Econometric Analysis and the Study of Economic Growth: A Skeptical Perspective

Steven N. Durlauf

Department of Economics, University of Wisconsin

This Draft: May 30, 2000

Forthcoming in Macroeconomics and the Real World, R. Backhouse and A. Salanti, eds., Oxford: Oxford Univesity Press. I thank the National Science Foundation, John D. and Catherine T. MacArthur Foundation, and Romnes Trust for financial support. I also William Brock, and Rodolfo Manuelli for useful conversations on the issues in this paper. Andros Kourtellos and Artur Minkin have provided excellent research assistance. All errors are mine.

1. Introduction

This essay has an explicitly provocative purpose – namely to argue that the econometric component of the new economic growth literature has done little to adjudicate leading growth questions. While I believe this empirical literature has identified a number of interesting data patterns and led to the development of a number of important new econometric methods (both achievements exemplified in the work of Quah (1996a,b,1997)), the vast body of empirical studies has shed remarkably little light on the main theoretical issues which have emerged since the seminal papers of Romer and Lucas. While largely a critique of existing empirical practice, my hope is that the essay will prove to be suggestive of alternative ways to think about growth data.

The empirical growth literature with which I am concerned consists of cross-section (and sometimes panel regressions) of aggregate growth rates on various country level variables.[1] A generic form of this type of regression is

$$g_i = X_i\gamma + y_{i,0}\beta + \epsilon_i \qquad (1)$$

where g_i is real per capita growth in economy i over a given time period, X_i is a vector of country-specific controls, $y_{i,0}$ is the real per capita income of the country at the beginning of the period over which growth is measured, and ϵ_i an unexplained residual. The singling out of initial income as a

[1]The use of panel regressions has been motivated primarily by a desire to permit the researcher to eliminate country-specific fixed effects by differencing. Durlauf and Quah (1998) criticize this approach. The arguments made in this paper apply equally well to cross-sections and panels.

determinant of growth is done in order to allow testing for the convergence hypothesis, which we describe below. This type of regression has been employed to study an immense range of growth theories.

In its initial incarnations, the equation was employed to determine whether the neoclassical Solow-Swan growth model well describes cross-country growth dynamics; the most persuasive paper of this type is Mankiw, Romer and Weil (1992) which limited the variables in X_i to those which are directly implied by the Solow model, namely population growth and human and physical capital savings rates. A fair description of this early literature is that it showed that observed growth dynamics for the time period of approximately 1950 to 1990 are in a number of respects qualitatively consistent with the predictions of the neoclassical model.

Perhaps the most important claim of this literature is the consistency of observed data with the convergence hypothesis. By convergence, this literature refers to the hypothesis that β is less than zero. A negative β in the standard cross-country growth regression implies, after controlling for various country-specific factors, that countries with lower initial incomes grow, on average, faster than those with higher initial incomes. The convergence hypothesis has been treated as a critical test of neoclassical versus endogenous growth theories, since the Solow model predicts a negative β. In turn, some endogenous growth theories, which typically are based on global increasing returns to scale, imply that β should be positive. Further, the idea that, other things equal, rich countries grow more slowly than poor ones, has the substantive implication that current per capita income differences across nations can be expected to narrow.[2]

The second stage of the empirical growth literature has been less concerned with convergence and tests of the validity of the neoclassical

[2]See Bernard and Durlauf (1996) for a discussion of different statistical notions of convergence and Galor (1996) for a discussion of the convergence-related implications of different growth theories.

model. Rather, this literature has been interested in which variables should be included in X_i. As theories have proliferated as to why factors ranging from democracy to trade openness to geography should influence cross-country growth patterns, evidence in support of these theories has been sought by the determination that empirical proxies for these theories, when included as elements of X_i, are statistically significant; Durlauf and Quah (1998) survey the range of candidate explanations.

2. Basic problems

While generating a range of interesting findings on patterns in the cross-country data, there are strong reasons to question whether the empirical growth literature has done much to adjudicate causal explanations of growth. This is a sweeping claim, of course, and can only be fully defended through a detailed study of the various empirical papers which purport to argue for a particular explanation. However, there do seem to be a set of general problems with the empirical growth literature which imply that attempts at structural or causal inference have not been successful.

i. causality versus correlation

The many candidate variables in the growth literature are typically endogenous. It is hard to think of reasons why savings rates, political characteristics, inequality and the like would *not* be determined jointly with growth rates. Therefore, any movement from the observation of statistical significance of a variable coefficient in a regression to claims about causality is not warranted. This criticism has the status of a "folk" objection in the growth literature, in that it has been made so many times that it is not clear to whom it can be first attributed. However, frequency of complaint does not invalidate the basic point. While there have been some very limited efforts to deal with endogeneity, it seems fair to say the literature has largely ignored this question, relying on vague suggestions (eg. Barro (1999) pg. 12) that timing can uncover the difference between causality and correlation. This of course, if nothing more than the *post hoc ergo propter hoc* fallacy which periodically reappears in economic debates (cf. the monetarist/Keynesian debates between James Tobin and Milton Friedman, Tobin (1970)).

At first glance, it might seem that there is an obvious solution to the simultaneity problem I have described – namely the use of instrumental variables to proxy for the various regressors X_i. Such a solution is more problematic than one might think. Suppose that there exists a set of instruments which are strictly exogenous, in the sense that they are determined outside of the system which determines growth and the other regressors in eq. (1), and that these instruments causally determine those regressors. Formally, what I mean is that suppose for a given choice of regressors X_i and eq. (1) we could identify an additional set of regressors Z_i such that

$$X_i = Z_i A + \eta_i \tag{2}$$

where Z_i is a deterministic function $F(\cdot)$ of some underlying factors ξ_i which are determined outside the socioeconomic system as constituted over the period of study.

$$Z_i = F(\xi_i) \tag{3}$$

The components of Z_i can be thought of as variables such as longitude or latitude, land area of country i, or language. This system would seem well designed for using Z_i to instrument for some or all of the elements of X_i.

However, even in this "ideal case" it is unclear that the Z_i's are valid instruments for the regressors in the context of equation (1).

The reason that Z_i's may not be valid instruments is that growth theories are "open-ended," by which I mean that the claim that one variable influences growth does not typically imply that other variables do not. So, a theory which says that democracy positively influences growth will usually have no implication for whether or not geography influences growth. The regressor error ϵ_i is the cumulation of these unmodelled factors, and cannot be given any further interpretation. Even if Z_i is determined outside the socioeconomic system, this does not imply it is uncorrelated with ϵ_i. Therefore, for an instrument to be valid for eq. (1), one has to assume that it is uncorrelated with ϵ_i, which would require the very strong assumption that it is uncorrelated with all the omitted growth determinants.

To see how this may be problematic, consider the paper by Frankel and Romer (1996) on trade and growth. (I single this paper out as it well illustrates the pitfalls of instrumental variables use in the context of thoughtful study.) Frankel and Romer argue that since trade openness is clearly endogenous, it is necessary to instrument such a variable in the standard cross-country regression if one is to consistently estimate the trade coefficient. In order to do this, they use a geographic variable, i.e. area, as an instrument. However, is it plausible that country land size is uncorrelated with the omitted growth factors in their regression? The history and geography literatures are replete with theories of how geography affects political regime, development, etc. For example, larger countries may be more likely to be ethnically heterogeneous, leading to attendant social problems. Alternatively, larger countries may have higher per capita military expenditures, which means relatively greater share of unproductive government investment and/or higher distortionary taxes. My point is not that any of these links is necessarily salient, but that the use of land area as an instrument presupposes the assumption that the correlations between land size and all omitted growth determinants are in total negligible. It is difficult to see how one can defend such an assumption when these omitted growth determinants are not specified.

Notice that this open-endedness is different from what one found in rational expectations models, at least with respect to instrumental variables. In the rational expectations literature, econometric implementation of a model is typically done by constructing a variable which equals the difference between some quantity realized at date t and the optimal forecast of that quantity at $t-1$, under the assumption that the model is true. Since an optimal forecast error must be orthogonal to all information available at the time the forecast is made, any variables which are realized before t are available as instruments.

Now, one might argue that the distinction I have drawn between the logic of growth regressions and the logic of rational expectations econometrics is unfair, in that in each case one is actually using instruments under the assumption that the model under analysis is correctly specified, which means that the instruments are uncorrelated with some error-ϵ_i in the case of growth regressions and a forecast error in the case of rational expectations. However, the key difference is that the rational expectations specification contains the total economic reasoning of a model in a way which the growth regressions do not, given the open-endedness of growth theories.

ii. model specification

A second problem with the empirical growth literature has been a lack of attention to the implications of various growth theories for the specification of the empirical models used to compare growth theories. The use of linear cross-country regressions is usually justified by appealing to a linearization of the Solow growth model, or by the assumption that the aggregate production function is Cobb-Douglas. Rather than direct testing of neoclassical versus endogenous growth alternatives, the literature has used the results of this linear regression to argue in favor of one or the other theory. Hence the interest in the sign of β.

The relevant issue in assessing growth theories using such regressions is whether alternatives to the Solow model will be revealed using such a functional form. The problem is that for many endogenous growth theories, the linear regression (1) is a misspecification of the growth process. Therefore, in assessing growth theories via the regression, one needs to ask how data generated by an endogenous growth model will appear when passed through a misspecified linear model. It is far from clear that the salient statistical implications of a given endogenous growth model will be preserved in such an exercise. Bernard and Durlauf (1996) give the following example of what might happen. For models of multiple steady states such as Azariadis and Drazen (1990), assume that all countries within a sample have identical characteristics X_i, which means that growth dynamics can be approximated as

$$g_i = (y_{i,0} - y_i^*)\beta' + \epsilon_i \tag{4}$$

where y_i^* is the steady state for country i and $\beta' < 0$. This model exhibits local convergence, in the sense that countries associated with the same steady state will exhibit Solow-type dynamics. Assume further that there are a discrete number of these steady states and that some countries are associated with each. This is a case where there is no economically interesting form of convergence as different initial incomes can produce different limiting income levels.

If one runs a regression of the form (1), then there is an omitted variable, $-y_i^*$, when the data are in fact generated by (4). To see how this affects interpretations of convergence, consider the value of the β coefficient in the regression

$$g_i = c + y_{i,0}\beta + \epsilon_i \tag{5}$$

(which is a special case of (4) when all countries have the same steady state). Expressed in terms of population moments, when eq. (4) is the correct specification, β equals

$$\beta = \frac{\beta' cov(y_{i,0} - y_i^*, y_{i,0})}{var(y_{i,0})} = \beta'(1 - \frac{cov(y_i^*, y_{i,0})}{var(y_{i,0})}). \tag{6}$$

This expression illustrates two things. First, $\beta < 0$ in the standard cross-country growth regression is compatible with a model with multiple steady states, since one can easily identify joint distributions y_i^* and $y_{i,0}$ where this is so (recall that β' was assumed to be negative). Second, the degree of global convergence observed via the misspecified regression can exceed that which occurs locally, if $cov(y_i^*, y_{i,0}) < 0$. Intuitively, if countries with lower initial incomes have relatively higher steady states (in the sense of how far they need to move to achieve their steady states), this will create spurious evidence of convergence.[3]

What this example is meant to illustrate is that by failing to consider the properties of cross-country regressions when estimating using data from non-neoclassical models, one can easily draw spurious inferences concerning the validity of the neoclassical model. Endogenous growth theories, as a corollary of their capacity to produce multiple steady states or explosive growth trajectories, embody deep nonlinearities in their structures. Regressions such as (1) which do not explicitly model these nonlinearities

[3]This example was chosen for algebraic simplicity; it can easily be generalized to allow for countries with different initial conditions to converge to different steady state growth rates.

are capable of producing results which appear supportive of the neoclassical model.

iii. discrimination between growth theories

A third problem with the empirical growth literature is the multiplicity of theories. In their 1998 survey of the empirical growth literature, Durlauf and Quah (1998) found over 90 different variables that have been proposed for cross-country growth regressions of approximately 100 observations; since that survey a number of new variables have appeared. Each of these variables is, in my judgment, at least somewhat plausible ex ante as at least a partial determinant of growth.

This large number of candidate variables relative to available data is naturally a worry. Unfortunately, there do not exist any satisfactory ways to let the data determine which variables are selected. At one level, this is a statement about the statistics literature, and not about empirical economics per se. While there are a plethora of procedures which have been proposed for variable selection, (see Miller (1990) for a survey of classical methods and Beauchamp and Mitchell (1988) and George and McCulloch (1993) for Bayesian approaches), it is fair to say that no consensus exists as to appropriate procedures.

This being said, the empirical growth literature has not even approached this level of sophistication. There are two well known efforts which attempt to determine which variables robustly "matter" in growth regressions: Levine and Renelt (1992) and Sala-i-Martin (1997). Each of these papers in essence attempts to identify which possible growth determinants, as defined as a fixed universe of variables, is robustly statistically significant, where robustness is defined as remaining significant in the presence of various other possible variables. Specifically, a large set of regressions are run with different combinations of variables; robustness of a given variable is evaluated relative to the distribution of coefficients and standard errors generated for the variable. In the case of Levine and Renelt, Leamer's (1983) extreme bounds analysis is employed, which identifies variable significance as fragile if the range of different point estimates is too great across the different regressions which are run. In the case of Sala-i-Martin, a variable is robust if the estimate of its associated coefficient is statistically significant for at least 95% of the regressions.

While these papers both represent useful exercises, there are clear limitations to each. One problem is the sensitivity of the robustness criteria to which alternative variables are included; see Durlauf and Quah (1998) for an extensive discussion. To give one example, suppose I augment the regressor set in Sala-i-Martin's original analysis with a number of silly (from the perspective of a growth regression) variables. Suppose further that after adding these variables, the percentage of regressions where a given variable coefficient is statistically significant is increased. Does this increased percentage of significant coefficients mean that the variable is more robust than before? It would given Sala-i-Martin's evaluative criterion for robustness, although for the thought experiment just described it clearly should not.

More generally, these procedures are based upon the assumption that the statistical significance of individual variables should be used to determine model selection – an assumption that is hard to justify (and is at odds with most of the model selection literature in statistics.) These types of procedures make it possible for one to reject a set of variables as nonrobust, due to high collinearity, even though the exclusion of all of them substantially degrades the explanatory power of the regression.

However, even accepting these exercises at face value, there is the more general question of whether robustness is a desirable criterion for assessing growth explanations. To see why this may well not be true, notice that robustness in these two papers ultimately means that a variable is statistically significant in the presence of combinations of other variables. For this to be true, it must be the case that the variable is not too highly correlated with others in the variable set under study. But why would one expect this lack of correlatedness to hold for the sorts of variables which are candidate growth theories? Among the candidate explanations for growth differences, various authors have studied democracy (e.g. Barro (1996), inequality (e.g. Alesina and Rodrik (1994)), corruption (e.g. Mauro (1995)), trade policy (e.g. Levine and Renelt (1992)) and a large number of alternative measures of a society's social, political and economic status. What is important about such variables is that none of them can be thought of as being determined independently of aggregate growth. In other words, the same processes which determine growth also will determine the social and political landscape of a country. Hence one would expect these various measures both to correlate with growth in a bivariate or regression sense, but also to correlate with each other and hence not be robust in the sense which we have described. Notice that this last point is independent of whether there is simultaneity bias in a cross-country growth regression. Each of these variables can causally determine growth, yet be highly correlated with one another. Therefore, for many of the candidate growth explanations, robustness is an implausible property.

***iv.* heterogeneity**

A final conceptual difficulty concerns the assumption of an invariant statistical model in cross-country growth analysis. As conventionally specified, cross-country regressions presuppose that the analyst believes that the coefficients of the exercise are interpretable as country invariant relationships between some control variable and growth. The existence of such a relationship is far from obvious. What does it mean to say that one expects a one percent change in school enrollment to have the same effect on the growth rate in the United States as Botswana? I realize that at some level, one can criticize any empirical exercise in social science for this reason. The interpretability of any statistical exercise will always be a matter of degree. However, when one is dealing with aggregates as complex as entire countries, it is difficult to see what in terms of causal structure it is that one learns from the standard cross-country exercises which fail to treat parameter heterogeneity as fundamental.[4]

This failure to control for parameter heterogeneity makes any predictive inference from the standard regression models problematic. Does one really think that a marginal change in the level of democracy (assuming a meaningful scale even exists) on growth in a country is independent of its level of education? Unless one understands how individual countries differ from the average behaviors described by the standard cross-country regression, it is impossible to make sensible claims about the effects of political and social change on growth.

This type of worry is more than a theoretical concern, as a number of analyses have demonstrated that there is a deep form of heterogeneity in the data. For example, Durlauf and Johnson (1995) employed a classification algorithm known as a regression tree to divide the countries in the Heston-

[4]Periodically, one finds a cross-country study in which the basic regression (1) is modified so that the coefficient for a particular variable is allowed to take on one of a small set of different values; this is done by sorting countries into groups so that the coefficient is required to be constant within groups. For example, countries with per capita incomes above the median are allowed to have one coefficient for some regressor and countries with per capita incomes below the median are allowed to have another. These analyses are generally quite unsatisfactory. For one thing, the reported country groupings are typically *ad hoc* with few (if any) alternative groupings explored, leaving one worried that the reported results are not robust. Further, allowing for different coefficients for different groups of countries for one variable seems question begging. Why one variable and not others? Logically, evidence of parameter heterogeneity for one variable has implications for how one analyzes all the parameters. What is lacking in the literature are systematic analyses of parameter heterogeneity.

Summers data set into groups, each of which obeys a common statistical model. The idea of a regression tree (see Breiman et al (1984) for an exhaustive treatment) is that one treats a data set as consisting of observations drawn from a set of linear models. The algorithm is designed to sort the data into groups of observations each of which obeys the same model. This sorting is done on the basis of a set of control variables. The supports of these variables are partitioned into intervals, these intervals define the groupings of observations. For example, Durlauf and Johnson use initial literacy and initial income as controls, so that country groupings are made on the basis of finding observations which fall into common intervals for these two variables.

Further evidence has been adduced by Canova (1999) in an important recent paper.[5] In this analysis, the logic of the regression tree is preserved in that one seeks to identify groups of countries obeying a common statistical model. However, unlike Durlauf and Johnson, Canova allows the model coefficients to be random for each country within a grouping. This randomness permits intragroup as well as intergroup heterogeneity in the growth process and is particularly appealing in cross-country contexts.

Papers such as Durlauf and Johnson and Canova as well as historical studies such as Landes lead to a common conclusion. Heterogeneity is a key feature of national experiences, so that even if one is willing to consider a common model for uniting these experiences, the parameters of the model are very likely to differ across countries.

It is worth noting that concern over heterogeneity need not lead to agnosticism about empirical work. Much of James Heckman's fundamental contributions to empirical microeconomics (see Cameron and Heckman (1998) and Heckman, Lochner and Taber (1998) for recent examples) has

[5]As will be seen in the next section, Canova (1999) is very much in the spirit of the directions I recommend for future empirical research.

been based on the development of constructive ways of dealing with observed and unobserved heterogeneity. The plausibility and importance of heterogeneity is no less in the case of aggregate economies.

3. What is to be done?

These basic concerns about growth empirics touch on many foundational issues concerning the use of statistical analysis in economics. These criticisms come from very different bases, and therefore have very different implications for econometric practice.

For some of the problems with current econometric practice in growth, the implications seem obvious. If the goal of an empirical exercise is to establish causality, then it is necessary to do two things. First, one can consider statistical solutions for regression endogeneity through the use of instruments, model transformation or any of many well established techniques. Second, and harder, one must make a plausible argument that the instruments employed do not, through omitted growth explanations, correlate with the regression error in the growth equation. This may be facilitated if the number of instruments available exceeds the number of variables which need to be instrumented, since one can then construct tests of instrument validity (i.e. lack of correlation of the instruments with the growth regression error). Even in this case, the plausibility of a claim of instrument validity based upon a statistical test needs to assessed in light of the openendness of growth theories. If the goal of an econometric exercise is to assess the neoclassical growth model relative to an endogenous growth competitor, then the exercise must include the endogenous growth specification if the comparison is to be interpretable.

For other problems, there is no obvious solution. For example, my reading of the statistics literature is that variable selection continues to be a very problematic area of analysis. Further, there also needs to be a recognition that there are limits to what econometric analyses can do in contexts as complicated as growth. With approximately 40 annual observations on approximately 100 countries, the high covariance of plausible growth explanations means that statistical analysis alone is very unlikely to be able to fully discriminate across the plethora of growth explanations.

In addition, two modifications in current econometric methodology warrant exploration.

***i.* greater eclecticism in empirical work**

Growth is an obvious area where econometric analysis should be supplemented by historical work to a much greater extent than has been done. Studies such as Landes (1998) contain an immense amount of information, one should also say wisdom, albeit not in the form of formal statistical statements. My belief is that historical analyses of this type can provide strong priors for issues such as variable selection and the parameterization of cross-country parameter heterogeneity. In order to do this formally, it might well make sense to employ Bayesian estimation methods whereby evidence from history is modelled as restrictions on a prior. Even if this integration of qualitative and quantitative analysis does not prove to be workable, there is no question that qualitative studies can supplement econometric analyses when it comes to the assessment of growth theories, which is after all the purpose of the empirical exercises.

***ii.* greater reliance on modern statistical methods**

At the same time, it seems clear that the empirical growth literature can be augmented by greater attention to modern developments in statistics. To give the example to which my current thinking is devoted, the issue of parameter heterogeneity is an active area of research in statistics. The literature on nonparametrics has proposed ways to explicitly model regression parameters which smoothly depend on data. Hastie and Tibshirani (1993), for example, explicitly consider ways to analyze regression models of the type

$$g_i = \gamma_1(Z_1)x_{i,1} + \dots + \gamma_K(Z_K)x_{i,K} + \beta(Z_{K+1})y_{i,0} + \epsilon_i \tag{7}$$

where Z_k is a regressor specific set of variables and $\gamma_i(\cdot)$'s and $\beta(\cdot)$ are arbitrary (subject to smoothness restrictions) variable-specific functions. By choosing initial conditions such as per capita income and the literacy rate among adults as elements of Z_i, one can in principle estimate how stages of development influence growth determinants through their influence on the effects of other growth variables.

A related approach of this type is known as empirical Bayes estimation (cf. Maritz and Lwin (1989)). In the empirical Bayes approach, one conceptualizes the observables R_i and model parameters θ_i as draws from a joint probability measure $\mu(R,\theta)$. The goal of the data analysis is the determination of the posterior probability measure for the model's parameters given the data, i.e. $\mu(\theta \mid R)$. This posterior probability can be written, using the definition of conditional probability as

$$\mu(\theta \mid R) = \frac{\mu(R \mid \theta)\mu(\theta)}{\mu(R)} \tag{8}$$

From the perspective of most econometric practice, the difficulty with Bayesian analysis is how to specify $\mu(\theta)$, which is the prior distribution of the model parameters.

In empirical Bayes, the key assumption is that each of the θ_i's is assumed to represent a separate draw from the marginal distribution $\mu(\theta)$.

This means that one can think about recovering $\mu(\theta)$ from the data. Hence, empirical Bayes methods do not need to rely on the specification of a subjective prior, but rather on the idea that each observation embeds a draw of the parameters from the prior. To see how this recovery can work, notice that the relationship between the prior and the other probabilities in this model can be written as

$$\mu(R) = \int_{\theta} \mu(R \mid \theta)\mu(\theta) \tag{9}$$

Since $\mu(R)$ and $\mu(R \mid \theta)$ (this latter being the likelihood function) are estimable, one can see that it is reasonable to ask whether $\mu(\theta)$ (which in principle can be conditioned on data as well) can be recovered from this functional equation. One key question in empirical Bayes analysis is the determination of conditions under which this can be done.

I do not know whether either of these approaches will prove to be particularly useful in studying growth, (although my current research is involved with investigating both). The point for the growth literature, however, is that statistical tools exist for constructively dealing with heterogeneity in national growth experiences – tools which have yet to be adequately exploited.

4. Conclusions

The theoretical and empirical growth literature has experienced a remarkable renaissance over the last 15 years. Yet it appears that few of the substantive economic differences between theories have been adjudicated by growth empirics. Greater attention to issues of model specification, heterogeneity, and nonformal sources of information all seem essential if progress is to be made in this direction.

Many of the criticisms which have been made may seem to be generic problems with empirical work in economics as well as other social sciences. Questions of functional form specification, legitimacy of instruments, and parameter heterogeneity can be raised for virtually any study. In fact, the eminent statistician and probability theorist David Freedman has written a number of very strong attacks on empirical social science research through a careful study of the unjustified assumptions which I believe he would say generically underlie social science studies (e.g. Freedman (1981,1991)). My own criticisms are certainly examples of the sorts of objections he has made. Where we differ is that I do not believe that empirical work should be assessed in black and white terms. Assumptions are more or less plausible, not either justified or unjustified. Hence, criticisms of empirical work need to do more than simply identify assumptions. What I believe to be true, and hope has been illustrated in this discussion, is that in the case of the growth empirics, is that issues such as the employment of instrumental variables and heterogeneity are likely to be particularly serious.

This essay has also suggested that Bayesian approaches might well prove to be especially useful in growth contexts. The justifications given for this have not followed traditional Bayesian concerns about coherence and the like, but has rather been based on a pragmatic interest in uncovering more information about the growth process.[6] And to be fair, some of the discussion may reflect my own limited knowledge of Bayesian ideas. Nevertheless, methodological eclecticism seems more than warranted when dealing with issues as complicated as growth.

[6]Indeed, many orthodox Bayesians object to the name "empirical Bayesian."

References

Alesina, A. and D. Rodrik, (1994), "Distributive Politics and Growth," Quarterly Journal of Economics, 109, 465-490.

Barro, R., (1996), "Democracy and Growth," Journal of Economic Growth, 1, 1-27.

Barro, R., (1999), "Inequality, Growth, and Investment," NBER Working Paper no. 7038.

Bernard, A. and S. Durlauf, (1996), "Interpreting Tests of the Convergence Hypothesis," Journal of Econometrics, 71, 161-172.

Breiman, L., J. Friedman, R. Olshen, and C. Stone, (1984), Classification and Regression Trees, New York: Chapman and Hall.

Cameron, S. and J. Heckman (1998), "Life Cycle Schooling and Educational Selectivity: Models and Evidence," Journal of Political Economy, 108.

Canova, F., (1999), "Testing for Convergence Clubs in Income Per-Capita: A Predictive Density Approach, mimeo, University of Pompeu Fabra, Spain.

Durlauf, S. and P. Johnson, (1995), "Multiple Regimes and Cross-Country Growth Behavior," Journal of Applied Econometrics, 10, 365-384.

Durlauf, S. and D. Quah, (1998), "The New Empirics of Economic Growth," working paper no. 6422, National Bureau of Economic Research and forthcoming, Handbook of Macroeconomics, J. Taylor and M. Woodford, eds., Amsterdam: North Holland.

Frankel, J. and D. Romer, "Trade and Growth: An Empirical Investigation," (1999), NBER Working Paper no. 5476.

Freedman, D., (1981), "Some Pitfalls in Large Econometric Models: A Case Study," Journal of Business, 54, 479-500.

Freedman, D., (1991), "Statistical Models and Shoe Leather," in Sociological Methodology, P. Marsden, ed., Cambridge: Basil Blackwell.

Galor, O., (1996), "Convergence? Inferences from Theoretical Models," Economic Journal, 106, 1056-1069.

George, E. and R. McCulloch, (1993), "Variable Selection via Gibbs

Sampling," Journal of the American Statistical Association, 88, 881-889.

Heckman, J., L. Lochner, and C. Taber, (1998), "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents," Review of Economic Dynamics, 1.

Landes, D., (1998), The Wealth and Poverty of Nations, New York: W. W. Norton and Company.

Leamer, E., (1983), "Let's Take the Con Out of Econometrics," American Economic Review, 73, 31-43.

Levine, R. and D. Renelt, (1992), "A Sensitivity Analysis of Cross-Country Growth Regressions," American Economic Review, 82, 942-963.

Lucas, R., (1988), "On the Mechanics of Economic Development," Journal of Monetary Economics, 22, 3-42.

Mankiw, N. G., D. Romer and D. Weil, (1992). "A Contribution to the Empirics of Economic Growth," Quarterly Journal of Economics, CVII, 407-437.

Maritz, J. and T. Lwin, (1989), Empirical Bayes Methods, New York: Chapman and Hall.

Mauro, P., (1995), "Corruption and Growth," Quarterly Journal of Economics, CX, 681-713.

Mitchell, T. and J. Beauchamp, (1988), "Bayesian Variable Selection in Linear Regression (with discussion), Journal of the American Statistical Association, 83, 1023-1036.

Miller, A., (1990), Subset Selection in Regression, New York: Chapman and Hall.

Quah, D., (1996a), "Convergence Empirics Across Economies with (Some) Capital Mobility," Journal of Economic Growth, 1, 95-124.

Quah, D., (1996b), "Empirics for Growth and Economic Convergence," European Economic Review, 40, 1353-1375.

Quah, D., (1997), "Empirics for Growth and Distribution: Polarization, Stratification, and Convergence Clubs," Journal of Economic Growth, 2, 27-59.

Romer, P., (1986), "Increasing Returns and Long Run Growth," Journal of Political Economy, 94, 1002-1037.

Sala-i-Martin, X., (1997), "I Just Ran Two Million Regressions," American Economic Review, Papers and Proceedings, 87, 178-83.

Tobin, J., (1970), "Money and Income: Post Hoc Ergo Propter Hoc?" Quarterly Journal of Economics, 84, 301-317.