels of economic development to be associated with higher elasticities of substitution. Instead, the theorem tells us how production surfaces are reshaped by changes in the elasticity of substitution, when the old and new surfaces are constrained to match certain quantities at a particular benchmark point.

A more general lesson is that, if a researcher thinks that the elasticity of substitution may vary, or that it can be an engine of growth, the analysis will probably need to be located within an explicit structural model in which the production

technology is endogenously determined. Such a model could then relate the parameters of the technology to a set of invariant (“deep”) structural parameters. In that case, because the appropriate variation of the technology parameters can be deduced from the structure of the model, the issue of normalization does not arise.

4 Conclusions

The above discussion of normalizing CES production functions can be summarized as follows. The various recent papers on this topic have rightly drawn attention to the potential importance of CES technologies. They have also helped to show that, when a researcher simulates a model using a CES technology, holding the distribution parameter fixed while varying the elasticity of substitution is not innocuous. If one technology parameter varies, the interpretation of others is likely to vary. A sensitivity analysis then needs to be carried out in several dimensions, not one. Alternatively, when a researcher knows the factor shares that obtain at a particular capital-output ratio, or that might hold in a steady-state equilibrium, the normalization literature provides a straightforward way to calibrate the distribution parameter so that it remains consistent with the data as the elasticity of substitution varies. These are all useful contributions to our understanding of CES technologies, which should have a significant impact on the future development of the literature.

There is a danger, however, that the literature on normalization has moved too far beyond this. In particular, normalizing CES technologies does not reduce the dimension of the underlying parameter space without raising significant problems of interpretation. The normalization approach is not enough to allow a meaningful comparison of economies that differ “only” in the elasticity of substitution, nor to permit the claim that labor productivity is increasing in the elasticity of substitution, nor that increases in the elasticity will inevitably form an engine of growth.

A final conclusion is that, if theoretical claims are to be made in a model where the elasticity of substitution can vary, ideally these claims should be made within a structural model that relates *all* the CES technology parameters to invariant structural parameters. Comparative statics can then be carried out with respect to the structural parameters, avoiding the need for normalization, or the arbitrary choice of a benchmark point at which a family of surfaces are required to be tangential to one another. The use of a structural model seems the only way to eliminate the problem rightly identified by the normalization literature, which is that changing the value of one parameter within a production technology will typically alter the economic interpretation of others.

References

- [1] Acemoglu, Daron (2003). "Labor- and Capital-Augmenting Technical Change," *Journal of the European Economic Association*, 1, 1-37.
- [2] Antràs, Pol (2004). "Is the U.S. Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution," *Contributions to Macroeconomics*, 4(1), Article 4.
- [3] Arrow, Kenneth J., Chenery, Hollis B., Minhas, Bagicha S. and Solow, Robert M. (1961). "Capital-labor substitution and economic efficiency," *Review of Economics and Statistics*, 43(3), 225-250.
- [4] Bentolila, Samuel J. and Saint-Paul, Gilles (2006). "Explaining Movements in the Labor Share," *Contributions to Macroeconomics*, 3(1), Article 9.
- [5] Blanchard, Olivier J. (1997). "The medium run," *Brookings Papers on Economic Activity*, 2, 89-158.
- [6] Caselli, Francesco and Coleman, Wilbur John (2006). "The World Technology Frontier," *American Economic Review*, 96(3), 499-522.

- [7] Chirinko, Robert S. (2008). “ σ : The long and short of it,” *Journal of Macroeconomics*, Elsevier, vol. 30(2), pages 671-686, June
- [8] de La Grandville, Olivier (1989). “In quest of the Slutsky diamond,” *American Economic Review*, 79(3), 468-481.
- [9] Duffy, John and Papageorgiou, Chris (2000). “The Specification of the Aggregate Production Function: A Cross-Country Empirical Investigation,” *Journal of Economic Growth*, 5, 83-116.
- [10] Kamien, Morton I. and Schwartz, Nancy L. (1968). “Optimal ‘Induced’ Technical Change”, *Econometrica*, 36(1), 1-17.
- [11] Klump, Rainer and de La Grandville, Olivier (2000). “Economic growth and the elasticity of substitution: two theorems and some suggestions,” *American Economic Review*, 90(1), 282- 291.
- [12] Klump, Rainer, McAdam, Peter and Willman, Alpo (2008). “The long-term sucCESs of the neoclassical growth model,” *Oxford Review of Economic Policy*, 23(1), 94-114.
- [13] Klump, Rainer and Preissler, Harald (2000). “CES production functions and economic growth,” *Scandinavian Journal of Economics*, 102(1), 41-56.
- [14] Klump, Rainer and Saam, Marianne (2008). “Calibration of normalised CES production functions in dynamic models,” *Economics Letters*, 98, 256-259.
- [15] Miyagiwa, Kaz and Papageorgiou, Chris (2007). “Endogenous aggregate elasticity of substitution,” *Journal of Economic Dynamics and Control*, 38, 2899-2919.
- [16] Papageorgiou, Chris and Saam, Marianne (2008). “Two-level CES Production Technology in the Solow and Diamond Growth Models,” *Scandinavian Journal of Economics*, 110(1), 119–143.

- [17] Rutherford, Thomas F. (1995). “Constant Elasticity of Substitution Functions: Some Hints and Useful Formulae”, manuscript, University of Colorado.
- [18] Turnovsky, Stephen J. (2002). “Intertemporal and intratemporal substitution, and the speed of convergence in the neoclassical growth model,” *Journal of Economic Dynamics and Control*, 26, 1765 – 1785.