



Poor targeting: A gridded spatial analysis of the degree to which aid reaches the poor in Africa

Ryan C. Briggs

Department of Political Science, Virginia Tech, 531 Major Williams Hall (0130), 220 Stanger St, Blacksburg, VA 24061, United States



ARTICLE INFO

Article history:

Accepted 11 October 2017

Available online 20 November 2017

Keywords:

Foreign aid

Aid targeting

Africa

World Bank

African Development Bank

Poverty

ABSTRACT

The effects of most kinds of project aid decline over distance, so in order for aid to help the poor it must be targeted to the places where the poor live. This paper examines if aid from the World Bank and African Development Bank flows to the relatively poor within African countries. The unit of analysis is an approximately 50 km × 50 km cell that is tiled across the continent to form a grid of roughly 10,500 cells. I aggregate geotagged aid from each donor into the grid cells in three ways: a binary measure marking if a cell received any aid, the count of aid projects per cell, and each cell's dollar value of aid. I operationalize poverty at the grid cell level in five ways: light at night, mean travel time from the cell to a major city, distance from the centroid of the cell to the recipient's capital city, and cell-level estimates of child malnutrition and infant mortality. I test for the influence of each poverty variable in models that control for the population within each cell and include recipient country fixed effects. Aid flows to richer rather than poorer cells.

© 2017 Elsevier Ltd. All rights reserved.

If foreign aid is to alleviate poverty, then it should flow to the places where poor people live. This general logic underpins attempts to get donors to send more aid to poorer countries,¹ and for many kinds of aid this logic holds just as strongly within countries as it does across countries. While analyses of aid targeting have traditionally been cross-national (e.g. [Alesina & Dollar, 2000](#); [Nunnenkamp & Thiele, 2006](#)), recent work has begun to examine within-country aid targeting. Prior research has shown that, at a sub-national level, aggregate foreign aid does not target poverty in China ([Zhang, 2004](#)), India ([Nunnenkamp, Öhler, & Sosa Andrés, 2016a](#)), Kenya ([Briggs, 2014](#)), Malawi ([Nunnenkamp, Sotirova, & Thiele, 2016b](#)), or across a number of countries in Africa ([Öhler & Nunnenkamp, 2014](#); [Briggs, 2017](#)).² The present article extends this research by examining aid targeting by the World Bank (WB) and African Development Bank (ADB) at a high level of spatial detail across the continent of Africa. It finds that aid does not flow to poorer people within countries. Rather, aid appears to flow to the places that hold the relatively rich.

E-mail address: ryancbriggs@vt.edu

¹ For example, the Commitment to Development Index by the [Center for Global Development \(2015\)](#) favors donors that send more aid to poorer countries.

² [Nunnenkamp et al. \(2016a\)](#) find evidence of needs-based targeting when looking at some individual sectors such as health and sanitation and [Jablonski \(2014\)](#) reports limited evidence of needs-based targeting in Kenya along with clear evidence of politically-oriented aid targeting.

1. Sub-national aid targeting

The goal of aid, at least from the point of view of the two donors under study in this paper, is to reduce poverty. The World Bank's President Jim Yong Kim describes the mission of the Bank "as ending extreme poverty by 2030 and boosting prosperity among the poorest 40 percent in low- and middle-income countries" ([Clemens & Kremer, 2016](#), p. 60).³ The ADB is less direct, but notes that their goal is to "promote sustainable economic growth and reduce poverty in Africa." ([African Development Bank Group, 2014](#)). Not only do multilateral donors like the WB or ADB want to use aid to reduce poverty, they are thought to be uniquely well positioned to do so because they are more shielded from political influence than are bilateral donors ([Rodrik, 1996](#); [Martens, Mummert, Murrell, & Seabright, 2002](#)). Evidence on cross-national aid targeting supports this claim, as multilateral donors tend to send more aid to poorer countries than do bilateral donors ([Maizels & Nissanke, 1984](#); [Dollar & Levin, 2006](#); [Nunnenkamp & Thiele, 2006](#)).

While cross-national aid targeting has received more attention than sub-national aid targeting in both research and policy circles, this probably owes more to data availability than to theory. While

³ Poverty reduction has been a major goal at the Bank for some time. For example, the goal of eradicating absolute poverty was stated by McNamara in 1973 ([Clemens & Kremer, 2016](#), p. 60). The idea that aid should reduce poverty is also clearly laid out in [Assessing Aid \(World Bank, 1998](#), p. 38) and reiterated for the Sustainable Development Goals on the WB's web page ([World Bank, 2015](#)).

accessible sub-national aid data are fairly new,⁴ the theoretical argument for targeting aid to poorer places within countries is not. In fact, the argument for sending aid to poorer places within countries is essentially the same as the argument for sending aid to poorer countries.⁵ The argument is as follows: If the goal of aid is to help the poor and if aid provides a good or service that is geographically bound, then aid should be targeted to where poor people live. Thus, money that funds a global public good like vaccine research can be spent anywhere. Aid for national-level programmes should be directed to countries with poorer people. Similarly, aid projects that provide club or local public goods such as health clinics, roads, village electrification, clean water, or schools should be targeted to the places within countries where poorer people live.⁶ Thus, there is no theoretical reason for working at the level of countries when examining the degree to which donors target aid to the poor. If aid is intended to reduce poverty, then it should be targeted to poverty according to the geographic scope of its anticipated effect.⁷

The issue of subnational aid targeting is more important if rich and poor people are spatially segregated within countries. This is the case in Africa, where “inequalities between regions as well as urban and rural areas are large” (Bigsten, 2016, p. 6). These differences extend beyond income. For example, “skilled health personnel attend 83% of births in urban areas” of Burundi but only 16% in rural areas (Sahn & Stifel, 2003, p. 583). Not only do “living standards in rural areas lag far behind those in urban areas” but there is also no evidence of convergence (Sahn & Stifel, 2003, p. 591).⁸ The issue of spatial inequality would be blunted if the poor were able to easily move to richer parts of countries, but the poor infrastructure that entrenches spatial inequalities also makes migration difficult. In Africa, rates of internal migration from poorer rural areas to richer urban areas are growing but are low by global standards (De Brauw, Mueller, & Lee, 2014).⁹ Thus, while aid targeting is important in general, it is especially important if one wants to reach the very poor in Africa.

Given that multilateral donors like the WB and ADB are more insulated from politics and send more aid to poorer countries than most other donors, one might expect that they would also be sensitive to poverty within countries. This should especially be the case for project aid, which is fairly easy for donors to target to specific places (Briggs, 2017). Somewhat counter-intuitively then, research on sub-national aid targeting by these donors has typically found that aid does not target poorer subnational regions within recipient countries (Öhler & Nunnenkamp, 2014; Briggs, 2017).¹⁰

⁴ Standardized, sub-national aid data have only recently started to become available thanks to the geocoding efforts of AidData.

⁵ One can also draw a distinction between sending aid to countries with more poor people rather than countries that are poorer on average (Kanbur & Sumner, 2012). For the purpose of the present paper, this distinction is immaterial.

⁶ The World Bank regularly provides this kind of aid. For example, a pamphlet from the International Development Association, the concessional side of the World Bank, notes “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, 2014, p. 11).

⁷ It is possible that an aid-funded good like a road could be built some distance from where poor people live but could still eventually improve some economic outcome for the poor due to second-order effects. However, the same argument applies to cross-national aid targeting and is usually dismissed. For example, aid that aims to improve a national-level variable like the institutional environment in a country could be spent on a richer country that neighbors a poorer one and, if it works, then the poorer one may benefit from having a richer neighbor.

⁸ The importance of spatial inequalities in Africa is also noted in Beegle, Christiaensen, Dabalen, and Gaddis (2016) and van de Walle (2009).

⁹ Between 1990 and 2000, the population-weighted rural-urban migration rate in sub-Saharan Africa was only about 1% (De Brauw et al., 2014, p. 34) I’d like to thank a reviewer for raising the issue of internal migration.

¹⁰ Investigations of sub-national aid targeting by only the WB find similar results (Zhang, 2004; Nunnenkamp et al., 2016a).

2. Contributions to the literature

The present paper builds on prior research by making one theoretical intervention, one methodological contribution, and two empirical improvements. The theoretical intervention is the simple but often overlooked point that aid cannot help the poor unless it both works *and reaches the places where poor people live*. This implies that even if we have certain knowledge that aid improves the lives of the people that get aid, we cannot claim that aid is helping any specific group (such as the poor) unless we also know that group is in fact receiving aid. This point is especially important when studying aid to Africa, as spatial inequalities on the continent are high and migration from poorer to richer areas within countries is relatively low. Where aid goes within countries is thus quite important, as it tells us who can benefit from aid (and who cannot).

The paper’s methodological contribution follows from the prior theoretical point. The question “Within countries, does foreign aid flow to the places where poorer people live?” is fundamentally descriptive, and so I answer it descriptively. As such, I am not concerned with identifying the effect of poverty on aid allocation. Instead, the goal is to describe the spatial relationship between aid and poverty at a high level of detail and over a large number of countries. The descriptive focus of this work implies that omitted variable bias, including bias caused by spatially correlated variables, is not a concern. While this analysis will not reveal the causal effect of poverty on aid, it will reveal if poorer people are more likely to receive aid than richer people. It may well be that aid does not flow to poorer people because of variables that correlate with poverty rather than because aid is actively avoiding poverty itself. However, given the descriptive question of the paper, controlling for such variables may produce misleading results. Put differently, aid can help the poor only if it reaches the poor—and from this point of view it does not matter if the mechanism causing it to reach the poor is something other than their poverty.

To see this clearly, consider a program whose goal is to give cash to the poorest people in a country and imagine that we have a dataset of all people in the country, their income, how much money they received from the program, and whether or not the person is co-ethnic with the president of the country. It could be the case that the bivariate relationship between how much cash a person received and their income is negative but that this relationship flips to become positive when controlling for co-ethnicity, which itself is positive.¹¹ This could occur if aid targeting is influenced by both co-ethnicity and income but that the president’s co-ethnic group is much poorer than most other groups in the country. This situation is depicted graphically using simulated data in Fig. 1. The solid line shows the relationship between benefit payout and income across all people. The dashed lines show the relationship between benefit payout and income within co-ethnics and within non-coethnics.

This example highlights that if one wants to answer a descriptive question like “within countries, are poorer people receiving more aid than richer people?” then adding a list of standard control variables can produce misleading results. In the present example, if we condition on co-ethnicity then we might incorrectly conclude that more benefits go to richer people. Control variables can be useful if one is interested in estimating the causal effect of poverty on aid.¹² However, control variables are not useful in all situations and their value depends on the specific question being asked.

¹¹ This example is essentially a restatement of Simpson’s paradox (Simpson, 1951).

¹² However, in order to estimate a causal effect in this way from observational data one needs to make exceptionally strong assumptions. For a good summary of the pitfalls of simply adding control variables in an effort to recover causal effects, see the discussion and citations in Samii (2016).

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات