



Munich Personal RePEc Archive

Evaluating Aid Effectiveness in the Aggregate: Methodological Issues

Carl-Johan, Dalgaard and Henrik, Hansen

Danida's Evaluation Department

March 2009

Online at <https://mpra.ub.uni-muenchen.de/23025/>
MPRA Paper No. 23025, posted 04 Jun 2010 10:20 UTC

EVALUATING AID EFFECTIVENESS IN THE AGGREGATE: METHODOLOGICAL ISSUES



Evaluating Aid Effectiveness in the Aggregate: Methodological Issues

March 2009

Carl-Johan Dalgaard, Department of Economics, University of Copenhagen
Henrik Hansen, Institute of Food and Resource Economics, University of Copenhagen

*Acknowledgement: We are grateful to Arvind Subramanian for sharing the data from Rajan and Subramanian (2008).

Disclaimer: The views expressed are those of the authors and do not necessarily represent the views of the Ministry of Foreign Affairs of Denmark. Errors and omissions are the responsibility of the authors.

Contents

Executive Summary.....	3
1. Introduction	4
2. What is “Aid Effectiveness” at the aggregate level?.....	6
2.1. Should the focus rather be on development?	7
2.2. Should the focus rather be on poverty alleviation?	11
3. The basic empirical approach to assessing aid effectiveness: Regression analysis	13
4. The cause and effect problem.....	17
4.1. Simple and multiple regression analysis	17
4.2. Two rules colliding: The identification problem	19
4.3. Is it really a problem?	25
4.4. A possible solution to the cause and effect problem.....	28
5. An illustration of the regression approach.....	32
6. Yet another specification problem: The impact may vary	42
6.1. The Effect depends on the characteristics of the recipients	46
6.2. The Effect depends on characteristics of the donors.....	46
6.3 The regression solution: more instruments.....	47
7. Concluding remarks	52
Literature.....	53
Annexes.....	56
Annex 1. The Cause and Effect Problem.....	56
Annex 2. Reverse causality: the impact of growth on aid.....	58
Annex 3. Instrumental variables estimation.....	59
Annex 4. The Data	61

Executive Summary

The purpose of the present Evaluation Study is to discuss the methodological problems researchers are facing in gauging the impact of aid on economic growth. The discussion is non-technical and aimed at an audience without much prior knowledge in the fields of macroeconomics and econometrics.

The paper provides insights into the following questions:

1. Why do economists view “aid effectiveness” as synonymous to asking whether aid increases growth in income per capita?
2. Why is it difficult to determine the macroeconomic impact of foreign aid on economic growth?
3. How is it, in principle, possible to solve the difficulties present in evaluating aggregate aid effectiveness?

A companion study surveys recent research on the topic, with reference to the methodological problems laid out in the present paper.

Key points:

- The objective of macroeconomic research on “aid effectiveness” is to gauge the impact of foreign aid on growth in GDP per capita. This choice of focus is appropriate for theoretical as well as practical reasons.
- Statistical modelling, mainly based on regression analysis, is the key methodological approach.
- Basic regression analysis *cannot* answer the question if foreign aid is effective in the sense that it increases the growth of GDP per capita.
- To elicit information about the *impact* of aid, application of more advanced regression techniques is required.
- Application of the more advanced regression techniques requires quantitative information which is in practise very difficult to obtain.
- There is considerable uncertainty as to whether it is reasonable to assume that the impact from aid on growth *is the same* in every country. Unless the researcher gets it right, the results from the analysis of aid effectiveness are likely misleading.

1. Introduction

Evaluation and *impact* are words used more frequently than *development* and *poverty* when major donors meet and discuss foreign aid. Although this may seem cynical to many who care about the poor people of the world, it is natural to ask if giving aid does any good, and this is what the evaluation of the impact of foreign aid is all about. The impact of aid has been discussed, and disputed, since the start of the major aid programmes in the late 1950s and early 1960s. The discussion is still ongoing and, today, the debate appears at many levels from highly technical analyses in academic journals over more popular arguments in bestselling books to brief articles and editorials in newspapers and even short sharp shocks on web-pages and blogs.

Often, the popular views on aid are polarized and stated as one-liners. The critics of aid will contend that ‘aid does not work—it is wasted’ while the supporters assert that ‘aid works—it should be doubled’. In popular writings, such as the bestselling books *The End of Poverty* by Jeffrey Sachs (2005) and *The White Man’s Burden* by William Easterly (2006) there are attempts at giving more nuanced pictures and documentation supporting the statements but when it comes to ‘hard evidence’ of the economy-wide impact of foreign aid the documentation is somewhat blurred.¹ The more technical discussions in the academic journals are primarily based on statistical analyses. Surprisingly to many, even when researchers look at the same data they can come up with quite different answers to the same basic question: does foreign aid flows increase economic growth?²

Discussions and disagreements are common in most fields of economics, in particular within development economics. So in this respect the aid effectiveness debate is not special. In fact, within the academic circles in economics, all aspects of economic growth are debated. Two other, well-known, areas of heated contention are the pros and cons of trade liberalization and of financial liberalization. Popular discussions about trade liberalization can be found in, for example, Bhagwati (2004) and Stiglitz (2006) while Mishkin (2006) and Stiglitz (2003, 2006) provide illustrations of the debate about financial liberalization.

Understanding how economists analyze data is important if one wants to come to grips with the aid effectiveness discussion. However, the statistical models used in the analyses are unfamiliar to most aid practitioners making them unable to judge if a particular study of aid effectiveness tackles the statistical problems in an appropriate way. The main purpose of this evaluation study is, therefore, to introduce the reader to the statistical problems encountered by researchers in their analyses of aid effectiveness at the aggregate level. A companion evaluation

¹ The fuzziness is on both accounts: Jeffrey Sachs is pro while William Easterly is con in the ‘aid works’ discussion.

² See, for example Burnside and Dollar (2000) and Dalggaard and Hansen (2001) who analyze exactly the same data.

study (Dalgaard and Hansen, 2009³) discusses and evaluates recent studies of aid effectiveness at the aggregate level using the present study to form a methodological benchmark for comparisons.

The study is organized as follows. In section 2 we explain why economists view “aid effectiveness” as synonymous to asking whether aid increases growth in income per capita. In section 3 we briefly introduce the idea of looking at data using regression analysis while section 4 focuses on some specific problems that leads researchers to get ‘wrong answers’ when they use the simple regression method known as ordinary least squares. In section 5 we introduce a more advanced regression method, called two-stage least squares, which is useful when researchers wish to find the causal impact of aid on economic growth rather than the mere correlation between the two, which is the result one gets when applying the ordinary least squares method. The importance of the choice of method is illustrated, using a real life data set, in section 6. In section 7 we briefly discuss some additional problems that arise when the effectiveness of aid depends, systematically, on either recipient or donor country characteristics. These added complexities are also illustrated using the same data as in section 5. Finally, section 8 offers some concluding remarks. For the interested reader, we have gathered some short mathematical presentations in three annexes. The material in the annexes is not, in any way, essential for understanding the main issues.

³ The companion study will be published as an Evaluation Study by the Evaluation Department of Danida in 2009.

2. What is “Aid Effectiveness” at the aggregate level?

Existing cross-country differences in GDP per capita (average income) almost defy comprehension. In 2000 the average income in Burundi was roughly 100 US\$. Meanwhile, the average American citizen’s income was roughly 35,000 US\$. This comes out to a per capita income difference of a factor of 350! Admittedly, this number overestimates the difference in purchasing power that the levels of income imply since 100\$ will buy many more goods and services in Burundi than what it would be feasible to obtain if the sum was spend in the States. Hence, the above common currency comparison of GDP per capita overestimates the true difference in *living standards*. At the end of the day the relevant metric for cross-country inequality is not how much income differs as such. Rather it is how much *consumption possibilities* differ.

Hence, to perform a more accurate comparison, suppose we were to ask how many hours it would take the two representative citizen’s to earn money enough to buy identical goods in their respective countries; say, 2000 calories worth of sweet potatoes.⁴ For simplicity, suppose the two citizen’s both work 24 hours per day, 365 days a year, to earn an annual income of 100\$ and 35,000\$, respectively. Factoring in the calorie contents of a gram of sweet potatoes (roughly 1), and local (producer) prices of sweet potatoes in 2000 (roughly 147 US\$/tonne in Burundi, and 337US\$/ tonne in the US), we find that it would take the average person in Burundi about 29 hours to work up the required income. By contrast, the average US citizen would only have to work for 0.2 hours, or a mere 11 minutes. This difference in “time to earn” is equivalent to a difference in income per capita, measured in terms of purchasing power over calories from a particular food stable, of $29/0.2 = 151$. Hence this simple *purchasing power parity adjustment* of income has reduced the GDP per capita difference in a major way, from a factor of 350 to about 150. Nevertheless, even after this adjustment the difference in living standards is truly remarkable.

In light of this simple fact, it is no wonder that economists are keen on discovering ways of elevating GDP per capita in the poorest places around the world. In theory, a means to this end could be foreign aid. Indeed, in economics, the question of whether aid is “effective” is usually viewed as synonymous to asking whether foreign aid increases growth in GDP *per capita*. Specifically, the object of interest is always GDP per capita, *adjusted* for purchasing power differences.⁵ Hence, aid is viewed as “effective” if it increases average living conditions over

⁴ Why sweet potatoes? Because sweet potatoes make up roughly 20% of the diet in Burundi. Hence, this is an item which is actually quite important for subsistence in this country. In any case, this is just an example.

⁵ This doesn’t mean economists are examining GDP per capita in terms of sweet potatoes, as in the example. Rather so-called PPP GDP per capita has been constructed which involves deflating income by price indices involving *many*

time. That is, if aid increases the growth rate of purchasing power adjusted GDP per capita. This choice of focus is sometimes criticized for being misplaced, or at least much too narrow.

One line of criticism is that the appropriate measure of effectiveness is whether aid fosters *development*, rather than the mere expansion of material wealth. There is merit to this complaint. After all an often used synonym for “foreign aid” is “development aid”. This would suggest that policymakers (at least) tend to have broader objectives in mind when they decide to disburse aid.

Another line of criticism starts from the observation that economic growth in *GDP per capita* measures the expansion of *average* income. Ultimately, the argument goes, it is more important to study *poverty*. That is, whether aid is able to decrease the number (or fraction) of people in a population that are living below some minimum income threshold. The distinction is real in the sense that a country may grow in terms of average income without much improvement in the living conditions of the poorest. If the personal income distribution is getting more unequal during the growth process, this could be the result. Accordingly, income per capita is arguably not fully satisfactory as a “measure of success” since it does not take the country specific distribution of resources into account.

Below we discuss these two views before we, in the sections to follow, lay out the methodological issues involved in examining the impact of foreign aid on the evolution of average living standards.

2.1. Should the focus rather be on development?

Development practitioners and academics from branches such as agronomy, geography, sociology and anthropology often accuse development economists of being awfully narrow minded in their preoccupation with national income measures such as the gross domestic product (GDP), and the growth of the national income. In particular, when we are dealing with development aid the relevant target surely needs to be development, which is a much broader concept than (growth of) national income.

Few economists would disagree with this view and over the past 60 years development economists, jointly with philosophers, have formulated theories and concepts concerned with development and the quality of life. Some examples are the basic goods approach of John Finnis, the basic needs approach of Paul Streeten and Des Gasper, the prudential values theories of James Griffin and, of course, the capability approach of Amartya Sen.⁶ All these

comparable goods at the same time, and taking the composition of consumption into account. Still, the principle of the adjustment is along the lines of the example.

⁶ Qizilbash (2002) discusses the differences and common ground among the four approaches.

approaches are concerned with the quality of human lives and they recognize that it has many dimensions.

In relation to evaluations of aid effectiveness one problem with these broad and inclusive theories of development lies in questions of how to measure the various dimensions of the quality of life. Finnis' basic goods theory includes 'life', 'play', 'knowledge' and 'sociability' while Griffin's prudential values include things such as 'accomplishment', 'understanding', 'enjoyment' and 'deep personal relationships'.

Sen generally avoids specifying a list of capabilities. Nevertheless, Sen, and the capabilities approach, has had a profound influence on the construction of the Human Development Index (HDI), which has been an integral part of the Human Development Reports since their inception in 1990. The first report ("Concept and Measurement of Human Development") specified three aspects of the quality of life to be enhanced by development: longevity, knowledge and 'command over resources to enjoy a decent standard of living'. (Human Development Report, 1990). In practise, longevity is measured by life expectancy at birth; knowledge by the literacy rate, while purchasing power adjusted real GDP per capita is used as a stand-in for 'command over resources'.

Anand and Sen (2000) note that the use of 'command over resources', and the income measure used as its stand-in (proxy), is meant to capture other basic capabilities not already included in the measures of longevity and education. However, they also stress the importance of including a measure of income, per se, in the HDI:

"Having an income is not, of course, comparable with being educated or living long, which are valued for their own sake. Having an income-related control over purchasable commodities can scarcely be intrinsically valuable. Nevertheless, in an indirect way – both as a proxy and as a causal antecedent – the income of a person can tell us a good deal about her ability to do things that she has reason to value. As a crucial means to a number of important ends, income has, thus, much significance even in the accounting of human development. While something is lost in terms of 'purity' in not sticking only to variables such as life expectancy and being educated which are valuable in themselves, a major practical gain is made in indirectly extending the coverage to take note of various capabilities that people do value intensely and which cannot be adequately reflected in figures of life expectancy and literacy." (Anand and Sen, 2000, p. 100)

Hence, one can surely argue that even though growth of national income is not a synonym for development, it is an indicator of an essential part of the quality of human life.

In addition to its independent status as an important indicator of development, national income per capita also has a close association with other indicators of human wellbeing. This is illustrated in Figure 1 which depicts the association between the components of the HDI using

data from the Human Development Report 2007/2008. The Index has three components in total. The first, for longevity, is life expectancy at birth while the second, for knowledge, is a combination of the adult literacy rate and the gross enrolment rate (share of children at each level of schooling actually attending school). The third component, for “command over resources”, is GDP per capita adjusted for differences in purchasing power, here, converted into daily income to ease the understanding of the enormous differences.




FIGURE 1. The association between purchasing power adjusted GDP per capita and other components of the Human Development Index.

Data Source: Human Development Report 2007/2008, Human Development Indicators, Table 1.

Figure 1 highlights that the individual components of the Human Development Index are mutually highly correlated. Hence, as is well known, there is a strong tendency for people in richer countries to live longer and be better educated.




FIGURE 2. The association between purchasing power adjusted GDP per capita and Gender Empowerment (GEM). The GEM index measures women’s economic participation and decision making, political participation and power over economic resources. Larger values mean greater empowerment. See the data source for explanation of the construction of the index.

Data Source: Human Development Report 2007/2008, Human Development Indicators, Table 1 and Table 29.

In a broader perspective, per capita GDP is correlated with essentially any indicator of the various dimensions of development that has been put forward. As another example, Figure 2 illustrates the strong association between per capita GDP and the ‘Gender Empowerment Measure’ calculated in HDR 2007/2008 revealing another stylized fact; gender equality is generally increasing with rising national income.

The strong correlation between per capita GDP and other development indicators is often used as an argument in favour of looking (only) at GDP and its growth rate. Interestingly, Anand and Sen (2000) turns this argument on its head by asking why one should not simply look at life expectancy or literacy instead of GDP. After all, life expectancy and literacy are direct measures of human wellbeing whereas GDP per capita is only an indirect measure. As the three variables are highly correlated, looking at GDP per capita may not add much information. This is a reasonable argument. The problem is however that while economists have a well-established tool box for analysing the growth process, the same cannot be said for other aspects of human

wellbeing. As Robert Lucas Jr. noted awhile ago, after reviewing the basic theoretical framework that economists' often use to study the development process:

“It seems universally agreed that the model I have just reviewed is not a theory of economic development. Indeed, I suppose this is why we think of “growth” and “development” as distinct fields, with growth theory defined as those aspects of economic growth we have some understanding of, and development defined as those we don’t.” (Lucas, 1988, p. 13).

Hence, when economists are faced with a choice between analysing economic growth and, say, life expectancy, they almost inevitably opt for analysing growth because the profession has developed a rich framework for this kind of analysis. In addition, as documented above, rising income levels do seem to be narrowly connected to more direct measures of development.⁷

2.2. Should the focus rather be on poverty alleviation?

Turning to the question if we should focus on poverty instead of average income (GDP per capita) in aid effectiveness analyses, it is obvious that the focus of many foreign aid initiatives is that of poverty alleviation. Further, the first millennium development goal is to “*halve, between 1990 and 2015, the proportion of people whose income is less than \$1 a day*” (www.un.org/millennium-goals/). Hence, it may seem more relevant to focus directly on measures of poverty rather than on growth in average income. Indeed, from this (policy) perspective the critique of studies that focus on growth in GDP per capita has merit.

At the same time, one may observe that a strong focus on within country income inequality, in the context of poor countries, represents an example of an inability to “see the forest for the trees”. Consider Figure 3, which shows purchasing power adjusted GDP per capita per day for 24 of the poorest countries in the world. As is apparent, most of these countries are located in Sub-Saharan Africa. Moreover, measured on a daily basis it is clear that many of the poorest countries are hovering around the two dollar a day threshold.

⁷ In some dimensions, however, there is debate as to whether affluence *brings* development. An example is democracy. The correlation between income and democratic right is strong and positive. But while there is a long tradition, going back to Lipset (1959) of believing economic prosperity also brings political reforms, this remains an area of controversy. In the cases discussed above (e.g., gender empowerment) there are well developed theories to suggest growth lead to “development”. It should be noted, however, that causality may equally well work in the opposite direction. To anticipate a theme developed below in the context of aid effectiveness research: correlations tell us nothing about cause and effect between the two variables in question. Regardless, in the present case the main point is that GDP per capita is the best single “marker” of development available. Moreover, as emphasized above, the framework for analysing the growth process is well developed and largely commonly accepted within the economics profession.




FIGURE 3. Purchasing power adjusted GDP per capita per day in 24 of the poorest countries. Data Source: Human Development Report 2007/2008.

One way of thinking about these numbers is as the daily living standards of the peoples of, say, Sierra Leone *in the absence of any income inequality within the country*. Hence, if we were to (as a mental experiment) even out all difference in living standards within Sierra Leone, every person in the country would end up around the 2 dollar per day subsistence boundary. The poor living conditions in Sierra Leone are therefore *not* simply a consequence of an unequal distribution of income. If poverty is to be reduced in Sierra Leone there is only one way in which this is feasible: by fostering growth in average income. A very similar point can be made in the context of the other countries in the figure.

This is not to say that growth in average income inevitably will improve the living standards of the poorest people within the poorest countries. But what we have to face up to is the simple realization that economic growth is a necessary condition for lasting reductions in poverty, whichever way we choose to measure the latter. It is not possible to eliminate poverty in the poorest places around the world unless growth in GDP per capita is (re-)vitalized. Consequently, it is natural to examine whether indeed foreign aid has been able to do so.

In what follows we, therefore, focus on the relationship between aid flows and growth of GDP per capita.

3. The basic empirical approach to assessing aid effectiveness: Regression analysis

The most basic way of analyzing the association between two variables of interest is by plotting them against each other in a so-called cross-plot. Hence, as a point of departure, Figure 4 plots data on aid and growth.⁸ More specifically, Figure 4 depicts the average, percentage, ratio of aid to GDP from 1970 to 2000 and the average annual growth rate of GDP per capita during the same period, for 78 countries. The average annual growth rate in GDP per capita for each country is calculated as $100 * [(y_{2000} / y_{1970})^{1/30} - 1]$ where y is GDP per capita and the subscript indicate the year. The aid to GDP ratio is the average of the annual ratios from 1971 to 2000.

A reasonable question is why one would focus on 30 year averages, rather than averages over shorter periods of time. The answer is that the long-run average tends to “iron out” short run fluctuations in growth and aid; in the short run (say at a yearly frequency) the data on aid and growth can be quite far from the long-run trend because random events, such as weather conditions for agricultural production, are influential in a given year. The impact of variation in rainfall and other short-run fluctuations will be smoothed out when we use averages over several years. Nevertheless, it is worth pointing out that there is no objective criterion that inevitably recommends taking averages over three or four decades; shorter averages (down to say 4 or 5 years) may be sufficient to expose long-run patterns. Still, for present purposes the 30 year average will do; cross plots of the average aid-to-GDP ratio and average annual growth rate in GDP per capita always tend to look like Figure 4 regardless of the choice of base period and the length of the average involved.

Since the early 1970s, plots like Figure 4 have appeared in numerous scholarly books, journal articles and government reports analyzing aid effectiveness. There are several things one may take away from the figure. First, one may note that there is a lot of variation in terms of how much aid various countries received during the 30 year period. At one end of the spectrum we find a country like Guinea-Bissau (GNB) where foreign aid accounted for nearly 30 percent of GDP, on average. Meanwhile, in Nigeria (NGA)—to name another African country—aid accounted for less than half a percent of GDP on average. For most countries aid constitutes a fairly low fraction of their GDP. The median level of aid is just below 3 percent.⁹ To put the latter number into perspective one may observe that the contribution to GDP from agriculture in Denmark accounted for about 3 percent in 2000. Hence, for the “typical” aid receiving nation aid is just about as important, in accounting for GDP, as the production of agricultural foods is, in a rich place like Denmark. As some of the 78 countries received quite high aid

⁸ The data is tabulated in Annex 4.

⁹ This means that half of the 78 countries received less than 3 percent in aid (as a fraction of GDP) while the other half received more than 3 percent.

inflows during the period, the mean aid-to-GDP ratio is somewhat higher than the median, at 5.5 percent.

Second, from the figure we also learn that most of the 78 countries became richer during the 30 year period; the median growth rate in the sample is just below 1.4 percent. Still, this rate does fall short of the average growth rate for most rich countries like Denmark where the average growth rate was around 2 percent during the same period. Hence, the *relative* income difference between the typical aid recipient in the figure, and the richest places on earth, tended to grow from 1970 to 2000. Moving away from the median we may observe that the *absolute* level of income per capita actually fell from 1970 to 2000 in 17 of the countries in the sample.

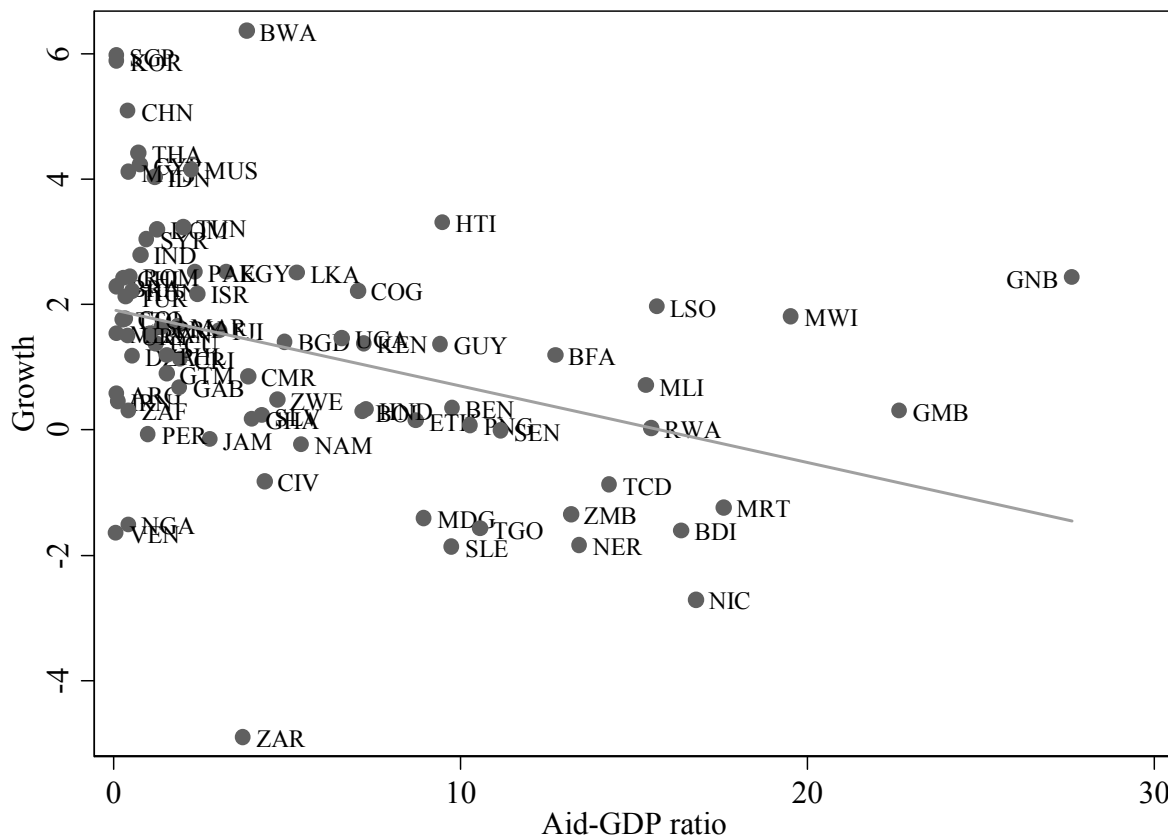


FIGURE 4. The association between the aid-to-GDP ratio 1970-2000 average, and average growth in GDP per capita (PPP) 1970-2000. See the text for explanations of the calculations. *Notes:* The line in the figure is a simple regression line estimated by ordinary least squares (OLS). Individual countries are identified by a three-letter ISO code which is unique. See Annex 4 for country codes and country names.

Data Source: Rajan and Subramanian (2008).

What most people take away from plots like Figure 4 is of course the negative association between the aid-to-GDP ratio and the growth rate in GDP per capita. That is, countries who

receive more foreign aid are on average growing more slowly. However intuitive, this “visual inspection approach” has its shortcomings and sometimes the approach does not answer the questions we have in mind. For starters, it is hard to tell if the association, in a meaningful statistical sense, is systematically negative or not.

Regression analysis is a statistical technique that allows a resolution of the latter problem in the sense that it allows us to look at the “strength” of the apparent negative association. For instance, employing regression analysis we can determine how to ‘best’ draw a straight line through the observations in Figure 4. This exercise amounts to specifying a linear relationship between aid and growth. The linear relationship is expressed mathematically as

$$Growth = a + b \cdot Aid,$$

where “ b ” is the slope of the line and “ a ” is the growth rate when no aid is given to a country. Subsequently, one may ask whether the slope of the line is positive, negative or zero, and how much confidence we should have in the association being systematic.

The regression describing the linear relationship between aid and economic growth is depicted in Figure 4. Statistically speaking the linear association is significant, which means it can be thought of as a reasonably strong association and not just a random coincidence. The statistical confidence we have in this result is high. In fact, the association is so strong that with 99% probability we reject the hypothesis that the two variables are unrelated. Hence, there is a statistically strong negative association between the development in living standards and aid disbursements, measured as a fraction of the recipient countries’ GDP.

The *economic* strength (as opposed to the *statistical* strength) of the association can also be gauged invoking the regression analysis, since we determine the slope of the line. In the present case, the slope, b , is -0.12. It is important to understand what this means.

Suppose we are observing two countries, A and B, and that the only knowledge we have about the two countries is that A receives 1 percentage point more aid than B along with the slope estimate, $b = -0.12$. Suppose next that we are asked what the *expected growth difference* is for these two countries. The answer is that we expect B to grow at a rate that is 0.12 percentage points higher than A. Naturally, this amounts to be taking the slope estimate (the numerical size of b) very seriously. That is, we need to assume the line in Figure 4 is an adequate description of A and B. Looking at the figure we know this may be problematic; some countries are far from the straight line. Still, as long as the only information we have pertains to aid flows, this is our best prediction.

Notice that the above statement does not involve words like “affects”, “leading to”, or “explaining”. In the most basic form, regression analysis *does not* allow us to say what created the link between aid and growth. This fact is crucial for understanding most of the aid effectiveness debate and this is why the next section discusses this issue in detail, and the refinements of basic regression analysis that—under some circumstances—allows us to attach a causal interpretation to regression results.

4. The cause and effect problem

This section falls in four subsections.¹⁰ To begin, we briefly explain why it is important to move beyond simple scatter plots like Figure 4 when looking at the association between two variables such as aid and economic growth. This takes us from simple regression to multiple regression analysis in order to deal with an issue called “omitted variable bias” in the econometric literature.

Subsequently, we lay out the tricky problem associated with interpreting regression coefficients, such as those recovered through simple and multiple regression analysis. Specifically, we discuss the problem of bi-directional causality, which arises when the amount of aid disbursed to countries has an impact on their growth rate and, at the same time, the growth rate affect the amount of aid a country receives. Bi-directional causality leads to the so-called “identification problem” in econometrics.

It is sometimes argued that the problem of bi-directional causality is more apparent than real in the context of aid effectiveness research.¹¹ Therefore, we next explain exactly why this problem *is* something serious aid effectiveness research needs to deal with.

Finally, against this background, we lay out one approach econometricians have developed in order to deal with the identification problem: instrumental variable estimation.

4.1. Simple and multiple regression analysis

When looking at patterns in the data, like the one depicted in Figure 4, it is natural to wonder about its interpretation. Some analysts quickly jump to the conclusion that it reflects a *casual* relationship: a high aid-to-GDP ratio *causes* low growth. If this is the true state of affairs there is good reason to seriously reconsider aid giving. However, there are several other reasons why a negative association between growth and aid could arise in the data.

For starters, it is possible that some other intervening variable could account for the association. To see how this works, suppose for a moment that aid does not affect growth, and that growth does not affect aid. Hence, there is no causal relationship between aid and growth. Next, imagine donor agencies have agreed to focus the lion’s share of all aid efforts on Sub-Saharan Africa (SSA). Not, suppose, because SSA is a poor region but simply because of its (strategic,

¹⁰ Good formal introductions to the statistical issues dealt with in the present section are given in most introductory econometric text-books. Two popular books are Wooldridge (2008) and Stock and Watson (2007). The introductory textbooks require, however, an understanding of probability theory and statistics at, at least, a high school level. A good introduction to probability and statistics, in Danish, is given in Malchow-Møller and Würtz (2003).

¹¹ See Griffin and Enos (1970) for an early discussion.

say) location. Moreover, consider the possibility that growth in GDP per capita just happens to be lower in countries located in SSA compared to other developing countries. If both propositions are true, aid and growth will be negatively related *even though aid and growth are actually not causally related to one another*. The association is accounted for by the interrelationship between geographical location (SSA), aid donations, and growth.




FIGURE 5. Aid and growth with geographical location as an intervening factor.

Note: The scatter plot depicts hypothetical data in which there is no direct association between aid and growth. African countries have high aid-to-GDP ratios and low growth while non-African (developing) countries have low aid-to-GDP ratios and high growth..

Figure 5 illustrates the point. If countries in Sub-Saharan Africa receive larger aid flows (relative to GDP) and at the same time have lower growth rates than countries in other continents then there is a strong tendency for countries outside Sub-Saharan Africa to cluster in the North-West corner (low aid, high growth) while Sub-Saharan African countries cluster in the South-East corner (high aid, low growth), resulting in a negative association between aid and growth when we look at all countries in the figure.

Naturally, the geographical location of developing countries is not the only possible underlying factor we need to take into account when we try to assess the aggregate effectiveness of foreign

aid. As a guiding principle one should include *all* factors that may have an impact on both economic growth and the allocation of aid when we use regression analysis to assess the impact of aid on growth. At the same time our regression models must be kept reasonably simple if we are to learn anything from them. Hence, almost inevitably, certain determinants of growth will have to be ignored in the analysis. The problem is choosing which to ignore. Since perceived aid effectiveness will be affected by this choice, as we have just seen, it is naturally a contested issue. Indeed, it represents one explanation for the abundance of aid effectiveness studies in existence.




FIGURE 6. Aid and growth determined simultaneously by an aid effectiveness rule and an aid allocation rule: “bi-directional causality”.

4.2. Two rules colliding: The identification problem

Interestingly, though, the biggest problem in evaluating aid effectiveness, by way of regression analysis, arises because politician’s direct aid flows to countries where the resources are perceived to be most needed: the poorest countries. If observed aid flows are distributed according to GDP per capita it becomes virtually impossible to interpret the association between aid and growth that we recover from simple regression analysis. To see why we reconsider the significant negative association between aid and economic growth, found in Figure 4, in the presence of an active aid allocation policy by which aid is directed towards the poorest countries in the world.

Consider Figure 6, which illustrates a possible way to think about the relationships between aid and growth in a country. The figure contains two lines—indicating two ‘rules’. On the one hand, we may hypothesize that more foreign aid increases the growth rate of GDP; this is captured by the upward sloping line that we will refer to as “the aid effectiveness rule”:

$$Growth = a + b \cdot Aid + x.$$

Notice that we allow growth to be affected by other factors beyond aid. These growth drivers are collected in the variable ‘ x ’, which determines the location of the line in Figure 6. For example, we would expect x to be low for a country in Sub-Saharan Africa, while it is high for a non-SSA country. Furthermore, one may imagine that when, say, the level of education raises the variable x increases and the line shifts upwards yielding faster growth.

The slope of the line, b , reflects the *impact* of aid on growth. Hence, if we wish to learn about “the effectiveness of aid” this is the slope one would like to estimate or *identify*. In Figure 6 we *assume* aid increases growth. This is certainly not the impression one is left with after studying Figure 4. Nevertheless, the present illustration will still lead to a “picture” akin to Figure 4 as will be seen.

The other line in Figure 6 captures the aid allocation policy by which a country receives less aid from the donors when it becomes richer. We call this line “the aid allocation rule”:

$$Aid = c + d \cdot Growth + z.$$

Aid allocation is also affected by other things than growth; these aid attractors are collected in the variable ‘ z ’. An example could be child mortality: higher child mortality translates into a higher value for z , which shifts the line upwards, resulting in higher aid levels for all possible growth rates. The slope of the line, d , reflects the *impact* of growth on the amount of aid a country receives.

Taken together the two lines provide an interpretation of how aid and growth is determined in a particular country, during a particular period. Specifically, the actual growth rate and amount of aid received is found as the intersection between the two lines. If the level of aid and the growth rate are determined by these two rules in every aid receiving country this will have profound impact on the interpretation of the data depicted in Figure 4. To see how, we consider some examples.

First of all, we fix the parameters of the aid effectiveness rule and the aid allocation rule. Specifically, we let the two constant terms, a and c , be equal to zero while b is 0.1 and d is -10. Then the rules become

$$\begin{aligned} Growth &= 0.1 \cdot Aid + x \\ Aid &= -10 \cdot growth + z \end{aligned}$$

Based on the two equations, ‘observations’ for aid and growth are completely determined by the growth driver, x , and the aid attractor, z . In Table 1 we show hypothetical outcomes for three countries, A, B and C for which the growth driver is the same while the aid attractor differs across countries. This means that the countries have exactly the same aid allocation rule while the location of the aid effectiveness rule varies.

Table 1: Aid and growth outcomes for three hypothetical countries having the same value for the growth driver

Country	Growth driver x	Aid attractor z	Resulting aid flow aid	Resulting growth $growth$
A	1	11	0.50	1.05
B	1	21	5.50	1.55
C	1	31	10.50	2.05




FIGURE 7. An example of an observed association between aid and growth when the aid allocation rule differs across countries while the aid effectiveness rule is the same. The aid and growth observations from Table 1 are plotted in Figure 7 in order to show the aid and growth information in the same way as in Figure 4. We cannot see the two rules but the hypothetical data illustrate a very clear linear relationship between aid and growth and the regression line indicated in the Figure has a slope of 0.1, which is equal to the aid effectiveness

parameter, b . In fact, we are actually tracing out the common aid effectiveness rule by combining the data points.

In Table 2 we consider another set of countries, A', B' and C'. The three countries have exactly the same aid effectiveness and aid allocation rules as before, but they differ in the values of their growth drivers and aid attractors. The three new countries share the same value of the aid attractor while the values of the growth drivers differ. Thereby the countries have a common location of the aid allocation rule while the location of the aid effectiveness rules varies.

Table 2: Aid and growth outcomes for three hypothetical countries having the same value of the aid attractor

Country	Growth driver x	Aid attractor ξ	Resulting aid flow aid	Resulting growth $growth$
A'	0	21	10.50	1.05
B'	1	21	5.50	1.55
C'	2	21	0.50	2.05




FIGURE 8. An example of an observed association between aid and growth when the aid effectiveness rule differs across countries while the aid allocation rule is the same.

The observations from Table 2 are plotted in Figure 8. What emerges is a completely different picture of the relation between aid and growth. The relationship is linear but the slope of the

combining regression line is -0.1 in Figure 8 compared to 0.1 in Figure 7. The slope we observe in Figure 8 is actually the inverse of the slope of the aid allocation rule: $1/d = 1/(-10) = -0.1$. The reason is that in Figure 8 the aid effectiveness rule varies while the aid allocation rule is common for the three countries and what we see is the aid allocation rule (turn the page counter clock wise) not the aid effectiveness rule.




FIGURE 9A. An example of an observed association between aid and growth when both the aid effectiveness rule and the aid allocation rule differs across countries.

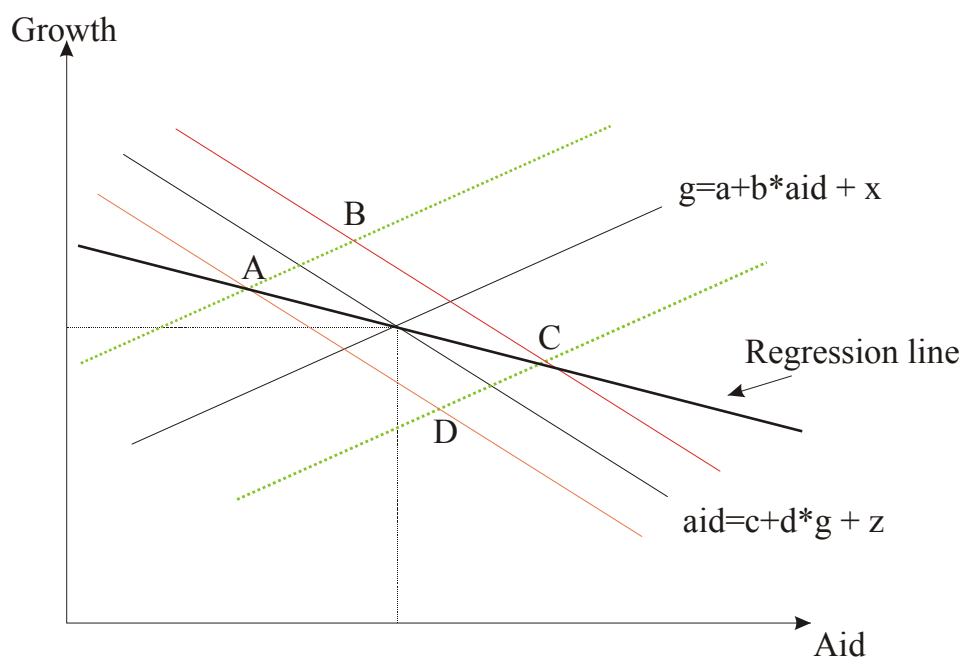


FIGURE 9B. The observed association between aid and growth when both the aid effectiveness rule and the aid allocation rule vary across countries.

Moving a little closer to the real world we know that every country will have their own value of both the growth drivers, x , and the aid attractors, z . So, the actual observations will be scattered around as in Figure 4. To illustrate, in Figure 9A we plot observations for the nine hypothetical countries that can be observed when we combine the three values of x (0, 1, 2) and the three values of z (11, 21, 31) from tables 1 and 2. The relationship between aid and growth will be positive or negative, depending on your choice of angle. However, the regression line in Figure 9A has a slope of zero, indicating no systematic relationship between aid and growth.

Figure 9B illustrates the general problem of having two rules generating the actual data. We have a lot of different, parallel, aid effectiveness rules (one for each country in a particular period) and just as many different, parallel, aid allocation rules. For the illustration we assume there is a maximal and a minimal value of the growth drivers, x , leading to a maximal and a minimal aid effectiveness line. Likewise we assume a maximal and a minimal level for the aid attractor, z , giving a maximal and a minimal aid allocation line. Since each data point, according to this simple model, is thought to reflect a point of intersection between the two rules, all data observations would appear in the area ABCD. If we use ordinary least squares regression to assess the link between aid and growth, the resulting line would tend to go through the points A and C. Crucially, observe that the slope of this line will not be equal to either the aid effectiveness rule or the aid allocation rule—no matter how much data we get. The line will always be a mix of the two rules. But notice that if the variation in the aid effectiveness rule is larger than the variation in the aid allocation rule then the regression line will be closer to the

aid allocation rule. The opposite is also true; if the variation in the aid allocation rule is larger than the variation in the effectiveness rule we get closer to the aid effectiveness rule.¹²

In essence, the examples show that if our data is the outcome of two rules then just observing the data points we have no way of knowing if we are estimating one rule or the other.¹³ The bottom line is that a regression coefficient cannot be interpreted as reflecting a causal impact of aid on growth (or growth on aid, for that matter). This is what economists' call "the identification problem".

Figure 10 summarizes the possible interpretation of (any kind of) correlation between foreign aid and economic growth. In the example from Section 4.2 "geography" is an "intervening variable", whereas the bi-directional link between aid and growth, illustrated by the two separate lines in Figures 6-9, is captured by the arrows connecting the growth and aid "boxes" in the figure."

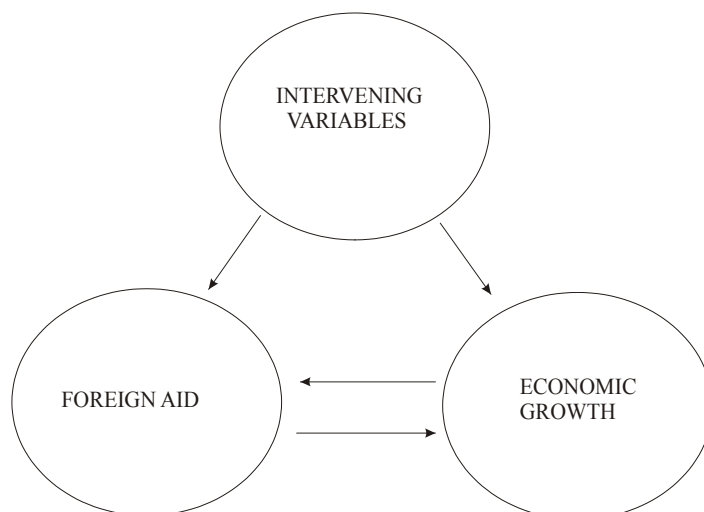


FIGURE 10. Possible reasons for a correlation between foreign aid flows and economic growth.

4.3. *Is it really a problem?*

At times one may come across research where the identification issue is not recognized, or, is "winked away". In the best of cases there is an argument in favour of ignoring the problem. If so the argument is that the notion of "bi-directional" causality, leading to the identification problem, is a fallacy. After all, there is little (if any) evidence that aid is given to the countries that *grow* at the slowest speed. To be sure, there is ample evidence that aid is given predominantly to the *poorest* countries. Consequently, if the analysis focused on the link between

¹² Technically speaking the two curves could also be shifting around due to statistical disturbances, which affect the individual curves. Hence, these shifts need not reflect differences in other determinants of growth and aid.

¹³ For a mathematical description, see Annex 1.

aid and the *level* of income the identification issue would be very real. However, the argument goes, there is likely *no* reverse causality problem between aid and *growth* of income. In other words, the aid allocation rule in Figure 6 does not exist. If true, all a researcher needs to do is to include the level of income per capita as an intervening factor in the growth regression and the problem is solved. Unfortunately, this reasoning is flawed.

Understanding this point is critical because we, in effect, dismiss a large part of about 40 years of scholarly research on the topic. Most aid effectiveness analyses before 1995 did not take bi-directional causality into account (see Hansen and Tarp, 2000). Therefore, to prove the point, we proceed in small steps.

To simplify the exposition assume aid has no causal impact on economic growth. That is, assume the slope, b , in the aid effectiveness rule, that we are trying to find, is zero.

Next, consider two countries that are identical with respect to GDP per capita and aid flows in 1970. That is, they are of equal size, equally rich and they receive the same amount of aid. Imagine (for now) the two countries receive a *constant* flow of aid, measured in real US dollar per capita, each year from 1970 to 2000. Now, let's assume one country (A, say) is hampered by problems leading to zero growth in GDP per capita, on average from 1970 to 2000 while the other country (B) is doing better, experiencing an average growth in GDP per capita of two percent a year during the same period of time.

While we thus know that the two countries receive exactly the same amount of aid per capita, and that this translates into the same share of aid-to-GDP in 1970, it is also clear that they will not have the same aid-to-GDP ratio in 2000 (or in any other year after 1970 for that matter). If both countries have an aid-to-GDP ratio of five percent in 1970, then country A will experience an *average* ratio of exactly five percent over the period, as the aid flows are constant and the average annual growth rate is zero. However, the annual aid-to-GDP ratio will not be constant in country B; it will be declining because the aid flow is constant while GDP per capita is increasing. A few calculations show that when the average growth rate is two percent a year, the average aid-to-GDP ratio, from 1970 to 2000, is roughly four percent in country B. Hence, the fastest growing country will have the lowest observed aid-to-GDP ratio. Yet, in this simple example the aid-to-GDP ratio is low precisely because the country is growing rapidly; not because aid is harmful.

This is a general result: In a world with constant aid flows in per capita terms and different growth rates in GDP per capita, we are faced by a “virtual” allocation rule showing a negative

association between aid (to GDP) and growth.¹⁴ Note that the larger the differences in the growth rates, the steeper the slope of this allocation rule.

But aid flows are clearly not constant over time. In fact, it is a widespread finding that, once intervening factors are controlled for, aid per capita decreases with the *level* of GDP per capita. The question is how this will impact on the aid-*growth* association that we observe in the data.

Again, we simplify to illustrate the effect in a transparent way. So, imagine the donors make their aid allocation decisions collectively once every decade, starting in 1970. Further, assume the donors follow an allocation rule saying that when GDP per capita is increased by one percent in a country, aid per capita is cut by one percent.

Now, reconsider countries A and B. They have the same aid flows per capita in the 1970s because they have the same GDP per capita in 1970. The initial aid-to-GDP ratio is five percent. But due to the differing growth rates—zero percent in country A and two percent in B—they will not receive the same aid flows from 1980 onwards. In fact, in 1980 country B is more than 20 percent richer than country A. (22 percent to be precise). This means that in the 1980s aid per capita will be cut back by 20 percent in country B. Country A, by contrast, will receive the same amount of aid since its income per capita level is the same. Note, however, that as country B is both richer and receives less aid per capita in 1980, the aid-to-GDP share is now more than 40 percent lower in country B compared to country A. This pattern is amplified after 1990 since country B is now almost 50 percent richer than country A (which still is at the 1970 level of both countries), implying aid flows to B are cut by 50 percent compared to the 1970 level. In 2000 country B is about 80 percent richer than country A because of the difference in the average growth rates of two percent per year. When we calculate the *average* aid-to-GDP share over the 30 years, 1970 – 2000, country A will have the same share as in the case of constant disbursements: five percent. In country B we have three decades of (step-wise) declining aid flows coupled with persistent growth in GDP per capita; the average aid-to-GDP ratio is just above three percent for country B; lower than the four percent average calculated with constant aid flows.

The general point to take away from this illustration is that the negative association between the average aid-to-GDP ratio and *growth* becomes more pronounced when donors are active in periodic reallocation of aid guided by the *level* of GDP per capita. Without reallocation of aid funds the fast growing country (B) obtained an average aid-to-GDP ratio of four percent; with reallocation the ratio shrinks to about three percent.

¹⁴ One may view the link as “virtual” since it is not a consequence of donors *deliberately* allocating lower aid flows to fast growing countries, or poor countries for that matter.

Aid allocation is of course not a decision made at collective meetings for all donors once every decade, but a mix of decisions made by individual donors at varying intervals. This will, however, only tend to strengthen our result, which establishes a strong *causal* influence of growth on aid.¹⁵

The identification problem, as discussed in Section 4.2, is therefore very real. Accordingly, to elicit information about the impact of aid on growth we need to deal with it; it cannot reasonably be winked away.

4.4. A possible solution to the cause and effect problem

Ideally one can imagine a simple fix to the identification problem; randomized trials. Such experiments, which are akin to the testing procedures used when new drugs are evaluated, have many supporters at the project level.¹⁶ Evaluating aid at the country level we need to imagine random helicopter drops of aid money across the globe. In this universe any correlation between aid and economic growth could not be ascribed to correlations with other variables (the intervening variables problem), nor reverse causality.

Naturally, this is not really an option (politically, at least): few people would find it reasonable for aid to be disbursed to middle income countries, like Argentina, rather than the most in need, just to elicit information about whether it works or not. And this could be the outcome of the random trial. Hence, absent a good experiment economists have to rely on statistical methods to try to resolve the problem at hand. A way forward is what econometricians call “instrumental variables estimation”.

To understand the basic logic of the approach, consider Figure 11, which is a slightly augmented version of Figure 10: the new feature being the box labelled “instrument”. What this box is supposed to encompass is a factor which affects aid disbursements. Yet, it is not just any determinant of aid flows. Notice that the box involving “instrument” is not connected to *any* other box, aside from aid. Hence the instrument does not affect growth (above its impact via aid), it is independent of other factors (intervening variables) that explain growth, and it is not itself explained by growth.

¹⁵ For a mathematical proof of this point, see Annex 2. A proof of the reverse causality result in a much more general setting is given in Dalgaard, Hansen and Tarp (2004).

¹⁶ For a good discussion about randomized trials in an aid and development context, see Banerjee (2007).

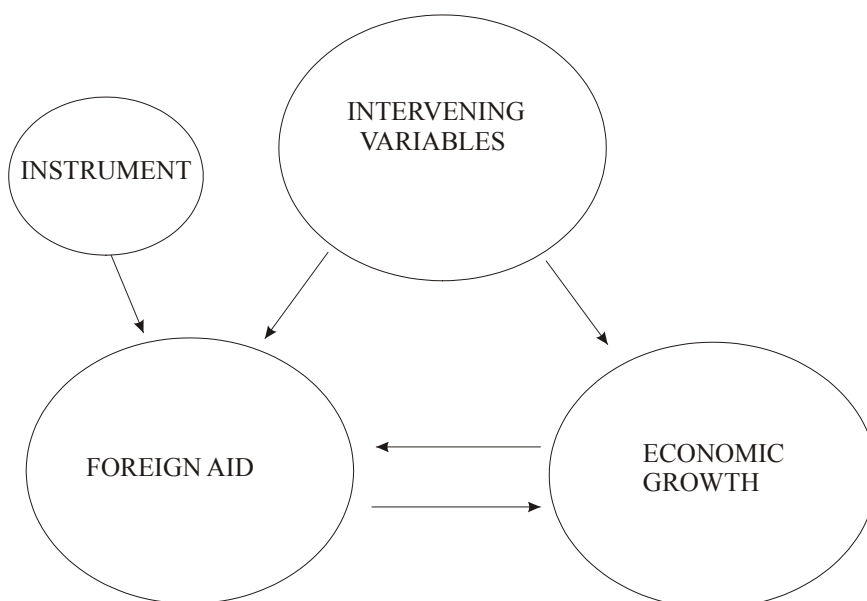


FIGURE 11. Teasing out the causal effect from aid to growth using instruments.

To see how this works, suppose for a moment that we have located such a variable. To fix ideas let us call it “size”, measuring the size of countries by, say, the population size. We can now perform the following “statistical experiment”. First we figure out how much of the cross-country differences in aid levels we can explain with the country size variable. That is, we quantify the strength of the arrow between the instrument (country size) and aid in Figure 11. Second, we take the amount of aid that size—and only size—can motivate in each country in the world, and look at the association with observed growth rates in these countries. If a significant association prevails we say that this is because aid *affects* growth.

Why must this be the case? If there is a link between size-generated-aid and growth rates we know this correlation cannot be explained away by reverse causality: growth in GDP per capita does not explain the size of a country (cf. the absence of any arrow from growth to “instrument” in Figure 11). Hence, reverse causality is ruled out. Next, the association cannot be explained away by intervening variables either, since country size is not related to these variables (cf. the absence of any link between “instrument” and intervening variables in Figure 11). Finally, country size does not explain growth directly (cf. the absence of any arrow from “instrument” to growth in Figure 11). If all of this is true, the *only* interpretation of the association which is left to us is that aid *affects* growth. The other options have been ruled out. Figure 12 provides a geometric interpretation of this “instrumental variable method”.




FIGURE 12. The Geometry of instrumental variable estimation. *Notes:* In the text we refer to “z” as country size. The bottom panel shows the link between country size (z) and aid allocation. The top panel concerns the link between aid and growth.

In the bottom panel we find the instrument (i.e., “size”) depicted on the vertical axis, denoted ‘z’. The line in the bottom panel reflect that size is related to foreign aid: a lower size goes along with a higher aid-to-GDP ratio. As size varies we get variation in the aid-to-GDP ratio. The average association between country size and aid is given by the “regression line” drawn in the bottom part of the figure. In the top panel we have illustrated the two aid-growth rules; aid effectiveness and aid allocation. The key thing to observe is that when country size varies we

get variation in the allocation rule which is independent of the effectiveness rule—such that the allocation rule moves while the effectiveness rule is fixed. This is exactly the situation in which the data will trace out the aid effectiveness rule as depicted in Figure 7. Consequently, we have thereby identified the impact of aid on growth statistically.¹⁷

Despite the apparent simplicity, the major problem with this procedure should be clear: it is hard to find a determinant of aid flows which fulfils the requirements for it to act as an instrument for aid. To see how the experiment could fail, consider the following line of reasoning. Suppose smaller countries do in fact receive more aid per capita and as a fraction of GDP per capita, as assumed above. This makes population size a candidate for being an instrument. However, suppose the population size affects growth *directly*. That is, imagine population size has an impact on growth above and beyond its potential impact via aid. In particular, imagine “size is good for growth”.¹⁸ If so, we are back to square one. A negative association between aid (as explained by country size) and growth may now be taken to imply either a negative impact of aid flows on growth, or, that big countries are growing faster due to their size, and simultaneously receive less aid. In terms of Figure 12, the problem can be seen as that of the instrument shifting *both* lines around, rather than just the aid allocation rule. In this case the procedure does not trace out the slope of the aid efficiency rule. We have thus not solved the problem of identifying the impact of aid on growth.¹⁹

These difficulties have been a main impetus for the aid effectiveness debate over the last decade or so. We take up another reason for the debate later in this study. But first, in the next section, we provide an illustration of how estimates change, when we try to deal with omitted variable bias and, in particular, the identification problem.

¹⁷ A mathematical demonstration of this point is found in Annex 3.

¹⁸ There are several reasons why this could be the case. For instance, in larger countries there may be greater scope for division of labor which stimulates growth; an idea that goes back to Adam Smith’s “Wealth of Nations”.

¹⁹ Another problem arises if the association between the instrument and aid is weak. This problem shows up as a very steep regression line in the bottom part of Figure 12. If the line is steep there is virtually no variation in the aid allocation rule in the top panel making it impossible to determine the slope of the aid effectiveness rule.

5. An illustration of the regression approach

If indeed there is a severe identification problem to deal with we should be able to illustrate the consequences of taking it into account using actual data. This section provides such an illustration. The data we use are from a recent study by Rajan and Subramanian (2008) and we begin the illustration by presenting the regression line in Figure 4. Expressed as an equation this line reads as

$$\begin{aligned} \textit{Growth} &= 1.92 & - & 0.12 \cdot \textit{Aid} \\ &(0.31) & & (0.04) \end{aligned}$$

The two parameters of the equation are estimated using ordinary least squares. Below the parameter estimates we show the standard errors of the estimates in parentheses. The standard errors give an indication of the precision of our estimates and, thereby, of the confidence we can have in the specific values. In general, small standard errors indicate precise parameter estimates and, as a rule of thumb, we say that if a parameter divided by its standard error is above 2, in absolute terms, then the parameter is statistically significant. This means that we consider the parameter to be different from zero whereby the variables we study are considered to be systematically related. In the equation above the slope coefficient divided by its standard error is -3, leading us to the conclusion of Section 3, that the negative association between aid and growth is highly significant.

Turning to the possible impact of intervening factors we use the example of Section 4 and question if (some of) the negative association is driven by a geographical factor. In Figure 13 we show the same cross-plot as in Figure 4, this time highlighting the geographical location of the countries. In the sample of 78 countries, 32 are located in Sub-Saharan Africa while 46 are non-SSA countries.




FIGURE 13. The association between the aid-to-GDP ratio and average growth in GDP per capita (PPP) 1970-2000. Highlighting countries in Sub-Saharan Africa (circles) and other countries (diamonds).

Data Source: Rajan and Subramanian (2008).

Obviously, using real life data, we do not get a complete split like the one indicated in Figure 5 but there is a clear tendency in the data: countries in Sub-Saharan Africa tend to get more aid and have lower growth rates. The tendency is also illustrated by the very different average growth rates and aid-to-GDP ratios in the two country groups, given in Table 3. On average, the countries in Sub-Saharan Africa received foreign aid amounting to ten percent of their GDP, while their average growth rate was only 0.2 percent. In contrast countries outside that region only received aid in the order of 2.4 percent of their GDP while their average annual growth rate was almost two percent per year. As discussed in Section 4, this difference in averages may be an underlying factor for the negative association between aid flows and growth.

	Average of all 78 countries	Average of 32 countries in Sub-Saharan Africa	Average of 46 countries outside Sub-Saharan Africa
Growth	1.2	0.2	1.9
Aid-to-GDP ratio	5.5	10.0	2.4

TABLE 3. Average growth rates and Aid-to-GDP shares inside and outside Sub-Saharan Africa

This assertion can be investigated using multiple regression analysis. We simply augment the regression above with a variable that takes the value 1 for all countries in Sub-Saharan Africa and the value 0 for other countries. (The variable is denoted *SSAfrica*). The resulting regression is

$$\begin{aligned}
 \text{Growth} = & 2.10 - 0.06 \cdot \text{Aid} - 1.22 \cdot \text{SSAfrica} \\
 & (0.27) \quad (0.06) \quad (0.68)
 \end{aligned}$$

Geographical location appears to explain quite a bit of the negative association between aid and growth. When we include the information about location in the regression the slope of the line is halved, from -0.12 to -0.06, and the coefficient divided by its standard error is now only -1, which is clearly less than 2 in absolute value, leading us to conclude that—‘conditional on geographical location’—the association is not statistically significant.²⁰

The new result can be illustrated in two ways. A very common way is to show that geographical location is actually moving the aid-growth line, as illustrated by the inclusion of ‘x’ in the figures in Section 4. Figure 14 shows how we can draw two parallel fitted lines in the cross-plot, one for countries in Sub-Saharan Africa and one for other countries. As the lines are parallel the slopes are equal (-0.06). The intercept for ‘other countries’ is 2.10 while the intercept for the Sub-Saharan African countries is 0.88 (= 2.10-1.22).

²⁰ If we divide the coefficient upon the geography variable by its standard error we get -1.8, which is also less than 2 in absolute value. Hence, we must be very careful in drawing too strong conclusions based on this regression.




FIGURE 14. The association between the aid-to-GDP ratio and average growth in GDP per capita (PPP) 1970-2000. With regression lines for Sub-Saharan African countries and other countries.

Data Source: Rajan and Subramanian (2008).

Figure 14 explains how we should interpret the regression above but it does not really tell us why the slope changes when we include information about geographical location. For this we need another cross-plot, given in Figure 15. In this cross-plot the observations are derived from the original data by subtraction of averages. Specifically, for the 32 countries in Sub-Saharan Africa we subtract the group averages; 10.0%, from the aid-to-GDP ratio and 0.2% from the growth rates. For the 46 other countries we subtract their averages; 2.4%, from the aid-to-GDP ratio and 1.9% from the growth rates. By subtraction of these averages we bring the countries in the sample ‘on par’ which is what is meant when we say that we ‘control for’ the effect of geographical location.




FIGURE 15. The association between the aid-to-GDP ratio and average growth in GDP per capita (PPP) 1970-2000 when group averages are subtracted.

Data Source: Rajan and Subramanian (2008).

Based on the data in Figure 15 we can again draw the best fit line to illustrate the linear association between aid and growth, which is now *conditional* on geographical location. The slope of the line in Figure 15 is -0.06 —exactly the slope in the regression equation above. The intercept is zero by construction because we have removed the means of the two variables. The important message to take away from Figure 15 is that when we include other variables in a regression, we ‘control’ for the effect of the variables by generating new observations. The new observations are in some sense more ‘true’ than the original because they are not blurred by the intervening factors. In general, we can control for as many intervening factors we like. This can always be done by generating new observations showing the conditional association between aid and growth.

Moving forward towards an instrumental variable regression, we first condition on a few more intervening factors. Specifically, we add a measure of legal institutions, capturing protection of property rights, bureaucratic quality, absence of corruption and so forth, and the logarithm of

the value of GDP per capita in 1970.²¹ In total our regression has four explanatory variables and an intercept.

$$\begin{array}{rcccccc}
 \textit{Growth} & = & 11.17 & -0.09 \cdot \textit{Aid} & -1.97 \cdot \textit{SSAfrica} & +9.49 \cdot \textit{Institutions} & -1.78 \cdot \textit{GDPcap}(1970) \\
 & & (2.31) & (0.05) & (0.56) & (1.86) & (0.29)
 \end{array}$$

In this regression, all factors, except for aid, have significant slopes measured by the absolute size of the slope divided by standard error. The effects of the three intervening factors are in line with economic reasoning: growth is lower in Sub-Saharan Africa, as already discussed; “good institutions” are associated with higher growth rates, and initially richer countries tend to have lower growth in the subsequent years—everything else being equal.²²

The conditional association between aid and growth is again negative, with a slope of -0.09, which is nicely between the -.12 we found in the simple regression and -.06 that came out when we only condition on Sub-Saharan Africa. Variation in the estimated slope of this order of magnitude is quite common and it shows that aid flows are correlated with the other factors in the regression model. Again, we find that the coefficient divided by its standard error is less than 2 in absolute value whereby we conclude that the conditional association between aid and growth is statistically insignificant. A cross-plot of the aid and growth observations, in which variation due to geography, institutions and initial richness has been removed, is given in Figure 16. The slope of the regression line in the figure is -0.09 as reported in the regression above and the intercept is zero by construction.

We can now turn to the instrumental variable regression. Rajan and Subramanian (2008) have constructed an instrument, partly based on country size but also on other factors, that should not have a direct impact on growth.²³ By the instrumental variable regression approach we must first establish the linear association between aid and the instrument and, further, we must ensure that the association is not driven by the intervening factors already included in the model. This is done via a regression, known as the first-stage regression in the econometric literature. The regression has aid as the dependent variable while the instrument and the intervening factors are the explanatory variables:

²¹ How to measure the quality of legal institutions is an area of research in itself. We use a measure which is very popular within economics. The data is commercial and known as the international country risk guide (ICRG). See <http://www.prsgroup.com/ICRG.aspx> for a description of the data and Rajan and Subramanian (2008) for a description of the transformation of the ICRG data in the regression analysis.

²² The last result is known as “conditional convergence” in the growth literature. See Barro (1991) for a discussion and early evidence of the phenomenon.

²³ At this stage we will not go into detail with the construction of the instrument, nor whether it is plausible or not. But one may observe that it derives directly from a recent scholarly article (Rajan and Subramanian, 2008), which is forthcoming in a leading economics journal. Hence, the instrument has been viewed as sufficiently convincing to appease several anonymous reviewers and an editor of the journal. In our companion report, Dalggaard and Hansen (2008), we discuss the study in more detail.




FIGURE 16. The association between the aid-to-GDP ratio and average growth in GDP per capita (PPP) 1970-2000 when geography, institutions and initial GDP per capita is controlled for.

Data Source: Rajan and Subramanian (2008).

$$\begin{aligned}
 \text{Aid} = & 32.22 + 0.62 \cdot \text{Instrument} + 1.97 \cdot \text{SSAfrica} - 3.00 \cdot \text{Institutions} - 3.76 \cdot \text{GDPcap}(1970) \\
 & (4.95) \quad (0.12) \qquad (1.03) \qquad (3.79) \qquad (0.76)
 \end{aligned}$$

In the first-stage regression it is not the size of the parameters but rather the statistical ‘strength’ of the association that is important. This strength is typically measured by the value of the coefficient divided by the standard error, exactly as when we evaluate the significance of the regression parameters. The first stage regression above shows that the instrument has a strong linear association with aid, also when we take account of the intervening factors as the coefficient divided by its standard error is just above 5. The conditional association between aid and the instrument can be shown graphically just as the association between aid and growth in

Figure 16, because the first stage regression is just a standard regression. The association is shown in the bottom part of Figure 17.




FIGURE 17. The association between the instrument and aid-to-GDP ratio 1970-2000 when geography, institutions and initial GDP per capita is controlled for. (First-stage regression, bottom part). And the association between the filtered aid-to-GDP ratio and average growth in

GDP per capita (PPP) 1970-2000 when geography, institutions and initial GDP per capita is controlled for (Second-stage regression, top part).

Data Source: Rajan and Subramanian (2008).

Now, we can use the regression line in the cross-plot to generate new aid data that are unaffected by growth; “exogenous” in the econometric terminology. Technically, we can move all 78 aid observations in Figure 17 horizontally to be aligned on the regression line.

The new aid observations are now compared with realized growth rates. Specifically, the association between the growth and aid is obtained by regressing the new aid observations on growth in a final regression, called “the second-stage”. The regression results are follows:

$$\begin{array}{rcccccc}
 \textit{Growth} & = & 5.11 & +0.13 \cdot \textit{Aid} & -3.02 \cdot \textit{SSAfrica} & +10.9 \cdot \textit{Institutions} & -1.19 \cdot \textit{GDPcap}(1970) \\
 & & (3.01) & (0.06) & (0.53) & (1.98) & (0.36)
 \end{array}$$

The cross-plot of aid and growth, after controlling for the three intervening factors and reverse causality, is in the top part of Figure 17, and is comparable to the illustration in Figure 12. The linear association is now positive with a slope of 0.13—quite the opposite of the simple association—and the slope is significant in the sense that the coefficient divided by its standard error is larger than 2. Hence, based on these new regression results we would say that an exogenous, permanent, increase in aid flows to a country, amounting to one extra percent of GDP, *causes* an *increase* in the average annual growth rate of about 0.13 percentage points.

In sum, this exercise illustrates that the choice of regression approach matters. Shifting from the simple regression framework to instrumental variable regression modifies our impression of aid effectiveness. The findings can be summarized as follows:

1. Regressions using ordinary least squares (OLS) of aid and growth will be misleading unless *all* relevant common determinants of growth and aid are taken into account.
2. Regressions using the method known as Ordinary Least Squares (OLS) will confound the impact of aid on growth with the opposite chain of causality—which is an allocation rule that captures a negative impact of growth on aid flows. The resulting linear association may be positive or negative, but it is not a causal relationship from aid flows to growth rates.
3. Instrumental variable regressions, such as the method known as Two-Stage Least Squares (TSLS), can in principle allow the researcher to disentangle the cause and effect of foreign aid. If aid is properly instrumented we can estimate the aid effectiveness rule.

While the results of our illustration support an upward sloping aid efficiency rule, yielding a positive impact of aid on growth, one should *not* leap to such a conclusion. In part because the reverse causality problem is not the only problem researchers in this literature face (cf. Section

6). For a critical assessment of what the evidence on the topic shows—and does not show—including the study by Rajan and Subraminian (2008), we refer to the companion paper, Dalgaard and Hansen (2009).

6. Yet another specification problem: The impact may vary

Up to this point we have assumed the aid effectiveness rules have equal slopes for all countries. This assumption is a natural starting point in any regression analysis. Moreover, the assumption should not be taken literally. We do *not* require, nor expect, the effectiveness of aid to be exactly the same in all countries at all times. When we interpret regression models like the ones presented in Section 5, the idea is that the slope coefficient (properly identified and estimated) represents *an average impact* of aid on economic growth. The underlying (unknown) impacts in each country may vary, but—and this is the crucial point—the variation cannot be related to any of the intervening factors in either the effectiveness or the allocation rule. If this condition is met, estimates such as those in Section 5 are in fact good estimates of the average impact. However, if the impact of aid varies from country to country in a systematic fashion the analysis may not be capturing the average impact of aid on growth.




FIGURE 18. Two aid effectiveness rules with slopes depending on an intervening factor.

To illustrate, consider Figure 18, which shows effectiveness rules for two countries. Country A, which could be a country in Sub-Saharan Africa, has a low value of the intervening factor, x , leading to a low level of the aid effectiveness rule. Country B, outside Sub-Saharan Africa, has a

higher value of the factor ‘ α ’ and, thus, a higher level of the effectiveness rule. In the figure, the two effectiveness rules have both different levels and slopes. The slope of the effectiveness rule for country B is steeper than the slope of country A’s rule: aid is ‘more effective’ in B than in A in the sense that an equal increase in aid in the two countries will cause a larger growth effect in country B than in country A. Even if we condition on geographical location, as in Sections 4 and 5, we will not end up with a single (conditional) effectiveness rule for the two countries. Further, the average slope is convoluting information that may be crucial for our understanding of the effectiveness. The lacking information is, in effect, an omitted variable problem as the one discussed in Section 4.1.

Another example is given in Figure 19, in which we have three countries. First, consider the case with a single effectiveness rule (line \mathcal{A} in the figure) and three allocation rules. In this case the rule can be traced from the intersections as indicated by the three stars. This is the story from Section 4 (see Figure 7).




FIGURE 19. Aid effectiveness rules and aid allocation rules for three countries.

Now, assume, instead, that the effectiveness rules for the three countries have different slopes, indicated by the three lines, A , B and C . The variation in the slopes is not related to a factor that changes the level of the curves. Hence, the situation is not as in Figure 18, and as such the

different slopes need not pose a problem—until we notice how the slopes are related to an underlying factor that changes the location of the allocation rules: big slope coefficients (high impact countries) are linked to low allocation rules whereby the allocation rules are negatively related to the slopes of the effectiveness rules. The resulting observations are the three dots.

Imagine again how we only have the observations, not the rules. Connecting the dots leads to a very flat regression line, which is not the average slope of the effectiveness rules. Interestingly, if the slopes of the effectiveness rules are related to intervening factors, as in the figure, we again have the omitted variable problem presented in Section 4.1.

To deal with the challenge posed by Figures 18 and 19 we need to specify and estimate the relationship between the slopes and the intervening factor. The solution is illustrated in Figure 20, which is like Figure 11, but there is a new box of intervening factors, affecting the “effect” directly.

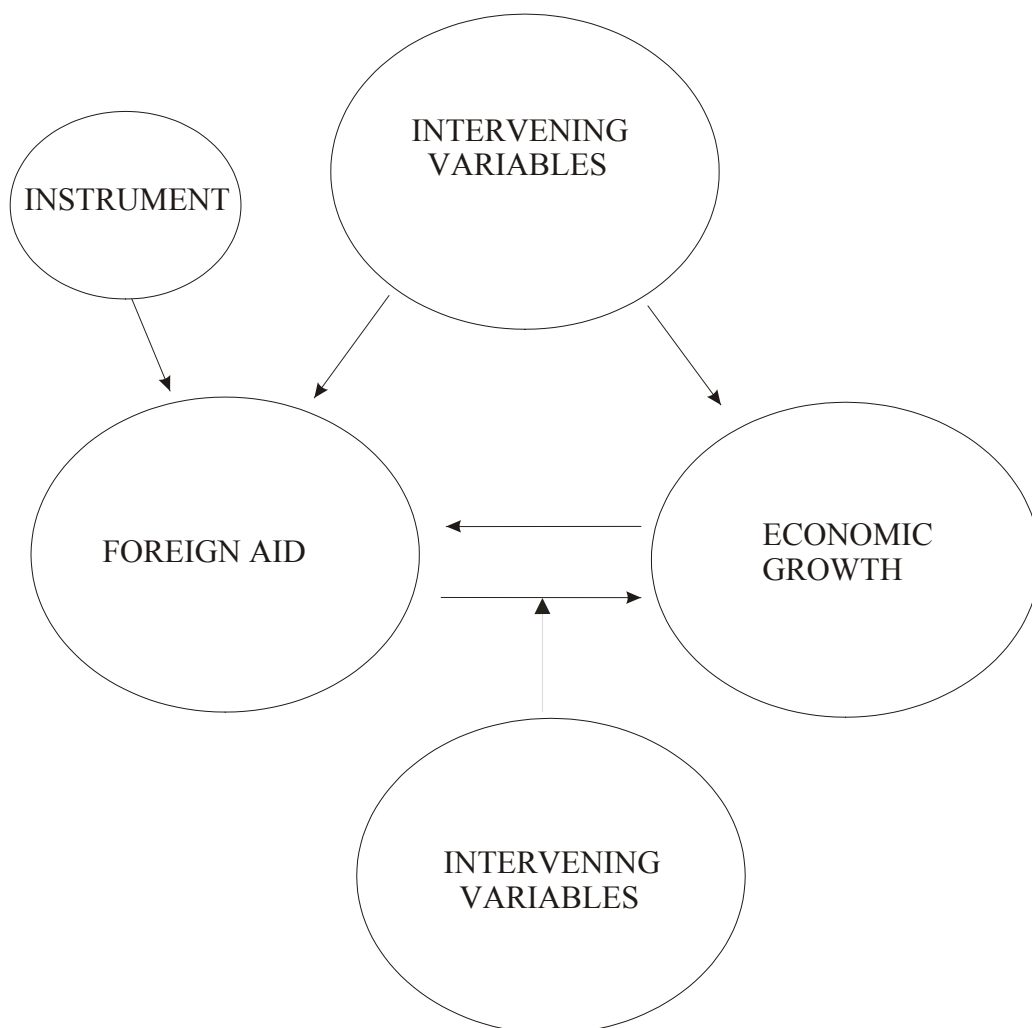


FIGURE 20. Possible reasons for a correlation between foreign aid flows and economic growth and for varying efficiency across countries.

Mathematically, the problem can be illustrated as a missing variable and a missing equation: we need an equation for the slope of the effectiveness rule in addition to the rule itself.

$$\begin{aligned} \text{Growth} &= a + \text{Effect} \cdot \text{Aid} + x \\ \text{Effect} &= b + f \cdot w \end{aligned}$$

Here, the first equation specifies the location of the effectiveness rule with ‘ x ’ as an intervening variable. We no longer have a unique slope parameter ‘ b ’ in the equation. Instead, the impact varies, and is labelled “*Effect*”. The second equation specifies the slope of the effectiveness rule (*Effect*), and here the intervening variable is ‘ w ’, which can be related to variables in either the effectiveness rule or the allocation rule.

While we have observations for the growth rate, aid and the intervening factors, we will never directly observe the slope of the effectiveness rule. Fortunately, we might be able to cope nevertheless since “*Effect*” can be substituted into the effectiveness rule. The result is a new, more complex, rule

$$\text{Growth} = a + b \cdot \text{Aid} + f \cdot (w \cdot \text{Aid}) + x$$

This is a multiple regression equation. But this equation features a new variable, $w \cdot \text{Aid}$, representing “a non-linear effect of aid on growth”. Hence, we can no longer speak of ‘*the* impact of aid’. Rather, we need to incorporate the intervening variable, w , and based on this value, the effectiveness of aid is found as “*Effect*” in the equation above.

During the past decade aid effectiveness researchers have spend considerable mental power generating, and discussing, ideas about systematic relationships between the slope of the effectiveness rule and intervening factors. Broadly speaking, the ideas can be classified into two broad groups: (i) the impact of aid varies due to recipient characteristics or (ii) the impact varies due to donor characteristics. Naturally, these two sets of explanations for a country specific impact of aid are not mutually exclusive.

In the companion study, Dalgaard and Hansen (2009), we discuss some of the empirical analyses within each of the groups in detail. Therefore, we will only sketch the main ideas below before moving to a brief discussion of the added challenges and requirements in terms of estimation strategies.

6.1. The Effect depends on the characteristics of the recipients

When researchers look at determinants of growth and development it is quite common to ask if the impacts of the main determinants vary systematically with other factors. Durlauf and Johnson (1995), for example, demonstrate that the impact of investments differ across countries at different stages of development. Another relevant example is given in Block (2001) and in Masajala and Papageorgiou (2008) who argue forcefully that the growth process in Africa exhibit unique characteristics compared to the process observed in the rest of the world.

Within the aid effectiveness literature, Burnside and Dollar (2000) argue that a sound macroeconomic policy is beneficial for growth and, at the same time, it increases the impact of aid on growth—thereby providing a double bonus. Despite being somewhat discredited (see, among others, Easterly, Levine and Roodman, 2004) the basic idea in Burnside and Dollar of looking at important characteristics in the recipient countries has received broad support in donor organizations and several later studies have looked at characteristics other than macroeconomic policies. Quite generally, these ideas are what we illustrate in Figure 18.

6.2. The Effect depends on characteristics of the donors

If we turn to look at the donors, it is clear that having more than one donor may also create systematic variation in the effectiveness of aid disbursements across the recipient countries. The many donors have different allocation policies which may be important in at least three respects. First, the composition of the aid disbursements varies across recipient countries, in part, because donors have different ‘tastes’ and ‘traditions’ regarding projects and the sectors they choose to support. Second, the relative importance of poverty orientation (or GDP per capita) vis-à-vis other development (and political) factors varies across donors whereby the recipient countries may actually face allocation rules with different slopes.²⁴ A third issue is donor harmonization. The problem being that aid, in most poor countries, is disbursed as hundreds of separate donor-managed projects. For example, in Vietnam in 2002, 25 bilateral and 19 multilateral donors and about 350 international NGOs were operating more than 8000 different development projects (Knack and Rahman, 2007).²⁵

To illustrate how these issues may affect aid effectiveness we describe some details of the aid composition problem. To this end, it is useful to begin by examining aid disbursements by sector. Table 4 displays the relative distribution of official development assistance (ODA) from all DAC donors in 2000 to 2006. Most observers, unfamiliar with the topic at hand, might be surprised to see that humanitarian aid (including disaster relief) accounts for a relatively modest

²⁴ See Dalgaard (2008) for a discussion of the implications for the growth process.

²⁵ Within aid agency circles the harmonization problem is one of the more popular explanations for low aid impact. This is vividly illustrated by the Rome Declaration on Harmonization (Rome, February 25, 2003) and the Paris Declaration on Aid Effectiveness (Paris, March 2, 2005).

fraction of total disbursements (7 percent in 2006). At the other end of the spectrum we find aid for “social infrastructure”, which encompasses expenditure related to health, education, water supply and sanitation. These kinds of expenditure account for more than 30 percent of the total disbursements. Finally, one may note that initiatives related to debt has increased considerably over the last 5-6 years and is to a considerable extent responsible for the expansion in total aid disbursements over the period.

	2000	2001	2002	2003	2004	2005	2006
Social infrastructure	32	32	34	31	37	31	33
Economic Infrastructure	17	15	12	9	17	11	11
Production sectors	7	9	7	5	6	5	5
Multi sector and cross-cutting initiatives	8	7	7	8	6	6	6
General Programme Assistance	7	7	5	5	3	3	4
Debt	8	10	13	23	11	27	22
Humanitarian Aid	5	5	6	6	7	8	7
Total (\$ million)	44980	42335	49885	69883	74401	96411	98532

TABLE 4. Distribution of Official Development Aid by Selected Sectors (percentages): 2000-2006.

Notes: (1) Disbursements only refer to DAC donors. (2) Dollar amounts in current prices. (3) Percentages do not sum to 100, as a few posts have been suppressed.

Source: Creditor Reporting System (OECD)

Now, imagine that countries who receive large amounts of aid tend to receive larger fractions in the form of aid for social infrastructure. And, by contrast, suppose countries that receive comparatively small amounts of aid get it in the form of infrastructure investments. Assume further, that infrastructure investment generate more growth than social infrastructure, being more directly geared towards economic benefits. In this scenario we have a picture like Figure 19, whereby the regression formulation of sections 4 and 5 would tend underestimate the impact of aid because the composition of aid is an omitted factor.

6.3 The regression solution: more instruments

In principle, the added complexity introduced by allowing the slope of the effectiveness rule to vary systematically with an intervening factor poses no problem, as long as the researcher is able to specify the association between the slope and the intervening factor, just as in the equation above. However, the new variable in the regression model ($w \cdot Aid$) is as problematic as the aid variable in itself because the bi-directional causality is affecting all terms in the model in which the aid variable appears. This means that one needs to find an instrument for the new variable, which satisfies all the requirements of an instrument we discussed in Section 4. We illustrate the procedure by an example.

Following the line of reasoning in Burnside and Dollar (2000) and related studies one may speculate if having good institutions increases the effectiveness of foreign aid. To flesh it out, one could note that ‘quality of the bureaucracy’ is part of good institutions and it is not difficult to imagine how better bureaucracy (and control of corruption) could ease the life of aid managers and, possibly, even increase the number of successful projects. If this holds true there could well be a positive association between institutions and aid effectiveness, as illustrated by Figure 18.

To formulate a regression model we specify ‘*Effect*’ as being related to our measure of institutions

$$Effect = b + f \cdot Institutions$$

And we substitute this into the effectiveness rule

$$Growth = a + b \cdot Aid + f \cdot (Institutions \cdot Aid) + x$$

Where ‘*x*’ contains the conditioning variables we described in Section 5.

We already have an instrument for aid (simply called ‘instrument’) but we need an additional instrument because ‘*Institutions \cdot Aid*’ is affected by the bi-directional causality. We choose the simplest way to obtain an instrument, by mimicking what we do to the aid variable itself. Hence, we construct an instrument by multiplying institutions and the instrument, ‘*Institutions \cdot Instrument*’. This way of constructing instruments is quite common in aid effectiveness studies.²⁶

From this point the steps of analysis are parallel to the Two Stage Least Squares regression in Section 5. We will, therefore, jump straight to the result

$$\begin{array}{rcl}
 Growth & = & 5.36 & +0.04 \cdot Aid & +0.20 \cdot (Institutions \times Aid) \\
 & & (3.01) & (0.06) & (0.62) \\
 & & -3.08 \cdot SSAfrica & +10.1 \cdot Institutions & -1.17 \cdot GDPcap(1970) \\
 & & (0.62) & (1.93) & (0.39)
 \end{array}$$

Comparing this regression to the last regression in Section 5 we find only very small changes in the estimated coefficients on the control variables while the ‘aid effectiveness coefficient’ has

²⁶ But probably also quite problematic, since part of the “full” instrument (i.e., the “institutions” component), obviously, is not exogenous.

changed substantially. The main reason for the change is that it is no longer easy to interpret the coefficient. The best way of thinking of the results is actually to insert the estimated coefficients into the equation for “*Effect*” and then interpret that function

$$\begin{aligned} \textit{Effect} &= 0.04 + 0.20 \cdot \textit{Institutions} \\ &(0.06) \quad (0.62) \end{aligned}$$

From this equation we learn that the slope is 0.2 and, further, that the slope divided by its standard error is much less than 2, whereby we conclude that the association between *Effect* and *Institutions* is not statistically significant (in the current model). Some researchers look at both the slope and the intercept in the Effect-equation and, based on the coefficient-standard-error rule of thumb they conclude that the impact of aid on growth, as such, is zero because none of the coefficients are (individually) significant. This conclusion is not correct, as can be seen from the regression model in Section 5: in the model with a constant slope we find a statistically significant impact of aid on growth.

Although it appears straight forward to extend the analysis of Section 5 to regression models in which the effect of aid varies systematically with an underlying factor there are serious, hidden, problems. Strictly speaking, one should not have too much confidence in the conclusion about institutions and the effectiveness of aid based on the analysis in this section because by adding another variable and another instrument we end up in the situation in which the association between the instruments and aid is weak. In some sense, the problem arises because we are asking too much of the data.

This may be illustrated by looking at some of the cross-plots underlying the estimated coefficients. Starting with the simple cross-plots of growth versus the two aid variables, in Figure 21, it is clear that multiplying “*Aid*” and “*Institutions*” mainly results in a rescaling of the aid variable. The left hand side plot in Figure 21 is the same plot as Figure 5, while the plot on the right hand side has institutions times the aid-GDP ratio on the x-axis. Looking briefly at the labels identifying the individual countries in the plots one notices that there is almost no ‘reshuffling’ of the data points—the two plots, by and large, illustrate the same relationship.




FIGURE 21. The association between the aid-to-GDP ratio and average growth in GDP per capita (PPP) 1970-2000 and between institutions times the aid-to-GDP ratio and average growth in GDP per capita. With regression lines added.

Data Source: Rajan and Subramanian (2008).




FIGURE 22. The association between the filtered aid-to-GDP ratio and average growth in GDP per capita (PPP) 1970-2000, and between the filtered institutions times the aid-to-GDP ratio and average growth in GDP per capita, when geography, institutions and initial GDP per capita is controlled for (Second-stage regression).

Data Source: Rajan and Subramanian (2008).

The two cross-plots in Figure 22 are illustrating the associations between growth and the two aid variables once they have been filtered through the instruments and we have controlled for the effect of geography, institutions and initial GDP per capita (in the same way as illustrated in Figure 17). Hence, the plots show the data for the second-stage regression. Notice the scale of the x-axes in the four cross-plots in Figures 21 and 22. As seen, there is, practically speaking, no variation in the filtered aid-to-GDP ratio and even less when it is multiplied by the variable for institutions.

The lack of variation in the filtered data when we specify the non-linear model of aid effectiveness is the reason why we cannot exclude the possibility of a systematic interaction between institutions and the effectiveness of aid: we cannot, meaningfully, reject an association when there is no variation in the data. Conclusions like this one are surprisingly common in empirical studies of growth across countries. Hence, even though the added complexity introduced by allowing the slope of the effectiveness rule to vary systematically with an intervening factor poses no problem, in principle, it does pose practical problems when data are scarce.

7. Concluding remarks

This Evaluation Study is concerned with evaluation of foreign aid at the aggregate level. As we explain, for practical purposes, evaluation of the aggregate impact of foreign aid has been synonymous with asking if the inflow of foreign aid, to the poor countries, increases the growth rate of income per capita (economic growth). The focus on economic growth is rooted in two causes. First of all, average income (GDP per capita) is a good indicator of important parts of human development—although most people would agree that it does not capture all aspects. Therefore, it makes sense to look at the impact of aid on average income. Second, economists have a long tradition for analyzing economic growth leading to a rich ‘toolbox’ in terms of theories and in terms of data.

The main purpose of the study is to introduce the reader to the statistical problems researchers encounter in their analyses of aid effectiveness at the aggregate level. The statistical problems are intimately linked to empirical problems in the social sciences and they have mainly been dealt with by the special branch of economists known as econometricians. Therefore, the type of problems and their solutions are not taught at introductory statistical courses in high schools or at universities. Further, within natural sciences (including life sciences), health and humanities, these special statistical problems are never taught in any course in statistics. This means that most people working with aid ‘on the ground’ are not familiar with the statistical requirements making them unable to judge if a study of aid effectiveness is tackling the statistical problems in an appropriate way. Frankly, even within the economics profession many lack the necessary skills—leading to unfortunate statements, by economists, about the impact of foreign aid on economic growth.

For years, statistical analyses, using data across many countries have shown that the association between foreign aid and growth of GDP per capita is negative. This has led some to conclude that aid is, at best, useless and, possibly, even harmful. In most cases, however, the studies forming the basis for such claims cannot be said to have evaluated *the causal link* from aid to economic growth.

When the statistical problems are used as a yardstick in judging if research studies (and other scholarly writings) of aid effectiveness are ‘trustworthy’ we end up with less than a dozen published works out of more than one hundred studies written during the past forty-some years. And even among the ‘trustworthy’ studies there is disagreement and ample room for improvement. In the companion paper (Dalgaard and Hansen, 2009) we present and discuss some of these studies.

Literature

Anand, S. and A. Sen, 2000. The income component of the Human Development Index. *Journal of Human Development*, 1, 83-106.

Banerjee, A. V., 2007. *Making Aid Work*. Boston review books.

Barro, R. J., 1991. Economic growth in a cross section of countries. *Quarterly Journal of Economics*, 106, p. 407-43.

Bhagwati, J., 2004. *In Defense of Globalization*. Oxford: Oxford University Press.

Burnside, C. and D. Dollar, 2000. Aid, policies and growth. *American Economic Review*, 90, 847-868.

Block, S., 2001. Does Africa Grow Differently? *Journal of Development Economics*, 65, 443-67.

Dalgaard, C-J., 2008. Donor policy rules and aid effectiveness, *Journal of Economic Dynamics and Control*, 32, 1895-1920.

Dalgaard C-J and H. Hansen, 2001. On Aid, Growth and Good Policies. *Journal of Development Studies*, 37, 17-41.

Dalgaard, C-J, H. Hansen and F. Tarp, 2004. On the Empirics of Aid and Growth. *Economic Journal*, 114, 191-216.

Dalgaard, C-J. and H. Hansen, 2009. Evaluating Aid Effectiveness in the Aggregate: A critical assessment of the evidence. Evaluation Study. The Evaluation Department, Danida (Forthcoming).

Durlauf, S. and P. Johnson, 1995. Multiple Regimes and Cross-Country Growth Behaviour. *Journal of Applied Econometrics*, 10, 365-84.

Easterly, W., 2006. *The White Man's Burden: Why the West's Efforts to Aid the Rest Have Done So Much Ill and So Little Good*. New York: Penguin Press.

Easterly, W., R. Levine and D. Roodman, 2004. Aid, policies and growth: comment. *American Economic Review*, 94, 774-780.

- Griffin, K.B. and J.L. Enos, 1970. Foreign Assistance: Objectives and Consequences, *Economic Development and Cultural Change*, 18, 313-327.
- Hansen, H. and F. Tarp, 2000. Aid Effectiveness Disputed, *Journal of International Development*, 12, 375-398.
- Lipset, S., 1959. Some Social Requisites of Democracy; Economic Development and Political Legitimacy. *American Political Science Review*, 53, 69-105.
- Lucas, R. Jr., 1998. On the Mechanics of Economic Development. *Journal of Monetary Economics*, 22, 3-42.
- Malchow-Møller, N. and A. Würtz, 2003. *Indblik i Statistik—En Grundbog for Videregående Uddannelser*. København: Gyldendal.
- Masanjala W. and C. Papageorgiou, 2008. Rough and lonely road to prosperity: a reexamination of the sources of growth in Africa using Bayesian Model Averaging, *Journal of Applied Econometrics*, (forthcoming).
- Mishkin, F. S., 2006. *The Next Great Globalization. How Disadvantaged Nations Can Harness Their Financial Systems To Get Rich*. Princeton: Princeton University Press.
- Rajan, R. and A. Subramanian, 2008. Aid and Growth: What does the cross-country evidence really show? *Review of Economics and Statistics*, (forthcoming).
- Sachs, J., 2005. *The End of Poverty: Economic Possibilities for Our Time*. New York: Penguin Press.
- Stiglitz, J. E., 2003. *Globalization and Its Discontents*. New York: W.W. Norton & Company.
- Stiglitz, J. E., 2006. *Making Globalization Work. The Next Steps to Global Justice*. New York: Penguin Press.
- Stock, J.H. and M. W. Watson, 2007. *Introduction to Econometrics 2/E*. Addison-Wesley.
- UNDP, 1990. *Human Development Report 1990: Concept and Measurement of human Development*, Published by Oxford University Press for UNDP.
- UNDP, 2007. *Human Development Report 2007/2008: Fighting Climate Change: Human solidarity in a divided world*, Published by Palgrave Macmillan for UNDP.

Wooldridge, J. M, 2008. *Introductory Econometrics*. South Western College.

Qizilbash, M,, 2002. Development, Common Foes and Shared Values, *Review of Political Economy*, 14, 2002

Annexes

The three annexes below present some of the results of the main text using math rather than graphical illustrations. This allows a more formal description of the problems and solutions. More general introductions to the estimation problems can be found in Stock and Watson (2007) or Wooldridge (2008) among many others.

Following standard practice we use Greek letters (α , β , γ and δ) to denote unknown parameters while the letters v and u are used to indicate independent random variables. Further, we focus on ‘large sample results’ in the sense that we illustrate the properties of estimators, not the formulae as they are computed in a software program.

Annex 1. The Cause and Effect Problem

Consider the simple aid effectiveness rule with random disturbances v

$$Growth = \alpha_0 + \alpha_1 Aid + v$$

Where *Growth* is average annual growth in GDP per capita over some relevant time period, *Aid* is the aid-to-GDP ratio averaged over the relevant period, and v is the random disturbance, which has mean zero and variance σ_v^2 .

The parameter α_1 is the quantity the researcher would like to estimate. When he or she estimates the parameters of the regression model using OLS (as in Figure 5 in the text), the parameters are chosen such that the sum of the squared deviations between the fitted line and the data points is as small as possible. This OLS formula will always give a unique solution when we have at least two different data points. The OLS formula for α_1 is equivalent to selecting an estimator with the property that

$$\alpha_1^{OLS} = \alpha_1 + \frac{\text{cov}(Aid, v)}{\text{var}(Aid)}$$

in which $\text{cov}(Aid, v)$ is the covariance between aid and the random disturbance term and $\text{var}(Aid)$ is the variance of aid.

Since we never observe the disturbance term or the true parameter, we never get exactly this result when we apply the OLS formula to real data. But, the key insight from the formula is that

if (and only if) the covariance between Aid and the disturbance term, v , is zero, the procedure provides the “correct” answer to question: what is the impact of aid on growth?

Now, suppose we have the allocation rule by which growth affects aid.

$$Aid = \beta_0 + \beta_1 Growth + u$$

where the slope parameter, β_1 , is negative reflecting how countries that grow faster have lower average aid-to-GDP ratios. The disturbance term in the allocation equation, u , has mean zero and variance σ_u^2 .

Inserting the efficiency equation into the allocation equation, we can solve for Aid and express it as a function of the parameters of the two equations and of the two disturbance terms

$$Aid = \frac{\beta_0 + \alpha_0 \beta_1}{1 - \alpha_1 \beta_1} + \frac{\beta_1 v}{1 - \alpha_1 \beta_1} + \frac{u}{1 - \alpha_1 \beta_1}.$$

Using this expression, we can now find the covariance between aid and the disturbance term, v , and the variance of aid. Hence, we can find an exact expression for the OLS estimator, when we have this simple kind of bi-directional causality.

$$\alpha_1^{OLS} = \alpha_1 \left(1 - \frac{1}{1 + \frac{1}{\beta_1^2} \frac{\sigma_u^2}{\sigma_v^2}} \right) + \frac{1}{\beta_1 + \frac{1}{\beta_1} \frac{\sigma_u^2}{\sigma_v^2}}$$

The two special cases discussed in the main text are easily found from this general expression. When there is no disturbance in the effectiveness equation ($\sigma_v^2 = 0$) we estimate the ‘correct’ slope; β_1 . In addition we also find that when $\alpha_1 = 0$, whereby the aid allocation to the poor countries is randomly distributed around the mean β_0 , we also estimate the correct slope. (This has been the argument for using OLS in aid effectiveness analyses). In contrast, when there is no disturbance in the allocation equation, implying that $\sigma_u^2 = 0$, the estimated slope equals $1/\beta_1$, which is the inverse of the aid allocation slope parameter.

Finally, when the slope of the allocation equation is negative while the slope of the effectiveness rule is non-negative (i.e., positive or zero) the OLS estimator will always be less than the true slope: $\alpha_1^{OLS} < \alpha_1$. This is called the simultaneity bias.

Annex 2. Reverse causality: the impact of growth on aid

From the extensive literature on aid allocation the following equation has found considerable support

$$a_{it} = \delta_0 y_{it}^{\delta_1}$$

Where a_{it} is aid *per capita* in country i , year t while y_{it} is GDP per capita in the same country. The key result from the aid allocation literature, which matters in the present context, is that δ_1 , is negative such that richer countries tend to receive less and per capita, when other factors are equal.

Rearranging the equation above, we can derive the aid-to-GDP ratio for a country in a given year. This shows that the aid-to-GDP ratio is even more sensitive to income differences than aid per capita because GDP per capita is both having a negative impact on aid, and, it is the nominator of the expression:

$$Aid_{it} = \delta_0 y_{it}^{\delta_1 - 1}$$

Next, suppose GDP per capita grows over time, at a constant, country specific rate; $Growth_i$. Then GDP per capita at any given year can be found from the initial income and the growth rate

$$y_{it} = (1 + Growth_i)^t y_{i0}$$

Inserting this equation into the allocation equation above, yields

$$Aid_{it} = \delta_0 (1 + Growth_i)^{(\delta_1 - 1)t} y_{i0}^{\delta_1 - 1}$$

Finally, we need to find the average of the aid variable on the left hand side over some time period, T , which is 30 years in the main text. We obtain the simplest expression by taking the average of the logarithmically transformed equation, which is equivalent to using the geometric average instead of the standard arithmetic average.

The aid allocation rule for the average aid-to-GDP ratio becomes

$$\ln Aid_i = \beta_0 + \beta_1 Growth_i + z_i,$$

where

$$\beta_0 = \ln(\delta_0)$$

$$\beta_1 = (\delta_1 - 1)(T + 1) / 2$$

$$z_i = (\delta_1 - 1) \ln(y_{i0})$$

As β_1 is negative it follows that β_0 is negative; countries where growth is faster will end up with a lower average aid to GDP ratio, as illustrated in Figure 5. Notice also how β_1 increases with the length of the period over which the average is calculated.

Annex 3. Instrumental variables estimation

We reconsider the two equations for aid effectiveness and aid allocation introduced in Annex 1

$$Growth = \alpha_0 + \alpha_1 Aid + v$$

$$Aid = \beta_0 + \beta_1 Growth + z + u$$

There is a small difference; we have added another variable in the allocation equation: z . It is a determinant of aid, which crucially does not affect growth directly, nor is it affected by growth or by aid. As in Annex 1, we can substitute $Growth$ out of the aid allocation equation;

$$Aid = \gamma_0 + \gamma_1 z + w$$

where

$$\gamma_0 = \frac{\beta_0 + \beta_1 \alpha_0}{1 - \alpha_1 \beta_1}, \quad \gamma_1 = \frac{1}{1 - \alpha_1 \beta_1},$$

$$w = \frac{\beta_1 v + u}{1 - \alpha_1 \beta_1}$$

We can now find variation in Aid, which is not caused by growth as this is captured by the variation in z . This means that we can consider z as our instrument.

Instrumental variable estimation in the present model can be illustrated by inserting the filtered aid variable into the aid effectiveness equation. The filtered aid variable is given by

$$Aid^* = \gamma_0 + \gamma_1 z$$

The two unknown parameters in this equation can be estimated by OLS and we find that

$$\gamma_1^{OLS} = \gamma_1 + \frac{\text{cov}(z, w)}{\text{var}(z)} = \gamma_1$$

Where $\text{cov}(z, w)$ is the covariance between z and the composite disturbance term, w , and $\text{var}(z)$ is the variance of z . The latter equality in the equation comes from the requirement that z is not directly causing growth nor caused by growth or aid. This requirement implies that the covariance between z and the disturbance, w , is zero.

Using OLS estimates of the two parameters, β_0 and β_1 , we have observations for the filtered aid data Aid^* . This data, which is independent of variation in growth, is inserted in the aid effectiveness equation to replace the aid variable itself whereby we have a new disturbance term in the equation

$$Growth = \alpha_0 + \alpha_1 Aid^* + (\alpha_1 w + v)$$

The parameters in this equation can be estimated by OLS. The result is that

$$\alpha_1^{IV} = \alpha_1 + \frac{\text{cov}(Aid^*, \alpha_1 w + v)}{\text{var}(Aid^*)} = \alpha_1 + \frac{\gamma_1 \text{cov}(z, \alpha_1 w + v)}{\gamma_1^2 \text{var}(z)} = \alpha_1$$

As we use OLS twice to estimate the slope parameter in the aid effectiveness equation, the procedure is called Two Stage Least Squares (TSLS or 2SLS).

Annex 4. The Data

Country	Growt h	Aid	SSAfric a	Institution s	GDPcap(197 0)	Instrumen tplots	Code in
Algeria	1,183	0,537	0	0,529	8,141	2,244	DZA
Argentina	0,574	0,081	0	0,579	9,134	-0,025	ARG
Bangladesh	1,406	4,925	0	0,313	7,007	-1,172	BGD
Benin	0,346	9,757	1	0,376	6,998	8,467	BEN
Bolivia	0,289	7,171	0	0,381	7,823	4,913	BOL
Botswana	6,363	3,832	1	0,713	7,084	9,444	BWA
Brazil	2,287	0,092	0	0,682	8,194	-7,656	BRA
Burkina Faso	1,193	12,73	1	0,498	6,506	5,793	BFA
Burundi	-1,610	5	1	0,528	6,743	7,264	BDI
Cameroon	0,854	3,879	1	0,563	7,365	5,649	CMR
Chad	-0,872	5	1	0,554	7,073	7,373	TCD
Chile	2,426	0,289	0	0,646	8,475	2,604	CHL
China	5,086	0,409	0	0,641	6,703	-6,181	CHN
Colombia	1,777	0,313	0	0,565	8,058	0,146	COL
Congo, Dem. Rep.	-4,899	3,731	1	0,225	6,962	2,679	ZAR
Congo, Rep.	2,220	7,043	1	0,415	6,834	11,244	COG
Costa Rica	1,131	1,924	0	0,670	8,338	7,429	CRI
Cote d'Ivoire	-0,820	4,344	1	0,626	7,779	5,870	CIV
Cyprus	4,223	0,760	0	0,723	8,571	9,006	CYP
Dominican Republic	3,200	1,257	0	0,510	7,610	4,309	DOM
Ecuador	1,381	1,214	0	0,573	7,737	3,923	ECU
Egypt, Arab Rep.	2,511	3,254	0	0,551	7,586	0,792	EGY
El Salvador	0,229	4,262	0	0,372	8,329	5,361	SLV
Ethiopia	0,145	8,712	1	0,399	6,409	0,488	ETH
Fiji	1,589	3,044	0	0,611	8,141	9,900	FJI
Gabon	0,677	1,902	1	0,574	8,833	14,594	GAB
Gambia, The	0,299	4	1	0,568	7,014	10,156	GMB
Ghana	0,174	3,993	1	0,540	7,156	3,592	GHA
Guatemala	0,896	1,521	0	0,371	8,003	4,292	GTM
GuineaBissau	2,431	8	1	0,340	5,804	20,378	GNB
Guyana	1,365	9,425	0	0,379	7,796	9,203	GUY
Haiti	3,310	9,482	0	0,236	6,835	6,595	HTI
Honduras	0,323	7,279	0	0,424	7,529	6,293	HND

Country	Growt h	Aid	SSAfric a	Institution s	GDPcap(197 0)	Instrumen tplots	Code in
Hungary	2,215	0,533	0	0,788	8,589	2,542	HUN
India	2,791	0,765	0	0,591	6,978	-5,739	IND
Indonesia	4,029	1,178	0	0,484	6,992	-2,008	IDN
Iran, Islamic Rep.	0,458	0,121	0	0,486	8,561	0,523	IRN
Israel	2,172	2,413	0	0,756	9,087	5,981	ISR
Jamaica	-0,153	2,771	0	0,544	8,260	7,025	JAM
Kenya	1,386	7,221	1	0,582	6,711	2,943	KEN
Korea, Rep.	5,886	0,090	0	0,735	7,907	0,809	KOR
Lesotho	1,964	15,66	1	0,661	6,783	8,227	LSO
Madagascar	-1,405	8,948	1	0,476	7,150	5,074	MDG
Malawi	1,812	19,53	1	0,503	6,121	5,042	MWI
Malaysia	4,118	7	0	0,687	7,967	2,555	MYS
Mali	0,708	15,35	1	0,311	6,664	5,991	MLI
Mauritania	-1,235	17,58	1	0,406	7,540	11,091	MRT
Mauritius	4,155	3	1	0,704	8,295	8,860	MUS
Mexico	1,539	2,229	0	0,592	8,616	-2,148	MEX
Morocco	1,657	0,083	0	0,563	7,723	2,147	MAR
Namibia	-0,233	1,851	1	0,462	8,470	8,422	NAM
Nicaragua	-2,706	5,412	0	0,523	8,289	5,774	NIC
Niger	-1,837	16,79	1	0,514	7,326	6,903	NER
Nigeria	-1,514	13,43	1	0,428	7,015	-0,500	NGA
Pakistan	2,520	2	0	0,453	6,849	-0,793	PAK
Panama	1,538	0,420	0	0,410	8,249	7,835	PAN
Papua New Guinea	0,072	10,28	0	0,625	7,959	7,971	PNG
Paraguay	1,628	3	0	0,486	7,964	6,571	PRY
Peru	-0,070	1,381	0	0,438	8,452	1,670	PER
Philippines	1,191	1,001	0	0,407	7,782	3,253	PHL
Romania	2,448	1,519	0	0,516	7,628	1,200	ROM
Rwanda	0,032	15,49	1	0,387	6,787	7,553	RWA
Senegal	-0,011	3	1	0,487	7,394	6,900	SEN
Sierra Leone	-1,866	11,16	1	0,398	7,311	6,236	SLE
Singapore	5,972	1	0	0,859	8,571	6,790	SGP
South Africa	0,307	9,745	1	0,740	8,836	1,476	ZAF
		0,079	0				
		0,430	1				

Country	Growt h	Aid	SSAfric a	Institution s	GDPcap(197 0)	Instrumen tplots	Code in tplots
Sri Lanka	2,503	5,288	0	0,463	7,351	2,273	LKA
Syrian Arab Republic	3,038	0,949	0	0,491	7,406	5,114	SYR
Thailand	4,417	0,711	0	0,711	7,508	0,584	THA
Togo	-1,578	10,55 9	1	0,446	7,242	9,543	TGO
Trinidad & Tobago	1,765	0,255	0	0,616	8,792	8,497	TTO
Tunisia	3,234	2,000	0	0,541	7,851	5,840	TUN
Turkey	2,119	0,339	0	0,601	8,194	1,018	TUR
Uganda	1,458	6,580	1	0,368	6,409	3,299	UGA
Uruguay	1,502	0,398	0	0,564	8,721	5,216	URY
Venezuela, RB	-1,648	0,072 13,19	0	0,612	9,262	2,260	VEN
Zambia	-1,346	0	1	0,424	7,197	5,212	ZMB
Zimbabwe	0,477	4,711	1	0,545	7,675	4,422	ZWE

Source: Rajan and Subramanian (2008).

Note: Numbers in the table are rounded compared to the data used in the regression analyses.

Published by:
Ministry of Foreign Affairs of Denmark
Evaluation Department
Asiatisk plads 2
1448 Copenhagen K
Denmark
E-mail: eval@um.dk

The publication can be downloaded from:
www.evaluation.dk

ISBN: 978-87-7087-129-7
ISBN: 978-87-7087-130-3 (internet version)

Other Evaluation Studies:

- 2007/1: Assistance to Development Research
- 2007/2: Synthesis of Evaluations on Technical Assistance
- 2007/3: Evaluating the Impact of Rural Roads in Nicaragua
- 2008/1: Synthesis of Evaluations of HIV/AIDS Assistance
- 2008/2: Synergies between Bilateral and Multilateral Activities
- 2008/3: Assessing Multilateral Organisation Effectiveness
- 2008/4: Economic Empowerment of Women
- 2008/5: Public Private Partnership Programme