

The Effects of Foreign Aid in Sub-Saharan Africa

Robert Gillanders*
rgillanders@gmail.com

UCD School of Economics

September 28, 2010

Abstract

This paper revisits the aid effectiveness debate. A vector autoregression model is applied to a panel of Sub-Saharan African countries. This method allows one to analyse the impact of foreign aid on human development and economic development simultaneously. The results on the full sample indicate a small increase in economic growth following a fairly substantial aid shock. The size of the effect puts the result somewhere between the arguments of aid optimists and those of aid pessimists. Human development, as proxied by the growth rate of life expectancy, shows a small but positive response to aid shocks. Economic growth is found to respond more to aid shocks in groups defined by better economic policies, poor institutions and high aid dependence. Human development responds more to aid shocks in democracies and good institutional environments.

Keywords: Foreign Aid; Economic Development; Human Development; PVAR
JEL: F35, O11

*UCD School of Economics, Newman Building (Room G215), University College Dublin, Belfield, Dublin 4, Ireland. This work, which forms part of my Ph.D. thesis, was funded by the Irish Research Council for the Humanities and Social Sciences. I am indebted to both my supervisor, Karl Whelan, and Inessa Love, who kindly made her PVAR Stata program available to me. I also wish to thank all the participants at both the 2010 EDGE Jamboree and the 2010 ISNE Conference (Anna Gumpert and Markus Eberhardt in particular) for helpful comments and suggestions.

1 Introduction

For better or for worse, Official Development Assistance (ODA) is the main tool employed by the rich world in its attempts to promote prosperity in the developing world. Given the importance of this tool, it should not be a surprise that many scholars of development consider its impacts worthy of study. Taken as a whole, the literature on aid is vast. Many studies look at aid's effect on growth. Others look at the effect on human development, as proxied by variables such as the Human Development Index (HDI) or the infant mortality rate. The problem is that the aid literature largely focuses on one issue at a time.

Why is it important to look at both elements? We all know that while GDP per capita is a good measure of the overall state of a society (due to its high correlations with things we actually do care about), it is far from perfect. Amongst other factors, it ignores distributional issues, non-market activities and environmental quality. Economic development is surely a large part of the concept of "development" but there are other elements that matter. Other social sciences are at pains to make this point to economists, but I would argue that we already know this. Certainly, the small but growing empirical literature on aid and human development (which is outlined below) shows that economists are aware of the multi-dimensional nature of development.¹ The task then, is to examine both (broadly defined) elements of development simultaneously.

A panel vector autoregression model (PVAR) allows one to do just that. Another merit of this notoriously atheoretical approach is that, as both Easterly [2003] and Easterly, Levine, and Roodman [2004] point out, much of this literature suffers from a lack of a formal theory linking aid and growth to guide applied researchers in settling on an econometric specification. It also avoids the need for instrumentation strategies which have been criticised by some.²

Interest lies specifically with the impulse response functions (IRFs) obtained from the model. These will show the response path of economic growth and human development to a one time shock in foreign aid holding all other shocks at zero. To deal with concerns about biases that could arise from estimating the model without allowing for unobserved heterogeneity (i.e. fixed effects), the sample is restricted to countries in Sub-Saharan Africa. These IRFs serve as the baseline results and it is shown that these results are in line with GMM estimates which allow for country specific fixed effects. The sample is then split into groups defined by economic policy, institutional quality, democracy and aid dependence. This allows us to see if aid is more effective in certain environments.

¹Not to mention the work of Amartya Sen.

²See Deaton [2010] for a recent exploration of the potential problems.

Several important and policy relevant results emerge from the analysis. Firstly, aid shocks do seem to induce economic growth within the first few years following the shock. However, this response is less than 1% and is tempered by some later negative responses. Secondly, aid appears to have a small positive impact on human development, though the error bands do not allow us to rule out that the effect is negative or non-existent. In countries with good economic policy, the early impact of aid on economic growth jumps to about 2%, though this is once again tempered by some negative responses later on. In democracies, aid has a much larger impact on human development and a less pronounced effect on economic growth than in either the full sample or in autocracies. Relative to countries with good institutions, those with poor institutional environments make more use of aid in the economic sphere and less use of aid in terms of human development. Finally, countries which are more dependent on aid see a better response in terms of both economic and human development.

One of the main implications of these findings is that while aid does lead to economic growth, the impact is not transformative and to maintain even an extra 1% of growth would require frequent large injections of aid. Another is that while the response of GDP per capita growth is higher in some environments, it is not so much higher that donors could justify focusing all aid monies on countries with these characteristics. This is further supported by the fact that often it is the case that when aid fails to have an impact on one of the dimensions of development in one of the sub-samples, it impacts on the other.

The remainder of the paper proceeds as follows. Section 2 outlines the existing literature. Section 3 outlines the data and discusses issues of measurement. Section 4 describes the econometric approach in detail. Sections 5 and 6 present the baseline and grouped results respectively and Section 7 concludes.

2 The Existing Literature

There is an extensive literature that seeks to examine the impact of foreign aid on growth and a smaller one that looks at aid's impact on human development. Some studies have found positive effects, some conditional effects and others no effects. Given the wide range of findings and the debates that they have prompted, it is worth examining some of the relevant recent papers in brief.

2.1 Aid and Economic Development

There are many studies that have examined the effect of aid on growth. I limit myself here to the more recent work, both for brevity and because it adequately highlights the issues that will be examined in Section 6. Recent work has chiefly focused on conditional aid effectiveness.³

Burnside and Dollar [2000] reignited the aid effectiveness debate when they found that while aid has no effect on growth on average, aid works in a good policy environment. They include an aid*policy interaction term and find that it is statistically significant and robust to a number of specifications. This paper launched the debate on conditional aid effectiveness. Easterly, Levine, and Roodman [2004] recreate the Burnside-Dollar dataset and expand on it significantly. By following the approach of Burnside and Dollar to the letter, they find that not only is the crucial aid*policy coefficient insignificant but it has the opposite sign.

Easterly [2003] re-examines the issue in a different way. By employing Official Development Assistance (ODA) as the measure of aid as opposed to the measure used by Burnside and Dollar, Effective Development Assistance (EDA), Easterly finds that the aid*policy interaction term is no longer significant.⁴ He also varies the specification of good policy and again finds that the aid*policy term is insignificant. The crucial interaction term is also found to be insignificant by varying the definition of growth (Burnside and Dollar defined growth as real GDP growth over four years) to consider eight, twelve and twenty four years.

One of the best examples of support for the Burnside-Dollar result comes from Collier and Dollar [2002]. They expand the Burnside-Dollar dataset to include 349 aid-growth-policy episodes as opposed to the original 275. The other modification they make is to employ the World Bank's Country Policy and Institutional Assessment (CPIA) as their policy variable. Their findings agree with the Burnside-Dollar result.

However, Dalgaard, Hansen, and Tarp [2004] question the suitability of the CPIA for growth regressions. They include a climate*aid interaction term and find that it is negative and significant. The aid*policy term loses its significance once climate*aid enters the specification. They suggest that the climate variable may be a proxy for deep determinants such as institutions. This is a significantly different conditional effectiveness result and one which will be examined in Section 6. Given the sketch of the literature above, it will come as no surprise that the role good policy has to play will also be investigated.

Svensson [1999] examines whether civil and political liberties play any role in aid effectiveness. He finds that an aid*democracy interaction term is highly significant and that an

³Roodman [2007] provides an excellent overview of the entire literature.

⁴As many authors note, the two measures are highly correlated. The two are fully defined in the data section below.

aid*policy term à la Burnside and Dollar is insignificant. It is important to note that he considers, and rejects, the possibility that democracy is simply a proxy for good policies. This suggests we should add democracy to our list of conditional results to be examined.

This sample of the aid and economic growth literature is chiefly included to motivate the division of the sample into groups defined by economic policy, institutions and democracy in Section 6. However, it also shows that the traditional approach of running either standard cross-sectional or panel growth regressions augmented with aid and aid interaction terms leads to fragile results. Indeed, the title of Roodman [2007], “*The Anarchy of Numbers: Aid, Development, and Cross-Country Empirics*”, sums up the literature well. Each of the papers is a fine econometric work, and it may be that these factors do indeed matter, but the fragility of the results is undeniable.⁵ This is another motivating factor in employing the PVAR methodology – it offers a fresh perspective.

2.2 Aid and Human Development

While nowhere near as extensive as that which concerns itself with economic development, there is a small but growing literature that seeks to empirically assess the impact of foreign aid on human development.

Kosack [2003] finds that aid has a positive effect on HDI growth but only in democratic countries. His estimates also suggest that aid will have a negative effect on HDI growth in autocracies. Interestingly, he finds that democracy alone has a negative effect on HDI growth. He interprets these findings as implying that ‘more-democratic poor countries have, on their own, lower growth in the quality of life, but that aid to these countries may reverse this negative tendency’ (p6).

McGillivray and Noorbakhsh [2007] examine the impact of aid on the level of the HDI and allow conflict to enter the analysis. They find that aid alone has a negative impact on HDI scores but disagree with Kosack [2003] in that they do not find either a negative effect of democracy on the HDI or a positive aid*democracy interaction term.⁶ These two studies gives us a second reason to divide the sample along lines of democracy.

Using quantile regression, Gomane, Girma, and Morrissey [2005a] examine aid’s effects on human development as measured by both the HDI and the infant mortality rate. They argue that while aid might not have a direct impact on welfare, it may have an indirect one via pro-poor expenditures (PPE). By constructing a PPE index they find that aid has a positive

⁵Roodman examines seven leading papers in the aid-growth literature, (including Burnside and Dollar [2000], Collier and Dollar [2002] and Dalgaard, Hansen, and Tarp [2004]) and finds that each of them is susceptible to changes in the sample and in specification.

⁶Rather they find that aid and conflict have negative effects on human development and that aid is no more or less effective in conflict situations.

impact on welfare through public expenditure and that the effect is greater in countries with lower welfare. They also argue that good economic policy is not required for aid to be effective in promoting human development. A related paper, Gomanee, Morrissey, Mosley, and Verschoor [2005b], finds to the contrary that there is a direct impact of aid on human development and little evidence of an indirect effect via PPE. It is clear that there is nearly as much disagreement and tendency for conflicting results in the aid-human development literature as there is in the aid-growth literature.

There are of course more works that concern themselves with the effects of aid on economic and human development. However, the papers above provide a sufficient overview of the evidence and suggest that it may be illuminating to examine whether aid has different effects in groups of countries defined by economic policy, institutional quality and democracy.

Another division of the sample is inspired by Hansen and Headey [2009]. They utilise a PVAR model to examine the short-run macro effects of aid. They split the countries in their sample into those that are highly dependent on ODA and those that are not. In the context of the current work, it is plausible that there is some critical level of aid dependence under which aid cannot be effective. It is equally plausible that being overly dependent on aid results in macro distortions or poor governance.

3 Data

The data on yearly economic growth comes from the World Bank's *World Development Indicators* (WDI).⁷ While economic growth is a (relatively) concrete concept, the other two variables that will be used in this analysis are somewhat nebulous. Thus, it is good practice to devote a little time to defining exactly what the "aid" and "human development" variables used in the analysis actually are.

3.1 Measuring Aid

This paper uses the notation aid_{it} to represent total net ODA per capita. This includes flows from all donors (as measured by the OECD) to recipient i at time t . aid_{it} is measured in constant 2007 US dollars.⁸ The data come from the OECD Development Assistance Committee (DAC).⁹ In choosing how to measure aid, one must make two important decisions.

⁷The WDI can be accessed at <http://data.worldbank.org/indicator>

⁸Given that ODA contains concessional loans, these net flows can be negative if the repayment of loans are greater than new loans and grants.

⁹The data can be accessed at www.oecd.org/dac/stats/data

Firstly, **what counts as foreign aid?** This effectively boils down to the choice between ODA and EDA. ODA is defined on the OECD's website as:

“Grants or loans to countries and territories on the DAC List of ODA Recipients (developing countries) and to multilateral agencies which are: (a) undertaken by the official sector; (b) with promotion of economic development and welfare as the main objective; (c) at concessional financial terms (if a loan, having a grant element of at least 25 per cent). In addition to financial flows, technical co-operation is included in aid. Grants, loans and credits for military purposes are excluded. Transfer payments to private individuals (e.g. pensions, reparations or insurance payouts) are in general not counted.”¹⁰

EDA is an adjustment to ODA that replaces the official loans component with “the grant equivalents of official loans” and disregards grants that are tied to technical assistance.¹¹ Although this may seem like an important modification, the two are hugely correlated.¹² I choose to use ODA as the aim of the paper is to assess the impact of the West's efforts and ODA is the tool they employ.

The other issue one must decide on is **how to normalise the aid variable**. Many studies express aid as a fraction of GDP. In my case this would lead to a mechanical relationship between two of my main variables of interest - aid and economic growth. To avoid this, I divide yearly aid data by yearly population data from the Penn World Tables. Thus, aid_{it} is fully defined as total net ODA per capita.

3.2 Measuring Human Development

One also faces two issues when attempting to quantify human development (HD_{it}). The first is **finding an ideologically neutral measure**. While some of us would consider low levels of poverty or inequality to be good indicators of human development, others would disagree. Claiming either of these variables as good proxies for human development requires one to take an implicit ideological stance.¹³ Using democracy or measures of government accountability is similarly unsatisfactory.

The second problem is **the lack of sufficiently long time series data relating to human development**. The UN's Human Development Index (HDI) is equal parts economic

¹⁰See www.oecd.org/dac/glossary

¹¹Chang, Fernandez-Arias, and Serven [1999].

¹²Dalgaard and Hansen [2001] show that the correlation between the two (as a fraction of GDP) in nominal terms is 0.98 and between real EDA and nominal ODA it is 0.89. They also suggest that the difference between the two measures “seems to be a simple transformation” (p26).

¹³Of course, if one set out to examine the impact of ODA on inequality or poverty then this issue of an ideological bias does not arise.

prosperity, education and life expectancy and is largely free of ideology.¹⁴ However, the series is not of sufficient length to be usable in an analysis such as this.

I consider life expectancy to be an ideologically clean proxy for human development. While there are other aspects to human development, a long lifespan is essential to pursue many of them. The data series, which again comes from the World Bank's WDI, is of sufficient length and covers a sufficient number of countries. The specific variable, denoted as HD_{it} , is the growth rate of total life expectancy at birth (in years).¹⁵

The data allows for a balanced panel of 31 countries over the period 1973-2005. The main reason for the omission of 17 countries is GDP growth series that start too late. Some of these data poor countries have very small populations and others had particularly turbulent starts to their independence. The list of countries that are used in the analysis, along with summary statistics, can be seen in Appendix A.

4 Econometric Approach

This paper is concerned with the IRFs associated with the following PVAR model:

$$Z_{it} = \alpha_0 + \sum_{q=1}^p \beta_q Z_{it-q} + \epsilon_{it} \quad (1)$$

where Z_{it} is the vector $(aid_{it}, growth_{it}, HD_{it})'$, α_0 is a vector containing the constant terms, β_q is the matrix of coefficients for lag q and ϵ_{it} is distributed as $(0, \sigma_i^2)$ with $E(\epsilon_{it}\epsilon'_{is}) = 0$ for all $t \neq s$.

This model can be viewed as the most restrictive possible in that it imposes common slopes and common intercepts and can be estimated by Pooled Ordinary Least Squares (POLS). It is well known that POLS yields biased and inconsistent estimates if the true data generating process contains a fixed effect. To reduce such concerns, only countries in Sub-Saharan Africa are considered. All countries are of course different, but these countries should form a sufficiently homogeneous group to allow POLS estimates of the relationships between ODA, economic growth and human development to be meaningful. Such concerns should be further reduced by the division of the sample into the subgroups outlined above.

Less restrictive models have been estimated. The next most restrictive model allows for heterogeneity in the relationships by allowing for country specific intercepts. This modification leads to a dynamic panel data model with country fixed effects. Such a model can be

¹⁴Though it is somewhat arbitrary in that it gives equal weight to each component. For more detail see United Nations Development Programme [various years].

¹⁵I use the growth rate to minimize concerns about stationarity.

estimated via GMM techniques as is done by Love and Zicchino [2006]. This approach will be used as a check on the main results.

An even less restrictive variant of equation (1) allows for both country specific intercepts and slopes. Pesaran and Smith [1995] show that these sort of models can be estimated by applying the Mean Group Estimator (MGE). This is the method adopted by Hansen and Headey [2009]. However, Rebucci [2003] shows, using Monte Carlo simulations, that standard Fixed Effects estimation performs better than MGE in a PVAR with country specific slopes and intercepts as long as slope heterogeneity is not “relatively high”. In addition, the simulations indicate that for MGE to be a useful alternative, the time dimension of the panel must be longer than that which is to be found in most macroeconomic datasets.

As it is not the reduced form disturbances of equation (1) that we are interested in, we need a method to recover the pure structural shocks. The most common way to obtain orthogonalised shocks is to employ a Choleski decomposition which orders the shocks in a sensible way. The channel of influence I impose is the following:

$$aid_{it} \rightarrow growth_{it} \rightarrow HD_{it}$$

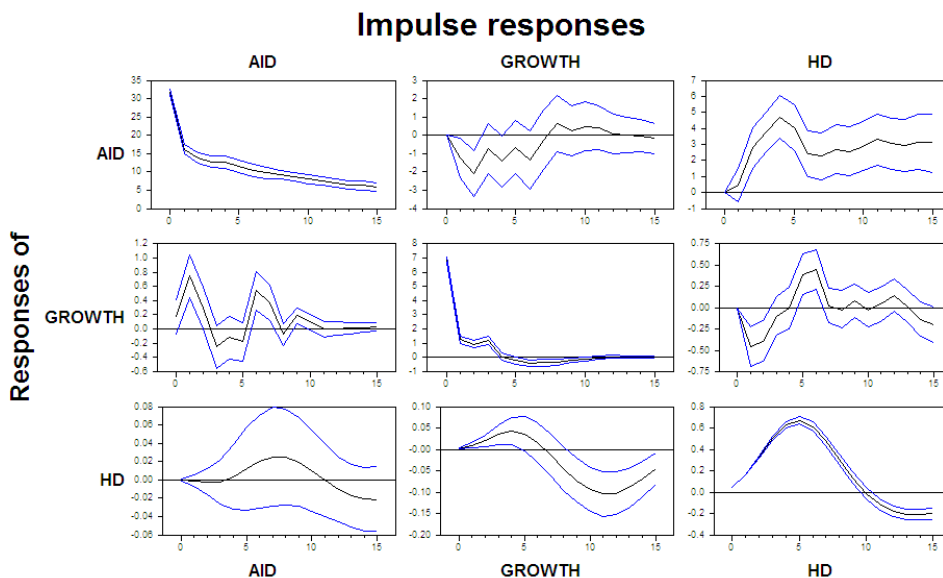
Aid shocks can impact on growth and human development contemporaneously; growth can impact on human development contemporaneously but only influences aid flows with a lag; human development only operates on the other two variables with a lag. The choice of ordering is a crucial factor in any VAR exercise.¹⁶ The rationale behind the chosen ordering is as follows:

- Donor countries (especially their bureaucrats) need time to observe and react to changes in the recipient country. Thus aid flows will only respond to changes in economic and human development with a lag.
- It takes time for increasing human development to translate into economic development. It is much more plausible that economic growth can have a contemporaneous effect on human development.

While this is only one of six possible orderings, it seems the most plausible and the sensitivity of the results to alternative Choleski decompositions is examined in Appendix B. To interpret any econometric results correctly, one needs some idea of how accurate they are likely to be. To this end, 95% confidence intervals are constructed using standard Monte Carlo methods. Finally, to determine the appropriate number of lags to include, that is the p in equation (1), I employ standard information criteria.

¹⁶The importance of the ordering increases with the correlation between the reduced form innovations. See Enders [2004] page 276.

Figure 1: Full Sample IRFs: POLS



Notes: Each IRF depicts the response of the row variable to a one standard error shock in the estimate of the column variable. The upper and lower lines are 95% error bands generated using Monte Carlo simulation with 2500 draws.

5 Full Sample Results

Both the Akaike information criterion and the Schwarz-Bayesian information criterion indicate that the model should include 8 lags. I will maintain this lag structure in what follows to make each set of results fully comparable.

5.1 POLS

The IRFs obtained from the full sample of 31 countries can be seen in Figure 1. For completeness I present all the IRFs from the system though interest only really lies with the responses of the growth rates of GDP per capita and life expectancy to a shock in aid. Each IRF shows the response path of the variable in question to a one time positive one standard error increase in the estimate of the shock variable holding all other shocks at zero. The values of the shocks can be read off the Y-axis of the diagonal elements of Figure 1 at time zero. The aid shock is approximately \$33 per person.

It is, however, reassuring that the none of the other IRFs display inexplicable behaviour. Aid tends to decline following an economic growth shock and tends to increase following a shock to the growth of human development. It is plausible that when donors see an increase in economic growth they scale back aid but when they see some improvements in human development they attempt to consolidate these gains. Economic growth is slightly lower for a few periods following a shock to human development growth – perhaps due to people switching from economic to leisure or other non-market activities – though the effect quickly dies out. Finally, a shock to growth increases the growth rate of human development until about seven periods after the shock, at which point the growth of human development is lower than it would have been. This is perhaps due to increased pollution or other retarding influences that can arise from increased economic activity.

Turning to our main concerns, we can see that an aid shock at time zero has an immediate, but small, positive impact on GDP per capita growth. This positive impact increases to about 0.8% extra growth one period after the shock and the positive impact still exists in period two. These short run increases in growth can be explained as increases in government investment and perhaps consumption. The next three periods see growth that is lower than it would have been in the absence of the initial aid shock. A plausible explanation for this is that recipient governments treat the aid shock at time zero as a permanent increase and are taken by surprise when it proves to be transitory.¹⁷ After this we see a return to a positive response. This would fit with government investments on infrastructure (made with the initial aid shock) paying off. After the 10th step the initial shock has no more influence. The error bands are generally tight enough for us to have confidence in the estimated response path.

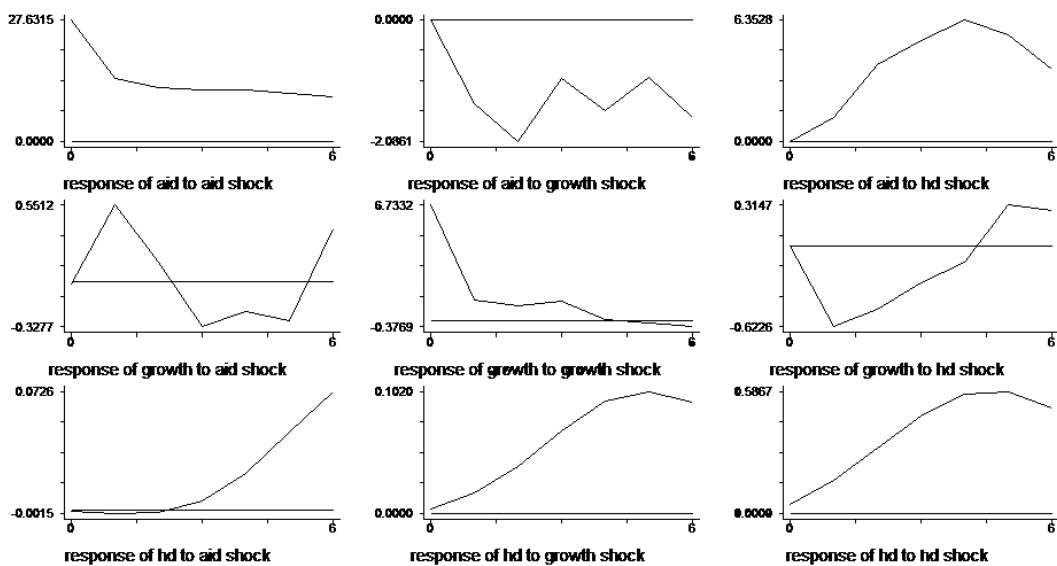
So is aid effective in terms of generating economic growth? Given the results of Figure 1, one would have to say yes. The positive responses to aid seem to outweigh the periods where growth is lower than it would have been. However, there are two important issues one should bear in mind:

1. The extra injection of \$33 per person is not cheap. Even at this level, the results don't fit the description of transformative.
2. Given that the effects die out by the 11th period, can we say that the induced growth is sustainable?

These results thus sit somewhere between the arguments of aid optimists and aid pessimists. Yes aid seems to be effective in that it does induce *some* economic development, but to make it a promising tool for eliminating global poverty it would need frequent large injections.

¹⁷They may not be entirely wrong in this as Figure 1 shows that the response of aid to its own shocks persists beyond fifteen periods, though it does drop rapidly to roughly a third of the value of the shock.

Figure 2: Full Sample IRFs: GMM



Notes: Each panel shows the response of the indicated response variable to a one standard deviation shock in the shock variable.

Results obtained using the PVAR program kindly made available to me by Inessa Love.

So are the results in terms of human development more supportive of aid advocates? Not really. The estimated response path is positive for the most part but very small in magnitude. That said, the human life span is measured on a (sadly) small scale so an increase in its yearly growth rate of even a fraction of a percent may be quite good, especially given that the response persists until the 10th period following the initial shock. However, the error bands allow for the possibility that the impact may be negative. The results in this respect are thus ambiguous.

5.2 Robustness: GMM

Before moving on to examine the responses in various sub-samples, it is important to check to see if the baseline results are robust to allowing for country specific effects i.e. if the POLS approach gives different results to a model that allows for fixed effects. To do this I employ the PVAR Stata program written by Inessa Love. For specific details of the procedure, I direct the reader to Love and Zicchino [2006]. Briefly, the estimator allows for country specific intercepts which are removed by forward mean differencing. The model is then estimated using GMM methods.

The results are presented in Figure 2. The IRFs only trace out the responses to six periods after the initial shock and use a slightly different value for the shocks, but we can immediately see that the responses to an aid shock are very similar to those obtained using simple POLS, both in the pattern and in magnitude. The only difference of note is that the aid shock is smaller yet it generates a larger, though still rather small, response in the growth rate of human development. Overall we can be confident that POLS is a satisfactory method and that the results presented in Figure 1 are valid.

6 Conditional Results

To investigate whether aid works in certain environments, that is conditionally, previous authors in this literature have included interaction terms in their specifications. While one could do likewise in the VAR framework, the interpretation of shocks in these multiplicative terms would be difficult and the extra terms would have to be made to fit into a Choleski decomposition. I take a simpler alternative approach and merely run the model over groups defined by certain characteristics.

6.1 Good Versus Bad Economic Policy

Given how influential the Burnside-Dollar result has been in donor circles, it makes sense to use the model to see if we can find evidence of aid working better in a good policy environment. One lesson to be taken from the literature is that often the definition of good policy is somewhat arbitrary in what is included and once it is changed the significant results disappear. With this in mind, I have chosen the most objective measure of economic policy that I have come across. The World Bank's *Doing Business* project collects information across a wide variety of aspects of the business environment and ranks countries on the overall ease of doing business. The full methodology can be found on the project's website.¹⁸ Here, it is sufficient to note that the surveys are as objective as possible and cover most aspects of a country's business environment. Thus, the ranking should provide a good proxy for overall economic policy. This is quite a different measure of policy from the weighted indices of macroeconomic indicators that are commonly used in this literature.

I use the most recent rankings as my measure. While this is open to criticism, the data only goes back to 2004 and each year has seen additional important variables added. I believe that the data can safely be regarded as a good measure of a country's stock of economic policies. I consider a country with a ranking in the top 100 to have good economic policy. Countries below this, admittedly arbitrary, cutoff form the bad policy sample. A benefit of

¹⁸<http://www.doingbusiness.org>

dividing the sample this way is that even if we would see countries move around somewhat in the rankings over time had we the data, such broad ranges make it more likely that they remain within the bounds I have set for good and bad policy.

Figure 3 presents the IRFs obtained from running the model on each of the policy groups. The top set are for the good policy sample and the bottom set for the bad policy sample. It will come as no surprise that only five countries made the cut for good policy.¹⁹ As mentioned above, I continue to include eight lags in the model over all samples.

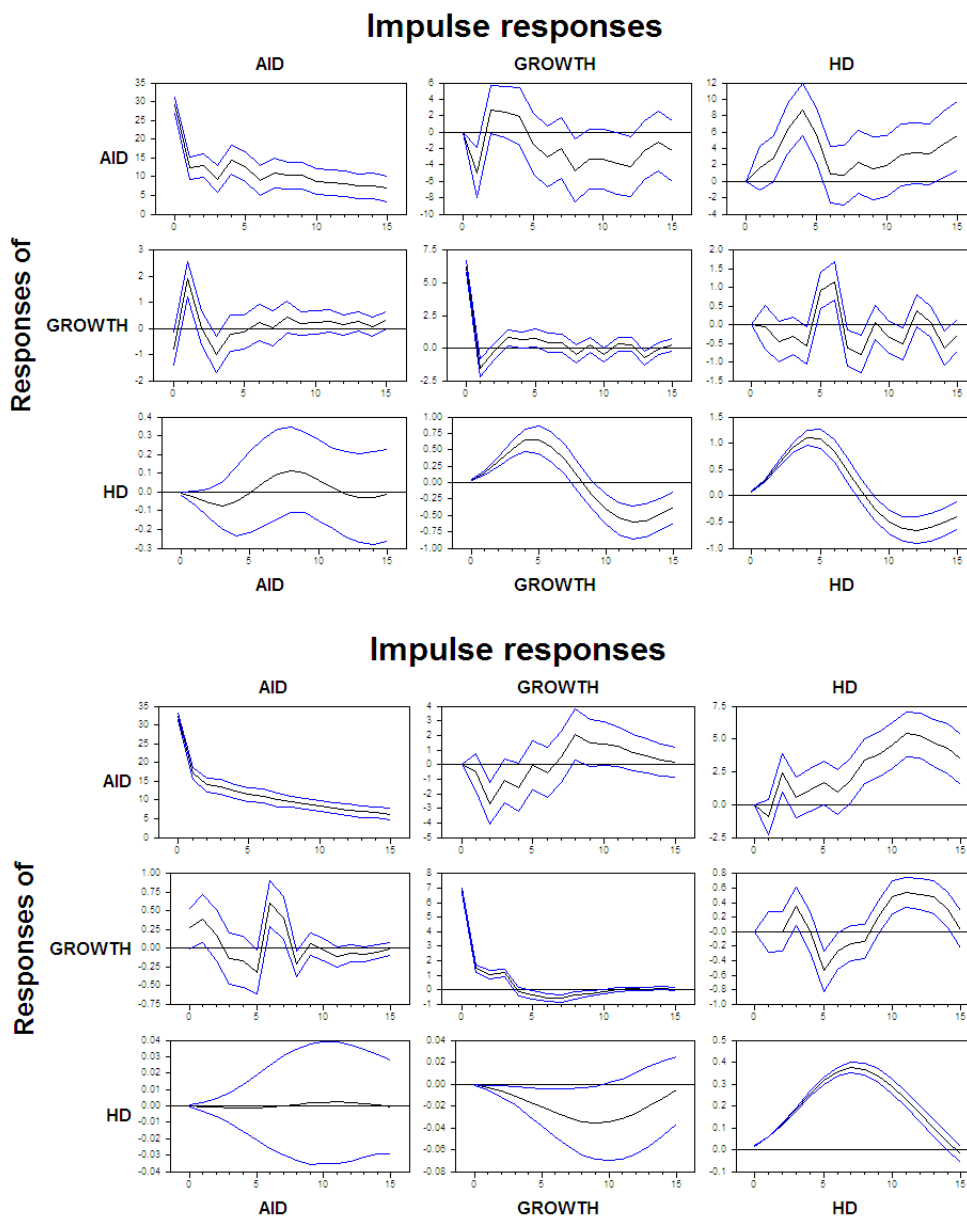
The results are not as precise as those of Section 5 but are striking nonetheless. It is worth noting that the magnitude of the aid shock in both cases are very similar. In the good policy group, economic growth is actually *lower* than it would have been at $t=0$ absent the aid shock. However, by $t=1$ growth is roughly 2% higher than it would have been. We see the same pattern of negative response that we saw in the full sample after $t=2$ with some marginal positive responses after $t=5$. While the error bands do not allow one to discount the possibility that these long lasting effects are negative, they were not present in the full sample results. The bad policy sample responses are markedly different and (not surprisingly given that these countries form the vast majority of the full sample) are similar to the full sample results. The level of response is smaller though and we see some long lasting negative responses that were absent in the full sample.

In terms of human development, the error bands are too broad to draw much inference. Certainly in the bad policy group the estimated effect is indistinguishable from zero for the most part with error bands evenly distributed on either side. In the good policy group the error bands are also too wide but we can see initial negative responses. After that we cannot say with any certainty what sign the effect takes but the positive responses seem to outweigh the negative ones.

These results suggest that, in terms of increasing GDP per capita growth, aid does work better in a good policy environment. However, I would argue that it still does not reach the level of transformative power that aid advocates claim. A 2% boost to economic growth is certainly impressive but once again it is not cheap and is tempered by some negative responses. On the other hand, the results certainly do not conform to the arguments of the most despairing of aid pessimists. Once again we see some beneficial impact of aid, even in countries with bad policy. In good policy environments we see a large early increase in economic growth alongside some long lasting, albeit quite modest, benefits.

¹⁹Appendix A shows which countries are in this and each subsequent sub-sample.

Figure 3: Good Vs. Bad Policy IRFs



Notes: Each IRF depicts the response of the row variable to a one standard error shock in the estimate of the column variable. The upper and lower lines are 95% error bands generated using Monte Carlo simulation with 2500 draws. The top set is for the sample composed of those countries which rank in the top 100 in the World Bank's ease of doing business ranking. The bottom set is for the sample composed of those countries which rank below 100 in the World Bank's ease of doing business ranking.

6.2 Democracy Vs. Autocracy

We saw in Section 2 that there is evidence that being a democracy matters for both dimensions of aid effectiveness. To examine this issue using my PVAR model, I utilise the measure of democracy created by the Polity IV project. This variable combines information on democratic and autocratic features of a country and places them on a scale from -10 (hardcore autocracy) to $+10$ (hardcore democracy). This suggests a simple division of the sample. Countries with an average Polity score (over the sample period) less than 0 are consigned to the autocratic group. Those with a score of 0 or higher form the democratic group.

Figure 4 shows the IRFs for both sub-samples. As was the case with the bad policy group, the autocracy group forms the vast majority of the full sample and it is therefore unsurprising that the IRFs are mostly very similar to the baseline results. The big difference is that we see no appreciable positive impact of aid on human development. Turning to the democratic sample, we see a relatively large response of human development to an initial aid shock. Not only is it relatively large, it is positive and persists to the 11th step. As the shocks are once again of very similar magnitudes, this result partially supports the findings of Kosack [2003] – aid seems to work in terms of promoting human development only in democratic countries but we cannot see any clear evidence of it retarding human development in autocracies.

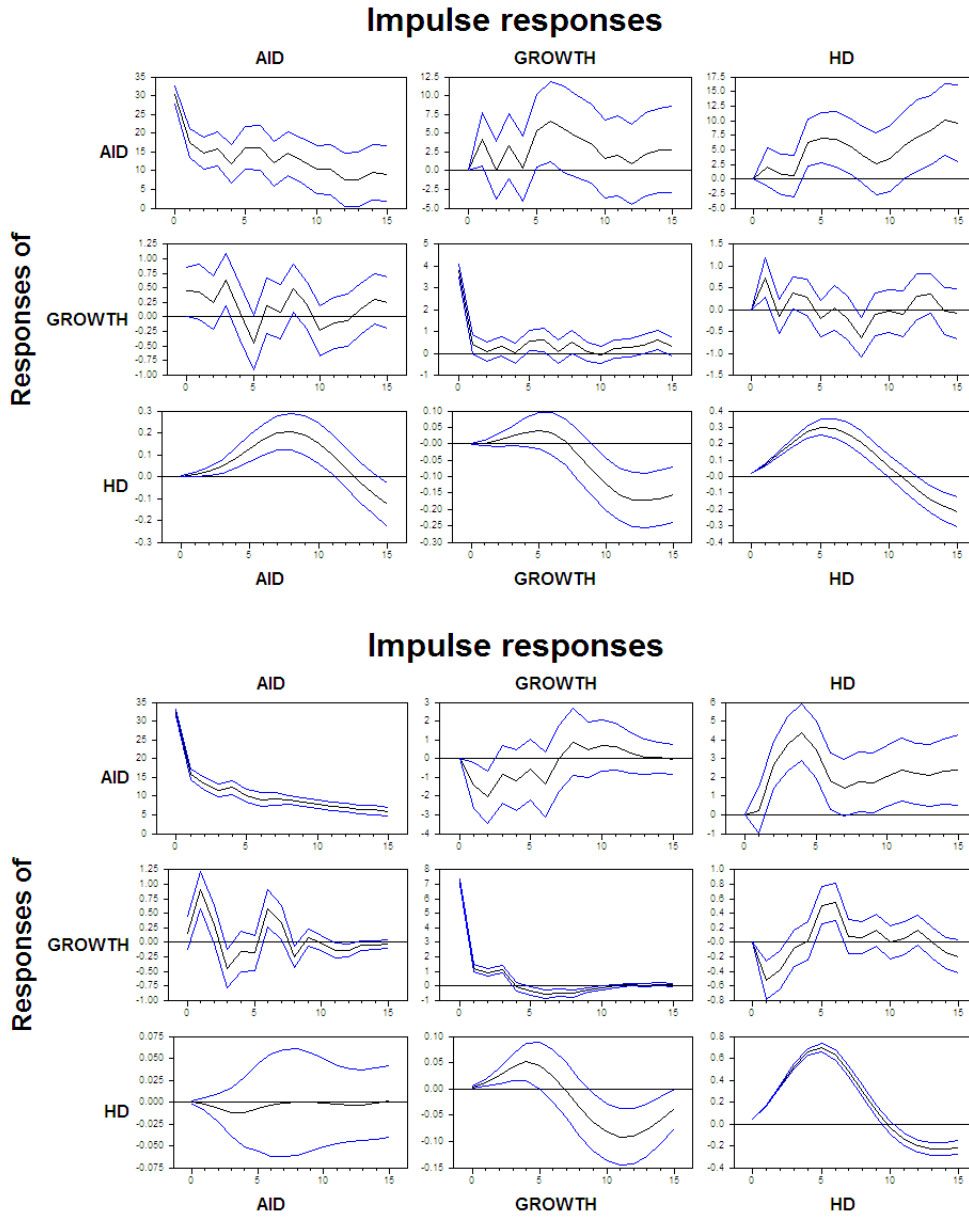
In terms of aid’s impact on economic growth, as noted above, the autocratic sample’s responses are very similar to the full sample’s. The major difference is the late negative responses that were absent in the full sample. The democratic countries response path has error bands that are too wide to draw any firm inference but the response does seem to be positive until the 5th period following the aid shock as opposed to the 3rd period in the autocratic sample. After that it is generally not possible to make any inference about the sign of the response. These results clearly do not conform with Svensson [1999]. Aid does not seem to be more effective in promoting economic growth in democratic environments.

6.3 Good Vs. Bad Institutions

The next division of the sample is inspired by Dalgaard, Hansen, and Tarp [2004], who, as we noted in section 2, argue that deep determinants of a society may play a big role in making aid effective. They point specifically to institutional quality. This has enormous intuitive appeal – good institutions should lead to a well ordered society which in turn should lead to an effective use of resources such as foreign aid.

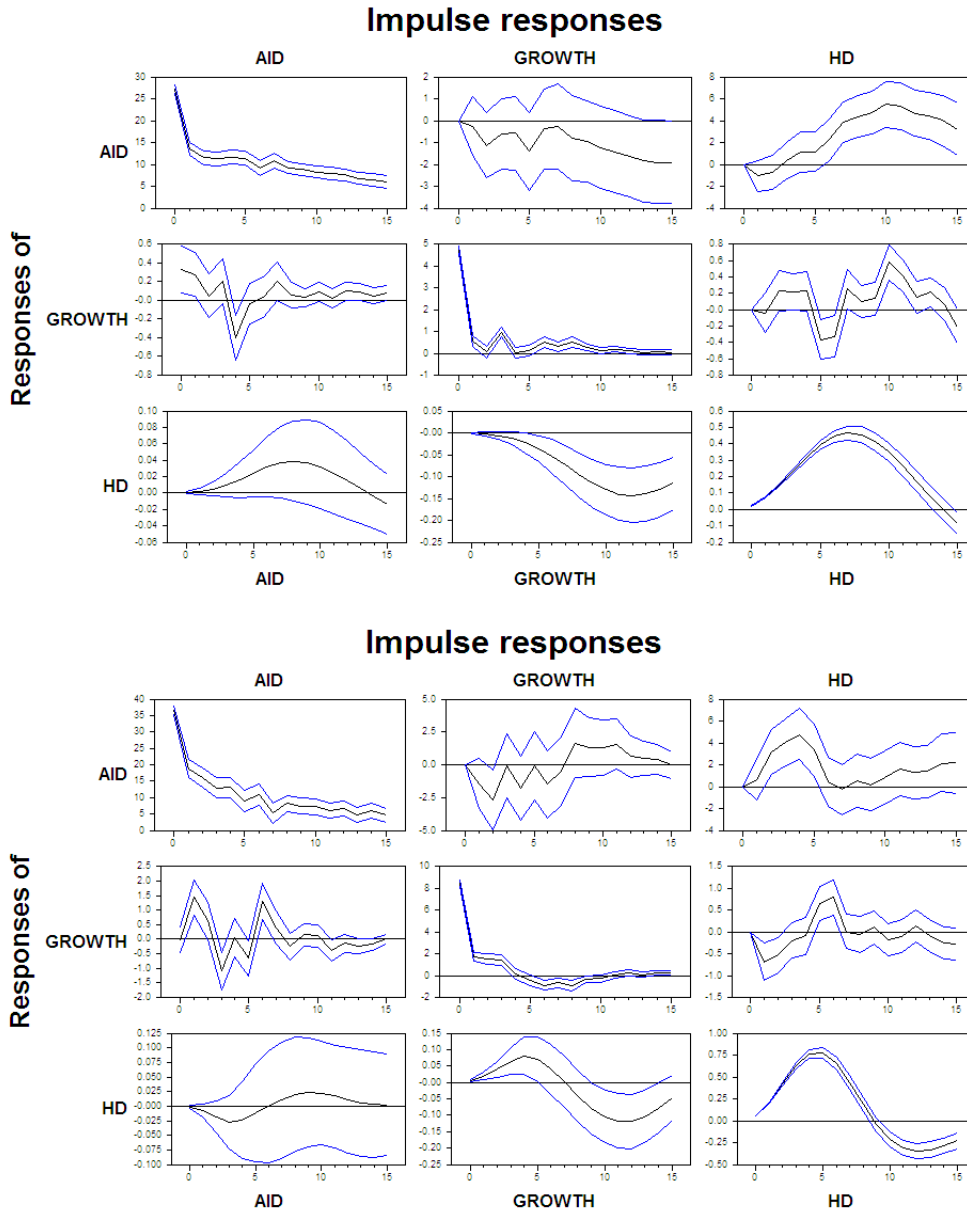
To investigate, I divide the sample using a standard measure of institutional quality. The Polity IV dataset contains a variable which quantifies the constraints on the power of the executive along a seven point scale. I define countries with an average value over the sample

Figure 4: Democracy Vs. Autocracy IRFs



Notes: Each IRF depicts the response of the row variable to a one standard error shock in the estimate of the column variable. The upper and lower lines are 95% error bands generated using Monte Carlo simulation with 2500 draws. The top set is for the sample composed of those countries with an average score of 0 or higher in the Polity index. The bottom set is for the sample composed of those countries with a negative average Polity score.

Figure 5: Good Vs. Bad Institutions IRFs



Notes: Each IRF depicts the response of the row variable to a one standard error shock in the estimate of the column variable. The upper and lower lines are 95% error bands generated using Monte Carlo simulation with 2500 draws. The top set is for the sample composed of those countries with an average score of 3 or higher in the constraint on executive power variable created by the Polity IV project. The bottom set is for the remaining countries.

period of three or higher as possessing good institutions. This is inescapably and unavoidably arbitrary to some degree, but the definition of point three on the scale suggests it is a valid cutoff point.²⁰

Figure 5 presents the results which suggest that good institutions do not lead to more effective aid in terms of economic growth.²¹ In fact, the IRF shows very little response of economic growth to aid. Bad institutions on the other hand, do seem to yield better aid-growth results in the early periods following the aid shock. Compared to the baseline results, the same aid shock gives an economic growth response that is about twice as big. As in the baseline case, this is tempered by some negative responses after period two: negative responses which are likewise larger than in the baseline case. We can also see that there are some marginal negative responses after $t=10$ that persist until $t=15$. Overall, the responses of economic growth to aid, both positive and negative, are larger in the bad institution sub-sample. The relative impact of aid is reversed when we turn to human development. The response in the good institutions sample is in line with the baseline results while in the bad institutions sample we see no appreciable effect.

Why might it be the case that bad institutional environments get more economic bang for their aid buck and less in terms of human development? One possibility is that as they are not as constrained by legislatures representing the popular will and thus can concentrate the aid monies in a few large scale national projects rather than having many small scale local ones. This is all speculation and highlights the biggest drawback of the PVAR approach – while we can see outcomes quite well, the mechanisms are completely obscured.

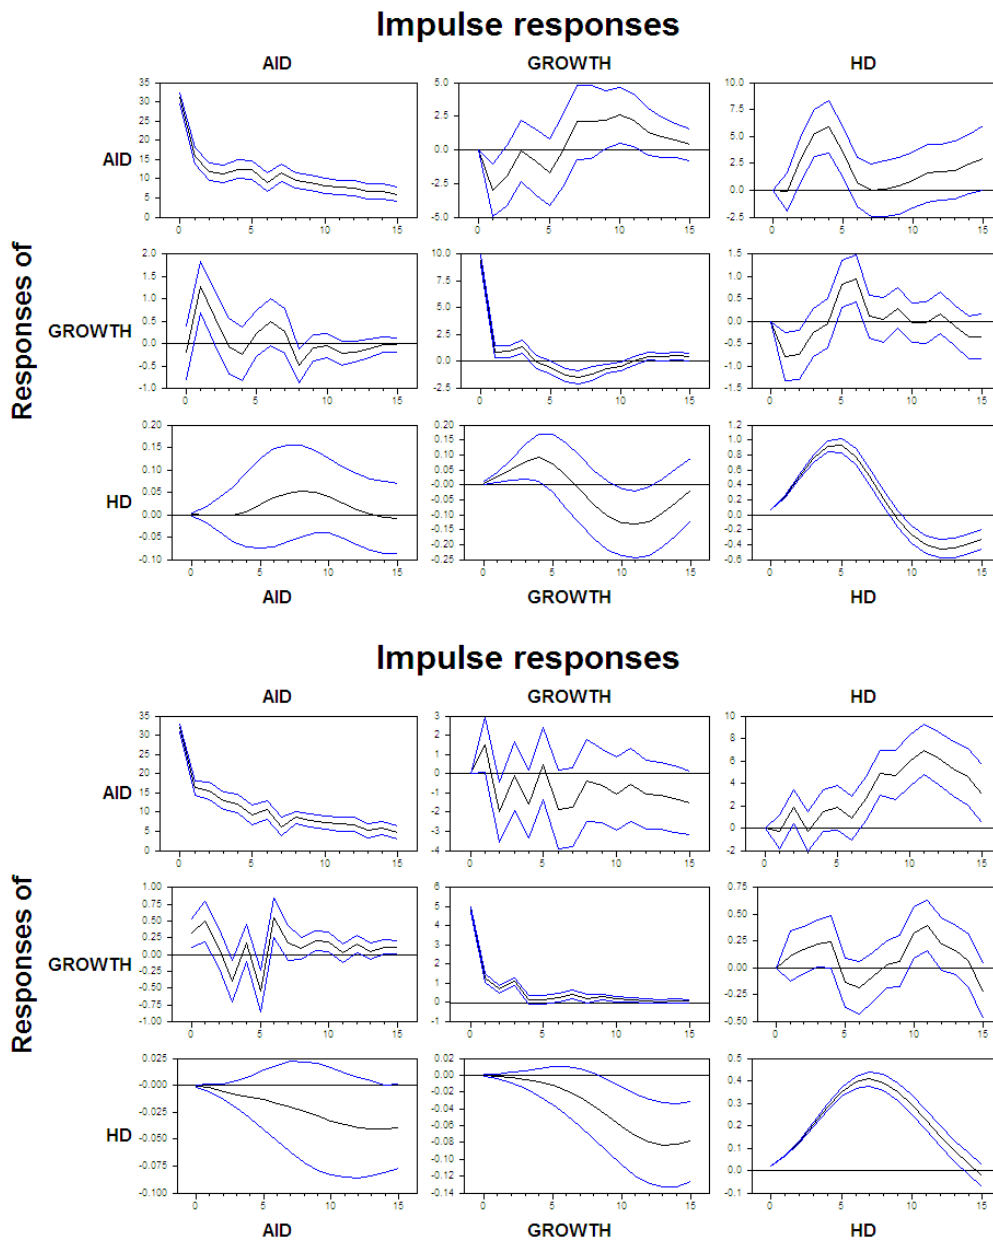
6.4 High Vs. Low Aid Dependence

The final division is along lines of aid dependency. There are no firm theoretical reasons to do so, but there are some plausible mechanisms that could relate aid effectiveness to the level of aid dependency. It could be that a certain level of aid relative to the size of the economy is needed for its effects to be substantial. On the other hand, it could be that high aid dependence fosters corruption, retards reform and induces other such harmful behaviours. There are also issues of absorption that may induce macroeconomic problems which could limit aid effectiveness. These positive and negative potential consequences of the level of aid dependence are not necessarily mutually exclusive. One effect may temper the other or each could operate on different time scales.

²⁰ “Slight to Moderate Limitation on Executive Authority: There are some real but limited restraints on the executive.” Marshall and Jagers [2008].

²¹This result is supported if one uses the World Bank’s Rule of Law (RL) variable as the measure of institutional quality. The Polity variable is preferred as RL is only available for relatively recent years. The results in terms of human development are quite different however, with neither group showing much response. Results available on request.

Figure 6: High Vs. Low Aid Dependence IRFs



Notes: Each IRF depicts the response of the row variable to a one standard error shock in the estimate of the column variable. The upper and lower lines are 95% error bands generated using Monte Carlo simulation with 2500 draws. The top set is for the sample composed of those countries with an average nominal aid to nominal GDP ratio of 15% or higher. The bottom set is for the remaining countries.

While a PVAR model cannot get at these mechanisms, it would be interesting if we could observe outcomes consistent with these arguments. There is no firm definition of aid dependence. I take any country with an average nominal aid to nominal GDP ratio of 15% or greater to be aid dependent. Once again this is somewhat arbitrary but 15% is a large slice of any economy and the mean across the full sample was approximately 12%.

Figure 6 shows the results. The top set are for the aid dependent countries and the bottom set for the remainder. For the most part the results in terms of aid and economic growth are very similar to what we saw in the full sample: an aid shock induces economic growth in the early periods in both cases followed by lower growth in the next few periods with a final positive spike before tapering off. There are some differences worthy of note. Compared to both the full sample and the less dependent sample, the level of the early positive responses in the highly dependent countries are somewhat higher and the subsequent negative responses are not as severe. Perhaps most interestingly, the responses after $t=10$ in the highly dependent sample are negative. In the baseline case, there was no further response after this point. In the low dependence group, we can see some small positive induced growth persisting until the 15th period. While the model offers no explanation for these differences, it is interesting nonetheless to see them in the data. The higher responses in the dependent sample suggests that the first “hypothesis” above may hold. Likewise the negative effects that become apparent only in the dependent sample could be consistent with behavioural changes (such as corruption or postponed reform).

With regards to human development, the results from the two groups are near polar opposites. In the highly dependent countries, an aid shock increases the growth rate of human development in a similar way as in the full sample, though the error bands once again allow for the possibility that the effect is negative or non-existent. In the other countries, aid shocks tend to lower the growth of human development. It is difficult to think of a reason why this is the case and I will refrain from speculation in this instance.

7 Conclusions

This paper has been an attempt to look at whether aid is effective using a different approach to that found in the existing literature. Given the range of contradictory findings on this question, no single work can claim to be definitive. However, the approach taken here is free from many of the difficulties encountered by the usual panel approach. In particular, the PVAR method does not rest on the oft maligned instrumentation strategies used in much of the previous work.

The results lie somewhere between the findings and arguments of aid pessimists and aid optimists. Aid does seem to work in terms of generating economic growth but not to an

extent that could be called transformative. The time path of the responses in most samples showed that the early success of aid is mitigated by some later negative responses before a recovery and eventual petering out. Evidence was also presented that suggests that this impact is more pronounced in countries with good economic policy. Aid does not seem to be more effective at promoting economic development in democracies but does appear to work better, at least initially, in what would commonly be regarded as poor institutional environments. Economies that are relatively dependent on aid show lower negative responses for the most part but small negative responses in later periods that one does not observe in the less dependent sample.

The results for human development were generally ambiguous but suggest that aid may induce small increases in the variable used as a proxy for human development, the growth rate of life expectancy at birth. This increase was larger and unambiguous in democracies and we also saw a fairly unambiguous positive response in good institutional environments. Aid dependent countries see positive responses while the rest show a negative impact.

The major failing of the PVAR method is that it offers no insight into what the underlying mechanisms generating these results might be. That said, much of this literature is based on ad-hoc speculation and econometric specifications. The benefit of the approach is that, as other VAR advocates point out, it lets the data speak for itself. In this case the data seem to be making two points:

1. Aid works in general but not all that well.
2. Aid seems to work better in terms of economic growth in some environments, but not to such an extent that it would be sensible to focus aid only on countries with those characteristics. It is often the case that when we see aid failing to stimulate economic growth in one of our groupings, we see it causing human development growth and vice versa.

A Summary Statistics and Sub-Sample Allocation

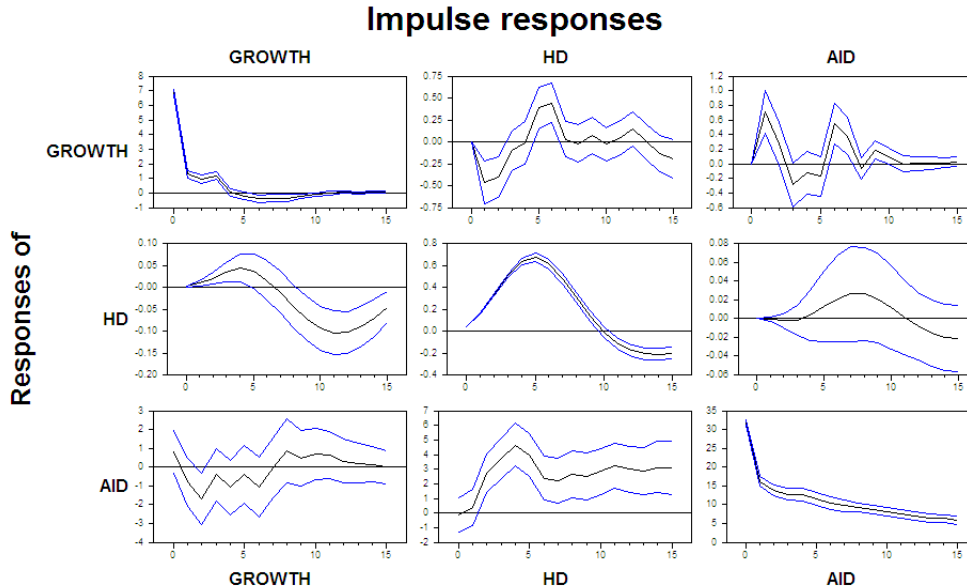
Table 1: Means of Variables Used

Country	Growth rate of Life Expectancy	Growth rate of GDP per capita	ODA per capita	ODA as a percentage of GDP	Doing Business Rank	Constraints on Executive Power	Polity Democracy Score
Benin	0.777	0.466	64.527	9.782	172	3	-1
Botswana	-0.273	6.017	172.829	5.911	45	6	7
Burkina Faso	0.635	1.849	66.388	12.909	147	2	-4
Burundi	0.330	-0.234	56.202	18.623	176	1	-4
Cameroon	0.236	1.141	57.082	4.013	171	2	-6
Central African Rep.	0.179	-1.220	78.300	12.041	183	2	-3
Chad	0.201	1.547	56.652	11.874	178	1	-4
Côte d'Ivoire	0.393	-1.497	57.644	4.103	168	2	-6
DR Congo	0.215	-3.783	26.721	8.449	182	1	-5
Gabon	0.631	0.784	167.456	1.790	158	1	-7
Gambia	0.849	0.567	119.769	20.230	140	4	3
Ghana	0.383	0.176	46.631	7.648	92	3	-2
Guinea-Bissau	0.663	-0.149	149.697	40.479	181	3	-3
Kenya	-0.052	0.292	44.148	6.366	95	3	-4
Lesotho	-0.317	3.016	102.086	17.163	130	3	-2
Liberia	0.696	-3.229	75.870	24.116	149	2	-3
Madagascar	0.837	-1.480	48.313	9.553	134	4	0
Malawi	0.630	-0.037	57.742	18.609	132	2	-4
Mali	0.671	1.478	79.681	16.654	156	3	-1
Mauritania	0.462	-0.226	223.875	22.284	166	3	-7
Niger	0.785	-1.421	67.676	13.665	174	3	-2
Nigeria	0.417	0.469	4.306	0.616	125	3	-2
Rep. Congo	-0.129	1.383	122.089	6.590	179	2	-5
Rwanda	0.247	1.562	75.300	18.995	67	1	-6
Senegal	0.716	-0.034	110.040	10.127	157	4	0
Sierra Leone	0.686	-0.226	50.235	15.077	148	3	-4
Sudan	0.596	1.961	47.863	4.874	154	2	-5
Swaziland	-0.268	2.596	94.583	4.823	115	1	-10
Togo	0.593	-0.414	69.909	9.487	165	1	-5
Zambia	-0.463	-1.498	111.299	15.371	90	3	-3
Zimbabwe	-0.882	-1.099	42.201	3.500	159	4	-2

Notes: Except in the case of the Doing Business rank, values are means over the period 1973-2005. The Doing Business data comes from the 2010 issue. To be included in the good policy sample a country must place in the top 100 in the 2010 Doing Business rankings. The democratic sample comprises countries with an average Polity democracy score of 0 or higher. The good institutions sample is populated by countries with an average constraints on executive power score of 3 or more. Aid dependent countries are defined as those with an average aid as a percentage of GDP in excess of 15%.

B Robustness: Alternative Choleski Decomposition

Figure 7: Full Sample IRFs: Alternative Choleski Decomposition



Notes: Each IRF depicts the response of the row variable to a one standard error shock in the estimate of the column variable. The upper and lower lines are 95% error bands generated using Monte Carlo simulation with 2500 draws.

As mentioned in the main text, the Choleski decomposition chosen is only one of six possible orderings. While I believe that the restrictions are valid, it is worth while to point out that the results are nearly identical if one uses the next most sensible, and quite different, ordering. In particular, let us endow donor agencies with quicker reflexes and restrict aid so that it requires at least one period to have an effect. The ordering is thus:

$$growth_{it} \rightarrow HD_{it} \rightarrow aid_{it}$$

Given that the recovery of the structural shocks hinges on the choice of ordering, this is an important check on the robustness of the results. Figure 7 shows that, apart from the imposed restrictions on contemporaneous impacts, the IRFs are for all intents and purposes identical to those found using the preferred ordering.

References

- The World Bank. *Doing Business 2010*. The International Bank for Reconstruction and Development / The World Bank, 2010a.
- The World Bank. World Development Indicators, 2010b.
- Craig Burnside and David Dollar. Aid, policies, and growth. *American Economic Review*, 90(4):847–868, September 2000.
- Charles C. Chang, Eduardo Fernandez-Arias, and Luis Servén. Measuring aid flows : a new approach. Policy Research Working Paper Series 2050, The World Bank, February 1999.
- Paul Collier and David Dollar. Aid allocation and poverty reduction. *European Economic Review*, 46(8):1475–1500, September 2002.
- C-J. Dalgaard and H. Hansen. On aid, growth and good policies. *The Journal of Development Studies*, 37(6):17–41, August 2001.
- Carl-Johan Dalgaard, Henrik Hansen, and Finn Tarp. On the empirics of foreign aid and growth. *Economic Journal*, 114(496):F191–F216, 06 2004.
- Angus Deaton. Instruments, randomization, and learning about development. *Journal of Economic Literature*, 48(2):424–55, June 2010.
- William Easterly. Can Foreign Aid Buy Growth? *Journal of Economic Perspectives*, 17(3): 23–48, Summer 2003.
- William Easterly, Ross Levine, and David Roodman. Aid, policies, and growth: Comment. *The American Economic Review*, 94(3):774–780, 2004.
- Walter Enders. *Applied Econometric Time Series*. Wiley, 2nd edition, 2004.
- Karuna Gomanee, Sourafel Girma, and Oliver Morrissey. Aid, public spending and human welfare: evidence from quantile regressions. *Journal of International Development*, 17(3): 299–309, 2005a.
- Karuna Gomanee, Oliver Morrissey, Paul Mosley, and Arjan Verschoor. Aid, government expenditure, and aggregate welfare. *World Development*, 33(3):355–370, March 2005b.
- Henrik Hansen and Derek Headey. The short-run macroeconomic impact of foreign aid to small states: An agnostic time series analysis, 2009.
- Stephen Kosack. Effective aid: How democracy allows development aid to improve the quality of life. *World Development*, 31(1):1–22, January 2003.

- Inessa Love and Lea Zicchino. Financial development and dynamic investment behavior: Evidence from panel VAR. *The Quarterly Review of Economics and Finance*, 46(2):190–210, May 2006.
- Monty G. Marshall and Keith Jagers. Polity IV Project: Political Regime Characteristics and Transitions, 1800-2008, 2008.
- Mark McGillivray and Farhad Noorbakhsh. Aid, conflict and human development. Working Papers 2007-03, Department of Economics, University of Glasgow, 2007.
- M. Hashem Pesaran and Ron Smith. Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, 68(1):79–113, July 1995.
- Alessandro Rebucci. On the heterogeneity bias of pooled estimators in stationary var specifications. IMF Working Papers 03/73, International Monetary Fund, April 2003.
- David Roodman. The anarchy of numbers: Aid, development, and cross-country empirics. *World Bank Economic Review*, 21(2):255–277, May 2007.
- J. Svensson. Aid, growth and democracy. *Economics and Politics*, 11(3):275–297, November 1999.
- United Nations Development Programme. Human Development Report , various years.