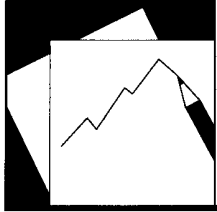


Growth: Now and Forever?



WP/14/117

IMF Working Paper

Growth: Now and Forever?

Giang Ho and Paolo Mauro

INTERNATIONAL MONETARY FUND

IMF Working Paper

African Department and Strategy, Policy and Review Department

Growth: Now and Forever?

Prepared by Giang Ho and Paolo Mauro*

Authorized for distribution by Paolo Mauro and Catherine Pattillo

July 2014

This Working Paper should not be reported as representing the views of the IMF.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

Abstract

Forecasters often predict continued rapid economic growth into the medium and long term for countries that have recently experienced strong growth. Using long-term forecasts of economic growth from the IMF/World Bank staff Debt Sustainability Analyses for a panel of countries, we show that the baseline forecasts are more optimistic than warranted by past international growth experience. Further, by comparing the IMF's *World Economic Outlook* forecasts with actual growth outcomes, we show that optimism bias is greater the longer the forecast horizon.

JEL Classification Numbers: O47

Keywords: growth forecast, optimism bias, forecast horizon, debt sustainability analysis

Authors' E-Mail Addresses: gho@imf.org; pmauro@imf.org

* We gratefully acknowledge helpful comments from Andy Berg, Prakash Loungani, Chris Papageorgiou, Cathy Pattillo, and participants in a seminar in the IMF's African Department. Special thanks to Abebe Selassie for encouraging us to delve into this topic.

Contents	Page
I. Introduction	3
II. Empirical Approach and Preliminary Analysis.....	5
A. Data of GDP Growth Outturns and Forecasts.....	5
B. Estimation of Growth Autocorrelation Coefficients	6
III. Comparing Long-run Growth Forecasts with Recent Growth Outcomes	10
IV. Correlates of Forecast Errors	19
V. Conclusions.....	23
References.....	25
 Tables	
1. Average Per Capita Growth Rates	6
2. Simple and Rank Correlations of Growth Rates Across Periods.....	7
3. Estimates of Growth Persistence for 10-year Horizon	8
4. "Explaining" Forecasts Errors.....	22
 Figures	
1. Persistence of Per Capita Growth Rates, 1950–2010	9
2. Persistence of Per Capita Growth Rates by Income, 1970–2010	9
3. Forecast vs. Past Per Capita Growth, 10-Year Horizon.....	11
4. Forecast vs. Past Per Capita Growth, 20-Year Horizon.....	12
5. Forecast vs. Past Per Capita Growth, 3-Year Horizon.....	13
6. Forecast vs. Past Per Capita Growth, 5-Year Horizon.....	13
7. Optimism vs. Reversion to Mean.....	15
8. Forecast vs. Past Growth for Non-fragile LICs, 10-Year Horizon	16
9. Forecast vs. Past Growth for Non-fragile LICs, 20-Year Horizon	16
10. Corrected Forecast vs. Past Growth, Half Stdev. Shock, 10-Year Horizon	18
11. Corrected Forecast vs. Past Growth, Half Stdev. Shock, 20-Year Horizon	18
12. Rolling Forecast Error by Forecast Horizon, 1990-2012.....	20
13. Mean and Median Forecast Error by Forecast Horizon.....	21

I. INTRODUCTION

Optimism bias and wishful thinking about the future are well documented human tendencies.¹ A specific manifestation of optimism bias is the overestimation of the relevance of recent positive outcomes when predicting future outcomes.² Economic growth forecasts are no exception, and the problem may well become more pronounced for longer term horizons. Indeed, although the ex-post empirical association between a country's economic growth rate in a given decade and in the following one is weak (Easterly and others, 1993), forecasters often predict continued rapid economic growth into the medium and long term for countries that have recently experienced strong growth. Drawing on these observations, Pritchett and Summers (2013) have recently argued that, for example, most medium- and long-term economic growth forecasts available at the time of writing for China and India—where growth has been exceptionally high for more than a decade—are overly optimistic. These authors suggest that longer term forecasts for these two countries fail to take into account the “reversion to the mean” effect, whereby exceptional performance tends to dissipate.³ Their argument applies more generally, however, and is worth investigating on the basis of a broader data set and longer-term projections. It is also worth analyzing potential systematic biases in forecasts regarding countries whose recent growth has been unusually weak or negative.

In this paper, we gauge the degree of optimism bias—and the extent to which the persistence of strong growth may be overestimated—in economic forecasts at horizons of increasing length. We are especially interested in projections made over longer-term horizons; thus, we analyze economic growth forecasts for horizons of up to twenty years, which we draw from

¹ Thaler (2000) reports that on the first day of class all MBA students expect that their grades will be above the median. Sharot (2012) shows that most individuals display optimism bias when estimating their own chances of success in various aspects of personal life. She finds that an optimistic attitude can lead not only to greater happiness, but also to objectively better outcomes on average, though it can also lead individuals to incur excessive risks.

² This tendency has long been familiar to psychologists: people usually expect that a sequence of events generated by a random process will resemble the essential characteristics of that process even when the sequence is short; they see trends and patterns even when the sequence is truly random (“excessive extrapolation”) (Tversky and Kahneman, 1974). Such tendencies are also well documented in investment allocation choices. For example, purchases of mutual funds are unduly influenced by recent good performance, even though performance shows no persistence (Patel and others, 1991). Moreover, employees whose firms experienced the best stock performance over the last ten years allocate a higher share of their discretionary contributions to their 401(k) accounts to their own company's stock than is the case for employees of firms that experienced the worst stock performance (Benartzi, 2001).

³ Galton (1886) first documented this phenomenon, by showing that children of tall parents tend to be shorter than their parents. Throughout the present paper, we use the expression “reversion to the mean” rather than “regression to the mean” to simplify exposition given that we also estimate empirical relationships using regression analysis.

the debt sustainability analysis (DSA) exercises routinely undertaken by IMF and World Bank teams for a large sample of countries.

Projecting a country's economic growth into the medium term and beyond is notoriously difficult. At the same time, getting the growth projections wrong has major adverse consequences. For example, overestimating future economic growth implies underestimating the government debt-to-GDP ratio that will be reached at the end of the projection period (in the absence of corrective policy measures). As a result, either the country will end up with a higher-than-expected debt ratio, which could result in a debt crisis, or future policymakers will have to tighten fiscal policy abruptly—with disruptive consequences—at a later stage. To illustrate the magnitude of the potential impact, consider a country that under current policies and assuming a given growth rate will have a stable government debt-to-GDP ratio projected at, say, 50 percent of GDP twenty years from now. In the event that average economic growth turns out to be 0.5 percentage point lower for the next twenty years, and assuming that a one percentage point decline in GDP results in a higher deficit by 0.4 percentage point of GDP, then the debt-to-GDP ratio at the end of the decade would be above 90 percent.⁴ Indeed, historically, several debt crises have been brought about by declines in economic growth that were likely unanticipated by policy makers (Easterly, 2001 and 2013).

Several studies have analyzed the accuracy (and biases) of one- or two-year ahead economic growth forecasts (e.g., Timmermann, 2007). Of particular relevance here is the result that forecasters seem to have an especially difficult time predicting turning points in the economic cycle. For example, Loungani and Juhn (2002) show that two-thirds of economic recessions in a large sample of countries remained undetected by April of the year in which the recession occurred. Few systematic studies have analyzed longer-term economic forecasts. Frenkel (2011) analyzes optimism bias at horizons of up to three years in a sample of 33 countries (for the most part, advanced economies). Pritchett and Summers (2013) analyze five-year forecasts drawn from the IMF's *World Economic Outlook* (WEO). Broader reviews of IMF forecasts (including, to some extent, forecasts over horizons up to five years) are periodically undertaken (see, for example, Independent Evaluation Office, 2014, including background studies and cited works). In what follows, we consider forecasts over horizons of up to twenty years from DSA exercises; in addition, we provide a more in-depth analysis of the ex-post differences between outcomes and WEO forecasts at the five-year horizon.

⁴ For this illustrative exercise, we assume that government fiscal policy is not tightened in response to the decline in economic growth; rather, the automatic stabilizers are allowed to operate fully.

Our key findings are the following:

- Longer horizon economic growth forecasts are more optimistic than would be warranted by projecting future growth on the basis of a panel autocorrelation model (in which growth is regressed on its lag).
- Projections are optimistic all round—both for countries whose recent growth has been above average (failure to consider reversion to the mean) and for countries with below-average past growth (projected acceleration is greater than would be indicated by reversion to the mean).
- Optimism bias is greater the longer the projection horizon.
- Comparing projections with actual outcomes, optimism bias tends to be greater for (i) longer horizons; (ii) forecasts made while output is below trend; (iii) countries about to enter an IMF-supported program.

The remainder of this paper is organized as follows. In the next section, we describe the data and empirical approach; in addition, using actual data, we estimate how quickly countries' growth rates revert to the sample mean, a key ingredient for our analysis to follow. In Section III, we compare long-run growth forecasts from the IMF/World Bank debt sustainability analyses with forecasts obtained by applying the estimated degree of reversion to the mean from actual data, thereby documenting optimism bias. In Section IV, to delve deeper into the correlates of optimism bias, we turn to growth forecasts from the WEO database for horizons of up to five years, and compare them with actual outcomes. In Section V we conclude by speculating on possible factors underlying optimism bias in forecasts of long-run economic growth.

II. EMPIRICAL APPROACH AND PRELIMINARY ANALYSIS

As a preliminary step, we use actual, ex-post data on economic growth to estimate the empirical relationship between growth in one period and growth in the next. This will allow us to derive forecasts that take fully into account the “reversion to the mean” effect as observed in past data. In Section III, we will then compare previously published growth forecasts with the most recent observations of growth, as well as with forecasts prepared on the basis of the empirical relationship we estimated in this section.

A. Data on GDP Growth Outturns and Forecasts

All data throughout this paper refer to annual real GDP growth. We use three main data sets. The first consists of actual data for 1950–2012 (subject to data availability) for 188 countries from the WEO database as of December 2013 (supplemented by per capita real GDP growth rates from the Penn World Table version 8, not PPP-adjusted; least-squares growth rates using the methodology explained at <http://data.worldbank.org/about/data->

[overview/methodologies](#)). The second consists of the WEO projections (April vintages of 1990-2012) for the following five years, for 188 countries. And the third consists of the projections underlying the debt sustainability analyses jointly prepared by IMF and World Bank teams, for the vintages beginning in 2006. These analyses are routinely prepared for low-income countries at least once a year for countries that have an IMF program, otherwise every time the country has an Article IV consultation with the IMF (typically on an annual or biennial cycle). The sample for this third data set consists of 70 countries, for a combined total of 308 projection paths over 20-year horizons. Population growth projections are from the United Nations (medium fertility specification). Throughout the paper, we define low-income countries as those eligible to borrow from IDA (International Development Association—the part of the World Bank that helps the world’s poorest countries); middle income countries as those classified as such by the World Bank at the beginning of each period or decade we consider; and high-income countries as the residual. The list of fragile countries—those facing severe development challenges including weak institutional capacity, poor governance, and political instability—follows the standard IMF definition as in IMF, 2011.

B. Estimation of Growth Autocorrelation Coefficients

We begin by using actual data on economic growth for a large panel of countries during 1950–2010 to estimate the extent to which past growth helps to predict future growth. This will allow us to establish empirically the speed with which economic growth “reverts to the mean” in actual data.

Average growth rates of per capita income are reported in Table 1 for high income, middle income, and low income economies. As is well known, average growth did not display consistent differences across income groups over the past few decades (when such groups are appropriately defined on the basis of data at the beginning of each period). In other words, there is no evidence of unconditional convergence in per capita incomes—see, for example, Barro (1990).

Table 1. Average Per Capita Income Growth Rates
(percent per annum)

	1970s	1980s	1990s	2000s
High income	1.6	1.7	2.0	1.5
Middle income	2.7	1.0	1.9	3.0
Low income	2.0	0.7	1.2	3.1

Sources: *World Economic Outlook*, IMF, and Penn World Tables.

Moreover, the (simple and rank) correlation coefficients between individual countries’ growth rate in one decade with the previous decade are low, ranging between 0.0–0.5 depending on the sample and period under consideration (Table 2). These coefficients are remarkably similar to those reported for a slightly different sample, more than two decades

ago, by Easterly and others (1993). Such low correlation coefficients are a preliminary indication that past growth performance provides limited predictive power for future growth.

Table 2. Simple and Rank Correlations of Growth Rates Across Periods

	1970s with 1980s		1980s with 1990s		1990s with 2000s	
	Simple	Rank 1/	Simple	Rank	Simple	Rank
All countries	0.34	0.32	0.37	0.41	0.24	0.29
All countries, non-oil 2/	0.43	0.41	0.48	0.48	0.31	0.35
High income	0.33	0.31	0.02	0.09	0.24	0.40
Middle income	0.27	0.24	0.53	0.46	0.36	0.30
Low income	-0.07	-0.03	0.13	0.18	0.20	0.38

Sources: *World Economic Outlook*, IMF, and Penn World Tables.

1/ Spearman rank correlation

2/ Using WEO definition of fuel exporters

To get a more precise gauge of persistence in countries' per capita income growth rates, we estimate the autocorrelation coefficient for per capita income growth rates in panel regressions for 1950–2010. We first conduct the estimation for economic growth at the annual frequency; we then repeat the exercise for average economic growth rates over two-year periods (that is, we regress average growth, for example, in 1982–83 on average growth in 1980–81, etc...); and so on, up to regressing average growth for a twenty-year period on average growth for the previous twenty-year period.

To use a formal expression, we thus estimate several panel regressions for different horizons denoted by k , as follows:

$$g_{t,t+k} = \alpha + \beta_k g_{t-k-1,t-1} + \gamma \ln y_{t-1} + \varepsilon_t$$

in which $g_{t,t+k}$ represents the average growth rate between year t and year $t+k$ (where the horizon $k = 0, 1, 2, \dots, 9, 10$); $\ln y_{t-1}$ denotes the logarithm of the level of per capita real GDP in year $t-1$ to allow for a convergence effect whereby poorer countries tend to grow faster, other things equal; and ε_t is the error term. The autocorrelation coefficient, β_k , is an estimate of how persistent growth is from a k -year period to the next.

In our estimates, we find that past growth has limited predictive power for future growth, at all horizon lengths considered. Using data at the 10-year horizon, the β coefficient amounts to 0.2–0.4 depending on the specification, with relatively small standard errors (Table 3). Our estimated β coefficients are also close to those originally estimated by Easterly and others (1993).⁵

⁵ This is impressive robustness: the addition of twenty years of data since the original estimates by Easterly and others (1993) does not alter the results much.

We obtain broadly similar coefficients using various techniques: overlapping observations (with Newey-West standard errors to correct for autocorrelation in the error terms); GDP-weighted regressions that assign more weight to larger countries (whose growth rates may be more persistent over time, possibly because their economies tend to be more diversified and less likely to be overwhelmed by individual natural disasters); non-overlapping observations (each decade's growth separately); and pooled decadal growth rates with fixed effects for each decade.

Table 3. Estimates of Growth Persistence for 10-year Horizon^{1/}

	Panel, overlapping observations 2/			Non-overlapping observations 3/					
	Baseline	Baseline w/ convergence	GDP- weighted	1970s	1980s	1990s	2000s	Pooled, FE	Pooled w/ convergence, FE
Average past 10-year growth	0.31 [0.04]***	0.30 [0.05]***	0.31 [0.00]***	0.38 [0.12]***	0.32 [0.08]***	0.29 [0.07]***	0.22 [0.10]**	0.31 [0.04]***	0.33 [0.05]***
Log initial GDP per capita		0.03 [0.07]							-0.12 [0.09]
Constant	1.26 [0.12]***	1.00 [0.59]*	1.33 [0.01]***	1.21 [0.41]***	0.32 [0.26]	1.49 [0.22]***	2.28 [0.26]***	1.41 [0.28]***	2.35 [0.74]***
Observations	4,853	4,853	4,853	106	142	142	142	585	585
R-squared	0.11	0.11	0.12	0.12	0.12	0.14	0.06	0.18	0.18

1/ Estimation sample consists of 142 countries over 1950-2010 period. Significance at *10%, **5%, and ***1%.

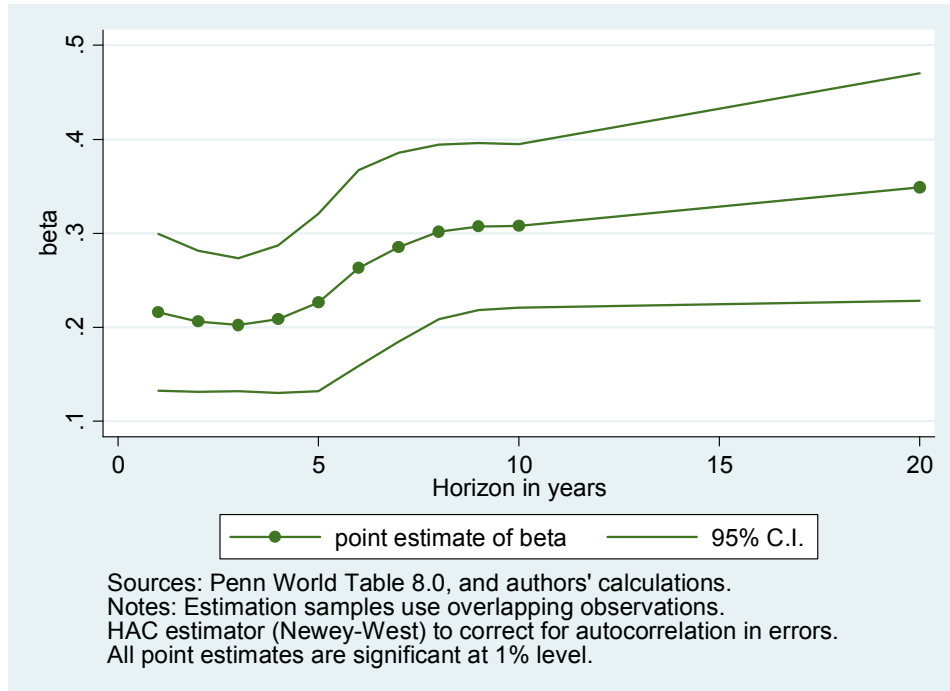
2/ Newey-West standard errors in brackets.

3/ Average growth for each decade is regressed on average growth rate over the past decade. The last specification pools all time periods and includes decade fixed effects.

To illustrate how these coefficients would be used to derive a forecast of economic growth (in the absence of additional information), suppose that a particular country with average per capita income grew by 6 percent over the past decade. Consider the baseline panel regression with overlapping observations (first column): on that basis, the baseline expectation would be for the same country to grow by 3.1 percent on average during the next decade (the constant plus beta times last decade's growth rate: $1.26 + 0.31 * 6 = 3.12$), with a 95 percent probability that the growth rate would be in the range 2.7–3.5 percent.

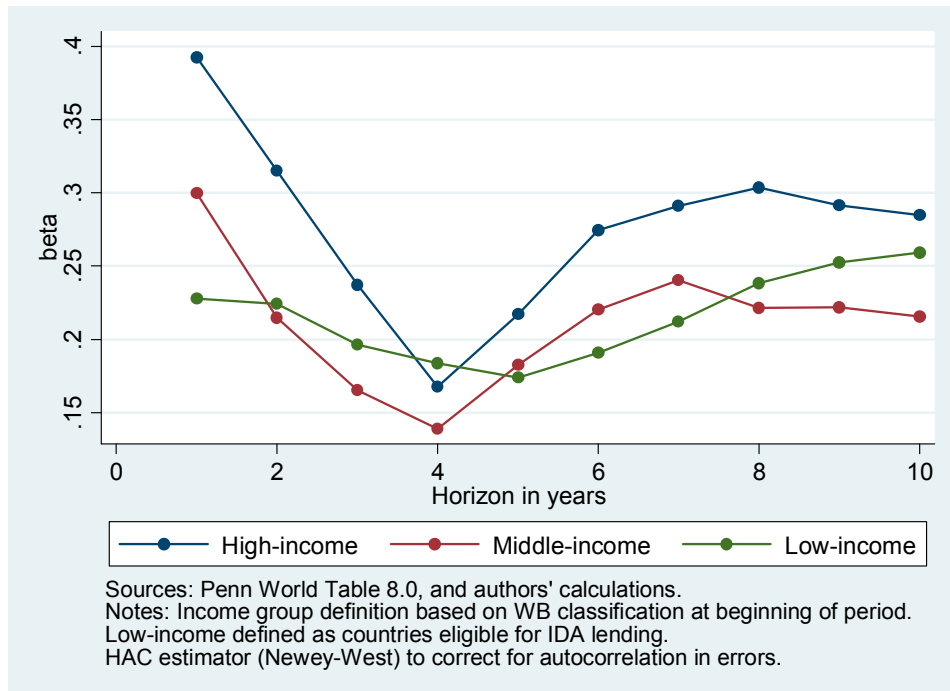
The autocorrelation coefficient, β_k , is generally low but its exact value depends on the horizon k over which forecasts are made. Considering the whole panel of countries for which we have data available, beginning in 1950, at the one year horizon, the autocorrelation coefficient is estimated at 0.2; as the horizon lengthens, the autocorrelation coefficient declines slightly, for horizons of up to three years; it then rises to a maximum of 0.35 for horizons of twenty years (Figure 1).

Figure 1. Persistence of Per Capita Income Growth Rates, 1950–2010



To explore in greater detail how these patterns differ across income groups, we use data beginning in 1970, so as to have a sufficient number of observations for each income group (Figure 2).

Figure 2. Persistence of Per Capita Growth Rates by Income, 1970–2010



Over this estimation period, the “U” pattern in the figure is more pronounced. Persistence is generally higher for high-income economies than for middle-income or low-income economies.⁶ In part, this may stem from the high-income economies’ greater economic diversification and economic size (which reduces the impact of natural disasters on the economy as a whole) and the lower frequency of severe political conflicts. In all cases, persistence declines up to horizons of four to five years, and then rises again for the longer horizons. A possible interpretation of the “U” pattern is that it reflects the dominance of cyclical output fluctuations at horizons of four to five years. This would be consistent with the observation that it is most visible for the high-income economies, whereas cyclical fluctuations are less relevant for middle- or low-income economies.

III. COMPARING LONG-RUN GROWTH FORECASTS WITH RECENT GROWTH OUTCOMES

We can now turn to a comparison of growth forecasts for the medium and long run as jointly prepared by IMF and World Bank teams with those that one could obtain by drawing on the regression results from the previous section. The point here is not that one should use the mechanistic projections based on the regressions estimated above, ignoring the additional in-depth information available to individual country teams. Rather, by documenting a systematic discrepancy between individual country teams’ projections and those obtained mechanistically, it is to alert forecasters to be aware of systematic biases, while retaining valuable insights from in-depth, country-specific information. We return to the interpretation of the results and its implications for improving projections later in the paper. To report the results, we broadly follow the approach used by Pritchett and Summers (2013) and apply it to our longer-term forecasts.


We first focus on growth projections over 10-year horizons. Figure 3 plots the 10-year per capita growth forecasts underlying all DSA vintages for low-income countries issued between 2006 and 2013 against the average per capita growth rates over the previous decade.⁷ As the DSA growth projections are made on the basis of total growth, we convert them into per capita terms using population projections from the UN. To get a visual summary of the information provided, the chart displays several lines representing different

⁶ Higher volatility and lower persistence of growth rates in non-advanced economies compared with advanced economies are consistent with previous studies. For example, Pritchett (2000) discusses the “mountains, cliffs, and plains” that characterize developing economies’ growth patterns. Aguiar and Gopinath (2007) argue that large shocks to trend growth drive the business cycles in developing economies, as opposed to the transitory fluctuations around a stable trend observed in more advanced economies. Berg and others (2012) show that in historical data, the average duration of “growth spells” (defined as periods of high and sustained growth) tends to be considerably shorter for countries in Sub-Saharan Africa and Latin America compared to industrial economies and emerging Asia, and that across countries there were at least as many occurrences of “growth down-breaks” as “growth up-breaks.”

⁷ For example, a country with a DSA published in 2008 will be represented by its growth forecast for 2008–17 and its past growth computed over 1998–2007.

benchmarks. First, the 45-degree line (black line) representing “naïve” forecasts in which the future would look exactly like the past. Second, the line of best fit through the IMF/World Bank staff forecasts of economic growth (“fitted line,” green line). Third, the forecasts based on our baseline estimates of growth persistence for all countries in the previous section (“reversion to the mean,” red line).⁸ With the green line visibly above the red line, IMF/World Bank country teams predict better growth performance in the future than would be implied by a mechanistic projection based on the model discussed above. The slope of the red line, our β estimate at the 10-year horizon, is 0.31 (standard error 0.04), lower than the slope of the green line, 0.37. The intercept (forecast growth corresponding to zero growth in the previous decade) of the red line is 1.26, lower than the intercept of the green line, 2.2.

Figure 3. Forecast vs. Past Per Capita Growth, 10-Year Horizon (percent)



⁸ Using the estimated “reversion to the mean” coefficient for our sample of low-income countries would produce even lower persistence forecasts, and thus larger optimism bias than in the figures above.

To illustrate the difference, consider a country that grew over the past decade at an average rate equal to the mean per capita growth rate of the sample considered (2.4 percent). Corresponding to 2.4 percent past growth, the forecast is 3.1 percent according to the green line and 2.0 percent according to the red line. That is, the typical DSA forecast would predict this country to grow at an average rate of 3.1 percent in the next decade, compared to 2 percent predicted by our reversion to the mean framework, so that the optimism bias in this example would be 1.1 percentage points. The bias tends to get larger for countries that have experienced more rapid growth in the past.

At the 20-year horizon, optimism bias becomes even greater (Figure 4). Conversely, the bias is visibly lower at the 3-year and 5-year horizons (Figures 5 and 6, respectively). Indeed, the data suggest a tendency for the bias to increase as the horizon becomes more extended—a finding that we'll return to in Section IV.

Figure 4. Forecast vs. Past Per Capita Growth, 20-Year Horizon (percent)




Figure 5. Forecast vs. Past Per Capita Growth, 3-Year Horizon (percent)

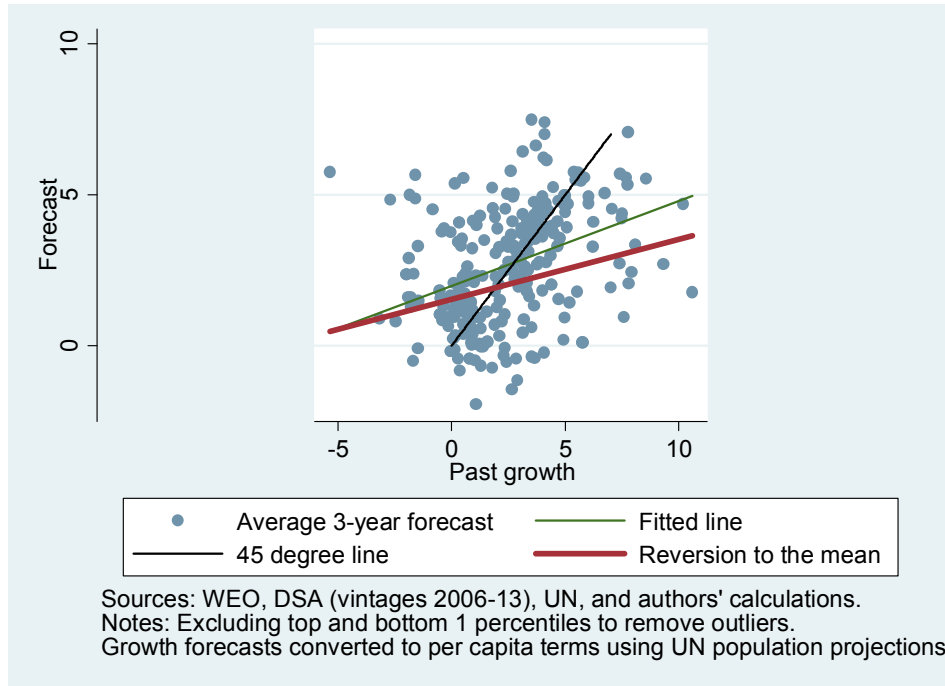



Figure 6. Forecast vs. Past Per Capita Growth, 5-Year Horizon (percent)

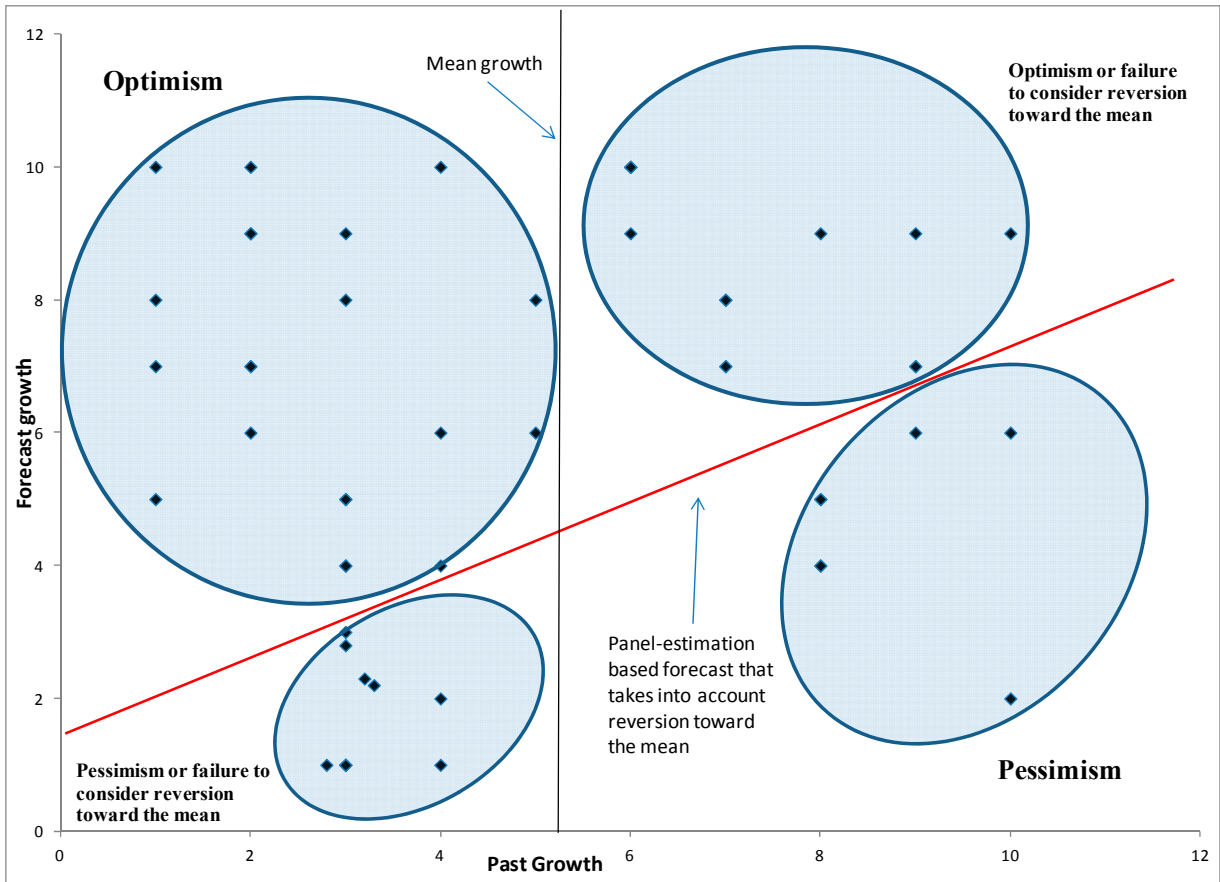


To interpret the empirical findings outlined above, it may be helpful to consider the following illustrative diagram, Figure 7 (in which the data points are fictitious). It is instructive to consider separately the cases of countries whose growth rate over the previous decade was relatively high (points to the right hand side of the “mean growth” vertical bar), and then of those countries whose past growth rate was unusually low (to the left hand side of the “mean growth” vertical bar). The red line represents forecasts that take into account the reversion to the mean phenomenon, on the basis of the estimated regressions above. The dots represent combinations of forecasts underlying the DSA exercises (vertical axis) and past growth (horizontal axis).

Points above the red line represent DSA forecasts that are more optimistic than those that would be produced, for a given past growth rate, on the basis of the regressions reported in Section II.B. Conversely, points below the red line represent overly pessimistic forecasts by comparison to those produced by the regressions. Points in the top-right quadrant are those for which the actual DSA forecasts are more optimistic than those produced taking into account the “reversion to the mean” effect; this is the quadrant emphasized by Pritchett and Summers (2013). However, for countries that experienced unusually weak growth, failure to consider the “reversion to the mean” effect would place them in the bottom-left quadrant. Points in the top-left quadrant are those for which past growth was relatively weak and the DSA forecasts are even more optimistic than would be implied by reversion to the mean. Turning to the data, this is a highly populated quadrant. The bottom-right quadrant represents cases in which past growth was rapid but forecasters expect to fall below average. In practice, this is a scarcely populated quadrant.

On the whole, the results suggest that the failure to consider reversion to the mean is asymmetric: when past growth was strong, forecasters expect it to continue; but when past growth was weak, forecasters expect it not only to revert to the mean, but to exceed the mean. Thus, our results suggest optimism bias throughout—for both countries that grew rapidly in the recent past (observationally equivalent to failure to consider reversion to the mean) and for countries that grew slowly in the recent past (where forecasters expect better growth than would be implied by reversion to the mean).

Figure 7. Optimism vs. Reversion to Mean



The empirical observation that forecasts tend to display optimism bias remains strong when fragile low-income countries are excluded from the estimation sample, to rule out instability of economic growth that results from special conflict episodes (Figures 8 and 9).

Figure 8. Forecast vs. Past Growth for Non-fragile LICs, 10-Year Horizon (percent)





Figure 9. Forecast vs. Past Growth for Non-fragile LICs, 20-Year Horizon (percent)



Given the observed optimism bias in the baseline forecasts by IMF/World Bank country teams, are there ways in which the adverse consequences of such bias can be mitigated? In the context of the debt sustainability analyses for which the long-term growth forecasts are prepared, it is already common practice for IMF/World Bank teams to assess whether the debt remains sustainable under several “adverse scenarios” prepared using a methodology common to all countries. One such scenario consists of assuming that the economic growth rate in each year going forward will be lower than in the baseline scenario by one standard deviation of past growth. To gauge the extent to which this helps to temper optimism bias, we plotted the combination of past growth and IMF/World Bank teams’ baseline forecasts minus one standard deviation of the growth rate for the country in question. This led to overcorrection—imparting excessive caution—especially for countries that experienced slow or negative growth in the recent past.

Thus we experimented with a smaller degree of correction and found that subtracting half a standard deviation from the growth forecast is about “right,” in the sense that it lowers the forecast to a level consistent with that implied by our “reversion to the mean” framework. To show this, we plot the fitted line through these data points (green line). For the 10-year horizon (Figure 10), the red and green lines essentially touch each other for countries that have been growing rapidly in the past, and the red line is only slightly above the green line for countries that have experienced slow or negative growth. For the 20-year horizon (Figure 11), the two lines are essentially indistinguishable. Thus, a more appropriate baseline seems to be about half way between the current forecasts and the “low-growth” scenario imposing a one standard deviation correction on the current baseline growth forecast.

Figure 10. Corrected Forecast ($\frac{1}{2}$ Std. Dev.) vs. Past Growth, 10-Year Horizon (percent)





Figure 11. Corrected Forecast ($\frac{1}{2}$ Std. Dev.) vs. Past Growth, 20-Year Horizon (percent)



IV. CORRELATES OF FORECAST ERRORS

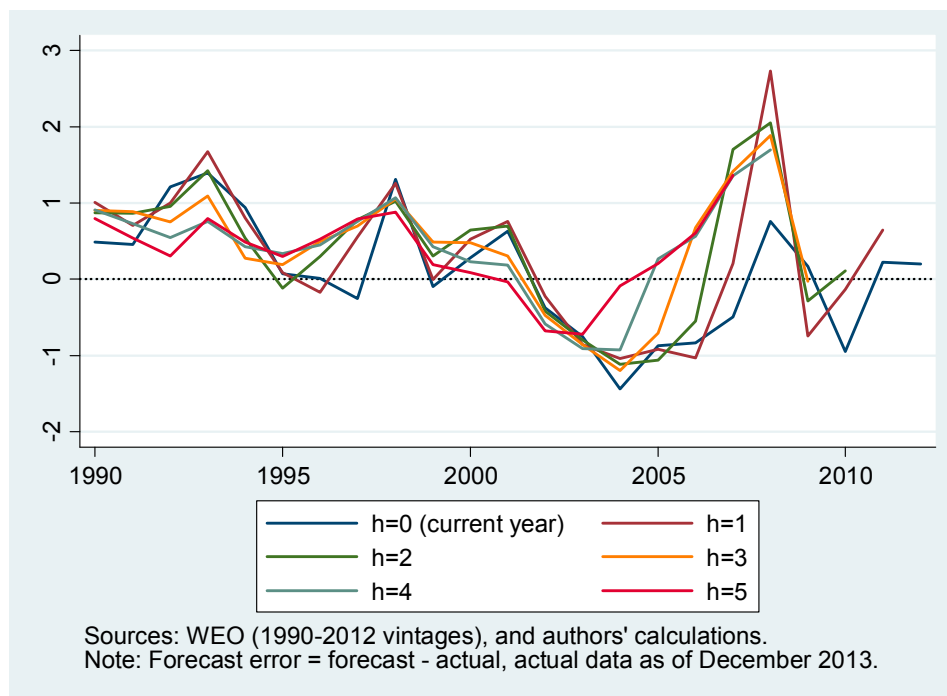
Thus far, we have made the case that today's forecasts look optimistic by comparison with past actual growth performance. But what if there are good reasons to expect that the future will be better than the past? We cannot assess that claim directly. However, as we show in this section, we have had a pretty consistent record of forecasts that turned out to be optimistic by comparison with actual outcomes. To do so, we turn to five-year forecasts from the WEO database, because twenty-year forecasts are available to us only beginning in 2006, so that it is too soon to assess outcomes.

We use the April published version of all forecast vintages between 1990 and 2012 reported by all 188 member countries. Forecasts in the WEO database are available for different horizons, ranging from current year to five years ahead. We define forecast errors as forecasts minus outcomes, so that a positive forecast error signifies ex-post optimistic bias.

A preliminary look at the data shows that the average forecast errors across countries tend to be positive (overly optimistic) for the majority of the 1990-2012 period, except for a few years of strong actual growth in the mid 2000s (Figure 12).⁹ The optimism bias was exacerbated during the global financial crisis, consistent with the documented challenges involved in predicting economic downturns and recessions (Loungani and Juhn, 2002). In the subsequent analysis, we discard the forecasts made in 2008 and subsequent years; this allows us to focus on projections made in relatively "normal" periods, and it gives us the same number of forecast error observations across all horizons.

⁹ Using median forecast errors (to mitigate the effect of large outliers) gives a qualitatively similar picture.

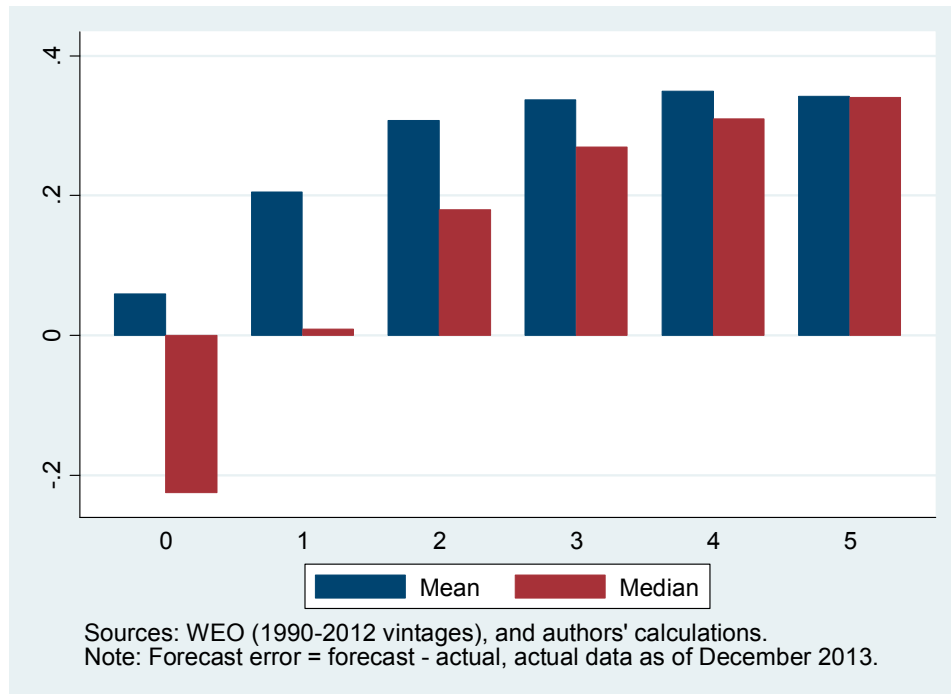
Figure 12. Rolling Forecast Error by Forecast Horizon, 1990-2012
(percentage points, annual average)



The tendency for growth forecasts to be overly optimistic has been documented (at least for short horizons) in several studies (e.g., Timmerman, 2007). However, what has rarely been discussed is the extent to which optimism bias varies with the horizon over which forecasts are made. Do the larger uncertainties associated with longer horizons give more room for forecasters to engage in wishful thinking about future growth prospects? Frankel (2011) documented for a sample of 33 advanced economies that government growth forecasts for budgetary planning purposes tend to be more optimistic at the 3-year horizon than at shorter horizons. We show that this finding applies more generally to our broader data set of a large number of countries and time periods, as well as longer horizons.

As a first step, we examine the pattern of average forecast errors for different horizons (Figure 13). On average, forecasts of current-year growth rate (made around April of the same year) have a positive bias of 0.06 percentage point. The bias increases to 0.2 percentage point for next-year forecasts, 0.3 percentage point for 2-year ahead forecasts, and averaging about 0.34 percentage point for horizons from three to five years. Using median instead of mean forecast errors reduces somewhat the size of the positive bias, but the pattern with respect to horizon (more optimistic at longer horizons) remains unchanged.

Figure 13. Mean and Median Forecast Error by Forecast Horizon (percentage points)



Next, we test to see whether the increase in optimism bias for longer horizons is statistically significant, controlling for a variety of country characteristics that may systematically influence the forecast errors. The exercise is also useful for exploring the sources of optimism bias, which might point to potential improvements in future projections. Specifically, we estimate a set of panel regressions in which the dependent variable is the WEO forecast errors (defined as above—forecast minus actual outcome) for the 188 countries in the WEO database over the 1990-2012 period. The forecast horizons range from zero (current year forecast) to five years ahead. We report the results in Table 4.

Table 4. “Explaining” Forecast Errors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1-year horizon	0.23 [0.04]***	0.23 [0.04]***	0.23 [0.04]***	0.23 [0.04]***	0.23 [0.04]***	0.23 [0.04]***	0.23 [0.04]***	0.23 [0.04]***
2-year horizon	0.30 [0.05]***	0.31 [0.05]***	0.31 [0.05]***	0.31 [0.05]***	0.31 [0.05]***	0.31 [0.05]***	0.31 [0.05]***	0.31 [0.05]***
Medium-term horizon (3-5 year)	0.33 [0.08]***	0.32 [0.08]***	0.33 [0.08]***	0.58 [0.08]***	0.63 [0.06]***	0.64 [0.06]***	0.77 [0.06]***	0.64 [0.06]***
Positive output gap			-0.29 [0.07]***	-0.24 [0.08]***	-0.19 [0.07]***	-0.20 [0.07]***	-0.19 [0.07]***	-0.19 [0.07]**
Negative output gap			0.27 [0.04]***	0.39 [0.05]***	0.40 [0.05]***	0.40 [0.05]***	0.39 [0.05]***	0.40 [0.05]***
Positive output gap*MT horizon				-0.10 [0.04]***	-0.07 [0.04]*	-0.07 [0.04]*	-0.07 [0.04]*	-0.07 [0.04]*
Negative output gap*MT horizon				-0.24 [0.04]***	-0.23 [0.04]***	-0.23 [0.05]***	-0.23 [0.05]***	-0.23 [0.05]***
Fuel exporters					-0.86 [0.37]**	-0.83 [0.37]**	-0.79 [0.38]**	-0.82 [0.47]*
Fuel exporters*MT horizon					-0.49 [0.27]*	-0.50 [0.27]*	-0.54 [0.28]*	-0.50 [0.27]*
Imminent Fund program						0.55 [0.20]***	0.49 [0.21]**	0.46 [0.21]**
Imminent Fund program*MT horizon						-0.11 [0.12]	-0.04 [0.13]	-0.11 [0.12]
Low-income							0.30 [0.19]	
Low-income*MT horizon							-0.34 [0.12]***	
AFR								0.56 [0.53]
APD								-0.23 [0.41]
EUR								0.20 [0.40]
WHD								0.18 [0.37]
Constant	0.06 [0.10]	0.58 [0.31]*	0.56 [0.31]*	0.44 [0.30]	0.51 [0.30]*	0.41 [0.30]	0.29 [0.31]	0.23 [0.52]
Observations	20,649	20,649	20,649	20,649	20,649	20,649	20,649	20,649
R-squared	0.001	0.036	0.082	0.089	0.099	0.101	0.102	0.107
Year FE	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Source: *World Economic Outlook*, IMF.

Notes. The number of observations is 20,649 in all columns (188 countries, 23 forecast vintages for 1990-2012, six possible horizons from zero to five years; subject to data availability). Robust standard errors (in brackets) clustered at country level. Significance at *10%, **5%, and ***1%. The dummy for medium-term forecast horizon equals one for 3-year, 4-year, and 5-year horizons. The reference (omitted) group is current year forecasts. “Positive output gap” equals the output gap if output is above trend, zero otherwise. “Negative output gap” equals the absolute value of the output gap if output is below trend, zero otherwise. Output gap estimated using the Hodrick-Prescott filter and measured in the year prior to when forecast is made. “Imminent Fund program” takes the value of one if an IMF-supported program is in place the year after the forecast is made. “Low income” and “Fuel exporter” dummies defined in Section II.A. AFR, APD, EUR, and WHD refer to regional dummies for countries in the Sub-Saharan African, Asia-Pacific, European, and Western Hemisphere departments at the IMF. Year FE indicates fixed effects for each individual year during the sample period.

We find that forecast errors are positively and significantly related to the length of the forecast horizon. In column 1, current year forecasts display a bias of 0.06 percentage points, not significantly different from zero (column 1); the bias becomes statistically significant and rises with lengthening horizons: 0.23 percentage point at one year, 0.30 percentage point at

two years, and 0.33 percentage point at three to five years ahead. The finding of increasing bias as the horizon lengthens is robust to changes in specification (columns 2-8).

Forecasts made when the estimated output gap is negative (i.e., in “bad times”) in the previous year, tend to be more optimistic (column 3). In other words, when there is a recession, the forecasters often underestimate the persistence of weak or negative growth, expecting the recovery to come too soon or to be stronger than is often the case. This optimism bias is more prominent in nearer-term forecasts (column 4). Conversely, forecasts made when the estimated output gap is positive in the previous year tend to be more pessimistic, with a larger bias as the horizon lengthens.

Next, we add country characteristics such as economic structure (fuel exporting or not), IMF engagement status (whether the country is about to have an IMF-supported program), income status (low income or not), and geographic region. A few interesting patterns emerge: (i) optimism bias is more pronounced in growth forecasts made prior to a country entering an IMF-supported program; (ii) forecasts made for resource-exporting economies were often too pessimistic during the sample period considered, especially at the longer horizons; and (iii) there are no significant differences in average forecast errors for countries at different income levels or in different regions.

V. CONCLUSIONS

In this paper, we have provided evidence that long-term economic growth forecasts embed sizable optimism bias, to an even greater extent than for nearer-term forecasts. For countries that have experienced rapid economic growth in recent years, optimism bias may be viewed as a tendency to underestimate the importance of the “reversion to the mean” phenomenon. In other words, forecasters seem to overestimate the persistence of rapid economic growth and to give much greater weight to a country’s recent past performance than would be warranted on the basis of the estimated ex-post persistence of economic growth in large samples of countries. Conversely, for countries that have experienced low or negative growth in recent years, optimism bias manifests itself in expectations that growth will not simply revert to the mean but rather will exceed the mean. One could view this as an asymmetry in the failure to consider reversion to the mean, giving rise to all-round optimism.

In this concluding section, we offer some tentative thoughts on the possible factors underlying such over-optimism. It is well known, from the seminal work of Easterly and others (1993), that what the economics profession considers as the most likely “fundamental” determinants of economic growth (such as actual and perceived levels of institutional quality, prudence of macroeconomic policies, educational attainment, and so on) are persistent, whereas economic growth itself is not persistent. Faced with a country that has been growing strongly above the mean, and broadly similar “fundamentals,” forecasters would be hard pressed to justify why they expect that economic growth in the period ahead will be lower than in recent years. Forecasters could of course refer to international experience, as we do in

this paper, and argue that “we don’t know why, we don’t know when, but economic growth eventually reverts to the international mean.” However, the counterargument that it would be inappropriate to change the forecast if the underlying growth factors are unchanged is likely to be rather powerful. Conversely, consider the case of a country whose economic growth has been negative or weak for a few years. Usually such weak performance can be attributed, at least in a proximate sense, to an economic or political crisis, or even a civil war. The forecaster is unlikely to be able or willing to assume that similar crises would recur. Forecasts are rarely constructed as a weighted average of a scenario in which there is no civil war, and a scenario in which a civil war occurs; rather, they are usually made under the implicit assumption that there will be no overwhelmingly negative shock, even though such shocks have occurred in the past and could well occur again (we just don’t know when and what form they would take). It is more likely that analysts make their forecasts assuming the absence of overwhelmingly adverse shock, by their nature hard to predict or quantify.

Turning to the more specific determinants of the extent of over-optimism, based on our analysis of forecasts versus actual outturns, we find an empirical association between an imminent IMF-supported program and over-optimism. Thus, special care must be taken to eschew over-optimism for countries that are about to enter an IMF-supported program, though it is worth emphasizing that the results include programs that went “off-track,” meaning where policy measures were not undertaken as envisaged, which may help to explain why growth turned out to be below what was originally envisaged.

One way to mitigate the potentially adverse implications of over-optimism is to ensure that scenarios (including the public debt dynamics) remain sustainable when economic growth is assumed to be lower than the baseline projection. In the IMF/World Bank DSA exercises, this is already done through a “low-growth” scenario in which economic growth is assumed to be lower than in the baseline by one standard deviation of the growth rate in the country in question. Our results show that a more appropriate baseline would be approximately in the middle between the “low growth” scenario and the baseline forecasts currently produced.

References

- Aguiar, Mark and Gita Gopinath, 2007, “Emerging Markets Business Cycles: The Cycle is the Trend,” *Journal of Political Economy*, Vol. 115, pp. 69-102.
- Benartzi, Shlomo, 2001, “Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock,” *Journal of Finance*, Vol. 56, No. 5, pp. 1747–64.
- Berg, Andrew, Jonathan Ostry, and Jeromin Zettlemeyer, 2012, “What Makes Growth Sustained?” *Journal of Development Economics*, Vol. 98, No. 2, pp. 149-166.
- Easterly, William, 2001, “Growth Implosions and Debt Explosions: Do Growth Slowdowns Explain Public Debt Crises?” *Contributions to Macroeconomics*, Vol. 1, No. 1, Berkeley Electronic Press Journals.
- Easterly, William, 2013, “The Role of Growth Slowdowns and Forecast Errors in Public Debt Crises,” in Alberto Alesina and Francesco Giavazzi, editors, *Fiscal Policy after the Financial Crisis*, University of Chicago Press and National Bureau of Economic Research.
- Easterly, William, Michael Kremer, Lant Pritchett, and Lawrence H. Summers, 1993, “Good Policy or Good Luck? Country Growth Performance and Temporary Shocks,” *Journal of Monetary Economics*, Vol. 32, No. 3, pp. 459–483.
- Frankel, Jeffrey A., 2011, “Over-optimism in Forecasts by Official Budget Agencies and its Implications,” *Oxford Review of Economic Policy*, Vol. 27, No. 4, pp. 536–562.
- Galton, Francis, 1886, “Regression towards mediocrity in hereditary stature,” *The Journal of the Anthropological Institute of Great Britain and Ireland*, Vol. 15, pp. 246–263.
- Independent Evaluation Office of the IMF, 2014, “IMF Forecasts: Process, Quality, and Country Perspectives,” www.ieo-imf.org.
- International Monetary Fund, 2011, “Macroeconomic and Operational Challenges in Countries in Fragile Situations.”
- Loungani, Prakash and Grace Juhn, 2002, “Further Cross-Country Evidence on the Accuracy of the Private Sector’s Output Forecasts,” Vol. 49, No. 1, *IMF Staff Papers*, pp. 49–64.

- Patel, Jayendu, Richard Zeckhauser, and Darryll Hendricks, 1991, “The Rationality Struggle: Illustrations from Financial Markets,” *American Economic Review*, Vol. 81, pp. 232–236.
- Pritchett, Lant, 2000, “Understanding Patterns of Economic Growth: Searching for Hills among Plateaus, Mountains, and Plains,” *The World Bank Economic Review*, Vol. 14, No. 2, pp. 221-250.
- Pritchett, Lant, and Lawrence H. Summers, 2013, “Asiaphoria Meet Regression to the Mean,” Harvard University.
- Sharot, Tali, 2012, *The Optimism Bias*, Vintage.
- Thaler, Richard H., 2000, “From Homo Economicus to Homo Sapiens,” *Journal of Economic Perspectives*, Vol. 14, No. 1, pp. 133–141.
- Timmermann, Allan, 2007, “An Evaluation of the World Economic Outlook Forecasts,” *IMF Staff Papers*, Vol. 54, No. 1, pp. 1–33.
- Tversky, Amos and Daniel Kahneman, 1974, “Judgment under Uncertainty: Heuristics and Biases,” *Science*, New Series, Vol. 185, No. 4157, pp. 1124–31.