5-2012

Targeting maps: An asset-based approach to geographic targeting

Corey Lang

University of Rhode Island, clang@uri.edu

Christopher B. Barrett

Felix Naschold

Follow this and additional works at: https://digitalcommons.uri.edu/enre_facpubs

Part of the [Agricultural and Resource Economics Commons](#)

Terms of Use

All rights reserved under copyright.

Citation/Publisher Attribution

Lang, C., Barrett, C. B., & Naschold, F. (2013). Targeting Maps: An Asset-Based Approach to Geographic Targeting. *World Development*, 41, 232-244.

Available at: <http://dx.doi.org/10.1016/j.worlddev.2012.06.006>

This Article is brought to you for free and open access by the Environmental and Natural Resource Economics at DigitalCommons@URI. It has been accepted for inclusion in Environmental and Natural Resource Economics Faculty Publications by an authorized administrator of DigitalCommons@URI. For more information, please contact digitalcommons-group@uri.edu.

Targeting maps: An asset-based approach to geographic targeting^{*}

Corey Lang, Christopher B. Barrett and Felix Naschold
May 2012 revised version

Abstract

Proper targeting of policy interventions requires reasonable estimates of the benefits of the alternative options. To inform such decisions, we develop an integrated approach stemming from the small-area estimation literature that estimates the marginal returns to a range of assets across geographically defined subpopulations. We create a series of maps that can be overlaid with traditional poverty maps to identify strong candidate areas for intervention, though an efficiency/equity tradeoff sometimes exists. We apply our method using recent Ugandan data. Results are consistent with independent empirical findings and suggest asset specific transfer schemes would improve with a spatially targeted strategy.

JEL classification: R12, O2, C15, I32

Keywords: geographic targeting, assets, poverty maps, spatial variation, Uganda

^{*} We appreciate helpful discussions with Nancy Johnson, GIS data assistance from John Owuor and Ugandan data advice from Thomas Emwanu. Useful comments were received from seminar participants at Cornell University, the International Food Policy Research Institute, the International Livestock Research Institute (ILRI) and the World Bank and two anonymous referees. Financial support was provided by ILRI and the National Science Foundation Expeditions grant in computational sustainability (number 0832782). Any remaining errors are the authors' sole responsibility. Email addresses: clang@mail.uri.edu (Lang), cbb2@cornell.edu (Barrett), fn23@cornell.edu (Naschold).

1. Introduction

Improved targeting of development interventions has long been recognized as central to increasing the impact from poverty reduction efforts. However, effective targeting requires reasonable estimates not just of who or where the poor are, but also of where the returns to various programs are likely to be highest. Put differently, targeting concerns “what” and “where” questions every bit as much as the more familiar “who” questions. No means currently exist, however, for estimating and comparing expected benefits across space and across alternative interventions, much less of linking such estimates to the spatial distribution of poverty. In this paper, we develop a method that, first, estimates the marginal returns to a range of assets, allowing returns to vary by household and by geography and, second, maps the estimated marginal returns to the various assets, creating a visual tool that can inform the targeting decisions of an asset transfer scheme.

This paper’s motivational and methodological starting point is poverty mapping. Elbers et al. (2003) pioneered a technique that combines detailed, nationally representative household survey data with national census data to estimate poverty rates at fine levels of disaggregation for an entire country. Once estimated, the poverty rates for the different regions of a country can be used to create a poverty map, a visual representation of the spatial distribution of poverty.¹ This simple tool is popular and widely used by governments, NGOs and donors in low-income countries to guide poverty reduction efforts.²

Although poverty maps can facilitate policy discussions, they offer no explicit recommendation as to the best means of alleviating poverty. If a government is trying to reach a specific welfare target such as the Millennium Development Goals, poverty maps can at best guide the government to regions with high poverty rates. They do not, however, inform the critical subsequent choice of what exactly the government should do in that region.

Targeting maps address this crucial shortcoming of poverty maps by answering two general questions: 1) for a given region, which asset building activity will have the largest marginal gross-benefit? and 2) for a given type of asset building activity, in which regions are the

¹ The resulting poverty estimates have also been used to investigate the causes of poverty (Kam et al. 2005, Okwi et al. 2007) or its consequences (Demombynes and Ozler 2005).

² Tarozzi and Deaton (2009) and Elbers et al. (2008) evaluate the validity of poverty mapping methods using census data that include income measures. Tarozzi and Deaton argue that useful information is contained in the poverty estimates, but standard errors are too small and assumptions spatial homogeneity are too strong. We partially address the spatial homogeneity concern by including location-specific interaction terms in our model.

marginal gross-benefits largest? Good answers to either or both of these questions can improve the efficacy of targeted, asset-based development programs. Answers to the first question are paramount for those wishing to cut poverty by the most efficient means possible. The second question appeals to groups interested in investments of a specific type, such as Heifer International in building livestock holdings or The Nature Conservancy in safeguarding natural resources. With scarce resources available to finance transfers, targeting maps can help identify where poverty reduction efforts are likely to generate the most bang-for-the-buck.

This approach takes as given the desirability of geographic targeting. The idea of geographic targeting is to determine a subset of geographic regions most in need and then transfer benefits first (or only) to individuals within the chosen regions. While there are several methods of targeting aid, such as a proxy-means tested targeting, community-based targeting, categorical or indicator targeting and self-targeting, the empirical evidence suggests that geographic targeting is particularly effective for poverty alleviation (Coady et al. 2004, Baker and Grosh 1994) and is easier and less expensive to monitor and administer than other methods (Bigman and Fofack 2000).

The major disadvantages to geographic targeting are that non-poor individuals living in targeted regions receive benefits (leakage) and poor individuals not living in targeted regions do not receive benefits (undercoverage). One remedy that is routinely applied is to combine geographic targeting with additional targeting tools to limit leakage. Coady et al. (2004) survey 122 targeted transfer programs and find the mean number of targeting tools used is more than two; for example, Mexico's celebrated PROGRESA/Oportunidades program uses four (Coady 2006). A second solution is to target more finely partitioned regions. As regions become increasingly disaggregated, within region heterogeneity decreases and targeting performance increases (Elbers et al. 2007, Baker and Grosh 1994).

In this paper, we build on the proven successes of geographic targeting and propose an enhanced, asset-based approach. We explore the possibility of transfers from an entire range of private and public assets, such as livestock, mobile phones, means of transportation, and access to roads or microfinance institutions. Our focus on assets stems from the importance of a household's asset portfolio in determining the nature, extent and persistence of poverty and vulnerability (Moser 1998, Ellis and Freeman 2004, Adato et al. 2006). Further, if and where

poverty traps exist, asset transfers may push households beyond an asset poverty threshold and allow them to engineer their own escape from income poverty (Carter and Barrett 2006).

While in-kind transfers can appear paternalistic, as they constrain household choice in ways that cash transfers do not, there are several reasons why an asset-based approach could perform better than a monetary approach.³ First, imperfect markets can make it difficult for households to procure desired assets; this is a common rationale for in-kind food or seed aid in many remote or disaster-affected regions. Second, in-kind transfers may stick to the targeted households better than cash because of the well-established endowment effects associated with physical goods but not with cash. For example, the findings of Hoffmann et al. (2009) suggest that in-kind transfers of mosquito nets would result in greater use of the nets than would equivalent cash transfers. Third, some assets – especially public goods such as paved roads – are not readily available for private purchase. Fourth, in-kind transfers often enjoy greater political support than do monetary transfers. Further, monetary transfers, due to their ready divisibility, may also be subject to a high rate of social taxation compared to a lumpy asset, perhaps undoing efforts to control leakage. And in practical terms, governments and charitable organizations routinely make in-kind transfers so improving the efficacy of such interventions is desirable even if one believes cash transfers generally preferable.

The targeting maps tool introduced in this paper improves the information set guiding geographic targeting of in-kind transfers. Given substantial spatial heterogeneity in poverty incidence and its causes (Emwanu et al. 2007, Okwi et al. 2007, Kam et al. 2005), there is little reason to believe that any single transfer form is best suited for all places in a country. Likewise, asset valuation is inevitably spatially heterogeneous, given the place-specificity of many complementary inputs – e.g., agro-ecological conditions that affect livestock value, economic activity that affects the returns to transportation infrastructure. If poverty and the returns to assets both vary markedly across space for a variety of geographic, institutional, policy and technological reasons, then it is desirable to exploit the predictable component of such variation in targeting asset-based development interventions. Previous research has found considerable intra-regional variation in expected returns to different development investments, such as high yielding seed varieties and roads, in Africa and Asia (Fan and Chan-Kang 2004). By customizing

³ Currie and Gahvari (2008) review the debate over monetary versus in-kind transfers, though mainly from the perspective of developed countries.

asset-based interventions to specific geographic areas, significant gains could be made in cost-effectively addressing poverty. Our approach integrates spatially-explicit estimation of the marginal benefits to multiple assets into a single framework such that inter-asset comparisons of expected marginal benefits can be made for each region and linked to spatially-explicit poverty estimates.

While poverty maps offer a ranking of areas based on need, targeting maps rank areas in terms of the size of marginal benefits. This presents the possibility that high returns may not correspond to need, and thus a tradeoff between equity and efficiency is necessary. However, this tradeoff is present regardless of whether or not targeting maps are used. Targeting maps help to quantify the tradeoff, but also highlight synergies between equity and efficiency. In this paper, we do not judge which targeting schemes are best, we merely provide flexible empirical tools that can help inform the process with the preferences of the policy maker guiding the process. Ultimately, we envision the targeting maps output being used as one of several components, including poverty maps and local knowledge, informing a targeted asset transfer plan.⁴

The method of creating targeting maps, detailed in Section 2, involves several distinct steps similar to those involved in creating a poverty map. Using detailed household survey data and spatially explicit environmental and infrastructure data, we apply multivariate regression and bootstrapping techniques to estimate the returns to various assets and to determine how the estimated returns vary across space. We then project the parameter estimates onto the broader national census data and calculate the marginal returns as a function of projected estimates and current household asset holdings, while simultaneously estimating household-specific poverty status, this latter output very similar to conventional poverty mapping. Finally, we aggregate the estimated marginal returns across households for small geographic areas and, using Geographic Information Systems (GIS), generate maps of both the magnitude and scope of estimated benefits as well as a poverty map.

In Section 2, we also discuss limitations of the methodology, largely centered around issues of endogeneity. Our estimation strategy necessarily ignores bidirectional causality between assets and welfare and unobserved household heterogeneity, both of which could bias estimates. This is a serious concern, but one that is unfortunately unavoidable in any analysis that

⁴ Local knowledge could include customs, norms, local government, supplier behavior, corruption within government, etc.

tries to answer the questions posed above. There is no feasible way to estimate marginal returns to many assets across a large geographical space with ironclad identification. We submit that an explicit, albeit clearly imperfect decision tool is better than none at all and thus that targeting maps deliver useful information that can improve the efficacy of development interventions. While it is impossible to argue a purely causal relationship, understanding how households' asset portfolios and local environment covary with their welfare can nonetheless provide useful insights to inform development interventions. Given the considerable policy and operational importance of the questions targeting maps address, this tradeoff is attractive. Perhaps future research can ameliorate this shortcoming.

We illustrate our approach using Ugandan household survey and census data. The data are discussed in Section 3. The results, discussed in Section 4, are encouraging; estimated and projected marginal benefits to asset transfers seem reasonable and show remarkable variation across space. Our results identify promising areas to target as well as indicate key assets to use in a geographic targeting scheme. Further, our results are consistent with recent Uganda-specific research regarding transportation infrastructure (Lall et al. 2009, Raballand et al. 2009). Our findings reinforce the value of geographic targeting and the importance of spatial analysis.

2. Method

We estimate average expected marginal household-level benefits to various assets across geographically defined subpopulations. In the context of this paper, assets are taken as anything whose stock can affect a household's welfare.⁵ We classify assets along two dimensions: private vs. public and targetable vs. non-targetable. Private and public goods follow standard definitions; public goods are non-rival and non-excludable and private goods constitute the rest. We delineate targetable from non-targetable assets based on whether an asset's quantity, quality or existence can be changed by an intervention. This classification results in four categories: private targetable assets (e.g., livestock holdings, literacy, land holdings), public targetable assets (e.g., source of potable drinking water, access to health clinics, road access), private non-targetable assets (e.g., education level of household head, gender of household head) and public non-

⁵ For now, we remain general about the measurement of welfare, although we use expenditure data to illustrate our method.

targetable assets (e.g., rainfall, temperature). Our method estimates the returns to all types of assets, but ultimately we are only interested in those that are targetable.

The minimum data necessary to create a targeting map are a nationally representative household survey and a census taken at about the same time. In the first step of our analysis, we compare the data available in the household survey and the census to generate a set of variables that are common to both data sets, such as demographic variables, livestock holdings and durable goods. We restrict the data in this way because we must use a regression specification for the survey data that is replicable in the census for all independent variables. Additional environmental or public good variables can and should be added when available to supplement both the survey and census data.

The second step is to use the survey data to estimate the relationship between household welfare and asset holdings, which include the variables selected in the first step as well as relevant environmental and public good variables. We assume that household welfare is a function of asset holdings and location-specific asset returns.⁶ We remain agnostic about the functional form of the asset returns equations and model the relationship between welfare and asset holdings using a second order flexible functional form. For household i in location c , we can write the general model as:

$$\ln w_{ic} = A_{ic}' R_A(A_{ic}, \bar{A}_c, B_c, Y_{ic}, Z_c) + B_c' R_B(A_{ic}, B_c, Y_{ic}, Z_c) + Y_{ic}' R_Y(A_{ic}, B_c, Y_{ic}, \bar{Y}_c, Z_c) + Z_c' R_Z(A_{ic}, B_c, Y_{ic}, Z_c) + \delta' X_{ic} + \varepsilon_{ic} \quad (1)$$

where w_{ic} = household welfare

A_{ic} = private, targetable assets

$\bar{A}_c$ = location - specific means of A_{ic}

B_c = public, targetable assets

Y_{ic} = private, non - targetable assets

$\bar{Y}_c$ = location - specific means of Y_{ic}

Z_c = public, non - targetable assets

X_{ic} = additional controls⁶

$R_j(\bullet)$ is a vector of returns to asset type $j = A, B, Y, Z$ and is the object of estimation. The functional form of asset returns allows the expected returns to each asset to depend on the stock

⁶ This specification can be thought of as permanent or structural income (Carter and May 2001, Adato et al. 2006, Naschold and Barrett forthcoming).

⁷ The place specific means, $\bar{A}_c$ and $\bar{Y}_c$, are derived from the census, eliminating sampling error.

of every other asset. For example, the returns to a head of cattle may depend on the household head's level of education, the average number of cattle owned in that region, the existence of a nearby livestock market and/or local precipitation levels. Place-specific asset means are only interacted with household levels of the same variable (i.e., average cattle holding is interacted with each household's cattle holdings, but not with each household's pig holdings or mobile phone ownership). Further, we assume the error term is composed of a location component and a household-specific component:

$$\varepsilon_{ic} = \eta_c(M_c) + \mu_{ic} = \gamma' M_c + \mu_{ic} \quad (2)$$

where $M_c = [\bar{A}_c, B_c, \bar{Y}_c, Z_c]$.

Our principal goal in this second step in constructing the targeting map is to accurately estimate the coefficients in the welfare-asset relationship. With all interactions included in Equation (1), the specification will include more than $N(N+3)/2$ right-hand-side variables, where N is the combined number of assets in A , B , Y , and Z . With so many variables, the likelihood of a spurious relationship is high, which would adversely affect the out-of-sample prediction.

With that in mind, we use stepwise iterative deletion (with a threshold p -value of 0.05) to drop variables from the specification.⁸ This in turn can lead to other problems, specifically an important variable for an asset return function or potentially even an entire asset function could be deleted erroneously based on randomness. To alleviate this concern, we bootstrap the whole process 200 times. For each of the 200 iterations, we bootstrap the sample of households from the survey and then estimate Equation (1) using stepwise iterative deletion.⁹ The regression uses weighted least squares (weighted by population expansion factors) with errors clustered at the enumeration area level.

Having thus estimated the shape of asset returns, in the third step we project the estimated coefficients from the first stage regressions onto the census data. Ultimately, however, we are not

⁸ This practice is common in poverty mapping (Okwi et al. 2006, Emwanu et al. 2007, Demombynes et al. 2007).

⁹ Poverty mapping methods often partition the data into the smallest regions for which the survey data are statistically representative and run regressions for each of those regions separately. For example, Okwi et al. (2006) and Emwanu et al. (2007) split Ugandan data into nine strata and Demombynes and Ozler (2005) split South Africa into nine provinces. The idea behind this step is to allow coefficient estimates to vary over space. In contrast, we pool all survey data into a single regression. While in our method coefficient estimates themselves do not vary over space, asset returns can vary via the large number of place-specific interaction terms. Our motivation for this choice is to explicitly take into account the influence of place-specific characteristics on asset returns. If in contrast the geographic scope of regressions was limited, the variation in some variables, especially the place-specific variables such as climate, would necessarily also be very limited leading to biased and inconsistent parameter estimates.

interested in the coefficient point estimates, but in the expected marginal household-level return for a given targetable asset, k:

$$\frac{\partial E[\ln w_{ic}]}{\partial A_{ick}} = A_{ic} \cdot \frac{\partial \hat{R}_A(\bullet)}{\partial A_{ick}} + B_c \cdot \frac{\partial \hat{R}_B(\bullet)}{\partial A_{ick}} + Y_{ic} \cdot \frac{\partial \hat{R}_Y(\bullet)}{\partial A_{ick}} + Z_c \cdot \frac{\partial \hat{R}_Z(\bullet)}{\partial A_{ick}} \quad (3)$$

For each iteration of the bootstrap, we project the coefficient estimates onto the census data and calculate the derivatives for all targetable assets.¹⁰ Combining iterations, we calculate the mean estimated marginal return for each household.

We then aggregate households over geographically defined areas and calculate statistics fundamental to the final product. First, we compute the mean and standard error of the expected marginal benefits for every geographic area and determine which areas have average marginal benefits (AMB) that are statistically significantly greater than zero (at the 5% level).¹¹ The estimated average marginal returns and their statistical significance inform essential questions about the expected magnitude of average benefits associated with specific asset transfers in particular areas. Second, we calculate the proportion of households with positive expected marginal returns for every geographic area, which reflects the scope of benefits from specific asset transfers in particular areas.

Finally, using GIS, we generate maps that display and enhance the results. Unlike with poverty mapping, no one map can summarize all of the results. This product requires a series of maps. One map can display the most beneficial asset, as judged either by the highest expected average marginal returns of any asset or the highest proportion of positive expected marginal returns of any asset, for each geographic area. This map would address question one above: for a given region, which asset building activity will have the largest marginal gross benefit? Then, maps can be made for each asset, showing either the expected average marginal returns or the proportion of households with positive expected marginal returns to that asset for each geographic area. These maps would address question two above: for a given type of asset

¹⁰ Elbers et al. (2003) use a complex simulation procedure to account for the effect of the distribution of residuals on the poverty estimates. Because our focus is the derivative of Equation (1), our method departs with theirs in this respect. We still model the error, as shown in Equation 2, to accurately control for covariates.

¹¹ The most common way to estimate the error surrounding the poverty estimates is to use parametric bootstrapping (Elbers et al. 2003, Demombynes et al. 2007). Parametric bootstrapping projects coefficient estimates onto census households by taking random draws from the distribution defined by a single set of regression coefficient estimates and their associated covariance matrix. The poverty statuses of individual households are then averaged by geographic areas. This process is repeated many times to obtain a distribution of each area's poverty. We choose instead to bootstrap the first stage estimation in order to reduce bias in the estimates, since our method puts a greater premium on the regression coefficient estimates themselves.

building activity, in which regions are the marginal gross benefits to such an investment highest? Two estimated objects, two broad targeting questions, and many assets make for a large number of maps, each catering to a different audience or targeting question. Combined with poverty maps that are naturally generated in the third step, one then has a powerful, visual set of tools for informing the geographic targeting of asset-based poverty reduction interventions.¹²

2.1. Methodological Concerns

There are unavoidable shortcomings to this approach. First, this is a partial equilibrium analysis that cannot account for general equilibrium effects. Substantial, large-scale asset transfers could affect prices, in which case the estimated marginal benefits would be inaccurate. For example, if too many cattle were transferred into an area, the market price of milk might decline and the benefits of owning a cow become less than estimated. However, this could go both ways, as substantial asset transfers could also lead to positive externalities, as would likely be the case with mobile phones or transportation infrastructure or any private asset characterized by (positive) network or technological externalities. We assume that aggregate asset transfers will typically be marginal in magnitude and therefore that partial equilibrium assumptions suffice.

In the introduction, we began to discuss endogeneity concerns. The first such concern is the dual causality between welfare and assets. Does an asset increase a household's welfare or does an increase in welfare cause a household to invest in an asset? Clearly both are plausible, and we cannot separate the two effects.

The second source of endogeneity bias comes from unobserved heterogeneity. Current asset holdings are not randomly distributed; households choose them based in part on information not available to the policy analyst. Households that perceive large returns to an asset due to such unobservables will invest in that asset, while low return households will not. This will likely bias our estimates of marginal returns upwards. However, this bias is attenuated when households face constraints on their investment patterns (i.e., credit and savings constraints, missing markets for desired assets), as is often the case in low-income countries.

¹² If poverty estimates generated using a traditional poverty map method are preferred, one could just as easily combine our marginal benefit estimates with those separately estimated poverty rates.

Additional bias could arise due to using imperfect proxies for welfare as the dependent variable. In our illustrative application, we use expenditure to proxy for welfare. This is the best available choice in the Uganda data, as in many other instances. But it is still incomplete, especially when thinking about asset investments. Some assets are acquired not because they will produce more current expenditure, but because they enhance welfare in some other way or at some future date. For example, some livestock may be held for risk prevention or social status. Further, expenditure can be correlated with asset holdings either positively (one must spend to acquire assets) or negatively (selling assets generates income which increases expenditure).

Collectively, these concerns imply that the cardinality of estimates could be biased, which would affect inter-asset comparisons and cost-benefit analysis. Unfortunately, we are unable to estimate the magnitude of the bias using only cross-sectional, observational data, and thus do not know the extent or magnitude of these effects. But at the very least, our estimates have ordinal significance for comparing the benefits of a single asset across regions. This in itself would have operational value as there exist aid organizations and government ministries that deal in only one asset and have to make intervention siting decisions routinely.

To summarize this sub-section, we deem it important to call attention to the unavoidable shortcomings of the targeting maps method. But we caution against throwing the policy analysis baby out with the statistically imperfect bathwater. We are confident that this method generates meaningful information to help fill an important void that currently plagues development policymaking and programming. Interventions today are typically planned in the absence of any empirical estimates of marginal benefits that permit comparison across space or transfer forms. Despite our method's admitted imperfections, it is a substantial improvement over the status quo.

3. Data

We apply our method using the 2002 Ugandan National Household Survey, the 2002 Ugandan Population and Housing Census and the 2002 Ugandan Community Survey, all administered by the Ugandan Bureau of Statistics (UBOS). The household survey and census are stratified by four regions (Central, East, North, and West) and an urban-rural split. For the purposes of this paper, we restrict our attention to rural households only (5,648 households in the survey and nearly 4.4 million in the census), due to their greater reliance on natural capital and the greater likelihood of spatially heterogeneous asset returns. The hierarchy for Ugandan

administrative units, from largest to smallest, is nation, district, county, sub-county, and parish.¹³ Parishes contain less than 1,000 households on average and are roughly one-fifth the population of a sub-county. There are one or two enumeration areas (EA) per parish. The household survey clustered observations at the EA level and randomly sampled (usually) ten households within the EA.

We use per adult male equivalent expenditure as our key measure of welfare. The private asset variables come from the household survey and the census.¹⁴ We use the census, the community survey and several GIS layers to create location specific public asset variables.¹⁵ From the census, we calculate measures of population density and ethnic diversity, as well as average asset holdings at the parish level. The community survey includes information on roads, market access and microfinance access. In addition, we use GIS to derive variables such as average distance to urban areas, average distance to freshwater and average annual rainfall and temperature, among others. Data layers for urban areas and water locations were provided by the International Livestock Research Institute (ILRI). Weather data were downloaded from www.worldclim.org at a resolution of 30 arc-seconds. These geographic variables are aggregated at the sub-county level, due to limitations with the GIS software.¹⁶ Table 1 lists each asset variable used in the analysis, gives summary statistics for each from the survey sample and census, and defines the variables, if warranted. Cattle and chicken are the most common livestock held. Human capital is low with, on average, five years of education for the household head and less than half of the household literate. Mobile phone ownership stands at just three percent; as a result, estimates of the marginal returns to phones are likely to be inapplicable to current Uganda, given rapid mobile phone uptake in the intervening period. The statistical support is sufficiently similar for the two datasets, supporting our out-of-sample prediction.

¹³ Table 1 in the online appendix lists how many administrative units of each type exist and the average and median number of households in each unit, for the rural areas of Uganda.

¹⁴ As stated above, we are constrained to only use variables that appear in both the census and the survey. There are several instances where potentially informative variables (e.g., mosquito net coverage of all household members) could not be included due to this limitation. This underscores the importance of planning and coordinating between household surveys and censuses.

¹⁵ Due to the incomplete coverage of the Community Survey at the parish level, these variables are aggregated to both the parish and sub-county level, and the sub-county value is joined with the household data when the parish value is unavailable.

¹⁶ Due to the small area of some of the parishes and the relatively larger size of the weather raster data, the zonal statistics could not be calculated for all parishes.

Table 2 in the online appendix lists all of the additional variables used as controls (i.e., the matrix X in Equation (1)).

In addition to the numerical comparability of the data, geographic comparability is important. The online appendix gives statistics and maps illustrating that the survey data is spatially well dispersed and provides excellent geographic coverage. Thus, we have confidence that our first stage estimates effectively represent many different geographies.

4. Results

As a first step in analyzing the results, we determine the appropriate level of aggregation for the expected marginal returns. In standard poverty mapping exercises, there is a tradeoff between geographic aggregation and precision (Elbers et al. 2003). The goal is to aggregate households into the smallest possible geographic area without sacrificing precision, which enables inter-regional comparison.

We aggregate derivatives and calculate mean marginal benefits and standard deviation of benefits for all targetable assets at three different administrative levels: county, sub-county, and parish. Table 2 gives the estimated mean standard deviations for each targetable asset. Clearly, as the area of aggregation grows so does the standard deviation. This finding contrasts with the standard inverse relationship found in poverty mapping due to the difference in our method, which first estimates household level marginal returns via simulation and then aggregates over geographic areas. Our estimates are a composite of ordinary imprecision plus inter-household variation. As the geographic scale grows, more inter-household heterogeneity is introduced and the standard deviation increases. The empirical findings unequivocally indicate that parish is the appropriate level of aggregation for our estimates.

Relevant summary statistics for estimated marginal benefits are presented in Table 3. In addition to the targetable assets, Table 3 presents statistics for the non-targetable asset household head years of education for purposes of comparison. The table begins by presenting national averages and then presents results related to between- and within-parish heterogeneity. Column 1 gives the mean of the estimated average marginal benefit (AMB) for all rural parishes. The magnitudes of estimated AMB seem reasonable. For example, motorized vehicles are more than eight times as valuable as bicycles, and livestock generally offer low returns. However, the AMB estimates of mobile phones seems inflated, which is likely a result of the fact that at the time of

the survey and census mobile phones were a scarcely owned asset and thus likely a luxury good, biasing the estimated benefits of ownership upward. In addition, several assets (goats, pigs, microfinance and road access) have mean returns less than zero. Column 2 gives the proportion of parishes with AMB greater than zero. These results tend to mirror the AMB in Column 1 in that assets with large average returns tend to offer positive returns in many parishes. While negative returns are counterintuitive, we attribute them to imprecise point estimates and substantial between-parish heterogeneity. Column 3 gives the standard deviation of AMB (between parishes), which is frequently as large or larger than mean AMB. For instance, the standard deviation of AMB across space of goats is 15 times larger than the mean. It is standard to view such large variation as undesirable, but we view this spatial variance positively and think it underscores the value of this work. Microfinance access stands out as having negative estimated returns and relatively little variation across space; this is likely a function of selection bias since microfinance services are frequently targeted to poor areas.

Columns 4 through 6 of Table 3 examine the extent and importance of within-parish heterogeneity. As discussed above with respect to Table 2, we have an estimate of the standard deviation of estimated marginal benefits across households for each parish. We can then categorize an asset as having statistically significantly greater than zero AMB. While this is an abuse of standard statistical language, it captures well the idea that policy makers may be most interested in targeting areas with consistently large returns. Column 4 gives the proportions of parishes with statistically significant AMB for each asset. There tends to be a strong correlation between the proportion of positive returns the proportion of significant returns. The returns to literacy are an outlier in this respect, as 92% of parishes have positive AMB, but only 0.3% have significant returns, which suggest substantial within-parish heterogeneity. This implies that a geographically targeted literacy program may be inefficient. Column 5 gives the mean of the estimated AMB conditional on the AMB being statistically significantly greater than zero. It is these numbers that are indicative of the possible returns that could be achieved within-parish heterogeneity is additionally used as a condition for a geographically targeted asset transfer scheme. Column 6 offers an additional measure of within-parish heterogeneity that may be useful for practitioners. It gives the mean of the proportion of households in each parish with expected marginal benefits greater than zero. Again, there tends to be a correlation between this and the other measures of estimated benefits. This statistic captures the scope of benefits, as well

as the probability of inefficient allocation of an asset. Given the substantial within-parish heterogeneity, Column 6 additionally points to the value of using multiple targeting tools to improve efficacy.

Statistics can only give a sense of the results' value. The more interesting results are the spatial distribution, heterogeneity and patterns of the estimated benefits. Figure 1 plots the estimated average marginal returns that are statistically significantly greater than zero for cattle, chickens, bicycles, and road access at the parish level. We see pockets of high returns, like those in the northwestern Uganda for cattle, as well as clear spatial patterns, such as the graded decline of marginal benefits to bicycles as one moves further interior from the Northeast border. For each asset, a considerable portion of the country does not exhibit statistically significant estimated returns, reflecting both relatively large standard errors and several negative point estimates. It is reasonable that some returns are actually negative because we estimate marginal returns comprehensively, including areas that are completely unsuitable for certain assets.¹⁷ While only 11% of parishes exhibit statistically significant estimated returns for road access, we see that those significant returns are clustered in the South-central and Southern part of the country, Uganda's most urbanized parts.

Figure 2 plots the proportion of households with estimated marginal returns greater than zero for cattle, chickens, bicycles, and road access. The maps mostly reinforce the information displayed in Figure 1. Areas with high significant returns also have a large proportion of households with positive expected returns (e.g., cattle in the northwest). While the near-monochromatic map for bicycles offers little information in terms of geographic targeting, it suggests that outside of the urban areas in South-central Uganda, positive marginal benefits are near universal.

Next, we identify which asset offers the largest estimated benefits for each parish and map the results in Figure 3. Motor vehicles and mobile phones dominate these maps, which is not surprising given the large value and expense of motor vehicles and the scarcity and luxury status of mobile phones at the time these data were collected. However, preferred asset maps can be created with any set of assets desired, excluding those that are infeasible (e.g., due to expense) or about which there exists doubt as to the validity of the estimates. To this end, we also generate maps of the assets with maximum returns, limited to livestock assets only, in the bottom panel of

¹⁷ Fan and Chan-Kang 2004 and Kam et al. (2005) also find negative estimated returns to assets in some areas.

Figure 3. Cattle, goats and chickens all have substantial presence on these maps indicating that each species is valuable, but differentially across space, resulting in geographically heterogeneous preferences.

Beyond looking at the spatial distribution of estimated marginal benefits of an asset, we examine how those benefits relate to existing holdings of that asset and to the poverty headcount rate by parish.¹⁸ The first two columns of Table 4 give the correlations of AMB conditional on significance and the proportion of households with positive marginal returns with mean asset holdings. A negative correlation between benefits and holdings suggests untapped potential, perhaps indicating the presence of a market failure. If positive, on the other hand, then asset investments are already in line with returns, but further investments could still improve welfare. The results suggest that goats, chickens, and motorized vehicles may be difficult for households to procure in areas of high estimated returns. On the other hand, the results for mobile phones, literacy, and road access suggest the presence of positive network externalities.

The second set of columns in Table 4 examines how the estimated returns correlate with poverty, which is particularly relevant for designing the distribution of aid as it reveals prospective tradeoffs and synergies between the objectives of efficiency (i.e., maximizing total expected benefits) and equity (i.e., targeting the poor). The results indicate that poverty reduction and efficiency goals align for the assets cattle, motorized vehicles, and bicycles, whereas there may be a tradeoff between equity and efficiency for the assets goats, pigs, chicken, and road access. To reiterate what was discussed in the introduction, we are not implying that our methods can judge which assets are best for use in a targeting program. Ultimately it is the policy makers that must decide which assets are best and which tradeoffs between equity and efficiency they are willing or unwilling to make. We are merely hoping to provide tools and empirical estimates that will enable an informed decision.

The analysis thus far has centered on estimated marginal gross returns; information about the costs of supplying different assets has been conspicuously absent. In order to address this deficiency and to enable explicit benefit-cost comparisons (albeit simplistically and

¹⁸ The poverty headcount rate is the percentage of the population that is poor. In Uganda, a household is deemed poor if their estimated monthly expenditure falls below the expenditure thresholds set by Emwanu et al. (2007). As a check on our method, we compare our poverty estimates to those previously estimated for Uganda using the same data from Emwanu et al. (2007), who estimated the poverty headcount rate at the sub-county level. The correlation between the two estimated poverty headcount rates is 0.85; the rank correlation is 0.83. The poverty map created using our method is shown in Figure 2 of the online appendix.

incompletely), we compare estimated benefits with actual costs for all livestock assets.¹⁹ Costs are based on the mean price of livestock purchased or sold, as reported in the household survey (costs of other assets are unavailable in the data). The cost data do not include the marginal costs of maintaining stocks; total costs of acquiring and holding an animal would be higher. Because we are unsure over what time horizon the stream of benefits would accrue and what discount rate is appropriate, we report only the expected increase in expenditure for a single month. Table 5 presents the findings.

While crude and simplistic, our approach underscores the considerable marginal returns to investment in rural Uganda. Although only pigs pass a cost-benefit comparison outright, the other livestock assets would surely pass if the timeframe was extended in accordance with an animal's expected lifespan. For instance, chicken would pass with a time horizon of three months, and cattle would pass for time horizons of 6.7 years, just two years in areas with expected returns on the high end of the distribution.²⁰

While detailed exploration of the behavioral and institutional reasons for these findings is beyond the scope of this methodological paper, the results clearly underscore apparent underinvestment in productive assets in rural Uganda. Targeting maps of this sort can help development agencies identify best bet forms for asset transfers in a specific area, given such apparent underinvestment. Such targeting maps are perhaps especially useful for geographic targeting of a specific asset transfer program (e.g., livestock or bicycles), since the costs of provision typically vary only modestly across space for a given asset.

As the final step in illustrating the potential utility of targeting maps, we detail a hypothetical chicken transfer program. We choose chickens because they are inexpensive, do relatively well in the cost-benefit comparison and one can easily imagine a development agency implementing such a scheme. We select candidate parishes based on the following three criteria: 1) expected AMB greater than 0.005 and statistically significant, 2) at least 80% of households have positive expected marginal benefit to chickens, and 3) a poverty headcount rate greater than 50%. A total of 58 parishes meet these criteria and are mapped in Figure 4. Of those, we highlight two parishes that show particular promise for this sort of development intervention, Itojo parish in the southwest and Ating parish in the northeast, based on high levels of both

¹⁹ The expected household marginal benefit was calculated using the approximation for log-linear models that a marginal return of β would increase the household's expenditure by $\beta\%$.

²⁰ These calculations assumed a 5% annual discount rate.

expected returns and poverty. This sort of simple – and very useful –geographic targeting guidance can be easily repeated for any asset included in the estimation, as well as different criteria for selection.

Having now presented the results, we now compare our results using recent empirical research on road access in Uganda. Raballand et al. (2009) examine the roads investment strategy and find that the goal of extending road access to within 2 km of every household is misguided and that larger gains exist in improving and maintaining existing roads. Lall et al. (2009) estimate a locational choice model for industry and find strong agglomeration forces that suggest investment in rural infrastructure is unlikely to benefit the rural poor in terms of job creation. Our results are consistent with these findings. Figures 1 and 2 and the correlations in Table 4 indicate that the benefits to additional road investments are largest primarily surrounding urban areas.²¹ Given the differences in methods and data between our work and that of Raballand et al. (2009) and Lall et al. (2009), the consistency in results suggests that our method generates sensible results. We found no other comparable empirical studies of spatial distribution of returns in Uganda against which we could compare our results.

5. Conclusions

This paper presents a novel method that has the potential to enhance the efficacy of geographically targeted asset transfer schemes. We add to the substantial literature on small area estimation, moving beyond estimating poverty so as to begin to identify the best means of alleviating it. Development agencies and government ministries need to know not only where the poor reside, but also what forms of transfers are most likely to help move them out of poverty. Our method first estimates the marginal returns to various assets and then creates a series of maps that can address a variety of questions regarding the magnitude and scope of benefits and the efficient spatial allocation of development programs. The results produced using Ugandan data are promising; estimated and projected asset returns seem reasonable and show substantial variation across space. When combined with a simultaneously generated poverty map, a potentially powerful geographic targeting tool emerges.

²¹ Figure 3 in the online appendix visually explores the relationship between population density and marginal returns to road access using maps and finds similar results.

While the maps and other results produced in this paper serve mainly to demonstrate the potential usefulness of this method, our hope is that the method can be eventually implemented in development programming. We envision targeting maps being part of a suite of tools, complementing local knowledge and the well-established use of poverty maps, that policy makers employ to design asset transfers in low-income countries. The results presented highlight the possibility that policy makers may face tradeoffs between equity and efficiency, that is targeting the areas most in need versus the areas with the largest returns. Our results illustrate how these goals can be balanced, but importantly these tools are flexible and targeting maps can be designed with specific needs in mind.

Continued work with additional inputs is needed to complement targeting maps. First, even if a policy maker has a targeting map in hand, there are still unanswered questions about the net benefits to and final effects of various asset transfers. We addressed some of these concerns with a limited benefit-cost analysis. A more thorough analysis for all assets with more precise information on procurement and maintenance costs, as well as asset lifespan, is a natural and straightforward exercise for agencies intending to implement a transfer scheme using targeting maps as an input.

Second, targeting maps are not an end in themselves. They estimate marginal returns, which is only an intermediate step to an end goal of poverty reduction. A natural extension of the targeting maps method is to use panel data to determining the expected impacts of an asset transfer program on poverty (or on other outcome variables of interest). Further, optimization algorithms could be constructed to maximize expected poverty reduction given a fixed budget and spatial constraints to transfers (e.g., due to logistical concerns). Lastly, we hope that the promise of these methods might also help encourage organizers of household surveys and censuses to better coordinate future questionnaires with poverty maps and targeting maps in mind.

References

- Adato, M., Carter, M. R. & May, J. (2006). Exploring poverty traps and social exclusion in South Africa using qualitative and quantitative data. *Journal of Development Studies*, 42(2), 226-247.
- Baker, J. L. & Grosh, M. E. (1994). Poverty Reduction Through Geographic Targeting: how well does it work? *World Development*, 22(7), 983-995.
- Bigman, D. & Fofack, H. (2000). Geographical Targeting for Poverty Alleviation: An introduction to the special issue. *World Bank Economic Review*, 14(1), 129-145.
- Coady, D. (2006) The welfare returns to finer targeting: The case of the Progresa program in Mexico. *International Tax and Public Finance*, 13, 217-239.
- Coady D., Grosh, M., & Hoddinott J. (2004). Targeting Outcomes Redux. *World Bank Research Observer*, 19(1), 61-85.
- Currie, J. & Gahvari, F. (2008) Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature*, 46(2), 333-383.
- Demombynes, G. & Ozler, B. (2005). Crime and local inequality in South Africa. *Journal of Development Economics*, 76, 265-292.
- Demombynes, G., Elbers, C., Lanjouw, J. O., & Lanjouw, P. (2007). How good a map? Putting small area estimation to the test. World Bank working paper 4155.
- Easterly, W. & Levine, R. (1997). Africa's growth tragedy: Policies and ethnic divisions. *Quarterly Journal of Economics*, 112(4), 1203-1250.
- Elbers, C., Fujii, T., Lanjouw, P., Ozler, B., & Yin, W. (2007). Poverty alleviation through geographic targeting: how much does disaggregation help? *Journal of Development Economics*, 88, 198-213.
- Elbers, C., Lanjouw, J. O., & Lanjouw, P. (2003). Micro-level estimation of poverty and inequality. *Econometrica*, 71(1), 355-364.
- Elbers, C., Lanjouw, J. O., & Lanjouw, P. (2005). Imputed welfare estimates in regression analysis. *Journal of Economic Geography*, 5, 101-118.
- Elbers, C., Lanjouw, P., & Leite, P.G. (2008). Brazil within Brazil: Testing the poverty map methodology in Minas Gerais. World Bank Policy Research Working Paper 4513.
- Ellis, F & Freeman, A. (2004) Rural livelihoods and poverty reduction strategies in four African countries. *Journal of Development Studies*, 40(4), 1-30.

- Emwanu, T., Okwi, P. O., Hoogeveen, J. G., Kristjanson, P., & Henninger, N. (2007). *Nature, distribution and evolution of poverty and inequality in Uganda 1992-2002*. Uganda Bureau of Statistics and the International Livestock Research Institute.
- Fan, S. & Chan-Kang, C. (2004). Returns to investment in less-favored areas in developing countries: a synthesis of evidence and implications for Africa. *Food Policy*, 29, 431-444.
- Hoffmann, V., Barrett, C. & Just, D. (2009) Do free goods stick to poor households? Experimental evidence on insecticide treated bednets. *World Development*, 37(3), 607-617.
- Kam, S., Hossain, M., Lal Bose, M., & Villano L.S. (2005). Spatial patterns of rural poverty and their relationship with welfare-influencing factors in Bangladesh. *Food Policy*, 30, 551-567.
- Lall, S.V., Schroeder, E. & Schmidt, E. (2009). Identifying spatial efficiency-equity tradeoffs in territorial development policies: Evidence from Uganda. World Bank Policy Research Working Paper 4966.
- Moser, C. (1998). The asset vulnerability framework: Reassessing urban poverty reduction strategies. *World Development*, 26(1), 1-19.
- Naschold, F. & Barrett, C.B., Forthcoming. Do short-term observed income changes overstate structural mobility? *Oxford Bulletin of Economics and Statistics*.
- Okwi, P.O., Ndeng'e, G., Kristjanson, P. Arunga, M., Notenbaert, A. Omolo, A., Henninger, N., Bensen, T., Kariuki, P., & Owuor, J. (2007). Spatial determinants of poverty in rural Kenya. *Proceedings of the National Academy of Sciences of the United States of America*, 104 (43), 16769-16774.
- Raballand, G., Macchi, P., Merotto, D., & Petracco, C. (2009). Revising the roads investment strategy in rural areas: An application for Uganda. World Bank Policy Research Working Paper 5036.
- Tarozzi, A. & Deaton, A. (2009). Using Census and Survey data to estimate poverty and inequality for small areas. *Review of Economics and Statistics*, 91(4), 773-792.

Tables

Table 1: Summary statistics of all asset variables for rural Uganda

	Survey		Census	
Number of households	5648		4376978	
Monthly household expenditure (Ugandan Shilling)	118147		-	
	Mean	St. dev.	Mean	St. dev.
Private, Targetable assets				
Cattle (head)	1.74	12.61	1.19	12.49
Goats (head)	0.33	3.76	1.00	7.22
Pigs (head)	0.09	1.15	0.15	1.24
Chicken (head)	1.87	24.75	2.37	16.93
Land ownership (1=yes)	0.29	0.45	0.16	0.36
Motor vehicle ownership (1=yes)	0.04	0.19	0.03	0.17
Bicycle ownership (1=yes)	0.47	0.50	0.35	0.48
Mobile phone ownership (1=yes)	0.03	0.16	0.03	0.16
Proportion of household literate	0.46	0.29	0.45	0.32
Public, Targetable Assets				
Microfinance access (1=yes)	0.79	0.41	0.79	0.41
Road access index	1.09	0.28	1.11	0.30
Private, Non-targetable assets				
Household head education (years)	5.06	3.74	4.55	3.85
Public, Non-targetable Assets				
Population density (per sq. km)	289.2	454.2	396.9	875.5
Ethnic diversity of parish	0.28	0.26	0.29	0.27
Existence of market in parish (1=yes)	0.57	0.50	0.56	0.50
Average distance to an urban area in parish (km)	15.7	10.8	16.0	11.4
Average distance to freshwater in parish (km)	1.98	3.58	1.85	3.16
Average annual temperature (°C)	21.83	2.01	21.86	2.02
Average annual total precipitation (mm)	1227.5	181.6	1224.5	182.9
Average precipitation in driest month (mm)	34.1	15.4	34.3	16.1

Notes: Distance is measured as Euclidean, or straight-line, distance. Motor vehicle ownership equals one if a household owns either a car or motorcycle. Ethnic diversity is calculated (as in Easterly and Levine 1997) as the probability that two people of different ethnicity meet if randomly matched. Microfinance access is derived from the Community Survey and equals one if at least one community within a parish indicated having access to microfinance services. Road access index is derived from the Community Survey, in which community respondents rate their local roads as 0 = "no roads", 1 = "seasonal roads" and 2 = "all weather roads". Responses are averaged from all communities within a parish to form the index.

Table 2: Mean standard errors of estimated average marginal returns at different levels of geographic aggregation

Asset	County	Sub-county	Parish
Cattle	0.004	0.003	0.003
Goats	0.011	0.009	0.008
Pigs	0.072	0.062	0.052
Chicken	0.003	0.002	0.002
Motorized vehicle	0.174	0.153	0.134
Bicycle	0.043	0.029	0.022
mobile phone	0.155	0.126	0.108
Proportion of household literate	0.535	0.531	0.524
Microfinance access	0.070	0.061	0.053
Road access	0.155	0.092	0.076

Table 3: Summary statistics for estimated marginal benefits to assets

Asset	Mean AMB	Proportion of parishes with positive AMB	Standard deviation of AMB	Proportion of parishes with significant AMB	Mean AMB conditional on significance	Mean PROP
	(1)	(2)	(3)	(4)	(5)	(6)
Cattle	0.004	88.0%	0.005	41.4%	0.007	81.5%
Goats	-0.001	50.0%	0.016	17.3%	0.021	49.0%
Pigs	-0.077	9.3%	0.063	0.9%	0.086	15.2%
Chicken	0.000	56.4%	0.003	30.5%	0.003	57.3%
Motorized vehicle	0.522	99.8%	0.241	96.2%	0.536	99.3%
Bicycle	0.063	87.7%	0.064	73.2%	0.086	87.1%
mobile phone	0.401	97.9%	0.148	91.5%	0.422	97.6%
Proportion of household literate	0.203	92.0%	0.172	0.3%	0.923	62.1%
Microfinance access	-0.068	2.0%	0.032	-	-	2.3%
Road access	-0.079	32.1%	0.183	11.2%	0.200	33.4%
Head education	0.009	95.6%	0.006	0.1%	0.027	66.1%

Notes: AMB stands for average marginal benefit and PROP stands for proportion of households with expected positive marginal benefits. “-” indicates a missing value.

Table 4: Correlation of estimated average marginal benefits that are significantly greater than zero and proportion of households receiving a positive expected benefit with average asset holdings and the poverty rate

Asset	Asset holdings		Poverty rate	
	Average significant benefit	Proportion positive benefit	Average significant benefit	Proportion positive benefit
Cattle	-0.02*	0.06*	0.05*	0.2*
Goats	-0.12*	-0.2*	-0.04	-0.59*
Pigs	0.6*	-0.16*	-0.37*	-0.37*
Chicken	-0.04*	-0.24*	-0.04*	-0.01*
Motorized vehicle	-0.36*	-0.57*	0.24*	0.11*
Bicycle	-0.21*	0.15*	0.04	0.25*
mobile phone	0.29*	0.06*	-0.1	0.08*
Proportion of household literate	0.01*	0.94*	0.25*	-0.8*
Microfinance access	-	-0.1*	-	0.11*
Road access	0.39*	0.74*	-0.24*	-0.36*

Notes: * indicates significance at 5% level. “-” indicates a missing value.

Table 5: Simplified cost-benefit analysis

Asset	Cost	Expected marginal monthly benefit	
		Median	95th percentile
Cattle	214,112	3,074	8,851
Chicken	3,308	1,216	2,864
Goats	16,301	10,200	18,490
Pigs	19,788	43,631	78,401

Notes: All numbers are in Ugandan Shillings. Expected benefits are only calculated for parishes that have significantly greater than zero average marginal benefits.

Figures

Figure 1: Examples of maps of estimated average marginal returns that are significantly greater than zero for the given asset.

Figure 2: Examples of maps of proportion of households with estimated positive marginal return for the given asset.

Figure 3: Maximum asset returns

Figure 4: Sample targeting exercise

Notes: Parishes are selected by meeting three criteria: 1) estimated poverty rate greater than 50%, 2) expected AMB greater than 0.0005, and 3) Proportion of households with expected positive marginal benefits greater than 80%.