Does Foreign Aid Target the Poorest?

Ryan C. Briggs*

ryancbriggs@vt.edu

Abstract

This paper examines the extent to which foreign aid reaches people at different levels of wealth in Africa. I use household surveys to measure the sub-national distribution of a country's population by levels of wealth and match this information to data on the location of aid projects from two multilateral donors. Within countries, aid disproportionately flows to regions with more of the richest people. Aid does not favor areas with more of the poorest people. These findings violate the stated preferences of the multilateral donors under study, suggesting that the donors either cannot or are not willing to exercise control over the location of aid projects within countries. The results also suggest that aid is not being allocated effectively to alleviate extreme poverty.

Forthcoming at *International Organization*

*Virginia Tech, Department of Political Science. I wish to thank Senior DHS Specialist Shea Rutstein for advice on working with the DHS surveys and Doug Nicholson at AidData for advice on working with the data from AidData. I would also like to acknowledge useful comments from Maya Berinzon and from presentations at CIDE and MPSA, 2015. This paper appeared as AidData Working Paper 13. Any errors are my own.

“We may not build every road in these countries, but we’re going to build the roads that are going to increase the incomes of the poorest.” –Dr. Jim Yong Kim, President of the World Bank¹

“...foreign aid is a process by which poor people in rich countries help rich people in poor countries.” –Peter Bauer²

The World Bank’s (WB) mission is literally carved in stone at its Washington Headquarters: “Our Dream is a World Free of Poverty.” Foreign aid, in various forms, has been a part of this mission since the founding of the Bank. The link between aid and poverty alleviation was reiterated in the late 1990s when a major Bank report plainly stated that “the main aim of aid is to reduce poverty.”³ Since 2000, this focus has been sharpened by the first Millennium Development Goal of halving the global rate of extreme poverty and the proportion of people who suffer from hunger. In April 2013, with an eye towards a new set of development goals, the WB committed itself to ending extreme poverty and promoting shared prosperity, defined in terms of income growth in the poorest 40% of people in each country.⁴ Poverty reduction is equally a founding and enduring goal of The African Development

¹Tavernise 2014.

²Bauer 1976, 115.

³World Bank 1998, 38.

⁴World Bank 2014a.

Bank (ADB), whose mission is to “promote sustainable economic growth and reduce poverty in Africa.”⁵

Multilateral donors like the WB and ADB are known to be able to direct a great deal of their resources to the world’s poorest countries, resisting the political pressures that skew bilateral aid away from the poorest countries and towards, for example, former colonies.⁶ While multilateral donors are able to target aid to poorer *countries*, we have much less of an understanding about the distribution of their aid within countries. This paper directly answers this question by examining how aid targets people at various levels of wealth within a diverse sample of African countries. An analysis of sub-national aid targeting across many countries helps us understand the extent to which multilateral donors are able to exercise control over aid allocations inside recipient countries in order to realize their stated preferences for a pro-poor allocation of aid. The question of if aid targets poverty within countries is practically quite important as well, as project aid will alleviate poverty more effectively when it is allocated to places where poor people live.

The question of how aid targets sub-national wealth is answered using a two-year sample of geolocated project aid commitments from two multilateral donors to 17 African countries containing a total of 195 regions. By analyzing project aid commitments from multilateral donors, the paper tests for the presence of donor control over aid in a most likely case. This is because

⁵African Development Bank Group 2014.

⁶Maizels and Nissanke 1984; Alesina and Dollar 2000; Dollar and Levin 2006; Nunnenkamp and Thiele 2006. For an overview of the research examining donor behavior, see Neumayer 2003b. Multilateral donors are more resistant to these pressures than bilateral donors, but they are not completely immune (e.g. Dreher, Sturm, and Vreeland 2009).

aid commitments are recorded before implementation, because project aid is a form of aid that is especially easy for donors to target, and because multilateral donors are known to have uniquely apolitical and poverty-sensitive preferences for aid allocation. The analysis includes a demonstration of how to produce reliable, cross-nationally comparable estimates of the sub-national distribution of each country's population by their level of wealth using household surveys. This approach allows us to move from using regional averages of wealth to understanding the unique effect of having more poor or rich⁷ people living in a region. The analysis reveals that aid flows disproportionately to regions with more of the richest people rather than to regions with more of the poorest. It also expands the sub-national study of aid targeting away from single country case studies, and in doing so it demonstrates that there are durable cross-national patterns in sub-national aid allocation in Africa. Finally, it shows that even in this most-likely case, donors do not seem to exercise a great deal of control over aid targeting within countries.

1 Argument

Multilateral donors are uniquely good at directing their aid to poor countries, and the WB and ADB are among the most poverty sensitive of the multilateral donors.⁸ Multilateral donors likely target a larger share of their aid to

⁷In the context of this paper, “poor” or “rich” are always relative terms denoting the poorest or richest people within a country. They are not absolute judgments on levels of wealth.

⁸Maizels and Nissanke 1984; Dollar and Levin 2006.

poor countries than bilateral donors because they have a mission to use aid to reduce poverty and because they have voting arrangements that prevent any one stakeholder country's government from forcing its preferences on all issues. This voting structure encourages aid to flow to areas where stakeholder preferences overlap, implying that it should be relatively easy to send multilateral aid to poor countries but harder to send aid to a country that is a close ally of only one stakeholder country.⁹ Such impediments to political targeting do not exist for bilateral donors, and bilateral aid allocations are more heavily skewed by political factors.¹⁰

Once aid reaches a low-income country, it is generally used to either increase economic growth or directly improve the lives of poor people through the provision of goods or services.¹¹ Aid for the former is typically aimed at either increasing investments, such as infrastructure, or changing economic policies that are thought to stifle growth. Due to almost intractable causal identification issues, there is no consensus on the effect of foreign aid on growth.¹² Bracketing this admittedly large issue, the most common current result from this literature is that under some circumstances aid seems to

⁹On the institutional arrangements that insulate multilateral donors from politics, see Rodrik 1996; Martens, Mummert, Murrell et al. 2002.

¹⁰Alesina and Dollar 2000; Dollar and Levin 2006. In the words of Neumayer, regarding bilateral aid allocation “there is little doubt that economic, political, and sometimes military-strategic interests of donors play a significant and sometimes dominating role for practically all donors...” (Neumayer 2003a, 102).

¹¹Most aid is intended to increase growth and improve the lives of poor people. However, any single project typically aims to directly achieve only one of these two goals while the other goal is expected to come about as a result of attaining the first.

¹²Easterly 2003; Roodman 2007.

have a small, positive effect on growth.¹³ While economic growth is clearly important, it does not always reduce poverty. Consider, for example, that fifteen years of impressive economic growth in Africa has had only a small effect on poverty rates and a little under half of all Africans still live below a \$1.25 a day poverty line.¹⁴

However, if the goal of aid is to reduce poverty, then rather than trying to increase growth and then hoping that growth helps the poor, aid could be used to simply provide the poor with goods and services directly. This could be done through private goods provision, like cash transfers, or, more typically, through the provision of local public goods such as roads, schools, or health clinics.¹⁵ Aid for local public goods can be especially valuable because even when communities get wealthier, they may still struggle to provide local public goods due to collective action problems.¹⁶ However, the benefit of these kinds of goods declines as one moves away from them—a health clinic built near you is useful, a clinic built far away is less useful—so a necessary condition for this kind of aid to help the poor is that local public

¹³Hansen and Tarp 2001; Dalgaard, Hansen, and Tarp 2004; Clemens, Radelet, Bhavnani et al. 2012. For a dissenting view see Doucouliagos and Paldam 2009.

¹⁴The World Bank 2013, 14.

¹⁵The World Bank explicitly provides aid for these purposes. A pamphlet describing the work of the International Development Association, the concessional side of the World Bank that provides aid to 40 countries in Africa, states “When the poorest are ignored because they’re not profitable, IDA delivers. IDA provides dignity and quality of life, bringing clean water, electricity, and toilets to hundreds of millions of poor people” (World Bank 2014b, 11).

¹⁶These collective action problems are usually harder to overcome in more ethnically diverse communities (Miguel 2004; Habyarimana, Humphreys, Posner et al. 2007; Lieberman and McClendon 2013). This makes aid even more useful in sub-Saharan Africa, the world’s most ethno-linguistically diverse region.

goods must be built where poor people live.

While the logic for using aid to provide local public goods is compelling, it is not obvious why donors would choose to provide these goods by funding discrete projects rather than by providing recipient governments with forms of aid that have lower overhead, such as budget support, and then allowing the recipient government to allocate and build these goods themselves.¹⁷ One common explanation is that project aid is used because it gives donors increased control over their aid.¹⁸ This may explain why recipients with worse governance tend to receive WB loans that are targeted more precisely within countries.¹⁹ The fact that donors are ostensibly in control of (especially project) aid allocations, combined with the simple fact that the donors are the ones with the resources that recipients want, has led many scholars to assert that foreign aid should be an exceptionally apolitical and developmental resource. One version of this claim is that aid is a uniquely “scrutinized” resource and so it is able to produce public goods in environments that otherwise produce private goods and pork.²⁰ Another version is that control over aid by foreign donors means that aid should be largely immune to political influence by recipient politicians.²¹ Even studies that find that aid has negative effects on governance typically assume that aid is controlled by donors and that any negative effect of aid must occur through fungibility.

¹⁷This argument is similar to that of Briggs 2014.

¹⁸Morrison 2012, 60.

¹⁹Winters 2010.

²⁰Collier 2006.

²¹van de Walle 2007.

For example, Morrison,²² writes that aid may harm recipient governance like oil because it “is a highly fungible resource” and Smith²³ argues that aid is like oil “to the extent that it is fungible.” Fungibility can be a problem for donors trying to control aid,²⁴ but before worrying about fungibility it is worth considering if recipients can directly influence aid targeting.

While standard accounts of aid bargaining emphasize donor control, recent work has shown that donors sometimes fail to prevent recipients from influencing sub-national aid targeting. For example, Jablonski and Briggs have both shown that recipients skew sub-national aid allocations according to local political incentives in Kenya.²⁵ Domestic politics in recipient countries has also been shown to influence the location of aid-funded electrification in Ghana.²⁶ In the most directly relevant case study for the present paper, Abdulai and Hulme have demonstrated that aid for poverty relief in Ghana was not spent in the country’s poorest regions but rather in its richest regions.²⁷ While these studies draw on detailed sub-national data, their single country focus means that it is difficult to ascertain the extent to which recipient control over aid is a general phenomenon. I address this issue by examining sub-national aid targeting across many recipient countries. Based

²²Morrison 2009, 110.

²³Smith 2008, 780.

²⁴Feyzioglu, Swaroop, and Zhu 1998; Devarajan, Rajkumar, and Swaroop 2007. See Altincekic and Bearce 2014 for an argument that the fungibility of aid is lower than is commonly presumed.

²⁵Jablonski 2014; Briggs 2014.

²⁶Briggs 2012.

²⁷Abdulai and Hulme 2015.

on the public preferences of the donors under study, I expect that *if donors have strong control over aid targeting then project aid should flow to sub-national regions with more of the poorest people*. Not finding such targeting is thus evidence against accounts of aid bargaining that stress donor control.

It should be noted that this paper studies *commitments* for *projects*, and so it does not examine a representative or random sample drawn from the population of all disbursed aid from these donors. The use of commitment data rather than a retrospective analysis of where aid actually went will likely bias the results towards overstating how much aid actually reaches the poor. This is because analyses of aid commitments implicitly assume that there is no leakage between the committed plan and the final outcome of a project. From a donor's point of view, this is a best case scenario. To give one example of how commitments can misrepresent actual aid allocations, in Malawi in the 1990s a large amount of aid for education was essentially turned into cash by members of the regime through contract fraud.²⁸ The commitment data says that the aid was used to fund school construction, but in practice it was used to fund the incumbent party's election campaign. If one assumes that corruption or mismanagement is more likely to direct aid away from the poorest rather than towards the poorest, then any analysis of aid commitments is likely to overstate the extent to which aid actually reaches the poor. The use of data on only project aid will also likely lead to bias if one is trying to draw conclusions about the degree to which aid in general is controlled or targeted by donors. This is because project aid, unlike other

²⁸Briggs 2015.

forms of aid, is given for discrete, specific purposes like building a road in a specific district. This makes project aid relatively easy for donors to control. By analyzing only project aid, I maximize the information available to donors about where aid is going as well as their ability to precisely target aid within recipient countries. These biases imply that the results from the present paper likely overstate both the extent to which aid *in general* is targeted to the poor and the extent to which aid actually reaches poor people.

Sub-national, cross-national work in Africa is complicated by a lack of reliable data,²⁹ and so one additional contribution of the present paper is to highlight how one can use household surveys to produce estimates of the spatial distribution of wealth groupings across regions. This approach complements existing methods for measuring sub-national wealth. One increasingly popular approach to measuring the wealth of regions is to measure the intensity of night-time light. Hodler and Raschky, for example, use this method to show that birthplaces of leaders tend to see better economic outcomes than other areas.³⁰ While this approach is useful, as it measures the presence of economic activity and a working electric grid at a fine level of detail, it says little about the wealth mixture of individuals in a region. From space, a region with a large number of rich people and large number of poor people may look the same as a region with only a large rich population. The approach used in this paper complements strategies like measuring light at night by estimating the spatial distribution of wealth quintiles across regions within

²⁹Jerven 2013.

³⁰Hodler and Raschky 2014.

countries. Rather than simply working with the brightness of light or even average levels of wealth per region, the approach in this paper allows one to parse out the unique effects of richer and poorer people on aid allocation at a sub-national level.

In sum, this paper compares sub-national measures of wealth against geo-tagged aid projects to measure the degree to which foreign aid is targeted to poor people inside countries. It will be shown that aid targets the rich rather than the poor. This finding holds in minimal regressions and after controlling for a number of potentially confounding factors. These results suggest that donors do not exercise strong control over sub-national aid targeting. They also suggest that most aid is not currently able to reduce extreme poverty, as it does not actually reach the very poor.

2 Data

I examine the relationship between wealth and aid using data on the sub-national distribution of aid projects and sub-national measures of wealth created from comparable household surveys. The dependent variable is a measure of aid projects per region.³¹ The data on aid projects comes from AidData³² and includes all projects from the ADB and WB that were ap-

³¹The shapefiles for the ADM1 regions under study come from the UN Food and Agriculture Organization's (FAO) Global Administrative Unit Layers (GAUL) dataset.

³²Strandow, Findley, Nielson et al. 2011.

proved in 2009 and 2010.³³ I included all projects that could be geolocated with a level of precision at the regional (ADM1) level or better and merged both years of data into one cross-sectional dataset. The final dataset includes about 1,400 geolocated projects across Africa.³⁴ Both donors spend on broadly similar sectors. The largest shares of the WB’s aid goes to health and social services, followed by agriculture, transportation, and energy. The ADB’s project aid goes to energy, then transportation, social, and agriculture.³⁵ The biggest sectoral differences between the donors are the WB’s emphasis on health and social services and the ADB’s focus on energy.

The empirical analyses measure aid per region in three ways. The first measure is the region’s share of a country’s total number of aid projects. This measure implicitly considers all projects to be equal in terms of value. The second measure weights projects by their cost. This cannot be done in a straightforward manner because each geolocated project is typically part of a much larger project and the dataset only reports the cost of the larger project. For example, a national electrification project could involve many unique, small, geolocated projects across a country but the dataset only includes the cost of the high-level, national project. To calculate costs per geolocated project, I assume that the cost of a project is split equally across all of its

³³The ADB data is only available for the period 2009/2010 and so the WB data were trimmed to this time period so that they would match.

³⁴Unless otherwise noted, when I speak of “projects” or “geolocated projects,” I am speaking of the unique subprojects that make up the lowest level of information in the AidData dataset. Technically, these are not projects themselves but rather sub-projects that are a part of larger projects.

³⁵The two donors do not use identical coding criteria for their sectors and so the categories are not directly comparable.

sub-projects. This second approach then measures each region’s share of a country’s total dollar value of projects. Both of these initial measures range from zero to one. The third measure is the natural log of the total dollar value of each region’s projects.³⁶

Due to a lack of good regional data on wealth (discussed below), the analysis will be limited to 17 countries. While this limitation is unfortunate, it represents a large widening of scope over the single country focus of the current sub-national aid targeting literature. The total number of aid projects ranges from a high of 247 in Tanzania to single digits in Benin, Guinea, and Niger. Figure 1 shows the number of aid projects per country for the 17 countries in the sample.

2.1 Wealth Data

The key independent variables are drawn from measures of wealth quintiles created by the Demographic and Health Surveys (DHS).³⁷ The DHS has carried out nationally representative and cross-nationally comparable household surveys across Africa for decades. Regional information on wealth from the DHS can be matched to regional data on aid allocations, and this process allows us to examine if areas with more poor (or relatively rich) people receive

³⁶The log transformation of the total value of aid is the preferred operationalization of the aid variable as the dollar value of aid per region exhibits strong positive skew. I report results from all three transformations for transparency. Technically, the logged aid variable is the natural log of the cost per region variable after 0.1 has been added to each region’s total to remove zeros.

³⁷DHS implementing partners and ICF International. Demographic and Health Surveys 2003–2008.

Figure 1: Number of sub-projects per country.

more aid.

The sample is limited to include all Sub-Saharan African countries that had a DHS survey between 1999 and 2008 (inclusive). This ensures that the survey data are relatively close to, but always prior to, the 2009/10 period of the aid data. The sample was then restricted to countries that had DHS surveys that sampled within the country's standard regional boundaries. This procedure leaves 17 countries.³⁸ The sample includes small countries like Lesotho and large ones like the Democratic Republic of the Congo. It includes countries that were colonized by the French, English, Belgians, and Portuguese. It includes populous countries like Nigeria and ones with much smaller populations such as Sierra Leone. In all, it is a diverse sample of African countries.

The DHS has constructed a wealth index for each country in the sample. The index is calculated at the household level and is based on respondents' answers to questions about their assets, such as if they own a bike or a radio, and the quality of their housing. Based on these answers, the DHS groups each respondent into a wealth quintile based on the wealth of their household.³⁹ To make these figures useful for the present paper, they have been recalculated so that they show the estimate of the fraction of each wealth quintile that lives in each region in a country.⁴⁰ This reveals the relative

³⁸Technically, we are left with 17 countries after this procedure and after dropping any countries that did not receive new aid commitments from the WB or ADB in 2009 or 2010.

³⁹For a more in-depth discussion of the construction of the wealth index, see Rutstein and Johnson 2004.

⁴⁰Appendix A includes more information on the construction of this variable.

wealth composition of each region. We can also compare across quintiles because they are equally sized by definition. Within the same country, 10% of the richest quintile will represent the same number of people as 10% of the poorest quintile. In its focus on the location of people grouped by wealth quintile, this study moves away from looking at simple regional averages of wealth and towards taking the underlying distribution of wealth seriously.

An example can serve to clarify the measure and demonstrate the value of tabulating the shares of wealth quintiles over regions. Table 1 shows how the share of wealth quintiles breaks down across regions in Kenya. Each column sums to 100% and represents one fifth of Kenya's total population. While the table does not tell us how many people live in Kenya, it reveals the fraction of each wealth quintile that lives in each region. For example, about a third of Kenya's richest people live in Nairobi. At the other end of the spectrum is Kenya's North Eastern province, which is quite poor and is the site of the Dadaab refugee camp and the Millennium Village of Dertu.⁴¹ However, while North Eastern is poor, it is also lightly populated. This mixture of poverty and low population comes through in the table as the region holds 10% of Kenya's poorest quintile and very small shares of all richer quintiles. This can be contrasted with Rift Valley, which has many people spread across all of the wealth quintiles. Rift Valley is unique in this regard, as most of Kenya's regions skew either towards the rich or the poor. While North Eastern and Nyanza skew towards the poor, Nairobi, and to a lesser extent Central, skew

⁴¹For an illuminating discussion of the challenges facing Dertu and the Millennium Villages Project, see the account by Munk 2013.

towards the rich.

Table 1: Kenya’s distribution of wealth quintiles

	Poorest	2nd Poorest	Middle	2nd Richest	Richest
Central	1.2%	10.9%	19.5%	24.7%	11.9%
Coast	11.4%	4.9%	6.4%	6.5%	11.6%
Eastern	12.5%	18.8%	23.4%	22.5%	6.4%
Nairobi	0.0%	0.0%	0.0%	0.8%	36.3%
North Eastern	10.3%	1.7%	1.1%	0.7%	0.5%
Nyanza	19.4%	23.7%	15.0%	9.8%	8.5%
Rift Valley	32.1%	21.9%	17.8%	26.7%	21.5%
Western	13.0%	18.2%	16.8%	8.3%	3.3%

The segregation of the wealthy into regions that are geographically distinct from the poorer quintiles is a general phenomenon. Table 2 shows the mean of the correlations between wealth quintiles across the 17 countries that make up the sample.⁴² The poorest quintiles of the population correlate highly with each other, implying that places with lots of the poorest people also hold many people from the second poorest segment of the population. This pattern holds all the way to the second richest quintile. Thus, if we see favoritism to the poorest, it could partially reflect aid that is also being targeted to the people from the second-poorest or middle quintiles that live in the same region. The richest group, however, does not match this trend. Aid going to the richest quintile is benefiting the relatively rich far more than anyone else because the richest people are the most likely to live

⁴²A correlation table was calculated separately for each country and then the mean of each cell over all countries is presented in Table 2. This reveals the cross-national average of the country-level correlations between the wealth quintiles.

in relatively homogeneous, wealthier regions. While simply descriptive, the empirical demonstration of this level of geographical segregation is a contribution to our understanding of the sub-national distribution of wealth in Africa.

Table 2: Mean correlations between wealth quintiles over 17 African countries

	Poorest	2nd Poorest	Middle	2nd Richest	Richest
Poorest	1				
2nd Poorest	0.72	1			
Middle	0.51	0.82	1		
2nd Richest	0.27	0.47	0.67	1	
Richest	-0.21	-0.19	-0.11	0.27	1

The results above come from manipulations of household surveys, not census figures, and one may be concerned that this approach produces unreliable estimates of population shares. To address this concern, we can calculate the expected share of the total population per region from the DHS data and compare this against the population shares from census data. Calculating the expected total population share per region from the DHS data requires multiplying each wealth quintile share for a region by 20 and summing the resulting numbers. This approach to validation is useful because it tests a core assumption of the paper, namely that each wealth quintile represents an equal fifth of a country’s population. For example, to calculate Central Province’s expected share of Kenya’s total population one would multiply each percentage in the row representing Central in Table 1 by 20 and then sum the numbers to arrive at 13.6%. The expected shares are then compared against regional population share figures taken directly from Kenyan censuses from 1999 and 2009. These comparisons are done for two

Figure 2: Validating the DHS measure against census data.

countries with reliable censuses, Ghana and Kenya, and the results are plotted in Figure 2. The results confirm that the DHS estimates are able to produce a good representation of sub-national population distributions.

To briefly recap, if donors control aid targeting then the location of aid projects should broadly reflect the stated preferences of the donors giving the aid and so skew to regions with more poor people. Further, the data under study are especially likely to reveal such control because they are project aid commitments from multilateral donors. In order to examine sub-national distributions of wealth, I introduced a method for measuring the distribution of people across regions by their relative level of wealth and I validated this measure against census data. The analysis below leverages this new variable, and the fact that the poorest segments of the population positively correlate with each other while the richest live alone, to examine the extent to which sub-national aid distributions reflect the stated, pro-poor preferences of the multilateral donors under study. The next section presents the formal analysis of how project aid and wealth overlap.

3 Analysis

The analysis explains how aid is allocated across 195 regions inside 17 countries, and all regressions use country fixed effects. Country fixed effects are important to the estimation, as they allow us to examine only the factors that vary across regions within countries. The fixed effects also are necessary for the wealth quintile variables to make sense, as all wealth quintiles

are equally sized within countries but not across countries.⁴³ The dependent variable is measured in three ways: as each region’s share of a country’s total dollar value of aid, as each region’s share of a country’s total number of aid projects, and finally as the natural log of the total dollar value of aid to each region. The key independent variables are the share of each wealth quintile that resides in each region.

The initial results are presented in Figure 3 and show the bivariate relationships between the share of people in each wealth quintile and the region’s level of aid.⁴⁴ The dependent variables in the left panel the figure shows each region’s share of the total value of aid per country and each region’s share of the total number of projects. The right panel uses the natural log of the total value of aid per region and the natural log of each wealth share variable.⁴⁵ The unconditional relationship between aid and wealth is descriptive, but it is important and reveals that aid is not targeting the poorest. In general, regions tend to receive more aid when they have more people in higher wealth quintiles. This effect holds across all three ways of measuring aid.

The remaining analyses work with the same dependent variables but only

⁴³The fixed effects are also important because the cut points in the wealth quintiles are country-specific. The poorest quintile in Nigeria, for example, likely does not have the same level of wealth as the poorest quintile in Malawi.

⁴⁴The “bivariate” OLS regressions used to produce Figure 3 include one wealth variable at a time and country fixed effects. The figure shows point estimates and 95% confidence intervals. Standard errors are clustered on countries. All regressions include 17 countries and 195 regions.

⁴⁵The wealth share variables range from 0 to 1 and exhibit positive skew. I added 0.001 to the variables when taking the log to remove zeros. The right panel shows the preferred specification, as the fit of the models is improved by the addition of the logs. This is demonstrated graphically in Appendix B.

Figure 3: Bivariate (fixed effects) results.

include the share of the richest and poorest quintiles as independent variables. Aside from the fixed effects, they also include each region's share of the country's total area, the region's share of conflict in the country, and a dummy variable marking if the region holds the country's capital. All of the independent variables are either fixed (e.g. area) or were measured before aid commitments were made in 2009 and 2010. The area control accounts for the fact that the optimal number of many aid-funded goods may increase with the area of a region. For example, all else equal, longer roads are needed in larger regions. The region's share of conflict controls for the fact that donors might direct aid away from regions that have worse security. This is measured as each region's share of the total number of battles in the two years leading up to 2009.⁴⁶ Finally, capital cities hold many of the wealthiest people and aid recipients may want to skew aid or other resources towards capital cities because unrest in capital cities is uniquely dangerous to recipient governments.⁴⁷ The analyses control for the presence of a capital city in a region using a dummy variable. The results are shown in Table 3.

The key result is that while the effect of the fraction of the poorest people on aid is not statistically significant, the effect of the richest is consistently large and significant. This is the case when regional aid is measured as a fraction of the total cost of all projects, when aid is measured as a fraction of the total number of all projects, or when aid is measured as the log of a

⁴⁶Battles are specific instances of violent conflict between two politically organized armed groups. There is no minimum level of violence required for an event to be a battle. For more details, see Raleigh, Linke, Hegre et al. 2010.

⁴⁷Bates 1981.

Table 3: Main Results

	(1) Value	(2) Projects	(3) ln(Value)
% of Poorest	0.29 (0.175)	0.25* (0.131)	
% of Richest	0.61** (0.218)	0.66*** (0.214)	
ln(% of Poorest)			0.10 (0.086)
ln(% of Richest)			0.72*** (0.204)
% of Battles	-0.10 (0.083)	-0.06 (0.081)	
Battles			-0.00 (0.002)
% of Area	0.39 (0.227)	0.45** (0.207)	
ln(Area)			0.32*** (0.087)
Capital	0.03 (0.068)	-0.03 (0.056)	1.13 (0.765)
Fixed Effects	Yes	Yes	Yes
Number of countries	17	17	17
Number of regions	195	195	195
R-squared	0.23	0.25	0.24

Robust standard errors clustered on countries in parentheses

*** p<0.01, ** p<0.05, * p<0.1

region's total value of aid (and the wealth share variables are also logged). Even after controlling for area, violence, and capital city bias, aid is still skewed towards the rich. There is no evidence that violence or the presence of a capital city has any direct influence on aid.⁴⁸ The effect of the richest segment is also much larger than 0.2, which would be its expected value if aid was allocated equally by population.⁴⁹

3.1 Robustness

The above results hold across a large number of robustness checks, all of which are fully reported in the appendix and only briefly mentioned here. The robustness tests include: dropping all countries that received fewer than five aid projects or had fewer than five sub-national regions, adding a control for within-region inequality, the use of models that are designed for dependent variables that are bound at zero and one, regressions with observations weighted so that each country (rather than each region) contributes equally to the results, tests where I sequentially exclude one country at a time from the models to examine if the results are being driven by individual countries, tests to see if aid is more likely to flow to regions where the recipient country's sitting president was born, selection models where I examine how the

⁴⁸The battles variable is also insignificant and does not change the results if it logged (after adding 0.1) or if it is entered as a dummy variable where it takes a value of one if there were any battles in the region in the past two years and zero otherwise.

⁴⁹Note that the wealth variables measure the share of a population quintile, which itself represents 20% of the population. The coefficient could plausibly be larger if people from the poorer quintiles tended to live in the same regions as people from the richest quintiles, but this generally does not happen (see Table 2).

wealth variables influence the probability of a region receiving any aid, and disaggregated analyses where I see if the two donors use different targeting strategies. In total, the main text and the appendix present 93 tests of the influence of the poorest and richest wealth quintiles on aid. The effect of the richest was significant ($p < 0.05$) in 92 (99%) of the tests and the effect of the poorest was significant in 3 (3%) of the tests. While these figures do not come from an exhaustive specification search, they demonstrate that the core results hold across a wide variety of variable transformations, dataset modifications, and modeling choices.⁵⁰ Aid does not flow disproportionately to regions with more of the poorest people.

While most of the robustness tests simply confirm the paper’s core results, a few findings warrant mention. First, regions where the recipient’s president were born are not more likely to receive aid. While this result is unexpected, it is in line with other work suggesting that co-ethnics are not always favored by African presidents⁵¹ or that co-ethnics are not evenly favored across types of goods.⁵² The selection models reveal that regions with more of the richest are more likely to receive at least one aid project. There is no significant effect for the poorest. The results by donor show that neither donor targets aid to the poorest and both target aid to the richest. The main difference between the donors is that it seems that the ADB’s skew to the rich is largely driven by a bias to regions with capital cities while the WB’s is not.

⁵⁰The inspiration for this approach is Leamer 1985.

⁵¹Kasara 2007.

⁵²Kramon and Posner 2013.

In interpreting all of the results, one should keep in mind the correlational structure of the data (see Table 2). Across the 17 countries under study, the poorer quintiles tend to cluster together while the richest tend to be more segregated. This implies that when aid provides local public goods to the regions with more of the richest people, as it is doing in these countries, it is missing the regions that hold not only the poorest, but also large portions of the bottom half of society.

Finally, it is important to stress that all of the results in this paper are produced using model-based, as opposed to design-based, empirical strategies.⁵³ While this implies that the causal effect of poor or rich people on aid is not identified, this should not overshadow the importance of the above correlations. Aid is flowing to the rich instead of the poor, and this is not due to factors such as the size of regions, the location of capital cities, internal security, or ethnic targeting. Further, the coefficient for the effect of the richest on aid in the preferred specification barely changes when moving from the bivariate model graphed in the right panel of Figure 3 to the results with full controls shown in Table 3, model 3. When conditioning on observables has such a minimal effect on a coefficient it suggests that the unobserved variables would have to be significantly more confounding than the all of observed variables in order to overturn the result.⁵⁴ The next section turns to the implications of these findings.

⁵³Sekhon 2009; Dunning 2010.

⁵⁴This logic is presented in Bellows and Miguel 2009; González and Miguel 2015.

4 Discussion

The first implication of the results concerns the question of why aid might be expected to influence recipient governance similarly to oil. The general answer is that aid may act like oil because both are tax free resources.⁵⁵ However, one important difference between aid and oil is that aid may be much more “encumbered,” meaning scrutinized and controlled by the donor, than oil.⁵⁶ If aid is highly encumbered, then it should only harm recipient governance if it is also highly fungible. If aid is not encumbered, however, then there is good reason to expect it to directly harm recipient governance and the question of fungibility is not particularly relevant. While the theory suggesting that aid will be highly encumbered is compelling, this paper found that sub-national aid allocations directly contradict the stated preferences of the donors. This suggests that donors are either not willing or not able to encumber their aid. If this is happening to project aid, which is relatively easy to monitor and target, then it should be more likely with other forms of aid. Similarly, bilateral donors, who have much stronger incentives to use aid as a bribe, should be less likely to encumber their aid than poverty-conscious multilateral donors. Thus, unencumbered aid is probably quite common.

While this paper did not probe the mechanisms preventing donors from targeting their aid to the poor, they are worth mentioning. First, it is possible that donors do not actually care where aid goes within countries. For

⁵⁵Moore 2001; Morrison 2009, 2012; Djankov, Montalvo, and Reynal-Querol 2008. For dissenting views on the link between aid and declines in governance, see Bermeo 2016; Jones and Tarp 2016.

⁵⁶Collier 2006; van de Walle 2007; Altincekic and Bearce 2014; Bermeo 2016.

example, a World Bank staff member in Tanzania stated that “The World Bank provides financing. And the client, as we call it, or the government then basically decides where they want to spend the money.”⁵⁷ If this sentiment is common, then it is no surprise that aid does not follow donor preferences. Second, it is possible that donors care about sub-national aid allocation but that recipients strategically approach donors and manipulate information to produce their desired distribution.⁵⁸ In this case, a recipient’s strategic action can allow it to defeat attempts at donor control. Finally, donors may not have good information on where the poor live, and so they cannot convert their abstract preference for pro-poor aid allocations into concrete preferences of levels of aid over regions in specific countries. If this is the case, then better information should result in aid allocations more closely aligning with donor preferences. Which of these mechanisms best explains how donors lose control over aid targeting is a question for future research.

The present paper showed that there are robust cross-national patterns in sub-national aid targeting, with aid being directed to regions with more of the richest people. This finding is consistent with previous work showing that aid is insensitive to measures of regional need like average rates of infant mortality or malnutrition.⁵⁹ Given that project aid typically provides local public goods, these results suggest that aid is not being allocated effectively to alleviate extreme poverty.

⁵⁷Masaki 2014, 8.

⁵⁸For an example of this kind of strategic action, see Briggs 2014.

⁵⁹Öhler and Nunnenkamp 2014. For results showing that World Banks aid is largely insensitive to regional poverty in India, see Nunnenkamp, Öhler, and Sosa Andrés 2016.

The finding that aid targets the richest people within countries also has implications for choices about which countries should receive aid. Recent work has shown that a majority of the world’s poorest people currently live in middle, rather than low, income countries.⁶⁰ This could lead one to conclude that poverty-sensitive donors should direct more aid to these middle income countries with many poor people. While sensible, this assumes that it is trivially easy for donors to target aid to the poorest people within countries. This seems to not be the case. In general, we should probably be skeptical of strategies for cross-national aid allocation that hinge on donors successfully encumbering their aid.

While ethnicity often influences the distribution of resources in Africa, the present paper found no evidence that aid from these donors favors the president’s home region. Instead, these results support an explanation that hinges on the role of wealth in influencing distributive politics.⁶¹ A strong focus on ethnicity may blind us to other politically relevant groupings. Future work can look more closely at the relative influence of ethnic and economic factors and the conditions under which governments are more sensitive to one factor over the other.

Finally, this article presented a new way of measuring the relative spread of wealth across regions within countries. This approach complements other approaches, such as measuring sub-national levels of light at night, by al-

⁶⁰Kanbur and Sumner 2012; Sumner 2012.

⁶¹Explanations for resource distributions that stress the importance of wealth in influencing politics fit well with institutional explanations for economic growth, such as Acemoglu and Robinson 2012.

lowing one to parse out the unique influence of different wealth groupings. This measure is likely to be useful in countries where national statistical data are questionable or where sub-national measures are not easily comparable across countries.

5 Conclusion

This paper combined sub-national measures of wealth with geolocated project aid data across 17 African countries to show that aid does not favor the poorest. Rather, aid disproportionately flows to regions that hold more of a country's richest people. While the paper did not identify the causal effect of poor or rich people on aid allocations, it showed that the results are not being driven by factors such as a region's size, regional inequality, the presence of a capital city, internal security, or presidents favoring their home regions. In line with recent case study research on the politics of foreign aid, and opposed to work emphasizing the power of donors, the results suggest that donors do not exercise strong control over aid targeting within recipient countries.

References

- Abdulai, Abdul-Gafaru, and David Hulme. 2015. The Politics Of Regional Inequality In Ghana: State Elites, Donors And PRSPs. *Development Policy Review* 33 (5):529–553.
- Acemoglu, Daron, and James A Robinson. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Business.
- African Development Bank Group. 2014. About us. <http://www.afdb.org/en/about-us/>.
- Alesina, Alberto, and David Dollar. 2000. Who Gives Foreign Aid to Whom and Why? *Journal of Economic Growth* 5 (1):33–63.
- Altincekic, Ceren, and David H Bearce. 2014. Why There Should be No Political Foreign Aid Curse. *World Development* 64:18–32.
- Bates, Robert H. 1981. *Markets and States In Tropical Africa: The Political Basis of Agricultural Policies*. Oakland, CA: University of California Press.
- Bauer, Péter Tamás. 1976. *Dissent on Development*. Cambridge, MA: Harvard University Press.
- Bellows, John, and Edward Miguel. 2009. War and Local Collective Action in Sierra Leone. *Journal of Public Economics* 93 (11):1144–1157.
- Bermeo, Sarah Blodgett. 2016. Aid is Not Oil: Donor Preferences, Heterogeneous Aid, and the Aid-Democratization Relationship. *International Organization* 70 (1):1–32.

- Briggs, Ryan C. 2012. Electrifying the Base? Aid and Incumbent Advantage in Ghana. *The Journal of Modern African Studies* 50 (4):603–624.
- . 2014. Aiding and Abetting: Project Aid and Ethnic Politics in Kenya. *World Development* 64:194–205.
- . 2015. The Influence of Aid Changes on African Election Outcomes. *International Interactions* 41 (2):201–224.
- Clemens, Michael A, Steven Radelet, Rikhil R Bhavnani, and Samuel Bazzi. 2012. Counting Chickens when they Hatch: Timing and the Effects of Aid on Growth. *The Economic Journal* 122 (561):590–617.
- Collier, Paul. 2006. Is Aid Oil? An Analysis of Whether Africa Can Absorb More Aid. *World Development* 34 (9):1482–1497.
- Dalgaard, Carl-Johan, Henrik Hansen, and Finn Tarp. 2004. On The Empirics of Foreign Aid and Growth. *The Economic Journal* 114 (496):F191–F216.
- Devarajan, S., A.S. Rajkumar, and V. Swaroop. 2007. What Does Aid to Africa Finance? In *Theory and Practice of Foreign Aid*, edited by Sajal Lahiri, 333–355. Amsterdam: Elsevier.
- DHS implementing partners and ICF International. Demographic and Health Surveys. 2003–2008. Demographic and Health Surveys (various) [Datasets]. Calverton, Maryland: ICF International [Distributor].
- Djankov, S., J.G. Montalvo, and M. Reynal-Querol. 2008. The Curse of Aid. *Journal of Economic Growth* 13 (3):169–194.

- Dollar, David, and Victoria Levin. 2006. The Increasing Selectivity of Foreign Aid, 1984–2003. *World Development* 34 (12):2034–2046.
- Doucouliafos, Hristos, and Martin Paldam. 2009. The Aid Effectiveness Literature: The Sad Results of 40 Years of Research. *Journal of Economic Surveys* 23 (3):433–461.
- Dreher, Axel, Jan-Egbert Sturm, and James Raymond Vreeland. 2009. Development Aid and International Politics: Does Membership on the UN Security Council Influence World Bank Decisions? *Journal of Development Economics* 88 (1):1–18.
- Dunning, Thad. 2010. Design-Based Inference: Beyond the Pitfalls of Regression Analysis? In *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, edited by Henry E. Brady and David Collier. Lanham, Maryland: Rowman & Littlefield Publishers, 2nd edn.
- Easterly, W. 2003. Can Foreign Aid Buy Growth? *The Journal of Economic Perspectives* 17 (3):23–48.
- Feyzioglu, Tarhan, Vinaya Swaroop, and Min Zhu. 1998. A Panel Data Analysis of the Fungibility of Foreign Aid. *The World Bank Economic Review* 12 (1):29–58.
- González, Felipe, and Edward Miguel. 2015. War and Local Collective Action in Sierra Leone: A Comment on the Use of Coefficient Stability Approaches. *Journal of Public Economics* 128:30–33.
- Habyarimana, James, Macartan Humphreys, Daniel N Posner, and Jeremy M

- Weinstein. 2007. Why Does Ethnic Diversity Undermine Public Goods Provision? *American Political Science Review* 101 (4):709–725.
- Hansen, H., and F. Tarp. 2001. Aid and Growth Regressions. *Journal of Development Economics* 64 (2):547–570.
- Hodler, Roland, and Paul A Raschky. 2014. Regional Favoritism. *The Quarterly Journal of Economics* 995–1033.
- Jablonski, Ryan. 2014. How Aid Targets Votes: The Impact of Electoral Incentives on Foreign Aid Distribution. *World Politics* 66 (2):293–330.
- Jerven, Morten. 2013. *Poor Numbers: How We Are Misled by African Development Statistics and What to Do about It*. Ithaca, New York: Cornell University Press.
- Jones, Sam, and Finn Tarp. 2016. Does foreign aid harm political institutions? *Journal of Development Economics* 118:266–281.
- Kanbur, Ravi, and Andy Sumner. 2012. Poor Countries or Poor People? Development Assistance and the New Geography of Global Poverty. *Journal of International Development* 24 (6):686–695.
- Kasara, Kimuli. 2007. Tax Me if You Can: Ethnic Geography, Democracy, and the Taxation of Agriculture in Africa. *American Political Science Review* 101 (1):159–172.
- Kramon, Eric, and Daniel N Posner. 2013. Who Benefits From Distributive Politics? How the Outcome One Studies Affects the Answer One Gets. *Perspectives on Politics* 11 (2):461–474.

- Leamer, Edward E. 1985. Sensitivity Analyses Would Help. *The American Economic Review* 75 (3):308–313.
- Lieberman, Evan S, and Gwyneth H McClendon. 2013. The Ethnicity–Policy Preference Link in Sub-saharan Africa. *Comparative Political Studies* 46 (5):574–602.
- Maizels, Alfred, and Machiko K Nissanke. 1984. Motivations for Aid to Developing Countries. *World Development* 12 (9):879–900.
- Martens, Bertin, Uwe Mummert, Peter Murrell, and Paul Seabright. 2002. *The Institutional Economics of Foreign Aid*. Cambridge, UK: Cambridge University Press.
- Masaki, Takaaki. 2014. The Political Economy of Aid Allocation in Africa: Evidence from Zambia. Working Paper.
- Miguel, Edward. 2004. Tribe or Nation? Nation Building and Public Goods in Kenya Versus Tanzania. *World Politics* 56 (3):328–362.
- Moore, Mick. 2001. Political Underdevelopment: What Causes ‘Bad Governance’. *Public Management Review* 3 (3):385–418.
- Morrison, K.M. 2009. Oil, Nontax Revenue, and the Redistributive Foundations of Regime Stability. *International Organization* 63 (1):107–138.
- . 2012. What Can We Learn about the Resource Curse from Foreign Aid? *The World Bank Research Observer* 27 (1):52–73.
- Munk, Nina. 2013. *The Idealist: Jeffrey Sachs and the Quest to End Poverty*. New York: Random House.

- Neumayer, Eric. 2003a. The Determinants of Aid Allocation by Regional Multilateral Development Banks and United Nations Agencies. *International Studies Quarterly* 47 (1):101–122.
- . 2003b. *Pattern of Aid Giving: The Impact of Good Governance on Development Assistance*. London: Routledge.
- Nunnenkamp, Peter, Hannes Ohler, and Maximiliano Sosa Andrés. 2016. Need, merit, and politics in multilateral aid allocation: A district-level analysis of World Bank projects in India. *Review of Development Economics* .
- Nunnenkamp, Peter, and Rainer Thiele. 2006. Targeting Aid to the Needy and Deserving: Nothing but Promises? *The World Economy* 29 (9):1177–1201.
- Öhler, Hannes, and Peter Nunnenkamp. 2014. Needs-Based Targeting or Favoritism? The Regional Allocation of Multilateral Aid within Recipient Countries. *Kyklos* 67 (3):420–446.
- Raleigh, Clionadh, Andrew Linke, Håvard Hegre, and Joakim Karlsen. 2010. Introducing ACLED: An Armed Conflict Location and Event Dataset Special Data Feature. *Journal of Peace Research* 47 (5):651–660.
- Rodrik, Dani. 1996. Why Is There Multilateral Lending? In *Annual World Bank Conference on Development Economics, 1995*, edited by Michael Bruno and Boris Pleskovic, 167–193. Washington, DC: International Monetary Fund.

- Roodman, D. 2007. The Anarchy of Numbers: Aid, Development, and Cross-country Empirics. *The World Bank Economic Review* 21 (2):255–277.
- Rutstein, Shea Oscar, and Kiersten Johnson. 2004. DHS Comparative Reports No. 6: The DHS Wealth Index. *Tech. rep.*, Demographic and Health Statistics, <http://dhsprogram.com/pubs/pdf/CR6/CR6.pdf>.
- Sekhon, Jasjeet S. 2009. Opiates for the Matches: Matching Methods for Causal Inference. *Annual Review of Political Science* 12:487–508.
- Smith, Alastair. 2008. The Perils of Unearned Income. *The Journal of Politics* 70 (3):780–793.
- Strandow, Daniel, Michael Findley, Daniel Nielson, and Josh Powell. 2011. *The UCDP-AidData Codebook on Geo-Referencing Foreign Aid. Version 1.1*. Uppsala, Sweden: Uppsala University.: Uppsala Conflict Data Program.
- Sumner, Andy. 2012. Where do the Poor Live? *World Development* 40 (5):865–877.
- Tavernise, Sabrina. 2014. Head of World Bank Makes Ebola His Mission. New York Times. <http://www.nytimes.com/2014/10/14/science/a-bank-chief-makes-ebola-his-mission.html>.
- The World Bank. 2013. *Africa's Pulse*, vol. 7. The World Bank.
- van de Walle, Nicolas. 2007. Meet the New Boss, Same as the Old Boss? The Evolution of Political Clientelism in Africa. In *Patrons, Clients, and*

Policies, edited by Herbert Kitschelt and Steven I. Wilkinson, chap. 2, 50–67. Cambridge University Press.

Winters, Matthew S. 2010. Choosing to Target: What Types of Countries Get Different Types of World Bank Projects. *World Politics* 62 (3):422–458.

World Bank. 1998. *Assessing Aid: What Works, What Doesn't, and Why*. Washington, DC: The World Bank.

———. 2014a. Poverty Overview: Strategy. <http://www.worldbank.org/en/topic/poverty/overview#2>.

———. 2014b. The World Bank's Fund For The Poorest. <http://www.worldbank.org/ida/what-is-ida/fund-for-the-poorest.pdf>.

6 Appendix A: Making the Wealth Variables

This appendix first discusses how countries were selected into the sample in more detail and then discusses the construction of the wealth variables. The sample of countries was composed of every Sub-Saharan African country that had at least one regionally geolocated aid project from the WB or ADB in 2009 or 2010,⁶² and that had a DHS survey that:

1. Was published between 1999 and 2008,
2. Was constructed to allow for estimates at the regional level,
3. Included the wealth index, and
4. Used the country's standard ADM1 regions.

If a country had more than one such DHS survey in the 10-year window, the most recent one was selected. Most of the countries that were cut from the sample were cut because they either did not have any DHS surveys or they did not have one in the decade before aid was disbursed. Additionally, six countries were not considered because they received no new commitments of aid from the WB or ADB in 2009 or 2010. This selection process produced a sample of 17 countries. The countries and the dates of their DHS surveys are listed in Table 4.

DHS surveys with the wealth index also include information placing each household within one of five wealth quintiles. These quintiles are constructed from questions asking about ownership of various assets such as televisions, toilet facilities, or the type of flooring material. The quintiles are constructed so that they should reflect the respondent's placement within the de jure household population. This is different from the population of individuals surveyed because it includes people younger than 15 and older than 49.

To calculate shares of wealth quintiles across regions, I divided the number of surveyed households in a given wealth quintile in each region by the total number of surveyed households in that quintile. In essence, this approach takes advantage of the fact that all of the DHS surveys under examination at some point divide the country into ADM1s and then sample within regions

⁶²At least one project had to be geolocated to the regional level or better. All data on aid projects comes from AidData (Strandow, Findley, Nielson et al. 2011).

Table 4: DHS Survey Years

Country	Year of DHS Survey
Benin	2006
DRC	2007
Ethiopia	2005
Ghana	2008
Guinea	2005
Kenya	2003
Lesotho	2004
Malawi	2004
Mali	2006
Mozambique	2003
Namibia	2006
Niger	2006
Nigeria	2008
Rwanda	2007-2008
Sierra Leone	2008
Tanzania	2004-2005
Zambia	2007

with probability proportionate to population. All calculations were done while weighting the figures by both the probability of being sampled and de jure household membership. In practice, these “household membership weights” are constructed by multiplying the sample weight (typically hv005) by household size (typically hv012). This allows us to take into consideration the fact some small populations are oversampled and the fact that households vary in size (in ways that are not caught by the sampling because they include people below 15 and above 49). While the use of the weights is clearly best practice, as it corrects for oversampling and for different sizes of households (outside of the 15–49 year sampling frame), in practice the use of the weighting scheme leads to only small changes when compared to similar calculations without the use of the weights.

While a wealth share variable constructed in this way from the DHS surveys must include some random error, it also produces estimates of regional populations that are very close to national censuses, as was shown in Figure 2. This section presents tables that show the raw data behind Figure 2. Table 5 reproduces Table 1 but also includes three extra columns. The first extra column estimates the fraction of the total Kenyan population in each region from the DHS wealth quintile distributions. It does this by multiplying every percentage in each row by 20 and then summing the resulting numbers. This number is then compared to the regional population distributions from the Kenyan censuses of 1999 and 2009 (the DHS report was carried out in 2003). The DHS and census results align closely.

Table 6 repeats the same procedure for Ghana. As with Kenya, the censuses and the manipulated DHS quintiles produce similar results and there is no obvious bias. While Accra has a smaller population in the DHS output than in the Ghanaian census data, Nairobi is similar in the Kenyan census data and the DHS output. While the DHS figure for the rather poor Upper East is larger than the census figure in the Ghanaian data, the DHS figure for the similarly poor North Eastern is smaller than the Kenyan census data. No DHS estimate misses its nearest census by more than 1.5 percentage points and most differences fall much closer to 0. These similarities, as well as the good match between quintile distributions within countries and prior expectations (e.g. Nairobi is rich, North Eastern is poor and lightly populated, Rift Valley is populous), reinforce the utility and validity of this way of measuring the distribution of people according to wealth across regions within countries. This is significant, as it is difficult to construct nuanced, sub-national, and cross-nationally comparable measures of wealth in Africa.

Table 5: DHS Wealth Quintiles Compared to Kenyan Census Reports

	Poorest	2nd Poorest	Middle	2nd Richest	Richest	Total	Census (1999)	Census (2009)
Central	1.2%	10.9%	19.5%	24.7%	11.9%	13.6%	13.0%	11.4%
Coast	11.4%	4.9%	6.4%	6.5%	11.6%	8.2%	8.7%	8.6%
Eastern	12.5%	18.8%	23.4%	22.5%	6.4%	16.7%	16.1%	14.7%
Nairobi	0.0%	0.0%	0.0%	0.8%	36.3%	7.4%	7.5%	8.1%
North Eastern	10.3%	1.7%	1.1%	0.7%	0.5%	2.8%	3.4%	6.0%
Nyanza	19.4%	23.7%	15.0%	9.8%	8.5%	15.3%	15.3%	14.1%
Rift Valley	32.1%	21.9%	17.8%	26.7%	21.5%	24.0%	24.4%	25.9%
Western	13.0%	18.2%	16.8%	8.3%	3.3%	11.9%	11.7%	11.2%

Table 6: DHS Wealth Quintiles Compared to Ghanaian Census Reports

	Poorest	2nd Poorest	Middle	2nd Richest	Richest	Total	Census (2000)	Census (2010)
Ashanti	5.6%	18.3%	21.0%	24.6%	21.5%	18.2%	19.1%	19.4%
Brong Ahafo	11.7%	11.2%	10.6%	9.9%	3.1%	9.3%	9.6%	9.4%
Central	1.5%	13.5%	14.9%	11.7%	6.4%	9.6%	8.4%	8.9%
Eastern	6.4%	13.5%	13.7%	11.4%	5.3%	10.1%	11.1%	10.7%
Greater Accra	0.5%	1.8%	5.8%	18.1%	45.8%	14.4%	15.4%	16.3%
Northern	32.9%	9.6%	7.0%	4.2%	2.3%	11.2%	9.6%	10.1%
Upper East	21.2%	3.7%	1.4%	1.5%	1.8%	5.9%	4.9%	4.2%
Upper West	7.1%	3.0%	1.7%	1.3%	0.5%	2.7%	3.0%	2.8%
Volta	8.7%	12.7%	13.3%	7.3%	3.4%	9.1%	8.6%	8.6%
Western	4.3%	12.8%	10.6%	10.0%	9.9%	9.5%	10.2%	9.6%

7 Appendix B: Additional Information

This appendix holds additional statistical tables and robustness checks. It presents: summary statistics for the variables used in the regressions, a table showing that the richest people are more likely to live in regions that hold capital cities, scatter plots showing how logging the wealth share variables improves fit but does not significantly alter the results, an analysis replicating Table 3 but carried out on a smaller sample of countries, an analysis replicating Table 3 but including a control for regional inequality, three different analysis that take into account the fact that the dependent variables that are represented in percentages (share of total value of aid or share of total number of projects) are bound at 0 and 1, a replication of Table 3 that re-weights each region so that each country contributes equally to the analysis, a set of tests that runs the regressions from Table 3 while sequentially excluding each country in the sample, an analysis of ethnic aid targeting, a selection model that examines the factors that influence which regions get any aid, a comparison of poverty targeting across the two donors, a table that counts the total number of tests and the number of tests that found significant results for the poorest and richest wealth quintiles, and finally a map of the aid projects under study.

To start, Table 7 shows the summary statistics for variables included in the regressions. Variables that had true zeros or ones are expressed without decimal points. The first three variables are the dependent variables from the main analysis.

Table 7: Summary Statistics

Variable	Obs	Mean	Std. Dev.	Minimum	Maximum
% of Aid Value	195	0.0872	0.1574	0	1
% of Count of Projects	195	0.0872	0.1483	0	1
ln(Aid Value)	195	1.5928	2.6612	-2.3026	7.1862
% of Poorest	195	0.0872	0.0976	0	0.4959
% of Richest	195	0.0872	0.1289	0.0014	0.6962
ln(% of Poorest)	195	-3.2767	1.6087	-6.9078	-0.6994
ln(% of Richest)	195	-3.2480	1.3054	-6.0529	-0.3607
Capital	195	0.0872	0.2828	0	1
% of Battles	195	0.0769	0.1880	0	1
Battles	195	5.3436	21.2995	0	227
% of area	195	0.0872	0.0934	0.0002	0.5263
ln(Area)	195	9.9851	1.7336	4.0584	13.3468

Table 8 provides support for the argument, expressed in the main text, that the rich are more likely to live in regions with capital cities. The unit of observation is the region and the dependent variable is the region's share of the wealthiest people in a country (*% of Richest*). The regressions are estimated with OLS and include country fixed effects. The capital dummy is substantively large and statistically significant. In these countries, the capital dummy alone explains about half of the within-country variation in the location of the richest quintile of the population.

Table 8: Explaining Where the Wealthy Live

	(1)	(2)
Capital	0.30*** (0.050)	0.27*** (0.054)
ln(Area)		-0.01 (0.009)
Fixed Effects	Yes	Yes
R-squared	0.50	0.51
Number of Regions	195	195
Number of Countries	17	17

Robust standard errors clustered on countries in parentheses

*** p<0.01, ** p<0.05, * p<0.1

In footnote 45, I note that I take the log of the wealth share variables because they are skewed and that the logged variables produces the best fit. This is supported by Figure 4, which shows the bivariate relationship between the log of the total dollar amount of aid to each region and either the fraction of the poorest or richest quintile or the log of the fraction of the poorest or richest quintile. In all cases the variables are demeaned using the country-level means of the relevant x and y variables. This allows me to graphically show all data points in a scatter plot while still maintaining the logic of a country fixed effects regression. As noted in the text, aid does not flow to places with more of the poorest people but it does flow to areas with more of the richest. Both results hold with and without the log transformation of the quintile variables, though the fit is clearly better when the quintile variables are logged.

Figure 4: Scatter plots of bivariate results (demeaned).

Table 9 replicates Table 3 but drops any country that has fewer than five regions or received fewer than 5 aid projects during the two years under study. This drops Malawi, Sierra Leone, and Niger from the analysis. The results for the rich stay the same while the (already weak) results for the poor are weakened further.

Table 9: Main Analysis on Trimmed Sample

	(1) Value	(2) Projects	(3) ln(Value)
% of Poorest	0.28 (0.207)	0.20 (0.160)	
ln(% of Poorest)			0.10 (0.095)
% of Richest	0.61** (0.236)	0.61** (0.231)	
ln(% of Richest)			0.72*** (0.215)
Capital	0.03 (0.079)	-0.03 (0.062)	1.20 (0.837)
% of Battles	-0.06 (0.080)	-0.01 (0.073)	
Battles			0.00 (0.002)
% of Area	0.33 (0.269)	0.45* (0.244)	
ln(Area)			0.36*** (0.081)
Fixed Effects	Yes	Yes	Yes
Number of countries	14	14	14
Number of regions	180	180	180
R-squared	0.22	0.23	0.25

Robust standard errors clustered on countries in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 10 replicates Table 3 but includes a control for within-region inequality. The inequality measure is not typical, as I do not have absolute measures of wealth but rather a division of people into quintiles of the population according to wealth. The inequality control is thus a measure that captures the extent to which the region has more people at the first and fifth quintile relative to the middle quintile. The exact formula is: % of Poorest + % of Richest - 2 * % of Middle. While this measure makes sense on its own, in practice almost no regions have high shares of the poorest and richest and also low shares of people in the middle quintile. One of the only examples of this kind of inequality is Katanga in the DRC. Katanga has 17.5% of the poorest quintile, 17% of the richest quintile, and 7% of the middle quintile.

In general, the richest quintile tends to live in regions with below averages shares of the middle quintile while the poorest quintile tends to live in regions with above average shares of the middle quintile (see Table 2). This implies that the inequality measure is mostly being driven by regions that have high shares of the rich and low shares of the middle. These are usually capital regions.

Table 10: Main Analysis with Inequality Control

	(1) Value	(2) Projects	(3) ln(Value)
% of Poorest	0.30 (0.179)	0.26* (0.134)	
ln(% of Poorest)			0.10 (0.085)
ln(% of Richest)	0.67*** (0.186)	0.76*** (0.186)	
ln(% of Richest)			0.73*** (0.195)
Capital	0.04 (0.071)	-0.01 (0.056)	1.25 (0.927)
% of Battles	-0.09 (0.082)	-0.04 (0.080)	
Battles			-0.00 (0.002)
% of Area	0.36 (0.225)	0.41* (0.203)	
ln(Area)			0.31*** (0.079)
Inequality	-0.09 (0.121)	-0.14 (0.113)	-0.48 (1.333)
Fixed Effects	Yes	Yes	Yes
Number of countries	17	17	17
Number of regions	195	195	195
R-squared	0.23	0.26	0.24

Robust standard errors clustered on countries in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The dependent variables that measure regional aid as a share of total aid have censoring at 0 and 1, which could bias the results of the analysis. Table 11 examines the data using a random effects tobit model that takes this censoring into account and shows consistent favoritism to the rich and a lack of favoritism to the poor. The dependent variable in models one and two is each region's share of the country's total dollar value of aid and the dependent variable in models three and four is the region's share of the total number of projects. Models one and three use the full sample and models two and four drop countries with fewer than five projects or regions. The effect of the richest is similarly significant ($p < 0.05$), and the effect of the poorest is similarly insignificant, in tobit models with country dummies instead of random effects, and these models are reported in the 'count of all models' discussion in the robustness section of the main text. Code to produce the fixed-effects tobit models is present in the replication files but the results are not reproduced here.

Table 11: Tobit Models

	Share of Value		Share of Projects	
	(1)	(2)	(3)	(4)
	Main	Small Sample	Main	Small Sample
% of Poorest	0.21 (0.154)	0.19 (0.178)	0.16 (0.142)	0.11 (0.162)
% of Richest	0.66*** (0.136)	0.68*** (0.147)	0.70*** (0.125)	0.67*** (0.134)
Capital	0.02 (0.063)	0.02 (0.068)	-0.04 (0.058)	-0.04 (0.062)
% of Battles	-0.08 (0.071)	-0.03 (0.078)	-0.04 (0.065)	0.02 (0.071)
% of Area	0.36** (0.164)	0.35* (0.195)	0.42*** (0.151)	0.47*** (0.178)
Random Effects	Yes	Yes	Yes	Yes
Number of Regions	195	180	195	180
Number of Countries	17	14	17	14

Standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

The tobit model used above is not ideal as it assumes that I have a censored set of observations of a variable that is not logically or practically bound. In my case, the share of aid variables are truly bound at zero and one. Another approach is to use a generalized linear model with a logistic link function and binomial distribution. This has the effect of ‘condensing’ the predictions from the model into the range of zero and one (or, similarly, one can view the logistic link as ‘spreading’ the dependent variable out along the real number line for modeling). As with all of these robustness tests, I do not intend to heavily defend the assumptions of any one model. Rather, I am emphasizing that the core result holds across a wide variety of specifications and models. I use the same specification as models one and two in Table 3 but rather than using OLS I use a GLM with a logistic link. Standard errors are clustered on countries. For clarity, I present the results graphically in Figure 5 rather than in tabular form. The dependent variable in the top two panels is each region’s share of the total count of projects per country. The dependent variable in the bottom two panels is each region’s share of the total cost of projects per country. As is evident from the graphs, the effect of the richest is highly statistically significant ($p < 0.01$) and the effect of the poorest is not.

Figure 5: Relationship between wealth shares and aid shares.

Given that regions are the unit of analysis, countries with more regions will contribute more to the analysis if all regions are weighted equally. Table 12 presents results using weights that make each country count equally in the analysis rather than each region. This is done by giving each region a weight that is equal to 1 divided by the total number of regions in each country. The effect of the poorest is significant ($p < 0.05$) in only model 2 (share of projects). This result goes away if the percent of poorest variable is logged (not shown). The effect of the richest segment of the population on aid is consistently significant.

Table 12: Weighted Regions

	(1) Value	(2) Projects	(3) Ln(Value)
% of Poorest	0.35* (0.186)	0.31** (0.120)	
ln(% of Poorest)			0.13 (0.095)
% of Richest	0.66** (0.227)	0.81*** (0.232)	
ln(% of Richest)			0.73*** (0.131)
Capital	0.00 (0.075)	-0.08 (0.064)	0.82 (0.525)
% of Battles	-0.08 (0.106)	-0.05 (0.096)	
Battles			0.00 (0.002)
% of Area	0.39 (0.253)	0.47* (0.223)	
ln(Area)			0.23** (0.085)
Constant	-0.04 (0.036)	-0.06 (0.035)	1.70 (1.114)
Fixed Effects	Yes	Yes	Yes
Number of countries	17	17	17
Number of regions	195	195	195

Robust standard errors clustered on countries in parentheses
 Obs. (regions) weighted so each country contributes equally
 *** p<0.01, ** p<0.05, * p<0.1

I now show that the results are robust when individual countries are excluded from the sample. Figure 6 shows coefficients and 95% confidence intervals for % of Poorest and % of Richest for a set of regressions where each country in the dataset is sequentially excluded. This implies that the first estimate for each coefficient reports the result when Benin is excluded, the second reports results when dropping the DRC, and so on. The left pane is based on model one in Table 3 and the right pane is based on model two. Model three is presented on the following page. The figure examines if the results are sensitive to the exclusion of possibly outlying countries. The results for the richest segment of the population are always significant. The results for the poorest become significant at $p < 0.05$ in only model two if Namibia or Guinea are excluded.

Figure 6: Aid targeting to the poorest and richest, dropping one country at a time

Figure 7 is the same as Figure 6, but it is based on the preferred specification of model three (logged variables) in table 3. As before, each point estimate per coefficient corresponds to one regression and each excludes one country from the sample. The logged model is less sensitive to dropping countries. In no regression is the flow of aid to the poorest significantly different from zero. All regressions show significant effects for the richest, though when Nigeria is excluded the point estimate of $\ln(\% \text{ of Richest})$ drops to 0.45 and the p-value increases to 0.024. Across all of the manipulations in all of the models, the richest are always significantly favored with aid.

Figure 7: Aid targeting to poorest and richest, logged variables, dropping one country at a time

To examine ethnic aid targeting, I replicated Table 3 but added a dummy variable that took a value of one if the country’s president (in 2008) was born in the region,⁶³ I then run another model where I interact the dummy marking the president’s birth region with the share of the poorest and richest people to see if the poorest or richest people within the president’s home region are disproportionately favored with aid. As before, the richest people are consistently favored but sharing the president’s ethnicity—as proxied by being in a region that holds the president’s hometown—is not an important factor in explaining the location of aid projects.

Model 1 in Table 13 is similar to Model 1 in Table 3, but it includes the Pres. Birth dummy and has a smaller sample size. Model two interacts Pres. Birth with % of Poorest and % of Richest, which reveals if the rich or poor are favored more if they live in the President’s hometown. None of the hometown variables or interaction terms are statistically significant. Models 3 and 4 carry out the same analysis on the natural log of the total value of aid per region, and the results are similar. As before, the wealth variables are consistently important but sharing the president’s ethnicity—as proxied by being in a region that holds the president’s hometown—is not an important factor in explaining the location of aid projects. When the dependent variable is the share of the total number of projects instead of the share of the value of aid and models 1 and 2 in table 13 are estimated, the % of Richest p-value is consistently less than 0.05. Pres. Birth and % of Poorest are not statistically significant in these regressions. These results are included in the replication code and counted in the table that counts all tests, but are not presented here.

⁶³Using a dummy for the president’s home region has the benefit of reducing ambiguity around ethnicity (especially when some presidents come from mixed backgrounds), but it drops countries where the president was not born inside the country. I also drop countries where the president in 2008 lost power before the end of 2010. This results in Ghana, Guinea, and Zambia being dropped from this portion of the analysis. Ghana had an election in 2008 and John Kufuor of the New Patriotic Party (NPP) was term limited. Nana Akufo-Addo ran for the NPP and lost, giving Ghana its second electoral turnover. Lansana Conté, former president of Guinea, died in 2008. Rupiah Banda was president of Zambia from 2008 to 2011, but he was born in what is now Zimbabwe.

Table 13: Ethnicity Interactions

	(1) Value	(2) Interaction	(3) ln(Value)	(4) Interaction
Pres. Birth	0.12 (0.075)	0.09 (0.115)	0.67 (0.512)	1.34 (1.216)
% of Poorest	0.21 (0.205)	0.17 (0.220)		
Pres. Birth $\times$ % of Poorest		0.56 (0.430)		
% of Richest	0.41** (0.164)	0.39** (0.172)		
Pres. Birth $\times$ % of Richest		-0.40 (0.441)		
ln(% of Poorest)			0.07 (0.092)	0.07 (0.098)
Pres. Birth $\times$ ln(% of Poorest)				0.09 (0.201)
ln(% of Richest)			0.69** (0.235)	0.67** (0.247)
Pres. Birth $\times$ ln(% of Richest)				0.16 (0.299)
Capital	0.01 (0.070)	0.03 (0.072)	0.84 (0.844)	0.86 (0.867)
% of Battles	-0.14 (0.090)	-0.11 (0.094)		
Battles			-0.00 (0.002)	-0.00 (0.002)
% of Area	0.45* (0.227)	0.43* (0.237)		
ln(Area)			0.35*** (0.095)	0.35*** (0.092)
Fixed Effects	Yes	Yes	Yes	Yes
Number of countries	14	14	14	14
Number of regions	168	168	168	168
R-squared	0.22	0.24	0.24	0.24

Robust standard errors clustered on countries in parentheses

*** p<0.01, ** p<0.05, * p<0.1

On the suggestion of one reviewer, I now examine if the fraction of rich or poor people influences the likelihood of a region getting any aid. Accordingly, the dependent variable now takes a zero if a region received no aid and one if a region received at least one aid project. I shows results with and without control variables. While a (conditional) fixed-effects logistic regression is appropriate given the binary nature of the dependent variable, it also leads to separation. This means that 7 countries (66 regions) are dropped because all regions received at least one aid project.⁶⁴ To show robustness, I present results in Table 14 using a logit model and an OLS (linear probability) model. Regions with more of the richest people are consistently more likely to receive at least one aid project, and the results are substantively large (e.g. Models 1 and 2 have odds ratios above 2). The effect of the rich is similar but a little weaker in the linear probability models. The estimated effect of the poorest is almost exactly zero.

⁶⁴The countries dropped due to separation are: DRC, Kenya, Malawi, Mali, Rwanda, Sierra Leone, and Tanzania.

Table 14: Binary DV

	(1) Logit 1	(2) Logit 2	(3) LPM 1	(4) LPM 2
ln(% of Poorest)	0.05 (0.190)	-0.12 (0.230)	-0.01 (0.019)	-0.04 (0.024)
ln(% of Richest)	0.85*** (0.272)	0.95*** (0.308)	0.08*** (0.026)	0.08*** (0.026)
Capital		1.34 (1.452)		0.04 (0.079)
Battles		-0.03 (0.036)		-0.00 (0.000)
ln(Area)		0.58** (0.271)		0.05* (0.025)
Fixed Effects	Yes	Yes	Yes	Yes
Number of countries	10	10	17	17
Number of regions	129	129	195	195
R-squared			0.08	0.11

Robust standard errors clustered on countries in parentheses

*** p<0.01, ** p<0.05, * p<0.1

I now demonstrate that the main results are qualitatively similar when the two donors are analyzed separately. Figure 8 replicates the right (logged variable) portion of Figure 3, but analyzes the World Bank and African Development Bank resources separately. The results are generally similar, with the poorest two quintiles not being statistically significant and the top two having significant effects. Across both donors, aid is flowing to the richest and not the poorest.

Figure 8: Bivariate (fixed effects) results, by donor.

Table 15 shows that the degree to which each donor sends aid to the richest or poorest remains similar when most control variables are added. The dependent variable in models 1 and 2 is logged aid per region from the African Development Bank and models 3 and 4 show logged aid from the World Bank. Models 1 and 3 drop the capital city control, and models 2 and 4 include the control. Neither donor's targeting to the rich is due to the level of conflict in the region nor the size of the region. However, the table reveals that the African Development Bank's skew to the rich is primarily

due to capital city bias, while the World Bank's skew to the rich is not. While this difference in mechanism is interesting, neither donor has a pro-poor aid allocation within the countries under study.

Table 15: Tareting By Donor

	(1) ADB NoCap	(2) ADB	(3) WB NoCap	(4) WB
ln(% of Poorest)	0.51 (0.309)	0.64* (0.320)	-0.01 (0.096)	-0.03 (0.111)
ln(% of Richest)	0.77*** (0.227)	0.51* (0.245)	1.09*** (0.301)	1.14*** (0.342)
Capital		3.16*** (0.860)		-0.58 (1.245)
Battles	0.01** (0.006)	0.02** (0.006)	-0.01 (0.008)	-0.01 (0.008)
ln(Area)	-0.19 (0.300)	-0.05 (0.307)	0.52** (0.196)	0.49** (0.215)
Fixed Effects	Yes	Yes	Yes	Yes
Observations	182	182	163	163
Number of countries	16	16	13	13
R-squared	0.09	0.13	0.21	0.21

Robust standard errors clustered on countries in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 16 counts all of the times that I tested for the statistical significance of the relationship between either the poorest or richest wealth quintiles and aid. It also records how many times each test yielded a statistically significant ($p < 0.05$) result. The “Num. Test” column is not a count of the number of unique regressions but rather a count of the number of unique statistical tests for the effect of the richest or poorest quintile on aid (both wealth variables were tested an equal number of times). The models are listed in the order that they appear in the text and appendix. While the text and appendix do not present an exhaustive specification search, these results should increase our confidence in the relationships reported in the main text.

Table 16: Counts of Models and Significant Results

Model	Poor Sig.	Rich Sig.	Num. Tests
Bivariate	0	3	3
Table 3	0	3	3
Bivariate Scatter	0	1	1
Trimmed	0	3	3
Unequal	0	3	3
RE Tobit	0	4	4
FE Tobit	0	4	4
GLM Logit	0	2	2
Re-weighted	1	3	3
Drop 1, Share of Value	0	17	17
Drop 1, Share of Project	2	17	17
Drop 1, log Value	0	17	17
Pres. Interaction	0	6	6
Binary DV (selection)	0	4	4
Disag. Donor Bivariate	0	2	2
Disag. Donor	0	3	4
Sum	3	92	93
Percent Significant	3%	99%	

Finally, Figure 9 maps the regional-level (ADM1) boundaries of the countries in the sample (in black) and the location of many, but not all, of the aid projects used in the analysis. More specifically, the map plots all aid projects—covering both donors and both years under study—provided that the project could be geolocated at a level of precision that was better than the regional level (a level of precision of less than 4 in the AidData coding scheme). The analysis in the text uses all projects with a precision coding of less than 5, meaning that it include projects that were geolocated to a region but where the precise location of the project within the region is unknown. It makes little sense to plot these regionally-geolocated projects in the map and so they were dropped for this purpose only.

Figure 9: Regions in the Sample and Aid Projects.