



Department of Economics and Finance
Gordon S. Lang School of Business and Economics
University of Guelph

Discussion Paper 2018-01
(revised September 20, 2019)

Why Aid-to-GDP Ratios?

By:

Kurt Annen
University of Guelph
kannen@uoguelph.ca

Stephen Kosempel
University of Guelph
kosempel@uoguelph.ca



IMPROVE LIFE.

Department of Economics and Finance

Gordon S. Lang School of Business and Economics | University of Guelph
50 Stone Road East | Guelph, Ontario, Canada | N1G 2W1
www.uoguelph.ca/economics

Why Aid-to-GDP Ratios?

Abstract

Virtually all aid-growth regression studies normalize aid by dividing it by GDP. This paper questions the usefulness of this practice: First, there are no clear theoretical reasons for this normalization unless one assumes that donors allocate aid-to-GDP ratios. Second, using aid-to-GDP ratios introduces econometric problems that most likely introduce a downward bias for the aid-growth relationship. We illustrate this point by running simulations from a neo-classical growth model in which aid does not affect growth by construction but find strong negative and in some cases also positive correlations when using aid-to-GDP ratios. When we bring our analysis to the data focusing on the last 20 years of development aid practice, we find a robust positive and statistically significant relationship between aid and growth when using total aid, and we find a mostly negative and insignificant relationship when using the aid-to-GDP ratio.

Keywords: Aid Effectiveness, aid and growth, growth convergence

JEL Codes: O10, O19

1 Introduction

Virtually all aid-growth regression studies normalize aid by dividing it by GDP. We surveyed 48 papers with an aid-growth regression focusing on papers that appeared after 2000 and found that 92% of these papers use aid-to-GDP ratios as their aid measure.¹ This practice started with the very first aid-growth regression we are aware of. Papanek (1973) runs this regression with aid expressed as a percentage of GDP, however, his normalization choice is not discussed.² Papanek's paper was written in the spirit of the Financing-Gap theory based on the Harrod-Domar growth model, where a target growth rate can be achieved by meeting a target savings rate. Foreign aid, then, was seen as an instrument to prop up low domestic savings rates in poor countries. Since these rates are expressed as a percent of GDP, one may then be inclined to do the same for aid. Since then virtually all papers in that literature followed Papanek's lead.

Consistency with the existing literature is one reason to use the aid-to-GDP ratio in a growth regression, but there may be other reasons for this practice. For example, one may argue that a million dollars of aid may have a larger impact in a small economy as compared to a large economy. Aid-to-GDP ratios may better control for the size of an economy. However, this will depend on the underlying model one uses. For example, if aid finances an infrastructure project such as a bridge, then the impact of this bridge on an economy may depend more on how many trucks can cross the bridge in a given amount of time and less on whether the trucks crossing that bridge carry products with a high value added (high GDP) or low value added (low GDP). In this case, aid per capita may be the more appropriate normalization. Or, if aid finances pure public goods such as a knowledge transfer or state capacity building, then the impact may not depend on population size nor on the size of the economy altogether. In this case, total aid may be the appropriate normalization. But

¹The bulk of the remainder of the papers uses aid per capita.

²There are earlier regression studies, but they lumped aid together with other types of foreign resource flows. For a discussion of this literature see Papanek (1973) and Mosley (1980). Thus, to our knowledge Papanek (1973) is the first paper to run a proper growth-aid regression.

from an empirical perspective it is obvious that multivariate regression analysis allows the introduction of proper controls for such economy or population size effects without the need of melting such controls with the aid measure.

In this paper, we question the appropriateness of the practice to use aid-to-GDP ratios. There are no clear theoretical reasons to do so unless one assumes that donors allocate aid-to-GDP ratios instead of just aid as proposed by the Financing-Gap theory. We believe it is hard to justify such a theory. More importantly, to use aid-to-GDP ratios introduces econometric problems that most likely create a downward bias for the aid-growth relationship. We call it the *aid normalization bias* (ANB). We demonstrate this in several ways: First, based on existing econometric theory we show how such a bias can occur in an aid-growth regression. Here, we follow the work of Kronmal (1993), but adjust his presentation to better fit an aid-effectiveness study. Second, we illustrate the ANB by running simulations using a neoclassical growth model in which aid does not affect growth by construction but find strong negative and in some cases also positive correlations when using the aid-to-GDP ratio.³ Third, in the empirical part of the paper we present regression results using aid and per capita GDP growth data between 1995 and 2014 that support our conjectures from the simulation experiments. The empirical analysis focuses on years after the 1980s as this includes the post cold war era only where one would conjecture that aid has become less politicized. In addition, there is evidence that donor aid allocations during this time period have become more policy- and poverty-selective (Dollar and Levine, 2006; Annen and Knack, 2018). The results show a positive and statistically significant correlation between aid and growth when we use total aid as our independent variable, but we find no correlation or a significant negative correlation when using aid-to-GDP ratios. Furthermore, we find that our estimates are fairly robust across specifications and samples when using total aid, but change quite substantially in terms of the sign and significance of the aid coefficient when

³This section of the paper is related to work by Carter (2017) in that we use artificial data generated from a neoclassical growth model to test the performance of an aid-growth regression. The analysis in Carter (2017) focused on the distinction between transitory and long-run effects of aid, whereas the analysis in the present paper is focused on identifying problems related to the normalization of aid.

using aid-to-GDP ratios. Our analysis suggests that aid effectiveness studies should not use aid-to-GDP ratios.

The public debate on the usefulness of foreign aid is contentious with academics and other stake holders deeply divided on the issue. It is clear that aid-growth regressions contribute to this debate. For example, Easterly (2003) gives an insightful description of the tremendous impact the study by Burnside and Dollar (2000) had on policy makers and development practitioners. Similarly, Swanson (2015) in her Washington Post blog about the Nobel-winning economist Angus Deaton entitled “Why trying to help poor countries might actually hurt them” presents a scatter plot taken from Rajan and Subramanian (2008) that shows a statistically significant negative correlation between the aid-to-GDP ratio and income per capita growth. Why is the aid-to-GDP ratio used in such a plot? Our paper shows that such plots may be misleading. In fact, when replicating this plot with total aid, the slope coefficient remains negative but it is no longer statistically significant. Given the weight such regressions receive in policy debates, an analysis of the effects that the aid normalization choice may have on the estimated regression coefficients is important.

The remainder of the paper is organized as follows: Section 2 provides a review of the relevant literature. Section 3 presents simulation experiments of how aid-to-GDP ratios introduce a bias in aid-growth regressions. Section 4 presents aid-growth regressions focusing on the 20 years after 1995. Here, we want to test whether the conjectures derived from our simulation experiments show up in the data. Section 5 concludes by discussing the implications of our analysis for aid effectiveness studies and the public debate on the usefulness of foreign aid.

2 Literature Review

Although a large majority of aid effectiveness studies use the aid-to-GDP ratio, there are only a few papers that discuss the aid normalization choice. In fact, Easterly (2003) com-

ments on a lack of theoretical models on the aid growth relationship which help to pin down the specification needed for empirical analysis. In the current paper we contribute to the literature by demonstrating that the aid normalization choice matters; and, more importantly, we explain why.

The theoretical support for at least most of the early empirical work on aid effectiveness is provided by the Solow (1956) growth model. A presentation of the Solow model with aid is provided by Dalgaard and Erickson (2006), and in their study the aid-to-GDP ratio is treated as a fixed parameter. Although no explanation is provided for this restriction, it was probably imposed for technical convenience. If the aid-to-GDP ratio is fixed in the Solow model, then aid will have an effect on the level of income in the long-run, but not on the speed of convergence. In addition, with a fixed aid-to-GDP ratio a closed-form analytical solution exists for the transitional dynamics in the model. In Section 3 of our analysis we follow Dalgaard and Erickson (2006) by calibrating a version of the Solow model with aid, but we relax their restriction on having a fixed aid-to-GDP ratio. The goal of our calibration and simulation exercise is to test the performance of an aid-growth regression, whereas the model in Dalgaard and Erickson (2006) is calibrated to quantify the impact of aid on growth. In fact, in our simulations the model is constructed so that aid will have no impact on growth, nonetheless when an econometric analysis is performed on the artificial data it reveals a statistical relationship (or bias) between aid and growth when aid is measured as a fraction of GDP.

A critique of the Solow approach applied to study aid effectiveness is that savings behaviour is exogenously given, and therefore unaffected by inflows of foreign aid. In response to this critique, theoretical work on aid effectiveness has been extended to models that allow savings to be endogenous. For example, Dalgaard and Erickson (2006) and Annen and Kosempel (2009) quantify the impact of foreign aid within the Ramsey-Cass-Koopmans (RCK) model; and Dalgaard, Hansen, and Tarp (2004) and Dalgaard and Erickson (2006) study aid flows within the framework of a two period Diamond model. In all of these studies

foreign aid enters the theoretical framework in per capita units. There are probably two reasons for this. First, no matter the restrictions imposed on aid flows, these models tend not to have closed-form solutions. Therefore, assuming a constant aid-to-GDP ratio will not ease the burden of the calculations, as it did in the Solow model. Second, for these models to be structurally accurate it is important to correctly model the aid allocation mechanism that donors use. In these models agents solve dynamic optimization problems, and aid allocations both present and future will impact their consumption and savings decisions. If aid enters as a fraction of GDP, then this will reveal information about future aid levels that will be factored into the agents' decisions. For example, if development occurs rapidly and donors target aid-to-GDP ratios (as in the Financing-Gap theory), then aid also must grow rapidly. We believe it is hard to justify such a theory. Accordingly, these theoretical models use aid-per-capita instead of aid-to-GDP.

There are also a few papers that discuss the aid normalization choice on empirical grounds. For example, Fielding and Knowles (2011) find that the sign and significance of coefficients on foreign aid variables in crosscountry panel growth regressions are very sensitive to the way that aid is measured. Including terms in aid per capita instead of aid as a fraction of GDP completely changes the results. However, they do not argue in favour of a particular way of measuring aid as we do.

To our knowledge, Arndt, Jones, and Tarp (2015, p. 9) is the only empirical paper that provides an argument for using the aid-to-GDP ratio in aid effectiveness regressions as “[r]aw values are not informative due to differences in income and population between countries.” As argued earlier whether total aid values are informative or not will depend on the exact mechanism through which aid affects output in an economy. If aid finances a public good, then the use of total aid will be appropriate. In addition, if scaling matters — i.e. there are heterogeneous treatment effects — say — of an extra million dollar of total aid related to differences in income and population among countries — then one way to econometrically deal with such scale effects is to use an interaction term in addition to having total aid and

the scale variable separately in a regression.⁴ There is no need to melt variables that control for scale with the aid measure. Arndt, Jones, and Tarp (2015) also argue for using the aid-to-GDP ratio as other values make the interpretation of regression results more difficult, given that many macro measures are expressed as a fraction of GDP and given that the real cost for the provision of public goods increases with GDP. Convenience in terms of the interpretation of regression results is useful, but the main point of our paper is to show that this convenience comes at a cost: Using aid-to-GDP ratios most likely introduces a downward bias for the aid-growth relationship.

In an earlier study, Arndt, Jones, and Tarp (2010) use aid-per-capita as the dependent variable in the first stage of their IV analysis, but use aid-to-GDP in the second stage of the analysis. They claim that aid per capita should be used in the first stage “. . . as the dependent variable which accords closely with the explicit aid allocation rules used by donors, such as the World Bank”.⁵ It is not entirely clear to us, however, whether this characterization is accurate as the IDA’s allocation mechanism is performance based and allocations are typically not strictly related to population other than that aid caps are expressed in per capita terms (see IDA15, 2008; Annen and Knack, 2019). From that perspective, we believe that using total aid is plausible as aid allocation rules by donors typically are not strictly related to population nor to GDP.

In order to avoid the ANB, we propose to use the log of total aid instead of aid-to-GDP. Note that Juselius, Møller, and Tarp (2014) also point to the possibility that the use of ratios can “significantly influence the results unless the implied parameter restriction is empirically valid.” They show that for the sample of Sub-Saharan countries such a restriction is not valid. Accordingly, they use the natural log of total aid in their subsequent analysis. In our analysis we use the natural log of total aid, where we convert aid amounts into PPP adjusted

⁴For example, Kronmal (1993) proposes to use interaction terms in order to avoid the problem of spurious correlations when using ratio variables in regressions. In our empirical work, we did run regressions that include an interaction term between log of total aid and the log of population and found the interaction term not to be statistically significant.

⁵Arndt, Jones, and Tarp (2010, p. 12)

dollars in order to account for the real cost of providing public goods with aid funded flows. We recognize that this is not perfect as the basket of goods used for the calculation of PPPs is not representative of the basket of goods a donor would typically use for implementing aid projects. To take the natural logarithm is often done in time-series analysis with trending variable. Here, it addresses also potential non-linearity in terms of the scale of aid, which in our context is useful.

This is not the first paper to observe that “the scaling exercise” of dividing by the size of an economy can introduce econometric biases leading to spurious results. For example, Brunnschweiler and Bulte (2008) refer to problems when studying the relationship between natural resources and income per capita growth. In the resource literature, the ratio of resource exports to GDP is often used as a measure of natural resources. Brunnschweiler and Bulte (2008) point out that the ratio depends on economic policies, institutions and other factors that affect GDP, which then makes it endogenous. In order to avoid the bias, they use resource stocks per capita as their preferred measure of resource abundance.⁶ This channel of introducing endogeneity would also apply in our context.

There is also a literature in Sociology discussed in Kronmal (1993) that shows that using ratio variables can lead to spurious correlations. In this literature the ratios typically use population size as the denominator. In the present paper, we show that such spurious correlations emerge when studying the impact of aid on growth when using aid-to-GDP ratios. To demonstrate this consider the case of “one independent variable ratio,” which is analyzed in Section 4 in Kronmal’s paper. Specifically, consider the following model:

$$g = \beta_0 + Y^{-1}1_n\beta_{1/Y} + A\beta_A + \epsilon, \tag{1}$$

where n denotes the number of aid recipients, g is an $n \times 1$ vector of growth observations, Y is an $n \times n$ diagonal matrix containing values for our aid normalization variable (GDP),

⁶The work by Brunnschweiler and Bulte (2008) got subsequently criticized by van der Ploeg and Poelhekke (2010), who show that their resource measure also has endogeneity problems due to the fact of how the World Bank calculates resource stocks.

and A is an $n \times 1$ vector of aggregate aid values. Assume that the independent variables and the error term in (1) are uncorrelated. Note that aid and GDP enter separately into (1).

Consider now the alternative model given by

$$g = \alpha_0 + Y^{-1}A\alpha_{A/Y} + \varphi. \quad (2)$$

The regression equation (2) now uses an aid ratio variable, where aid values are divided by the normalization values in Y . The estimated OLS regression coefficient $\hat{\alpha}_{A/Y}$ in (2) equals

$$\hat{\alpha}_{A/Y} = (A'Y^{-1'}Y^{-1}A)^{-1}A'Y^{-1'}g. \quad (3)$$

Now if (1) is the “correct” specification but we estimate (2), maybe for reasons discussed earlier, then the estimated regression coefficient on aid-to-GDP can be derived by substituting (1) into (3):

$$\begin{aligned} \hat{\alpha}_{A/Y} &= (A'Y^{-1'}Y^{-1}A)^{-1}A'Y^{-1'}Y^{-1}\mathbf{1}_n\beta_{1/Y} + (A'Y^{-1'}Y^{-1}A)^{-1}A'Y^{-1'}A\beta_A + \\ &\quad + (A'Y^{-1'}Y^{-1}A)^{-1}A'Y^{-1'}\mathbf{1}_n\beta_0 + (A'Y^{-1'}Y^{-1}A)^{-1}A'Y^{-1'}\epsilon. \end{aligned}$$

The last two terms in this expression equal zero, because A and Y are exogenous as assumed earlier. We then obtain

$$\hat{\alpha}_{A/Y} = \hat{\beta}_{Y^{-1}, Y^{-1}A}\beta_{1/Y} + \hat{\beta}_{A, Y^{-1}A}\beta_A, \quad (4)$$

where $\hat{\beta}_{y,x}$ denotes a predicted regression coefficient of a regression of y on x . The key to see in (4) is that it is possible for $\hat{\alpha}_{A/Y}$ to be non-zero even when there is no relationship between aid and growth, i.e. $\beta_A = 0$. This happens when both $\hat{\beta}_{Y^{-1}, Y^{-1}A}$ and $\beta_{1/Y}$ are non-zero. If our normalization variable Y is GDP, then this is likely the case. Growth and GDP are correlated, certainly then, when controlling for initial GDP. As our experiments below will show, the sign of $\beta_{1/Y}$ depends on the GDP dynamics we are considering. For example, if

growth is a random draw every year, then growth and current income are positively correlated conditional on initial income. In this case, $\beta_{1/Y} < 0$ producing a negative ANB. On the other hand, if incomes converge, then average growth and income are negatively correlated. In this case, $\beta_{1/Y} > 0$ producing a positive ANB. Finally, note that if A and Y are not exogenous,⁷ then this endogeneity will also be channeled through $\hat{\alpha}_{A/Y}$.

3 Analysis and Simulation Experiments

To illustrate the econometric analysis above, we generate artificial datasets from models in which aid will have no actual effect on the level of GDP or investment. In the simulations, all aid is consumed. Therefore, the econometric methods used to analyze the artificial data should not reveal a relationship between aid and growth. If a relationship is detected, then this reveals a bias created by our data analytics.

3.1 Random Growth Experiment

Consider a set of 100 identical recipient countries, where for each recipient i the growth rate of GDP per capita every year is

$$g_{i,t} = \theta_{i,t} + \gamma_{i,t}, \tag{5}$$

where $\theta_{i,t}$ and $\gamma_{i,t}$ each are randomly drawn from a normal distribution:

$$\theta_{i,t} \sim iid N(\theta, \sigma_\theta^2), \tag{6}$$

$$\gamma_{i,t} \sim iid N(\gamma, \sigma_\gamma^2). \tag{7}$$

Therefore, growth amounts to the combination of two independent random draws for every recipient-year pair from two distributions that remain constant over time. We assume that $\theta_{i,t}$ is observable, whereas $\gamma_{i,t}$ is not. A component of growth is assumed to be observable to

⁷This would be true if donors allocate more aid to poorer countries.

identify the effects that an omitted variable has on the results of a growth regression, and particularly on the ANB.

Assume further that all recipients receive a random amount of aid, where aid for recipient i in year t equals

$$A_{i,t} = z_{i,t}A_0(1 + \phi)^t, \quad (8)$$

and the random component $z_{i,t}$ of aid is derived as follows:

$$\ln z_{i,t} \sim iid N(0, \sigma_z^2). \quad (9)$$

The parameters A_0 and ϕ are the common initial aid allocation and growth rate of aid, respectively.

Let these economies grow for 30 years and then calculate the average growth rate and the average aid-to-GDP ratio over these years. Parameter values used in this simulation exercise are not particularly important. Nonetheless, we choose values to match the following averages in our dataset:⁸ an average initial aid-to-GDP ratio of 8%; an average annual growth rate of aid of 4%; and an average annual growth rate of GDP per capita and standard deviation for aid recipients of 2.6% and 5%, respectively. We set $\theta = \gamma = 1.3\%$ and $\sigma_\theta^2 = \sigma_\gamma^2 = 2.5\%$ to match these numbers. All countries are assumed to have a common population growth rate of 1.4% per year, so that aid-to-GDP has no long run trend. Finally, we set aid volatility at $\sigma_z^2 = 0.53$, which is equal to the log of the average standard deviation across years of the ratios of total aid to average aid in a given year in our sample.

In this experiment, aid and income are determined by independent processes. Clearly, aid does not affect growth rates, so by construction aid should be uncorrelated with growth. However, Figure 1, which plots this relationship, shows a clear negative correlation between aid (measured by the aid-to-GDP ratio) and growth. The intuition for this is straightforward: Countries that have grown much relative to others in average – for random reasons – have a

⁸The countries included in our full sample are depicted in the scatter plot shown in Figure 3.




Figure 1: Aid-growth relationship with random growth

low aid measure as high growth translates into higher GDP, which reduces the aid-to-GDP ratio. Likewise, countries that have grown little relative to others in average, have a high aid measure as low growth rates translate into a lower GDP, and therefore a higher aid-to-GDP ratio. Notice that this experiment controls for initial income levels as typically done in aid-growth regressions because countries are assumed to be identical initially.

Table 1 confirms the negative relationship between the aid-to-GDP ratio and growth showing that this relationship is significant at the 1 percent level in Column (I). The results in Column (I) predict incorrectly that a one-percent point increase in aid reduces growth by a third of a percentage point. Note that by performing a few additional regressions we can reconstruct the estimated aid-to-GDP coefficient by using (4). The coefficient $\hat{\beta}_{Y^{-1}, Y^{-1}A} = 0.00045$ and $\hat{\beta}_{1/Y} = -712.65$. Multiplying the two yields -0.3207 which is close to the estimate of -0.3353 shown in Column (I). Column (II) runs a regression between growth and θ , as we have assumed that θ is observable. We observe a one-to-one relationship, which is expected. Increasing θ by one percentage point increases growth by one percentage point. Note also that the R-squared is 50% and exactly matches its expected value, as θ is constructed to produce half of the variation in growth rates. However, when we add the aid-

to-GDP ratio in Column (III) that coefficient now drops quite a bit, and we observe that the coefficient for the aid-to-GDP ratio is still negative and highly significant. Thus, some of the effect of θ on growth is channeled through our aid measure. The R-squared increases to 66%, which happens despite the fact that aid has no relationship with growth by construction. In addition, when not controlling for θ , the coefficient on the aid-to-GDP ratio is substantially larger in absolute value (Column I) as compared to when we control for θ (Column III). Thus, we can conclude that the failure to properly control for variables affecting growth increases the ANB. Finally, Column (IV) shows a regression that uses total aid instead of the aid-to-GDP ratio as a comparison. Here, the aid coefficient is not significant – which of course is expected – and the coefficient on θ and the R-squared have the expected magnitudes.

Table 1: Aid and Growth in Random Growth Experiment

	(I)	(II)	(III)	(IV)
Aid-to-GDP Ratio	-0.3353*** (0.0303)		-0.2244*** (0.0331)	
θ		1.0140*** (0.1016)	0.5866*** (0.1052)	0.9997*** (0.1044)
Total Aid				-0.0182 (0.0285)
Constant	5.7379*** (0.2882)	1.2490*** (0.1437)	3.9206*** (0.4120)	1.5799*** (0.5389)
N	100	100	100	100
R-squared	0.56	0.50	0.66	0.51

When repeating the random growth experiment a thousand times, we obtain an average coefficient of -0.281 and -0.174 for the aid-to-GDP ratio corresponding to Columns (I) and (III), respectively. The average coefficient for θ equals 0.9992 and 0.7706 corresponding to Columns (II) and (III), respectively.

Here, we assumed that aid is allocated randomly. But note that if poorer countries get more aid, as it is the case in reality (see Alesina and Dollar, 2000; Annen and Moers, 2017; Annen and Knack, 2018), then the negative correlation shown in Figure 1 is exacerbated as this increases the aid measure of low growth recipients and decreases it for high growth recipients. Thus, a more realistic aid allocation scenario in our experiment – in which donors are poverty selective – increases the negative bias between aid and growth.

3.2 Divergence Experiment

In the following we show that strong biases also emerge when growth happens as explained by a standard version of the Solow growth model. We work with this model because it has been used to provide the theoretical support for most empirical aid research. Variables are expressed in discrete time. Consider again a set of 100 recipient countries that are initially identical, that is, they have common initial endowments of capital K_0 , labour L_0 , and technology X_0 . Parameter values, with the exception of the savings rate s_i , are also identical. The only other difference between countries is in their aid allocation, which is randomly drawn from a distribution that is identical across countries and time. Here, aid is allocated the same as in the previous experiment. This part of the analysis is not essential but will guarantee that the observed relationship between aid and growth in the model is not perfectly linear. The equations that describe the model economics are (8), (9), and:

$$Y_{i,t} = (X_t L_t)^{1-\alpha} K_{i,t}^\alpha, \quad (10)$$

$$X_t = X_0(1+g)^t, \quad (11)$$

$$L_t = L_0(1+n)^t, \quad (12)$$

$$K_{i,t+1} = I_{i,t} + (1-\delta)K_{i,t}, \quad (13)$$

$$I_{i,t} = s_i Y_{i,t}, \quad (14)$$

$$C_{i,t} = (1-s_i)Y_{i,t} + A_{i,t}, \quad (15)$$

where all notation is as in the macroeconomics and development literature. Notice that in this setting all aid is consumed, and therefore by construction it will have no actual effects on capital investment or output. Economies will differ in their steady-state income levels, and these differences are fully captured by the savings rate. We assume that each country has a savings rate that is randomly drawn from a uniform distribution between s_L and s_H ,

$$\text{Prob}(s_i = s) = \frac{1}{s_H - s_L} \text{ for } s \in [s_L, s_H]. \quad (16)$$




Figure 2: Aid-growth relationship with diverging countries

We calculate growth rates and aid statistics for 100 artificial economies over a 30 year transition period. The model is calibrated using parameter values common in the literature or to match averages in our dataset of aid recipients: $\alpha = 1/3$, $g = 2.6\%$, $n = 1.4\%$, $\delta = 10\%$, $s \in [0.09, 0.50]$, $\phi = 4\%$, and $A_0/Y_0 = 8\%$. Furthermore, Y_0 is chosen such that the country with the average savings rate is in its steady state. All other countries will experience transitional growth. Aid volatility σ_z^2 is set at 0.53 as before. Notice that all the growth differences across countries happen because of transitional dynamics, as each artificial economy has the same steady-state growth rate $g + n$. Again, we find a negative correlation between our aid measure and growth, even though aid does not contribute to growth by construction as shown in Figure 2.

Column (I) in Table 2 confirms the strong negative relationship shown in Figure 2. This relationship is statistically highly significant: Increasing the aid-to-GDP ratio by one percentage point reduces growth by a little bit less than a third of a percentage point. Column (II) in this table adds the savings rate as a control. We know that in average, countries with a higher savings rate have a higher growth rate, because countries are initially identical and they only differ by their steady-state income level. This regression shows that the aid-to-

GDP ratio is negatively and significantly correlated with growth. The key insight here is that controlling for the steady state level of income, which here is defined by the savings rate s , does not remove the ANB. We find a negative and statistically significant relationship between aid and growth. Similar than in Table 1, we find that the ANB increases in the extent our regression fails to properly control for variables that affect growth as the aid coefficient in Column (I) is substantially larger than in Column (II). Also, some of the effect of savings on growth is channeled through the aid measure as the coefficient for savings in Column (II) is smaller than in Column (III). Noteworthy is also the high R-squared of 70% in Column (I), even though we know that aid does not affect growth by construction. The regression reported in Column (IV) confirms this: there is no correlation between total aid and growth.

Table 2: Aid and Growth with Divergence

	(I)	(II)	(III)	(IV)
Aid-to-GDP Ratio	-0.3176*** (0.0210)	-0.0338*** (0.0119)		
Savings Rate		5.7724*** (0.1980)	6.2362*** (0.1169)	6.2410*** (0.1189)
Total Aid				0.1321 (0.5139)
Constant	6.1242*** (0.2075)	1.6642*** (0.1669)	1.1997*** (0.0374)	1.1663*** (0.1353)
N	100	100	100	100
R-squared	0.70	0.97	0.97	0.97

3.3 Convergence Experiment

Consider now a different simulation exercise with recipients in a Solow growth setting, but this time recipients are identical except for their endowment of capital $K_{i,0}$. The initial capital stocks will be drawn from a uniform distribution between K_L and K_H .⁹ Again, we assume that growth happens as explained in a standard Solow model. In particular, we

⁹Values for these supports are chosen so that average growth rates for the artificial economies range between -2% and 10%, which is consistent with the majority of counties in our sample of aid recipients.

assume that all countries have an identical savings rate ($s=20\%$), which implies that they will all converge in income levels and growth rates. The economies in our model are allowed to grow for 30 years, where $s * Y$ is added to the capital stock in every year.

As before, economic growth happens because of transitional dynamics. But now initial income is perfectly correlated with growth, as initially poor countries will catch up to the income level in the wealthy countries. Here, a high aid measure, due to having low initial income, is associated with high growth, and this produces a strong positive correlation between our aid measure and growth that is statistically highly significant, as reported in Column (I) in Table 3. Running a regression conditional on initial income will not remove the correlation between our aid measure and income per capita growth. This time, this correlation is positive, as reported in Column (II). The coefficient for initial income is negative, which is expected because countries converge by construction. We confirm a similar finding as in the previous two tables, which states that the ANB increases in the extent the regression fails to properly control for variables that affect growth. The aid coefficient is substantially larger in Column (I) than in Column (II), which adds the control for initial income. Also, the convergence coefficient is smaller in Column (II) than in Column (III), which again suggests that some of the income effect on growth is channeled through the aid measure. Finally, Column (IV) confirms that there is no correlation between total aid and growth.

Table 3: Aid and Growth with Convergence

	(I)	(II)	(III)	(IV)
Aid-to-GDP Ratio	1.3018*** (0.0542)	0.7903*** (0.1109)		
Initial GDPpc		-0.0984*** (0.0192)	-0.2215*** (0.0103)	-0.2212*** (0.0104)
Total Aid				0.1665 (0.6540)
Constant	-5.3147*** (0.3201)	-0.6374 (0.9559)	6.0683*** (0.2092)	5.8232*** (0.9851)
N	100	100	100	100
R-squared	0.85	0.89	0.83	0.83

Notice that the last two simulations considered countries where growth is driven by

the process of transitional dynamics to a steady-state. If countries have different steady state growth rates then the analysis shown in Table 1 applies. For example, Bernanke and Gürkaynak (2001) show that total factor productivity rates vary considerably across countries and in fact are positively correlated with each country’s savings rate, which suggests that using aid-to-GDP measures produces a negative correlation between that aid measure and growth also in the long run.

To sum up, if economies do not converge in their steady-state growth rates, then using an aid-to-GDP measure produces a negative correlation between that aid measure and growth when aid does not affect growth rates. This scenario applies in the long-run. If growth is mostly determined by transitional dynamics, then the correlation between aid-to-GDP ratio and growth can be positive or negative depending on whether countries are “catching up” (positive correlation) or “diverging” (negative correlation). By “diverging” we mean countries that are more similar initially but then converge to their respective steady state so that income differences across countries increase. Evidence of increased in-between country income inequality (Bourguignon and Morrisson, 2002) points more to such a diverging pattern than a converging one in the data. Exploring such data is what we consider next.

4 Aid Effectiveness in the Last 20 Years

In this section we estimate aid-growth regressions as done in the large literature on aid effectiveness. The main purpose of this is not to identify a causal relationship between aid and growth but to provide evidence that the aid normalization choice matters. We also can show that the predictions derived from our simulation exercises show up in our aid growth regressions using actual data. Taken together, we interpret this as evidence that using aid-to-GDP ratios is likely to downward bias the results. In fact, in all our regressions we find a positive and often statistically highly significant relationship between aid and growth when we use the log of total aid; whereas when using aid-to-GDP ratios, results vary a lot across

specifications, samples, and estimation techniques with a significant positive relationship in some regressions and a significant negative one in others.

The aid data we obtain from the OECD’s DAC database (Table 2a) and our measure for total aid has been PPP adjusted. We include bilateral and multilateral aid as bilateral donors allocate a substantial amount of their aid via multilateral aid agencies such as IDA and other agencies (Milner, 2006; Rodrik, 1996; Annen and Knack, 2018). In our analysis we use the same aid measure as in Annen and Kosempel (2009). This is a measure of gross aid that removes debt forgiveness, emergency aid, and food aid. This brings our measure closer to a measure of what we consider “development aid” — intended for countries to grow. Emergency aid cannot be expected to increase growth rates (Annen and Strickland, 2017) and to somewhat a lesser extent the same can be said about food aid. Debt forgiveness introduces strong aid volatility as it is counted as aid at the time the debt is forgiven even though that debt has been incurred many years ago and the forgiveness may result in very little fiscal expansion. Fiscal expansion will of course depend on how much debt repayment occurred before the debt relief (Cassimon, Campenhout, Ferry, and Raffinot, 2015) and even with positive debt repayments, these yearly repayments will be of orders of magnitude smaller than the debt forgiven at a given moment.

Our regression analysis focuses on the years between 1995 and 2014. There are good reasons to do so as the 70s and 80s were characterized by economic turmoil in many aid recipient countries: The oil crisis first and the debt crisis next.¹⁰ At the end of the 80s and early 90s many recipient countries introduced economic and political reforms as many underwent structural adjustment programs supported by the IMF and the World Bank. There are also reasons to believe that aid at that point in time became less politicized because of the end of the Cold War. In addition, aid allocations became more policy- and poverty selective, particularly among multilateral agencies such as IDA or regional development banks (e.g Dollar and Levine, 2006; Knack, Rogers, and Eubank, 2011; World Bank, 2005). This

¹⁰For example, Clemens, Radelet, Bhavnani, and Bazzi (2012) show how the inclusion of the 70s changes panel regression results in RS, who dropped the 70s in their panel regressions.

Figure 3 shows a cross-section scatter plot between initial income per capita in 1995 and average growth between 1995 and 2014 for all aid recipients in our sample. We include all recipients in our analysis for which we have aid, GDP, and growth data with the exception of Equatorial Guinea and Liberia that both are outliers in our sample.¹¹ The two dashed lines in Figure 3 are drawn at the average value for initial income and growth respectively. The two lines divide the observations into four quadrants. There is no evidence of unconditional convergence among the poorest aid receiving countries. Countries are fairly evenly distributed in the four quadrants.

Table 4 provides some descriptive statistics related to the four quadrants in Figure 3. High-income-high-growth countries (Quadrant I) received a considerable amount (33%) of total aid. The countries in the two high-growth quadrants (I and IV) received the bulk of foreign aid with 72%. Again, this evidence cannot be used to reject the claim that aid increases growth. Countries in the two low-income quadrants (III and IV) received 57% of total aid, where about 68% of that went to high-growth countries. Aid appears to be (somewhat) poverty selective as the poorer countries received more than half of total aid. In terms of policy, we observe that higher income countries have higher policy levels, whereas the policy levels of countries in quadrants III and IV are similar. In particular, the higher growth performance among poorer countries seems unrelated to better policies as countries in this quadrant have essentially the same policy level than the countries in the low growth quadrant III. Aid per capita varies quite a bit across quadrants with quadrant III with 287 PPP adjusted USD having the largest per capita amount. Quadrant I with 92 USD has the smallest amount. We also observe that the aid-to-GDP ratio varies substantially across quadrants and it ranges between 0.8% in Quadrant I and 11.6% in Quadrant III. Some of this difference among low income countries is driven by the difference in growth performance as the ratio of aid per capita between Quadrant IV and III is smaller than the ratio of aid-

¹¹Equatorial Guinea had an average growth rate of 15% in our sample period and Liberia was at the heights of a civil war in 1995. Note that our econometric results are not sensitive to the exclusion of these outlier countries.

to-GDP between these two quadrants. The higher growth in Quadrant IV countries deflates the aid-to-GDP ratio for these countries.

Table 5: Univariate Cross-Section-OLS Regressions 1995–2014

Sample	Total Aid (log)			Aid-to-GDP		
	Full	Conv.	Div.	Full	Conv.	Div.
	(I)	(II)	(III)	(IV)	(V)	(VI)
Aid	0.31*** (0.08)	0.63*** (0.09)	0.00 (0.08)	-0.04*** (0.01)	0.14 (0.09)	-0.08*** (0.02)
Constant	0.68 (0.53)	-0.95* (0.51)	2.24*** (0.56)	2.80*** (0.18)	2.38*** (0.36)	2.74*** (0.25)
N	148	76	72	148	76	72
R-squared	0.07	0.27	0.00	0.03	0.09	0.16
F statistic	14.13	53.49	0.00	11.23	2.21	11.15

Dependent variable is per capita GDP growth (PPP adjusted). Recipient level cluster-robust standard error reported in parenthesis. Aid is measured in PPP adjusted USD. Full, Conv., and Div. refers to the full sample, the sample of converging and diverging countries respectively. Significance levels : * : 10 percent ** : 5 percent *** : 1 percent.

Our simulation experiment in Section 2 suggests that using aid-to-GDP ratios introduces a negative bias among diverging and a positive bias among converging countries. These biases show up in negative and positive correlations in aid-growth regressions when using aid-to-GDP ratios as our aid measure. Of course, in these experiments aid had no impact on growth by construction. With actual data, in contrast, recipients could use aid to finance capital projects, which may then affect growth rates.

Table 6 shows the results of univariate aid-growth regressions in a cross-section using averages in aid and growth between 1995 and 2014. The first three columns use the log of total aid measured in PPP adjusted USD for the full sample, and the samples of converging countries and diverging countries respectively. A country belongs to the sample of converging countries if it is located in quadrant II and IV in Figure 3. A country belongs to the sample of diverging countries if it is located in quadrant I or III in Figure 3. In all these three regressions we find a positive relationship between aid and growth, although the coefficient in the divergence sample is small. The coefficients in the full and convergence sample are each

significant at the 1% level. The next three columns report regression results using the same specification and samples but using the aid-to-GDP ratio as our aid measure. We find that the coefficient changes substantially across samples, with a significant negative correlation in the full and in the divergence sample and a positive but not significant coefficient in the convergence sample. We believe that this result connects well to our simulation experiments: A positive bias among converging and a negative bias among diverging countries. These regressions show that the aid normalization choice matters. We get a positive and highly significant relation between aid and growth when using total aid, and we get a negative and highly significant relation between aid and growth when we use the aid-to-GDP ratio, and the coefficient differences between the divergence and convergence sample are consistent with the directions of the ANB found in our simulation exercises. Since these regressions do not control for any variables that may affect growth, our simulation exercises suggest that the ANB in the regressions that use the aid-to-GDP ratio is large.

Table 6: Cross-Section-OLS Regressions 1995–2014

Sample	Total Aid (log)			Aid-to-GDP		
	Full	Conv.	Div.	Full	Conv.	Div.
	(I)	(II)	(III)	(IV)	(V)	(VI)
Aid	0.45*** (0.10)	0.16 (0.11)	0.25*** (0.08)	-0.08*** (0.02)	-0.02 (0.03)	-0.03*** (0.01)
Initial GDPpc (log)	-0.34* (0.18)	-1.16*** (0.17)	1.04*** (0.31)	-0.81*** (0.21)	-1.32*** (0.16)	0.91** (0.35)
Policy (WGI)	1.14*** (0.34)	0.27 (0.36)	0.56 (0.54)	0.89*** (0.33)	0.16 (0.30)	0.33 (0.52)
Constant	3.08* (1.72)	11.90*** (1.82)	-7.94*** (2.52)	10.10*** (1.90)	14.29*** (1.55)	-5.14* (3.00)
N	148	76	72	148	76	72
R-squared	0.15	0.48	0.46	0.12	0.47	0.44
F statistic	8.07	32.19	27.73	9.19	31.60	20.41

Dependent variable is per capita GDP growth (PPP adjusted). Recipient level cluster-robust standard error reported in parenthesis. Aid is measured in PPP adjusted USD. All regressions include period dummy variables, which are not reported. Full, Conv., and Div. refers to the full sample, the sample of converging and diverging countries respectively. Significance levels : * : 10 ** : 5 percent *** : 1 percent.

These univariate regressions clearly suffer from omitted variable bias. The most obvious

one relates to poverty selectivity, which occurs when donors allocate more aid to poorer countries. In the convergence sample, we know there is a negative correlation between initial income and growth, whereas in the divergence sample this correlation is positive. If aid indeed is given selectively depending on income per capita then the estimates in Table 5 are biased – an upward bias in the convergence sample and a downward bias in the divergence sample. In addition, to properly control for poverty-selectivity seems more important for the diverging sample as this sample includes initially poor countries that have not grown between 1995 and 2014. Another potential omitted variable bias may have to do with the fact that many donors allocate aid policy-selectively. And if better policies increase growth then some of the effect of aid on growth found in Table 5 should be attributed to better policies instead of aid.¹² For these reasons, Table 6 includes a control for initial income per capita and policy. For our policy variable we use the World Bank Governance Indicators (WGI). Our measure is calculated by taking the average between “Control of Corruption,” “Rule of Law,” and “Government Efficiency.” Note that these measures are expressed as z -scores. Table 6 reveals that we continue to find a positive relationship between aid and growth after controlling for policy and initial income per capita when we use total aid as our aid measure. We also continue to find a negative correlation between aid and growth when we use the aid-to-GDP ratio as our aid measure. Also, note that the aid-to-GDP coefficients between Table 5 and 6 change as predicted as they decrease in the convergence sample and increase in the divergence sample.

In all the subsequent tables shown in this paper, we analyze the aid growth relationship by dividing our data into 4-year panels. Averaging is done in order to smooth out business cycle shocks.¹³ We also include many more controls using a specification that comes closer to the one used in Burnside and Dollar (2000). For example, we add policy controls such as openness and inflation, and M2/GDP lagged. All this data is taken from the World

¹²In contrast, if policy-selective aid provides incentives for recipient countries to improve policies, then controlling for policy removes one channel through which aid affects growth.

¹³In Annen, Batu, and Kosempel (2016) we show that productivity shocks easily dominate aid shocks, which highlights the importance of averaging in order to smooth out business cycle shocks.

Table 7: OLS, 4-year panel with non-lagged aid

	Total Aid (log)			Aid-to-GDP		
	(I)	(II)	(III)	(IV)	(V)	(VI)
Aid	0.45*** (0.08)	0.55*** (0.10)	0.44*** (0.13)	-0.03 (0.04)	-0.03 (0.04)	0.04 (0.04)
Initial GDPpc (log)	0.01 (0.19)	-0.36 (0.24)	-0.44* (0.25)	-0.27 (0.24)	-0.70*** (0.26)	-0.57** (0.25)
Policy (WGI)	0.74** (0.30)	0.84** (0.38)	0.80** (0.38)	0.45 (0.31)	1.06** (0.41)	0.81** (0.40)
Life Expectancy		0.02 (0.04)	0.03 (0.04)		0.02 (0.05)	0.04 (0.05)
Openness		0.02*** (0.00)	0.02*** (0.00)		0.01*** (0.00)	0.02*** (0.00)
Inflation		-0.00 (0.01)	-0.00 (0.01)		0.00 (0.00)	-0.00 (0.01)
M2/GDP lagged		-0.02*** (0.01)	-0.02*** (0.01)		-0.02*** (0.01)	-0.02*** (0.01)
Political Stability (WGI)		0.23 (0.27)	0.33 (0.28)		-0.25 (0.25)	0.30 (0.29)
EAP		-0.84* (0.48)	-0.90* (0.48)		-0.45 (0.55)	-1.00* (0.52)
SSA		-0.96 (0.65)	-0.98 (0.67)		-1.19* (0.69)	-1.06 (0.69)
Population (log)			0.14 (0.12)			0.44*** (0.10)
Constant	-0.58 (1.90)	0.45 (3.02)	-0.56 (3.20)	4.73** (2.22)	7.48** (3.00)	-2.00 (3.88)
N	588	588	588	588	588	588
R-squared	0.14	0.21	0.21	0.09	0.15	0.19
F statistic	18.67	12.81	12.26	12.98	8.58	10.67

Dependent variable is per capita GDP growth (PPP adjusted). Recipient level cluster-robust standard error reported in parenthesis. Aid is measured in PPP adjusted USD. All regressions include period dummy variables, which are not reported. Significance levels : * : 10 ** : 5 percent *** : 1 percent.

Development Indicators. We also include a control for political stability, which we take from the World Governance Indicators. Finally, we include a Sub-Saharan Africa dummy and East Asia Pacific dummy to control for geography.

In Table 7 – our main table –, we observe that our aid measure is robust to these additional controls when we use total aid. When we use the aid-to-GDP ratio, we find that the aid coefficient changes across specifications but it is not significant with the additional controls. We also ran the aid-to-GDP ratio regressions with a square term as this is often done in studies that use the aid-to-GDP ratio. With a square term and the regression specification

used in Column (V), the aid coefficient decreases to -0.12 with a p -value of 0.108 and the coefficient on the square term is positive but not significant. Notice also that Column (III) and (VI) add the log of total population in order to control for the size of the economy. Again, we find that our aid coefficient hardly changes when we use total aid. In contrast, when using the aid-to-GDP ratio, the aid coefficient is now positive but not significant. Surprisingly, the coefficient on population is highly significant when we use the aid-to-GDP ratio and it is not significant when we use total aid. This is puzzling but may explain why population is often used as an instrument in IV estimation in aid effectiveness studies.¹⁴

A sensitivity analysis was conducted and revealed that our main result does not depend on whether we exclude or include the two outliers Equatorial Guinea and Liberia in our regression. Also, our results are robust to changes in the time period length when creating regression panels. For example, Burnside and Dollar (2000) uses 4 year periods whereas Rajan and Subramanian (2008) uses 5 year panels. When we repeated the regressions reported in our main Table 7, but with 5-year instead of 4-year averages, we observed that our results are robust to this change. Specifically, we found that the coefficient is almost the same when using total aid, whereas when using the aid-to-GDP ratio we still find that the coefficient is positive or negative depending on whether we control for population or not.

Note also that in an Appendix to this study we show replication results for two influential papers in the aid literature, namely Burnside and Dollar (2000) and Rajan and Subramanian (2008). These papers cover the years between 1970 and 1997, and 1961 and 2000, respectively. These replications show that the aid normalization choice affects the results. In all regressions we find larger and sometimes statistically significant aid coefficients when using the log of total aid, which is in line with our conjecture that using aid-to-GDP very likely introduces a negative bias.

¹⁴See Clemens, Radelet, Bhavnani, and Bazzi (2012) for an interesting discussion of this point.

4.1 Lagging Aid and Controlling for Fixed Effects

Clemens, Radelet, Bhavnani, and Bazzi (2012) argue that one should use lagged aid as a simple remedy for potential reversed causality bias as “current growth is likely to affect current aid.” However, lagging the aid variable can also help to reduce the ANB. With lagged aid-to-GDP, any shock that affects current growth rates will no longer show up in the aid measure as the GDP used in the aid measure is lagged by one period. Lagging the aid measure may reduce the aid normalization bias if growth shocks are uncorrelated across periods. However, if countries have persistent differences in their growth performance, then lagging the aid measure has limited use in reducing this bias. Additionally controlling for country fixed effects may then be a useful tool to further reduce this bias. As we have seen in our experiments in Section 3, the ANB increases in the extent to which we fail to properly control for *all* variables that affect economic growth. Of course, there are also other reasons why controlling for country fixed effects is useful, as pointed out by Clemens, Radelet, Bhavnani, and Bazzi (2012).

Table 8 reports the result of regressions that control for recipient country fixed effects and that use lagged aid as the key independent variable. We observe now that aid is positively and significantly related to growth in *all* the regressions, including the ones that use the aid-to-GDP ratio. We also observe, that controlling for the size of the population does no longer affect the aid-coefficient when using the aid-to-GDP ratio. Controlling for fixed effects and lagging the aid measure by one period produces a large change in the aid coefficient in the regression reported in Column (IV). This regression includes the least amount of controls. This coefficient changes from a statistically insignificant negative value of -0.03 to a statistically highly significant positive 0.08. This shift supports our predictions that we produced using our simulation experiments.

Note that to isolate the individual effects that lagging the aid measure and controlling for country fixed effects have on the estimated aid effectiveness coefficient it is necessary to perform separate regressions with lagged aid and OLS and non-lagged aid and fixed effects,

Table 8: FE, 4-year panel with lagged aid

	Total Aid (log)			Aid-to-GDP		
	(I)	(II)	(III)	(IV)	(V)	(VI)
Aid (lagged)	0.80*** (0.29)	0.77*** (0.27)	0.75*** (0.26)	0.08*** (0.03)	0.08*** (0.03)	0.07*** (0.03)
Initial GDPpc (log)	-6.53*** (0.83)	-6.61*** (0.75)	-7.04*** (0.86)	-5.96*** (0.79)	-6.08*** (0.74)	-6.48*** (0.86)
Policy (WGI)	1.14 (0.81)	0.31 (0.81)	0.10 (0.83)	1.24 (0.79)	0.35 (0.81)	0.18 (0.83)
Life Expectancy		0.01 (0.07)	0.05 (0.08)		0.03 (0.07)	0.06 (0.07)
Openness		0.03*** (0.01)	0.03*** (0.01)		0.03*** (0.01)	0.03*** (0.01)
Inflation		-0.02** (0.01)	-0.02** (0.01)		-0.02** (0.01)	-0.02** (0.01)
M2/GDP lagged		-0.02* (0.01)	-0.03** (0.01)		-0.03** (0.01)	-0.03** (0.01)
Political Stability (WGI)		0.42 (0.40)	0.46 (0.40)		0.44 (0.38)	0.48 (0.38)
Population (log)			-2.70 (1.86)			-2.39 (1.95)
Constant	55.28*** (7.49)	53.74*** (8.59)	97.63*** (32.69)	55.03*** (7.03)	52.26*** (8.31)	91.33*** (34.31)
N	597	597	597	597	597	597
R-squared	0.27	0.35	0.35	0.26	0.34	0.34
F statistic	22.42	16.14	15.32	22.60	18.59	17.25

Dependent variable is per capita GDP growth (PPP adjusted). Recipient level cluster-robust standard error reported in parenthesis. Aid is measured in PPP adjusted USD. All regressions include period dummy variables, which are not reported. Significance levels : * : 10 ** : 5 percent *** : 1 percent.

respectively. Although not reported here, these regression were performed and revealed that when the aid-to-GDP is used, its estimated impact on growth is positive only if we both lag aid and control for country fixed effects. Lagging the aid measure reduces the ANB when growth shocks are temporary and are uncorrelated across periods, whereas controlling for country fixed effects reduces the ANB associated with having missing explanatory variables that have persistent long run effects on growth performance. The combination of lagging aid and controlling for country fixed effect seems to reduce the ANB if we use the aid regressions with total aid as our reference point. However, we would like to emphasize that using lagged aid and country fixed effects may not be sufficient to remove the entire bias. Recall from our simulation experiments that some of the ANB remained even after controlling for all

variables affecting growth. Thus, we do not read these regression results as suggesting that using lagged aid and controlling for country fixed effects will completely resolve the issue raised in this paper. A result that we find worth highlighting is that the aid coefficient when using total aid is fairly robust to all the changes we have introduced so far. In contrast, results change substantially when using the aid-to-GDP ratio.

4.2 Aid and Growth in Low Income Countries

In this subsection we investigate whether our results hold up when reducing the sample to low-income countries. We define a country as a low-income country when the GDP per capita income of that country is below the median income in our full sample, which amounts to 5055 PPP adjusted USD in 1995. Hereby we effectively remove all Eastern European countries (among others) from our sample. We focus now on the low-income countries, many of them in sub-Saharan Africa, because the public debate on “Why trying to help poor countries might actually hurt them” often focuses on these countries. For example, Moyo (2009) argues that aid is detrimental to sub-Saharan African countries.

Table 9 shows the results. We observe that we continue to estimate a strong positive and statistically significant relationship between aid and growth when using total aid, and we continue to observe a statistically insignificant relationship between aid and growth when using the aid-to-GDP ratio. In these regressions, the coefficient is positive or negative depending on whether we include a control for population or not as before. Interesting is that the aid coefficient when using total aid increases quite a bit when changing from the full to the low-income country sample. They increase from 0.55 to 0.75 or from 0.44 to 0.63 depending on whether we include a control for population or not. Thus, our main result is not driven by the fact that the full sample includes richer countries, and in particular countries in Eastern Europe. We observe a strong positive and statistically significant relationship between aid and growth among low-income countries.

Table 10 repeats these regression but uses lagged aid and controls for country fixed effects

Table 9: OLS, 4-year panel with non-lagged aid

	Total Aid (log)			Aid-to-GDP		
	(I)	(II)	(III)	(IV)	(V)	(VI)
Aid	0.71*** (0.15)	0.75*** (0.19)	0.63* (0.34)	-0.06* (0.04)	-0.06 (0.05)	0.03 (0.05)
Initial GDPpc (log)	0.60 (0.42)	0.04 (0.45)	0.03 (0.44)	0.27 (0.48)	-0.32 (0.55)	0.12 (0.49)
Policy (WGI)	0.84 (0.57)	0.97 (0.65)	1.00 (0.65)	1.10* (0.59)	1.78*** (0.64)	1.20* (0.66)
Life Expectancy		0.02 (0.07)	0.03 (0.07)		0.02 (0.07)	0.04 (0.07)
Openness		0.01* (0.01)	0.01* (0.01)		0.00 (0.01)	0.01* (0.01)
Inflation		0.01 (0.01)	0.01 (0.01)		0.02* (0.01)	0.01 (0.01)
M2/GDP lagged		-0.04*** (0.01)	-0.04*** (0.01)		-0.04*** (0.01)	-0.04*** (0.01)
Political Stability (WGI)		0.30 (0.40)	0.34 (0.43)		-0.11 (0.39)	0.39 (0.43)
EAP		-1.06 (0.73)	-1.05 (0.73)		-0.86 (0.88)	-1.07 (0.79)
SSA		-1.32 (1.08)	-1.34 (1.09)		-1.73 (1.07)	-1.42 (1.13)
Population (log)			0.11 (0.28)			0.58*** (0.16)
Constant	-6.34 (3.85)	-2.07 (5.80)	-2.95 (6.56)	1.62 (4.15)	7.65 (6.24)	-7.36 (7.93)
N	295	295	295	295	295	295
R-squared	0.17	0.23	0.23	0.10	0.18	0.23
F statistic	8.48	7.35	7.04	5.58	5.40	7.46

Sample of aid recipients with an income per capita below the median of all aid recipients and excluding Equatorial Guinea and Liberia. Dependent variable is per capita GDP growth (PPP adjusted). Recipient level cluster-robust standard error reported in parenthesis. Aid is measured in PPP adjusted USD. All regressions include period dummy variables, which are not reported. Significance levels : * : 10 percent ** : 5 percent *** : 1 percent.

as done in Table 8 for the full sample. Similar to our results when using the full sample, we now find a statistically significant positive relationship between aid and growth in *all* the regressions. Compared to the full sample, the point estimates all increase, which suggests that the positive results found in the previous regressions are not driven by the fact that this sample also includes richer countries. Again, the change in aid coefficient reported in Column (IV) in the last two tables is large: From a significant negative coefficient of -0.06 when using current aid and no fixed effects to a highly significant positive coefficient of

Table 10: FE, 4-year panel with lagged aid

	Total Aid (log)			Aid-to-GDP		
	(I)	(II)	(III)	(IV)	(V)	(VI)
Aid (lagged)	1.27** (0.58)	1.24** (0.49)	1.16** (0.45)	0.10*** (0.04)	0.09*** (0.03)	0.08** (0.03)
Initial GDPpc (log)	-5.73*** (0.96)	-5.77*** (0.87)	-6.72*** (1.10)	-4.90*** (0.92)	-5.00*** (0.86)	-6.04*** (1.15)
Policy (WGI)	0.86 (1.18)	0.09 (1.15)	-0.16 (1.20)	1.04 (1.12)	0.21 (1.12)	-0.06 (1.18)
Life Expectancy		-0.02 (0.09)	0.06 (0.09)		0.02 (0.09)	0.10 (0.10)
Openness		0.04*** (0.02)	0.04*** (0.01)		0.04*** (0.02)	0.04*** (0.02)
Inflation		-0.04** (0.01)	-0.04*** (0.01)		-0.03** (0.02)	-0.03** (0.01)
M2/GDP lagged		-0.02 (0.02)	-0.03 (0.02)		-0.03 (0.02)	-0.03 (0.02)
Political Stability (WGI)		0.25 (0.59)	0.22 (0.58)		0.22 (0.58)	0.19 (0.57)
Population (log)			-5.37 (3.70)			-5.61 (3.97)
Constant	40.08*** (8.41)	38.86*** (9.10)	128.19** (62.25)	41.75*** (7.50)	38.75*** (8.82)	132.09* (67.03)
N	295	295	295	295	295	295
R-squared	0.25	0.33	0.34	0.25	0.32	0.34
F statistic	9.33	8.80	9.36	9.21	11.61	11.86

Sample of aid recipients with an income per capita below the median of all aid recipients and excluding Equatorial Guinea and Liberia. Dependent variable is per capita GDP growth (PPP adjusted). Recipient level cluster-robust standard error reported in parenthesis. Aid is measured in PPP adjusted USD. All regressions include period dummy variables, which are not reported. Significance levels : * : 10 percent ** : 5 percent *** : 1 percent.

0.10 when using lagged aid and controlling for country fixed effects. Column (IV) reports regression results of a regressions with the least controls, and our experiment suggests that in these regressions the ANB will be the largest. The large change in the coefficient between these two tables supports this finding.

5 Conclusions

The common practice in the aid effectiveness literature has been to normalize (or scale) aid by dividing it by GDP. This practice started with the first aid-growth regression that we

are aware of, and (perhaps) for consistency carried forward into the literature that followed. However, in the present paper, we show that using aid-to-GDP ratios introduces econometric problems that most likely bias the results against finding a positive effect of aid on growth - we have referred to this as the aid normalization bias (ANB). We demonstrate the presence of the bias in several ways: First, we generated artificial data from a growth model that was constructed so that aid had no actual effect on GDP growth. However, we found that aid-growth regressions performed on the artificial data displayed strong (usually negative) correlations between aid and growth, when the dependent variable was the aid-to-GDP ratio. A negative correlation (or bias) between growth and aid-to-GDP exists in the artificial data because an economy that grows quickly will tend have a relatively high income level, and therefore a relatively low aid-to-GDP level.

In our empirical analysis covering the years between 1995 and 2014 we confirm that our findings are robust to various alterations of the data and statistical analysis. The only time we find a strong positive correlation between the aid-to-GDP ratio and growth is when using a fixed effect estimator with lagged aid. However, even though lagging aid and controlling for country fixed effects may reduce the ANB, these two procedures very likely will not remove the entire bias. Our analysis shows that the aid normalization choice affects the results. The scaling procedure of dividing aid levels by the size of the economy, as has been typically done in the aid effectiveness literature, will create a bias against finding a positive effect of aid on growth. This finding has important policy implications, because the academic literature that has investigated the economic impact of foreign aid has been very influential. In future research, we recommend that aid effectiveness studies use total aid instead of aid-to-GDP.

References

ALESINA, A., AND D. DOLLAR (2000): “Who Gives Foreign Aid to Whom and Why?,” *Journal of Economic Growth*, 5, 33–63.

- ANNEN, K., M. BATU, AND S. KOSEMPEL (2016): “Macroeconomic effects of foreign aid and remittances: Implications for aid effectiveness studies,” *Journal of Policy Modeling*, 38(6), 1136–1146.
- ANNEN, K., AND S. KNACK (2018): “On the delegation of aid implementation to multilateral agencies,” *Journal of Development Economics*, 133, 295–305.
- (2019): “Better policies from policy-selective aid?,” Discussion Paper WPS8889, World Bank, World Bank Policy Research Paper.
- ANNEN, K., AND S. KOSEMPEL (2009): “Foreign Aid, Donor Fragmentation, and Economic Growth,” *The B.E. Journal of Macroeconomics (Contributions)*, 9(33), 1–30.
- ANNEN, K., AND L. MOERS (2017): “Donor Competition for Aid Impact, and Aid Fragmentation,” *The World Bank Economic Review*, 31(3), 708–729.
- ANNEN, K., AND S. STRICKLAND (2017): “Global Samaritans? Donor Election Cycles and the Allocation of Humanitarian Aid,” *European Economic Review*, 96, 38 – 47.
- ARNDT, C., S. JONES, AND F. TARP (2010): “Aid, Growth, and Development: Have We Come Full Circle?,” *Journal of Globalization and Development*, 1(2).
- (2015): “Assessing Foreign Aid’s Long-Run Contribution to Growth and Development,” *World Development*, 69, 6 – 18, Aid Policy and the Macroeconomic Management of Aid.
- BERNANKE, B. S., AND R. S. GÜRKAYNAK (2001): “Is Growth Exogenous? Taking Mankiw, Romer, and Weil Seriously,” *NBER Macroeconomics Annual*, 16, 11–57.
- BOURGUIGNON, F., AND C. MORRISSON (2002): “Inequality Among World Citizens: 1820-1992,” *American Economic Review*, pp. 727–744.
- BRUNNSCHWEILER, C. N., AND E. H. BULTE (2008): “The resource curse revisited and revised: A tale of paradoxes and red herrings,” *Journal of environmental economics and management*, 55(3), 248–264.

- BURNSIDE, C., AND D. DOLLAR (2000): "Aid, Policies, and Growth," *American Economic Review*, 90(4), 847–868.
- CARTER, P. (2017): "Aid econometrics: Lessons from a stochastic growth model," *Journal of International Money and Finance*, 77(Supplement C), 216 – 232.
- CASSIMON, D., B. V. CAMPENHOUT, M. FERRY, AND M. RAFFINOT (2015): "Africa: Out of debt, into fiscal space? Dynamic fiscal impact of the debt relief initiatives on African Heavily Indebted Poor Countries (HIPCs)," *International Economics*, 144, 29–52.
- CLEMENS, M. A., S. RADELET, R. R. BHAVNANI, AND S. BAZZI (2012): "Counting Chickens when they Hatch: Timing and the Effects of Aid on Growth," *The Economic Journal*, 122(561), 590–617.
- DALGAARD, C.-J., AND L. ERICKSON (2006): "Solow Versus Harrod-Domar; Reexamining the Aid Costs of the First Millennium Development Goal," IMF Working Papers 06/284, International Monetary Fund.
- DALGAARD, C.-J., H. HANSEN, AND F. TARP (2004): "On the Empirics of Foreign Aid and Growth," *Economic Journal*, 114, F191–F216.
- DOLLAR, D., AND V. LEVINE (2006): "The Increasing Selectivity of Foreign Aid, 1984–2003," *World Development*, 34(12), 2034–2046.
- EASTERLY, W. (2003): "Can Foreign Aid Buy Growth?," *Journal of Economic Perspectives*, 17(3), 23–48.
- EASTERLY, W., R. LEVINE, AND D. ROODMAN (2004): "Aid, Policies, and Growth: Comment," *American Economic Review*, 94(3), 774–780.
- FIELDING, D., AND S. KNOWLES (2011): "Dangerous Interactions: Problems in Interpreting Tests of Conditional Aid Effectiveness," *The World Economy*, 34(6), 972–983.

- IDA15 (2008): “Additions to IDA Resources: Fifteenth Replenishment,” techreport, World Bank, Report from the Executive Directors of the International Development Association to the Board of Governors.
- JUSELIUS, K., N. F. MØLLER, AND F. TARP (2014): “The long-run impact of foreign aid in 36 African countries: Insights from multivariate time series analysis,” *Oxford Bulletin of Economics and Statistics*, 76(2), 153–184.
- KNACK, S., H. F. ROGERS, AND N. EUBANK (2011): “Aid Quality and Donor Rankings,” *World Development*, 39, 1907–1917.
- KRONMAL, R. A. (1993): “Spurious Correlation and the Fallacy of the Ratio Standard Revisited,” *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 156(3), 379–392.
- MILNER, H. V. (2006): “Why Multilateralism? Foreign Aid and Domestic Principal-agent Problems,” in *Delegation and Agency in International Organizations*, ed. by D. G. Hawkins, pp. 107–139. Cambridge University Press, Cambridge.
- MOSLEY, P. (1980): “Aid, Savings, and Growth Revisited,” *Oxford Bulletin of Economics and Statistics*, 42(2), 79–95.
- MOYO, D. (2009): *Dead Aid: Why Aid is not Working and How there is a Better Way for Africa*. Douglas & McIntyre, Vancouver.
- PAPANEK, G. F. (1973): “Aid, Foreign Private Investment, Savings, and Growth in Less Developed Countries,” *Journal of Political Economy*, 81(1), 120–130.
- RAJAN, R. G., AND A. SUBRAMANIAN (2008): “Aid and Growth: What does the Cross-Country Evidence Really Show?,” *The Review of Economics and Statistics*, 90(4), 643–665.
- RODRIK, D. (1996): “Why is there Multilateral Lending?,” *NBER Working Paper*, (5160).
- SOLOW, R. (1956): “A Contribution to the Theory of Economic Growth,” *The Quarterly Journal of Economics*, 70(1), 65–94.

SWANSON, A. (2015): “Why trying to help poor countries might actually hurt them,” <https://www.washingtonpost.com/news/wonk/wp/2015/10/13/why-trying-to-help-poor-countries-might-actually-hurt-them>, The Washington Post (accessed April 18, 2017).

VAN DER PLOEG, F., AND S. POELHEKKE (2010): “The pungent smell of “red herrings”: Subsoil assets, rents, volatility and the resource curse,” *Journal of Environmental Economics and Management*, 60(1), 44–55.

WORLD BANK (2005): “Global Monitoring Report 2005,” Washington DC.

Appendix: A Replication of Two Influential Aid Studies

In this section we show that the aid normalization choice affects the results of two studies that has been very influential in the aid debate: Burnside and Dollar (2000) (BD) and Rajan and Subramanian (2008) (RS). We begin with the BD study, however, instead of using their original data we use the updated data set that Easterly, Levine, and Roodman (2004) (ELR) used in their critique of the BD study. ELR replicate BD based on several samples. For our replication we use their most comprehensive sample, which covers the years between 1970 and 1997. The results for this sample are reported on the last three rows in Table 2 in ELR (p. 777).

Column (I) in Table 11 shows the replication result with the Aid*Policy coefficient being identical to the one reported in Table 2 in ELR. Column (II) shows the same regression but without the aid-policy interaction term and using a reduced sample that eliminates the 20 observations with a negative aid value. Since our preferred aid measure is the log of total aid, we drop these observations. We confirm that there is no statistically significant correlation between aid and growth when using aid-to-GDP ratios in this reduced sample. The remaining three columns in Table 11 use the log of total aid instead of the aid-to-GDP ratio. We find in all three regressions a positive and statistically significant correlation between aid and growth. Column (III) uses the exact same specification as Column (II) but uses total aid instead of the aid-to-GDP ratio. The aid coefficient is now positive and significant at the 5% level. Column (IV) includes an aid-policy interaction term. We do not find any evidence that the correlation between aid and growth increases in policy. Finally, Column (V) includes the log of total population as an additional control. Combined with initial income per capita this measure controls for the size of an economy. Thus, the aid coefficient relates to the relationship between aid and growth for recipients holding income per capita and population constant. We find that our aid estimate does not change much compared to Column (III). Notice that this is expected when population size does not affect growth, a prediction many standard growth models make. Thus, we conclude that there is a positive and statistically significant relationship between aid and growth but this relationship does not change as policy levels change. Clearly, the aid normalization choice affects the result.

In order to get better sense of the data, it is useful to analyze Figure 4, which shows a scatter plot

Table 11: Aid and Growth (Burnside-Dollar Regression)

	Aid-to-GDP		Aid (log)		
	(I)	(II)	(III)	(IV)	(V)
Aid	0.08 (0.12)	0.15 (0.13)	0.38** (0.16)	0.50* (0.26)	0.33* (0.19)
Aid*policy	0.05 (0.06)			-0.08 (0.12)	
Log initial GDP per capita	0.02 (0.43)	-0.06 (0.47)	0.27 (0.51)	0.29 (0.51)	0.29 (0.50)
Ethnic	0.03 (0.70)	0.13 (0.74)	-0.28 (0.76)	-0.28 (0.76)	-0.51 (0.80)
Assassinations	-0.37 (0.25)	-0.38 (0.27)	-0.39 (0.26)	-0.39 (0.27)	-0.41 (0.26)
Ethnic*Assassinations	0.14 (0.65)	0.19 (0.67)	0.17 (0.66)	0.16 (0.66)	0.19 (0.67)
Sub-Saharan Africa	-1.47** (0.65)	-1.41** (0.66)	-0.94 (0.69)	-0.91 (0.70)	-0.76 (0.67)
Fast-growing E. Asia	1.58*** (0.52)	1.36** (0.53)	1.05* (0.54)	1.08** (0.54)	0.92* (0.55)
Institutional Quality	0.26** (0.13)	0.22* (0.13)	0.24* (0.13)	0.22* (0.13)	0.23* (0.13)
M2/GDP lagged	0.00 (0.01)	0.01 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
Policy	0.89*** (0.19)	1.09*** (0.15)	1.10*** (0.15)	2.65 (2.31)	1.12*** (0.16)
Log total Population					0.14 (0.18)
Constant	0.44 (3.10)	0.88 (3.41)	-8.25 (5.69)	-10.54 (6.81)	-9.66* (5.39)
N	356	336	336	336	336
R-squared	0.32	0.32	0.33	0.33	0.33
F statistic	15.84	15.68	15.67	14.58	14.94

Dependent variable is per capita GDP growth (PPP adjusted). Robust standard error reported in parenthesis. All regressions include period dummy variables, which are not reported. Column (I) uses the same sample as the “All developing countries including outliers” in Easterly, Levine, and Roodman (2004). All other columns use a reduced sample that removes observations with negative aid values. Aid is measured in PPP adjusted USD. Significance levels : * : 10 ** : 5 percent *** : 1 percent. Data Source: Easterly, Levine, and Roodman (2004).

of the data used in the regression reported in Table 11. Income per capita in 1970 is plotted on the horizontal axis and average growth between 1970 and 1997 on the vertical one. The dashed straight lines drawn at the mean level of income per capita and growth respectively divide the scatter plot in four quadrants. First, we observe that this plot for unconditional convergence does not show any evidence of convergence among aid recipients. If anything, these countries are diverging, which




Figure 4: Test for Unconditional Convergence (BD)

Table 12: Descriptive Statistics by Quadrant (BD)

	Q-I	Q-II	Q-III	Q-IV
Number of Countries	14	15	18	15
Share of Population	14%	7%	14%	65%
Share of Total Aid	16%	11%	22%	51%
Avg. Aid per Capita	22.7 USD	20.1 USD	21.4 USD	19.8 USD
Avg. Aid-to-GDP Ratio	0.8%	1.0%	2.8%	1.6%
Policy	1.6	1.2	1.1	1.7
Institutional Quality	5.3	4.3	3.9	4.0
Sub-Saharan Africa	7.1%	20.0%	83.3%	26.7%
Fast growing E. Asia	14.3%	0.0%	5.6%	13.3%

Data Source: Easterly, Levine, and Roodman (2004).

can be easily verified by a regression that shows a positive but not significant coefficient for initial income.

Table 12 produces descriptive statistics for the four quadrants displayed in Figure 4. We observe that the number of countries are fairly evenly distributed across quadrants, with the low-income-low-growth quadrant (III) having the highest number of countries. The population, in contrast is very unevenly distributed among the four quadrants with the low-income-high-growth quadrant (IV), with 65%, having the largest share by far. In contrast, quadrant III has mostly Sub-Saharan African countries (83%) that have small populations, which explains why this quadrant only accounts for 14% of the population among the aid receiving countries. The distribution of

total aid reflects the distribution of the population as the quadrant with the highest population receives the largest amount of aid (51%). However, when we calculate the average aid per capita between 1970 and 1997, we see that this number is quite similar across quadrants with the low-income-high-growth quadrant (IV) with 19.8 PPP adjusted USD per person receiving the lowest amount and the high-income-high-growth quadrant (I) with 22.7 USD receiving the largest amount. But the aid-to-GDP ratio, in contrast, shows stark differences: Among the low-income quadrants, the high-growth quadrant (IV) has an average aid-to-GDP ratio of 1.62%, whereas the low-growth quadrant (III) has a ratio of 2.75%. This difference is related to the difference in growth rates, which deflates the aid measure for high-growth countries and inflates this measure for low-growth countries. Notice also that the countries in the high growth quadrants received 67% of total aid, whereas the countries in the low-growth quadrants received 33%. This evidence fails to reject the hypothesis that aid is good for growth. We, however, observe that high-growth countries tend to have higher policy levels and better institutional quality, which also may explain part of the better growth performance. If aid is given policy selectively, then a simple correlation between aid and growth will overstate the impact of aid on growth.¹⁵ It seems clear to us that in the empirical framework used by BD it will be difficult to tease out aid effectiveness conditional on policy, when aid is given conditional on policy.

In conclusion, differences in the aid-to-GDP ratios would suggest a negative bias as this measure inflates the aid measure for low-growth countries and deflates it for high-growth countries. That the BD regressions do not find a positive relationship between aid and growth, then, is not that surprising. In contrast, when using total aid we find a positive and significant correlation between aid and growth after controlling for policy, institutional quality, and fast growing East Asia.

Another influential aid effectiveness study is Rajan and Subramanian (2008) (RS). Given the inconclusive results concerning the relationship between aid and growth in the existing literature, RS propose to start fresh and to examine this question in a comprehensive manner. This includes the development of an instrument that exploits the dyadic structure of the colonial past of most developing countries to control for potential endogeneity related to aid. In our replication, we focus

¹⁵Notice that controlling for policy may understate the impact of aid on growth if policy-selective aid leads to better policies in aid receiving countries.

Table 13: RS Cross-Section Regressions

	1961-2000				1991-2000			
	Aid-to-GDP		Total Aid		Aid-to-GDP		Total Aid	
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Aid	-0.06** (0.03)	0.06 (0.06)	-0.16 (0.12)	-0.19 (0.19)	-0.01 (0.05)	-0.40** (0.18)	0.25 (0.15)	0.73** (0.33)
GDPpc	-1.33*** (0.28)	-1.17*** (0.35)	-1.33*** (0.32)	-1.34*** (0.29)	-1.15** (0.56)	-2.23*** (0.63)	-0.90* (0.52)	-0.38 (0.58)
Policy	1.79*** (0.43)	1.62*** (0.60)	1.68*** (0.55)	1.68*** (0.50)	-0.16 (0.55)	-0.06 (0.66)	-0.19 (0.53)	-0.20 (0.49)
Life Exp.	0.02 (0.02)	0.06** (0.03)	0.03 (0.02)	0.03 (0.03)	0.15** (0.06)	0.04 (0.08)	0.16*** (0.06)	0.16*** (0.05)
Geogr.	0.35** (0.13)	0.53*** (0.17)	0.47*** (0.15)	0.47*** (0.14)	0.69* (0.41)	0.19 (0.38)	0.53 (0.41)	0.30 (0.41)
I.Qual.	3.94** (1.49)	4.56*** (1.53)	4.47*** (1.48)	4.50*** (1.41)	2.95 (3.15)	6.57** (3.23)	3.15 (2.96)	3.49 (2.75)
Infl.	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00*** (0.00)	-0.00** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)
M2	0.02 (0.01)	0.02 (0.01)	0.02 (0.01)	0.02* (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.01 (0.01)
B.Bal.	-0.01 (0.02)	0.02 (0.03)	-0.00 (0.02)	-0.00 (0.02)	0.20*** (0.06)	0.19** (0.09)	0.23*** (0.06)	0.26*** (0.05)
Revol.	-1.26** (0.51)	-1.14** (0.56)	-0.98* (0.53)	-0.94* (0.53)	-0.49 (0.65)	-0.34 (0.71)	-0.56 (0.63)	-0.74 (0.59)
Eth.Fr.	-0.10 (0.45)	0.72 (0.55)	0.29 (0.43)	0.28 (0.39)	1.74 (1.08)	-0.16 (1.27)	1.71 (1.03)	1.32 (0.98)
Constant	8.56*** (1.61)	4.65* (2.71)	8.76*** (2.27)	9.09*** (2.89)	-0.02 (4.00)	15.53** (7.13)	-4.20 (4.13)	-11.80** (5.78)
N	74	74	74	74	70	70	69	69
R-squared	0.77	0.66	0.75	0.75	0.63	0.35	0.63	0.60
F statistic	23.46		15.06		11.32		10.75	

This table reports replication results for Table 2 (OLS) and Table 4 (2SLS) in RS in the Columns labeled "Aid-to-GDP." The Columns labeled "Total Aid" report the results when using the log of total aid (PPP adjusted) instead of Aid-to-GDP. Dependent variable is per capita GDP growth (PPP adjusted). Robust standard error reported in parenthesis. All controls measured at initial level except for Geography, Institutional Quality, Revolutions, and Ethnic Fractionalization. Controls for sub-Saharan Africa and East Asia Pacific are included but not reported. Significance levels : * : 10 ** : 5 percent *** : 1 percent. Data Source: Replication File AidData.

on the cross-country regressions, where their instrumental variable approach arguably has most bite.¹⁶ We replicate their regression with the longest and shortest time span covering the years 1961–2000 and 1991–2000 respectively. The results are shown in Table 13.

For the long-run regressions covering the years between 1961 and 2000 we find a negative correlation between aid and growth when using the log of total aid, although this relationship is statistically not significant, both using OLS and 2SLS. RS find a negative and statistically significant relationship between aid and growth when using OLS. The scatter plot printed in their article and then re-reprinted in the blog post mentioned earlier on how aid hurts the poor was drawn from this regression. However, this relationship changes to a positive but statistically insignificant one, when using the colonial past of recipient countries as an instrument for aid. We observe that estimates change quite a bit when using aid-to-GDP ratios across estimation techniques but remain robust when using the log of total aid. For the shorter time span between 1991 and 2000, the aid normalization choice makes a difference. In the 2SLS regressions, the aid-coefficient is negative and

¹⁶In their panel analysis, this time-invariant instrument is of limited use (see Clemens, Radelet, Bhavnani, and Bazzi, 2012)

statistically significant when using aid-to-GDP and it is positive and statistically significant when using the log of total aid.¹⁷ The same applies to the OLS results although these coefficients are not statistically significant (the p -value for the aid coefficient in the OLS regression with total aid is just a bit over 10%).

¹⁷Note that here we are not able to exactly replicate the 2SLS results in RS for this time span, even though the result is very close. For example, the aid coefficient reported in Table 4 in RS equals -0.389 whereas ours is -0.40. There are also small differences for other coefficients.