Agriculture and economic growth in Africa revisited: A modified Okun's Law

Adusei Jumah and Robert M. Kunst

June 19, 2007

Abstract

In recent years, the relationship between agriculture and economic growth is being reexamined in the literature. The old development mantra-produce more food, feed more people-is giving way to a new call: Create more jobs, provide income to buy food. Because wages in agriculture cannot grow in the presence of an unlimited labor supply at the subsistence wage, the benefits of technological progress in agriculture in terms of improved productivity cannot be reaped, unless economic growth raises aggregate demand. In order to examine the validity of this argument, we consider a modification of Okun's Law in a vector autoregressive framework using data on several African countries. The model is evaluated by means of impulse response analysis and of various forecasting experiments.

1 Introduction

In recent years, the relationship between agriculture and economic growth is being reexamined in the literature. The old development mantra-produce more food, feed more people-is giving way to a new call: Create more jobs, provide income to buy food. Because wages in agriculture cannot grow in the presence of an unlimited labor supply at the subsistence wage, the benefits of technological progress in agriculture in terms of improved productivity cannot be reaped, unless economic growth raises aggregate demand. In order to examine the validity of this argument, we consider a modification of Okun's Law in a vector autoregressive framework using data on several African countries. The model is evaluated by means of impulse response analysis and of various forecasting experiments.

2 Statistical analysis of the data

2.1 Basic stylized facts

We use data on agricultural value added per worker and gross domestic product (GDP) at purchasing power parity—which represents total value added of the economy—for 29 African countries. The set of countries is quite heterogeneous and was dictated by data availability. As a general rule, we confine attention to countries where GDP was available for the full sample of 1975 to 2004 and where agricultural value added was available to 2003 without missing values. Logarithms of these two variables will be denoted by g and v in the following. Summary statistics on growth rates Δg and Δv for all countries are given in Table 1.

Table 1 reveals some basic stylized facts. GDP growth is positive on average for all countries except Congo (Zaire), one of the countries that were affected by civil war. Strong growth is observed for the booming economy of Botswana as well as for some more countries, such as Chad, Lesotho, and others. At the low end of the scale, we find countries with an average real growth around 1%, which is certainly not enough to catch up with the industrialized world in the longer run. For some countries, standard deviations of growth rates are sizeable and reflect the high volatility of the business cycle.

By contrast, almost half of the sample shows negative growth in agricultural value added per worker. For example, in Niger this indicator has declined by almost 2% per year. South Africa, Benin, and Cameroon are the economies with the fastest growth in agricultural productivity.

Figure 1 shows a scatter plot of the 29 observations on average growth rates. Measured correlation among the two variables is positive but low, slightly above 5%. Thus, while there appears to be some interdependence, strong economic growth does not necessarily coincide with strong productivity in agriculture.

Table 1: Descriptive statistics for growth rates.

	GDP		VA in agriculture	
	mean	st.d.	mean	st.d.
Benin	0.038	0.030	0.029	0.066
Botswana	0.085	0.070	0.015	0.103
Burkina Faso	0.037	0.033	-0.001	0.129
Burundi	0.020	0.050	0.006	0.080
Cameroon	0.030	0.063	0.028	0.051
Central Africa	0.008	0.046	0.013	0.028
Chad	0.051	0.140	0.001	0.130
Cote d'Ivoire	0.017	0.051	0.007	0.066
Congo	0.042	0.119	0.014	0.044
Congo (Zaire)	-0.014	0.058	-0.010	0.040
Gambia	0.038	0.062	-0.003	0.110
Ghana	0.030	0.037	-0.004	0.063
Guinea-Bissau	0.021	0.095	-0.005	0.120
Kenya	0.034	0.024	0.008	0.056
Lesotho	0.061	0.081	-0.003	0.110
Madagascar	0.013	0.047	-0.007	0.034
Malawi	0.032	0.054	0.012	0.137
Mauritania	0.029	0.052	0.005	0.102
Niger	0.020	0.058	-0.019	0.116
Nigeria	0.032	0.064	0.016	0.078
Rwanda	0.032	0.158	0.012	0.132
Senegal	0.028	0.039	-0.008	0.130
Seychelles	0.037	0.077	-0.012	0.102
South Africa	0.021	0.025	0.032	0.133
Sudan	0.042	0.058	0.017	0.119
Swaziland	0.045	0.058	-0.003	0.105
Togo	0.023	0.068	0.013	0.069
Zambia	0.012	0.043	-0.001	0.133
Zimbabwe	0.013	0.065	-0.005	0.138

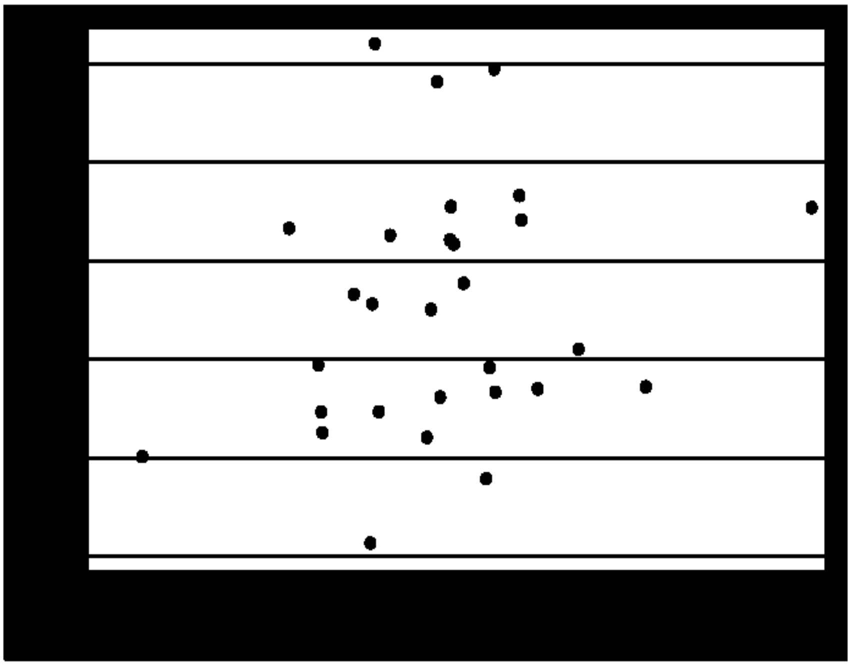

Figure 1: Scatter plot of time averages for Δv and Δg.

2.2 Time-series properties

Macroeconomic variables, such as GDP, are usually treated as first-order integrated (I(1)) in the literature, after some transformation to logarithms. This means that logarithmic differences or growth rates appear to be stationary, while original level data are not. This feature is often corroborated by unit-root tests. Cross-country panels, such as ours, offer a rich variety of panel unit-root tests that allow for various hypotheses, such as unit roots in some or all individual time series.

We subjected the two variables to the battery of unit-root tests that are offered by the software EViews. Results are shown in Table 2.

Table 2: Panel unit-root tests on log level series.

	GDP	VA in agriculture
Levin *et al*	-1.29	-2.02*
Breitung	1.39	-6.72**
Im *et al*	-1.79*	-1.61
ADF Fisher	92.03**	94.02**
PP Fisher	58.76	109.86**
Hadri	8.69*	10.78**

Note: All tests except the Hadri test assume unit roots as their null hypotheses. The Levin *et al.* and the Breitung test use stationarity of all individuals as their alternative, the next three tests use stationarity of some individuals as their alternative. Double asterisks mark rejection of the respective null at the 1% significance level, single asterisks at 5% . All tests use constants and linear trends as deterministic terms.

The results are not entirely conclusive. For the GDP series, results reflect some evidence on stationarity for some countries but the general impression supports the I(1) hypothesis. For the value-added variable, there is more support for stationarity. Given that the series are rather short and contain many irregularities that are not properly captured in the I(0)/I(1) framework, we will proceed with the differenced data. This decision is also suggested by the ultimate aim of forecasting, for which it is known that differencing is advantageous in borderline cases (see Clements and Hendry, 1999).

When the tests are applied to first differences of the data, all of them reject, unfortunately including the Hadri test, whose null hypothesis is stationarity. However, recent simulation evidence does not appear to favor the panel stationarity tests (for example, see Wagner and Hlouskova, 2007). Hence, we will use stationarity of the log differences as our working hypothesis.

If both g and v are I(1) variables, they may be cointegrated, a feature that can be exploited for prediction in error-correction models. There are now some procedures for panel cointegration tests available in the literature. However, due to the heterogeneity with regard to time-series properties across countries we prefer to consider individual cointegration events. We use the popular Johansen (1995) procedure for such cointegration tests, which has the special advantage that it is immune to whether component variables are I(0) or I(1).

The general impression from such cointegration tests is that g and v are typically not cointegrated. In 5 countries, the Johansen procedure yields a rank of 2, which would indicate that the observed level variables v and g are already stationary. In 7 further countries, we get a rank of 1, which would indicate cointegration, and in the remaining 17 cases we do not find cointegration.

The possibly system-stationary cases are Botswana, Kenya, Rwanda, the Seychelles, and Zimbabwe. The cointegrated cases are Central Africa, Cameroon, Ghana, Lesotho, Niger, Nigeria, and South Africa. While there is some evidence on a correspondence between political and economic stability of a country and evidence on cointegration, this issue does not have sufficient statistical support to be of practical interest. We will continue with the working hypothesis that g and v are I(1) and not cointegrated. Hence, our interest focuses on relationships among first differences of the two basic variables Δg and Δv.

3 Panel estimation

3.1 Panel regressions

Given our working hypothesis that both g and v are I(1) variables, our interest concentrates on estimating the basic regression relationship

$$\Delta v_{it} = a + b\Delta g_{it} + u_{it}$$

across the panel, that is, for $i = 1, \ldots, 29$ representing the country index, and for $t = 1975, \ldots, 2004$ representing the time index. Table 3 gives some estimation variants that differ with regard to assumptions concerning u_{it}, as it is customary in panel analysis (for example, see Baltagi, 2005).

In a pooled regression model, u_{it} is uncorrelated across both i and t. In a fixed-effects (FE) one-way model, $u_{it} = \mu_i + v_{it}$ incorporates a deterministic individual-specific level constant μ_i, while the remainder error is still uncorrelated across i and t. In a FE two-way model, $u_{it} = \mu_i + \lambda_t + v_{it}$ also incorporates deterministic time dummies.

Variants with autocorrelated v_{it} can also be considered, as estimation residuals typically indicate correlation effects across time, in violation of the original white-noise assumption. Alternatively, we also experimented with dynamic panel estimators that include lags of the dependent v_{it}. Unfortunately, the properties of these GMM-type estimators are rather complex and tend to become numerically unstable, particularly in panels like ours that are subject to local irregularities. Generally, estimates of b tend to be similar to the reported ones, hence we do not report dynamic estimates in detail.

Random effects (RE) specifications decompose u_{it} into two or three unobserved additive components, with μ_i perfectly correlated over time but independent across i, and λ_t independent over time but perfectly correlated across i. If the regressor Δg is assumed as stochastic, the RE models implicitly assume independence of all error components with the regressor, which is sometimes regarded as unrealistic.

An interesting result from Table 3 is that the most interesting coefficient estimate $\hat{b}$ is almost identical across estimators. In economic terms, the marginal elasticity of productivity in agriculture with respect to total output is around 0.6. The main reason for the limited variation of this coefficient estimate is that variation in panel effects μ and λ is small. Statistical tests fail to reject the pooled regression model against the alternative of one-way FE, and they only marginally reject one-way FE against the two-way FE with time effects. Similarly, Hausman tests fail to reject the RE model against the alternative of the FE model. Autocorrelation tends to be significant but also fails to modify estimates convincingly.

Table 3: Panel regression estimates of Δv on Δg.

	FE 1-way	FE 2-way	FE–AR 2-way	RE 1-way	RE 2-way
a	-0.012	-0.012	0.007	-0.011	-0.011
t_a	-3.78	-0.72	0.53	-3.44	-2.80
b	0.624	0.607	0.656	0.584	0.580
t_b	13.71	13.12	12.80	13.28	13.19
ρ			-0.289		

Note: ρ denotes the estimated autocorrelation in the remainder errors v_{it}.

We do not list further details on these specification tests. We will continue with the FE–AR two-way version, as it is the most general model. Individual effects correspond to the pattern that may be expected from the descriptive statistics in Table 1. The lowest effect is found for Botswana, where economic growth is impressive but productivity gains in agriculture are minor, and the largest effect is found for Benin. There is no visible trend pattern in time effects, indicating that the regression relation does not suffer systematic changes over time.

3.2 Granger causality

Viewed from the angle of policy recommendations, the identified panel relationship

$$\Delta v_{it} = 0.007 + 0.656 \Delta g_{it} + \text{error} \tag{1}$$

allows two possible interpretations. The first one would be that economic growth implies a less than one-for-one increase in agricultural productivity. A ten-percent real economic growth rate will only imply a six-percent growth in agricultural productivity, the remaining 4% maybe being absorbed by the increase in population or labor force.

However, we note that the relationship is static and not causal. Nothing prevents dividing by $\hat{\beta}$ and re-arranging so one obtains

$$\Delta g_{it} = 0.011 + 1.524 \Delta v_{it} + \text{error}. \tag{2}$$

This relationship would imply that any improvement in agricultural productivity entails a substantial effect on aggregate output. Keeping in mind that all effects are measured marginally in growth rates, one may conclude that a 10% rise in farming productivity would entail a 15% increase in real GDP. Policy implications would be totally different.

In order to obtain some insight on which of the two interpretations may be more plausible, we run Granger-causality tests on all individual countries. These are based on bivariate systems of Δg and Δv, with lag orders determined by AIC. Results are summarized in Table 4.

General impressions from Table 4 are that neither of the two causal directions dominates, that 14 out of 29 countries show no causal interaction for g and v at all, and that 8 countries indicate dynamic feedback. Only seven out of 29 cases show a clear causal ordering, three from v to g, and four from g to v. On the whole, the world is much more complicated than a simple causal interpretation of equations (1) and (2), and simplifying policy rules based on such relationships should possibly be avoided.

Figures 3 and 2 give an impression of the shock reaction in the vector autoregressions, as measured by non-orthogonalized impulse response functions that expand the inverse operator of the vector autoregressive polynomials. The intensity and even the direction of the reaction varies greatly across countries. However, it appears that reaction to g shocks in v tend to be slightly more pronounced than the reverse direction. This preponderance of the causal $g \rightarrow v$ reaction may conform to economic intuition.

Table 4: Granger causality among Δv and Δg for individual countries.

	$v \rightarrow g$	$g \rightarrow v$
Benin	-	-
Botswana	-	-
Burkina Faso	-	*
Burundi	***	-
Cameroon	-	**
Central Africa	-	*
Chad	***	**
Cote d'Ivoire	***	*
Congo	-	-
Congo (Zaire)	-	*
Gambia	-	-
Ghana	***	***
Guinea-Bissau	-	-
Kenya	**	-
Lesotho	*	-
Madagascar	***	*
Malawi	***	*
Mauritania	-	-
Niger	**	***
Nigeria	**	***
Rwanda	-	-
Senegal	**	**
Seychelles	-	-
South Africa	-	-
Sudan	-	-
Swaziland	-	-
Togo	-	-
Zambia	-	-
Zimbabwe	-	-

Note: One asterisk indicates significance at the 10% significance level, two asterisks at 5%, and three asterisks at 1% .

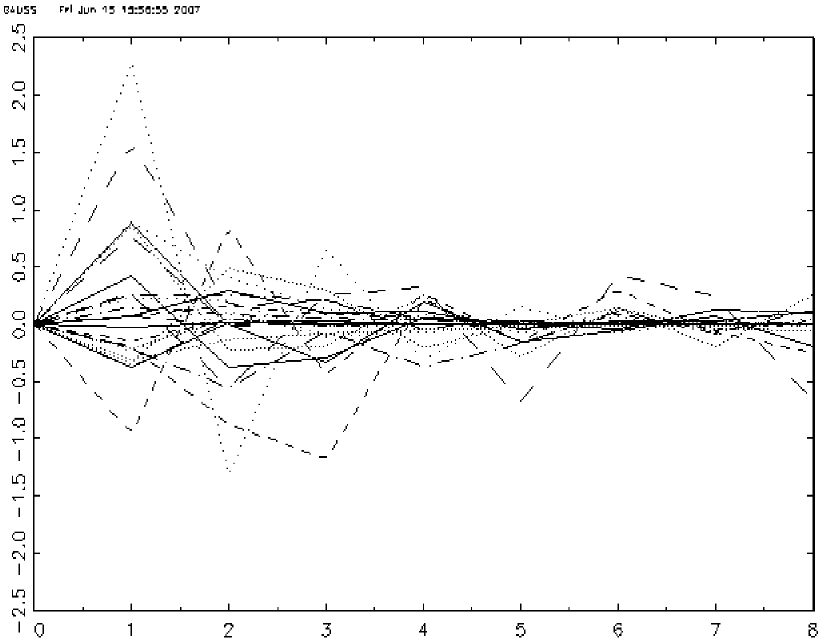

Figure 2: Non-orthogonalized impulse response functions for all 29 countries. Impulse variable is Δg and response variable is Δv.

Figure 3: Non-orthogonalized impulse response functions for all 29 countries. Impulse variable is Δv and response variable is Δg.

4 Prediction experiments

We are mainly interested in comparing time-series forecasts based on individual countries with forecasts based on the panel relationship represented by $\hat{a}$ and $\hat{b}$. Generally, we focus on predicting the last five observations, i.e. 2000 to 2004, using an out-of-sample design. Observations on v are not yet available for 2004 and, in some cases, are missing even for 2003. Therefore, forecast comparisons will be based on five observations for g and three to four observations for v.

4.1 Forecasts by individual vector autoregressions

In this section, we report two experiments. Firstly, vector autoregressions (VAR) are estimated for the sample ending in 1999 and starting in accordance with data availability, usually in 1975. Bivariate VAR models for $(\Delta g, \Delta v)$ are specified by a lag-order search using AIC. Then, forecasts are generated for 2000–2004 for individual countries. Summary measures, such as mean squared errors, necessarily mix up predictions of different step size, ideally 29 one-step and 29 two-step forecasts etc.

In a second prediction experiment, we update the VAR estimates after every step, such that we exclusively analyze one-step out-of-sample forecasts. Typically, we would expect that forecast errors are slightly lower in this second experiment.

Summary statistics are given in Table 5.

Table 5: Forecast accuracy measures for VAR predictions.

	g		v	
	RMSE	MAE	RMSE	MAE
VAR forecasts				
1 step	0.1032	0.0719	0.1333	0.0991
1–5 steps	0.1189	0.0814	0.1609	0.1122
AR+panel forecasts				
1 step	0.0694	0.0413	0.0697	0.0502
1–5 steps	0.1274	0.0789	0.1138	0.0967

4.2 Forecasts by exploiting the panel regression

The VAR forecasts serve as a convenient benchmark for checking the usefulness of the regression relationship that was presented in Section 3. They focus on dynamic interaction in individual country series but they do not account for the contemporaneous correlation among Δg and Δv. Given that a panel model

$$\begin{aligned}\Delta v_{it} &= a + b\Delta g_{it} + \mu_i + \lambda_t + u_{it}, \\ u_{it} &= \rho u_{i,t-1} + \varepsilon_t,\end{aligned} \tag{3}$$

holds, the conditional expectation of Δv_{it} is determined as

$$\mathrm{E}\left(\Delta v_{it}|I_{t-1}, \Delta g_{it}\right) = a + b\Delta g_{it} + \mu_i + \rho u_{i,t-1}. \tag{4}$$

The empirical implementation of this conditional expectation requires parameter estimates for a, b, μ_i, ρ and some approximation to the variables Δg_{it} and $u_{i,t-1}$. In fact, g_{it} appears to be available faster than v_{it}. For example, 2004 values for g_{it} are available at the time of writing, while 2004 values and even some 2003 values are missing for v_{it}. Hence, it may make sense to consider conditional prediction given observed g_{it}. However, particularly in the interest of compatibility with the results of the previous section, we focus here on predicting v_{it} given information dated $t-1$ exclusively. This implies that forecasts for g_{it} must be determined. We use the simple route of univariate autoregressions, with lag orders determined by AIC. In line with the results on Granger causality that we reported above, identified lag orders are typically low, zero or one in most cases. If zero was found, we nevertheless use a first-order autoregression for forecasting. Finally, the unobserved $u_{i,t-1}$ must be taken into account, for which we use the in-sample residual as an approximation.

It would be conceivable to use λ_{t-1} as an approximation for the unknown λ_t, and this may make sense if time effects express longer-run trend components across Africa. However, such trend behavior could not be found and all estimated time effects are small.

The lower part of Table 5 gives the results of this forecasting experiment. A short comparison with the upper part reveals that single-step forecasts have improved substantially. This discrepancy is probably not only due to the strength of the Okun-type panel regression but also to the poor performance of the bivariate vector autoregressions. The dynamic patterns across the two variables are infested with sizeable sampling variation, such that the univariate autoregression for g definitely outperforms the vector autoregression. The second step of the prediction procedure then accounts for the important simultaneous correlation of Δg and Δv that is ignored in the VAR models.

Prospects are less appealing for the multi-step experiment. Multistep forecasts based on pure univariate time-series models tend to return to estimated long-run means quickly, as the horizon increases. The quality of such non-informative predictions is close to the bivariate time-series model—with slight differences in ranking between the squared and the absolute loss criterion—that, in spite of its poor stability and parameter uncertainty, captures at least some of the longer-run business cycles. It is slightly surprising that the forecasts for agricultural productivity based on the relatively poor g forecasts and the panel coefficient estimates still outperform the time-series prediction, even at longer horizons. This again emphasizes the importance of the Okun-type relationship for forecasting.

References

[1] Baltagi, B. (2005) *Econometric Analysis of Panel Data*, 3rd Edition, Wiley.

[2] Clements, M., and Hendry, D.F. (1999) *Forecasting Non-Stationary Economic Time Series*. MIT Press.

[3] Johansen, S. (1995). *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*, Oxford University Press.

[4] Wagner, M., and Hlouskova, J. (2007) 'The Performance of Panel Cointegration Methods: Results from a Large Scale Simulation Study,' *IHS Economics Series* No. 210.