

Growth Is Good for the Poor: A Response to Lubker, Smith and Weeks

David Dollar and Aart Kraay
The World Bank
August 2002

Note: This response was prepared for submission to the Journal of International Development which published the critique of our paper. Despite repeated efforts we were unable to have the journal even acknowledge receipt of our response, let alone consider it for publication.

In a recent comment in this journal, Malte Lubker, Graham Smith and John Weeks (hereafter LSW) criticize our paper “Growth Is Good for the Poor” (Dollar and Kraay (2002)). In this response we show that their main criticisms are incorrect.

Our original paper is concerned with the relationship between growth in average incomes and growth in incomes of the poor, defined as those in the poorest quintile of the income distribution. We estimate a series of regressions of the logarithm of average incomes in the poorest quintile (y^P) on the logarithm of average incomes (y), and a set of control variables (X) in a panel of countries and years indexed by c and t , respectively:

$$(1) \quad y_{ct}^P = \alpha_0 + \alpha_1 \cdot y_{ct} + \alpha_2' X_{ct} + \mu_c + \varepsilon_{ct}$$

The error term consists of a country-specific time-invariant component, μ_c , and a term varying over both countries and years, ε_{ct} . Since average incomes in the bottom quintile are equal to the the first quintile share times average income divided by 0.2, Equation (1) is identical to a regression of the log of the first quintile share (Q1) on average income and a set of control variables:

$$(2) \quad \ln\left(\frac{Q1_{ct}}{0.2}\right) = \alpha_0 + (\alpha_1 - 1) \cdot y_{ct} + \alpha_2' X_{ct} + \mu_c + \varepsilon_{ct}$$

In our paper we argued that a very robust feature of the data in our sample of 92 countries over the past four decades is that the average elasticity of incomes of the poor with respect to average incomes, α_1 , is equal to one. From Equation (2), this is equivalent to the observation that, on average, the share of income accruing to the poorest quintile does not vary systematically with average incomes. In the absence of compelling evidence that either growth, or the variables that drive growth, are systematically associated with declines in the income share of the poorest, we conclude from our research that “growth is good for the poor”.

While LSW raise many issues in their paper with which we disagree, here we address three of their principal critiques in detail.

LSW argue that one should expect as a matter of arithmetic that the elasticity of incomes of the poor with respect to average incomes should be equal to one. LSW observe that in our dataset, averaging across all countries and years, incomes of the poorest quintile are roughly one-third of average incomes ($Y_{ct}^p = 0.315Y_{ct}$, LSW Equation (2e)). They then take logarithms of this average relationship to obtain $y_{ct}^p = -1.155 + y_{ct}$. Remarkably, they conclude from this arithmetic that econometric estimates of α_1 should by definition be equal to one.

This conclusion is incorrect. To see this, consider (as do LSW) the simple case where there are no additional control variables in the regression, i.e. $\alpha_2 = 0$. As is well known from basic econometrics, the ordinary least squares (OLS) estimate of α_1 is the *covariance* between log incomes of the poor and log average incomes, divided by the *variance* of log average incomes. Statistically, there is no reason a priori to expect that this ratio of covariance to variance should be equal to one. In fact, the only way to understand LSW's claim is to realize that in their simple arithmetic, they are *assuming* that the first quintile share is constant across countries and over time. It seems more useful to actually statistically investigate how the first quintile share varies across countries and over time, as we do in our paper.

LSW argue that our finding of an elasticity of one does not hold when the sample is divided into subgroups according to the level of income of the poor. In particular, LSW divide our data into three groups corresponding to low, medium, and high values of the left-hand-side variable in Equation (1). When they re-estimate Equation (1) by OLS in these subpopulations, they find estimates of α_1 that are statistically significantly lower than one in each of the three subgroups.

This claim is also rather remarkable, not least because it is clearly inconsistent with their previous (incorrect) assertion that α_1 must by definition be equal to one. More importantly, this claim is itself also incorrect. In fact, LSW have devoted a third of their paper to illustrating the elementary econometric mistake of stratifying the sample by the dependent variable. If the "true" model is given by Equation (1), it is well-known that OLS estimates in subsamples obtained by stratifying on the dependent variable will in general be biased towards zero. The formal proof of this is straightforward and is not reproduced here (see for example Chung and Goldberger (1984) for a reference). This is why in our paper we stratify only on the right-hand-side variables as part of our robustness tests.

A possibly more compelling illustration of this elementary point is given in Figure 1. We generate 1000 draws from a bivariate independent standard normal distribution. We label the first as "A", and construct the second as the sum of the two and label it "B". An OLS regression of B on A will by construction give a slope coefficient very close to 1, and this is exactly what we find in the top panel of Figure 1. The middle and bottom panels perform the same OLS regressions in the top and bottom half of the sample, ordered by the dependent variable, B. Not surprisingly, we find that in both subsamples the relationship between B and A is flatter than in the full sample, even though the "true" model generating the data features a slope of one. Rather than undermining our results,

LSW's empirical evidence in Table 2 of their paper is exactly what one would expect to find if the elasticity of incomes of the poor with respect to incomes were in fact equal to one.¹

LSW argue that our decomposition of the effects of policy on incomes of the poor into growth and distribution effects is logically inconsistent. In particular, LSW argue that since average incomes depend on policy, the only logical way to estimate the effects of policy on incomes of the poor is to directly regress incomes of the poor on policy variables alone (see their Equation (5)).

This third criticism is also both remarkable and incorrect. In our paper, we are interested in *separating* the effects of policy variables into their effects on average incomes (i.e. the “growth effect of policy”), and their effects on income distribution (i.e. the “distributional effect of policy”). This is why we need to estimate both Equation (1), which gives us the distributional effects of policy holding constant the level of income, and also a growth regression which gives us the effect of policies on average incomes. If we were to follow LSW's advice and estimate a reduced-form model of incomes of the poor as a function of policy alone, it would be impossible to statistically identify the “growth” and “distribution” effects separately, as should be obvious from their Equation (5). Finally, we leave it to readers more discerning than ourselves to reconcile the logical inconsistency between LSW's lengthy criticism our basic specification for being only a reduced form model (in Section 3 of their paper), and LSW's insistence that we should estimate a reduced-form model (in Section 5 of their paper).

In their concluding paragraph, LSW suggest that our work does not rise to professional standards of rigorous analysis. It is unfortunate that this critique is more properly applied to their comment than to our paper.

References

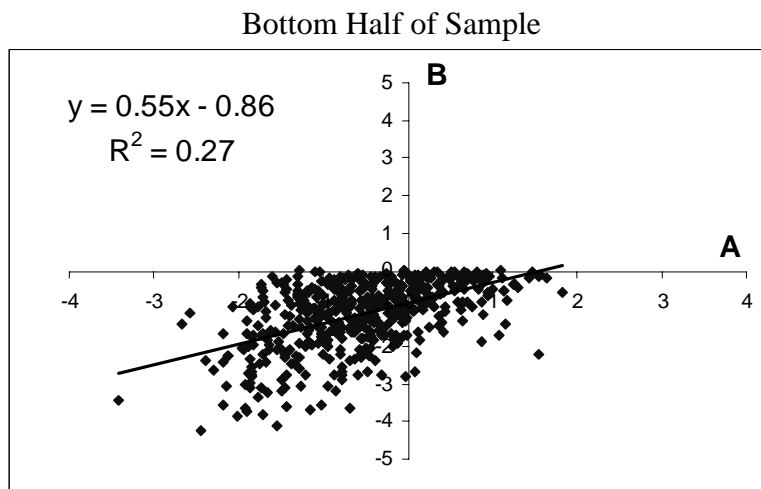


Dollar, David and Aart Kraay. “Growth Is Good for the Poor”. *Journal of Economic Growth*, forthcoming.

Lubker, Malte, Graham Smith, and John Weeks (2002). “Growth and the Poor: A Comment on Dollar and Kraay”. *Journal of International Development*. 14:555-571.

Chung, C.-F. and A. Goldberger (1984). “Proportional Projections in Limited Dependent Variable Models”. *Econometrica*. 52:531-534.

¹ The parallel between this simple example and our paper should also be clear. Since the logarithm of average incomes of the poor are the sum of the logarithm of the first quintile share and the logarithm of average incomes, and since we find in the data that the latter two variables are uncorrelated, we find a slope coefficient of one when we regress log average incomes of the poor on log average incomes.

Figure 1: Consequences of Stratifying on the Dependent Variable



Note: The variable “A” consists of 1000 i.i.d. draws from a standard normal distribution. The variable “B” is the sum of “A” and 1000 i.i.d. draws from another standard normal distribution independent of “A”. The top panel shows a regression of B on A, which by construction delivers a slope of 1. The middle and bottom panels illustrate the well-understood consequences of applying OLS to subsamples obtained by stratifying on the dependent variable, B.