

Trade, Growth, and Poverty: A Response to Nye, Reddy, and Watkins

**Aart Kraay
The World Bank¹
June 1, 2006**

In a recent paper (Dollar and Kraay 2004) we provided new empirical evidence in support of a positive and significant within-country relationship between changes in trade volumes and changes in per capita GDP growth rates. We also observed that there was little evidence in support of any significant relationship between changes in trade and changes in inequality. We concluded that, on average, expansions in trade raised growth as well as incomes of the poor. Nye, Reddy, and Watkins (2002) take issue with these findings.² Here we point out how their main critiques are either incorrect, or else based on a misleading mixture of misreading and misinterpretation of our findings.

Growth in Globalizers and Non-Globalizers

As a purely illustrative exercise we began our paper with a simple comparison of changes in growth rates in a group of "globalizing" developing countries economies with growth in a group of "non-globalizers". The former were selected based on their large proportionate increases in trade and/or declines in tariffs. We found that in most cases the globalizing countries also saw greater increases in growth. At the same time we recognized that it was difficult to draw conclusions from this simple set of comparisons, for all of the usual reasons. We then moved on to provide more systematic regression-based evidence of a significant partial effect of changes in trade on changes in growth. Finally we cited other work of ours that neither growth, nor changes in trade, tended to be significantly associated with changes in inequality.

NRW begin their critique by rehashing a number of points that we already acknowledged clearly in the paper. Changes in trade volumes are imperfect proxies for changes in trade policy. A correlation between changes in trade and changes in growth need not indicate a causal effect. There might be other factors correlated with changes in trade that also matter for changes in growth. We are glad that they have taken our caveats regarding our empirical work to heart. We are not sure however why it is necessary for them to repeat what we have already covered clearly in our paper, nor why this in any way constitutes a critique of what we have done.

¹ 1818 H Street NW, Washington, DC 20433, ddollar@worldbank.org, akraay@worldbank.org. The views expressed here are the authors and do not reflect those of the World Bank, its Executive Directors, or the countries they represent.

² Their critique was based on the working paper version of our paper, Dollar and Kraay (2001). The results in the published version (Dollar and Kraay (2004)) are sufficiently similar that we refer the ready to it instead. It should be noted that most of the points contained in this response were communicated privately to NRW in 2002, but to date they have not seen fit to revise their critique in order to remove the errors we have pointed out to them.

NRW next note that changes in growth between the 1980s and 1990s for the group of non-globalizers based on the tariff criterion are actually a bit higher than among the globalizers. They also argue that this comparison is more relevant than the one we made between the 1970s and 1990s (where the globalizers saw increase in growth while the non-globalizers saw massive declines in growth), since our tariff data refer to the mid-1980s and the mid-1990s. Factually correct. The question is whether it is particularly important. We think not, for two reasons.

First, in our Table 3, one can make six sets of comparisons between changes in growth between globalizers and non-globalizers: 1990s vs 1980s, and 1990s vs 1970s, for the three definitions of globalizers. In five of the six comparisons we find that the changes in growth among the globalizers are bigger than among the non-globalizers, both for simple, and population-weighted, averages. We think this regularity is at least as interesting as the single exception that NRW point out. And in any case, we repeat that this is intended as a purely illustrative exercise: our case for the growth effects of trade is developed more rigorously later with cross-country regressions.

Second, NRW's argument is premised on the idea that tariff rates in the mid-1980s are a poor proxy for tariff rates in the 1970s. Unfortunately we do not have systematic cross-country data that we could use to verify or disprove this claim, nor do NRW offer any evidence themselves. We can however get some suggestive evidence by looking at the more limited data we do have for the first half of the 1980s. Average tariff rates during the first half of the 1980s are available for only 47 developing countries, which is why we started our analysis in the second half of the 1980s where we had data for 73 countries. In any case, the correlation between average tariff rates in the first half of the 1980s and the second half of the 1980s is 0.81, which is quite high. Moreover, the average change in tariff rates between the first half of the 1980s and the second half of the 1990s for our group of tariff-based globalizers for which earlier data is available is a very substantial -23 percent (as opposed to -18 percent for the shorter period we studied). For the non-globalizers the tariff changes were very small over either period: the corresponding figures are -3 percent and -4 percent. Based on this we do not think it is unreasonable to suppose that our list of tariff-based globalizers would be very different if we had data for the 1970s. Unless NRW can provide systematic evidence to the contrary we do not find their critique to be at all compelling.

Levels and Changes

NRW continue their critique of our discussion of the globalizers with a very contradictory digression on levels versus changes. They correctly note that our groups of globalizers on average increased trade from very low bases, and lowered tariffs from quite high bases, and that even at the end of our sample the globalizers had lower trade ratios and higher tariff rates than the non-globalizers. They also correctly observe that these globalizers countries with lower trade ratios had higher growth in the 1990s than the non-globalizers. Based on this observation they leap bravely to the conclusion that our data supports the idea that countries that are more closed grow faster as a result.

This leap manages to miss the entire point of our paper -- that it is more informative to look at the within-country relationship between *changes* in trade and *changes* in growth precisely because this within-country variation is less likely to be tainted by various econometric problems than the cross-country variation. As we discuss the cross-country variation in trade is heavily influenced by geography which may also matter for growth. There are many candidate unobserved country-specific variables that do not change much over time that might drive both growth and trade, and differencing removes these variables. NRW surely understand these points, as they themselves repeat them in their paper. Yet they choose to selectively ignore them in order to try to score rhetorical points. We do not think this is very useful.

Cross-Country Regressions

NRW next turn to a critique of our more systematic regression-based evidence on trade and growth. Here their critiques are disappointingly generic and unsubstantiated. They correctly note that the partial correlation between trade and growth that we document could be driven by an omitted variable correlated with both trade and growth. As a general point this is fine, although not particularly insightful as it applies to every single empirical paper ever written. A critique along these lines only has content if NRW can produce such a specific omitted variable, and show that its inclusion in the regression overturns our results. Otherwise it is little more than idle speculation and does not deserve further response.³

Similarly, they note the completely generic possibility of reverse causation: faster growth could lead to higher trade. This is also hardly news: in fact it is a concern that we discuss at length in the paper. What they fail to provide is a coherent critique of our identification strategy that we think helps to address this concern. In particular, we use a fairly mild identifying assumption that shocks to growth in the 1990s do not affect trade in the 1970s. As with all identifying assumptions one can debate their value. But such debates have little content unless NRW can provide empirical evidence of a specific omitted variable that is correlated with trade in the 1970s and growth 20 years later that might invalidate our identifying assumption.⁴ Otherwise we are again left with little more than idle speculation.

Growth in Incomes of the Poor

NRW conclude by criticizing our empirical findings that on average changes in trade volumes do not appear to be correlated with changes in inequality, suggesting that

³ They do suggest that changes in institutions might be the relevant omitted variable. What they fail to note is that we have in this and its companion paper (Dollar and Kraay (2002a)) controlled for a variety of time-varying measures of institutional quality and found that the inclusion of this variable did not overturn our findings on trade.

⁴ Peculiarly in their footnote 19 they deliberately ignore results reported in the paper (and in more detail in the companion paper (Dollar and Kraay 2002a)) which show that trade in the 1970s is strongly correlated with subsequent changes in decadal trade ratios, thus satisfying the condition of instrumental relevance. They instead prefer to speculate vacuously that "there is no obvious economic rationale" for this correlation, while simultaneously ignoring its existence.

expansions in trade are on average distribution-neutral. They begin by propagating a fallacy common in uninformed discussions of econometric evidence: that a result can only be valid if it is accompanied by an R-squared equal to one. They take great exception to our passing remark that, based on our regressions, one should not expect a country like Burma to have a significant change in inequality if trade expanded. Their critique consists of observing that no country is right on the regression line in this relationship, and so this remark cannot be true. This is of course nonsense from an econometric viewpoint. An unbiased forecast of the marginal effect of trade on inequality is going to be just the estimated regression coefficient, which is implicitly what we were referring to. This does not mean that the effect will in fact be zero, it is just that zero is an unbiased forecast of the effect.

Moreover, they insist on quoting only very selectively from our remark on Burma in order to create the illusion of having ground to stand upon. In fact, what we actually wrote was the following quite qualified statement: "...what can we expect to happen when developing countries liberalise trade and participate more in the global trading system? *Obviously for a particular closed economy (say, Burma) we cannot predict with certainty what will happen. The specific outcome will depend on a whole host of factors (including the country's factor endowments, its location, and complementary policies put in place).* But we can make some qualitative predictions.....there is no reason to expect any large change in household income inequality." (Dollar and Kraay (2004), pp. 29-30, italics added for emphasis). It is hard to see how any fair reading of this could be interpreted as suggesting that we think that the effect of trade on inequality will in fact be zero in any given country, as NRW would have us believe.

NRW continue with another highly tangential argument. Rather than take at face value our empirical evidence that income of the poor tend to increase equiproportionately with average incomes, they argue that we somehow instead make the claim that this implies that absolute changes in incomes of the poor are equal to those of the rich. As a matter of simple arithmetic this is of course ludicrous: equiproportionate changes for the poor and the rich mean smaller absolute increases for the poor than the rich. That our results refer only to proportionate changes should also be abundantly obvious. From our abstract, where we write that "the increase in growth rates leads on average to proportionate increases in incomes of the poor", to our introduction where we summarize our results as "a one-to-one relationship between the growth rate of income of the poor and the growth rate of per capita income", to Section 2 where we write that "...growth in mean income is translated one-for-one into growth in income of the poor", to our conclusion where we write that "...the increase in growth rates ... therefore on average translates to proportionate increases in incomes of the poor" we are abundantly clear about what our results are. Our paper Dollar and Kraay (2002b) on which these findings are based is similarly littered with clear statements regarding proportionate changes.

It is therefore a mystery to us why NRW could possibly suggest that we fail to understand the implications of our own empirical work. If they think that a comparison of absolute changes in income between the rich and poor is of interest, this is fine and they are certainly entitled to explore this in their work. But we fail to see why such a line

of discussion has anything to do with our paper, much less that it constitutes some kind of critique.

They conclude by arguing that focusing on the effects of growth on incomes of the poorest quintile has nothing to do with poverty. At one level this is simply a question of terminology. There are enormous debates in welfare economics, and in other disciplines, as to what constitutes an appropriate social welfare function for evaluating policy. Poverty is often referred to as the fraction of people below a fixed poverty line, although Pritchett (2006) provides an interesting critique of why this can be inappropriate.

In any case, a more interesting question is whether our conclusions would be very different if we used a traditional poverty measure like the headcount. Our results suggest that growth is on average distribution-neutral, as increases in mean income tend not to be systematically correlated with both the Gini coefficient as well as the first quintile share. It follows immediately from this that growth will on average reduce headcount poverty, and this is a regularity that has been abundantly documented. In fact, a recent contribution Kraay (2006) has shown that in the long-run in a large sample of developing countries, fully 97 percent of the cross-country variation in changes in the headcount measure of poverty can be explained by growth, and that changes in inequality account for only 3 percent. NRW are of course free to define poverty however they want. But it seems unlikely to us that any sensible measure of poverty will not be consistent with the spirit of our results.⁵

In summary, we find NRW's critique of our paper baseless. It consists of highly selective quotation and interpretation of our paper, combined with entirely unsubstantiated speculation. We do not claim that our work is the last word on the growth or distributional effects of trade. To the contrary it is a small contribution in a very large literature that continues to be advanced by serious research, with some results supporting our work and others contradicting it. NRW's energies might be better spent developing their own serious contributions to this important area of research.

References

⁵ NRW do rather triumphantly cite an unpublished paper (Foster and Szekely (2001)) which looks at the relationship between the growth rate of average income and "generalized mean" income, where the latter is a geometric average of incomes. For negative values of the exponent in the geometric average, this measure puts greater weight on the lower incomes. Foster and Szekely (2001) document that the slope coefficient in a regression of growth in the generalized mean on growth in the simple average is lower the greater is the degree of bottom-sensitivity of the generalized mean. What NRW fail to note is that in all cases Foster and Szekely (2001) cannot reject the null hypothesis that the coefficient on average income growth is equal to one at conventional significance levels. For example, in their Table 1 they obtain a point estimate of 0.33 with a (huge) standard error of 1.5 for their most bottom-sensitive specification, and so cannot reject the null that this slope is equal to one. This approach is interesting in its own right and it is also perfectly consistent with our empirical results.

Dollar, David and Aart Kraay (2001). "Trade, Growth, and Poverty". World Bank Policy Research Department Working Paper No. 2615.

Dollar, David and Aart Kraay (2002a). "Growth is Good for the Poor". Journal of Economic Growth. 7(3):195-225.

Dollar, David and Aart Kraay (2002b). "Institutions, Trade, and Growth". Journal of Monetary Economics. 50:133-162.

Dollar, David and Aart Kraay (2004). "Trade, Growth, and Poverty". Economic Journal. 114(493):F22-F49.

Foster, James and Miguel Szekely (2001). "Is Economic Growth Good for the Poor? Tracking Low Incomes Using Generalized Means". Interamerican Development Bank Working Paper No. 453.

Kraay, Aart (2006). "When is Growth Pro-Poor? Cross-Country Evidence". Journal of Development Economics. 80(1):198-227.

Pritchett, Lant (2006). "Who Is Not Poor? Dreaming of a World Truly Free of Poverty". World Bank Research Observer. 21:1-23.