

Aid and Growth

New evidence shows that aid flows aimed at growth have produced results

Steven Radelet, Michael Clemens, and Rikhil Bhavnani

CONTROVERSIES about aid effectiveness go back decades. Critics such as Milton Friedman, Peter Bauer, and William Easterly have leveled stinging critiques, charging that aid has enlarged government bureaucracies, perpetuated bad governments, enriched the elite in poor countries, or just been wasted. They cite widespread poverty in Africa and South Asia despite three decades of aid, and point to countries that have received substantial aid yet have had disastrous records—such as the Democratic Republic of the Congo, Haiti, Papua New Guinea, and Somalia. In

their eyes, aid programs should be dramatically reformed, substantially curtailed, or eliminated altogether.

Supporters counter that these arguments, while partially correct, are overstated. Jeffrey Sachs, Joseph Stiglitz, Nicholas Stern, and others have argued that, although aid has sometimes failed, it has supported poverty reduction and growth in some countries and prevented worse performance in others. They believe that many of the weaknesses of aid have more to do with donors than recipients, especially since much aid is given to political allies rather than to support devel-

opment. They point to a range of successful countries that have received significant aid such as Botswana, Indonesia, Korea, and, more recently, Tanzania and Mozambique, along with successful initiatives such as the Green Revolution, the campaign against river blindness, and the introduction of oral rehydration therapy. In the 40 years since aid became widespread, they say, poverty indicators have fallen in many countries worldwide, and health and education indicators have risen faster than during any other 40-year period in human history.

Throughout this debate, however, most analysts have missed a critical point by treating all aid as if it were alike in its impact on growth. In a recent Center for

Students work at a computer in Abuja, Nigeria.





Global Development study, we try to rectify this gap by exploring the impact on growth of aid flows that actually are aimed at growth.

Three prevailing views on aid

Over the past three decades, three broad views have emerged on the relationship between aid and growth:

Aid has no effect on growth, and may actually undermine growth. There are several reasons why aid might not support growth. It can be wasted on frivolous expenses such as limousines or presidential palaces, or it can encourage corruption. It can undermine incentives for private sector production, including by causing the currency to appreciate, which weakens the profitability of tradable goods production (an effect known as “Dutch disease”). Similarly, food aid, if not managed appropriately, can reduce farm prices and hurt farmer income. Aid flows potentially can undermine incentives for both private and government saving. They can also sustain bad governments in power, helping to perpetuate poor economic policies and postpone reform.

This view has been supported by a range of empirical studies, mostly published from the early 1970s through the mid-1990s. While these studies have been influential, many are of questionable quality, especially using today’s research standards. For example, most assume only a simple linear relationship between aid and growth in which each new dollar of aid has exactly the same impact on growth as the first (eliminating the possibility of diminishing returns) and ignore possible endogeneity (in which faster growth might attract higher aid, or both might be caused by something else), among other issues. A recent paper by Raghuram Rajan and Arvind Subramanian (2005), which also assumes a simple linear relationship for most of its results, stands in sharp contrast to the bulk of recent research on the issue, as discussed below.

Aid has a positive relationship with growth on average (although not in every country), but with diminishing returns. Aid could support growth by financing investment or by increasing worker productivity (for example, through investments in health or education). It can bring new technology or knowledge, either imbedded in capital goods imports or through technical assistance. Several early studies found a positive relationship between aid and growth, but this strand of the literature took a significant turn in the mid-1990s when researchers began to investigate whether aid might spur growth with diminishing returns—that is, that the impact of additional aid would decline as aid amounts grew. Oddly, since economic theory and research had recognized the importance of diminishing returns on investment since the 1950s, research on aid and growth until the mid-1990s tested only a linear relationship, a specification that (surprisingly) persists in some studies even today.

Although they have received comparatively less popular attention, most of these studies (some published in top peer-reviewed journals) have found a strong aid-growth relationship, including research by Michael Hadjimichael and

colleagues at the IMF in the mid-1990s, and the work of Carl-Johan Dalgaard, Henrik Hansen, Finn Tarp, Robert Lensink, Howard White, and others between 1999 and 2005. These studies typically do not conclude that aid has always worked, but rather that, on average, higher aid flows have been associated with more rapid growth.

Aid has a conditional relationship with growth, helping to accelerate growth only under certain circumstances. The “conditional” view usually argues that aid effectiveness hinges on either recipient characteristics or donor practices.

Recipient country characteristics. World Bank researchers Jonathon Isham, Daniel Kaufmann, and Lant Pritchett opened this line of enquiry in 1995 by finding that World Bank projects had higher rates of returns in countries with stronger civil liberties. Craig Burnside and David Dollar followed with their influential study that concluded that aid stimulated growth in countries with good policies, but not otherwise. Others have proposed different characteristics that might affect the aid-growth relationship, including vulnerability to trade shocks, climate, institutional quality, political conflict, and geography. The statistical results of these studies tend to be fragile, however, and subsequent research has questioned some of the results.

“Although the impact of aid on growth diminishes as aid increases, in countries with stronger institutions or better health, more aid can be absorbed effectively.”

Nevertheless, the view that aid works best (or in a stronger version, aid works only) in countries with good policies and institutions has become the conventional wisdom among donors, partly based on this research and partly due to development practitioners’ own experiences. The appeal of this approach is that it can explain why aid seems to have supported growth in some countries but not others. This reasoning has had an enormous impact on donors, especially the multilateral development banks, and is the foundation of the U.S. Millennium Challenge Account (Radelet, 2003).

Donor practices. Multilateral aid might be more effective than bilateral aid, and untied aid is thought to have higher returns than aid tied to purchases in the donor country. Donors with large bureaucracies, heavy reporting requirements, or ineffective monitoring and evaluation systems probably undermine the effectiveness of their own programs. Two influential and overlapping views argue that aid is more effective when donors allow for greater “country ownership” or broader “participation” in setting priorities and designing programs (country ownership allows for the recipient country to have a stronger say in these decisions;

broader participation allows civil society and faith-based and nongovernmental organizations to have a voice alongside the government in these choices). These issues have been regularly debated and have begun to change donor practices, but have been subject to little systematic research.

The type of aid matters

New research has taken a different tack by exploring the idea that not all aid is alike in its impact on growth. This view suggests that most research on aid and growth is flawed regarding both substance and timing. On substance, almost all studies look at the relationship between total aid and growth, even though large portions of aid are not primarily directed at growth. For example, food and humanitarian aid are aimed primarily at supporting consumption, not growth, as is the provision of medicines, bed nets, and school books. And aid to support democracy or judicial reform is not primarily aimed at stimulating growth. These important aid-financed activities help improve recipient welfare by supporting basic consumption needs, developing political institutions, and strengthening health and education—but they are likely to affect growth only indirectly, if at all. By contrast, aid to build roads, bridges, or telecommunications facilities or to support agriculture and industry is more directly aimed at production and should be expected to accelerate growth. Given the range of likely impacts of different kinds of aid, it is not surprising that some research on aid and growth has shown a weak relationship.

On timing, most cross-country growth research uses panel data, with each observation (usually) corresponding to four years, but then investigates aid flows that cannot possibly affect growth in that period. Aid to support education and health, for example, may stimulate growth, but the impact is likely to take decades. One option for researchers is to use a longer time period, but there is a trade-off: the longer the time period, the harder it is to isolate the impact of aid (or other variables) on growth from other influences. Only a few studies have explored this idea, and most focus on specific countries. For example, one study found that household welfare in Zimbabwe was increased by “development aid” (such as infrastructure and agricultural extension) far more than by “humanitarian aid” (such as food aid and emergency transfers).

To remedy this weakness, our recent research focuses on the type of aid that is directed primarily at growth (Clemens, Radelet, and Bhavnani, 2004). We examined aid flows to 67 countries between 1974 and 2001 and divided aid into three categories:

(1) *Aid for disasters, emergencies, and humanitarian relief efforts, including food aid.* Here we find a negative simple relationship, since disasters simultaneously cause growth to fall and aid to increase. The recent tsunami undermined growth in Sri Lanka, and donors responded with more aid. In a simplistic growth regression, cases like this would show up as high aid with low or negative growth, making it appear that aid had a poor relationship with growth, an obviously misleading result.

Method matters

In conducting our empirical research, we focused on three issues: (1) ensuring the basic model was consistent with theory and evidence; (2) controlling for the possible two-way relationship between aid and growth; and (3) testing the results with a broad yet reasonable set of robustness checks.

At the core of our model is a nonlinear relationship between “early impact” aid and growth that allows for diminishing returns: each additional dollar of aid has a smaller impact on growth than the last. This specification is consistent with both theory and extensive evidence, but it is often overlooked in the literature, giving rise to the weak relationships found in some studies. We then control for a wide variety of other factors that influence growth, including income level, institutional quality, trade policy, inflation, budget deficits, life expectancy, location in the tropics, and the incidence of civil war. Our results show that each of these variables is strongly related to growth, with the exceptions of initial income and budget deficits.

A positive relationship between aid and growth does not prove causality. More aid could cause higher growth, but faster growth could attract more aid, or both could be caused by something else (such as a change in government). To control for potential endogeneity, we estimate the relationship using instrumental variables, using as instruments geopolitical variables and past aid flows that have been used in previous peer-reviewed journal articles. But since no instrumentation strategy is perfect, we also estimate the model using ordinary least squares with aid lagged one four-year period, and find essentially the same results.

We further test the robustness of the results by examining differences rather than levels, eliminating outlier observations, estimating the model using the generalized method of moments procedure, controlling for more or fewer variables, and examining alternative definitions for key variables, among other tests. The results hold firm across this array of tests, giving us confidence in the robustness of the results.

(2) *Aid that might affect growth, but indirectly and over a long period of time.* No one would expect aid aimed at environmental conservation or democratic reform to affect economic growth quickly, and certainly not over a four-year period. Similarly, aid to strengthen health and education is likely to affect labor productivity over many years, but not immediately (with some exceptions). In a standard cross-country growth regression, these observations are likely to appear as high aid and zero or very little growth, again weakening the results. As expected, we detect only a weak positive association between this “late-impact” aid and growth.

(3) *Aid aimed more directly to support growth relatively quickly.* Aid to build infrastructure—roads, irrigation sys-



tems, electricity generators, and ports—should affect growth rates fairly quickly. So should aid to directly support productive sectors, such as agriculture, industry, trade, and services. Aid that comes as cash, such as budget or balance of payments support, could be spent on a wide variety of activities, but to be conservative we assume it is directed at growth (to the extent it is not, our assumption would only weaken our results). For this “early impact” aid (which accounts for about half of all aid), it is perfectly reasonable for policymakers to expect, and for researchers to test for, a positive relationship with growth over a four-year period.

Early impact aid boosts growth

Our research shows a strong, positive, and causal effect of early impact aid on economic growth. The results exhibit diminishing returns, with larger amounts of aid having a progressively smaller impact. The estimated impact is nearly triple the magnitude found in other studies. We test the results over a very wide set of specifications and estimation techniques that control for other influences on growth, possible endogeneity, lags, and other factors. Throughout, the results remain strong and robust. We estimate the model over a four-year period, following the standard used in many studies, but we show (using lags) that the impact carries into a subsequent four-year period. We find no evidence that the effect is a short-run phenomenon that is later reversed. The results do not imply that aid has worked everywhere—it most definitely has not—but that, on average, growth-oriented aid has had a positive and significant impact on growth. The results underscore that the impact of early impact aid differs significantly from other types of aid (see chart).

How great is the effect of early impact aid on growth? Consider the mean observation, where early impact aid is 2.7 percent of GDP (roughly equivalent to where total aid is about 5.4 percent of GDP). Using our most conservative results, at the mean, a 1 percentage point of GDP increase in early impact aid produces an additional 0.31 percentage point of annual growth over the four-year period. With plausible assumptions about discount and depreciation rates (summing to 35 percent), we calculate that each \$1 in early impact aid yields \$1.64 in increased income in the recipient country in net present value terms. This country-level return roughly corresponds to a project-level rate of return of around 13 percent. For sub-Saharan Africa, we find that higher-than-average early impact aid raised per capita growth rates by about 1 percentage point over the growth that

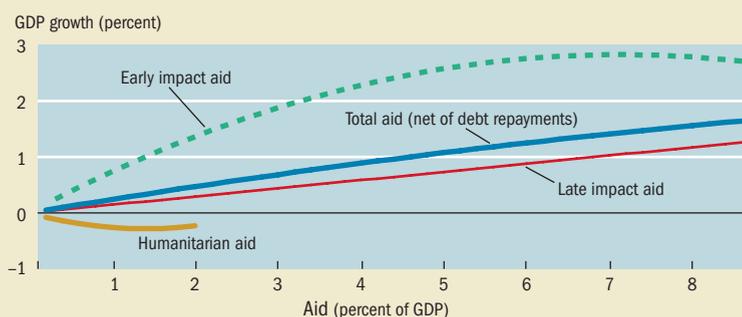
would have been achieved by average aid flows. This suggests that, while growth in sub-Saharan Africa has been disappointing, it would have been worse in the absence of this kind of aid.

What about the claim that aid works best in countries with good policies and institutions? To explore this idea, we looked at one of the most commonly used measures of institutional quality, drawn from the International Country Risk Guide. This index, which has been shown to be strongly correlated with growth, includes measures of the extent of corruption, rule of law, risk of expropriation or repudiation of contracts, and bureaucratic quality. We find some evidence that in countries with better institutions, the relationship between early impact aid and growth is stronger than otherwise. In addition, in countries with higher life expectancy (that is, better health), the aid–growth relationship is stronger than otherwise. But unlike other studies, we do not find that aid works only in countries with strong institutions or better health, and our results do not hinge on this interaction.

Are there limits on how much early impact aid typical recipients can absorb? The answer appears to be yes, but the maximum growth rate occurs on average when early impact aid represents 8–9 percent of GDP, more than three times the

Not all aid is alike

Some types of aid have a much bigger impact on growth than others—possibly why some studies examining aggregate aid find weak results.



Humanitarian aid	Early impact aid	Late impact aid
Disaster relief	Transport and storage	Government and civil society
Emergency aid	Communications	General environmental protection
Humanitarian relief	Energy generation and supply	Women in development
Food aid	Most banking and financial services	Health
	Business and other services	Education
	Agriculture, forestry, and fishing	Populations policies
	Industry, mineral resources, and mining	Water supply and sanitation
	Construction	Policy and administrative management
	Structural adjustment assistance	Support to nongovernmental organizations
	Budget support	Other social infrastructure and services
	Debt relief	

Source: Clemens (2004).

Note: All three curves are estimated using a similar model and include a nonlinear relationship between aid and growth (which is hard to detect visually in the curve for late impact aid). The curve for humanitarian aid is cut off at 2 percent of GDP because there are no data in our sample beyond this point, and to show an upward curve would be misleading. Although only the coefficient on early impact aid is statistically significantly different from zero, the weaker relationships for late impact and humanitarian aid do not necessarily mean these flows have no impact on growth, but rather that a different modeling technique is required to explore these relationships (which we leave for future research).

typical amount. As a rule of thumb, since early impact aid is slightly more than half of total aid on average, this implies that the maximum growth rate occurs when total aid reaches around 16–18 percent of GDP in the typical country. This does not mean that in any particular country, aid flows greater than this amount are necessarily a bad idea. Instead, this represents the typical pattern over the last 30 years—some countries can absorb more, and others less. Moreover,

“Policy discussions should not focus exclusively on determining the limits of aid on growth—but rather on how those limits can be expanded, and how aid can be made even more effective in supporting growth and development.”

we find that absorptive capacity depends to some extent on the quality of institutions and general health of the population. In countries with stronger institutions and higher life expectancy, the impact of early impact aid is stronger throughout, and more aid can be absorbed before reaching the maximum growth rate.

The results also suggest that aid is not fully fungible, at least in the sense that all aid is interchangeable. If this were true, different subcategories would show similar relationships with growth. Instead, we find that aid flows intended for different purposes have significantly different relationships with growth. It is more likely that aid is only partially fungible, not fully so, in accordance with several recent studies.

Going forward

The intense pessimism on aid effectiveness expressed by some analysts appears to be too strong: we find a positive, causal relationship between growth-oriented aid and growth. At the same time, no one should conclude that aid has always worked or that it cannot work better. There are many countries that have received substantial aid and have stagnated or worse, and much aid has been wasted, stolen, or otherwise used to support countries with poor governance. The evidence suggests, however, that on average aid that has been aimed at growth in fact has boosted growth.

Those who argue that aid works *only* in countries with good institutions overstate their case. It would be more accurate to say that aid works *better* in countries with strong institutions, but at times can be effective in other situations. Aid has helped support growth in Mozambique and Uganda over the past decade, even though policies and institutions were far from ideal, and aid has played an important role in stabilizing Sierra Leone since its cease-fire. Aid helped to support sustained growth and poverty reduction in

Indonesia during the Suharto regime—even in the 1970s and 1980s when institutions were weak, corruption was problematic, and policies were less than ideal.

We hasten to add that the weak relationship between late impact and humanitarian aid and growth over a four-year period should not be interpreted to mean that they are ineffective. Different modeling techniques are required to examine those questions, which we are exploring in subsequent research. Although (surprisingly) there is no systematic cross-country research on the relationship between health-oriented aid and health, there is evidence that at least some aid for health has been effective. For example, aid played an important role in supporting several large-scale successful health interventions, such as eradicating smallpox, significantly reducing the prevalence of polio and river blindness, and reducing the incidence of diarrheal diseases (Levine and others, 2004).

Finally, the evidence suggests that absorptive capacity constraints are real, but should not be seen as an immutable barrier to growth. Although the impact of aid on growth diminishes as aid increases, in countries with stronger institutions or better health, more aid can be absorbed effectively. This finding suggests that efforts to strengthen institutions and build human capital can increase returns to aid and help countries effectively absorb larger amounts of aid. Thus, policy discussions should not focus exclusively on determining the limits of aid on growth—but rather on how those limits can be expanded, and how aid can be made even more effective in supporting growth and development. ■

Steven Radelet is a Senior Fellow, Michael Clemens is a Research Fellow, and Rikhil Bhavnani is a Research Assistant at the Center for Global Development in Washington D.C.

References:

- Burnside, Craig, and David Dollar, 2000, “Aid, Policies, and Growth,” *American Economic Review*, Vol. 90. (September), pp. 847–68.
- Clemens, Michael, Steven Radelet, and Rikhil Bhavnani, 2004, “Counting Chickens When They Hatch: the Short-Term Effect of Aid on Growth,” *Center for Global Development Working Paper 44* (Washington: Center for Global Development).
- Isham, Jonathan, Daniel Kaufmann, and Lant Pritchett, 1995, “Governance and Returns on Investment: An Empirical Investigation,” *World Bank Policy Research Working Paper 1550* (Washington: World Bank).
- Levine, Ruth, and the “What Works” Working Group (with Molly Kinder), 2004, *Millions Saved: Proven Success in Global Health* (Washington: Center for Global Development).
- Steven Radelet, 2003, *Challenging Foreign Aid: A Policymaker’s Guide to the Millennium Challenge Account* (Washington: Center for Global Development).
- Rajan, Raghuram and Arvind Subramanian, 2005, “Aid and Growth: What Does the Cross-Country Evidence Really Show?” *IMF Working Paper 05/127* (Washington: International Monetary Fund).