Public Disclosure Authorized

Public Disclosure Authorized

WPS 2224

POLICY RESEARCH WORKING PAPER 2224

Growth Forecasts Using Time Series and Growth Models

Aart Kraay
George Monokroussos

It is difficult to choose the "best" model for forecasting real per capita GDP for a particular country or group of countries. This study suggests potential gains from combining time series and growth-regression-based approaches to forecasting.

The World Bank
Development Research Group
Macroeconomics and Growth
November 1999

Summary findings

Kraay and Monokroussos consider two alternative methods of forecasting real per capita GDP at various horizons:

- Univariate time series models estimated country by country.
- Cross-country growth regressions.

They evaluate the out-of-sample forecasting performance of both approaches for a large sample of industrial and developing countries.

They find only modest differences between the two approaches. In almost all cases, differences in median (across countries) forecast performance are small relative to the large discrepancies between forecasts and actual outcomes.

Interestingly, the performance of both models is similar to that of forecasts generated by the World Bank's Unified Survey.

The results do not provide a compelling case for one approach over another, but they do indicate that there are potential gains from combining time series and growth-regression-based forecasting approaches.

This paper — a product of Macroeconomics and Growth, Development Research Group — is part of a larger effort in the group to improve the understanding of economic growth. Copies of the paper are available free from the World Bank, 1818 H Street, NW, Washington, DC 20433. Please contact Rina Bonfield, room MC3-354, telephone 202-473-1248, fax 202-522-3518, email address abonfield@worldbank.org. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/research/workingpapers. The authors may be contacted at akraay@worldbank.org or gmonokroussos@worldbank.org. November 1999. (32 pages)

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the view of the World Bank, its Executive Directors, or the countries they represent.

Produced by the Policy Research Dissemination Center

Growth Forecasts Using Time Series and Growth Models

Aart Kraay

George Monokroussos

The World Bank

1. Introduction

In developed countries, a vast range of forecasting tools have been used to predict growth and other economic variables of interest. In contrast, growth projections for many developing countries are typically based on much more informal techniques. For example, both the World Bank and the International Monetary Fund rely largely on the informed judgement of their country economists to produce forecasts for internal and external use.[1] In this paper, we consider two simple formal models for forecasting growth in a large sample of developed and developing countries: univariate time series models estimated country-by-country, and cross-country growth regressions. The time series models constitute a useful benchmark which illustrates how well forecasts based on extremely limited information (only the history of per capita GDP itself) can perform. The growth regressions are of interest given the vast empirical literature which argues that a significant fraction of the cross-country and time series variation in longer-term growth rates can be explained by a fairly parsimonious set of explanatory variables. A natural question to ask is whether this popular empirical framework has any value for predicting future growth.

We consider the relative forecast performance of two straightforward models. Our time series model is very simple, and models (the logarithm of) real per capita GDP as following a first-order autoregressive process around a broken trend. We estimate this model country-by-country for 112 countries, for two time periods: 1960-1980, and 1960-1990. We then generate out-of-sample forecasts for the remaining years through 1997 based on these two information sets, and compare these forecasts with actual outcomes. Our growth model follows the vast empirical literature spawned by the neoclassical growth model. We estimate a dynamic panel regression of (the logarithm of) real per capita GDP on itself lagged five years, and a number of lagged explanatory variables which proxy for the steady-state of the neoclassical growth model and capture the effects of various policies on long-run growth: investment, population growth, trade openness, inflation, and the black market premium. We estimate this model using nonoverlapping quinquennial averages of data over the same two periods as for the time series model (although for a somewhat smaller sample of countries as dictated by data availability), and then generate forecasts for the remaining years in the sample which can be compared to actual outcomes. In order to benchmark the forecasts generated by these models against current practice, we also make some comparisons with long-term forecasts produced by the World Bank's Unified Survey in 1990. However, our primary interest is in the relative performance of the time series and growth models. [2]

---

[1] The World Bank's Unified Survey projections, and the IMF's World Economic Outlook projections are produced in this way. Both organizations also use large macroeconometric models: the World Bank's Global Economic Model (GEM) is used to produce forecasts appearing in the Bank's annual Global Economic Prospects publication, and the IMF maintains MULTIMOD for research and simulation purposes.

We assess the out-of-sample forecast performance of these models using standard summary statistics which capture their bias and mean squared error. These statistics suggest small median (across countries) differences in forecast performance of the alternative models, which vary with the forecast horizon. For example, there is some evidence -- consistent with our priors -- that the mean squared error of growth regression based forecasts is smaller at long forecast horizons (five years or more). However, these differences in median forecast performance are typically very small relative to the cross-country dispersion in forecast performance, casting doubt on the significance of observed "typical" differences. The relative performance of the alternative forecasting models is also very unstable over time within countries. We test for and do not reject the null hypothesis that the past relative performance of the growth model and the time series model in a particular country is independent of the future relative performance of the two models in that country.

These results indicate that neither forecasting model dominates, both across countries and within countries over time. Rather than attempt to choose a single "best" forecasting model, we instead ask whether there is value in combining the forecasts of alternative models. We implement forecast encompassing tests and find evidence that these approaches can "learn from each other", in the sense that the forecasts from both models are jointly significant in explaining actual outcomes. This is especially true at shorter horizons, and it suggests that there are potential benefits from combining these forecasts in some way to arrive at a superior overall method.

---

[2] For a more systematic assessment of the quality of World Bank forecasts, see Ghosh and Minhas (1993), and Verbeek (1999). Artis (1996) does the same for the IMF's short-term forecasts.

The remainder of this paper proceeds as follows. In the next section, we present the two models used to produce growth forecasts, and note the similarities and differences between them. In Section 3, we examine the cross-country performance of these forecasts using various summary statistics. In Section 4, we illustrate the results of our forecast encompassing tests, and consider whether a combined forecast can outperform either of the two alternatives. We also briefly consider whether the absolute performance of either model is adequate. Section 5 offers some concluding remarks.

2. Forecasting Models

In this section we describe the simple time series and growth models we use to forecast real per capita GDP in a large sample of developed and developing countries.

2.1. Time Series Forecasts

For each country, we estimate a very simple first-order autoregressive process around a linear trend, allowing for the possibility that the trend of the series changes once within the estimation period. In particular, we assume that the logarithm of real per capita GDP in country i at time t, y_{it}, is described by the following process:

$$(1) \qquad y_{it} = \rho_i \cdot y_{i,t-1} + \delta_{it} + \varepsilon_{it}$$

The trend term δ_{it} is a linear function of time, and both the slope and the intercept term may change at a date T within the estimation period, i.e. $\delta_{it} = \mu_i + \theta \cdot D_T^{\mu} + \beta \cdot t + \gamma \cdot D_T^{\beta}$, where D_T^{μ} is a dummy variable taking on the value 1 if t>T and zero otherwise, and D_T^{β} is a dummy variable taking on the value t-T if t>T and zero otherwise. The two dummy variables pick up a shift in the deterministic component of output that occurs in year T. The date of the trend break, T, is determined endogeneously, using the procedure of sequential Wald tests suggested by Vogelsang (1997).[3] At the estimation stage, we do not need to make strong assumptions about the properties of the error term. However, for the purposes of formal tests of model performance, it will be useful to assume that the error term is independent over time and is normally distributed with variance σ_i^2.

In order to evaluate the forecasting performance of this model, we divide the sample period in two at a particular year t. We then estimate Equation (1) using the data available until this year t, and then use the model to forecast the log-level of per capita

---

[3] However, we do not pre-test for a trend break, i.e. we allow for a trend break at time T even if this break is not statistically significant. There is some evidence that forecasts based on pre-tested models perform better than either of the alternative models that are being pre-tested (Diebold and Kilian (1999) perform Monte Carlo experiments, and Stock and Watson (1998) show this empirically in a large-scale comparison of many forecasting models of various macroeconomic aggregates for the United States). This suggests that the forecasting performance of both the time series model and the growth model might be improved by pretesting.

GDP for each subsequent year. In particular, if we divide the sample in two at year t, our forecast of per capita GDP for each subsequent year is:

$$(2) \qquad \hat{y}_{i,t+s|t} = \hat{\rho}_i^s \cdot y_{it} + \hat{\delta}_{i,t+s}$$

where $\hat{y}_{i,t+s|t}$ denotes the forecast of $y_{i,t+s}$ based on information available at time t and $\hat{\rho}_i$ and $\hat{\delta}_{i,t+s}$ are the parameter estimates for country i based on its data available through year t. Ignoring the uncertainty associated with the parameter estimates, i.e. assuming the parameters of the model are known, the corresponding forecast error is:

$$(3) \qquad e_{i,t+s|t} = \sum_{h=0}^{s-1} \rho_i^s \cdot \varepsilon_{i,t+s-h}$$

The variance of this error term can be used to construct the ex ante forecast confidence intervals associated with each forecast, which will depend on the autoregressive parameter, ρ_i , and the variance of the error term, σ_i^2. Replacing these with their estimates yields the usual ex ante forecast confidence intervals.[4]

Our data consists of a panel of 112 countries for which a complete time series on real per capita GDP adjusted for differences in purchasing power parity is available over the period 1960-1997.[5] We estimate this model twice for each country, once using data over the period 1960-1980, and once over the period 1960-1990. We then generate forecasts of real per capita GDP for the remaining years through 1997 for each country, and compare these forecasts with the actual realizations of per capita GDP for each country.

---

[4] In particular, a 90% forecast confidence interval extends $\pm 1.64 \cdot \sqrt{\sum_{h=0}^{s-1} \hat{\rho}_i^{2 \cdot s} \cdot \hat{\sigma}_i^2}$ around the forecast itself.