UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

TARGETED MICRO-HEALTH INSURANCE IN THE DEPARTMENT OF NOUNA, BURKINA FASO

Inaugural-Dissertation
zur Erlangung des Doktorgrades
des Fachbereichs Wirtschaftswissenschaften
der Ruprecht-Karls-Universität Heidelberg

vorgelegt von Michael Schleicher

Heidelberg
Juli 2017

****

Dedicated to my Parents

****

Acknowledgements

I wish to thank Stefan Klonner for his close supervision throughout my PhD. I am grateful for his open door, constant support, encouragement, and advice from which I and my work have benefited tremendously. I also want to thank Axel Dreher for very useful and encouraging comments towards the end of my studies. My appreciation extends to my colleagues from the Institute of Public Health and from the Centre de Recherche en Santé de Nouna without whom this work would be impossible. In particular, I want to thank Rainer Sauerborn for his initiative in establishing the Nouna health research project and for his enthusiasm to collaborate across disciplines. I am thankful to Aurélia Souares for invaluable background information and support in preparing my field visits. The field work would not have been possible without Ali Sié who welcomed me in Nouna. Furthermore, I am deeply grateful to Athanase Narangoro Pacere for his close and excellent collaboration in the field and beyond. I owe thanks to my colleagues in Heidelberg who created a positive and inspiring research atmosphere. I am very grateful to my family for all the opportunities provided, which eventually allowed me to follow my interests. Finally, I thank my friends and especially Bronja for their support, patience, and guidance on this journey.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Social policies that are designed directly to improve the well-being of the poor have become increasingly popular throughout the global South (Honorati et al., 2015). International organizations such as the World Bank regard such policies as a key factor to development, as is reflected by their push to increase the number of affordable social protection systems in low-income countries (World Bank, 2012). Similarly, the first goal of the UN's Sustainable Development Goals agenda includes the objective to ensure social protection for the poor and vulnerable and to increase access to basic services.[1] Social protection programs −also known as Social Safety Nets− go beyond normal emergency relief and look to to offer individuals effective protection against adverse economic shocks, encourage investments in human capital and improve people's ability to access jobs (Devereux and Sharp, 2006). Hence, social protection programs might work as a valuable complement to income-driven economic development and have the potential to break up poverty traps (Banerjee and Duflo, 2011). Alternatively, Sen (1988) provides a more intrinsic motivation by regarding social policies and economic development as similar means to reach fundamental achievements of individual well-being such as 'being in good health' or 'being educated'.

The outreach of social protection programs has increased remarkably over the last two decades and today they involve more than 1.9 billion people in roughly 130 low- and middle-income countries (Honorati et al., 2015). Most of these programs are funded as social policies by national or subnational governments, and there is a high degree of variation in their implementation (Banerjee and Duflo, 2011). Instruments are manifold, including cash transfers (both conditional and uncon-

[1] Retrieved from `https://sustainabledevelopment.un.org/topics/sustainabledevelopmentgoals` on June 23, 2017.

ditional in nature), public works or school feeding programs, and social insurance schemes. Today, almost every Latin American country runs a conditional cash transfer program, while the large rural populations in sub-Saharan African countries rely mostly on public works and unconditional cash transfer programs (World Bank, 2014). In sub-Saharan Africa, the number of unconditional cash transfers has increased remarkably, from 21 in 2010 to 40 in 2014 (Honorati et al., 2015). In addition to program outreach, substantial heterogeneity in program quality prevails. In low- and lower-middle income countries, on average, only one-quarter of the 20 percent poorest people are covered by a social protection scheme (Honorati et al., 2015). Figure 1.1 illustrates that program coverage of the poor varies substantially across both region and program type. Overall, no region is able to to cover more than 60 percent of the poor with any type of program. With about 40 percent of households covered by at least one social insurance scheme, the East Asia and Pacific region clearly outperforms the remaining four regions of the globe, which show an average coverage of about 8 percent.[2] On the other hand, South Asia and sub-Saharan Africa, where most of the global poor live, generally, show the severest gaps in coverage overall.

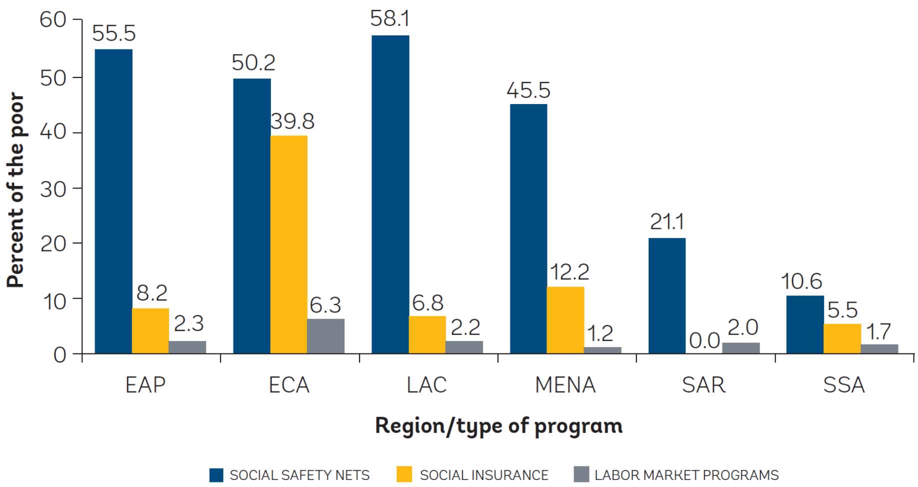

Figure 1.1: Coverage of social protection and labor (SPL) programs by region

Source: Taken from Honorati et al. (2015, p.46).

Notes: Poor households are defined as those in the poorest quintile of countries' respective consumption/income distribution. Aggregate statistics are based on countries with information on social safety net programs (105 countries). Regions: EAP = East Asia and Pacific; ECA = Europe and Central Asia; LAC = Latin America and the Caribbean; MENA = Middle East and North Africa; SAR = South Asia; SSA = Sub-Saharan Africa.

In my dissertation I shall focus on the two main reasons for low program coverage in low-income countries. The first reason lies in governments' limited abilities to accurately target the popula-

[2] Fairly high social safety net coverage in the first four regions (from the left to the right) is mainly driven by relatively high enrollment into cash transfer schemes (Honorati et al., 2015).

tion's poorest people in the first place. As a second reason, most of the social protection schemes are voluntary and face relatively low enrollment rates on the part of eligible households. These two problems are especially pronounced for welfare programs that are implemented in sub-Saharan Africa. A large fraction of the predominantly rural poor population in this area lacks access to elementary formal institutions and markets (Gertler and Gruber, 2002). It instead depends heavily on informal household- and community-based networks such that newly introduced social policies are met with skepticism (World Bank, 2004). In their recent book on safety nets in sub-Saharan Africa, Del Ninno and Mills (2015, p.3) observe that successful program implementation in this area is most severely hampered by ineffective targeting, which they broadly describe as "defining the rule and practice for allocating benefits to the most needy members of society". In line with this pattern, Monchuk (2013) observes that in many African countries safety net coverage of the poor is still low, heavily dependent on external donor support, and often temporary in nature.[3]

Apart from these implementation concerns, the amount of data availability and the emergence of sophisticated evaluation designs have remarkably improved our understanding about 'whether' and 'in which circumstance' social policies work, and work best (Banerjee and Duflo, 2011). The evolution of J-PAL, a network of scientists that focuses on randomized evaluations and policy outreach in the global South, reflects this trend. Established in 2003, J-PAL counted 240 ongoing or completed randomized evaluations by 2010 and has increased this number by a factor of 3.5 in the seven following years.[4] Latin America and the Caribbean, the pioneer region in embedding large-scale anti-poverty programs in experimental evaluation designs, accounted for one-third of all rigorously conducted impact evaluations in the world between 2000 and 2010. With 37 impact evaluations between 2010 and 2015, sub-Saharan Africa was shown to be the most-heavily evaluated region in recent years (Honorati et al., 2015). When reviewing the evidence, Honorati et al. (2015) further find that most social policies are effective in reducing consumption poverty and that recent impact evaluations confirm the positive effects on human development outcomes such as education, health and nutrition.

Given their huge potential for reduction of global poverty, the objective of improving coverage of social policies −especially in sub-Saharan Africa− appears even more pressing. The three parts of this dissertation evaluate different components of a social protection scheme in West Africa,

[3] On the other hand, there are signs that safety nets are taking hold as a core poverty reduction instrument in Africa as well. Monchuk (2013) finds that the implementation of systematic nation-wide safety nets features prominently on many governments' agendas. Rwanda, Kenya, and Tanzania emerged as a first set of countries which developed national social protection systems (Del Ninno and Mills, 2015)

[4] Retrieved from `https://www.povertyactionlab.org/` on June 24, 2017 and Banerjee and Duflo (2011).

specifically, a voluntary micro-health insurance that has started operating in northwestern Burkina Faso in 2004. Since 2007 poor households in the area can purchase health insurance at a 50 percent discount. I analyze both the underlying method applied to target poor households and the effects of the subsidy offer on health insurance take-up. The targeting procedure in the present context directly involved local communities to identify beneficiary households. While being widely applied in the field, only a few studies have systematically evaluated the accuracy or potential flaws of so-called Community-Based Targeting (CBT). Pricing of social insurance, on the other hand, is considered to be an important instrument to increase enrollment among poor liquidity-constrained households. For a relatively poor country in sub-Saharan Africa, this work thus provides novel empirical evidence on both the supply and demand side factors that are eminent to increase coverage of social policies.

A social program's primary targeting objective is to minimize the exclusion of poor households (exclusion error). Universal targeting achieves this goal but, at the same time, maximizes the number of non-poor households covered by the program (inclusion error). A limited program budget thus introduces a trade-off in the choice of the targeting quota (that is the share of households covered by the program). A higher targeting quota reduces the exclusion error, but comes at the cost of potentially smaller program effects, leaving a smaller benefit per capita for poor recipients (Del Ninno and Mills, 2015). The program's second crucial choice at the stage of targeting lies in the choice of a suitable method to identify the target group. In low-income countries, where administrative data on households' incomes ('means') are typically unavailable, targeting of welfare programs tends to rely on statistical procedures processing suitable proxies of households' means in a Proxy-Means Test (PMT) (Brown et al., 2016). Alternatively, targeting may be decentralized through Community-Based Targeting (CBT), where the choice of beneficiaries is delegated to local communities (Ravallion, 2003).[5] Both statistical and decentralized targeting methods have their pros and cons, so that from a policymaker's perspective each approach may be chosen on reasonable grounds. Existing studies mention the superior cost-effectiveness (Chambers, 1994b) and higher satisfaction rates (Alatas et al., 2012; Schüring, 2014) as two advantages of community-based over

[5] In addition, there are three more commonly applied targeting methods. First, categorical targeting determines beneficiaries with respect to one single −commonly survey-based− indicator (Sabates-Wheeler et al., 2015). Second, targeting based on self-selection introduces dis-incentives for the non-poor to enroll (Alatas et al., 2013b). Old-age pensions in South Africa (where 'age' is the relevant category variable) and a minimum wage used in India's national work employment guarantee scheme, NREGA, are examples for categorical targeting and self-selection, respectively. Third, geographic targeting uses the location of households' residences to assign them to geographical units that may differ in program eligibility or at least program insensitivity (Ravallion, 1993). Geographic targeting is commonly combined with both community-based (Alatas et al., 2012) and statistical (Bigman et al., 2000; Karlan and Thuysbaert, 2016) targeting.

statistical targeting methods. In addition, local participative assessments are found to consider more poverty dimensions than merely consumption (Alatas et al., 2012; Van Campenhout, 2007) and to improve local ownership and sustainability of the underlying program (Robertson et al., 2014). On the downside, it has been often argued that decentralization of political decision making is susceptible to capture by local elites (Conning and Kevane, 2002; Bardhan and Mookherjee, 2006). In contrast, PMT-based targeting is the more objective and replicable procedure, which can avoid potential principal-agent problems (Ravallion, 1993). Hence, if a central government aims to retain control over the targeting procedure, preference will likely be given to statistical methods. In a comprehensive review of targeted anti-poverty interventions throughout the global South, Coady et al. (2004b) find that statistical and decentralized targeting methods are similarly often employed. They conclude, however, that "(t)here is little documented evidence on community-based targeting as compared to other methods (Coady et al., 2004b, p.59)."

In my dissertation, I consider a program that has applied community-based targeting in communities of northwestern Burkina Faso to test empirically two main hypotheses of the aforementioned debate; (i) CBT's superior cost-effectiveness and (ii) its susceptibility to elite capture. In the first part of this dissertation (chapter 3), I conduct a comparative targeting accuracy assessment for 'my' CBT and four common methods of statistical targeting. Based on a rigorous cost-effectiveness analysis, this study provides the first empirical test of the prominent hypothesis that decentralized targeting outperforms statistical targeting when targeting accuracy is evaluated with respect to underlying program cost (Chambers, 1994b; Mayoux and Chambers, 2005; Coady et al., 2004b; Conning and Kevane, 2002).

The second part of my dissertation (chapter 4) employs self-collected information on the local targeting committees to test empirically for the presence of elite capture within the community-based targeting procedure. The anxiety of elite capture occurring in decentralized targeting is rooted in the general concern −among scholars and practitioners alike− that decentralization enables local leaders to dominate and corrupt community-planning and governance (Dasgupta and Beard, 2007; Bardhan and Mookherjee, 2006). An early argument for elite capture is that local leaders tend to dominate community decision making (Dreze and Sen, 1989), which might be facilitated both by severe power imbalances among local participants (Platteau and Gaspart, 2003) or assortative matching of elitist households into local committees (Arcand and Fafchamps, 2012; Baird et al., 2013). While statistical targeting procedures are not immune to manipulation neither (Bardhan and Mookherjee, 2005; Niehaus and Atanassova, 2013), there seems to be an especially widespread notion about the proneness of decentralized targeting to elite capture (Conning and Kevane, 2002). Nevertheless, the few existing empirical studies suggest that this perception might be unwarranted (Alatas et al., 2013a), especially within programs that encourage local participation without relying on constituted local governments (Mansuri and Rao, 2013).[6] One common manifestation of elite capture is favoritism which constitutes an in-group bias in the political leader's allocation decision (Bramoullé and Goyal, 2016).[7] My main interest lies in ethnic favoritism where coethnic households benefit from public policy decisions. At the central level, ethnic favoritism has been identified as a major driver of public policy distortions, especially in sub-Saharan African countries (Burgess et al., 2015; Dreher et al., 2016; De Luca et al., 2016). The empirical analysis in chapter 4 provides the first attempt to estimate the extent to which ethnic favoritism distorts public allocation decisions at the most decentralized level.

After having addressed two important issues that arise at the stage of poverty targeting, chapter 5 examines the effectiveness of a 50 percent price reduction in encouraging health insurance take-up among the poor. The expansion of access to formal health insurance in low-income countries is a high priority on the global, as well as on many national, poverty reduction agendas (see, for instance, the third Sustainable Development Goal of the UN's 2030 agenda). Most of the poor in low-income countries face enormous risks in falling ill and, at the same time, can only rely on limited informal means to deal with the corresponding income shocks (Gertler and Gruber, 2002). In response, Social Health Insurance schemes that are adapted to the local needs of the poor have emerged as a main cornerstone to improve universal health coverage (Bernal et al., 2016). This trend is reflected by the surge in voluntary subnational micro-health insurance programs over the last decade. Nevertheless, in spite of their improved accessibility, voluntary health insurance programs in low-income countries still face exceptionally low take-up rates, especially in sub-Saharan Africa (Carrin and James, 2005; Bocoum et al., 2017). According to the literature, several channels contribute to the lack of insurance demand by poor households in low-income countries, including liquidity-constraints, a lack of trust in formal institutions and financial illiteracy (see Eling et al., 2014 for a recent literature review). In line with these considerations, health insurers

[6] In their extensive review of about 500 studies on participatory development and decentralization, Mansuri and Rao (2013) distinguish between 'community development' and 'government decentralization' as the two main modes by which to induce local participation. They observe that many community-driven development projects explicitly incorporate collective decision making elements to circumvent existing local leaders.

[7] In addition to favoritism, elite capture is commonly measured in terms of direct personal enrichment by political leaders, reflected by the embezzlement of public resources (Ferraz and Finan, 2011; Beekman et al., 2014) or self-selection into beneficiary lists (Besley et al., 2012; Pan and Christiaensen, 2012).

have started to experiment with enrollment-encouraging innovations targeted towards the poor such as premium exemptions or discounts (see, national health insurance schemes in Colombia, Peru, China, and Georgia), information packages (see Vietnam's Health Care Fund for the Poor), nudging (see SMS reminders in the Philippines' National Health Insurance), or bundling with microcredit products (see SKS Microfinance in India). This surge in innovative program design provides a novel opportunity for careful impact evaluations to improve our understanding of health insurance demand in low-income countries. Chapter 5 provides novel quasi-experimental evidence on the impact of targeted health insurance subsidies on take-up.

In this dissertation, I follow three different empirical approaches to program evaluation. All of them represent methods that are applied in the absence of randomized field experiments. Instead, they are based on observational data and centered around a targeted health insurance subsidy intervention that was conducted in 2009 in Nouna, an administrative department in northwestern Burkina Faso. The underlying health insurance program was set-up in 2004, based on a collaboration between the *Centre de Recherche en Santé de Nouna* (CRSN) and the *Institute of Public Health* (IPH) from Heidelberg University. This partnership was also involved in extensive data collection activities and produced a rich amount of individual and household level information. In addition to a favorable data situation, in the way the community-based targeting was conducted and documented, it represents a nice social quasi-experimental design. Combining the targeting data with three independent data sources allows me to address three separate research questions. In what follows, for each chapter of my dissertation, I explain the methodological approach applied, briefly discuss favorable features of the corresponding empirical setting, and summarize the main findings.

Chapter 3 investigates which method, community-based or statistical targeting, targets consumption poor households more accurately. The empirical analysis combines beneficiary lists from the community targeting exercises with household survey data that include consumption as well as indicators commonly applied for statistical targeting. I compare community-based targeting with four frequently used statistical procedures, where my reference is a hypothetical targeting outcome based on survey consumption. In addition, I study how the targeting accuracy of the two families of methods compares across alternative specifications of statistical targeting, across rural and semi-urban sectors, and across community characteristics. Finally, I employ targeting cost data to evaluate the existence of a trade-off between CBT's lower program costs on the one hand and the econometric PMT's higher accuracy on the other hand.

All statistical targeting indices under consideration contain the commonality that they are calculated as weighted averages of potentially transformed proxy-means variables, while differing along three dimensions, the set of indicators, the way these indicators are transformed into proxy-variables, and the weights used to aggregate the proxy-variables into a single index. Accordingly, I distinguish between four types of statistical targeting. First, the *Econometric PMT* is based on a linear regression model, which typically employs a large number of indicators available in census data (Brown et al., 2016). The indicators are usually not transformed and the weights are obtained from a regression of consumption on proxy-means variables (Alatas et al., 2012; Filmer and Scott, 2012; Klasen and Lange, 2014). Second, the weights can be obtained from the joint distribution of the indicators themselves through Principal-Components Analysis, or PCA for short (Filmer and Pritchett, 2001). Such a PCA-based index is most frequently used to proxy a household's socio-economic status in the absence of consumption data and in economics, as well as demography, it is often referred to as *Asset index*. Third, I consider two scorecards which rely on a limited set of transformed indicators, India's *Below the Poverty Line* scorecard (Sundaram, 2003) and the *Poverty Scorecard Index*, a targeting tool popular among practitioners (Schreiner, 2015). Finally, I calculate a *Multidimensional Poverty Index* following Alkire and Santos (2010). In this approach all indicators are first transformed into binary deprivation indicators and the index equals a weighted deprivation count. As I only observe community-based targeting outcomes, statistical targeting is applied in a hypothetical fashion. In particular, I ask how statistical targeting performed in case one would have used the available survey data to implement it. Hence, my first evaluation approach is a comparative targeting accuracy assessment based on ex-post simulations of five statistical targeting methods and a consumption-based benchmark.

Relative to the literature, the empirical setting provides me with two favorable features. First, community-based targeting interventions are especially common in rural settings (Chambers, 1994a) and it has been argued that they are poorly suited for wealth assessments in more urbanized environments (Coady et al., 2004a). The availability of detailed community-based targeting outcomes for both rural and semi-urban communities allows me to test this argument quantitatively. Second, the targeting accuracy literature that compares statistical with community-based targeting (Alatas et al., 2012; Sabates-Wheeler et al., 2015; Karlan and Thuysbaert, 2016; Stoeffler et al., 2016) is relatively recent and variation in study designs challenges an easy comparison across studies. By considering a 'pure' CBT that is not combined with other targeting methods, my assessment is most closely related to the seminal work by Alatas et al. (2012).

The main findings from my comparative targeting accuracy assessment in chapter 3 are fourfold. First, community-based targeting is substantially less accurate than the econometric proxy-means test but more accurate than several other common statistical targeting methods. It is approximately as accurate as the best-performing statistical methods that do not rely on consumption data for the construction of weights, the asset index and the Below the Poverty Line scorecard. Second, relative to statistical methods, CBT performs slightly better in semi-urban than in rural areas. Third, the accuracy of community-based targeting significantly worsens with community size, which is not the case for any of the statistical methods. On the other hand, other community characteristics, such as economic inequality, have identical effects on the accuracy of CBT and the statistical methods. Finally, I find statistical targeting to be more cost-effective than community-based targeting only for very large transfer amounts, exceeding 75 US Dollars per household (not purchasing-power-parity adjusted, reference year 2014), which equals roughly one hundred times the daily per capita national poverty line. Hence, for the benefits usually encountered in welfare programs in low-income countries, community-based targeting is by far the more cost-effective method in the sub-Saharan African context studied here. Moreover I find that the less expensive statistical methods, which do not require consumption data for calibration, have no cost-effectiveness advantage when community-based and econometric targeting are available to the policy maker.

In practice, the targeting exercise evaluated here was undertaken by three local representatives who received a targeting mandate from their community. My analysis in chapter 4 investigates whether these local targeting committees disproportionally favored households from the same ethnic group when making a targeting choice. Such a test for elite capture makes necessary two types of information. First, one needs to observe the ethno-religious affiliation of the representatives to distinguish in-group from out-group households. Second, in an ideal data scenario, one would like to have wealth information that perfectly predicts the committee's unbiased targeting decision to identify the distortion due to the in-group bias. To meet the first requirement, I went to the study area in December 2015 to compile additional information on the representatives that allowed me to identify them in a population census dataset, which I have merged with the targeting data. In the merged dataset I observe the following three pieces of information at the household level;

(i) wealth ranks and beneficiary status, (ii) ethnicity and ethnic affiliation to the three local representatives, and (iii) a vector of socio-economic characteristics. With 50 variables, the latter covers six different wealth dimensions, which makes me confident in capturing a satisfactory amount of the variation in observable household wealth. Taken together, my empirical specification aims to identify distortions in the decentralized allocation decision due to ethnic favoritism, conditional on observable wealth and on possible wealth differences across ethnic groups. In this sense, the second evaluation is based on an OLS regression framework to identify ethnic favoritism under fairly weak identifying assumptions.

In addition to introducing a novel elite capture dimension, namely ethnic favoritism, the present community-based targeting design comes with four favorable features to test for this phenomenon.[8] First, it directly determines the political elite in the form of three locally appointed representatives per community, who were instructed to determine the set of beneficiary households. The choice of delegating authority to a set of representatives is common in participative development programs (Mansuri and Rao, 2013) and by observing households' ethnic affiliations I can directly determine the political leaders' in-group households without relying on more noisy self-reported network measures (Comola and Fafchamps, 2017). Second, the targeting design requires each representative to reveal independently his targeting preferences during the first step, which allows me to test whether my results are driven by an unintended information-based kind of favoritism. The latter hypothetically arises when representatives have superior unobserved information on coethnic rather than on non-coethnic households. This separation problem is common to most of the study designs in the elite capture literature but, so far, has received only little attention. Third, the program benefit of consideration is a locally targeted but centrally administered small-stake discount voucher. These two features reduce the scope for capture at the implementation stage, which is an important second source of elite capture that occurs after the targeting step (Alatas et al., 2013a; Niehaus and Atanassova, 2013). If at all, I would thus expect relevant allocation distortions to take place at the targeting stage only. Fourth, the application of population data allows me to identify ethnic majority groups and to calculate relevant community characteristics such as ethnic diversity.

My test finds that local targeting committees in semi-urban communities favor households of

[8]The literature on elite capture in decentralized welfare programs has so far considered favoritism with respect to relatives (Alatas et al., 2012, 2013a; Panda, 2015; Basurto et al., 2016), friends or coreligionists (Schüring, 2014), or politically connected households (Dasgupta and Beard, 2007; Caeyers and Dercon, 2012) but not along ethnic lines.

their own ethnicity. Evidence against a potential information story is given by the finding that local leaders similarly favor coethnic households in comparison with non-coethnic households that are ethnically represented in the committee. While precisely estimated, the magnitude of the distortion due to favoritism is modest. Expressed in terms of a targeting exclusion error, it leads to a number of wrongly excluded non-represented households that is less than ten percent of all beneficiary households. Finally, ethnic favoritism solely comes from allocation decisions made in semi-urban ethnically diverse communities. This latter finding confirms a well established cross-country relationship between ethnic diversity and political economy outcomes for the local level (Mauro, 1995; Easterly and Levine, 1997; Alesina et al., 1999; La Porta et al., 1999; Treisman, 2000; Alesina et al., 2003). In the present context, no evidence for favoritism is found with respect to religion, which is only weakly correlated with household ethnicity. Finally, my results suggest that the number of three representatives in relatively diverse semi-urban communities is too small to prevent ethnic minority discrimination.

My third analysis, in chapter 5, mainly relies on the merging of targeting and health insurance data and follows a quasi-experimental approach to program evaluation. The insurer's arbitrary choice of a 20 percent targeting quota introduces a discontinuity in the subsidy assignment rule and allows me to use a 'fuzzy' regression discontinuity design (RDD) to identify causal effects of the subsidy on insurance demand. Specifically, in each community the subsidy assignment is mainly based on a comparison of three wealth rankings (aggregation rule) and 'jumps' at the 20 percent wealth threshold. Within a sufficiently close neighborhood around the threshold I can plausibly assume that households are as good as randomized absent the subsidy (Angrist and Pischke, 2009). Employing the wealth ranking data, I mimic the aggregation rule, which is a very strong −albeit not perfect− predictor for final beneficiary status. My main empirical specification then estimates the intent-to-treat effect of being selected by the aggregation rule on take-up, which is a lower bound of the local average treatment effect. By exploiting variation in average community wealth, I also examine heterogeneous effects of health insurance pricing along the household wealth distribution. The data further allow me to estimate the effect of subsidy eligibility on both dynamic and intra-household demand of health insurance. In addition, I test for adverse selection into the health insurance by comparing average insurance claims of insured households just above and below the subsidy threshold. In the presence of adverse selection, one would expect, on average, that a price increase in the premium leads to a more risky pool of insured households when compared to the low price scenario. Thus, significantly different magnitudes of insurance claims across both price groups would lead me to reject the null hypothesis that there is no selection based on risk.

The fuzzy RDD rests on the identifying assumption that my outcome variable absent the intervention must be continuous in the wealth-ranking-based forcing variable (Lee and Lemieux, 2010). This is a much weaker assumption than, for instance, with propensity score matching studies, which are widespread in the health insurance evaluation literature. The quasi-experimental setting evaluated here also compares favorably to hypothetical willingness-to-pay-studies, which can suffer from hypothetical bias and in which results strongly depend on the experimental set-up (Stewart et al., 2002). In terms of intervention design, the 50 percent subsidy nicely complements existing experimental designs that consider effects of much more extreme (25 and 100 percent) price reductions.

Overall, I find an economically large effect of pricing on health insurance take-up for households that can be regarded as poor by national standards. Specifically, semi-urban households around the beneficiary cutoff increase their demand by almost a factor of five, from initially 7 to 34 percent, when being offered a 50 percent price reduction. In contrast, the subsidy is far less effective for rural households around the eligibility threshold, which can be regarded as ultra-poor in national terms. Extrapolation of Engel Curves shows that an upward move in wealth by one quantile is associated with an increase in take-up by a factor of 1.4 and 0.9 for households around the 20 and 80 wealth percentile, respectively. Nevertheless, I do not find evidence that the pricing effect on take-up changes with wealth. At the intensive margin, the subsidy offer increases coverage of adolescents relative to adults and narrows the enrollment gap between prime aged and elderly household members in favor of the former. Finally, I find no statistical differences in average insurance claims between subsidized and regularly insured households, which I take as evidence against adverse selection.

Chapter 2

Empirical Setting

2.1 The Local Context

Burkina Faso is a landlocked country in the Sahel region of West Africa, sharing borders with Mali in the north, Niger in the east and Benin, Togo, Ghana and Ivory Coast in the south (see Figure 2.1). The country's population amounts to almost 20 million people and is currently growing at an annual rate of three percent. Half of the population is below the age of 17 years and one-third lives in urban or semi-urban areas. Burkina Faso gained independence from France in 1960 and experienced first multi-party elections in the nineties. Former president Blaise Compaore governed the country for more than 25 years and resigned at the end of 2014, following an uprising against his attempt to change the constitution for his own advantage. The subsequent government interim-phase included a failed coup and led to rescheduled national elections in November 2015 (Central Intelligence Agency, 2017). Administratively, the country is divided into 13 regions which are subdivided into 45 provinces and 351 departments. A decentralization reform in 2006 introduced nationwide local elections, which were repeated in 2012 and 2016 (Lierl, 2017).

French, the official national language, is not widespread in rural areas. Instead, the three main local languages, Mooré, Dioula, and Fufulde facilitate communication across more than 60 different ethnic groups. With a population share of 50 percent, the Mossi people represent the clear majority, followed by the Fulani, the Gurma, the Bobo, the Gurunsi and the Senufo who account for 4 to 8.5 percent, each (Central Intelligence Agency, 2017). Overall, ethnic groups cohabit peacefully within the country, reflected by a reconcilable approach in renaming it from Upper

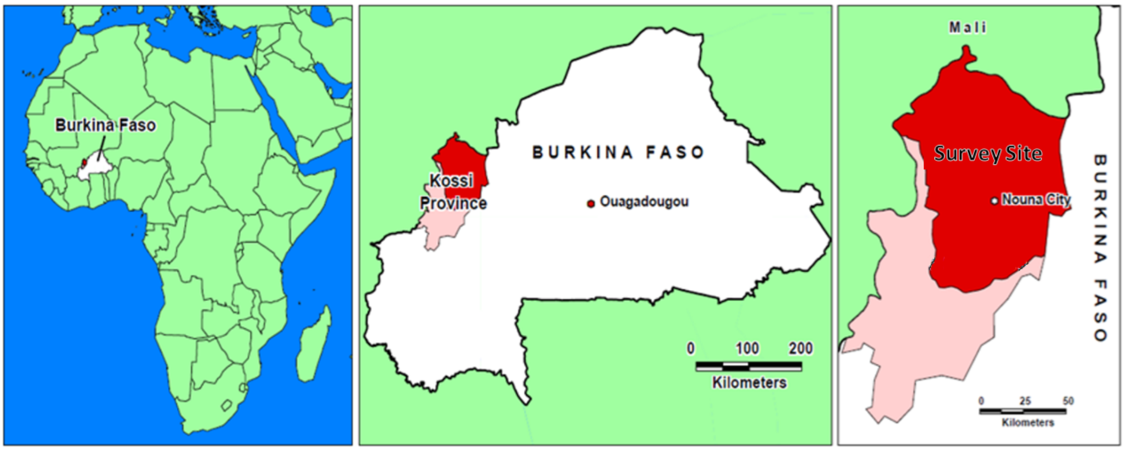

Figure 2.1: Location of the study site

Volta into Burkina Faso in 1984. Literally, the whole phrase means 'land of the upright people' and comprises components from the three main local languages (De Allegri, 2006). Muslims, catholics, animists and protestants account for about 60, 23, 7 and 7 percent, respectively.

As a low-income country, Burkina Faso has experienced a fall in its poverty rate from 46 percent in 2010 to 40 percent in 2014 (World Bank, 2017a; based on national poverty lines) and today ranks fifth from the bottom in the UNDP Human Development Index. The incidence of poverty in rural and urban areas is 48 and 14 percent, respectively. Human development indicators have improved over the last decade but infant and maternal mortality rates of 43 and 400 per 100,000 live births, respectively, are still above the average rate for sub-Saharan Africa (World Bank, 2017a). The country's weak health infrastructure is seen as an important impediment to development and is reflected by a relatively low density of primary health care centers (*Centre de Santé et de Promotion Sociale*; CSPS), the first contact points of the health system. Further supply problems in this regard are absenteeism and a strong urban-bias in the allocation of medical staff (Ministere de la Santé, 2011). In 2014, the total amount spent on health care per capita per year was $35, almost half of $60, the minimum amount recommended by the World Health Organization (WHO) (World Bank, 2017b). The Ministry of Health states that by 2009 there was only one doctor per 14,000 inhabitants and one nurse per 3,700 inhabitants (Ministere de la Santé, 2011). With no country-wide statutory health insurance in place and a negligible private health insurance market, Burkinabé, as the citizens of Burkina Faso are called, usually pay for health care at the point of service. About three quarters of private health expenditures are out-of-pocket expenditures which are mostly spent on drugs (Ministere de la Santé, 2011). By joining the WHO's Universal Health

Coverage Partnership in 2015, the government of Burkina Faso has recently committed itself to establish a social health insurance system to provide improved financial risk protection.

The study area of this dissertation is the administrative department *Nouna*, depicted as dark red area on the right-hand side map in Figure 2.1. Belonging to the *Kossi* province in northwestern Burkina Faso, it is located approximately 300 km from the capital, Ouagadougou, and 100 km from the boarder with Mali. At the time of this study, the department was inhabited by a population of about 80,000 individuals of whom two-thirds live in villages and one third in and around the town of Nouna, the only semi-urban area. The Marka or Dafing, the Bwaba, the Samo, the Mossi and the Peulh represent the five main ethnic groups in the area and Dioula is the local 'lingua franca'. The country's main four religions −Islam, Catholicism, Animism, Protestantism− are practiced in the study area with Islam being the dominant religion (De Allegri, 2006). The majority of inhabitants are farmers and primary subsistence crops include millet, corn, and sorghum (Fink et al., 2013). The Nouna department belongs to the *Boucle du Mouhoun* region, which is relatively poor in national terms. With a regional poverty rate of 60 percent, it is the largest regional contributor to national poverty incidence (INSD, 2015). Health shocks were found to be a major cause of poverty in the region (Belem et al., 2011). While informal risk sharing networks exist, illness is traditionally rather seen as an individual problem (Sommerfeld et al., 2002).[1]

The Kossi province is administratively equivalent to the *Nouna Health District*, which, at the time of our study, contained 34 primary health care centers and one district hospital within Nouna town (Fink et al., 2013).[2] The average distance to the next health care center amounts to about ten kilometers, which is above the national average of 7.2 (Robyn et al., 2012b). Since 1989, the Nouna department has been under demographic surveillance by the *Centre de Recherche en Santé de Nouna* (CRSN), one of the country's few research institutions that directly report to the Burkinabé Ministry of Health. Figure 2.2 shows that the area under surveillance covers a relatively sparsely populated area including 41 rural villages and Nouna town.

In addition to supply inadequacies, people in the study area show low demand for formal health care. A 2004 conducted household survey revealed that among those who reported a

[1] Informal insurance schemes have been also found to perform particularly poorly in insuring health shocks in other contexts. In Indonesia, consumption drops by 20 percent when a household member falls severely ill (Gertler and Gruber, 2002). In the Philippines, Fafchamps and Lund (2003) found that village-based solidarity networks perform worst in insuring non-fatal severe diseases.

[2] Health districts represent the most decentralized administrative unit of the country's health system. Each health district belongs to one of the 11 Regional Health Directorates and usually covers one district hospital and several health care centers. In the Nouna health district, the government is the only provider of formal health care (De Allegri, 2006).

recently experienced illness only 20 percent visited a health care facility center, in comparison to 50 percent who applied self-medication (Dong et al., 2008). Seasonality in both income-generating

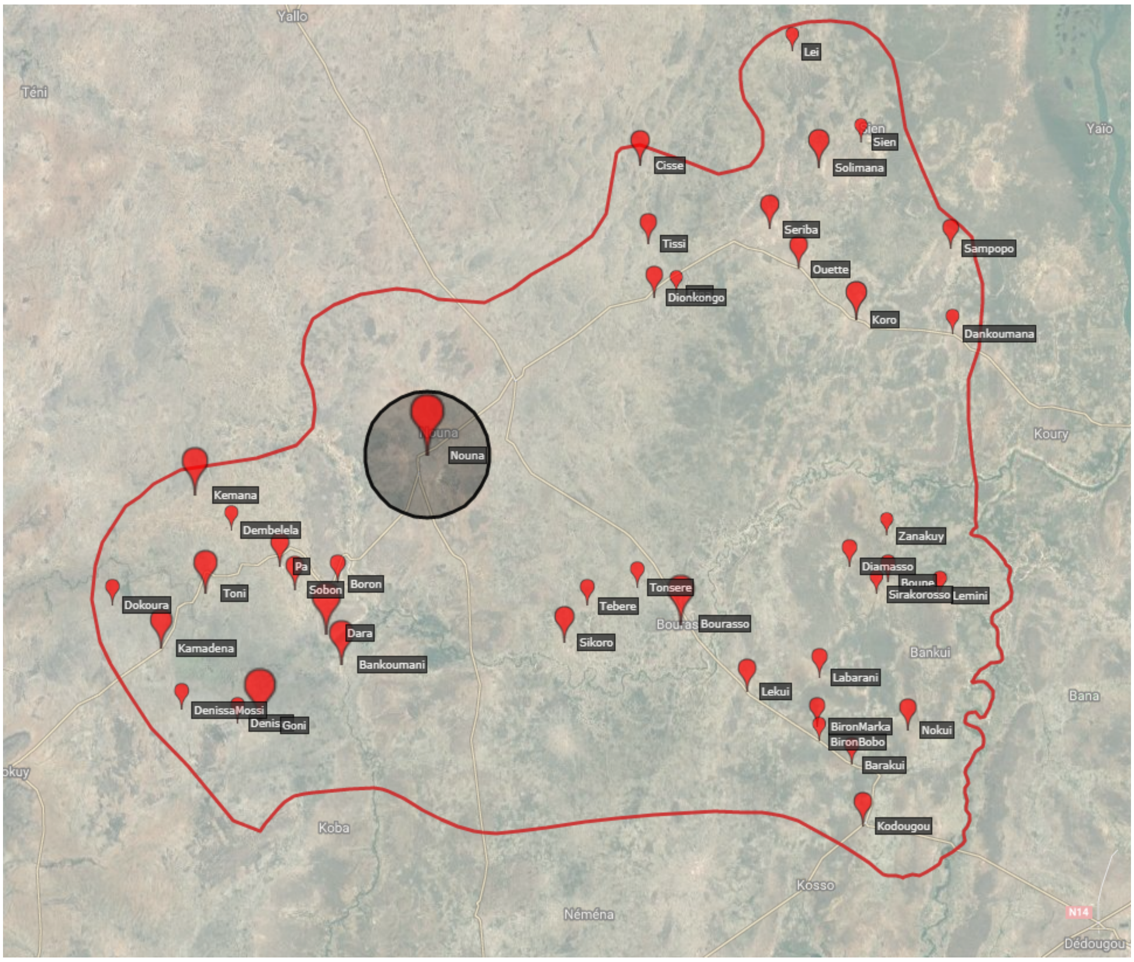

Figure 2.2: The survey site

Notes: Pins indicate village positions in the study area. Pin size corresponds with village size (except for Nouna town, the circled pin). Created by the author with GPS Visualizer.

activities and malaria transmission are seen as one explanation for the lack of demand. Malaria is the most frequently reported illness in the region and endemic during rainy season, between June and October (Fink et al., 2013). Based on administrative health care utilization data from the study area, Figure 2.3 illustrates this pattern. Between July and November the distribution of health care facility visits by insured individuals shows regular and slightly increasing peaks over time. It is also during rainy season when farmers heavily invest in agricultural inputs. Hence, this period is characterized by both high household liquidity constraints and an increased risk of falling ill (Sauerborn et al., 1996).

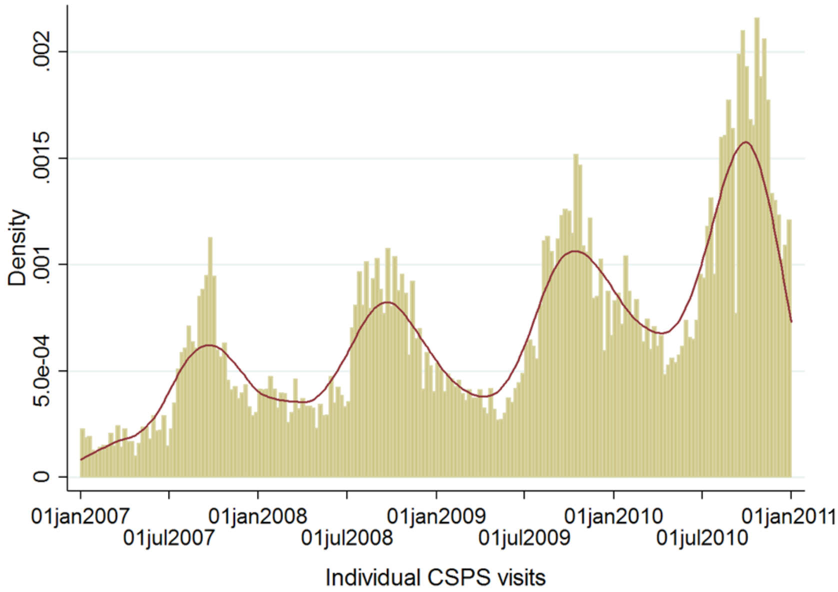

Figure 2.3: Histogram of health care facility visits over time

Source: Administrative data from the health insurer in the study area (see Section 2.3).

Notes: The graph depicts the distribution of individual health care visits over the years 2007 to 2010 for the subpopulation of insured individuals. In addition to the histogram, a kernel density estimate (epanechnikov) is added for smoothing. CSPS = Centre de Santé et de Promotion Sociale (health care facility center).

2.2 Targeted Subsidized Health Insurance in Nouna

After a decade of very limited success with ad-hoc health insurance programs in the country, in early 2000 the Burkinabé Ministry of Health was interested in promoting the establishment of a health insurance scheme at a larger scale. The idea was to first set up a pilot intervention at the department level in order to allow for a rigorous program evaluation and with the prospect of future expansion (De Allegri, 2006). Given its preexisting demographic surveillance system, the Nouna department was considered a suitable environment for such an intervention. As an additional advantage, the *Centre de Recherche en Santé de Nouna* (CRSN) had a longstanding partnership with the *Institute of Public Health* (IPH), a research institution from the University of Heidelberg, which was interested in a collaboration on the project (De Allegri, 2006). In 2004, the partnership introduced a voluntary social health insurance scheme, the *Assurance Maladie á Base Communautaire* (AMBC). It was randomly phased-in at the community-level over a three year period, such that by 2006 every household in the study area was able to buy health insurance.[3]

[3]This research collaboration already has produced a large amount of studies which are mainly published in the public health literature and deal with topics such as patterns of enrollment (Parmar et al., 2012, 2014) and

While the AMBC was restructured and transferred to a new implementing agency in 2012, my dissertation considers data from the years 2007 to 2011, a time when the initial health insurance was still in place. The 'original' AMBC was designed and implemented with the following features. Enrollment was voluntary and initially supposed to take place at the household level to limit adverse selection. Since this requirement further reduced already low insurance take-up, it was not fully enforced in the field and households could effectively buy health insurance at the individual level (Parmar et al., 2012; Fink et al., 2013). Enrolled households paid a one-time membership fee of CFA 200 and an annual flat premium of 1500 and 500 West African Franc (approximately 3 and 1 US dollar in 2009, respectively, not purchasing-parity-adjusted) per adult and child under the age of 16, respectively.[4] Payment could take place in installments and health insurance coverage started from the day of the last installment. To limit adverse selection, individuals who enrolled for the first time were subject to a three month waiting period before receiving insurance coverage. The benefit package included general and specialized consultation in one of the department's 13 health care centers or the department's hospital, covering essential and generic drugs, as well as the most important health care facility treatments (Fink et al., 2013).[5] Similar to many other social health insurances, the program could not work self-sufficiently but was subsidized by the Burkinabé Ministry of Health and international donors (Parmar et al., 2012).[6]

Despite of a seemingly affordable insurance premium, overall health insurance enrollment rates had remained far below expected levels and were especially low among poor households (Souares et al., 2010). In Figure 2.4, I have constructed household wealth percentiles by ranking all households in the study area with respect to an asset index value to illustrate its relationship with health insurance take-up in 2006.[7] The graph shows that take-up clearly increased in wealth and followed a convex form for the upper half of the wealth distribution. Figure 2.4 further includes kernel density estimates of the wealth variable by sector to illustrate that semi-urban households (the red dashed line) accounted for most observations in the upper half of the wealth distribution and

utilization (Dong et al., 2009), welfare effects (Fink et al., 2013; Schoeps et al., 2011), and health insurance product design (De Allegri et al., 2009).

[4] For comparison, the national poverty line is CFA 108,000 per person per year (IMF), which should roughly equal median consumption in the Nouna department. Thus, the adult premium can be expected to equal 1.5 to 2 percent of median annual per capita consumption expenditures.

[5] Specifically, the benefit package included comprehensive prenatal care, laboratory tests, inpatient hospital stays, X-rays, emergency surgery and transport by ambulance. It did not include treatment of teeth and eyes, addiction problems and neither HIV Aids or other chronic diseases.

[6] In 2007, the average benefit in prescriptions per insured individual was about twice as high as the average full premium paid (Parmar et al., 2012).

[7] The asset index is based on principal-components analysis (PCA) and since the asset data come from the 2009 census survey the index only provides an approximation of households' asset wealth in 2006.

vice versa for rural households (the blue solid line). After having observed especially low health insurance enrollment among poor households, in 2007 the insurer decided to offer a 50 percent premium discount to the poorest quintile of households in each community. To be precise, the insurer's program proposal to the ethical review committee of Burkina Faso states the intention to "identify the twenty percent poorest households (...) such that they could benefit from health insurance at lower prices" (Savadogo and Souares, 2006, p.2).

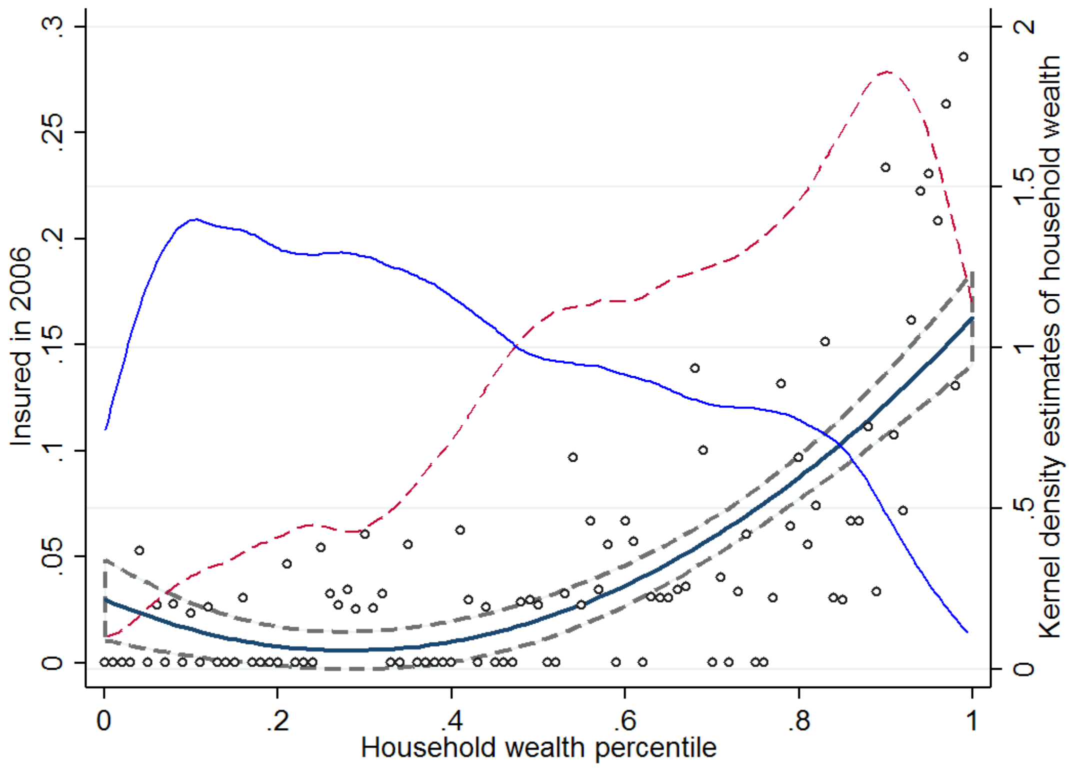

Figure 2.4: Health insurance demand and wealth before the subsidy intervention

Notes: Dots indicate households' average insurance enrollment status in 2006 by wealth percentile. Household wealth is based on an asset index that was constructed with principal-components analysis. The asset index employs asset data from the 2009 census survey and thus only provides an approximation of households' asset wealth in 2006. The solid fitted line is the predicted relationship between household wealth and insurance enrollment for a second-order polynomial regression model. The graph further includes kernel density estimates (epanechnikov) of household wealth for rural (blue solid line) and semi-urban (red dashed line) households.

For the targeting of households the insurer employed participatory wealth rankings, which were carried out at the level of villages and semi-urban neighborhoods during the first quarter of 2007, 2009 and 2011.[8] In each community, the procedure started with a publicly convened community meeting where the facilitators informed about the purpose of the meeting. People who have

[8] For the purpose of the targeting exercise, each of the seven administrative sectors of Nouna town was divided into up to four neighborhoods with similar numbers of households (see Figure A.1 in the Appendix for a map).

lived in the communities for long enough to have good information on their fellow households were especially encouraged to participate in the community meeting. On average, the community assembly consisted of ten participants and represented different types of local non-formal political leaders:

> They could be opinion leaders like the chief of the village, religious leaders, elders, etc. They could also be leaders of associations or groups like farming cooperative, artisan groups and associations of women (Savadogo et al., 2015, p.3).

The facilitators then initiated focus group discussions to elicit local criteria regarding poverty and wealth, which in turn were used to define three to four wealth categories (e.g. 'household has insufficient food' was mentioned as a common criterion for the lowest category). In Section 3.1, I present and discuss a systematic review of the extant literature on participatory wealth assessments and conclude that these three procedural steps − (i) summoning a community meeting and (ii) conduct focus group discussions, (iii) to elicit local wealth criteria − are commonly used in community-based targeting exercises. In contrast, I find that there is more variation across participatory methods when it comes to agency; targeting exercises are either carried out by the community as a whole, or by a small number of representatives (see Table 3.1 on page 34). The targeting exercise in my study required the community assembly to elect three local representatives by acclamation. Physically separated from the focus group and each other, each representative then had to generate a wealth ranking of all households in the community. They did so by first assigning each household to one of the previously defined wealth categories and, second, ordering all households within each category.

Finally, the facilitators applied a three-step-procedure to determine the set of beneficiary households. Given an intended community targeting share of 20 percent, the absolute number of beneficiary households in each community, b, was determined upfront and not disclosed. For each representative's wealth ranking the facilitators, in a first step, identified the corresponding target set by drawing all households with a wealth rank lower than $b + 1$. In a second step, facilitators preselected all households which appeared in at least two of the three representatives' target sets (i.e. they basically applied majority rule to the three wealth rankings). In communities where the number of these 'majority-selected' households, m, was below (33 communities) or above (8 communities) the intended number of beneficiary households, b, the facilitators instructed the local committee to fill-up remaining slots (if $b > m$) or to exclude excessive households (if $b < m$), respectively. Candidates for subsequent inclusion (exclusion) were those households which appeared in exactly one (two) of the three representatives' target sets and the choice was based on consultation between the three representatives. Overall, 15 percent of all households qualified as candidates for targeting by consultation, which on average determined 25 percent of all beneficiary households. The three Venn diagrams in Figure 2.5 illustrate the exact procedure, using a stylized example with a 20-households-strong community that targets $b = 4$ households. On average, the entire targeting exercise took half a day. As only five percent of beneficiary households actually redeemed their vouchers, the total amount of transferred benefits amounted to not more than 732 US dollar in 2009, which was covered by a German philanthropic organization. The finally distributed transfer amount clearly represents a low-stake scenario. Nevertheless, it is unlikely that representatives anticipated such a low take-up of the subsidy ex-ante, at the stage of targeting. The 50 percent subsidy offer itself represents a non-negligible transfer amount that roughly translates into one percent of median annual per capita consumption expenditures.

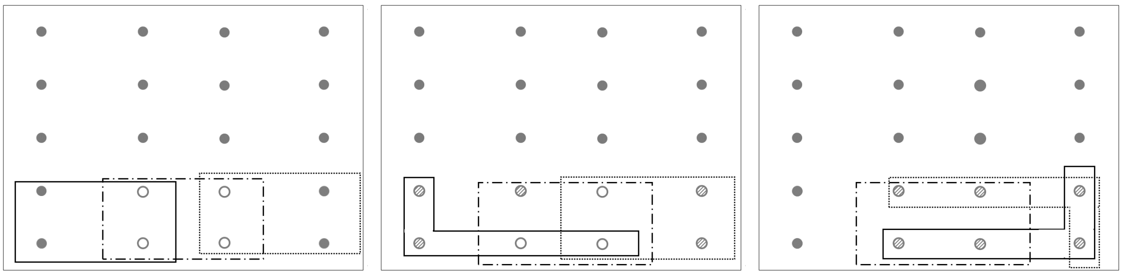

(a) $b = m$ (b) $b > m$ (c) $b < m$

Figure 2.5: Three scenarios when applying majority rule among three representatives

Notes: I consider a stylized example with a community inhabiting 20 households from which $b = 4$ households are to be targeted for benefits. m is the number of households targeted by majority rule. The set of circles framed by a box denotes the preliminary target set of one representative. Filled, shaded, and hollow circles refer to non-eligible, potentially eligible, and automatically eligible households, respectively.
Figure 2.5a presents the simplest case, where $b = m$, such that all beneficiary households are determined by majority rule. Figure 2.5b depicts a scenario with $b = 4 > 3 = m$, where one additional household has to be drawn from the set of shaded circles. In Figure 2.5c, where $b = 4 < 5 = m$, one shaded circle has to be excluded.

2.3 The Data

For the analysis, I have complemented original data from the community-based targeting exercises with additional self-collected information and merged this augmented dataset with three separate micro-level datasets. Figure 2.6 illustrates the merging procedure for the three final datasets. As the main implementing institution in the field, the CRSN was in charge of both surveying the area and running the health insurance. It used an initial population census in 2001 as sampling frame for subsequent data collection activities and followed the same assignment of unique identifiers across datasets. All my datasets are centered around the year 2009, when the CRSN conducted its second community-based targeting intervention. They differ with respect to coverage, the lowest observational unit, and data collection method and, in the following, I shall give a brief description of each dataset.

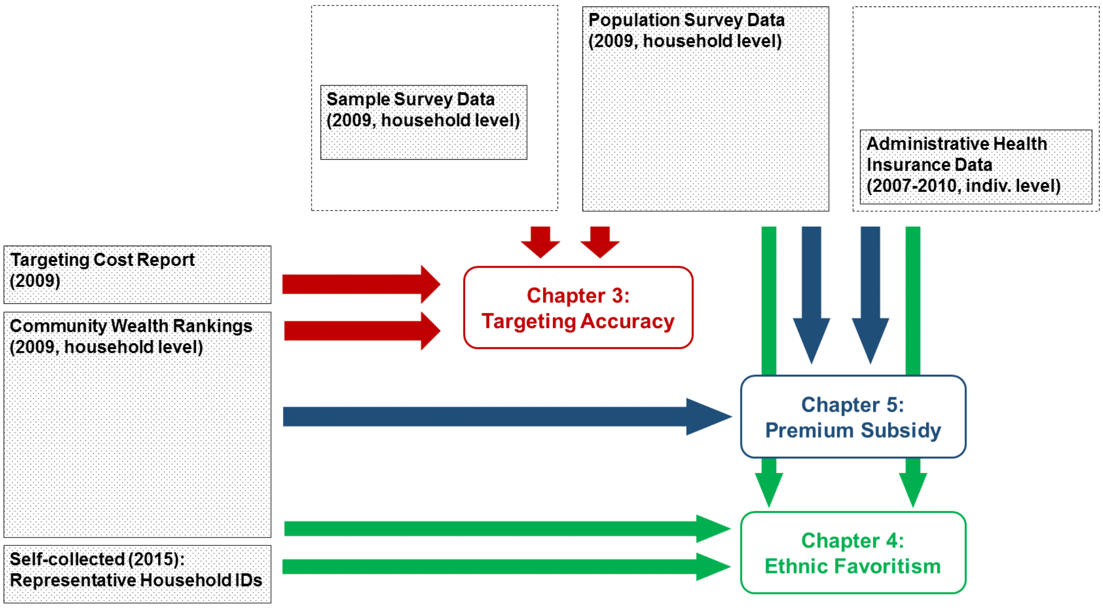

Figure 2.6: Merging of data sources to construct three final datasets

Notes: The flow chart illustrates the merging of six data sources into three final datasets that were used for the three analyses in my dissertation. The six independent data sources are reflected by rectangles and their shading roughly indicates population share covered.

My main dataset draws on two procedural outcomes of the community-based targeting (CBT) exercises. For the ranking task, each representative was given a pile of cards where each card included the name and survey ID of one resident household (see Figure A.2 in the Appendix). After a representative has produced a wealth ranking of all household-cards, the facilitators manually tagged each card with its respective wealth category and rank. From each card CRSN employees entered both the household's survey ID and wealth rank to construct the community-wealth ranking dataset. By the end of the targeting exercise facilitators wrote an official list containing the names and survey IDs of all finally eligible households in the community. CRSN employees used these lists to add each household's final beneficiary status to the wealth ranking data. In addition, two field visits in June and December 2015 allowed me to considerably enhance the CBT dataset in three ways. First, and most importantly, I reviewed and restructured the community wealth ranking cards from the semi-urban neighborhoods to include them for analysis. Second, I extracted additional information from the original beneficiary lists which allowed me to identify the survey IDs of the three targeting committee members in each community.[9] Third, after personal consultations with the CRSN staff I was provided with a cost report of the 2009 targeting intervention. To summarize, the final CBT dataset comprises three final wealth ranks and the final beneficiary status for each household in the survey area. In addition, it identifies the three representatives in each community and provides information on aggregated program cost.

Since all finally merged datasets of my analyses include the community-based targeting data (see Figure 2.6), I report summary statistics of this dataset in Table 2.1 and relegate descriptive results of the finally merged datasets to the corresponding chapters. Overall, the CBT data cover 36 rural and 22 semi-urban communities. The average community size is about 110 households and 20 percent of households were targeted in each community. On average, there is a sizable positive correlation of about 0.65 between the three informants' rankings as measured by the Spearman rank correlation coefficient, and this applies to rural and semi-urban communities alike. Nonetheless, if one defines an individual key informant's target group by his m lowest-ranked households, an unanimous agreement occurs for only 40 percent of beneficiary households. Table 2.1 also includes three measures of intra-community heterogeneity. Regarding economic inequality, which is captured by a Gini index for a weighted average of household assets, there is remarkably little heterogeneity across communities; the average standard deviation of four percentage points is just one tenth of the average level of 42 percentage points. The two indices of ethno-linguistic fractionalization, for ethnicity and religion, show that there is a good deal of diversity both within and across communities. In particular, ethnic diversity is very high in the semi-urban neighborhoods.

My second dataset is the *Nouna Household Survey*, a representative dataset that contains a wide range of socio-economic household information and covers approximately ten percent of the survey site's population. It involves a sampling methodology where households are randomly drawn from clusters of similar size (De Allegri et al., 2008). Data collection took place between September and November 2009.

My third dataset is based on the *Nouna Health and Demographic Surveillance System (HDSS)*,

[9]Specifically, each representative had to write his full name on the original beneficiary list. By employing the complete information of the population census data, the CRSN was able to match the representative names with the corresponding household identifiers and provided me with the latter.

Table 2.1: Community-based targeting and community characteristics

	Rural	Semi-urban
Community-based targeting		
Ranked households per community	114	93
	(87)	(39)
Targeted households per community	23	18
	(18)	(7)
Targeted households per community (share)	0.20	0.20
	(0.01)	(0.01)
Targeted by all 3 informants (share)	0.08	0.08
	(0.04)	(0.02)
Targeted by exactly 2 informants (share)	0.10	0.12
	(0.30)	(0.32)
Targeted by exactly 1 informant (share)	0.21	0.20
	(0.41)	(0.40)
Rank correlation between 3 informants	0.65	0.66
	(0.15)	(0.09)
Community characteristics		
Gini (Wealth)	0.42	0.43
	(0.05)	(0.03)
ELF (Ethnicity)	0.33	0.67
	(0.25)	(0.18)
ELF (Religion)	0.39	0.32
	(0.23)	(0.14)
Number of communities	36	22
Number of households	3655	2053

Notes: Standard deviations in parentheses. All sample means are calculated at the community level. A community is a village or urban sub-sector. *ELF* is the ethno-linguistic fractionalization index and measures the probability that two randomly drawn individuals belong to different ethnic or religious groups, respectively. *Gini* is the Gini index for an asset index obtained trough principal-components analysis from a total of 25 assets.

which updates vital events of all individuals living in the department on a three to four months basis (Schoeps et al., 2011).[10] For each individual 'entering' the survey system, it collects information about demographics, occupational choice and ethno-linguistic affiliation. It was established by an initial census in 1992 and complemented with a census survey component in 2009. The latter covers economic household information including dwelling conditions, agricultural output, and livestock as well as asset possession. For 2009 this dataset thus provides a wide range of demographic and socio-economic information, covering the whole population of households in the

[10] In the literature, demographic surveillance systems are also often regarded as vital registration systems. The main vital events covered by the HDSS in this study are events of birth, death and migration.

study area.

My fourth dataset is based on administrative records from the health insurance provider. It contains health insurance enrollment status, as well as information on utilization patterns of individuals who were insured at least once between 2007 and 2010. The first source of information is an annually updated enrollment register. Since the CRSN was in charge of both surveying the area and running the health insurance, each client from the enrollment register can be also identified in the surveillance system. This is reflected by the the insurance cards disbursed, which contain personal information, the insurance ID and survey identifiers (see Figure A.3 in the Appendix). Upon entering a health care facility, patients had to present their insurance card to receive free of charge treatment. To claim reimbursement from the health insurer, health care facility and hospital staff kept record of any such visit. Specifically, they wrote down the customer's insurance ID together with the date of sale as well as the name, amount and price of drugs obtained (see Figure A.4 in the Appendix). Seeing a doctor was a necessary condition for receiving drug prescriptions and consultation involved a capitation fee, which was also covered by the health insurance. For my dissertation, I have available data on customers' health care utilization within one of the 13 health care facility centers but not the district hospital.

Chapter 3

Community-based versus Statistical Targeting*

** Joint work with Stefan Klonner*

The first part of my dissertation investigates which method, community-based or statistical targeting, targets consumption-poor households more accurately. I compare community-based targeting with five frequently used statistical procedures, where my reference is a hypothetical targeting outcome based on survey consumption. Among other extensions, I employ targeting cost data to compare the cost-effectiveness of each procedure. Within the vast economic literature on targeting of welfare programs, my study contributes to the topic of targeting accuracy.[1] Until recently, the literature on this topic has followed a somewhat narrow approach, where the focus has been on one specific targeted anti-poverty program at a time and targeting accuracy is measured by the share of households meeting the program's targeting criteria in all beneficiary households (Ravallion, 2009). These studies mainly consider statistical targeting methods and usually compare a household's self-reported eligibility with hypothetical eligibility calculated from socio-economic household characteristics according to the program's eligibility rules.[2]

[1] Other prominent themes are leakage (Alatas et al., 2013b; Niehaus and Atanassova, 2013), elite capture (Alatas et al., 2012, 2013a; Panda, 2015), agency problems in decentralization (Galasso and Ravallion, 2005; Banerjee et al., 2014), and communities' poverty perceptions (Van Campenhout, 2007; Kebede, 2009; Alatas et al., 2012).

[2] Some prominent examples are Banerjee et al. (2007) for food distribution, housing and employment schemes in India, Skoufias et al. (2001) for Progresa in Mexico, Ahmed and Bouis (2002) for food subsidies in Egypt, Handa et al. (2012) for cash transfer programs in Malawi and Kenya, and Castañeda (2005) for Columbia's SISBEN. I describe this latter program in more detail in section 3.1. The review by Coady et al. (2004a) summarizes studies

A more recent set of studies has taken a broader approach to the topic of targeting accuracy by comparing alternative targeting methods in one empirical setting. This small but rapidly growing literature employs consumption as the reference and the variation in targeting methods comes either from alternative treatments (Alatas et al., 2012; Sabates-Wheeler et al., 2015) or hypothetical calculations with household-level survey data (Grosh and Baker, 1995; Filmer and Scott, 2012; Klasen and Lange, 2014; Karlan and Thuysbaert, 2016; Brown et al., 2016; Stoeffler et al., 2016). Within this recent comparative targeting accuracy literature one may distinguish between a first branch that compares various alternative statistical targeting methods with each other (Grosh and Baker, 1995; Filmer and Scott, 2012; Klasen and Lange, 2014; Brown et al., 2016), and a second branch, whose subject is the comparative assessment of community-based targeting. Among the latter, the seminal study is Alatas et al. (2012). These authors separately investigate the targeting accuracy of pure community-based targeting, as in the present analysis, and a hybrid method, where community-based targeting is combined with econometric targeting to identify the set of beneficiaries. While Karlan and Thuysbaert (2016) and Stoeffler et al. (2016) compare hybrid methods with selected statistical methods, this study is the first comparison of statistical targeting with pure community-based targeting after Alatas et al. (2012).[3]

The main contribution of my analysis is to merge the two so far disconnected branches of this recent comparative literature on targeting accuracy. I am first to compare the accuracy of pure community-based targeting with the four most prominent approaches to statistical targeting, including a comprehensive asset index, in one empirical setting. The second major contribution of this study is a detailed cost-benefit analysis covering various statistical methods as well as community-based targeting, an important topic, on which empirical evidence has been especially thin.[4] I use comprehensive targeting cost data, consider alternative cost scenarios, and also make a methodological contribution by quantifying the trade-off between costs of targeting and its benefits, in terms of poverty reduction. Third, by assessing the targeting accuracy in rural and semi-urban communities separately I am first to address the issue of targeting accuracy outside a village context. Finally, mine is the first comparative study of statistical versus pure community-

of 122 anti-poverty programs in 48 countries.

[3]Another related article is Sabates-Wheeler et al. (2015), who compare the targeting accuracy of pure community-based targeting with two forms of categorical targeting, where the target groups are households with high fractions of elderly and dependents, respectively. The empirical context is the Hunger Safety Net Programme in Kenya.

[4] In their recent book Del Ninno and Mills (2015, p.12) point out that "Trade-offs between the administrative costs of targeting and lower program costs are not well documented; further research is needed in this area." Karlan and Thuysbaert (2016) compare the costs, but not the cost-effectiveness, of a hybrid targeting method to the costs of two statistical targeting methods.

based targeting in a sub-Saharan African context, where community-based methods have been employed more frequently than anywhere else (Garcia et al., 2012; Handa et al., 2012).

3.1 Statistical versus Community-Based Poverty Targeting

Statistical targeting is a relatively recent but increasingly popular targeting tool in low-income countries (Coady et al., 2004a). In Latin America, statistical targeting has been used for large-scale cash-transfer programs in Mexico (Progresa/Oportunidades), Colombia (Familias en Acción), and Chile (PASIS and SUF). National food-subsidization programs such as those in Indonesia and Egypt use statistical targeting as well (Coady et al., 2004a; Ahmed and Bouis, 2002). Statistical methods are also popular among small-scale poverty reduction programs, where often only a small set of indicators is used. In practice, statistical targeting is often combined with a first-stage geographic targeting procedure (Coady et al., 2004a).[5]

Statistical targeting typically relies on self-reported, and sometimes validated, information on a household's demographic, occupational, and asset structure to calculate for each household in a population a wealth index, the approximate 'means' of a household.[6] A household is targeted if its index value falls short of a pre-specified cutoff, which may be defined in absolute terms or as a population quantile. The wealth index is calculated as a weighted average of the potentially transformed proxy-means variables and, in general, involves three choices; first, the set of indicators: given the high cost of data collection for entire populations, often indicators available from existing census data are used (Ravallion, 2009); second, the transformation of each indicator into a proxy-means variable, and third the index weights. I discuss four statistical methods which I consider along these lines and contrast them with community-based targeting.

Econometric proxy-means testing

This method typically uses a large set of proxy-means variables. The indicators are often obtained from census data and may or may not be transformed (Filmer and Pritchett, 2001; Klasen and

[5] Two examples are the Mexican Progresa program (Skoufias et al., 2001) and the national cash transfer program in Indonesia (Alatas et al., 2012).

[6] Such information is usually preferred over self-reported income or expenditures for several reasons. First, collecting detailed income or consumption data for an entire population is very costly. Second, both measures leave more room for strategic misreporting and can be hardly verified by the enumerator. Finally, income often suffers from considerable short-term fluctuations (Alatas et al., 2012).

Lange, 2014; Alatas et al., 2012). Weights are obtained from a regression of per capita consumption on the set of proxy-means variables. More precisely, regression coefficients are used as weights for the entire population. Hence, for a given household, its wealth index equals its predicted value of consumption in a linear regression sense. The data used for this exercise typically comes from a sample survey (Filmer and Scott, 2012; Klasen and Lange, 2014). This approach, hence, requires consumption data for at least a subset of households. When a program's purpose is to reduce consumption poverty, this approach is easily motivated by the fact that the resulting index is the best linear predictor of household consumption given the information available in a population census. Most of the recent comparative targeting accuracy literature involves this statistical targeting method, and the large-scale cash transfer programs in Mexico (Progresa) and Indonesia (BLT) are two prominent applications.

Asset index

For the asset index weights are obtained from the joint distribution of the proxy-means variables themselves. Specifically, principal-components analysis (PCA) is used to reduce a large set of proxy-means indicators to a small set of orthogonal linear combinations of the variables that best capture the variation in the original indicators. Following Filmer and Pritchett (2001), the first principal component is used as wealth index and its so-called factor loadings as weights.

The PCA-based index is most frequently used to proxy a household's socio-economic status in the absence of consumption data and has been particularly popular in health-related studies that rely on data from Demographic and Health Surveys (DHS). Within this discipline the index is often called 'wealth index' (Howe et al., 2009) or 'index of socio-economic position' (Wagstaff and Watanabe, 2003). No study in the recent targeting accuracy literature considers a comprehensive asset index.[7] While in general less popular in the area of targeting, I am aware of one prominent application, Columbia's *Sistema de Selección de Beneficiarios para Programas Sociales* (SISBEN), which has been in effect for more than twenty years. In this system, eligibility for various social programs relies on a wealth index with 13 proxy-means variables and PCA-based weights (Castañeda, 2005).

[7] Karlan and Thuysbaert (2016) consider a PCA-based index where five housing variables are aggregated into a housing index.

Scorecards

In comparison to the just discussed two statistical methods, scorecards typically rely on a smaller set of indicators. By means of a scorecard each indicator realization is transformed into an indicator score, which typically takes only integer values. The sum of indicator scores gives the wealth index, here called wealth score. The mapping of indicator realizations into indicator scores simultaneously delivers the transformation of each indicator and the weighting between indicators. A property the scorecard approach has in common with all other statistical methods considered here is that the final wealth index is additively separable in the individual indicators. The indicator scores are usually obtained using either regression techniques or common sense. I include the following two scorecard-based indices into my targeting accuracy analysis.

First, I consider the Poverty Scorecard Index (PSI). It was initially developed by a microlender in Bosnia-Herzegovina and primarily used to measure the microfinance institution's outreach to the poor and the institution's impact on customers' welfare. It has subsequently been managed on a global scale by Grameen Foundation and, lately, the non-governmental organization Innovations for Poverty Action (IPA), where it is called Progress out of Poverty Index (PPI). The index has also been used for targeting of anti-poverty interventions and is increasingly used in contexts other than microfinance, such as health and education (Schreiner, 2015; Alkire et al., 2015). According to IPA's 2014 report, the PSI is being used by more than 200 organizations for anti-poverty programs around the global South. Among them are the Bangladesh Rural Advancement Committee (BRAC), the Grameen Bank, the Ford Foundation, and the International Finance Corporation (Innovations for Poverty Action, 2014). Customized scorecards for 46 countries are available as of 2016. The selection of indicators is based on "statistics and judgment" and, similar to the econometric PMT, indicator scores are obtained from national expenditure surveys through regression techniques with consumption poverty as the dependent variable (Schreiner, 2015, p.556). In my analysis, I use the 2011 version of the PPI scorecard for Burkina Faso, which I have retrieved from IPA's website in January 2016.

Second, I calculate an index based on the Below the Poverty Line (BPL) scorecard, which was developed by an expert group commissioned by the Indian government in 2001 (Sundaram, 2003). The BPL criterion is used in India's public food distribution system and for several poverty alleviation programs administered by the Ministry of Rural Development. The indicator selection builds on India's 2001 census questionnaire. The scorecard includes thirteen indicators and a score of between zero and four is assigned to each realization. Its methodology has been vividly debated in the policy as well as the academic community (Sundaram, 2003; Banerjee et al., 2007; Alkire and Seth, 2013).

The Multidimensional Poverty Index

In this approach all indicators are first transformed into binary deprivation indicators and the relevant poverty index is a weighted deprivation count. The so-called Global MPI (Alkire et al., 2015) comprises ten indicators from three different dimensions of well-being, education, health, and standard of living. Its weights are equal across and within the three dimensions of well-being, such that the sum of all indicator weights within a dimension always equals one third. The MPI and scorecards have in common that they involve normative judgments regarding the selection and the transformation of indicators, as well as the choice of weights.

The Global MPI has been developed for the United Nations Development Programme. It is annually reported in the Human Development Report and is calculated for more than one hundred countries (Alkire and Santos, 2014). Its primary purpose is the measurement of multidimensional poverty in the developing world given common data constraints. In addition, Alkire and Santos (2010) argue that the methodology underlying the global MPI may also serve as a tool for the targeting of anti-poverty programs (see also Alkire et al., 2015). Along these lines, four recent studies compare targeting based on the MPI methodology with other targeting approaches, such as the BPL scorecard (Thomas et al., 2009; Alkire and Seth, 2013; Azevedo and Robles, 2013; Robano and Smith, 2013). I am aware that the MPI intends to capture a different, more multidimensional concept of poverty than the consumption benchmark. Due to its popularity in policy applications and its ambitions for targeting, I find it nonetheless of interest to include it in my comparative analysis to see how suited (or not) it is for proxying for consumption poverty.

Community-based targeting

In community-based targeting the choice of beneficiary households is delegated to local communities (Ravallion, 1993). The approach usually includes a so-called community wealth ranking and has earlier often been called *Rapid Rural Appraisal*, or RRA for short. According to Chambers (1994a), RRAs were pioneered in the late 1970's because of a growing discontent with statistical poverty assessments and, in particular, their relatively high costs. Since then, community wealth rankings have not only been used for poverty assessments (see, for instance, Devereux and Sharp, 2006; McGee, 2004; Van Campenhout, 2007) but have also emerged as a targeting tool.

Recent examples include small to medium-scale asset creation programs geared at the ultra-poor and funded by the Consultative Group to Assist the Poor (CGAP). Karlan and Thuysbaert (2016) analyze one such program in Honduras and Peru. Banerjee et al. (2007) investigate CBTs within the context of a similar asset-creation program in rural India. Community-based targeting is also sometimes used on a larger scale. In their cross-country analysis of targeted anti-poverty interventions, Coady et al. (2004a) find that, overall, community-based targeting is similarly often used as proxy-means testing, equally popular on all continents and especially wide-spread in very poor countries.

To the best of my knowledge, there is no structured summary of the procedural details of community-based targeting in the extant literature. Therefore, in Table 3.1, I review eighteen studies of CBTs, inclusive of the intervention preceding the one studied in this dissertation (Souares et al., 2010), which are sufficiently explicit regarding procedural details. Eleven exercises have been implemented in sub-Saharan Africa. All eighteen instances have in common that the targeting exercise involves the entire community, at least at an initial stage. They differ along five characteristics, which are set out in columns 1 to 5 of Table 3.1. First, most CBT exercises start with a public focus group discussion to elicit communal wealth and poverty perceptions, and sometimes also to define wealth brackets. Second, in most of the CBTs summarized in Table 1, all households of the community are assigned to the different wealth brackets. Third, in ten of the studies, a complete wealth ranking of households is undertaken by sorting households within each wealth bracket subsequently. Fourth, the outcomes of the wealth ranking exercise are used for targeting of a welfare program in two-thirds of the cases. Finally, there is variation regarding agency. In particular, the assignment of households to wealth brackets as well as the comprehensive ranking may be carried out either by the community as a whole or by a small number of elected local informants.

Table 3.1: Community wealth rankings and community-based targeting: procedural details

	Country	Study Population (Villages/ Households)	Focus Group Discussions	Wealth Brackets	Complete Wealth Ranking	Targeting	Number of Informants
			(1)	(2)	(3)	(4)	(5)
Latin America							
Takasaki et al. (2000)	Peru	8/300	*n.r.*	*n.r.*	YES	NO	3 – 4
Karlan and Thuysbaert (2016)	Honduras	40/1,060	YES	YES	NO	YES	*n.r.*
Karlan and Thuysbaert (2016)	Peru	40/1,007	YES	YES	NO	YES	*n.r.*
Asia							
Adams et al. (1997)	Bangladesh	55/1,637	*n.r.*	YES	NO	NO	5
Banerjee et al. (2007)	India	5/213	*n.r.*	YES	NO	YES	*n.r.*
Caizhen (2010)	China	1/473	*n.r.*	YES	YES	YES	ALL
Alatas et al. (2012)	Indonesia	640/5,753	YES	NO	YES	YES	ALL
Sub-Saharan Africa							
Scoones (1995)	Zimbabwe	1/21	YES	YES	NO	NO	ALL
Shaffer (1998)	Guinea	1/8	*n.r.*	YES	YES	NO	8
Temu and Due (2000)	Tanzania	12/300	*n.r.*	YES	YES	YES	6
Hargreaves et al. (2007)	South Africa	8/9,671	YES	YES	YES	YES	*n.r.*
Van Campenhout (2007)	Tanzania	4/877	*n.r.*	YES	YES	NO	1
Kebede (2009)	various	37/1,300	*n.r.*	*n.r.*	*n.r.*	NO	*n.r.*
Souares et al. (2010)	Burkina Faso	57/910	YES	YES	YES	YES	3
Handa et al. (2012)	Malawi	7/9,840	NO	NO	YES	YES	5
Robertson et al. (2014)	Zimbabwe	30/12,000	YES	YES	YES	YES	5
Sabates-Wheeler et al. (2015)	Kenya	48/5,108	*n.r.*	*n.r.*	YES	YES	*n.r.*
Stoeffler et al. (2016)	Cameroon	15/4,300	YES	YES	NO	YES	*n.r.*

Notes: n.r.: Not reported, ALL: Whole community

3.2 Empirical Approach

Sample target sets

In the main analysis I construct and compare seven different sets of target households. First, the set of households actually targeted by the communities. I denote the corresponding set of target households in community c in my sample by T_c^{CBT}. For my sample, the remaining six hypothetical target groups are constructed from the household survey data as follows. Let n_c^{CBT}denote the number of sample households targeted by the community-based method in community c. To construct T_c^m, the hypothetical target set based on statistical method m in community c, I first sort all sample households in c by the wealth index of method m and select the n_c^{CBT} households with the lowest index values.[8] The aggregate sample target set is the union set $T^m = \cup_{c=1}^{C} T_c^m$, where C denotes the number of communities. Following the recent comparative targeting accuracy literature (Alatas et al., 2012; Filmer and Scott, 2012; Karlan and Thuysbaert, 2016; Klasen and Lange, 2014; Brown et al., 2016; Stoeffler et al., 2016), I take consumption as the benchmark wealth index and denote by T^{CON} the benchmark target set. I assess the targeting accuracy of method m in terms of the overlap of T^m with T^{CON}. In the following, I make precise how I calculate the five wealth indices for statistical targeting in detail. For considerations of space, estimation outputs and calculations are relegated to the Appendix (see Table B.1 for details). Table 3.2 provides an overview of the statistical targeting methods regarding indicators and weights.

Econometric PMT

For the econometric PMT I define a dummy variable $Elig_{ci}^{CON}$ equal to one if household i in community c is targeted according to my benchmark. I then estimate the linear regression model

$$Elig_{ch}^{CON} = \alpha_c + \beta x_{ch} + u_{ch},$$

where x_{ch} is a vector of proxy variables, β is the corresponding coefficient vector, α_c are community fixed effects, and u is a stochastic error term. I conduct this regression analysis separately for rural and semi-urban households (see Table B.2 in the Appendix for the regression output). In a

[8] To illustrate, consider a village with 100 households from which 20 have been targeted by the community-based method. Suppose that, from this village, the household survey includes 10 households with 3 sample households targeted and 7 not targeted. While the population targeting set contains 20 percent of all households, all seven of my sample target sets contain precisely 3 sample households from that village.

Table 3.2: Targeting methods

Set	Description	Number of indicators	Transformation of indicators	Weighting of indicators
Benchmark				
T^{CON}	Household head's monthly consumption expenditures (one month recall)	1	None	n.a.
Statistical targeting				
Linear regression				
$T^{\widehat{CON}}$	Econometric proxy-means test	44	None	OLS
Principal component analysis				
T^{PCA}	Asset index	44	None	PCA
Scorecards				
T^{PSI}	Poverty scorecard index (PSI)	9	Ordered categorical	hybrid
T^{BPL}	Below the poverty line (BPL) scorecard	12	Ordered categorical	equal
Counting of deprivations				
T^{MPI}	Multidimensional poverty index	9	Binary	Global MPI
Community-based targeting				
T^{CBT}	Households identified by three local informants	n.a.	n.a.	n.a.

Notes: n.a.: not applicable; OLS: ordinary least squares regression of consumption, transformed into a dummy variable, on the set of all indicators; PCA: principal component analysis; hybrid: implicit weights for the different indicators as chosen by the creators of the Burkina poverty scorecard; hybrid: implicit weights for the indicators obtained by predicting consumption from the scorecard indicators using a national household survey (Schreiner, 2015, p.556); equal: equal weight is given to each transformed indicator; Global MPI: Weighting scheme follows that of the global MPI.

second step, index values are calculated for each household as $\widehat{Elig}_{ch}^{CON} = \widehat{\beta} x_{ch}$. Third, for each community, households are sorted with respect to the index value and the n_c^{CBT} lowest ranked households are assigned to $T_c^{\widehat{CON}}$.

Asset Index

I apply principal component analysis (PCA) to obtain weights that derive from the joint distribution of the proxy-means variables themselves. My asset index employs the same set of indicators as the econometric PMT. To focus on within-community differences across households, I first subtract community means from each indicator. Following Filmer and Pritchett (2001), I take the first principal component as wealth index and the corresponding factor loadings as weights (see Appendix Table B.3 for details).

Poverty Scorecards

I first discuss the Poverty Scorecard Index (PSI). I use the 2011 Burkina Faso poverty scorecard from IPA's website[9] to calculate indicator scores and to construct the hypothetical target set T^{PSI}. Of its ten indicators I omit one which is not covered in the household survey, 'Does the household own a bed or mattress?'. As this indicator accounts for not more than three percent of the full score, I am confident that this omission will not threaten its overall performance. Table B.4 of the Appendix provides details about this scorecard.

Second, I employ the Below the Poverty Line scorecard as reproduced in Alkire and Seth (2013) to construct the respective hypothetical target set T^{BPL}. For reasons of data availability I have to omit one of the thirteen original indicators, 'food security'. This is a limitation of my study as food security represents a potentially important and the only directly consumption-related variable in the scorecard. Nevertheless, with an indicator weight of 7.5 percent I am confident that its absence does not considerably affect the BPL's overall targeting performance. In addition, I substitute the original indicators 'availability of normal wear clothing', 'type of indebtedness', 'reason for migration from household', and 'preference for assistance' with 'drinking water source', 'type of risk coping', 'emigration incidence', and 'use of transfers received', respectively (see Appendix Table B.5 for details).

The Multidimensional Poverty Index

I build on the Global Multidimensional Poverty Index (MPI) by Alkire and Santos (2010). For reasons of data availability, I make two modifications regarding the health dimension. I omit the nutrition indicator and substitute the child mortality indicator by the incidence of a recent severe health shock, which is recorded in my sample survey with a recall of one month.[10] I adjust the weights for this dimension accordingly (see Appendix Table B.6 for details).

Targeting accuracy

I take the benchmark target set T^{CON} and assess the targeting accuracy of CBT and the statistical targeting methods in terms of the mean targeting error. The latter is defined as the proportion of households, which are erroneously classified as either poor or non-poor (Alatas et al., 2012).[11]

[9] See `http://www.progressoutofpoverty.org/country/burkina-faso`.

[10] With an indicator weight of about 16 percent, the omission of nutrition in the MPI likely has a more severe effect on targeting performance than it has for the BPL scorecard. I shall be cautions in interpreting the results and comment on potential implications in subsection 3.4.

[11] The Targeting Differential (TD) (Galasso and Ravallion, 2005) and the Coady-Grosh-Hoddinott (CGH) index (Coady et al., 2004a) represent two additional popular targeting accuracy measures. When assuming constant benefit amounts per household, both measures can be expressed as monotone transformations of the mean targeting error (MTE) because, within my framework, the share of beneficiary households equals the share of "poor" households. First, the Targeting Differential is the difference between the share of the poor and the non-poor

To be precise, when there are H_c sample households in community c and n is the total number of observations in the data set, $n = \sum_{c=1}^{C} H_c$, the mean targeting error (MTE) for method m is calculated as

$$MTE_m = \frac{1}{n} \sum_{c=1}^{C} \sum_{h=1}^{H_c} \left[\begin{array}{c} \mathbb{1}\{ \text{ household } ch \text{ is in } T^{CON} \text{ and not in } T^m \}+ \\ \mathbb{1}\{ \text{ household } ch \text{ is in } T^m \text{ and not in } T^{CON} \ \} \end{array} \right], \tag{1}$$

$$m = \{\widehat{CON}, PCA, PSI, BPL, MPI, CBT\},$$

where $\mathbb{1}\{\}$ denotes the indicator function. The mean targeting error is the sum of two types of errors. An exclusion error (false negative) occurs when consumption-poor households are not targeted by the targeting method under consideration. Conversely, non-poor households which are targeted by the method under consideration contribute to an inclusion error (false positive). As a benchmark for comparison, I also calculate the mean targeting error when households are targeted at random. For the sample targeting probability of 24 percent, the probability for erroneous targeting under random targeting is $0.76 \cdot 0.24 + 0.24 \cdot 0.76 = 0.37$.

When I compare two alternative targeting procedures, A and B, the object of interest is the difference in the mean targeting error. To conduct statistical inference, I estimate the regression equation

$$Err_{chm} = \gamma + \delta \cdot \mathbb{1}\{m = B\} + u_{chm},$$

where Err_{chm} is the targeting error of observation ch with procedure m, the term in brackets in equation 1, and u is a stochastic error term. Procedure A defines the reference category and the least squares estimate $\widehat{\delta}$ equals the difference between the mean targeting errors of procedures B and A. The data set for this exercise has $2n$ observations as every household appears twice, once with procedure A and once with procedure B. I cluster standard errors at the household level because only the differences $Err_{chB} - Err_{chA}$ can safely be assumed to be statistically independent.

3.3 The Data

For the empirical analysis I match a cross-sectional household survey data set with data from the community-based targeting exercises (see Figure 2.6 one page 22). The merged dataset contains

participating in the program, $TD = 1 - \frac{MTE}{q} \in [-1, 1]$, where q is the share of targeted households. Second, the Coady-Grosh-Hoddinott index is the amount of resources transferred to the poor over the total amount transferred by the program, $CGH = \frac{2q - MTE}{2q^2} \in [0, 1]$.

566 households, for which summary statistics are set out in Table 3.3. Households are relatively large and literacy rates low. Agriculture is the predominant activity in villages and for half of the semi-urban households. Livestock possession is wide-spread, especially in the rural sector. Targeted households appear to be slightly oversampled with a targeting share of 0.24 in comparison to 0.20 in the population (see Table 2.1 on page 24).

As reference variable for the subsequent targeting accuracy analysis, I use the value of non-durable items purchased by the household head during the thirty days preceding the interview, as recorded in the household sample survey. I shall point out here that my consumption variable does not include the value of self-produced consumption items or purchases of durable (low-frequency) consumption items. The household survey makes no attempt to record the value of the household's entire consumption. From my experience in the field, however, I think that household head expenditures are a good proxy of mean per capita consumption. I also conduct a robustness check with an alternative consumption measure exploiting the fact that the survey includes high-frequency consumption expenditures of each adult member of a household. When I use the sum of these expenditures and divide by the number of household members as a proxy for per-capita consumption, all of my results remain qualitatively unchanged. Because of my impression that household heads' responses are more reliable, I only present the results with household-head consumption in the main text. The respective additional results (Table B.7) as well as an illustration that depicts the distribution of my reference variable by sector (Figure B.1) are available in the Appendix.

Table 3.3: Household survey summary statistics

	Rural	Semi-urban
Community-based targeting		
Number of ranked households	15.9	12.4
	(8.9)	(5.5)
Number of targeted households	3.8	3.0
	(2.3)	(1.7)
Share of targeted households	0.24	0.24
	(0.43)	(0.43)
Consumption		
Monthly household head expenditures (CFA)	13,709	29,667
	(22,032)	(43,798)
Demographics		
Household size	9.38	8.75
	(6.32)	(5.38)
Household head literate (incidence)	0.33	0.38
	(0.47)	(0.49)
HH head occup. non-agric. (incidence)	0.16	0.50
	(0.37)	(0.50)
Asset possession (incidences)		
Bullock	0.50	0.40
	(0.50)	(0.49)
Goat or sheep	0.84	0.60
	(0.37)	(0.49)
Motorbike	0.19	0.33
	(0.39)	(0.47)
Bicycle	0.91	0.94
	(0.29)	(0.24)
Number of communities	36	22
Number of households	354	212

Notes: Standard deviations in parentheses. Monthly household head expenditures are based on a one month recall.

3.4 Results

Targeting accuracy

Table 3.4 reports mean targeting errors as well as exclusion and inclusion errors by sector. For the rural sector MTEs are set out in the first column. Mean targeting errors range from 16 to 36 percentage points, which amounts to a reduction of the random MTE between 57 and three percent. The econometric PMT has by far the lowest MTE with less than every sixth household wrongly classified. The difference to the next-best method, the asset index, is nine percentage points, which is statistically different from zero at the one percent significance level. Econometric targeting reduces the MTE in a statistically significant fashion relative to any of the other techniques (see Appendix Table B.8). On the other end, across sectors, the Poverty Scorecard Index and the MPI deliver only marginal and statistically insignificant (at the five percent level) improvements relative to random targeting, while the asset index performs well in villages and semi-urban neighborhoods alike. Averaged across sectors, it is the best performing statistical method not involving consumption data for obtaining weights, closely followed by the BPL scorecard, whose performance is not statistically different at conventional levels. Regarding community-based targeting, there are two salient findings emerging from Table 3.4. First, averaged across the two sectors (that is across columns 1 and 4), CBT is about as accurate as the two best performing statistical methods that do not require consumption data, the asset index and the BPL scorecard. With mean targeting errors roughly twice as large as those of the econometric PMT, the average performance of these three procedures is not statistically different from each other at conventional levels. Second, in relation to the competing statistical targeting methods, CBT performs somewhat better in the semi-urban neighborhoods than in the villages. In the semi-urban areas, CBT has a slightly smaller mean targeting error than the asset index and the BPL scorecard. The two differences fail to be significantly different from zero, however, at conventional levels.

Table 3.4 also contains exclusion and inclusion errors, and the corresponding random targeting errors as reference. By construction, the number of erroneously included households always equals the number of erroneously excluded households. Accordingly, the values in columns 2 and 5 equal the mean targeting error divided by two times the sample targeting share of 0.24. The mean inclusion errors are a multiple of the respective exclusion errors, where the factor of proportionality

Table 3.4: Targeting errors (in percent)

	Rural			Semi-urban		
	(1) Mean Targeting Error	(2) Mean Exclusion Error	(3) Mean Inclusion Error	(4) Mean Targeting Error	(5) Mean Exclusion Error	(6) Mean Inclusion Error
Econometric PMT	15.3	32.1	10.0	9.4	19.6	6.2
	(1.9)	(5.1)	(1.8)	(2.0)	(5.6)	(1.9)
Asset index	24.3	51.2	15.9	25.5	52.9	16.8
	(2.3)	(5.5)	(2.2)	(3.0)	(7.1)	(3.0)
Scorecards						
Below the Poverty Line	28.2	59.5	18.5	23.6	49.0	15.5
	(2.4)	(5.4)	(2.4)	(2.9)	(7.1)	(2.9)
Poverty Scorecard Index	35.6	75.0	23.3	34.0	70.6	22.4
	(2.5)	(4.8)	(2.6)	(3.3)	(6.4)	(3.3)
Multidimensional Poverty Index	31.6	66.7	20.7	33.0	68.6	21.7
	(2.5)	(5.2)	(2.5)	(3.2)	(6.6)	(3.3)
Community-based targeting	27.1	57.1	17.8	22.6	47.1	14.9
	(2.4)	(5.4)	(2.3)	(2.9)	(7.1)	(2.8)
Random targeting error	36.4	76.0	24.0	36.4	76.0	24.0
Number of households	354	84	270	212	51	161

Notes: Standard errors in parentheses.

is the sample targeting share, s say, divided by one minus s. I will return to the exclusion errors in the cost-benefit analysis.

In Table 3.5 I decompose targeting errors along the consumption distribution. In particular, I calculate exclusion errors separately for extremely poor and moderately poor households, and inclusion errors for households around the distribution's median as well as for relatively affluent households. I define the expenditure classes such that the shares of extremely and moderately poor households are roughly equal and sum up to the sample targeting shares of the community-based targeting exercises, 24 percent (see Table 3.3). The other two expenditure brackets contain the complementary sets of households and are defined such that they are roughly of equal size; for example in the rural subsample the affluent and around median expenditure brackets roughly contain the first and second 38 percent wealthiest households as measured by consumption, respectively. As a consequence, the mean exclusion and inclusion errors in Table 3.4 are the arithmetic means of the respective consumption-bracket-wise errors in Table 3.5.[12] While the point estimates suggest that, relative to the BPL scorecard, community-based targeting is more accurate in identifying extremely (moderately) poor households in rural (semi-urban) areas, none of the differences between the asset index, the BPL scorecard and the CBT are statistically significant at conventional levels. When averaged across sectors, the four best-performing procedures all have in common that the extremely poor are identified more accurately than the moderately poor, as would be expected. This pattern is statistically significant for the econometric PMT, the asset index and CBT at the five percent significance level. In this regard, community-based targeting is no different from the three best-performing statistical procedures.

Table 3.5: Targeting errors by consumption expenditure quantiles

	Rural				Semi-urban			
	(1) Extremely Poor	(2) Moder. Poor	(3) Around Median	(4) Affluent	(5) Extremely Poor	(6) Moder. Poor	(7) Around Median	(8) Affluent
Dependent variable:	Exclusion Error		Inclusion Error		Exclusion Error		Inclusion Error	
Econometric PMT	9.4	46.2	12.5	7.1	9.5	26.7	5.8	6.7
	(5.2)	(7.0)	(2.8)	(2.3)	(6.6)	(8.2)	(2.5)	(2.9)
Asset index	40.6	57.7	20.8	10.3	33.3	66.7	25.6	6.7
	(8.8)	(6.9)	(3.4)	(2.7)	(10.5)	(8.8)	(4.7)	(2.9)
Scorecards								
Below the Poverty Line	62.5	57.7	20.8	15.9	28.6	63.3	23.3	6.7
	(8.7)	(6.9)	(3.4)	(3.3)	(10.1)	(8.9)	(4.6)	(2.9)
Poverty Scorecard Index	87.5	67.3	24.3	22.2	71.4	70.0	30.2	13.3
	(5.9)	(6.6)	(3.6)	(3.7)	(10.1)	(8.5)	(5.0)	(4.0)
Multidimensional Poverty Index	65.6	67.3	22.9	18.3	61.9	73.3	27.9	14.7
	(8.5)	(6.6)	(3.5)	(3.5)	(10.9)	(8.2)	(4.9)	(4.1)
Community-based targeting	50.0	61.5	22.9	11.9	38.1	53.3	22.1	6.7
	(9.0)	(6.8)	(3.5)	(2.9)	(10.9)	(9.3)	(4.5)	(2.9)
Random targeting error	76.0	76.0	24.0	24.0	76.0	76.0	24.0	24.0
Number of households	32	52	144	126	21	30	86	75

Notes: All figures are in percent. Standard errors are in parentheses. The expenditure classes are defined such that for each community the shares of 'extremely poor' and 'moderately poor' households are roughly equal and sum up to the sample targeting share of the community-based targeting exercise. Analogously, for each community the sample shares of 'around median' and 'affluent' households sum up to one minus the community's sample targeting share of the CBT.

[12] The numbers of households in columns 1 and 2, 3 and 4, 5 and 6, and 7 and 8, respectively, are not exactly equal because of communities where the CBT sample target set or its complement contains an odd number of households. In that case, I have chosen to allocate the median household of the consumption sample target set to the moderately poor group and the median household of the complementary group to the around-median group. Changing this rule does not affect the results substantively.

Weights and indicators in statistical targeting

As mentioned earlier (see Table 3.2), statistical targeting methods vary along three dimensions, the set of indicators, the transformation of indicators into what may be called proxy-means variables or indicator scores, and the weighting. In section 3.4 I have found that the five different statistical indices I consider result in very different targeting errors. In this section, I explore to what extent these differences in targeting accuracy can be attributed to the sets of indicators on the one hand and the weighting schemes on the other. Accordingly, Table 3.6 reports MTEs for various hypothetical specifications of statistical indices by sector (see Appendix Tables B.8 and B.9 for statistical differences). The shaded cells contain MTEs of the five statistical methods considered in the previous analyses. Ten more indices are constructed by varying the five existing statistical indices along two dimensions, the set of indicators (in columns), and the weighting method (in rows).

Table 3.6: Mean targeting errors for alternative sets of indicators and weights

	Set of proxy-means variables			
	(1) MPI(#9)	(2) PSI(#9)	(3) BPL(#12)	(4) All(#43)
Weights				
Rural (N=354)				
Original	31.6	35.6	28.2	.
PCA	37.9	24.9	27.7	24.3
Econometric				
Log(expd)^{OLS}	27.7	22.0	24.3	18.1
Elig(expd)^{OLS}	26.6	21.5	24.9	15.3
Semi-urban (N=212)				
Original	33.0	34.0	23.6	.
PCA	40.6	29.2	23.6	25.5
Econometric				
Log(expd)^{OLS}	26.4	23.6	22.6	16.0
Elig(expd)^{OLS}	24.5	18.9	18.9	9.4

Notes: All figures are in percent. The shaded cells contain mean targeting errors for the statistical indices used in the main analysis. Further hypothetical indices are constructed by varying the five existing statistical indices along two dimensions, the set of indicators included (columns) and the method for obtaining the weights (rows). 'Original Weights' refers to the weights used in the previous main analyses.

For the role of variable selection it is interesting to compare the first three columns, where the total number of variables is similar. In the context evaluated here, the variables employed by the MPI work relatively poorly for targeting consumption-poor households. Note that both the MPI and the BPL potentially suffer from missing information on nutrition or food security, respectively. By observing a fair indicator performance of the BPL I am nonetheless confident that missing information does not mainly drive the poor performance of the MPI. Across sectors and weighting methods the MPI has the highest targeting error. Turning to weights, the Poverty Scorecard Index has the potential to outperform the MPI as well as the BPL-based scorecard when its weights are appropriately modified. The broad picture emerging for these three approaches involving normative choices of indicators, transformations and weights is that the PSI is strong regarding the set of variables but poor regarding the choice of implicit weights. The MPI performs poorly regarding both, while for the BPL scorecard both the set of indicators and the choice of weights are fair.

Finally, for the weights of the econometric PMT, I find that the choice of the dependent variable in the regression delivering the weights makes a non-negligible difference. When the logarithm of consumption is used (as in Alatas et al., 2012, and others) rather than the dummy variable 'Eligible by consumption', the MTE increases from 9 to 16 percent ($\alpha = 0.05$) for the semi-urban sector. This suggests that there are important non-linearities in the true regression function that are better dealt with by my dichotomous specification of the consumption variable.

Community characteristics and targeting accuracy

Drawing on population survey data (see Figure 2.6), I construct four community-level characteristics, community size, economic inequality, and indices of ethnic and religious heterogeneity. My measure of economic inequality is a Gini coefficient calculated for a PCA-based asset index involving 25 commonly used census variables. For measures of heterogeneity, I calculate two indices commonly known as ethno-linguistic fractionalization, which equals the probability that two randomly drawn individuals from the same community belong to different ethnic or religious groups, respectively. Table 2.1 contains summary statistics of these variables and Table 3.7 the estimation results. Since there are only 37 and 22 observations for rural and semi-urban communities, respectively, I pool the two sectors for this exercise. For ease of interpretation, all four explanatory variables of interest have been standardized.

Table 3.7: Mean targeting errors and community characteristics

	Dependent variable: Mean targeting error (in percent)					
	(1)	(2)	(3)	(4)	(5)	(6)
Targeting method:	Econo-metric	Asset index	Scorecards BPL	PSI	MPI	Commun.-based
Community Size	2.19	1.76	2.74	3.14	2.37	4.90**
	(1.76)	(2.05)	(2.14)	(3.00)	(2.23)	(2.07)
Gini (Wealth)	6.20**	7.62***	6.45**	3.12	5.05*	5.68*
	(2.77)	(2.82)	(2.83)	(3.14)	(3.02)	(3.18)
ELF (Ethnicity)	4.21**	7.92**	2.38	9.68**	1.83	5.91
	(1.94)	(3.52)	(2.93)	(4.18)	(3.20)	(3.77)
ELF (Religion)	-3.15	-7.17**	-3.93	-1.29	-2.72	-7.97***
	(1.91)	(2.79)	(2.39)	(2.87)	(2.96)	(2.95)
Rural sector (dummy variable)	10.38*	8.69	6.09	14.70	0.74	6.92
	(5.61)	(7.49)	(5.36)	(9.05)	(6.24)	(8.24)
Constant	5.42	20.03***	20.94***	24.05***	29.79***	19.91***
	(3.84)	(4.93)	(3.40)	(6.03)	(4.20)	(4.99)
Number of communities	58	58	58	58	58	58
F-test for joint significance	0.00	0.01	0.05	0.21	0.46	0.00
R^2	0.22	0.23	0.17	0.13	0.09	0.20

Notes: * p < 0.1, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses. Each community is one observation. The dependent variable is the sample mean targeting error in percent. All explanatory variables are standardized. Regarding the rural sector dummy, the reference category is semi-urban community.

I find that the targeting error of community-based targeting increases significantly with community size. Increasing community size by one standard deviation, that is by 72 households or 70 percent, increases the mean targeting error by 5 percentage points ($\alpha = 0.05$). The point estimates for the statistical targeting methods are also positive but substantially smaller, statistically different from neither zero nor the CBT estimate. For both statistical and participatory targeting I find a statistically significant negative relationship between economic inequality and targeting accuracy, and all methods respond very similarly to changes in community heterogeneity as measured by ethnic and religious fractionalization. To summarize, only community size appears to put community-based targeting at a small disadvantage relative to statistical targeting.

Cost-benefit analysis

Given the superior targeting accuracy of the econometric proxy-means test and partly also the asset index over community-based targeting, I compare costs and benefits of these three methods. Mayoux and Chambers (2005, p.283) state that "a key advantage of participatory methods is their cost-benefit in rapidly bringing together information and knowledge from many participants." In the same vein, the meta-studies of Coady et al. (2004b) and Conning and Kevane (2002) attribute the lower administration costs of CBT to the wage differential between external enumerators and community agents. When a welfare program's intention is to reduce poverty and CBT is cheaper but at the same time less accurate than statistical targeting, there is a trade-off and the relatively inexpensive CBT will be more cost-effective than statistical targeting for programs with relatively small transfer amounts, while the opposite holds for large transfer amounts. This is precisely what Alatas et al. (2012) find in their Indonesian context (Table 5, columns 1 and 2, of their Online Appendix). In the present analysis I will calculate transfer-amount thresholds for the competing targeting methods. In this context, I make two innovations. First, on the cost side, I consider alternative scenarios regarding the availability of data for statistical targeting. Second, on the benefit side, I derive explicit formulae linking exclusion errors from the estimations to poverty reduction instead of relying on numerical poverty simulations (as in Ravallion, 2009; Alatas et al., 2012; Klasen and Lange, 2015).

I use cost information from the 2009 community-based targeting intervention and implementation costs for the statistical methods based on data collection campaigns in 2010 (Lietz et al., 2015).[13] All figures are inflated to 2014 CFA (African Financial Community) francs using the consumer price index of Burkina Faso and converted to 2014 US dollars using the 2014 average market exchange rate of 526 Francs per dollar. Total implementation costs for CBT amount to \$2,373. For the two statistical methods I consider three cost scenarios. First, I assume that census and household survey information are freely available and only data processing costs of \$5,761 for the econometric PMT and \$2,665 for the asset index accrue. The difference between the two amounts reflects the extra work required to process the consumption survey data for the econometric PMT. In addition, the second scenario takes into account the data collection costs for the household consumption survey of \$41,899, which is needed to calibrate the econometric

[13] I focus on direct implementation cost and do consider neither opportunity cost of the survey respondents nor the communities' focus group participants.

PMT. Hence I calculate a total cost of \$47,660 for the econometric PMT, while the cost of the asset index remains unchanged. In the third scenario, for both statistical indices, I add the cost of collecting the census data of \$36,053, amounting to total costs of \$83,713 and \$38,718 for the econometric and asset index, respectively. To these fixed targeting costs I will add the aggregate benefits paid to beneficiary households as variable costs, to obtain the total cost of a targeted welfare program.

Turning to the benefits, I will employ the so-called distributional characteristic (Coady and Skoufias, 2004), in which the aggregate transfers received by initially consumption-poor households are the social benefit of a targeted program. Given the purpose of the community targeting exercise studied here, to identify the 20 percent poorest households in each community, I use this as the definition of poverty. This approach is equivalent to considering the change in a welfare function, where the consumption of each initially poor household enters with a weight of one and all other households with a weight of zero. Many recent papers on targeting accuracy take the change in poverty indices of the FGT family as the social benefit of a targeted welfare program and rely on numerical poverty simulations to quantify a program's effect on poverty (Ravallion, 2009; Alatas et al., 2012; Klasen and Lange, 2015). This approach, in contrast, has the advantage that there is a single metric linking the targeting accuracy of a specific targeting method to the social benefit of a welfare program, namely the exclusion error. I think that this explicit relationship instead of less transparent simulations is a substantial advantage for the understanding of the link between targeting accuracy and poverty reduction.

To fix ideas, I denote by B the average benefit per consumption-poor household in response to a targeted welfare program which transfers t dollars to each eligible household and relies on a specific targeting method. Since, in my framework, the probability that a consumption-poor household is a beneficiary equals one minus the exclusion error I have that

$$B = (1 - E)t,$$

where E denotes the exclusion error of the targeting method under consideration. The cost of such a program per eligible household will be denoted by

$$C = t + TC,$$

where TC denotes the fixed targeting costs per eligible household, i.e. the total targeting costs divided by the number of beneficiary households. Consolidating the two equations, one obtains

$$B(C; E, TC) = (1 - E)(C - TC). \tag{3.1}$$

As pointed out above, this benefit is proportional to the reduction in the poverty-gap index, where the factor of proportionality is independent of C, E and TC. Following Ravallion (2009), I am interested in which targeting method delivers the greatest benefit given a budget for the total cost per eligible household, C.

I now compare the benefits of transfer programs involving community-based targeting with econometric and asset-index based proxy means tests, respectively. The econometric PMT is always more cost-effective for programs with a large transfer benefit, and hence total costs, because, as C tends to infinity, the limit of the benefit-to-cost ratio approaches one minus the targeting method's exclusion error. For low total costs, in contrast, it is solely the fixed targeting cost TC that matters for cost-effectiveness. For all three cost scenarios regarding statistical targeting, CBT always accrues less than half the targeting cost of the two statistical methods, implying that it is the most cost-effective method for anti-poverty programs with small transfer amounts.

Table 3.8 contains transfer-amount thresholds for pairwise comparisons of the three targeting methods. When only data processing costs accrue (column 1), community-based targeting is the most cost-effective method for transfer amounts of up to 1.95 and 6.68 dollars in rural and semi-urban communities, respectively. When all data collection costs are taken into account for the two statistical methods, these figures increase to 99 and 133 dollars, respectively. It is also interesting to compare the two statistical procedures with each other. Recall that employing principal components does not require the use of consumption data. Accordingly, in cost scenarios 2 and 3, the econometric PMT is more cost-effective only for relatively large transfer amounts, of about 55 and 100 dollars in rural and semi-urban areas, respectively. Figure 3.1 contains plots of the benefits as functions of total costs for the three methods for the scenario where all data collection and processing costs are calculated for the statistical methods. Accordingly, for rural areas (upper panel), asset-index targeting is most cost-effective for only a small intermediate cost range, around one hundred dollars per eligible household, while CBT (econometric targeting) is the method of choice for programs with a small (very large) budget. In the semi-urban areas (lower panel), where CBT targets more accurately than the asset index, the picture looks qualitatively similar, except that the upper envelope is formed by CBT and the econometric PMT exclusively. For the two sectors taken together I find little scope for the relatively less expensive statistical methods, which employ no consumption data for the calibration of index weights, as the upper envelope of the cost-benefit frontier is largely formed by community-based and econometric targeting. This argument applies also to the BPL scorecard, which can be expected to be somewhat cheaper −but at the same time less accurate− than the asset index.

Table 3.8: Cost-benefit analysis

Method A delivers higher benefits than method B for program costs smaller than ...		Cost scenarios for statistical targeting		
		(1) Data processing and no data collection	(2) Data processing, consumption data collection	(3) Data processing, full data collection
A	B			
Rural				
CBT	*Econometric PMT*	5.74	62.14	110.69
CBT	*Asset index*	1.95	1.95	99.31
Asset index	*Econometric PMT*	8.11	99.77	117.80
Semi-urban				
CBT	*Econometric PMT*	6.68	74.77	133.37
CBT	*Asset index*	always	always	always
Asset index	*Econometric PMT*	5.11	56.17	74.20

Notes: All figures are in 2014 US dollars, not PPP adjusted. The maximum cost per eligible household for which method A is more cost-effective than method B is calculated by solving the equation $B(C; E^A, TC^A) = B(C; E^B, TC^B)$ for C, where B is the benefit per eligible households (see equation 3.1). In column 1 we consider only data processing costs of the census for the asset index ($1.33 per eligible household) and data processing costs of the census and the consumption survey for the econometric PMT ($2.88 per eligible household). In column 2 we consider data processing costs of the census for the asset index ($1.33 per eligible household), and data processing costs of the census and the consumption survey as well as data collection costs of the consumption survey for the econometric PMT ($23.83 per eligible household). In column 3 we consider data collection and processing costs of the census for the asset index ($19.36 per eligible household), and data collection and processing costs of the census and the consumption survey for the econometric PMT ($41.86 per eligible household). For community-based targeting we consider the CBT implementation and data processing costs of $1.19 per eligible household throughout.

To put these figures into perspective, the effective average benefit per eligible household in the present intervention, a discount on the premium of a health-insurance police valid for 24 months, amounts to $1.28, which implies a total cost of $3.35 per eligible household (benefit of $1.28 plus targeting cost of $2.07). I conclude that, among the targeting procedures considered here, CBT was indeed the most cost-effective method for targeting consumption-poor households - even though CBT's targeting cost of $2.07 per eligible (or consumption-poor) household amounts to more than three times the average transfer benefit received by a consumption-poor household (\$0.64, one minus CBT's average exclusion error of 0.5 times \$1.28). Given this seeming disproportion, would an untargeted subsidy have been more cost-effective? As there are no targeting costs for such a program and, by construction, twenty percent of households are poor in my application, the benefits received by poor households equal one fifth of the total costs of the program. This implies that a universal program is always most cost-effective for very small program budgets - because no fixed costs accrue. For a uniform transfer, a universal program is more cost effective than CBT up to cost thresholds of \$2.18 and \$1.89 in the rural and semi-urban sector, respectively. Given the total cost per beneficiary household of \$3.35, it appears that community-based targeting has indeed been the most cost-effective method in the context considered here.

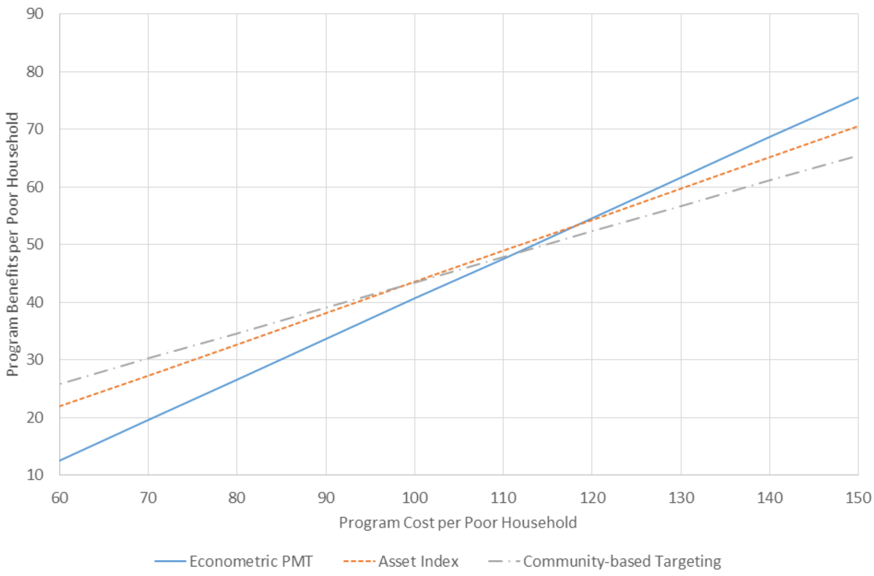

(a) Rural

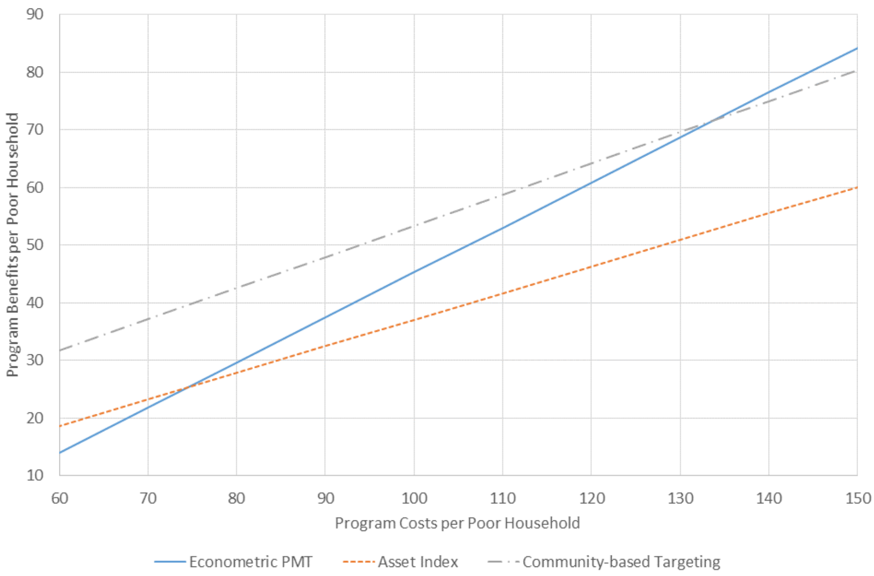

(b) Urban

Figure 3.1: Costs and benefits of statistical and community-based targeting

Notes: All figures in current local prices (in CFA) are inflated to 2014 local prices using the consumer price index of Burkina Faso and converted to US dollars using the 2014 average exchange rate. There is no adjustment fur purchasing power.

3.5 Discussion

In the following I shall summarize my findings and make explicit how they contribute to the existing literature. First, regarding the performance of various statistical targeting methods, I confirm the common and little surprising finding that the econometric PMT is by far the most accurate method. My findings are partially in accordance with Filmer and Scott (2012), who find no statistical differences when comparing the asset index with other common statistical indices that do not involve consumption data for the calibration. In this setting, the asset index performs similarly well as the BPL scorecard, but significantly better than the poverty scorecard and the multidimensional poverty index. Second, regarding CBT and statistical targeting methods, my targeting accuracy results are very similar to the results obtained in a large field experiment in Indonesia, where Alatas et al. (2012) find that the econometric PMT is about ten percentage points more accurate than the CBT ($\alpha = 0.10$).

Third, my finding of the CBT's good performance in semi-urban neighborhoods is novel. CBT initially emerged from so-called rapid rural appraisals and has so far predominantly been applied in rural settings (Chambers, 1994a). Coady et al. (2004a) expect the method to perform worse in urban sectors, where anonymity is greater and hence the information advantage of local community members smaller. My finding suggests that communities in rural and semi-urban areas follow different poverty concepts and that survey consumption is more correlated with the latter.

Fourth, the finding that decentralized targeting is less accurate in bigger communities is in line with Alatas et al. (2012). Karlan and Thuysbaert (2016) find that in comparison with two statistical procedures the CBT's cost advantage increases with community size as it incurs high fixed cost. Taken together, these two findings point towards a trade-off between higher targeting errors and lower costs when implementing community-based targeting in larger communities.

Finally, findings from the cost-benefit analysis demonstrate the trade-off between CBT's lower program costs on the one hand and the econometric PMT's higher accuracy on the other. Even if there is much anecdotal evidence for CBT's relative cost advantage over statistical targeting methods, there are very few studies including cost data (Alatas et al., 2012; Karlan and Thuysbaert, 2016). In the present context, where I consider an inexpensive decentralized expert assessment, community-based targeting is more cost-effective than any of the statistical methods. The accuracy gains of the econometric PMT outweigh the CBT's cost advantage only for very large transfer amounts and there is little scope for less expensive statistical methods, such as the asset index or scorecards. For the average transfer in my application, participatory targeting is the method of choice. But even for larger, hypothetical, transfers CBT dominates in this African context. To illustrate, first, we consider the Indonesian unconditional cash transfer program investigated by Alatas et al. (2012). Per individual in an eligible household, the program's annual cash transfer equals about 1.4 percent of Indonesia's per capita GDP in 2008. Translated to Burkina, this figure amounts to about \$70 per household and year in 2014 US dollars. Depending on the overhead cost available in addition to the funds earmarked for transfer benefits, a program of this size may or may not exceed the threshold of around \$125 below which community-based targeting is the most cost-effective method. On the other hand, if benefits are granted to eligible households for more than one year, econometric targeting will be the method of choice. A Burkinabé cash transfer program geared at improving schooling and access to health care for children in poor families, the Nahouri Cash Transfers Pilot Project, employed sophisticated econometric proxy-means testing (Akresh et al., 2014). Implemented between 2008 and 2010, it involved transfers of about \$160 per targeted household over the course of two years. Clearly, for such a program, econometric targeting would be the method of choice in the context I have studied.

Chapter 4

Ethnic Favoritism in Decentralized Targeting

In this chapter, I test for ethnic favoritism in the decentralized targeting choice of the local committees. The merged dataset allows me to observe a household's ethnic and religious affiliation to each member of the targeting committee and, at the same time, to control for a wide range of dimensions in observable wealth. My empirical specification aims to identify a bias in the decentralized targeting decision due to ethnic favoritism, conditional on observable wealth and on possible wealth differences across ethnic groups. With this analysis I directly contribute to three strands of literature.

The first strand of literature studies elite capture within decentralized targeted welfare programs and can be decomposed into two sub-strands. A first branch considers participative or community-driven allocation decisions and does not find evidence for elite capture (Dasgupta and Beard, 2007; Alatas et al., 2012; Schüring, 2014).[1] A second branch looks at targeted welfare programs at the local government level and finds mixed results. Indian village committees tend to capture Below-the-Poverty-Line benefit entitlements for themselves and close relatives (Besley et al., 2012; Panda, 2015), while more powerful political elites do not worsen pro-poor targeting in West Bengal communities (Bardhan and Mookherjee, 2006). Political leaders in sub-Saharan

[1] Alatas et al. (2012) randomly vary the committee composition within community-based targeting of Indonesia's national cash transfer program and do not find a higher targeting error for more elitist committees. Dasgupta and Beard (2007) analyzes a community-driven poverty alleviation program and do not find that benefits were allocated differently when elites were in charge. Schüring (2014) conducts controlled lab-in-the-field experiments in 25 Zambian villages and finds no evidence for favoritism in the benefit allocation decisions.

Africa favor politically connected (Pan and Christiaensen, 2012, in Tanzania; Caeyers and Dercon, 2012, in Ethiophia) and kin households (Basurto et al., 2016, in Malawi) when making an allocation choice on private in-kind benefits. For several Indonesian welfare programs, Alatas et al. (2013a) link both sub-strands when comparing elite capture across formal and non-formal leaders. Overall, they find a very moderate extent of elite capture, only exercised by formal leaders at the implementation but not at the targeting stage. With my study on ethnic favoritism I introduce a novel and important elite capture dimension to this literature. The ability to empirically distinguish between intentional and information-driven favoritism provides a methodological contribution and the analysis of a community-driven welfare program in a sub-Saharan African country considers a novel context.

The second strand of literature investigates ethnic favoritism in distributional policies and has exclusively considered central-level decision making, with a particular focus on sub-Saharan African countries. Substantial favoritism −amounting to a doubling in channeled resources− is found with respect to the president's coethnic districts in Kenya (Burgess et al., 2015, on the allocation of roads) and political leader's birthplaces in 47 sub-Saharan African countries (Dreher et al., 2016, on the allocation of Chines aid).[2] In their comprehensive assessment of ethnic favoritism, De Luca et al. (2016) cover 140 multi-ethnic countries and consider observations from 1992 to 2013. Among other things, they find that the current leader's ethnic homelands enjoy up to 10 percent more intense nighttime light and that ethnic favoritism is not only prevalent in sub-Saharan Africa, but can be considered a global phenomenon. In contrast to Burgess et al. (2015), they find institutional quality to only moderately reduce the extent of ethnic favoritism. I contribute to this literature by shifting the empirical assessment of ethnic favoritism in distributional policy decisions to the most decentralized level.

Third, my analysis contributes to a rapidly growing literature on the role of ethnic diversity in explaining socio-economic outcomes. A number of empirical cross-country studies has found that ethnic fractionalization is associated with lower growth rates (Easterly and Levine, 1997), lower public good contributions (Alesina et al., 1999; Desmet et al., 2016), and higher corruption (Mauro, 1995; La Porta et al., 1999; Treisman, 2000; Alesina et al., 2003). A more recent set of studies primarily investigates public good provision outcomes and channels for the subnational

[2] See also two studies that employ Demographic Health Survey data and find ethnic favoritism to correlate with lower health and education outcomes (Kramon and Posner, 2012, in Kenya; Franck and Rainer, 2012, for 18 African countries).

level and finds mixed evidence.[3] To the best of my knowledge, Olken (2006) provides the only subnational capture-related study and finds that ethnically diverse areas show a higher likelihood to be missing rice within Indonesia's public distribution system. By studying a well-documented decentralized targeting program for the sub-Saharan African context, I complement Olken (2006) with a more disaggregated analysis (that identifies the nature of leakage) and with evidence for a new geographical context. In addition, the choice of my outcome variable provides novel insights on the potential mechanisms of the 'ethnic diversity relationship'. Ethnic favoritism can be seen as a manifestation of discriminatory preferences, where individuals attach a higher utility to the welfare of coethnic than to non-coethnic individuals (Becker, 1957). Such preferences can explain lower public good contribution (Habyarimana et al., 2009) or higher corruption levels (Mauro, 1995) in ethnically diverse communities. Within a literature that is relatively ambiguous about the channels through which ethnic diversity might worsen political economy outcomes (Habyarimana et al., 2007) my study thus contributes with a qualitative test of the plausibility of the prominently hypothesized 'preferences mechanism'.

4.1 Empirical Approach

With my empirical analysis I aim to identify possible distortions in the decentralized allocation decision, conditional on observable wealth and on wealth differences across ethnic and religious groups. The distortion of interest arises when two equally needy households are treated differently because of their ethnic or religious affiliations to the local representatives. I start this section by presenting my main empirical specification and discussing identification concerns. Following this, I set out how I test for favoritism across community fractionalization and derive a measure to quantify the distortion in the social outcome due to favoritism.

[3] For the low-income context, the adverse relationship between public good provision and ethnic diversity is confirmed for school and water provision in rural Kenya (Miguel and Gugerty, 2005), for community-maintained infrastructure projects in Northern Pakistan (Khwaja, 2009), and a cooperation game in rural Malawi (Dionne, 2015). No relationship is found for public good provision in Sierra Leone (Glennerster et al., 2013) and Tanzania (Miguel, 2004). A positive relationship is suggested for government spending in Zambian districts (Gisselquist et al., 2016) and a public good game in urban neighborhoods of Kampala (Schuendeln, 2013). Subnational studies that consider high- or middle-income countries include Alesina and La Ferrara (2000), Alesina et al. (2004), Glaeser and Saks (2006), Alesina et al. (2014), Mavridis (2015), Swee (2015), Algan et al. (2016).

Testing for favoritism

Primarily, I want to test for favoritism in the committee's final allocation choice. The latter occurs when households which are ethnically affiliated to at least one committee member are favored over non-represented households. While this definition does not fully match the more narrow notion of ethnic favoritism, where political leaders favor coethnic households exclusively (Burgess et al., 2015; Bramoullé and Goyal, 2016), it considers an outcome that likely reflects the preference for ethnic favoritism. Given that committee members share an interest in favoring *coethnic* households, the choice to favor all *ethnically represented* households provides a plausible strategy. Specifically, favoritism of ethnically represented households can be understood as the representative's best response, leading to an equilibrium allocation that contains sufficiently many coethnic households from each representative. Similarly, elected representatives belong to the community's elite (see section 2.2) and as such they likely share a common identity (Platteau, 2004). Ethnic favoritism with respect to committee-affiliation then would reflect a natural way to act in the common interest of the elite.

In the data I observe one final allocation decision per household and consider the following regression equation:

$$\begin{aligned} \mathbb{1}\{\text{Targeted}\}_{cher} =& \alpha_e + \mu_r + \phi_c + \mathbf{X}\text{'}\gamma + \beta_1 \cdot \mathbb{1}\{\text{Ethnicity represented}\}_{cher} \\ & + \beta_2 \cdot \mathbb{1}\{\text{Religion represented}\}_{cher} + u_{cher}, \end{aligned} \tag{4.1}$$

where $\mathbb{1}\{\textit{Targeted}\}_{cher}$ equals one if household h in community c with ethnicity e, and religion r is entitled to the benefit.[4] A household h with at least one ethnically or religiously affiliated committee member is indicated by $\{\text{Ethnicity represented}\}_{cher}$ and $\{\text{Religion represented}\}_{cher}$, respectively. α_e, μ_r, ϕ_c are ethnicity, religion, and community fixed effects, respectively, and $\mathbf{X}$' is a vector of 50 socio-economic control variables. Given the empirical specification, my identifying assumption may be stated as follows: A household's 'unobservable' wealth must not be systematically different when living in a community where it is affiliated to the local committee, in comparison to a community where it shows no such affiliation. For simplicity, I estimate a linear probability model and relegate alternative regression specifications to my robustness checks

[4] In addition to ethnicity, it is of interest to also consider favoritism along religious lines. It primarily depends on the local context and, thus, is an empirical question whether favoritism takes place through the one or the other channel. In the local context considered here, ethnic and religious affiliation are not strongly correlated and I show in section 4.4 that the exclusion of one of the two dimensions does not qualitatively change the main results. For ethnically and religiously mixed households I assign the household head's affiliation.

in Section 4.4. Economically, equation 4.1 is of main interest as it refers to the final allocation decision. Econometrically, however, this empirical specification is prone to three potential threats from selection on unobservables (the following considerations mainly refer to the ethnic dimension, but they equally apply for religion).

First, I cannot rule out that the targeting committee processed wealth information that is unobservable in the data. When unobservable household characteristics correlate with ethnic affiliation, such a scenario can bias my ethnic favoritism coefficient. This is likely to occur when committee members have superior unobserved information on coethnic households. Let us assume that, for instance, ethnic affiliation allows to better observe a household's poverty dimension which is not observable to the researcher. In such a scenario, a positive $\hat{\beta}_1$ might include both an intended preferential treatment of coethnic households but also an unintended information-driven kind of favoritism. The direction of the latter is not clear; a more precise wealth assessment of coethnic in comparison to non-coethnic households could lead a representative to either overvalue or undervalue the wealth of non-coethnic households. Nevertheless, the resulting outcome would reflect a deviation from a hypothetical allocation that is based on my set of observable wealth variables. To address this separation problem between an 'intended' and an 'information-based' type of favoritism, I exploit the fact that the targeting design in the context of my study is based on the aggregation of three individual, independently determined beneficiary lists. All households were ranked by three representatives. Thus, I observe three (preliminary) allocation decisions per household, such that the number of observations equals three times the number of sample households. At the same time, the affiliation status for a given household might vary. Household h is either (i) affiliated to representative i, (ii) not affiliated to representative i but affiliated to at least one other committee member, or (iii) without any affiliation to the committee. Consequently, I can test whether representatives only favor coethnic households −probably due to better information− or whether they also favor non-coethnic represented households. While evidence for the latter does not rule out the occurrence of 'information-based' favoritism overall, it can be taken as evidence for the existence of an 'intentional' favoritism.[5] The corresponding

[5]In the local elite capture literature, surprisingly few studies address information-based identification concerns. For the set of studies that tests for capture and redirection of benefits towards relatives (Alatas et al., 2012, 2013a; Panda, 2015; Basurto et al., 2016), friends (Schüring, 2014), or politically connected households (Dasgupta and Beard, 2007; Caeyers and Dercon, 2012) I have come across two. First, Caeyers and Dercon (2012) state that an in-group information advantage should disproportionally pay-off in larger communities. They do not find a significant interaction effect with community size and take this as evidence in favor of the 'intended' favoritism channel. Second, for the local allocation of agricultural input and food subsidies, Basurto et al. (2016) test whether chiefs' allocation choice incorporates additional local information on household's returns to subsidies.

regression equation is

$$\begin{aligned}\mathbb{1}\{\text{Targeted}\}_{cheri} &= \alpha_e + \mu_r + \phi_c + \mathbf{X}\text{'}\gamma \\ &+ \beta_1 \cdot \mathbb{1}\{\text{Same ethnicity}\}_{cheri} + \eta_1 \cdot \mathbb{1}\{\text{Different ethnicity, represented}\}_{cheri} \\ &+ \beta_2 \cdot \mathbb{1}\{\text{Same religion}\}_{cheri} + \eta_2 \cdot \mathbb{1}\{\text{Different religion, represented}\}_{cheri} + u_{cheri},\end{aligned} \tag{4.2}$$

where $\mathbb{1}\{\textit{Targeted}\}_{cheri}$ indicates whether household h in community c with ethnicity e, and religion r is targeted by representative i, with $i \in \{1, 2, 3\}$. $\{\text{Same ethnicity}\}_{cheri}$ in equation 4.2 equals one if household h has the same ethnic affiliation as representative i, while $\{\text{Different ethnicity, represented}\}_{cheri}$ indicates a household h which ethnicity is represented by the committee, but not by representative i. The baseline group is a household not being ethnically represented in the committee and the same applies for religion. A statistically significant effect for both η_1 and η_2 should reduce concerns about an information story, especially, when I cannot reject $H_0 : \beta_j = \eta_j, j \in \{1, 2\}$. I cluster standard errors at the household-level.

Second, there remains a risk that my estimates are driven by the one-fourth of decisions which were not based on the aggregation of the three independently created household wealth rankings, namely the majority rule. These allocation decisions on how to fill-up remaining or exclude excess households were based on consultations within the committee (see Section 2.2). At this stage, a distinction between local-information-processing and discriminatory preferences is not possible either. Additionally, and in contrast to the majority rule, within-committee consultation allowed the representatives to share their information. The existence of information aggregation could lead to an even improved assessment of represented households' wealth and has the potential to bias my estimates further. One natural way to test whether consultation-based allocation decisions drive the effects is to estimate equation 4.1 with the subsample of households considered for committee consultation. In all communities where the share of majority-selected households was below (above) the intended targeting share, all those households which have been targeted by exactly one (two) representative(s) were considered for consultation.

Finally, it is worth noting that the community-based targeting design is based on the endogenous formation of local committees. Such an empirical setting is common to most of the local elite capture studies that consider participative targeting exercises and distributional decisions within constituted local governments. To my knowledge, only Alatas et al. (2012) and Alatas et al. (2013a) provide two novel exceptions for the Indonesian context by introducing random variation in targeting committee composition.[6] Hence, even if considering a common decentralized targeting setting, endogenous committee formation can complicate the interpretation of my results. Specifically, a coefficient estimate that indicates the presence of ethnic favoritism could be induced by the committee election procedure. One evident case is reciprocity; a community assembly that assigns the targeting mandate favorably to ethnically connected representatives (that is ethnic majority representatives) could increase the likelihood that representatives in turn 'respond' with a similarly biased targeting choice. In this case it is difficult to distinguish an 'intended' from a 'reciprocal' type of favoritism. Nevertheless, in comparison with targeting choices made by formal elites (such as in Bardhan and Mookerjee, 2006; Besley et al., 2012; Caeyers and Dercon, 2012; Pan and Christiaensen, 2012; Alatas et al., 2013a; Basurto et al., 2016), two features of the present targeting design clearly hamper the room for favoritism based on reciprocal relationships. First, public election of representatives likely reduced the incentive to suggest clearly self-serving candidates. Second, representatives were given a one-shot targeting mandate that neither involved private benefits nor farther reaching political powers. Thus, apart from reputation, the benefits from accomplishing the mandate are small and so are the reasons to act reciprocally. Nevertheless, it is still possible that elected representatives also held other political offices and used this one-shot intervention as a mean to maintain their network of reciprocal relations. That I am not able to observe such kind of committee information is a limitation of this study. It is nevertheless reasonable to consider my results as an upper bound of ethnic favoritism in comparison to a hypothetical allocation that takes place under a randomly selected local committee.

Favoritism and community characteristics

I am further interested in heterogeneous effects of ethnic and religious favoritism. In particular, I test whether political leaders capture and redirect relatively more benefits in communities that are more diverse. To measure ethnic and religious diversity, I apply the widely-used index of ethno-linguistic fractionalization (ELF), which is a decreasing transformation of the Herfindahl-Hirschman concentration index (Hirschman, 1964). Due to a limited total number of communities I construct a fractionalization indicator variable, which is one for the fifty percent least fractionalized communities and zero otherwise.[7] For both ethnicity and religion I add the corresponding

[6] For a collective community exercise, Alatas et al. (2012) produce differently composed groups by randomly varying the venue or time of the exercise, while Alatas et al. (2013a) randomly delegate targeting authority either to a group of formal or informal political local leaders.

[7] In Section 4.4 I show that results are not driven by community outliers which are located close to the median-based cutoff and that the results remain qualitatively unchanged when interacting the two indicator variables with

community fractionalization indicators and interact them with the favoritism variables. When augmenting equations 4.1 and 4.2 in this way, I obtain the following two regression equations:

$$\begin{aligned}\mathbb{1}\{\text{Targeted}\}_{cher} = \alpha_e + \mu_c + \phi_c + \mathbf{X}'\gamma \\ + \delta_1 \cdot \mathbb{1}\{\text{Ethnicity represented}\}_{cher} + \delta_2 \cdot \mathbb{1}\{\text{Religion represented}\}_{cher} \\ + \lambda_1 \cdot \{\text{Ethnicity represented}* \text{ Ethn. homogenous community}\}_{cher} \\ + \lambda_2 \cdot \{\text{Religion represented } * \text{ Rel. homogenous community}\}_{cher} + u_{cher}.\end{aligned} \tag{4.3}$$

$$\begin{aligned}\mathbb{1}\{\text{Targeted}\}_{cheri} = \alpha_e + \mu_c + \phi_c + \mathbf{X}'\gamma \\ + \delta_1 \cdot \mathbb{1}\{\text{Ethnicity represented}\}_{cher} + \delta_2 \cdot \mathbb{1}\{\text{Religion represented}\}_{cher} \\ + \lambda_1 \cdot \{\text{Ethnicity represented}* \text{ Ethn. homogenous community}\}_{cher} \\ + \lambda_2 \cdot \{\text{Religion represented}* \text{ Rel. homogenous community}\}_{cher} + u_{cheri}.\end{aligned} \tag{4.4}$$

Note that the application of community fixed effects eliminates the community fractionalization indicators and that my identifying assumption is further relaxed by the introduction of the interaction terms. The empirical specifications in equations 4.3 and 4.4 allow me to test whether ethnic or religious favoritism statistically differ across high- and low-fractionalized communities. For instance, for the effect of being ethnically represented on the probability of being targeted in equation 4.3, I consider two coefficients of interest. First, δ_1 is the effect of being an ethnically represented household who lives in an ethnically high-fractionalized community. Second, $\delta_1 + \lambda_1$ gives the same effect for households living in ethnically low-fractionalized communities. For religion, the two corresponding coefficients of interest in equation 4.3 are δ_2 and $\delta_2 + \lambda_2$, and the same logic applies for equation 4.4. In line with existing studies from the elite capture literature, I shall also examine favoritism across a community's degree of economic inequality.

Measuring the distortion to the social outcome due to favoritism

I finally examine how the favoritism-based allocation bias deteriorates the overall targeting accuracy. The latter is commonly assessed in terms of the mean targeting error, which is the proportion of households which are erroneously excluded from or included in the beneficiary list (section 3.2 on page 37). I shall focus on the exclusion error, which occurs when a household whose ethnicity is not represented in the committee fails to receive benefit entitlements because of favoritism.

the rank in ethnic or religious fractionalization.

Specifically, my objective is to calculate the proportion of not ethnically represented households which are excluded from the beneficiary set due to favoritism (reflected by $\beta_1 > 0$ in equation 4.1). When I express this figure as a fraction of the total number of beneficiary households, I have

$$\frac{1-p}{q}\Big(\Pr(Targeted = 1|R = 0, \beta = 0) - \Pr(Targeted = 1|R = 0, \beta > 0)\Big),$$

which I call the 'distortion to the set of beneficiary households' or 'exclusion error'. It can either be defined with respect to ethnic or religious favoritism. In the above expression, $Targeted$ indicates whether a household is eligible to receiving benefits, R is an indicator for being represented in the committee, p is the corresponding fraction of represented households in my application, and q is the share of beneficiary households, or the targeting quota. Notice that $(1-p)\cdot\Pr(Targeted = 1|R = 0, \beta = 0)$ is the fraction of not represented households in the population when there is no favoritism. Hence, $(1-p)$ times the difference in parentheses is the fraction of not represented households in the population excluded from the eligibility set due to favoritism. Divided by q, I have the population counterpart of the number of excluded households which are not represented, divided by the number of households in the beneficiary set.[8] Furthermore,

$$\Pr(Targeted = 1|R = 0, \beta = 0) = E_{X|R=0}[\gamma X]$$
$$\Pr(Targeted = 1|R = 0, \beta > 0) = E_{X|R=0}[\gamma X] - p\beta.$$

Hence, the distortion equals

$$\frac{(1-p)p}{q}\cdot\beta, \tag{4.5}$$

where p and q can be obtained from the data's sample means, and β from the LPM estimation. In my data the fraction in the latter expression is bounded by 1.33, as the maximum of the numerator is 0.25 (for $p = 0.5$) and the targeting share is 0.188. For both ethnicity and religion, Table C.2 in the Appendix provides the respective values for the fraction in equation 4.5 when considering all, high-fractionalized, and low-fractionalized communities.

[8] In equation 4.1, I assume a data generating process where a household being targeted ($Targeted = 1$) is a Bernoulli trial with the probability depending on observables, including representation status (R). I consider a population model, $\Pr(Targeted = 1|X, R) = \gamma X + \beta(R-p) = \gamma X + \beta\widetilde{R}$, where X is a vector of wealth characteristics observed by the researcher, including village dummies and ethnicity fixed effects, drawn from some multivariate distribution. $\widetilde{R}$ is the representation status normalized by subtracting its population mean. The normalization won't affect any of the parameters of interest that are estimated, in particular β, when X includes a constant term.

4.2 The Data

For empirical analysis, I combine the community-based targeting dataset with the population survey data and the health insurance register (see Figure 2.6 on page 22). For the year 2009 the merged dataset allows me to observe a household's ethnic or religious affiliation, whether and how it was targeted for private benefits as well as its ethnic or religious affiliation to the targeting committee. In addition, I observe a wide range of socio-economic household characteristics that are likely to explain a fair amount of the committee's 'unbiased' allocation decision. Given its universal coverage, I augment the dataset with information on how ethnic and religious groups are distributed within communities. From the initial set of 36 rural villages and 22 semi-urban neighborhoods I had to drop eight villages where I could not successfully identify all three representatives.

Figure 4.1 presents average community population shares for six ethnic and five religious groups that are present in the study area. Religion, in the lower panel, is very similarly distributed across rural and semi-urban communities. On average, Muslims and Catholics account for more than 80 and 96 percent of the rural and semi-urban community populations, respectively. In contrast, rural-urban differences exist for ethnicity, with ethnic groups being more uniformly distributed in semi-urban than in rural communities. On average, the Dafin and the Bwaba account for more than 75 percent of the rural but for less than 50 percent of the semi-urban community population. The Samo are very widespread in Nouna town but almost not existent in villages. See Table C.1 in the Appendix for more detailed descriptive information on the distribution of groups.

The upper panel in Table 4.1 presents alternative measures of ethnic and religious fractionalization at the community-level. According to all measures, and in line with the pie charts in Figure 4.1, fractionalization is relatively low except for ethnic groups in semi-urban communities. According to the ethno-linguistic fractionalization index (ELF), the average probability of randomly drawing two ethnically different households is twice as high in semi-urban than in rural communities, while no such pattern is observed with respect to religion. The two bottom panels in Table 4.1 report the composition of the local targeting committees which, overall, may be regarded as fairly representative of the community population. Higher ethnic fractionalization in semi-urban communities is reflected by more heterogeneous committees. About two-thirds of rural and one-third of semi-urban committees are occupied by three representative from the same ethnic group,

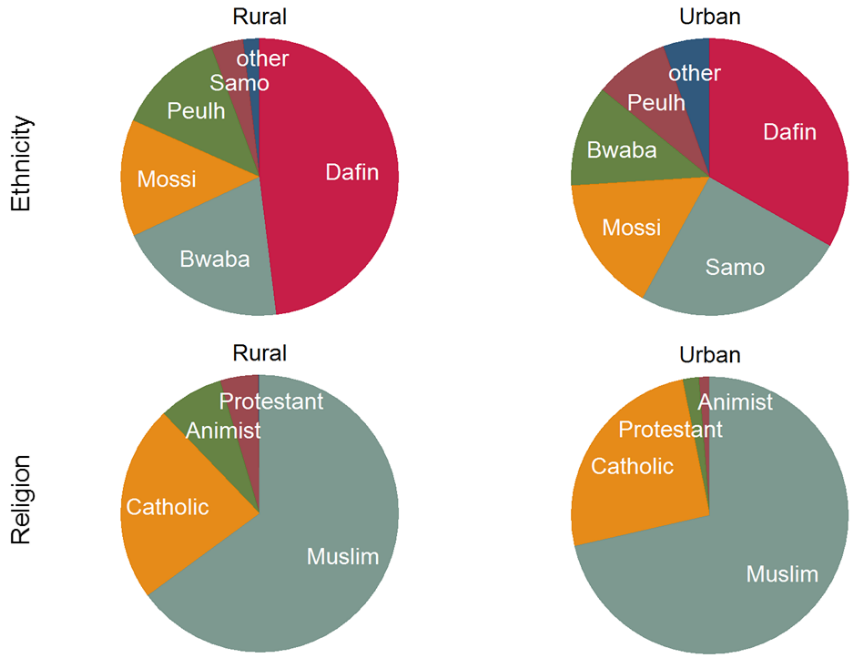

Figure 4.1: Average community population shares by sector

Notes: Pie charts depict average community population shares of six ethnic and four religious groups in the study area by sector.

respectively. The slightly higher religious fractionalization of rural communities is associated with a higher likelihood to observe a non-homogeneous committee, when compared to semi-urban communities. Lower fractionalization is associated with a higher average number of representatives belonging to the majority ethnic or religious group. The opposite holds for representatives from minority groups and representatives from ethnically or religiously mixed households.

Overall, I use a set of 49 socio-economic indicators covering demographics (7), occupational choice (11), dwelling conditions (7), livestock (8), non-productive assets (13), and health insurance preferences (3). In addition, I construct a variable that represents a household's position in the overall distribution of households' approximated landholdings.[9] Table 4.2 separately presents summary statistics for all, representative, and beneficiary households by sector. Overall, rural (column 2) and semi-urban (column 5) households are similar with respect to demographic characteristics and show only slight differences in education and migration patterns.[10] On average, semi-urban

[9]In particular, from the population census, I first took the households' statements on their agricultural output during last harvest, measured in kilogram. For each product the FAO gives the approximate cropland required per kilogram. By multiplying the latter with the output amount I calculated the approximate plot size used by each household. Based on this measure, I grouped households into deciles, for rural and semi-urban subsamples, separately.

[10]Merging the CBT data with the population survey data led to a loss of 735 rural and 400 semi-urban households. Most rural households are lost due to the exclusion of eight rural villages, where no information on the representatives could be obtained. Lost households in Nouna town are distributed evenly across neighborhoods.

Table 4.1: Community and committee characteristics: ethnicity and religion

	Ethnicity		Religion	
	Rural	Semi-urban	Rural	Semi-urban
A. Fractionalization at the community-level				
Number of groups per community	3.79	5.41	3.18	3.00
Average share of majority group	0.78	0.43	0.74	0.79
Average share of minority group	0.17	0.04	0.19	0.04
Ethno-linguistic fractionalization index (ELF)	0.32	0.67	0.36	0.32
B. Fractionalization at the committee-level				
Homogeneous committees	0.64	0.27	0.46	0.77
Majority committees	0.29	0.55	0.46	0.23
Heterogeneous committees	0.07	0.18	0.07	0.00
C. Average number of representatives				
... from majority group	2.43	0.82	2.07	2.41
... from minority group	0.04	0.36	0.43	0.05
... with mixed households	0.14	1.41	0.57	0.45

Notes: Panel A reports mean values calculated at the community level. The ethno-linguistic fractionalization index (ELF) measures the probability that two randomly drawn individuals belong to different ethnic or religious groups, respectively. In each community a local targeting committee with three representatives was formed. Panel B and C report mean values at the committee level.

households are less often occupied in the agricultural sector, show more sophisticated dwelling characteristics, higher possession incidences for assets, and higher health insurance utilization than rural households, while the opposite is true for livestock possession. When comparing the unrestricted sample with the subsample of political leaders, for the latter one observes higher mean values for a wide range of wealth-indicating variables (columns 1 and 4). Beneficiary households (columns 3 and 6), on average, show higher degrees of deprivation in all three wealth dimensions than the population sample. In rural villages, for instance, beneficiary households show possession incidences less than half of those of the population sample for one-sixth of dwelling characteristics and roughly one-third of livestock and asset indicators.

The community-based targeting design of consideration deliberately involves the participation of community members with political influence (Savadogo et al., 2015). Participants of the publicly summoned community meetings were instructed to propose three representative who, in turn, made a choice on the set of beneficiary households. Hence, within each community there is endogenous committee formation. While missing information on political power relations of these

Since the timing of CBT data collection does not perfectly overlap with the population survey period, matching of households performed worse in Nouna town, which has a much higher degree of in- and out-migration than the average rural village.

representatives, with the following two equations I set out to analyze their ethno-linguistic and socio-economic characteristics:

$$\mathbb{1}\{\text{Representative}\}_{her} = \alpha_e + \mu_r + \phi_c + u_{her}, \tag{4.6}$$

$$\begin{aligned} \mathbb{1}\{\text{Representative}\}_{her} = \phi_c &+ \gamma_1 \cdot Demographics_{her} + \gamma_2 \cdot Occupation_{her} \\ &+ \gamma_3 \cdot Housing_{her} + \gamma_3 \cdot Livestock_{her} + \gamma_3 \cdot Assets_{her} + u_{her}. \end{aligned} \tag{4.7}$$

The unit of observation is a household h, the left-hand side dependent variable is one for representative households and standardized by community size. First, in equation 4.6, I analyze whether representatives are more likely to come from certain ethnic groups by regressing the dependent variable on a set of ethnicity (α_e), religion (μ_r) and community (ϕ_c) fixed effects. I am mainly interested in the single intercepts of $\widehat{\alpha_e}$, the ethnicity fixed effects, which also capture across-group variation in all other socio-economic characteristics (the base ethnic group is 'other'). Second, in equation 4.7, I check whether the mandated political leaders are representative in terms of socio-economic characteristics or whether they are, for instance, outstandingly wealthy. To this end, I have applied principal-components analysis to construct socio-economic aggregates from my five sets of census variables in Table 4.2, and take a household's standardized rank in the corresponding PCA-based index as explanatory variable. I present regression results of equation 4.6 and 4.7 in panel A and B of Table 4.3, respectively. The 134 representatives identified in the dataset account for about 3 percent of all sample households. Overall, I find evidence that all semi-urban households not coming from the ethnic group 'other' have a 3 to 10 percentage points higher likelihood of being an elected representative. Overall, Peulh households show the highest coefficient for semi-urban communities. The second panel shows that semi-urban representative households are relatively well-equipped with livestock when compared to the average normal household.

Table 4.2: Census survey summary statistics by sector and sample specification

	Rural			Semi-urban		
Sample specification:	Represen-tatives (1)	All HHs (2)	Benefi-ciary HHs (3)	Represen-tatives (4)	All HHs (5)	Benefi-ciary HHs (6)
Demographics						
Household size	8.70	8.24	5.76	9.60	7.85	5.87
Age of household head	47.47	50.56	53.29	52.40	52.90	56.91
Household had is prime age	0.70	0.63	0.52	0.58	0.57	0.42
Share of children below age of 16	0.41	0.39	0.32	0.37	0.35	0.29
Share of prime age household members	0.40	0.38	0.37	0.44	0.44	0.43
Share of elderly household members	0.08	0.13	0.23	0.12	0.14	0.23
Any death of prime age member (2007-09)	0.04	0.04	0.05	0.04	0.05	0.05
Occupational choice						
HH head's graduation	73.12	78.63	85.66	48.40	57.56	73.03
HH head is literate	0.25	0.18	0.12	0.53	0.35	0.16
HH head is no employed in agric.	0.16	0.16	0.24	0.36	0.54	0.56
HH head never emigrated	0.05	0.05	0.08	0.02	0.03	0.03
HH head never imigrated	0.11	0.15	0.19	0.05	0.09	0.10
HH head migrated into study area	0.37	0.33	0.28	0.13	0.21	0.21
Highest graduation within HH	93.31	94.95	94.56	86.65	86.58	90.13
Number of literate HH members	7.05	6.45	4.49	12.55	9.42	6.33
Number of non-agric. employed HH members	6.28	5.61	3.73	8.35	6.39	4.49
Number of emigrated HH members	0.53	0.56	0.54	0.78	0.47	0.40
Number of imigrated HH members	1.59	1.65	1.28	2.16	1.62	1.28
Dwelling characteristics						
Good roof (better than tile)	0.27	0.23	0.09	0.53	0.52	0.29
Good wall (more than straw)	0.04	0.03	0.01	0.15	0.18	0.07
Good toiled toilet (not open field)	0.28	0.24	0.16	0.93	0.88	0.86
Number of living rooms	4.08	3.81	2.84	4.76	4.54	3.61
Light with electricity	0.00	0.00	0.00	0.38	0.39	0.15
Not cooking with wood	0.00	0.02	0.01	0.00	0.05	0.02
Good water connection	0.41	0.34	0.35	0.62	0.56	0.50
Livestock possession (incidences)						
Bullock	0.61	0.56	0.24	0.56	0.36	0.21
Sheep	0.42	0.40	0.21	0.55	0.39	0.25
Goat	0.62	0.63	0.45	0.27	0.21	0.14
Donkey	0.59	0.53	0.26	0.64	0.46	0.33
Horse	0.00	0.01	0.00	0.00	0.01	0.01
Pig	0.22	0.26	0.21	0.16	0.17	0.13
Chicken	0.94	0.83	0.68	0.67	0.61	0.48
Rabbit	0.03	0.01	0.01	0.02	0.02	0.02
Assets possession (incidences)						
Bike	0.99	0.90	0.77	0.89	0.92	0.84
Motorbike	0.41	0.28	0.07	0.38	0.38	0.11
Car	0.00	0.00	0.00	0.00	0.03	0.00
Wagon	0.52	0.49	0.19	0.60	0.48	0.33
Radio	0.82	0.75	0.56	0.91	0.83	0.62
TV	0.15	0.09	0.03	0.29	0.33	0.10
Mobile phone	0.41	0.26	0.10	0.64	0.64	0.40
Refrigerator	0.00	0.00	0.00	0.11	0.08	0.01
Sewing machine	0.04	0.04	0.01	0.07	0.10	0.05
Mill	0.00	0.01	0.00	0.00	0.01	0.00
Gun	0.14	0.12	0.06	0.09	0.06	0.00
Modern kitchen	0.00	0.00	0.00	0.02	0.07	0.01
DVD player	0.03	0.02	0.01	0.16	0.16	0.03
Health insurance						
Any HH members enr once betw 2007-09	0.06	0.04	0.07	0.13	0.09	0.15
Any HH members enr > once betw 2007-09	0.08	0.04	0.03	0.16	0.09	0.08
Number of communities	28	28	28	22	22	22
Number of households	79	2920	549	55	1653	305
Share of all households	0.03	1.00	0.19	0.03	1.00	0.18

Notes: The table reports mean values from the merged dataset across three subpopulations and by sector.

Table 4.3: Representatives' characteristics

	Being elected representative*100 / Community size					
	Panel A: Ethnic differences			*Panel B: Socio-economic Differences*		
	Pooled (1)	Rural (2)	Semi-urban (3)	Pooled (4)	Rural (5)	Semi-urban (6)
Ethnic group						
Dafin	0.05***	0.08*	0.03**			
	(0.02)	(0.04)	(0.01)			
Bwaba	0.04**	0.05	0.04**			
	(0.02)	(0.04)	(0.02)			
Mossi	0.05***	0.05	0.06**			
	(0.02)	(0.04)	(0.03)			
Peulh	0.04*	0.01	0.10***			
	(0.02)	(0.04)	(0.03)			
Samo	0.07***	0.11	0.06***			
	(0.02)	(0.07)	(0.02)			
Indices of socio-economic position (PCA)						
Demographics				0.02	0.02	0.00
				(0.02)	(0.02)	(0.01)
Occupation				0.00	-0.04	0.03
				(0.02)	(0.03)	(0.02)
Housing				0.01	0.00	0.02
				(0.02)	(0.02)	(0.03)
Livestock				0.01	-0.02	0.05**
				(0.03)	(0.03)	(0.02)
Assets				-0.01	0.02	-0.03
				(0.02)	(0.03)	(0.02)
Landholding decile				0.00	0.00	0.00
				(0.00)	(0.00)	(0.00)
Controls (50)	NO	NO	NO	NO	NO	NO
Community FEs	YES	YES	YES	YES	YES	YES
Ethnicity FEs	YES	YES	YES	NO	NO	NO
Religion FEs	YES	YES	YES	NO	NO	NO
Number of households	4573	2920	1653	4573	2915	1635
Number of representatives	134	79	55	134	79	55

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. Indices for demographics, occupation, housing, livestock, and assets are obtained through principal-components analysis from a set of 7, 11, 7, 8, and 13 variables, respectively (see Table 4.2). Landholding decile: a household's position in the overall distribution of households' approximated landholdings, based on households' self-reported statements on agricultural output in the last season.

4.3 Results

Ethnic and religious favoritism

Estimation results for my main favoritism specifications from section 4.1 are presented in Table 4.4 on page 73. For the ease of interpretation, I have multiplied the dependent variable by 100, if not stated otherwise and report pooled as well as sector-wise results. Columns 1−6 and 7−12 depict individual and committee allocation decisions, respectively. According to Panel A, representatives similarly favor coethnic households in comparison to ethnically represented non-coethnic households. Even though the effect for the latter is about two percentage points greater than for the former, I cannot reject the hypothesis that both coefficients are equal at conventional levels. I thus interpret them as being similarly large and take it as evidence against the argument that my estimates are primarily driven by an unintended information-based favoritism. Overall, these effects are clearly driven by and precisely estimated for semi-urban communities, while no effects are found for religious affiliation. In addition, the results from Panel A suggest that representatives in semi-urban communities carefully take into account the targeting preferences of the joint committee.

Given the majority-based aggregation rule such collusive or strategic behavior likely reinforces the extent of ethnic favoritism at the committee level, for which Panel C sets out corresponding estimation results. The weakly significant effect of ethnic favoritism by rural representatives is not transmitted to the final allocation decision but instead vanishes. On the other hand, strong ethnic favoritism in individual choices made by semi-urban representatives is reinforced at the committee-level, where being an ethnically represented household increases the likelihood for receiving benefits by 6.40 percentage points, statistically significant at the one percent level. This gives a distortion measure of 0.083, indicating that more than eight percent of the total number of beneficiary households have been erroneously excluded due to favoritism.

Panels B and D in Table 4.4 present heterogeneous effects with respect to community fractionalization for the individual and committee allocation choice, respectively. Overall, I do not find any evidence for religious favoritism after controlling for the degree of religious diversity. Further, it is evident that the main action of ethnic favoritism comes from ethnically high-fractionalized semi-urban communities. That is reflected by large and highly significant coefficient estimates in the first row of column 6 and 12 (which is $\hat{\delta_1}$ from equation 4.3 and 4.4). For both individual and committee decisions, the estimated effect roughly doubles when I only consider high fractionalized semi-urban communities. Being ethnically represented in a high fractionalized semi-urban community increases the likelihood for receiving benefits by more than 11 percentage points, which translates into a relative number of erroneously excluded households of almost 15 percent (see Table C.2 in the Appendix for values applied for the exclusion error calculation).

Given that semi-urban communities, on average, are more ethnically diverse than rural communities, these findings put into question whether the difference in rural and semi-urban environments matters at all for identifying ethnic favoritism, or whether ethnic diversity is the eminent factor. My illustration in Figure 4.2, where I have ranked all communities (rural and semi-urban alike) by their ELF index value suggests the latter. It shows that 20 out of the 25 most diverse communities are semi-urban and that only one semi-urban community has a below-the-median ELF rank.

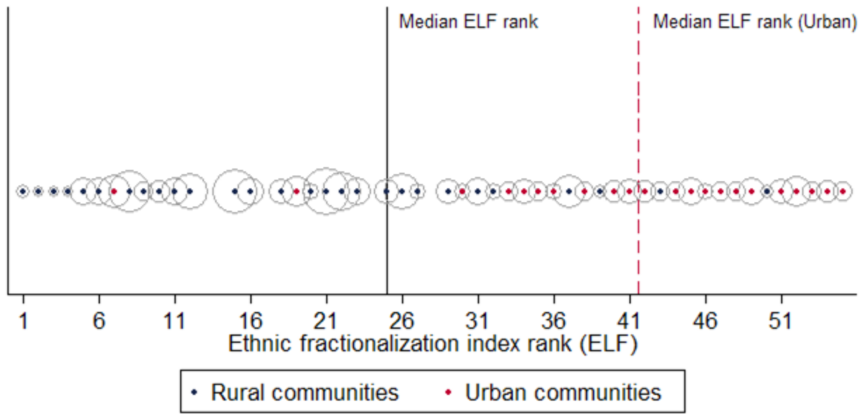

Figure 4.2: Rural and semi-urban communities, ordered by ethnic fractionalization

Notes: The graph ranks all communities (rural villages and semi-urban neighborhoods alike) according to their value of ethno-linguistic fractionalization (ELF). The community with the lowest ELF index is ranked first. Grew circles indicate a community's relative population size. Urban refers to the subsample of households from semi-urban communities.

Regression results from Table 4.4 make me confident that I do not confuse the intention to favor ethnically represented households in the final allocation decision with some unintentional information-driven favoritism. Additionally, I want to test whether my estimates are driven by consultation-based decisions to subsequently include additional or exclude excessive households, which took place at the committee level for one-fourth of beneficiary households. Even if mainly motivated by identification concerns, this exercise adds additional general insights by exploring whether and how an elite-capture outcome changes with the degree of discretion that is granted to the targeting committee. To this end, I estimate equation 4.1 only with households that qualified as candidates for committee consultation and present results in Panel A of Table 4.5 on page 76. My reduced sample covers 38 communities, indicating that in seven rural and five semi-urban communities all beneficiary household were selected by majority rule. I do not find evidence that ethnic favoritism in semi-urban communities is primarily driven by decisions where representatives were given some room for discretion. Given a relatively high targeting share of about 0.27 in this subsample, coefficient estimates are similarly large when compared to the unrestricted specification in Panel C of Table 4.4.

In addition to ethno-linguistic fractionalization, it is reasonable to assume that benefit capture is associated with a community's degree of economic inequality. This notion is based on the argument that higher economic inequality reflects a power advantage for local elites over non-elites and that the elites' ability to capture resources is increasing in the latter (Bardhan and Mookherjee, 2006). Thus, in what follows I estimate whether ethnic favoritism systematically differs with a community's observed inequality in wealth. My measure of economic inequality is the Gini coefficient calculated for an aggregated asset index based on principal-components analysis involving 28 census variables on dwelling characteristics, livestock and asset possession (see Table 4.2). My inequality variable, again, is based on a sample split at the Gini coefficient median value, and equals one for values above the median. The resulting regression equation is equal to equation 4.3, after substituting the two indicators {*Ethnic homogenous Community*} and {*Religious homogenous Community*} with the indicator {*Non-egalitarian Community*}. Estimation results in Panel B of Table 4.5 demonstrate that ethnic favoritism is strongest in the 50 percent most egalitarian semi-urban communities, suggesting a negative relationship between ethnic diversity and economic inequality. This is confirmed by Figure C.1 in the Appendix, which depicts the relationship between economic inequality and fractionalization for semi-urban communities. Coefficient estimates for religious favoritism provide suggestive evidence that religiously represented households are discriminated in semi-urban communities with high economic inequality levels; being religiously represented by the committee decreases the chances for benefit entitlements by 6.19 percentage points.

Table 4.4: Test for ethnic and religious favoritism

	Being targeted poor by representative*100						Being finally targeted poor*100					
	Panel A			*Panel B*			*Panel C*			*Panel D*		
	Pooled	Rural	Urban	Pooled	Rural	Urban	Pooled	Rural	Urban	Pooled	Rural	Urban
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Ethnicity												
HH ethnicity is reprsn. (in fract. commun.)				4.93***	3.06	8.85***	4.46***	1.38	6.40***	5.95***	1.71	11.49***
				(1.42)	(1.95)	(2.32)	(1.57)	(2.49)	(2.20)	(1.86)	(2.69)	(2.94)
Represented and coethnic	3.32***	3.19*	3.67**									
	(1.26)	(1.92)	(1.80)									
Represented but not coethnic	3.61***	1.94	5.73***									
	(1.39)	(2.10)	(1.94)									
H0: Ethn. coefficients are equal (p-value)	0.78	0.39	0.17									
Ethnicity represented * Low fractionalization				-4.65**	-1.66	-9.41***				-4.92*	-2.85	-10.98***
				(2.28)	(3.67)	(3.20)				(2.91)	(4.67)	(4.06)
HH ethnicity is reprsn. in homog. commun.				0.28	1.40	-0.56				1.03	-1.13	0.51
p-value				0.88	0.70	0.81				0.68	0.80	0.87
Religion												
HH religion is reprsn. (in fract. commun.)				-1.03	-0.67	-1.10	-0.65	0.47	-1.87	-1.09	0.13	-2.03
				(1.24)	(1.54)	(2.33)	(1.58)	(2.04)	(2.81)	(1.67)	(2.16)	(3.01)
Represented and coreligious	-0.37	0.16	-1.21									
	(1.20)	(1.51)	(2.17)									
Represented but not coreligious	-2.17	-1.71	-1.88									
	(1.38)	(1.62)	(2.88)									
H0: Rel. coefficients are equal (p-value)	0.05	0.07	0.74									
Religion represented * Low fractionalization				3.29	7.15	-0.34				5.17	10.18	1.02
				(3.01)	(5.02)	(3.93)				(4.07)	(7.35)	(5.07)
HH religion is reprsn. in homog. commun.				2.26	6.49	-1.44				4.07	10.31	-1.01
p-value				0.44	0.19	0.70				0.30	0.15	0.84
Controls (50)	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Community Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Ethnicity Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Religion Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Number of observations	13650	8745	4905	13650	8745	4905	4550	2915	1635	4550	2915	1635
Number of households	4550	2915	1635	4550	2915	1635	4550	2915	1635	4550	2915	1635
Number of communities	50	28	22	50	28	22	50	28	22	50	28	22
Number of key informants	150	84	66	150	84	66	150	84	66	150	84	66

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the household level in parentheses. For ethnicity and religion, H0 tests whether $\hat{\beta}_1 = \hat{\eta}_1$, and $\hat{\beta}_2 = \hat{\eta}_2$, respectively (see equation 4.2). Urban refers to the subsample of households from semi-urban communities. Low fractionalization: equals one for the fifty percent least fractionalized communities ('homog. commun.' in the table above) and zero otherwise ('fract. commun.' in the table above). The fractionalization measure is the ethno-linguistic fractionalization Index (ELF).

Capture of program benefits

As an extension to the main favoritism analysis, I want to test for two more dimensions of benefit capture. First, and similar to existing elite capture studies (Bardhan and Mookherjee, 2006; Besley et al., 2012; Pan and Christiaensen, 2012; Beekman et al., 2014; Schüring, 2014), I test whether political decision makers reallocate a disproportionally large share of program benefits to themselves. Such behavior would be reflected by self-selection into the beneficiary lists, conditional on observable wealth. Note that the targeting design of consideration naturally limits the extent of capture to the extensive margin, as each household could only redeem one voucher. Nevertheless, representatives were allowed to preselect their own households for benefits and that is what about 17 percent of them did. The ratio is considerably higher in rural communities and, overall, 13 percent of representatives belong to the set of beneficiary households. In the same spirit as equation 4.1, I estimate the following regression equation

$$\mathbb{1}\{\text{Targeted}\}_{cher} = \alpha_e + \mu_r + \phi_c + \mathbf{X}'\gamma + \beta_1 \cdot \mathbb{1}\{\text{Representative}\}_{cher} + u_{cher},$$

where the dependent variable is the committee's final targeting decision and $\{\text{Representative}\}_{cher}$ indicates whether household h provides a member from the local targeting committee. I include the common set of fixed effects and control variables. As Panel C in Table 4.5 sets out, I do not find evidence for benefit capture by representatives.

Second, I am interested to which extent favoritism of ethnically represented households allows for benefit capture by the communities' ethnic majority groups. As shown in section 3.1, community-driven programs differ in terms of agency. In two-thirds of the studies that are sufficiently explicit about their underlying community-based targeting procedure, villagers appoint representatives to determine the final targeting set, while the remaining schemes involve the community as a whole (see Table 3.1 on page 34). One major public choice argument in favor of maximizing the set of representatives is minority protection (Tullock, 1959), which often comes at higher institutional cost (Auriol and Gary-Bobo, 2012). I test for benefit capture by ethnic majority groups within a program that uses a relatively small number of representatives. By doing so, I provide descriptive insights on the optimal number of representatives in targeted welfare schemes. As Table 4.1 on page 65 shows, the small number of representatives makes ethnic or religious minority representation very unlikely. The opposite is true for majority groups and it is likely that ethnic favoritism corresponds with a situation of benefit capture by the ethnic majority group. To test for such a phenomenon I estimate

$$\begin{aligned}\mathbb{1}\{\text{Targeted}\}_{cher} =& \alpha_e + \mu_r + \phi_c + \mathbf{X}'\gamma + \beta_1 \cdot \mathbb{1}\{\text{Ethnic majority}\}_{cher} \\ &+ \beta_2 \cdot \mathbb{1}\{\text{Religious majority}\}_{cher} + u_{cher},\end{aligned} \tag{4.8}$$

and

$$\begin{aligned}\mathbb{1}\{\text{Targeted}\}_{cher} = & \alpha_e + \mu_c + \phi_c + \mathbf{X}'\gamma \\ &+ \delta_1 \cdot \mathbb{1}\{\text{Ethnic majority}\}_{cher} + \delta_2 \cdot \mathbb{1}\{\text{Religious majority}\}_{cher} \\ &+ \lambda_1 \cdot \{\text{Ethnic majority}* \text{ Ethn. homogenous community}\}_{cher} \\ &+ \lambda_2 \cdot \{\text{Religious majority} * \text{Rel. homogenous community}\}_{cher} + u_{cher},\end{aligned} \tag{4.9}$$

where the dependent variable reflects the committee's final allocation choice, while $\{\text{Ethnic majority}\}_{cher}$ and $\{\text{Religious majority}_{cher}\}$ equal one for a household h belonging to the community's ethnic or religious majority group, respectively. Remaining regression components are equal to my former specifications.

Regression results for equation 4.8 and 4.9 are set out in the first and second block of Panel D, respectively. They indicate that favoritism of ethnically represented households coincides with a situation in which ethnic majorities capture disproportionally more private benefits than non-majority households. On average, and conditional on observable wealth, they are five percentage points more likely to receive program benefits, which gives an exclusion error of more than six percent. In line with my favoritism results, benefit capture by ethnic majority groups is about 75 percent higher when one only considers highly fractionalized semi-urban communities, which amounts to an exclusion error of about 11 percent.

Table 4.5: Additional specifications for ethnic and religious favoritism

	Being finally targeted poor∗100									
	Panel A: Favoritism with room for discretion		*Panel B: Favoritism and econ. inequality*		*Panel C: Benefit capture by representatives*		*Panel D: Benefit capture by ethnic and religious majority groups*			
	Rural	Urban	Rural	Urban	Rural	Urban	Rural	Urban	Rural	Urban
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Representative's household					-6.62	-1.46				
					(5.02)	(7.85)				
Ethnicity										
HH ethnicity is reprsn. (in fract./egalit. commun.)	16.00*	11.11	3.96	9.61***						
	(9.14)	(8.67)	(4.49)	(3.04)						
Ethn. majority (in fractionalized commun.)							-1.03	5.02**	-0.60	8.71***
							(2.50)	(2.06)	(2.66)	(3.00)
Ethn. treatment * Low fractionalization									0.23	-7.98*
									(4.42)	(4.32)
Ethn. treatment * High inequality			-3.46	-7.14*						
			(4.76)	(4.32)						
Ethn. treatment in homog./non-eglat. commun.			0.50	2.47					-0.38	0.73
p-value			0.85	0.43				.	0.93	0.80
Religion										
HH religion is reprsn. (in fract./egalit. commun.)	6.44	16.57	2.13	5.66						
	(6.25)	(11.82)	(3.47)	(3.94)						
Relig. majority (in fractionalized commun.)							-0.97	-1.65	-0.57	-1.20
							(1.97)	(2.45)	(2.06)	(2.62)
Rel. treatment * Low fractionalization									-5.00	-0.88
									(5.89)	(4.98)
Rel. treatment * High inequality			-2.44	-11.86***						
			(4.14)	(4.44)						
Relig. treatment in homog./non-eglat. commun.			-0.31	-6.19*					-5.57	-2.08
p-value			0.90	0.06					0.33	0.66
Controls (50)	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Community Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Ethnicity Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Religion Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Number of observations	526	188	2915	1635	2893	1021	2915	1635	2915	1635
Number of households	526	188	2915	1635	2915	1635	2915	1635	2915	1635
Number of communities	21	17	28	22	28	22	28	22	28	22
Number of key informants	18	3	84	66	84	66	84	66	84	66

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. Estimations in Panel A use the subsample of households that qualified for subsequent inclusion into or exclusion from the beneficiary list. Urban refers to the subsample of households from semi-urban communities. The fractionalization measure is the ethno-linguistic fractionalization Index (ELF). The inequality measure applied in Panel B is the Gini index for an asset index obtained through principal-components analysis from a total of 28 assets. Low fractionalization: equals one for the fifty percent least fractionalized communities ('homog. commun.' in the table above) and zero otherwise ('fract. commun.' in the table above); Ethnic and religious majority; equals one for households that belong to the biggest ethnic or religious group of their community; High inequality: equals one for the fifty percent most unequal communities ('non-eglat. commun.' in the table above) and zero otherwise ('egalit. commun.' in the table above).

4.4 Robustness Checks

If ethnic and religious affiliation are highly correlated it is problematic to interpret my effect as working only through the channel of the former. In such a case, separate estimation of ethnic and religious favoritism likely captures a similar degree of the allocation bias due to favoritism. To check for multicollinearity, I separately test for ethnic and religious favoritism in the committee decision and present results in Panel B of Table C.3 in the Appendix. I find that testing for ethnic or religious favoritism alone does not change the results. As a second robustness check, I consider an alternative model specification. As Panel C shows, the pattern of my results remain qualitatively equal when estimating a Probit instead of a Linear Probability model.

In what follows, I check robustness of the heterogeneous effects with respect to ethnic fractionalization. For the set of 22 semi-urban communities, Figure C.2 depicts community-wise coefficient estimates of ethnic favoritism (particularly, $\hat{\beta}_1$ in equation 4.1) ordered by their rank in ethnic fractionalization. My fractionalization indicator, {Ethn. homogenous Community}, is one for communities below the median rank, represented by the vertical dashed line in Figure C.2. Even though community-wise coefficients are not precisely estimated, this figure provides two insights. First, it is reassuring to not observe any bunching of outlier coefficients close to the threshold. Hence, the heterogeneous results are unlikely driven by the choice of applying a median-based sample split. Second, the community which is ranked third qualifies for an outlier. When excluding this community and rerunning the main estimations I do not find significant changes in the results. A second robustness check involves the estimation of equation 4.3 with an alternative measure of fractionalization, the share of a community's ethnic or religious majority group. As I expect a strong negative correlation between ELF values and the share size of a community's majority (see Table 4.1), I adjust my fractionalization indicator {Ethn. homogenous Community} such that it equals one if community c has an above-the-median majority share. Panel D in Table C.3 shows that the community classification based on this alternative measure gives very similar estimation results when compared to the ELF classification in Panel D of Table 4.4. Finally, Table C.4 sets out results from an estimation of equation 4.3, where I have substituted the community-specific diversity identifier {Ethn. homogeneous community$_c$} with the actual rank in fractionalization {Ethn. fractionalization rank$_c$}. As expected, one only observes a precisely estimated interaction effect for ethnicity in semi-urban communities.

4.5 Discussion

I shall close this chapter by summarizing and discussing my two main results. First, I find evidence for ethnic favoritism by semi-urban elites, leading to an exclusion error of about 8 percent wrongly excluded households, as a share of all beneficiary households. My effects appear modest when compared to the distortions caused by ethnic favoritism at the central level, where Burgess et al. (2015) find a twofold increase in road investments for coethnic districts. When compared to other elite capture dimensions, my results are mainly in line with the existing literature. Alatas et al. (2012) consider a low-stake cash-transfer program in Indonesia and find no evidence for elite capture. For the same context, the authors also analyze a high-stake scenario (Alatas et al., 2013a) for which they can distinguish between two types of political leaders as well as between the targeting and distribution step. They find modest elite capture by formal elites at the stage of benefit distribution. With an exclusion error of 4 percent wrongly excluded households their effect is half as strong as ours. In addition, my results correspond with the extent of elite capture found for two programs in sub-Saharan Africa, where local governments make an allocation choice on private in-kind benefits. Favoritism of politically connected households in Ethiopia (Basurto et al., 2016) and kin households in Malawi (Caeyers and Dercon, 2012) lead to exclusion errors of 5 and 15 percent, respectively. Finally, and in line with my results, Schüring (2014) finds no evidence for religious favoritism in Zambia.

Second, my results indicate that community characteristics matter. I find that ethnic favoritism is increasing in communities' ethnic fractionalization, thereby confirming a large body of cross-country studies on the adverse relationship between ethnic fractionalization and political economy outcomes. While this empirical finding is powerful and robust across many studies −especially for public good outcomes at the country level− less is known about the underlying channels. One prominent hypothesis in this regard states that people would contribute relatively less to a welfare-enhancing public good in an ethnically diverse community, assuming they give higher weights to the utility of coethnic community-members (Habyarimana et al., 2009). Similarly, for corruption, Mauro (1995) expects bureaucrats to be relatively more willing to demand (higher) bribes from non-coethnics such that bribe demand should be higher in ethnically diverse environments. Based on the idea that ethnic favoritism reflects the existence of discriminatory preferences, my results support the plausibility of the prominently hypothesized 'preference mechanism'. Following Mauro (1995)'s argument, I would also expect ethnic favoritism −as one manifestation of elite capture− to increase with ethnic fractionalization. This is different to two recent studies, which have found that lower public good contributions in urban Ugandan neighborhoods (Habyarimana et al., 2007) and rural Tanzanian villages (Miguel and Gugerty, 2005) is mainly driven by the inability to enforce social sanctions within more diverse environments. As far as capture by political leaders is considered, three studies find that there is higher perceived corruption in more ethnically fractionalized countries (Mauro, 1995; La Porta et al., 1999; Treisman, 2000). At the subnational level, my results confirm Olken (2006), who finds significantly higher leakage in publicly distributed rice for ethnically diverse areas in Indonesia. Interestingly, he finds no such effect for religious fractionalization, neither. Overall, with my study I confirm that "one of the most powerful hypotheses in political economy" (Banerjee et al., 2005, p.639) persists when looking at one important manifestation of elite capture at the most decentralized level.

Chapter 5

Subsidized Health Insurance for the Poor: Effects on Health Insurance Demand

After I have analyzed two important aspects of poverty targeting in chapter 3 and 4, I now turn to the impact evaluation of the actual targeted intervention; the allocation of a premium subsidy to encourage health insurance enrollment among poor households. My empirical strategy exploits a discontinuity in the underlying community-based targeting procedure to identify causal effects of the subsidy offer on health insurance demand. This analysis directly contributes to two strands of literature.

First, I contribute to the literature on health insurance demand. To the best of my knowledge, only five studies rigorously evaluate price interventions aimed at expanding health insurance uptake in low-income countries.[1] Recently, Capuno et al. (2016) and Wagstaff et al. (2016) have evaluated social health insurance interventions that randomly varied the price of and information on the insurance product in the Philippines and Vietnam, respectively, and find significant enrollment effects. Thornton et al. (2010) find that Nicaraguan workers are more likely to buy

[1] The majority of studies on the relationship between price and insurance take-up relies on hypothetical willingness to pay studies which are presented elsewhere. These can suffer from hypothetical bias and results strongly depend on the experimental set-up (Stewart et al., 2002). Several authors have studied the price elasticity of demand for health insurance in the U.S. and mainly find fairly price-inelastic demand for health insurance (Chernew et al., 1997; Blumberg et al., 2001; Gruber and Washington, 2005; Royalty and Hagens, 2005; Finkelstein et al., 2012).

health insurance when being initially offered six months of free coverage and King et al. (2009) find considerable take-up effects for an intervention that offers improved health care supply at lower prices in Mexico. For a social health insurance in Cambodia, Levine et al. (2016) play a lottery with community-members who participated in an information campaign to offer them health insurance at two different prices. Their subsidized package offered participants five months of free health insurance and six months of health insurance at a 50 percent discount (leading to an effective premium discount of 80 percent) and increases enrollment from about 7 to 50 percent.[2] My study design comes with two favorable features when compared to the existent literature. First, it establishes exogenous variation in health insurance pricing only and does not combine the latter with an additional treatment, such as information (Thornton et al., 2010; Capuno et al., 2016) or supply side incentives (King et al., 2009). Second, the 50 percent subsidy from the context of my study nicely complements the two remaining study designs from the South-East Asian context, that consider the effect of a much smaller (25 percent subsidy in Wagstaff et al., 2016) and a much bigger (six months at price zero in Levine et al., 2016) 'pure' price reduction.[3] My main methodological contribution to this literature is the analysis of a pricing effect on health insurance take-up along the household wealth distribution. In my view, investigating this relationship within the same empirical setting is crucial because the inability to afford an insurance premium has been identified as a main obstacle to insurance take-up by poor households (Jakab and Krishnan, 2004; Giné et al., 2008). Furthermore, I am first in evaluating the effect of pricing on households' health insurance demand at the intensive margin, which contributes to the sub-literature on intra-household allocation of health inputs. Finally, this study provides novel evidence on the price elasticity of health insurance demand for the sub-Saharan African context, where average health indicators perform especially poorly and health insurance outreach faces enormous challenges (World Bank, 2014).

Regarding the second directly related literature, my study provides a rigorous test for adverse selection into health insurance in a low-income context. Social health insurance often uses the insurance's pooling mechanism as a market-based mean to achieve and maximize financial health protection for the poor (Yao et al., 2015). In this sense, evaluating the extent of adverse selection

[2] Another related study in this regard is Fischer et al. (2016), who randomly offer six different health insurance contracts to households in Pakistan in order to evaluate patterns of adverse selection.

[3] All three study designs suffer from minor limitations in generalizability in different ways. Wagstaff et al. (2016) look at a sampled sub-population of non-poor households which are self-employed or without a long-term contract, while Levine et al. (2016) randomize prices only among those households who attend the insurer's information campaign meeting. While my sampling frame considers the whole universe of households living in the study area, the RDD approach only allows me to identify causal effects for households that are close to the beneficiary cutoff.

is regarded important from both an efficiency and equity perspective; it informs not only on a scheme's financial self-sufficiency but also on its potential to cover the neediest and thus most risky individuals in the population. The literature on adverse selection in health insurance markets of high-income countries finds mixed results (Cohen and Siegelmann, 2010). Even though one also expects Ackerlof's (1970) main theoretical assumptions about the mechanisms of adverse selection to hold in a low-income setting, such environments still could produce different results. First, there is less capacity for ex-ante risk screening of clients as social health insurance often offers only one very simple product (Brau et al., 2011). Second, studies in high-income countries usually focus on switching behavior of clients between different insurance plans, while health insurance markets in the global South commonly offer only a choice between buying and not buying health insurance (Cohen and Siegelman, 2010). Third, risk-seeking individuals might be relatively more likely to buy health insurance in low- than in high-income settings as it constitutes a new and largely unknown product in the former (Polimeni and Levine, 2011).

The literature mainly considers two empirical approaches to test for adverse selection in health insurance markets (see Yao et al., 2015 for a recent review). The first approach examines the correlation between individuals' decision to buy health insurance and health care utilization. Since the choice on the latter might be affected by insurance status, such an approach cannot disentangle adverse selection from moral hazard (Chiappori and Salanie, 2000). Alternatively, one can address this separation problem by estimating the relationship between ex-ante health status and health insurance take-up. Studies that apply this second empirical approach for low-income countries, however, usually do not observe ex-post health care utilization and thus fail to quantify the financial consequences of adverse selection (Fischer et al., 2016).[4] Following a more recent alternative approach, I exploit exogenous health insurance price variation and employ health care utilization data among insured individuals to analyze whether the pool of insurees that bought health insurance at a higher price, on average, shows higher utilization levels. While this approach takes care of moral hazard and obtains a cost-relevant measure of adverse selection, exogenous price variation also allows for differences in unobservables. For the low-income context, Polimeni and Levine (2011) and Fischer et al. (2016) are the only studies which apply such a 'price test' to study adverse selection into a newly established insurance in Cambodia and Pakistan, respectively. I contribute to this literature by considering a setting where price variation takes place

[4]To my knowledge, Banerjee et al. (2014) provide the only exception in this regard. The empirical context is a newly introduced bundled health insurance by SKS Microfinance in Karnataka, India.

for a health insurance product that has been in place for three years already and by adding novel evidence for the sub-Saharan African context.

5.1 Empirical Approach

My identification strategy exploits a discontinuity in the premium subsidy allocation rule which is based on community wealth rankings. I explain in the following how I set up my fuzzy regression discontinuity design and discuss necessary identifying assumptions, before I introduce my empirical specification and outcome variables.

A fuzzy regression discontinuity design

In the local context of my study, community-based targeting of subsidy-eligible households mainly relies on an observable aggregation rule that is subject to the three representatives' wealth rankings. As I have set out in Section 2.2, three quarters of beneficiary households are selected by applying 'majority rule' to the three representatives' wealth rankings. By calculating a household's median wealth rank and using the community-specific beneficiary cutoff rank, I can mimic the majority rule as follows: A household is eligible by majority rule when its median wealth rank is below the beneficiary cutoff.[5]

Thus, with a household's median wealth rank as forcing variable, the majority rule gives a discontinuity in eligibility assignment at the community cutoff rank. Since not all finally eligible households are determined via majority rule, 'being eligible by majority rule' is a strong but not a perfect predictor for subsidy eligibility. In particular, I observe a fair fraction of 'non-complier' households which is subsidy-eligible but only mentioned by one representative. In some rare cases, where a small number of majority-eligible households had to be excluded, I also observe some 'non-compliers' which are not on the beneficiary list even though eligible by majority rule. Hence, my 'fuzzy' regression discontinuity design (RDD) uses the discontinuity in the majority rule as an instrument to predict subsidy-eligibility in a first-stage.[6] For households with a median wealth rank sufficiently close to the community's beneficiary cutoff, my identification strategy then allows

[5] To see this, consider community c with 50 households, a beneficiary share of 20 percent and a corresponding cutoff rank of 10. Assume further, that representatives $\{A, B, C\}$ assign wealth ranks $\{4, 9, 13\}$ and $\{6, 11, 12\}$ to household h and j, respectively. Then, household h is eligible by majority rule since its median wealth rank of 9 is below the beneficiary cutoff of 10, while household j is not.

[6] In Section 5.3 I run such a first-stage regression to demonstrate the relevance of my instrument. As the allocation of subsidy entitlements was the sole purpose of the community-based exercise, I do not expect the majority rule to work through a different channel than the subsidy entitlements itself (exclusion restriction).

me to estimate the effect of being offered a 50 percent premium discount on insurance take-up and other outcomes. My main identifying assumption is that the outcome variable absent the intervention must be continuous in the household median wealth rank.[7] For continuity to hold, individuals must not be able to manipulate the forcing variable, nor to precisely sort around the discontinuity threshold (Angrist and Pischke, 2009; Lee and Lemieux, 2010). In Section 5.4, I shall discuss the validity of my identifying assumption and present corresponding robustness tests.

Empirical specification

In order to obtain a lower bound of the average treatment effect (around the cutoff), I estimate the intent-to-treat effect (ITT) of being subsidy-eligible by majority rule on health insurance take-up, and utilization. As suggested by Lee and Lemieux (2010), I estimate a local linear regression model with rectangular kernels

$$\begin{aligned} y_{ch} = &\alpha_c + \beta \cdot \mathbb{1}\{EligByMajority\}_{ch} + \gamma \cdot MedianRank_{ch} \\ &+ \delta \cdot \mathbb{1}\{EligByMajority\}_{ch} * MedianRank_{ch} + u_{ch}, \end{aligned} \tag{5.1}$$

where $\mathbb{1}\{EligByMajority\}_{ch}$ is my instrumental variable, indicating whether household h from community c has been preselected for the subsidy entitlement by majority rule or not. As discussed in the former subsection, this treatment variable equals one for households with a median wealth rank below the beneficiary cutoff. If not stated otherwise, I include community fixed effects (α_c) to improve the precision of the estimation. My forcing variable is a household's median wealth rank ($MedianRank_{ch}$) which is centered at the cutoff, CDF-transformed and thus bounded between -0.2 and 0.8. I include the forcing variable as well as the interaction effect between the instrumental and forcing variable to allow for different slopes at both sides of the beneficiary cutoff. Since both the application of majority rule and the determination of final beneficiary status are community specific, I do not expect any benefit from including community fixed effects in the first-stage regression and thus leave them out. The dependent variable y_{ch} in my first-stage regression is an indicator for final eligibility status. Then, for households in a sufficiently small interval around the beneficiary cutoff, the ordinary-least square coefficient estimate $\hat{\beta}^{elig}$ gives the effect of being eligible by majority rule on final eligibility status. When y_{ch} is one of my outcome variables, $\hat{\beta}^{outc}$ gives the intent-to-treat effect of being majority-eligible for the subsidy on the

[7] To be precise, due to the aggregation rule it is even sufficient if the outcome variable is continuous in one of the three household wealth ranks.

outcome of interest. In addition, the Wald estimator, $\frac{\hat{\beta}^{outc}}{\hat{\beta}^{elig}}$, gives the local average treatment effect (LATE) of being eligible on taking up the insurance (Lee and Lemieux, 2010). Note that, economically, the LATE provides a more meaningful interpretation than the ITT. Nevertheless, I will show that the ITT is a very strong instrument for eligibility, which can be considered as a reasonable lower bound of the LATE.

Table 5.1 lists the outcome variables of my analysis. When the outcome variable is health insurance take-up, I consider the following three specifications. First, I estimate the ITT and LATE on household take-up at the extensive margin, with an outcome variable that indicates whether any member in household h has bought health insurance in the year 2009 or 2010. As an extension, I also estimate the ITT effects on take-up incidence across both years (that is: Has any household member bought health insurance at least once between 2009 and 2010?) and test whether subsidy-eligibility in 2009 affects the choice to renew health insurance into the year 2010. Second, to estimate the ITT effect at the intensive margin, I separately estimate equation 5.1 with corresponding subsets of household members (for instance, for male and female household members) and test for statistically significant differences across subgroup ITT estimates.

Third, I exploit wealth differences across semi-urban communities to estimate heterogeneous effects of the subsidy offer. I construct the community wealth variable in three steps. In a first step, I build a principal-components asset index that is based on household dwelling characteristics and asset possession indicators. I then sort all semi-urban households in the study area by the asset index to obtain household wealth percentiles. In a third step, I take out the average wealth percentile for the two households located around each community's beneficiary cutoff, w_c. Thus, w_c provides a measure for community wealth evaluated at the beneficiary cutoff and expressed in wealth percentiles of the semi-urban area's household wealth distribution.[8] The corresponding regression equation to estimate heterogeneous ITT effects is

$$\begin{aligned} y_{ch} =& \alpha + \beta \cdot \mathbb{1}\{EligByMajority\}_{ch} + \gamma \cdot MedianRank_{ch} \\ &+ \delta \cdot \mathbb{1}\{EligByMajority\}_{ch} * MedianRank_{ch} \\ &+ \pi \cdot \mathbb{1}\{EligByMajority\}_{ch} * \tilde{w}_c + \lambda \cdot \tilde{w}_c + u_{ch}, \end{aligned} \tag{5.2}$$

where $\tilde{w}_c$ is the community wealth percentile centered at the median of w_c. For the ease of interpretation, I do not include community-fixed effects (inclusion does not change the results).

[8] As I estimate local causal effects, I define w_c this way and do not take a more general measure for community wealth such as the mean asset index value.

Then, $\hat{\beta}$ is the ITT effect of subsidy-eligibility for a median-wealthy household and $\hat{\pi}$ gives the additional change in the ITT effect for moving up the wealth distribution by one percentile. Hence, equation 5.2 allows for the estimation of price elasticities of insurance demand across wealth. In addition, $\hat{\lambda}$ gives the relationship between wealth and health insurance demand such that equation 5.2 allows me to extrapolate Engel Curves at two different prices. The latter should be only drawn and interpreted for the range of wealth percentiles one observes in w_c.

Table 5.1: Set of outcome variables

Outcome variable	**Cross-section**	**Round**
Enrollment: Extensive margin		
Any HH member bought health insurance in the year 2009 or 2010	universe	2009/10
Any HH member bought health insurance at least once betw. 2009 and 2010	universe	2009
Any HH member bought health insurance in the year 2010	HHs insured in 2009	2010
Enrollment: Intensive margin		
Individual bought health insurance in the year 2009 or 2010	female vs. male	2009/10
Individual bought health insurance in the year 2009 or 2010	adolescents vs. adults	2009/10
Individual bought health insurance in the year 2009 or 2010	prime aged vs. elderly	2009/10
Individual bought health insurance in the year 2009 or 2010	core vs. extd. family	2009/10
Health care utilization		
Total household per capita number of annual health care facility visits	insured HHs	2009/10
Total household per capita amount of annual health insurance payouts	insured HHs	2009/10

Notes: Cross-section: indicates whether the whole population (universe) or a distinct subpopulation of households is covered; Round: indicates which survey round was used for the estimation.

From a health insurance perspective, risky individuals are more likely to fall ill or less likely to recover quickly and, thus, should show higher health care utilization rates. When the insurer is not able to screen for different risk types or has only one single health insurance product at offer, selection of risky types into the health insurance (adverse selection) drives up the insurer's cost in the long-term. As a consequence, the insurer has to lift up the premium, which crowds out less risky individuals with a lower willingness to pay and leads to market failure (Akerlof, 1970). In contrast to this efficiency-based perspective, adverse selection into health insurance can also be understood as a favorable feature to achieve universal health coverage, a concept that is advocated by the World Bank and aims towards a reality where every needy individual might receive health services without financial hardship (Wagstaff et al., 2016). I test for adverse selection into health insurance by analyzing whether the pool of insurees that bought health insurance at the normal price, on average, comprises higher health risk than the group that bought health insurance at the subsidized price. My outcome variable reflects health care utilization, a cost-relevant measure for the insurer's self-sufficiency that is driven by a household's risk to fall or stay ill. Specifically, I estimate equation 5.1 for the subsample of households which has been partially or fully insured at least once between 2009 and 2011. The two outcome variables, y_{ch}, are the total household per capita number of ambulance visits and the total household per capita value of annual health insurance payouts. The latter provides a lower-bound cost measure as it contains consultation capitation fees and prescribed drugs but excludes costs for treatments that were provided in the department hospital.

5.2 The Data

The empirical analysis in this chapter combines the community-based targeting data with population survey data and administrative data from the health insurance records (see Table 2.6 on page 22). Table 5.2 presents summary statistics for some key variable of the combined dataset by sector. I use the wealth ranking data to construct the forcing variable which is household h's median wealth rank. Descriptive statistics of this dataset are set out in Table 2.1 on page 24.

I construct the main outcome variables of interest from the administrative insurance data that is depicted in Panel A of Table 5.2. I consider insurance enrollment between 2009 and 2010 at the extensive (household level) and intensive (individual level) margin. While semi-urban communities show an average individual enrollment rate that is twice as high as in rural communities, the latter shows a higher percentage increase over years. By 2010, about one-sixth and one-fourth of all households had at least one member insured in rural and semi-urban communities, respectively. In line with the enrollment patterns, utilization rates increased from 2009 to 2010 and insurance payouts are on average 55 percent higher in semi-urban than in rural communities.

As Panel B illustrates, the average household consists of about 12 members, of which about six members are aged below 16 years, 5 are prime aged and 1 is above the age of 55 years. The average annual household mortality incidence is 0.2 and the average number of within-household deaths between 2009 and 2012 is 6.20 and 3.85 in rural and semi-urban communities, respectively. Finally, census survey data in Panel C of Table 5.2 illustrates that livestock possession is widespread in the study area and even more pronounced in rural communities.

Table 5.2: Summary statistics

	Rural		Semi-urban	
	2009	2010	2009	2010
A. Insurance enrollment (Admn. insurance data)				
Individual enrollment incidence	0.05	0.07	0.11	0.13
Households with at least one enrolled member	0.10	0.14	0.23	0.26
Individuals' number of ambulance visits	0.12	0.20	0.22	0.28
Households' number of ambulance visits	3.06	4.84	4.60	6.42
Amount of households' ambulance payouts	2,952.5	5,190.0	5,013.7	7,627.1
B. Demographics (Vital registration system)				
Household size	11.56	11.57	11.69	12.02
Household-share of 0-16 year-olds	0.53	0.53	0.47	0.46
Household-share of 16-55 year-olds	0.38	0.38	0.43	0.44
Household-share of >55 year-olds	0.09	0.09	0.10	0.11
Incidence of any within-household death	0.19	0.21	0.20	0.20
Number of within-household deaths over 2009/12	6.20	6.20	3.85	3.85
C. Livestock (Population census survey)				
Household owns any big livestock	0.74		0.65	
Household owns any small livestock	0.97		0.91	
Number of big livestocked household owns	5.30		3.11	
Number of small livestocked household owns	23.47		12.26	
Number of individuals	28,569	28,446	17,496	17,580
Number of households	2,791	3,128	1,506	1,884
Number of communities	28	28	22	22

Notes: The table presents mean values from the merged dataset for three different categories. Mean values in Panel A and B are calculated at the individual or household level. Asset information in Panel C is recorded at the household level.

5.3 Results

In addition to result tables, I illustrate the main local linear regression results with scatter plots, depicting the relationship between the outcome and the forcing variable. As suggested by Lee and Lemieux (2010), for all the RDD scatter plots, I calculate average outcome values by household median wealth percentile. The relative number of observations per percentile is expressed by bubble size. The plots also indicate the beneficiary cutoff value at zero and depict the fitted local linear regression line with a 95 percent confidence interval. This graphical presentation allows for immediate statistical inference. The coefficient estimate $\hat{\beta}$ reflects the distance between both fitted lines at the cutoff and is precisely estimated at conventional levels as long as the confidence intervals on both sides of the cutoff do not overlap.[9] As my identifying assumption only holds for a sufficiently small neighborhood around the beneficiary cutoff, I separately estimate equation 5.1 and 5.2 with household samples I have trimmed towards 20, 15, and 10 median-wealth-rank percentiles around the cutoff. Given an underlying non-linear functional form, the bias in my local linear approximation of the 'true' local effect increases in window size. As estimation accuracy also increases in window size, the optimal choice of window size is subject to a trade-off between a smaller bias at the expense of a loss in accuracy (Lee and Lemieux, 2010). The comparison of ITT estimates across window sizes provides a common check for the robustness of my results. Large variation in the estimates' magnitude across windows can be either driven by power issues or by a relatively bad fit of my local approximation for large windows. In either case I should be cautious when interpreting corresponding results. While my regression tables include different window specifications, I display my RDD scatter plots for only one selected window size. Summary statistics and estimation results from the study area thus far suggest that rural and semi-urban communities provide very different living environments, for instance, reflected by significant differences in asset wealth (see Figure 2.4). Hence, I am interested how health insurance pricing affects outcomes in both rural and semi-urban sectors, separately, and shall present my results in this fashion.

First-stage regression

For both rural and semi-urban communities, Figure 5.1 illustrates that the wealth-ranking-based majority rule is a strong predictor for finally receiving the subsidy entitlement. The scatter plots confirm that almost all majority-eligible households finally received the subsidy offer. In contrast, those households at the right-hand side of the cutoff, which are not majority-eligible but were nevertheless subsequently drawn into the beneficiary list after consultation, introduce some 'fuzziness' in the subsidy assignment and motivate the use of majority rule as an instrument. As a relevant instrument, being eligible by majority-rule increases the chances for subsidy-eligibility by about 85 percentage points in rural and semi-urban communities alike. In Figure 5.1 I depict this effect for a one-decile-window around the cutoff, while local linear regression estimates in Table

[9] There can be slight deviations between statistical inference from the graphical presentation in comparison to the output tables that come due to the fact that the confidence interval option in Stata does not allow for robust standard errors.

D.1 in the Appendix confirm the robustness across different window sizes.

Pricing and health insurance demand

Figure 5.2 and Table 5.3 set out the ITT effect of being subsidy-eligible by majority rule on households' health insurance take-up at the extensive margin. If not stated otherwise, my result interpretations in this subsection refer to a 15 median-wealth-rank percentile window around the cutoff (h=0.15). I find an economically large effect for semi-urban communities which is statistically significant at the one percent level. Being offered health insurance at a 50 percent discount increases average health insurance demand by a factor of almost 5, from initial 7 percent to almost 34 percent (column 5). According to the left-hand side graph in Figure 5.2, there is no such effect for rural communities. When taking the exact coefficient estimates from Table 5.3, for a 15 and 10 percentile window I find a price elasticity of health insurance demand of −1.3 and −1.5, respectively.[10] Table 5.4 sets out three additional specifications for the semi-urban

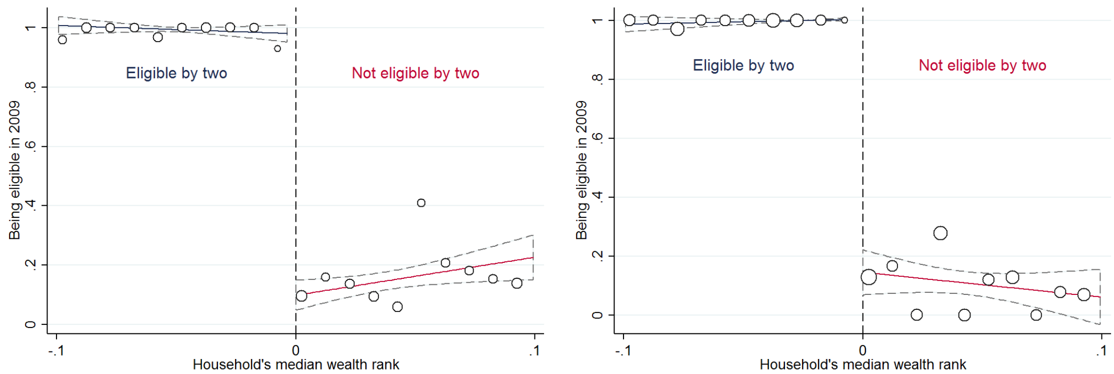

(a) Rural (b) Urban

Figure 5.1: First-stage regression

Notes: Dots represent average outcome values by household median wealth percentile. Dot size indicates the relative number of observations per percentile. The forcing variable on the horizontal axis is based on three community wealth rankings, centered at the beneficiary cutoff, CDF-transformed and bounded between −0.2 and 0.8.

subsample. The Wald estimator in the first panel gives the local average treatment effect for eligible households on the probability to take up health insurance. The LATE (Table 5.4, column 2) is about five percentage points higher than the ITT (Table 5.3, column 5), which reflects a moderate increase of roughly 15 percent. Hence, the ITT provides a reasonable lower bound

[10] I compute the price elasticity as the jump in demand relative to average insurance demand at the cutoff. For a 15 percentile window this gives $\frac{0.268}{0.5\cdot((0.268+0.07)+0.07)} = 1.31$.

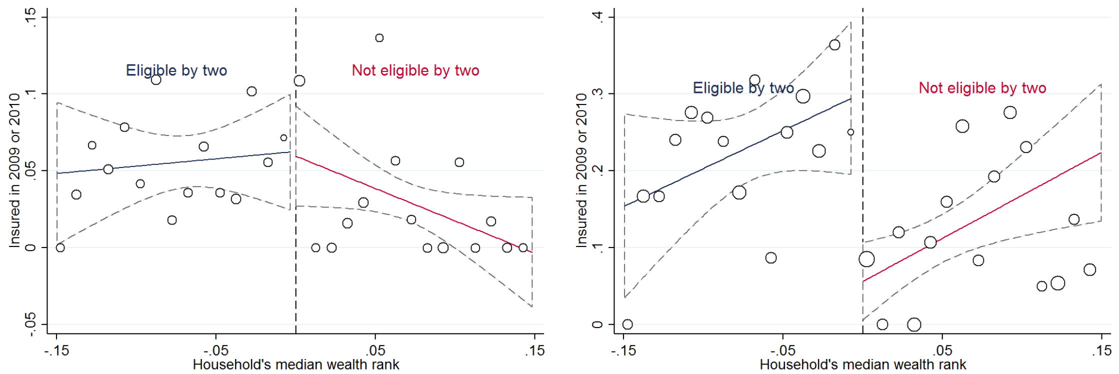

(a) Rural (b) Urban

Figure 5.2: Intent-to-treat effect of majority-eligibility on health insurance enrollment

Notes: Dots represent average outcome values by household median wealth percentile. Dot size indicates the relative number of observations per percentile. The forcing variable on the horizontal axis is based on three community wealth rankings, centered at the beneficiary cutoff, CDF-transformed and bounded between −0.2 and 0.8.

approximation of the economic effect I am finally interested in. Second, I check whether being offered a 50 percent discount affects the household's choice to insure at least one member over a two-years period. Relative to the pooled local linear regression from Table 5.3, I find a slightly smaller increase in take-up of a factor of 3.9 (column 5) reflecting a price elasticity of −1.2. For households that have bought health insurance in 2009, I finally estimate the ITT of the subsidy offer on the decision to extend health insurance coverage into the next year. Not surprisingly, this subsample shows a relatively high average 2010 enrollment rate of about 70 percent in a 20 percentile window around the cutoff. When keeping only households that were already insured in 2009, the sample size decreases considerably and challenges robust identification of the ITT effect. Even for a 20 percentile window around the beneficiary cutoff one observes only 31 households to the cutoff's right, that is an average 1.4 households per community. Given these concerns, I take the ITT estimate of 0.36 in column 7 of Table 5.3 rather as suggestive evidence for a positive effect of subsidy eligibility on the likelihood of extending health insurance into the next year.

So far, I have focused on enrollment effects at the extensive margin, where I considered a household's decision to insure at least one of its members. In addition, Table 5.5 sets out results for enrollment effects at the intensive margin. To test whether the price discount leads to selection into health insurance by certain household member groups, I separately estimate the ITT effect on enrollment for different pairs of complement household subsets and check for significantly different ITT coefficient estimates across both equations. For statistical inference, I consider the

Table 5.3: Intent-to-treat effect of majority-eligibility on enrollment

	Insured in 2009 or 2010					
	Rural (N=5692)			Semi-urban (N=2690)		
	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1
	(1)	(2)	(3)	(4)	(5)	(6)
Majority-eligible in 2009	0.008	0.007	0.018	0.238***	0.268***	0.324***
	(0.027)	(0.032)	(0.040)	(0.064)	(0.074)	(0.094)
Median wealth rank	-0.288*	-0.351	-0.337	0.441	1.043**	1.668*
	(0.156)	(0.248)	(0.485)	(0.280)	(0.489)	(0.854)
Majority-eligible * Median wealth rank	0.320	0.397	0.666	0.842	0.362	0.818
	(0.257)	(0.371)	(0.632)	(0.620)	(0.920)	(1.666)
Constant w./o comm FEs	0.053	0.056	0.059	0.090	0.070	0.053
Mean of insured	0.041	0.041	0.046	0.157	0.168	0.181
Community FEs	YES	YES	YES	YES	YES	YES
Observations (number of households)	2088	1675	1143	1038	809	548
Observations left of cutoff	987	862	603	507	420	292
Observations right of cutoff	1159	871	598	570	428	295

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. The corresponding number of observations is given in the bottom panel, while the header contains the overall number of observations.

corresponding test statistics in conjunction with the subgroups' baseline insurance rates.[11] When, for instance, subgroup A shows both a significantly higher ITT effect and an initially lower or equal insurance rate as subgroup B, I take this as evidence for selection into health insurance by households members of subgroup A. I find evidence for selection by household age groups (columns 7 to 14). Overall, the subsidy offer encourages households to enroll younger household members who are aged below 16 (column 8 vs. 10). Within the group of adults, elderly household members, on average, show twice as high initial enrollment rates as prime aged individuals but significantly lower enrollment effects of the subsidy (column 12 vs 14). Thus, the discounted price offer enhances convergence in health insurance enrollment among these two adult groups and, on average, puts adolescents at an advantage over adults. I do not find evidence for selection within sex (columns 4 to 6) and only suggestive evidence for selection by extended households members (columns 15 to 18).

[11] Otherwise it is hard to tell whether a larger coefficient actually provides evidence for a statistically significantly larger *relative* effect on take-up.

Table 5.4: Different enrollment specifications for semi-urban communities

	Insured in 2009 or 2010			Insured in 2009 and/or 2010			Renewed ins. in 2010		
	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Majority-eligible in 2009				0.263***	0.285***	0.333***	0.364*	0.487**	0.737***
				(0.073)	(0.085)	(0.108)	(0.217)	(0.220)	(0.236)
Finally eligible (instrumented)	0.279***	0.317***	0.383***						
	(0.075)	(0.087)	(0.111)						
Median wealth rank	0.645**	1.327**	2.085**	0.384	0.885	1.339	2.358	4.438**	7.449***
	(0.316)	(0.535)	(0.932)	(0.371)	(0.611)	(1.111)	(2.056)	(2.151)	(2.218)
Eligibility * Wealth rank	0.621	0.047	0.348	1.009	0.535	1.094	-1.020	-3.037	-4.179
	(0.586)	(0.868)	(1.570)	(0.714)	(1.065)	(1.908)	(2.433)	(2.507)	(2.844)
Constant w./o comm FEs	0.054	0.028	0.002	0.118	0.102	0.089	0.727	0.649	0.549
Mean of insured	0.157	0.168	0.181	0.190	0.202	0.216	0.808	0.828	0.867
Community FEs	YES	YES	YES	YES	YES	YES	YES	YES	YES
Observations (number of households)	1038	809	548	526	411	278	78	64	45
Observations left of cutoff	507	420	292	259	215	149	49	42	30
Observations right of cutoff	570	428	295	287	216	149	31	24	17

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. Estimations refer to the subsample of semi-urban households. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff.

Table 5.5: Selection into health insurance in semi-urban communities by household member type

	Being insured in 2009 or 2010																	
			Panel A: Sex				Panel B: Age group								Panel C: Family type			
	All (N=23102)		Male (N=11649)		Female (N=11453)		Non-adults (age < 16) (N=10772)		Adults (age ≥ 16) (N=12330)		Prime age (16 ≤ age ≤ 55) (N=10007)		Elderly (age > 55) (N=2323)		Nuclear family (N=14127)		Extended family (N=8975)	
	h=0.2	h=0.15	h=0.2	h=0.15	h=0.2	h=0.15	h=0.2	h=0.15	h=0.2	h=0.15	h=0.2	h=0.15	h=0.2	h=0.15	h=0.2	h=0.15	h=0.2	h=0.15
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
Majority-eligible in 2009	0.109**	0.135**	0.105**	0.124**	0.109**	0.139**	0.148***	0.181***	0.062	0.079	0.121**	0.150**	0.077	0.096*	0.062	0.084	0.128**	0.145**
	(0.049)	(0.059)	(0.052)	(0.061)	(0.053)	(0.063)	(0.051)	(0.061)	(0.052)	(0.062)	(0.055)	(0.065)	(0.048)	(0.056)	(0.064)	(0.076)	(0.051)	(0.059)
Median wealth rank	0.154	0.529*	0.142	0.484	0.162	0.569*	0.281	0.665**	0.001	0.341	0.114	0.507	0.101	0.451	0.048	0.446	0.154	0.418
	(0.189)	(0.315)	(0.208)	(0.356)	(0.190)	(0.312)	(0.186)	(0.309)	(0.218)	(0.367)	(0.225)	(0.377)	(0.183)	(0.305)	(0.210)	(0.311)	(0.212)	(0.368)
Majority-eligible * Median wealth rank	0.273	0.081	0.291	0.054	0.234	0.040	0.442	0.382	0.034	-0.289	0.292	0.117	0.154	-0.124	-0.167	-0.504	0.524	0.339
	(0.453)	(0.670)	(0.479)	(0.678)	(0.484)	(0.739)	(0.454)	(0.682)	(0.497)	(0.724)	(0.510)	(0.758)	(0.442)	(0.652)	(0.561)	(0.858)	(0.466)	(0.697)
Constant	0.012	0.001	0.012	-0.005	0.057	0.038	0.039	-0.022	0.084	-0.009	0.050	0.031	0.003	-0.019	-0.025	-0.022	0.059	0.064
	(0.031)	(0.029)	(0.051)	(0.061)	(0.067)	(0.072)	(0.059)	(0.028)	(0.054)	(0.048)	(0.067)	(0.079)	(0.036)	(0.038)	(0.017)	(0.013)	(0.079)	(0.066)
H0: Equal coefficient			0.942	0.818			0.024	0.012			0.003	0.000			0.033	0.325		
Mean of insured	0.094	0.099	0.088	0.093	0.100	0.105	0.095	0.102	0.093	0.096	0.079	0.082	0.143	0.144	0.089	0.096	0.103	0.104
Community FEs	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Number of individuals	7392	5830	3697	2942	3695	2888	4056	3200	3336	2630	4231	3352	6497	5108	2147	1681	4614	3619

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the household level in parentheses. Estimations refer to the subsample of semi-urban households. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. H0 reports the p-values from a t-test for equal coefficients across a pair of estimation equations: for instance, the two values 0.942 and 0.818 in Panel A indicate that ITT effects for male and female household members are not statistically different from each other at conventional levels for h=0.20 and h=0.15, respectively. The following household member pairs add up to the total number of individuals in the first column: male and female, non-adults and adults, nuclear and extended family members. Prime age and elderly household members add up to the total number of adult household (columns 9 and 10) members .

Drawing Engel Curves: Health insurance demand and wealth

In line with existing studies on the determinants of health insurance demand, my descriptive results indicate a strong positive association between household wealth and insurance take-up (see Figure 2.4). In addition, Figure 5.3 shows that the subsidy effect puts the demand of semi-urban poor households on par with the wealthiest households in the study area. To test for heterogeneous effects of health insurance pricing along household wealth, I exploit wealth differences across the 22 semi-urban communities. Since the participative wealth rankings cannot be compared across communities, I use census survey information on households' dwelling characteristics and asset possessions to identify a household's position in the area's asset wealth distribution. Specifically, my community wealth variable w_c gives the wealth percentile of the 'cutoff household' living in community c. Given across-community variation in average wealth, I expect 'cutoff households' to show different wealth percentiles across communities.

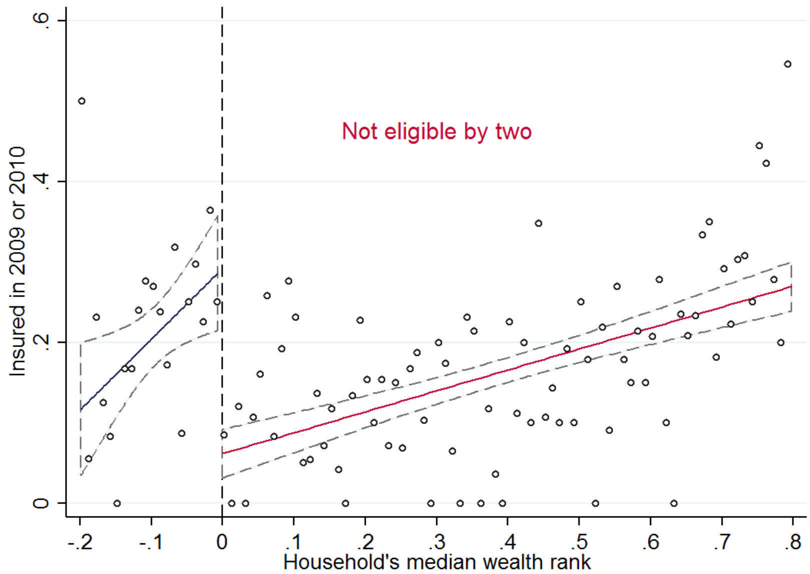

Figure 5.3: Take-up effect along the whole distribution of semi-urban households

Notes: Dots represent average outcome values by household median wealth percentile. Dot size indicates the relative number of observations per percentile. The forcing variable on the horizontal axis is centered at the beneficiary cutoff, CDF-transformed and thus bounded between -0.2 and 0.8.

The two bottom rows in Table D.2 in the Appendix indicate that most of semi-urban 'cutoff households' are positioned within the range of the 20 and 35 percentile of the semi-urban area's household wealth distribution. For this wealth spectrum the results suggest a substantial wealth effect on health insurance demand. Moving up one decile in the household wealth distribution is associated with an increase in demand by 4.15 percentage points (coefficient $\hat{\lambda}$ in equation 5.2). On the other hand, I cannot reject the hypothesis that differently wealthy households respond differently to health insurance pricing (coefficient $\hat{\pi}$ in equation 5.2). The relatively large increase in the standard errors of my estimated interaction effect (when compared to the linear effect), however, suggests a power problem (due to only limited variation of my wealth variable at the community level) to be a likely reason for this non-result.[12]

The results in Table D.2 motivate an Engel Curve illustration. Figure 5.4 depicts the relationship between health insurance demand and asset wealth for the semi-urban subpopulation. The gray dashed line shows the estimated (and smoothed) density in household eligibility status by wealth percentile. It indicates that only a tiny fraction of the wealthiest 40 percent households has received the subsidy entitlement. I regard this fraction as sufficiently low to interpret the corresponding fitted quadratic regression line as a reasonable Engel Curve approximation at high prices. For the poorest sixty percent, on the contrary, I observe health insurance demand for both subsidy-eligible and non-eligible households. For the subset of households residing within the range of percentile 20 to 35, I can exploit across-community variation in wealth (green hollow dots) to disentangle the choices made under both price scenarios.[13] Consequently, the blue line in the lower part reflects the strong wealth correlation of 0.415 absent the subsidy. Extrapolation indicates that households at the 20 percentile increase their demand by a factor of 1.4 when moving up to the 40 wealth percentile. For households around the 60 and 80 wealth percentile, a one-quantile upward jump in wealth is associated with a smaller increase in take-up of 67 and 90 percent, respectively. The linear pricing effect in Table D.2 (coefficient $\hat{\beta}$ in equation 5.2) is reflected by the overall upward shift of the red line. The slope of this upper line is not estimated precisely, which I take as evidence for similar levels of price elasticity along the wealth distribution.

Adverse selection

Table 5.6 sets out results from the adverse selection test. For this I have estimated local linear regression equation 5.1 with the subsample of ever-insured households between 2009 and 2010. This restriction trims the sample down to about 400 and 450 households in rural and semi-urban

[12] To illustrate the interpretation of my coefficients in equation 5.2, let us assume a statistically significant interaction effect. For a 15 percentile window, the latter then would roughly indicate a 50 percent decrease in the price elasticity of health insurance demand for moving up one wealth percentile.

[13] The extrapolation of both curves in this range of the distribution is based on the point I observe for the median community and the corresponding slopes that are provided by the coefficients estimates in Table D.2.

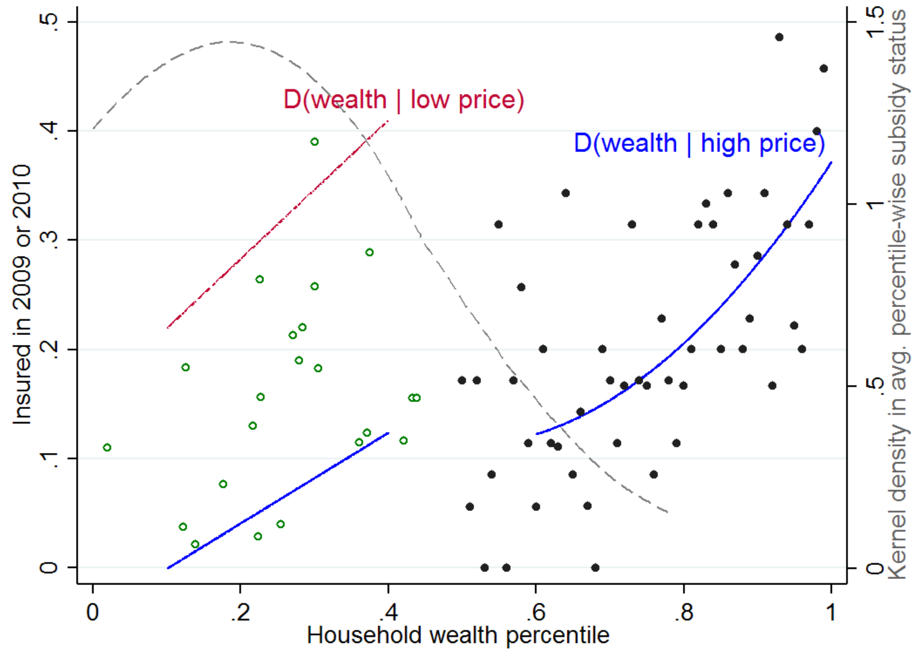

Figure 5.4: Extrapolation of Engel Curves for semi-urban households

Notes: The horizontal axis depicts household wealth percentiles based on an asset index. Hollow dots indicate average enrollment rates by community asset-wealth percentile at cutoff. Filled dots indicate average enrollment rates by asset-wealth percentile. The grey dashed line is the density of average percentile-wise subsidy eligibility, obtained with kernel density estimation that applies a bandwidth of 0.2 (see right-hand side vertical axis). The extrapolation of Engel Curves is based on a second-order polynomial fit for the relationship between asset-based household wealth and insurance status (right-hand side curve: D(wealth|high price)) and on regression results in Table D.2 (left-hand side curves).

communities, respectively, from which about one-fourth are located within the 20 percentile window around the cutoff. The first two panels reflect the extensive margin of adverse selection by considering the total household per capita number of health care facility visits as an outcome variable. For both rural and semi-urban communities my ITT estimate is not statistically significant at conventional levels for any window size such that my test cannot reject the hypothesis that there is no self-selection based on risk. For rural communities I observe negative coefficients which vary considerably with window size. In contrast, semi-urban communities show positive coefficients which are more precisely estimated and less sensitive to window choice. Nevertheless, my empirical design suffers from a power problem as it would require true ITTs of size one to obtain precisely estimated coefficients. The latter would reflect an unrealistically large effect, namely a threefold increase in average health care visits due to being offered a 50 percent price discount. Hence, coefficient estimates for the 20 and 15 percentile window only provide suggestive evidence that a lower price draws in higher risk in semi-urban communities.[14] Being offered health insurance at a discounted price more than doubles the average number of per capita health care visits from 0.53 to 1.32, for a 20 percentile window. The third panel in Table 5.6 considers ITT effects at the intensive margin of health care utilization in semi-urban communities. The corresponding outcome variable is the average annual per capita amount in payouts received by an insured household through drug prescriptions. With a threefold increase in payouts the coefficient estimates suggest an even stronger extent of selection at the intensive margin, even though not statistically significant at conventional levels (see also Figure D.1 in the Appendix).

Table 5.6: 'Price test' for adverse selection into health insurance

	Number of avg. annual per capita visits						Avg. annual p.c. payout amount		
	Rural (N=404)			Semi-urban (N=447)			Semi-urban (N=447)		
	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Majority-eligible in 2009	-0.123	-0.464	-0.177	0.789	0.765	1.152	843.741	900.153	1454.769*
	(0.727)	(0.848)	(1.036)	(0.493)	(0.601)	(0.758)	(563.674)	(683.161)	(824.022)
Median wealth rank	-8.886	-12.695	-11.856	2.692	3.452	5.484	2091.727	3250.665	4214.724
	(6.187)	(9.731)	(15.158)	(2.725)	(4.277)	(6.608)	(2460.224)	(3394.086)	(3896.860)
Majority-eligible * Median wealth rank	9.480	12.663	25.851**	-7.490	-9.026	-1.720	-4584.964	-5326.087	9567.475
	(6.643)	(8.929)	(12.064)	(7.402)	(9.146)	(9.743)	(6979.630)	(9552.132)	(13957.447)
Constant w./o community FEs	1.326	1.382	1.337	0.531	0.460	0.305	459.454	407.948	311.821
Mean of outcome variable	1.138	1.094	0.998	1.089	1.083	0.900	1006.163	1021.420	881.529
Community FEs	YES	YES	YES	YES	YES	YES	YES	YES	YES
Observations (number of households)	88	71	55	142	112	79	174	136	92
Observations left of cutoff	51	43	32	87	71	51	104	85	59
Observations right of cutoff	40	31	26	60	46	33	78	59	41

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. The outcome is health care facility utilization from insured households at the extensive (number of average annual per capita visits) and intensive (Average annual per capita health insurance payouts in terms of prescribed medicine) margin.

5.4 Robustness Checks

Internal validity

My main identifying assumption is that the outcome variable absent the intervention must be continuous in the household median wealth rank. Then the variation in the treatment in a neighborhood of the threshold is 'as good as randomized' (Angrist and Pischke, 2009). For continuity to hold, individuals must neither be able to manipulate the forcing variable, nor to precisely sort

[14] Given the small number of observations, I am even less confident about the results for the very small window.

around the discontinuity threshold (Lee and Lemieux, 2010).

The following three features make sorting very unlikely. First, the final targeting choice is based on community-specific multidimensional targeting criteria and an aggregation rule that relies on three individual wealth assessments. For a household, the nature and timing (targeting criteria were revealed just before the wealth rankings took place) of this stepwise procedure make it extremely difficult to correctly understand and anticipate the community's welfare function. Second, wealth assessments rely on local knowledge and not on one-shot survey questions. The former is likely to cover a longer period and should be less prone to short-term strategic signaling by households. Third, the targeting decision followed a relative concept of poverty and facilitators did not disclose the beneficiary cutoff ranks beforehand. For successful sorting households thus would have to master three challenges at once: (i) directly anticipating the community's welfare function, (ii) correctly assessing their own relative position within the household wealth distribution, and (iii) shifting their position below an undisclosed beneficiary cutoff through successful strategic signaling of deprivation. This makes it very unlikely that households were able to precisely sort around the cutoff.

A more credible thread to the RDD's validity might arise with elite capture when representatives manipulate the targeting procedure to favor themselves or crony households. As discussed in section 4.1, such a behavior could be both intrinsically (representatives dislike for non-cronies) or extrinsically (reciprocal relationships with cronies) motivated and it could apply for different types of cronies such as friends, family members or households from the same ethnic group. I argue that extrinsically motivated elite capture is unlikely to occur at all. To demand preferential treatment, cronies have to convince at least two representatives at the same time. In addition, cronies have to act very foresighted as physical isolation of the representatives during the ranking process prevents direct collusion. Even if a household colluded successfully, the representatives' ability to produce a preferential ranking would be hampered by the fact that they did neither know the cutoff rank nor the wealth rankings from the two other representatives. The last argument equally applies to a scenario where representatives have an intrinsic interest in favoring crony households. While my analysis in chapter 4 provides evidence for the existence of a moderate allocation bias due to ethnic favoritism, the features of the underlying targeting design make such a bias unlikely to disproportionally occur in the cutoff's neighborhood.

Placebo and balancing tests

First, I run a placebo test by estimating my main local linear regression equation with enrollment prior the subsidy intervention as a dependent variable. Specifically, I estimate the ITT of being majority-eligible for the subsidy in 2009 on a household's insurance status in 2007 and 2008. Based on the first targeting round in 2007, households in these two years already got assigned a 'baseline' eligibility status. Even though the 2009 targeting update followed the same design, I expect substantially different community wealth rankings due to differently composed targeting committees and changes in household wealth over time. If household characteristics were truly continuous in the household median wealth rank, subsidy eligibility must not affect the decision to buy health insurance prior to 2009. That is exactly what I observe in Figure 5.5. In contrast to my main result in Figure 5.2 on page 92, the placebo test reveals very imprecisely estimated negative coefficient estimates with a magnitude of about 3 percentage points for rural and semi-urban communities alike (see also Table D.3 in the Appendix). As a second robustness check,

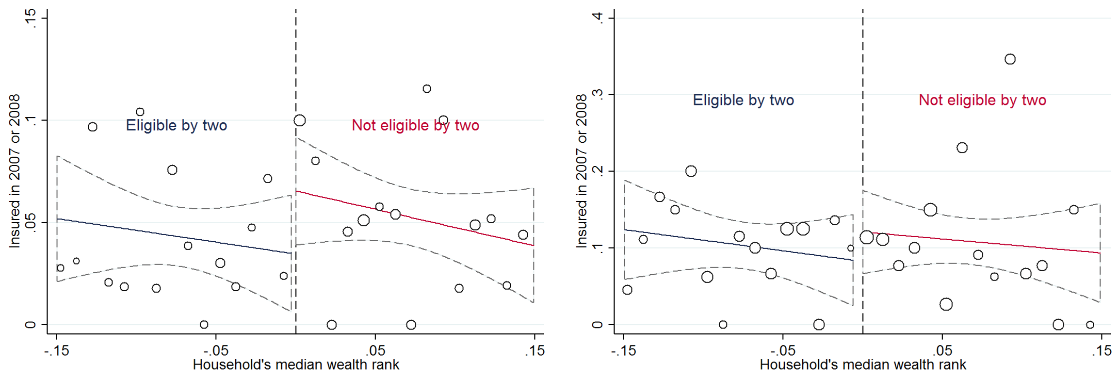

(a) Rural (b) Urban

Figure 5.5: Subsidy-eligibility in 2009 and baseline enrollment in 2007 to 2008

Notes: Dots represent average outcome values by household median wealth percentile. Dot size indicates the relative number of observations per percentile. The forcing variable on the horizontal axis is based on three community wealth rankings, centered at the beneficiary cutoff, CDF-transformed and bounded between −0.2 and 0.8.

I test whether subsidy-eligibility in 2009 has an effect on household outcomes from the same year which should not be affected (McCrary, 2008). This is similar to a test whether households are balanced across treatment and control groups within an experimental setting. My outcome variables come from the vital registration survey data and comprise two variables on each of the following dimensions; demographics, ethno-linguistic information and occupational choice. Figure

5.6 shows the corresponding local linear regression results for the six different outcomes. It is reassuring and confirmed by regression results in Table D.4 in the Appendix that no outcome variable shows a statistically significant jump at the beneficiary cutoff.[15] Taken together, results from my two tests in combination with the conceptual consideration from the former subsection make me very confident that my discontinuity design is sufficiently valid to estimate the causal effect of subsidy-eligibility on the outcomes of interest.

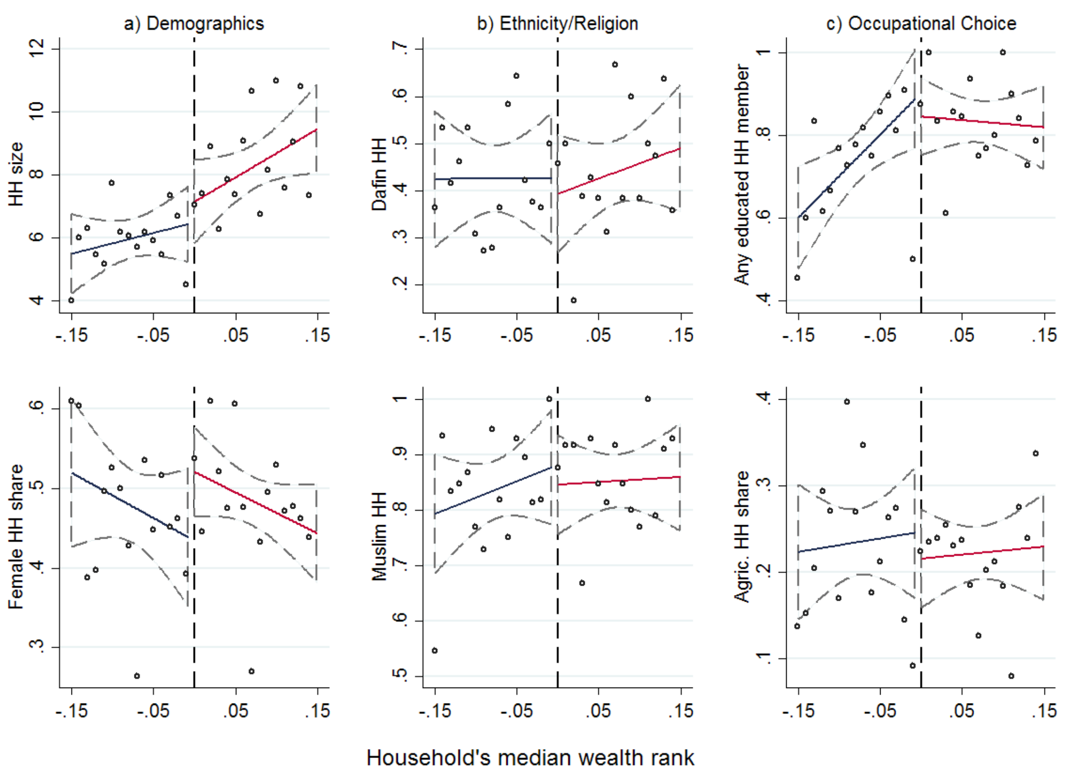

Figure 5.6: Balancing tests for six variables in semi-urban communities in 2009

Notes: Dots represent average outcome values by household median wealth percentile. Dot size indicates the relative number of observations per percentile. The forcing variable on the horizontal axis is based on three community wealth rankings, centered at the beneficiary cutoff, CDF-transformed and bounded between −0.2 and 0.8.

[15] As a robustness check for individual manipulation Lee and Lemieux (2010) further propose to examine the density of the forcing variable in order to check for a suspicious high density on the eligibility-side of the cutoff. Yet, the targeting design always selects a predetermined set of beneficiary households, namely the poorest quintile in each community, such that the design does not allow for such a robustness check.

Local polynomial regression

As a final robustness test, I check how sensitive my results respond to a change in the underlying functional form. Specifically, I estimate the following local polynomial regression model with distinct second-order polynomials to the left and the right of the cutoff

$$\begin{aligned} y_{ch} =& \alpha_k + \beta \cdot \mathbb{1}\{EligByMajority\}_{ch} \\ &+ \gamma \cdot MedianRank_{ch} + \delta \cdot \mathbb{1}\{EligByMajority\}_{ch} * MedianRank_{ch} \\ &+ \phi \cdot MedianRank_{ch}^2 + \lambda \cdot \mathbb{1}\{EligByMajority\}_{ch} * MedianRank_{ch}^2 + u_{ch}, \end{aligned} \tag{5.3}$$

where y_{ch} equals one if any member of household h living in community c has bought health insurance in the current year and where the quadratic term allows for a non-linear relationship between health insurance take-up and the household median wealth rank. The local polynomial regressions in Figure D.2 in the Appendix provide very similar results than the linear specification. While there is still no detectable jump at the rural beneficiary cutoff, I observe a ITT effect in semi-urban communities which is statistically significant and of similar size than in Figure 5.2. Table D.5 in the Appendix reveals that ITT estimates from both my linear and quadratic specification converge to a very similar value when trimming down the window size. Given that the non-linear model provides a relatively good fit of the underlying data generating process, this is what one expects.

5.5 Discussion

I will discuss my four main findings in the following. First, for a relatively poor sub-Saharan African country my main finding is that health insurance pricing matters and can have substantial effects on take-up. Specifically, for semi-urban households that are placed around the 20 percentile wealth threshold the 50 percent subsidy offer, on average, increases take-up by a factor of five. While providing novel evidence for the sub-Saharan African context my results nicely complement the existing literature on health insurance demand and pricing from Asian and Latin-American countries. To my knowledge, there are only three very recent studies that evaluate health insurance demand under an exogenously introduced new price that is not zero. For two rural districts in Vietnam, Wagstaff et al. (2016) evaluate an intervention that aims to encourage enrollment with a 25 percent premium subsidy and, overall, find no evidence for measurable take-up effects. Capuno et al. (2016) evaluate enrollment effects for a health insurance intervention in Vietnam where a premium subsidy is combined with an extra information leaflet. Very similar to the setting considered here, their subsidy amounts to a 50 percent reduction and initial average take-up is below ten percent. Nevertheless, their effect on take-up is modest when compared to ours, amounting to a 37 percent increase from 8.4 to 11.4 percent. Fischer et al. (2016) offer social health insurance in Pakistan at four randomized prices and find price elasticities in the range of -0.6 and -1.6. Finally, two studies consider a six-months offer of a fully subsidized health insurance in Cambodia (Levine et al., 2016) and Nicaragua (Thornton et al., 2010) and observe very different responses in uptake.[16]

A second important set of findings addresses the heterogeneity in enrollment effects. Results substantially differ between semi-urban and rural households around the beneficiary cutoff which can be regarded as poor and ultra-poor by national standards, respectively. While the former show a very price elastic health insurance demand, the subsidy is little effective for rural households. This is in line with the study by Capuno et al. (2016) which finds larger enrollment effects in urban relative to rural areas in the Philippines. Motivated by the question whether rather remoteness or poverty explains the lack in demand by rural households, I took a closer look at the relationship between insurance take-up and household wealth. When extrapolating Engel Curves for health insurance demand among semi-urban households I find a very strong relationship between enrollment and wealth. In addition, the 50 percent discount puts poor households in my study area on pair with the wealthiest households, while price elasticity does not seem to differ across household wealth groups. To the best of my knowledge, only Thornton et al. (2010) rigorously evaluate heterogeneous enrollment effects with respect to wealth and do not find significantly different patterns. With respect to other outcomes, wealthy households have been found to show a relatively strong increase in health care utilization (Wagstaff et al., 2009 in China; Kuuire et al., 2016 in Ghana) and more pronounced reductions in health care expenditures (Bernal et al., 2016 in Peru; Woldemichael et al., 2016 in Rwanda) when being offered social health insurance. In addition, household wealth has been also found an important predictor for rainfall insurance demand in rural India (Giné et al., 2008; Cole et al., 2013)

[16] For Cambodia, Levine et al. (2016) find a sixfold increase in demand from seven to about 50 percent, leading to a price elasticity that is very close to the one I have found in my study. As control households were simultaneously offered one month of fully subsidized health insurance, their effective six-months subsidy amounts to a 80 percent increase in the number of fully subsidized months. Thornton et al. (2010) consider the subpopulation of Nicaraguan informal workers who started at a fairly high initial enrollment rate of twenty percent and increase take-up only modestly, by a factor of 1.65.

Third, by exploiting the fact that health insurance purchase effectively took place at the individual level, I can evaluate the intent-to-treat effect of pricing on the intra-household allocation of health insurance coverage. Adolescents show similar insurance coverage than adults at pre-intervention levels and I find that the subsidy offer puts them at an advantage. Within the group of adults, a lower premium leads to favoritism of prime aged over elderly household members. These findings complement the small and recent literature that analyses the intra-household allocation of health insurance in low-income countries. The main finding from this literature is that household heads show a preference for insuring themselves and their spouses in comparison to other household members (Govender et al., 2014 in South Africa; Panda et al., 2014 in India). In addition, such a disaggregated analysis helps to better understand and reduce circumstances leading to unequally distributed health inputs within households. Prominent studies in this area deal with unequal nutritional intake (Rosenzweig and Schultz, 1982; Pitt et al., 1990) and the effect of negative income shocks on within-household mortality rates (Rose, 1999).

Fourth, my test for adverse selection rejects the hypothesis that a higher price draws in higher risk but, if anything, suggests the opposite. This finding confirms two former analyses from the present study area (De Allegri et al., 2006; Parmar et al., 2012) and differs in comparison to the two existing studies that run an adverse selection 'price test' in a low-income setting.[17] The first study by Fischer et al. (2016) on social health insurance in Pakistan finds evidence for substantial adverse selection, which increases in the premium amount but can be abated by bundling enrollment at the household level. Their latter finding might offer a partial explanation for the different result I observe in my study, where the insurer clearly encouraged enrollment by households and not individuals (Fink et al., 2013). Within a health insurance that makes enrollment at the household level mandatory, Polimeni and Levine (2011) find evidence for substantial adverse selection. The five-months-offer of free health insurance might led to considerable selection of low-risk households when compared to the 50 percent subsidy from my study. While context-specific differences will always complicate comparisons across studies, I conclude that more research of this type is needed to better link the occurrence of adverse selection with the underlying health insurance design.[18]

[17] De Allegri et al. (2006) and Parmar et al. (2012) employ household sample survey rounds between 2004 and 2007 to estimate the relationship between a household's current enrollment and health status and find weak or no statistically significant estimates, respectively.

[18] Clearly, evidence from 'ex-ante correlation' studies should be carefully included into the picture. Polimeni and Levine (2011), show that their 'ex-ante correlation' test for adverse selection gives a much lower bound for its existence than the 'price test'. Overall, this branch of studies provides evidence against (Dror et al., 2005; Nguyen and Knowles, 2010; Banerjee et al., 2014; Capuno et al., 2016) and in favor (Wagstaff et al., 2009; Wang et al., 2009; Wagstaff et al., 2016) of adverse selection into social health insurance.

Chapter 6

Conclusion

Social protection programs are becoming increasingly popular around the globe and most governments in low-income countries regard them as a crucial tool to achieve sustainable poverty reduction (Coady et al., 2004a; World Bank, 2014). In practice, however, many programs still fail to cover large parts of their poor populations, which can be attributed to both supply and demand side obstacles. Due to inaccurate targeting, on the supply side, governments often miss large parts of the intended target group when offering program benefits. Even if targeting is successful, on the demand side, entitled households are often reluctant to enroll into social protection programs. This state of development is especially applicable to the context of sub-Saharan Africa. In international terms, the region has experienced the most rapid expansion of anti-poverty programs over the last five years, but also performs worse in covering poor households (Honorati et al., 2015). Programs are often implemented in an ad-hoc fashion (Monchuk, 2013) and limited government capacity severely undermines effective targeting (Del Ninno and Mills, 2015). Furthermore, it has proven especially challenging to encourage enrollment among the large parts of rural and remotely located people in this region (World Bank, 2014).

My dissertation has evaluated a social protection scheme that was implemented in Burkina Faso, one of the poorest countries in sub-Saharan Africa. Within this context, I have analyzed an intervention that targeted poor households by offering them health insurance access at a reduced price. Detailed information on the underlying targeting procedure, which I merged with several micro-level datasets, allowed me to analyze both the intervention's effectiveness in encouraging enrollment and the performance of the underlying targeting method. The program evaluated here applied community-based targeting, a commonly employed approach that delegates authority of the targeting choice to the community (Alatas et al., 2012). Within my dissertation I have addressed three different research questions in the realm of social policy implementation. First, I have asked whether community-based or statistical targeting is more accurate in targeting consumption poor households and whether the picture changes in consideration of the programs' targeting costs. Since there exists a well-established concern about the threat of elite capture in decentralized programs, in the second part of my dissertation I have addressed the question whether the local targeting committees in the present context have favored ethnically connected households. The last part concerns program uptake by poor households and asks whether pricing is an effective instrument to encourage households' demand of voluntary micro-health insurance. While I have summarized and discussed the results in the respective chapters, I shall close this work with some farther reaching concluding remarks.

In line with existing studies, my comparative targeting accuracy analysis finds that program coverage might vary significantly with the underlying targeting method. Ceteris paribus, a higher targeting accuracy implies less benefits 'leaking' to non-poor households and more benefits that are available for the poor (Coady et al., 2004a). Nevertheless, my data confirm the general perception that targeting program costs also vary substantially across methods. "Targeting is ultimately meant to be a cost-control exercise" (Del Ninno and Mills, 2015, p.12) and a comprehensive cost-benefit comparison requires targeting program costs. This study provides the first such assessment and confirms common practice from the field; programs that have limited capacity or distribute small benefits find it far more cost-effective to employ community-based targeting exercises. The finding that econometric targeting is excessively expensive when census data are not readily available further motivates the design of hybrid targeting approaches.[1] Given that a general census is typically carried out no more often than every ten years, targeting based on census data will become less accurate the more outdated the underlying data. Community-targeting exercises, on the other hand, may be repeated on a revolving basis at a moderate cost and, in this way, keep track of poverty transitions of households over time. This argument further suggests that revolving community-based targeting might be particularly suited for quickly-evolving environments. In the present study area, for example, the community-based targeting exercise has been carried out

[1] Hybrid targeting approaches that combine community-based and statistical targeting exercises are already common practice (Coady et al., 2004a). Alatas et al. (2012), Karlan and Thuysbaert (2016), and Stoeffler et al. (2016) analyze a hybrid CBT where communities predetermine a household target set that is subsequently verified and refined through the application of survey-based information. Nevertheless, these hybrid methods are relatively expensive (Alatas et al., 2012) and so far have been not motivated on cost-benefit grounds.

three times between 2007 and 2011.

In addition, it has been argued that program cost might increase considerably when "local elites dominate and corrupt community-level planning and governance" (Dasgupta and Beard, 2007, p.230). Local elites might directly embezzle program benefits at the implementation stage (Ferraz and Finan, 2011; Beekman et al., 2014) or, instead, manipulate and bias the allocation choice at the targeting stage (Alatas et al., 2012). For the local context of my study, I have found evidence for the latter, reflected by local targeting committees which favor households from their ethnic group. The effects reported here are in line with recent studies on alternative elite capture outcomes and, economically, they appear modest in comparison to results from the targeting accuracy literature. Specifically, my exclusion error due to ethnic favoritism of not more than ten percent compares favorably to the potential increases in exclusion errors induced by a 'wrong' targeting method choice. This is reflected by the exclusion error range of 25 to 60 and 20 to 75 percent found within the comparative targeting accuracy literature (Alatas et al., 2012; Stoeffler et al., 2016) and this study, respectively. It suggests that policymakers should be less concerned about elite capture when making a choice between decentralized and statistical targeting to maximize program coverage.

Furthermore, my test for ethnic favoritism provides a motivation for future research on the role of institutions. Burgess et al. (2015) analyze distributional policies at the central level and show that the existence of democratic institutions can remarkably reduce the extent of ethnic favoritism. At the local level, my review of the existing literature cautiously suggests that elite capture rather takes place within constituted local governments when compared to participatory programs. According to Mansuri and Rao (2013), the latter often explicitly incorporate checks and balances that limit the potential for capture by political leaders. The targeting design I have considered, for instance, involves three such features. First, publicly determined wealth criteria and a clearly communicated mandate for targeting likely reduced the representatives' incentive for selfish decision making. Second, the application of a straight-forward rule to aggregate three independent wealth rankings and the public disclosure of beneficiary lists at the end of each exercise improved the procedure's transparency. Finally, the targeting design restricted the room for within-committee collusion by enforcing the independent creation and aggregation of three wealth rankings. However, in order to subsequently include or dismiss households, the committee was given some room for collusion to decide on a minor set of beneficiaries. A comparison of estimates under both targeting regimes suggests differences in the targeting outcomes and motivates further research on the effects of institutional change on elite capture at the local level.[2]

This study also provides two insights that highlight the important role of institutional design within ethno-linguistically diverse environments. First, for a program that delegates the targeting choice to a small set of representatives, I find that ethnic favoritism can lead to a situation of 'ethnic minority discrimination'. In practice, community-driven welfare programs differ in their agency structure (see Table 3.1), suggesting the existence of a well-known public choice trade-off between the number of representatives and institutional cost (Auriol and Gary-Bobo, 2012). More representatives assure minority protection, but can also complicate decision making. These institutional costs were also considered in the design of the present targeting intervention, as the following quotation shows:

> (A) larger group could be difficult to control with a high risk for participants to become shy, leading to the creation of subgroups and to the deterioration of the conversation (Savadogo et al., 2015, p.3).

Comparing my results with Alatas et al. (2012) suggests the existence of such a trade-off for decentralized welfare programs. Their community-based targeting intervention is based on collective choice within the community and thereby ensures a maximum degree of representation. As a result, and in contrast to my finding, communities produce an allocation that actually favors ethnic minorities. Nevertheless, on the cost side they find that community-members become tired and significantly less accurate in their wealth assessments the longer the targeting exercise proceeds. Second, the finding that ethnic favoritism comes from ethnic diverse communities adds to a broader debate, reflected by two popular arguments from the sociology literature. Allport's (1954) contact hypothesis states that more diverse communities can achieve a permanent reduction of prejudice and conflict between different groups. In contrast, Blumer's (1958) conflict theory takes the view that more interaction with individuals of other groups is costly and generates greater antagonism. I take my results as suggestive evidence in favor of conflict theory even though I can only conjecture about the underlying channels. Further analysis could investigate the role of competition, for instance, by testing whether ethnic favoritism rather takes place in environments

[2] While, so far, the (local) elite capture literature does not provide a rigorous distinction of elite capture across institutional settings, two studies examined the role of the decision maker type. For a collective choice design among villagers, Alatas et al. (2012) find no significant differences in targeting outcomes across different randomly induced group compositions in Indonesia. For the same context, Alatas et al. (2013a) rigorously distinguish between targeting by informal and formal leaders and find only negligible differences.

where people follow competitive (e.g. service sector) instead of complementary (e.g. farming) working relationships.

In addition to cost-benefit concerns, the choice of a targeting approach likely affects successful program implementation in two more indirect ways. First, it may be called into question whether consumption should be the sole targeting objective. Alternatively, there might be considerable value added to the targeting process when communities' concepts of poverty are taken into account. Recent empirical evidence on communities' poverty perceptions shows that communities consider more dimensions than only consumption (Van Campenhout, 2007; Alatas et al., 2012). Kebede (2009) shows that poverty perceptions reflect local circumstances and Alderman (2002) finds that community assessments put more weight on chronic poverty. For the context considered in this study, an ongoing project shows that the community's poverty concept is multidimensional and puts relatively much weight on asset wealth and demographics when compared to mortality and education.[3] In addition, when considering the wealth criteria defined by the communities in the present targeting exercise, it is striking that communities define most of the criteria in terms of capabilities such as 'has insufficient food', 'has nothing' or 'is not able to solve problems by himself' (Savadogo et al., 2015). This fits well into Amartya Sen's capability approach (Sen, 1988) and supports the view that communities consider consumption rather as a means to an end. In this perspective, community-based targeting appears to be well-suited for translating deprivations in the space of capabilities into targeting outcomes.

Second, satisfaction with an implemented targeting method likely affects public acceptance of the underlying program and, thus, its overall success. In this regard, the participative procedure of community-based targeting has the potential to produce some benefits of itself. Since the inception of participatory appraisals, local control over the targeting process has been viewed as a desirable attribute of CBT, powerful enough to increase ownership and awareness, and foster institutional change (Chambers, 1994b). This view is supported by empirical evidence, which shows remarkably high approval rates by communities for decentralized targeting methods (Alatas et al., 2012; Robertson et al., 2014; Schüring, 2014). Savadogo (2017) confirms this picture for the community-based targeting intervention I consider in my dissertation. In a representative sample of 115 households, he finds that more than 85 percent approve of the targeting method.

Hence, it is possible that the relatively large enrollment effect I have found in my price intervention is also partly driven by the participative nature of the underlying targeting design. While

[3]Results can be made available on request.

my study design does not allow for an assessment of the latter relationship, it provides the novel finding that health insurance pricing matters and substantially influences take-up in sub-Saharan African countries. I also find suggestive evidence that the subsidy successfully encourages poor households to extend their membership into the next year. This is in line with insights from a 2006 household survey in our study area, where affordability was the most-stated reason for terminating the health insurance contract (Dong et al., 2009). While short-term encouragement is successful for at least some parts of the population, my study cannot address the common concern that subsidies may undermine individuals' willingness to pay an actuarially fair premium in the long-term. As another limitation, my results primarily speak to the subpopulation of households around the beneficiary cutoff and not to the whole population. Nevertheless, by exploiting heterogeneity in community wealth, I am able to provide some additional policy lessons. In the present context, progressive subsidy schedules are necessary to facilitate access to health care through pay-for insurance schemes. For ultra-poor households, only free health insurance seems to be a viable option. This latter finding confirms a recent trend in targeting poor households for fully subsidized social health insurance in sub-Saharan Africa (Mensah et al., 2010), Latin America (Sosa-Rubi et al., 2009; Miller et al., 2013; Bernal et al., 2016) and East Asia (Bauhoff et al., 2011).

My pricing analysis deliberately considers the role of liquidity constraints in explaining low enrollment into social protection schemes. Nevertheless, additional non-income factors contribute to low insurance demand in low-income countries, including a lack of trust in formal institutions and financial illiteracy (Eling et al., 2014). The availability of administrative health insurance records over the years 2007 to 2011 and a Geographic Information System at the household-level allows for the analysis of the two latter factors in the context considered here. Specifically, it is the subject of an ongoing project to test whether a household's current purchasing decision is affected by health insurance experience in its neighborhood. Following Karlan et al. (2014), who analyze an indexed rainfall insurance in Ghana, experience is proxied with the average annual payouts a household received from the health insurance in the former year.

Finally, an individual's decision to purchase health insurance depends not least on the perceived quality of services one finds in the respective health care facilities. Poor health care provision considerably reduces the value of and, thus, the willingness to pay for any health insurance scheme, even if optimally designed. Anecdotal evidence suggests that poor health service quality might be a considerable explanation for low health insurance demand in the present study area. According to Dong's (2009) survey appraisal, the second and third most often cited reason for not extending health insurance into the next year are dissatisfaction with the staff behavior and the services received in health care facilities, respectively. In line with this perception, the most recent large-scale health intervention in Burkina Faso aims towards the improvement of health care provision through a financial incentive scheme, commonly known as performance-based financing (PBF).[4] With financial and technical assistance from the World Bank, in 2014, the Burkinabé Ministry of Health rolled-out a pilot PBF that covers over one-third of the country's districts, including the study area of this dissertation. Performance indicators focus on maternal and child health treatments while a quasi-experimental intervention design allows for a comprehensive impact evaluation. Clearly, the PBF's effects on health care provision outcomes are of primary interest to the initiators. Nevertheless, a side-analysis should investigate the potential of such supply-side incentives for the expansion of health insurance coverage in the department of Nouna.

[4]In recent years, performance-based financing (PBF) has become a popular measure to improve health care provision by conditioning the amount of funding for a health care provider on performance. The latter is usually defined in outputs (also known as results-based financing) and often prioritizes maternal and newborn health. The World Bank is its biggest advocate and partly responsible for the rise in PBF-implementing countries from 3 in 2006 to 32 in 2013 (Eijkenaar et al., 2013).

Bibliography

Adams, A. M., T. G. Evans, R. Mohammed, and J. Farnsworth (1997). Socioeconomic stratification by wealth ranking: Is it valid? *World Development 25*(7), 1165–1172.

Ahmed, A. U. and H. E. Bouis (2002). Weighing what's practical: Proxy means tests for targeting food subsidies in Egypt. *Food Policy 27*(5), 519–540.

Akerlof, G. A. (1970). The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*(3), 488–500.

Akresh, R., D. De Walque, and H. Kazianga (2014). Alternative cash transfer delivery mechanisms: Impacts on routine preventative health clinic visits in Burkina Faso. In *African Successes, Volume II: Human Capital*, NBER Chapters, pp. 113–135. National Bureau of Economic Research, Inc.

Alatas, V., A. Banerjee, and R. Hanna (2012). Targeting the poor: Evidence from a field experiment in Indonesia. *American Economic Review 102*(4), 1206–1240.

Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, R. Purnamasari, and M. Wai-Poi (2013a). Does elite capture matter? Local elites and targeted welfare programs in Indonesia. Working Paper Series 13-008, Harvard University, John F. Kennedy School of Government.

Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, R. Purnamasari, and M. Wai-Poi (2013b). Ordeal mechanisms in targeting: Theory and evidence from a field experiment in Indonesia. Working Paper 19127, National Bureau of Economic Research.

Alderman, H. (2002). Do local officials know something we don't? Decentralization of targeted transfers in Albania. *Journal of Public Economics 83*(3), 375–404.

Alesina, A., R. Baqir, and W. Easterly (1999). Public goods and ethnic divisions. *The Quarterly Journal of Economics 114*(4), 1243–1284.

Alesina, A., R. Baqir, and C. Hoxby (2004). Political jurisdictions in heterogeneous communities. *Journal of Political Economy 112*(2), 348–396.

Alesina, A., A. Devleeschauwer, W. Easterly, S. Kurlat, and R. Wacziarg (2003). Fractionalization. *Journal of Economic Growth 8*(2), 155–194.

Alesina, A., C. Gennaioli, and S. Lovo (2014). Public goods and ethnic diversity: Evidence from deforestation in indonesia. NBER Working Paper 20504, National Bureau of Economic Research.

Alesina, A. and E. La Ferrara (2000). Participation in heterogeneous communities. *The Quarterly Journal of Economics 115*(3), 847–904.

Algan, Y., C. Hemet, and D. D. Laitin (2016). The social effects of ethnic diversity at the local level: A natural experiment with exogenous residential allocation. *Journal of Political Economy 124*(3), 696–733.

Alkire, S., J. Foster, S. Seth, M. E. Santos, J. M. Roche, and P. Ballon (2015). *Multidimensional poverty measurement and analysis*. New York, USA: Oxford University Press.

Alkire, S. and M. E. Santos (2010). Acute multidimensional poverty: A new index for developing countries. Human development report office background paper 2010/11, United Nations Development Programme.

Alkire, S. and M. E. Santos (2014). Measuring acute poverty in the developing world: Robustness and scope of the multidimensional poverty index. *World Development 59*(7), 251–274.

Alkire, S. and S. Seth (2013). Selecting a targeting method to identify BPL households in India. *Social Indicators Research 112*(2), 417–446.

Allport, G. W. (1954). *The nature of prejudice*. Cambridge, MA: Addison-Wesley.

Angrist, J. D. and J.-S. Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton and Oxford.

Arcand, J.-L. and M. Fafchamps (2012). Matching in community-based organizations. *Journal of Development Economics 98*(2), 203–219.

Auriol, E. and R. J. Gary-Bobo (2012). On the optimal number of representatives. *Public Choice 153*(3-4), 419–445.

Azevedo, V. and M. Robles (2013). Multidimensional targeting: Identifying beneficiaries of conditional cash transfer programs. *Social Indicators Research 112*(2), 447–475.

Baird, S., C. McIntosh, and B. Özler (2013). The regressive demands of demand-driven development. *Journal of Public Economics 106*(1), 27–41.

Banerjee, A., E. Duflo, R. Chattopadhyay, and J. Shapiro (2007). Targeting efficiency: How well can we identify the poor? Working Paper 3, Centre for Micro Finance.

Banerjee, A., M. R. Hanna, and J. Kyle (2014). Information is power: Identification cards and food subsidy programs in Indonesia. NBER Working Paper 20923, National Bureau of Economic Research.

Banerjee, A., L. Iyer, and R. Somanathan (2005). History, social divisions, and public goods in rural India. *Journal of the European Economic Association 3*(2-3), 639–647.

Banerjee, A. V. and E. Duflo (2011). *Poor economics: A radical rethinking of the way to fight global poverty.* New York: Public Affairs.

Bardhan, P. and D. Mookherjee (2005). Decentralizing antipoverty program delivery in developing countries. *Journal of Public Economics 89*(4), 675–704.

Bardhan, P. and D. Mookherjee (2006). Pro-poor targeting and accountability of local governments in West Bengal. *Journal of Development Economics 79*(2), 303–327.

Basurto, P., P. Dupas, and J. Robinson (2016). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi. Working Paper 23383, Stanford University.

Bauhoff, S., D. R. Hotchkiss, and O. Smith (2011). The impact of medical insurance for the poor in Georgia: A regression discontinuity approach. *Health Economics 20*(11), 1362–1378.

Becker, G. S. (1957). *The economics of discrimination.* University of Chicago press.

Beekman, G., E. Bulte, and E. Nillesen (2014). Corruption, investments and contributions to public goods: Experimental evidence from rural Liberia. *Journal of Public Economics 115*(1), 37–47.

Belem, M., J. Bayala, and A. Kalinganire (2011). Defining the poor by the rural communities of Burkina Faso: Implications for the development of sustainable parkland management. *Agroforestry systems 83*(3), 287–302.

Bernal, N., M. A. Carpio, and T. J. Klein (2016). The effects of access to health insurance: Evidence from a regression discontinuity design in Peru. Technical report, Tilburg University.

Besley, T., R. Pande, and V. Rao (2012). Just rewards? local politics and public resource allocation in South India. *World Bank Economic Review 26*(2), 191–216.

Bigman, D., S. Dercon, D. Guillaume, and M. Lambotte (2000). Community targeting for poverty reduction in Burkina Faso. *The World Bank Economic Review 14*(1), 167–193.

Blumberg, L. J., L. M. Nichols, and J. S. Banthin (2001). Worker decisions to purchase health insurance. *International Journal of Health Care Finance and Economics 1*(3), 305–325.

Blumer, H. (1958). Race prejudice as a sense of group position. *Pacific Sociological Review 1*(1), 3–7.

Bocoum, F., M. Grimm, R. Hartwig, and N. Zongo (2017). Nudging households to take up health insurance: Evidence from a randomized experiment in Burkina Faso. Discussion Paper Series 10744, IZA Institute of Labor Economics.

Bramoullé, Y. and S. Goyal (2016). Favoritism. *Journal of Development Economics 122*(1), 16–27.

Brau, J. C., C. Merrill, and K. B. Staking (2011). Insurance theory and challenges facing the development of microinsurance markets. *Journal of Developmental Entrepreneurship 16*(4), 411–440.

Brown, C., M. Ravallion, and D. Van De Walle (2016). A poor means test? Econometric targeting in Africa. Policy Research Working Paper 7915, World Bank, Washington, DC.

Burgess, R., R. Jedwab, E. Miguel, A. Morjaria, and G. Padro i Miquel (2015). The value of democracy: Evidence from road building in Kenya. *The American Economic Review 105*(6), 1817–1851.

Caeyers, B. and S. Dercon (2012). Political connections and social networks in targeted transfer programs: Evidence from rural Ethiopia. *Economic Development and Cultural Change 60*(4), 639–675.

Caizhen, L. (2010). Who is poor in China? A comparison of alternative approaches to poverty assessment in rural Yunnan. *The Journal of Peasant Studies 37*(2), 407–428.

Capuno, J. J., A. D. Kraft, S. Quimbo, C. R. Tan, and A. Wagstaff (2016). Effects of price, information, and transactions cost interventions to raise voluntary enrollment in a social health insurance scheme: A randomized experiment in the Philippines. *Health Economics 25*(6), 650–662.

Carrin, G. and C. James (2005). Social health insurance: Key factors affecting the transition towards universal coverage. *International Social Security Review 58*(1), 45–64.

Castañeda, T. (2005). Targeting social spending to the poor with proxy-means testing: Colombia's SISBEN system. Social Protection Discussion Paper Series 529, The World Bank, Washington, D.C.

Central Intelligence Agency (2017). The world factbook. Website. `https://www.cia.gov/library/publications/the-world-factbook/geos/uv.html` (visited on 29/06/2017).

Chambers, R. (1994a). Participatory rural appraisal (PRA): Analysis of experience. *World Development 22*(9), 1253–1268.

Chambers, R. (1994b). The origins and practice of participatory rural appraisal. *World Development 22*(7), 953–969.

Chernew, M., K. Frick, and C. G. McLaughlin (1997). The demand for health insurance coverage by low-income workers: Can reduced premiums achieve full coverage? *Health Services Research 32*(4), 453.

Chiappori, P.-A. and B. Salanie (2000). Testing for asymmetric information in insurance markets. *Journal of Political Economy 108*(1), 56–78.

Coady, D., M. Grosh, and J. Hoddinott (2004a). Targeting outcomes redux. *The World Bank Research Observer 19*(1), 61–85.

Coady, D., M. E. Grosh, and J. Hoddinott (2004b). *Targeting of transfers in developing countries: Review of lessons and experience.* Washington, DC: The World Bank.

Coady, D. and E. Skoufias (2004). On the targeting and redistributive efficiencies of alternative transfer instruments. *Review of Income and Wealth 50*(1), 11–27.

Cohen, A. and P. Siegelman (2010). Testing for adverse selection in insurance markets. *Journal of Risk and Insurance 77*(1), 39–84.

Cole, S., X. Giné, J. Tobacman, R. Townsend, P. Topalova, and J. Vickery (2013). Barriers to household risk management: Evidence from India. *American Economic Journal. Applied economics 5*(1), 104.

Comola, M. and M. Fafchamps (2017). The missing transfers: Estimating misreporting in dyadic data. *Economic Development and Cultural Change 65*(3), 549–582.

Conning, J. and M. Kevane (2002). Community-based targeting mechanisms for social safety nets: A critical review. *World Development 30*(3), 375–394.

Dasgupta, A. and V. A. Beard (2007). Community driven development, collective action and elite capture in Indonesia. *Development and Change 38*(2), 229–249.

De Allegri, M. (2006). *Why do people choose to enrol or not to enrol in community health insurance? The case of the Nouna Health District, Burkina Faso.* Ph. D. thesis, Ruprecht-Karls-Universität Heidelberg, Heidelberg.

De Allegri, M., B. Kouyaté, H. Becher, A. Gbangou, S. Pokhrel, M. Sanon, and R. Sauerborn (2006). Understanding enrolment in community health insurance in sub-Saharan Africa: A population-based case-control study in rural Burkina Faso. *Bulletin of the World Health Organization 84*(11), 852–858.

De Allegri, M., S. Pokhrel, H. Becher, H. Dong, U. Mansmann, B. Kouyaté, G. Kynast-Wolf, A. Gbangou, M. Sanon, J. Bridges, et al. (2008). Step-wedge cluster-randomised community-based trials: An application to the study of the impact of community health insurance. *Health Research Policy and Systems 6*(10), 6–10.

De Luca, G., R. Hodler, P. A. Raschky, and M. Valsecchi (2016). Ethnic Favoritism: An Axiom of Politics? CEPR Discussion Papers 11351, Center for Economic Policy Research.

Del Ninno, C. and e. Mills, Bradford (2015). *Safety nets in Africa: Effective mechanisms to reach the poor and most vulnerable.* Washington, D.C.: Africa Development Forum Series.

Desmet, K., J. Gomes, and I. Ortuño Ortín (2016). The geography of linguistic diversity and the provision of public goods. Discussion Paper Series DP11683, Centre for Economic Policy Research, London.

Devereux, S. and K. Sharp (2006). Trends in poverty and destitution in Wollo, Ethiopia. *The Journal of Development Studies 42*(4), 592–610.

Dionne, K. Y. (2015). Social networks, ethnic diversity, and cooperative behavior in rural Malawi. *Journal of Theoretical Politics 27*(4), 522–543.

Dong, H., M. De Allegri, D. Gnawali, A. Souares, and R. Sauerborn (2009). Drop-out analysis of community-based health insurance membership at Nouna, Burkina Faso. *Health policy 92*(2), 174–179.

Dong, H., A. Gbangou, M. De Allegri, S. Pokhrel, and R. Sauerborn (2008). The differences in characteristics between health-care users and non-users: Implication for introducing community-based health insurance in Burkina Faso. *The European Journal of Health Economics 9*(1), 41–50.

Dreher, A., A. Fuchs, R. Hodler, B. Parks, P. A. Raschky, and M. J. Tierney (2016). Aid on demand: African leaders and the geography of China's foreign assistance. AidData Working Paper No. 3, AidData, Williamsburg, VA.

Dreze, J. and A. Sen (1989). *Hunger and public action.* Oxford University Press on Demand.

Dror, D. M., E. S. Soriano, M. E. Lorenzo, J. N. Sarol, R. S. Azcuna, and R. Koren (2005). Field-based evidence of enhanced healthcare utilization among persons insured by micro health insurance units in Philippines. *Health Policy 73*(3), 263–271.

Easterly, W. and R. Levine (1997). Africa's growth tragedy: Policies and ethnic divisions. *The Quarterly Journal of Economics 112*(4), 1203–1250.

Eijkenaar, F., M. Emmert, M. Scheppach, and O. Schöffski (2013). Effects of pay for performance in health care: A systematic review of systematic reviews. *Health policy 110*(2), 115–130.

Eling, M., S. Pradhan, and J. T. Schmit (2014). The determinants of microinsurance demand. *The Geneva Papers on Risk and Insurance Issues and Practice 39*(2), 224–263.

Fafchamps, M. and S. Lund (2003). Risk-sharing networks in rural Philippines. *Journal of Development Economics 71*(2), 261–287.

Ferraz, C. and F. Finan (2011). Electoral accountability and corruption: Evidence from the audits of local governments. *American Economic Review 101*(4), 1274–1311.

Filmer, D. and L. H. Pritchett (2001). Estimating wealth effects without expenditure data – Or tears: An application to educational enrollments in states of India. *Demography 38*(1), 115–132.

Filmer, D. and K. Scott (2012). Assessing asset indices. *Demography 49*(1), 359–392.

Fink, G., P. J. Robyn, A. Sié, and R. Sauerborn (2013). Does health insurance improve health: Evidence from a randomized community-based insurance rollout in rural Burkina Faso. *Journal of Health Economics 32*(6), 1043–1056.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, et al. (2012). The Oregon health insurance experiment: Evidence from the first year. *Quarterly Journal of Economics 127*(3), 1057–1106.

Fischer, T., M. Frölich, and A. Landmann (2016). Adverse selection in low-income health insurance markets: Evidence from a large-scale RCT in Pakistan. Unpublished manuskript, University of Mannheim.

Galasso, E. and M. Ravallion (2005). Decentralized targeting of an antipoverty program. *Journal of Public Economics 89*(4), 705–727.

Garcia, M., C. G. Moore, and C. M. Moore (2012). *The cash dividend: The rise of cash transfer programs in sub-Saharan Africa*. Washington, D.C.: The World Bank.

Gertler, P. and J. Gruber (2002). Insuring Consumption Against Illness. *The American Economic Review 18*(3), 257–273.

Giné, X., R. Townsend, and J. Vickery (2008). Patterns of rainfall insurance participation in rural India. *The World Bank Economic Review 22*(3), 539–566.

Gisselquist, R. M., S. Leiderer, and M. Nino-Zarazua (2016). Ethnic heterogeneity and public goods provision in Zambia: Evidence of a subnational "diversity dividend". *World Development 78*(1), 308–323.

Glaeser, E. L. and R. E. Saks (2006). Corruption in America. *Journal of public Economics 90*(6), 1053–1072.

Glennerster, R., E. Miguel, and A. D. Rothenberg (2013). Collective action in diverse Sierra Leone communities. *Economic Journal 123*(568), 285 – 316.

Govender, V., J. E. Ataguba, and O. A. Alaba (2014). Health insurance coverage within households: The case of private health insurance in South Africa. *The Geneva Papers on Risk and Insurance-Issues and Practice 39*(4), 712–726.

Grosh, M. and J. L. Baker (1995). Proxy means tests for targeting social programs. Living Standards Measurement Study Working Paper 118, The World Bank, Washington DC.

Gruber, J. and E. Washington (2005). Subsidies to employee health insurance premiums and the health insurance market. *Journal of Health Economics 24*(2), 253–276.

Habyarimana, J., M. Humphreys, D. N. Posner, and J. M. Weinstein (2007). Why does ethnic diversity undermine public goods provision? *American Political Science Review 101*(4), 709–725.

Habyarimana, J., M. Humphreys, D. N. Posner, and J. M. Weinstein (2009). *Coethnicity: Diversity and the dilemmas of collective action.* Russell Sage Foundation Series on Trust.

Handa, S., C. Huang, N. Hypher, C. Teixeira, F. V. Soares, and B. Davis (2012). Targeting effectiveness of social cash transfer programmes in three African countries. *Journal of Development Effectiveness 4*(1), 78–108.

Hargreaves, J. R., L. A. Morison, J. S. Gear, M. B. Makhubele, J. D. Porter, J. Busza, C. Watts, J. C. Kim, and P. M. Pronyk (2007). Hearing the voices of the poor: Assigning poverty lines on the basis of local perceptions of poverty. A quantitative analysis of qualitative data from participatory wealth ranking in rural South Africa. *World Development 35*(2), 212–229.

Hirschman, A. O. (1964). The paternity of an index. *The American Economic Review 54*(5), 761–762.

Honorati, M., U. Gentilini, and R. G. Yemtsov (2015). *The state of social safety nets 2015.* Washington, DC: World Bank Group.

Howe, L. D., J. R. Hargreaves, S. Gabrysch, and S. R. Huttly (2009). Is the wealth index a proxy for consumption expenditure? A systematic review. *Journal of Epidemiology and Community Health 63*(11), 871–877.

Innovations for Poverty Action (2014). 2014 global report on poverty measurement with the progress out of poverty index. Report, Grameen Foundation, Washington, DC. `http://www.progressoutofpoverty.org/global_report` (visited on 7/13/2017).

INSD (2015). Rapport enquête multisectorielle continue (EMC) 2014: Profil de pauvreté et dinégalités. Technical report, Institut national de la statistique et de la démographie (INSD), Ougadougou, Burkina Faso. `http://www.insd.bf/n/contenu/enquetes_recensements/Enq_EMC/Profil_de_pauvrete_et_d_inegalite_en_2014.pdf` (visited on 6/28/2017).

Jakab, M. and C. Krishnan (2004). Review of the strengths and weaknesses of community financing. In A. Preker and G. Garring (Eds.), *Health Financing for Poor People*, Chapter 2, pp. 53–117. Washington, D.C.: World Bank.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics 129*(2), 597.

Karlan, D. and B. Thuysbaert (2016). Targeting ultra-poor households in Honduras and Peru. *The World Bank Economic Review 36*(3), 1–38.

Kebede, B. (2009). Community wealth ranking and household surveys: An integrative approach. *The Journal of Development Studies 45*(10), 1731–1746.

Khwaja, A. I. (2009). Can good projects succeed in bad communities?. *Journal of Public Economics 93*(7-8), 899–916.

King, G., E. Gakidou, K. Imai, J. Lakin, R. T. Moore, C. Nall, N. Ravishankar, M. Vargas, M. M. Tllez-Rojo, J. E. H. vila, et al. (2009). Public policy for the poor? A randomised assessment of the Mexican universal health insurance programme. *The Lancet 373*(9673), 1447–1454.

Klasen, S. and S. Lange (2014). Accuracy and poverty impacts of proxy means-tested transfers: An empirical assessment for Bolivia. Discussion Paper 164, Courant Research Centre, Göttingen.

Kuuire, V. Z., E. Bisung, A. Rishworth, J. Dixon, and I. Luginaah (2016). Health-seeking behaviour during times of illness: A study among adults in a resource poor setting in Ghana. *Journal of Public Health 38*(4), 545–553.

La Porta, R., F. Lopez-de Silanes, A. Shleifer, and R. Vishny (1999). The quality of government. *Journal of Law, Economics, and Organization 15*(1), 222–279.

Lee, D. S. and T. Lemieux (2010). Regression discontinuity designs in economics. *Journal of Economic Literature 48*(2), 281–355.

Levine, D., R. Polimeni, and I. Ramage (2016). Insuring health or insuring wealth? An experimental evaluation of health insurance in rural Cambodia. *Journal of Development Economics 119*(1), 1–15.

Lierl, M. (2017). Elections and embezzlement: Experimental evidence from Burkina Faso. Policy Research Working Paper 8067, World Bank, Washington, DC.

Lietz, H., M. Lingani, A. Sié, R. Sauerborn, A. Souares, and Y. Tozan (2015). Measuring population health: Costs of alternative survey approaches in the Nouna Health and Demographic Surveillance System in rural Burkina Faso. *Global Health Action 8*(1), 28330.

Mansuri, G. and V. Rao (2013). *Localizing Development: Does Participation Work?*. World Bank Policy Research Report.

Mauro, P. (1995). Corruption and growth. *The Quarterly Journal of Economics 110*(3), 681–712.

Mavridis, D. (2015, MAR). Ethnic diversity and social capital in Indonesia. *World Development 67*, 376–395.

Mayoux, L. and R. Chambers (2005). Reversing the paradigm: Quantification, participatory methods and pro-poor impact assessment. *Journal of International Development 17*(2), 271–298.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics 142*(2), 698–714.

McGee, R. (2004). Constructing poverty trends in Uganda: A multidisciplinary perspective. *Development and Change 35*(3), 499–523.

Mensah, J., J. R. Oppong, and C. M. Schmidt (2010). Ghana's national health insurance scheme in the context of the health MDGs: An empirical evaluation using propensity score matching. *Health Economics 19*(1), 95–106.

Miguel, E. (2004). Tribe or nation? Nation building and public goods in Kenya versus Tanzania. *World Politics 56*(3), 327.

Miguel, E. and M. K. Gugerty (2005). Ethnic diversity, social sanctions, and public goods in Kenya. *Journal of Public Economics 89*(11-12), 2325–2368.

Miller, G., D. Pinto, and M. Vera-Hernández (2013). Risk protection, service use, and health outcomes under Colombia's health insurance program for the poor. *American Economic Journal: Applied Economics 5*(4), 61–91.

Ministere de la Santé (2011). Plan national de developpment sanitaire 2011-2020. Rapport, Ministere de la Santé, Burkina Faso. `https://www.internationalhealthpartnership.net/fileadmin/uploads/ihp/Documents/Country_Pages/Burkina_Faso/Burkina_Faso_National_Health_Strategy_2011-2020_French.pdf` (visited on 6/15/2017).

Monchuk, V. (2013). *Reducing poverty and investing in people: The new role of safety nets in Africa*. Washington, DC: World Bank Publications.

Nguyen, H. and J. Knowles (2010). Demand for voluntary health insurance in developing countries: The case of Vietnams school-age children and adolescent student health insurance program. *Social Science & Medicine 71*(12), 2074–2082.

Niehaus, P. and A. Atanassova (2013). Targeting with agents. *American Economic Journal: Economic Policy 5*(1), 206–238.

Olken, B. A. (2006). Corruption and the costs of redistribution: Micro evidence from Indonesia. *Journal of Public Economics 90*(4), 853–870.

Pan, L. and L. Christiaensen (2012). Who is vouching for the input voucher? Decentralized targeting and elite capture in Tanzania. *World Development 40*(8), 1619–1633.

Panda, P., A. Chakraborty, D. M. Dror, and A. S. Bedi (2014). Enrolment in community-based health insurance schemes in rural Bihar and Uttar Pradesh, India. *Health Policy and Planning 29*(8), 960–974.

Panda, S. (2015). Political connections and elite capture in a poverty alleviation programme in India. *Journal of Development Studies 51*(1), 50–65.

Parmar, D., M. De Allegri, G. Savadogo, and R. Sauerborn (2014). Do community-based health insurance schemes fulfil the promise of equity? A study from Burkina Faso. *Health Policy and Planning 29*(1), 76–84.

Parmar, D., S. Reinhold, A. Souares, G. Savadogo, and R. Sauerborn (2012). Does community-based health insurance protect household assets? Evidence from rural Africa. *Health Services Research 47*(2), 819–839.

Parmar, D., A. Souares, M. de Allegri, G. Savadogo, and R. Sauerborn (2012). Adverse selection in a community-based health insurance scheme in rural Africa: Implications for introducing targeted subsidies. *BMC Health Services Research 12*(1), 181.

Pitt, M. M., M. R. Rosenzweig, and M. N. Hassan (1990). Productivity, health, and inequality in the intrahousehold distribution of food in low-income countries. *American Economic Review 80*(5), 1139–1156.

Platteau, J.-P. (2004). Monitoring elite capture in community-driven development. *Development and Change 35*(2), 223–246.

Platteau, J.-P. and F. Gaspart (2003). Disciplining local leaders in community-based development. Working paper, Centre for Research on the Economics of Develpoment, Washington, DC. `http://siteresources.worldbank.org/INTPUBSERV/Resources/platteau.pdf` (visited on 13/7/2017).

Polimeni, R. and D. I. Levine (2011). Adverse selection based on observable and unobservable factors in health insurance. Impact Analyses Series 10, University of California, Berkeley.

Ravallion, M. (1993). Poverty alleviation through regional targeting: A case study for Indonesia. In K. Hoff, A. Braverman, and J. E. Stiglitz (Eds.), *The Economics of Rural Organization*, Chapter 23, pp. 453–467. Oxford: Oxford University Press.

Ravallion, M. (2003). Targeted transfers in poor countries: Revisiting the trade-offs and policy options. Policy Research Working Paper 3048, The World Bank, Washington, D.C.

Ravallion, M. (2009). How relevant is targeting to the success of an antipoverty program? *The World Bank Research Observer 24*(2), 205–231.

Robano, V. and S. C. Smith (2013). Multidimensional targeting and evaluation: A general framework with an application to a poverty program in Bangladesh. IZA Discussion Papers 7593, Institute for the Study of Labor (IZA).

Robertson, L., P. Mushati, M. Skovdal, J. W. Eaton, J. C. Makoni, T. Crea, G. Mavise, L. Dumba, C. Schumacher, L. Sherr, C. Nyamukapa, and S. Gregson (2014). Involving communities in the targeting of cash transfer programs for vulnerable children: Opportunities and challenges. *World Development 54*(1), 325–337.

Rose, E. (1999). Consumption smoothing and excess female mortality in rural India. *Review of Economics and Statistics 81*(1), 41–49.

Rosenzweig, M. R. and T. P. Schultz (1982). Market opportunities, genetic endowments, and intrafamily resource distribution: Child survival in rural India. *American Economic Review 72*(4), 803–815.

Royalty, A. B. and J. Hagens (2005). The effect of premiums on the decision to participate in health insurance and other fringe benefits offered by the employer: Evidence from a real-world experiment. *Journal of Health Economics 24*(1), 95–112.

Sabates-Wheeler, R., A. Hurrell, and S. Devereux (2015). Targeting social transfer programmes: Comparing design and implementation errors across alternative mechanisms. *Journal of International Development 27*(8), 1521–1545.

Sauerborn, R., A. Nougtara, M. Hien, and H. J. Diesfeld (1996). Seasonal variations of household costs of illness in Burkina Faso. *Social Science & Medicine 43*(3), 281–290.

Savadogo, G. (2017). *Targeting the poorest households for subsidizing their premium for community health insurance in Nouna, Burkina Faso.* Ph. D. thesis, Karl-Ruprecht-Universität Heidelberg, Heidelberg.

Savadogo, G. and A. Souares (2006). Community wealth ranking. Auto-détermination des rangs de pauvreté par la communauté. Unpublished program proposal submitted to the ethical committee of Burkina Faso.

Savadogo, G., A. Souares, A. Sie, D. Parmar, G. Bibeau, and R. Sauerborn (2015). Using a community-based definition of poverty for targeting poor households for premium subsidies in the context of a community health insurance in Burkina Faso. *BMC Public Health 15*(84), 1–12.

Schoeps, A., S. Gabrysch, L. Niamba, A. Sié, and H. Becher (2011). The effect of distance to health-care facilities on childhood mortality in rural Burkina Faso. *American Journal of Epidemiology 173*(5), 492–498.

Schreiner, M. (2015). A comparison of two simple, low-cost ways for local, pro-poor organizations to measure the poverty of their participants. *Social Indicators Research 124*(2), 537–569.

Schuendeln, M. (2013). Ethnic heterogeneity and the private provision of public goods. *Journal of Development Studies 49*(1), 36–55.

Schüring, E. (2014). Preferences for community-based targeting - field experimental evidence from Zambia. *World Development 54*(3), 360–373.

Scoones, I. (1995). Investigating difference: Applications of wealth ranking and household survey approaches among farming households in southern Zimbabwe. *Development and Change 26*(1), 67–88.

Sen, A. (1988). The concept of development. In H. Chenery and T. Srinivasan (Eds.), *Handbook in Economics 9*, Volume 1 of *Handbook of Development Economics*, Chapter 1, pp. 9–26. Elsevier.

Shaffer, P. (1998). Gender, poverty and deprivation: Evidence from the Republic of Guinea. *World Development 26*(12), 2119–2135.

Skoufias, E., B. Davis, and S. De La Vega (2001). Targeting the poor in Mexico: An evaluation of the selection of households into PROGRESA. *World Development 29*(10), 1769–1784.

Sommerfeld, J., M. Sanon, B. A. Kouyate, and R. Sauerborn (2002). Informal risk-sharing arrangements (IRSAs) in rural Burkina Faso: Lessons for the development of community-based insurance (CBI). *The International Journal of Health Planning and Management 17*(2), 147–163.

Sosa-Rubi, S. G., O. Galarraga, and J. E. Harris (2009). Heterogeneous impact of the 'Seguro Popular' program on the utilization of obstetrical services in Mexico, 2001-2006: A multinomial probit model with a discrete endogenous variable. *Journal of Health Economics 28*(1), 20–34.

Souares, A., G. Savadogo, H. Dong, D. Parmar, A. Sie, and R. Sauerborn (2010). Using community wealth ranking to identify the poor for subsidies: A case study of community-based health insurance in Nouna, Burkina Faso. *Health and Social Care in the Community 18*(4), 363–368.

Stewart, J. M., E. OShea, C. Donaldson, and P. Shackley (2002). Do ordering effects matter in willingness-to-pay studies of health care? *Journal of Health Economics 21*(4), 585–599.

Stoeffler, Q., B. Mills, and C. del Ninno (2016). Reaching the poor: Cash transfer program targeting in Cameroon. *World Development 83*(1), 244–263.

Sundaram, K. (2003). On identification of households below poverty line in BPL Census 2002: Some comments on the proposed methodology. *Economic and Political Weekly 38*(9), 896–901.

Swee, E. L. (2015). Together or separate? Post-conflict partition, ethnic homogenization, and the provision of public schooling. *Journal of Public Economics 128*(1), 1–15.

Takasaki, Y., B. L. Barham, and O. T. Coomes (2000). Rapid rural appraisal in humid tropical forests: An asset possession-based approach and validation methods for wealth assessment among forest peasant households. *World Development 28*(11), 1961–1977.

Temu, A. E. and J. M. Due (2000). Participatory appraisal approaches versus sample survey data collection: A case of smallholder farmer's well-being ranking in Njombe District, Tanzania. *Journal of African Economies 9*(1), 44–62.

Thomas, B. K., R. Muradian, G. De Groot, and A. De Ruijter (2009). Multidimensional poverty and identification of poor households: A case from Kerala, India. *Journal of Human Development and Capabilities 10*(2), 237–257.

Thornton, R. L., L. E. Hatt, E. M. Field, M. Islam, F. Solís Diaz, and M. A. González (2010). Social security health insurance for the informal sector in Nicaragua: A randomized evaluation. *Health Economics 19*(1), 181–206.

Treisman, D. (2000). The causes of corruption: A cross-national study. *Journal of Public Economics 76*(3), 399–457.

Tullock, G. (1959). Problems of majority voting. *Journal of Political Economy 67*(6), 571–579.

Van Campenhout, B. F. (2007). Locally adapted poverty indicators derived from participatory wealth rankings: A case of four villages in rural Tanzania. *Journal of African Economies 16*(3), 406–438.

Wagstaff, A., D. Cotlear, P. H.-V. Eozenou, and L. R. Buisman (2016). Measuring progress towards universal health coverage: With an application to 24 developing countries. *Oxford Review of Economic Policy 32*(1), 147–189.

Wagstaff, A., M. Lindelow, G. Jun, X. Ling, and Q. Juncheng (2009). Extending health insurance to the rural population: An impact evaluation of China's new cooperative medical scheme. *Journal of Health Economics 28*(1), 1–19.

Wagstaff, A., H. T. H. Nguyen, H. Dao, and S. Bales (2016). Encouraging health insurance for the informal sector: A cluster randomized experiment in Vietnam. *Health Economics 25*(6), 663–674.

Wagstaff, A. and N. Watanabe (2003). What difference does the choice of SES make in health inequality measurement? *Health Economics 12*(10), 885–890.

Wang, H., W. Yip, L. Zhang, and W. C. Hsiao (2009). The impact of rural mutual health care on health status: Evaluation of a social experiment in rural China. *Health Economics 18*(1), 65–82.

Woldemichael, A., D. Z. Gurara, and A. Shimeles (2016). Community-based health insurance and out-of-pocket healthcare spending in Africa: Evidence from Rwanda. IZA Discussion Paper 9922, Institute of Labor Economics.

World Bank (2004). *Making Services Work for the Poor People.* World Development Report 2004. Washington, D.C.: The World Bank.

World Bank (2012). *Resilience, equity, and opportunity: The World Bank's social protection and labor strategy 2012-2022.* Washington, DC: World Bank.

World Bank (2014). *Risk and Opportunity: Managing risk for development.* World Development Report 2014. Washington, DC: The World Bank.

World Bank (2017a). Burkina Faso overview. Website. http://data.worldbank.org/country/burkinafaso (visited on 7/1/2017).

World Bank (2017b). World development indicators. Dataset. http://data.worldbank.org/country/burkina-faso?view=chart.com (visited on 7/1/2017).

Yao, Y., J. T. Schmit, and J. R. Sydnor (2015). The role of pregnancy in micro health insurance: Evidence of adverse selection from Pakistan. *Journal of Risk and Insurance 82*(4), 1–30.

Appendix

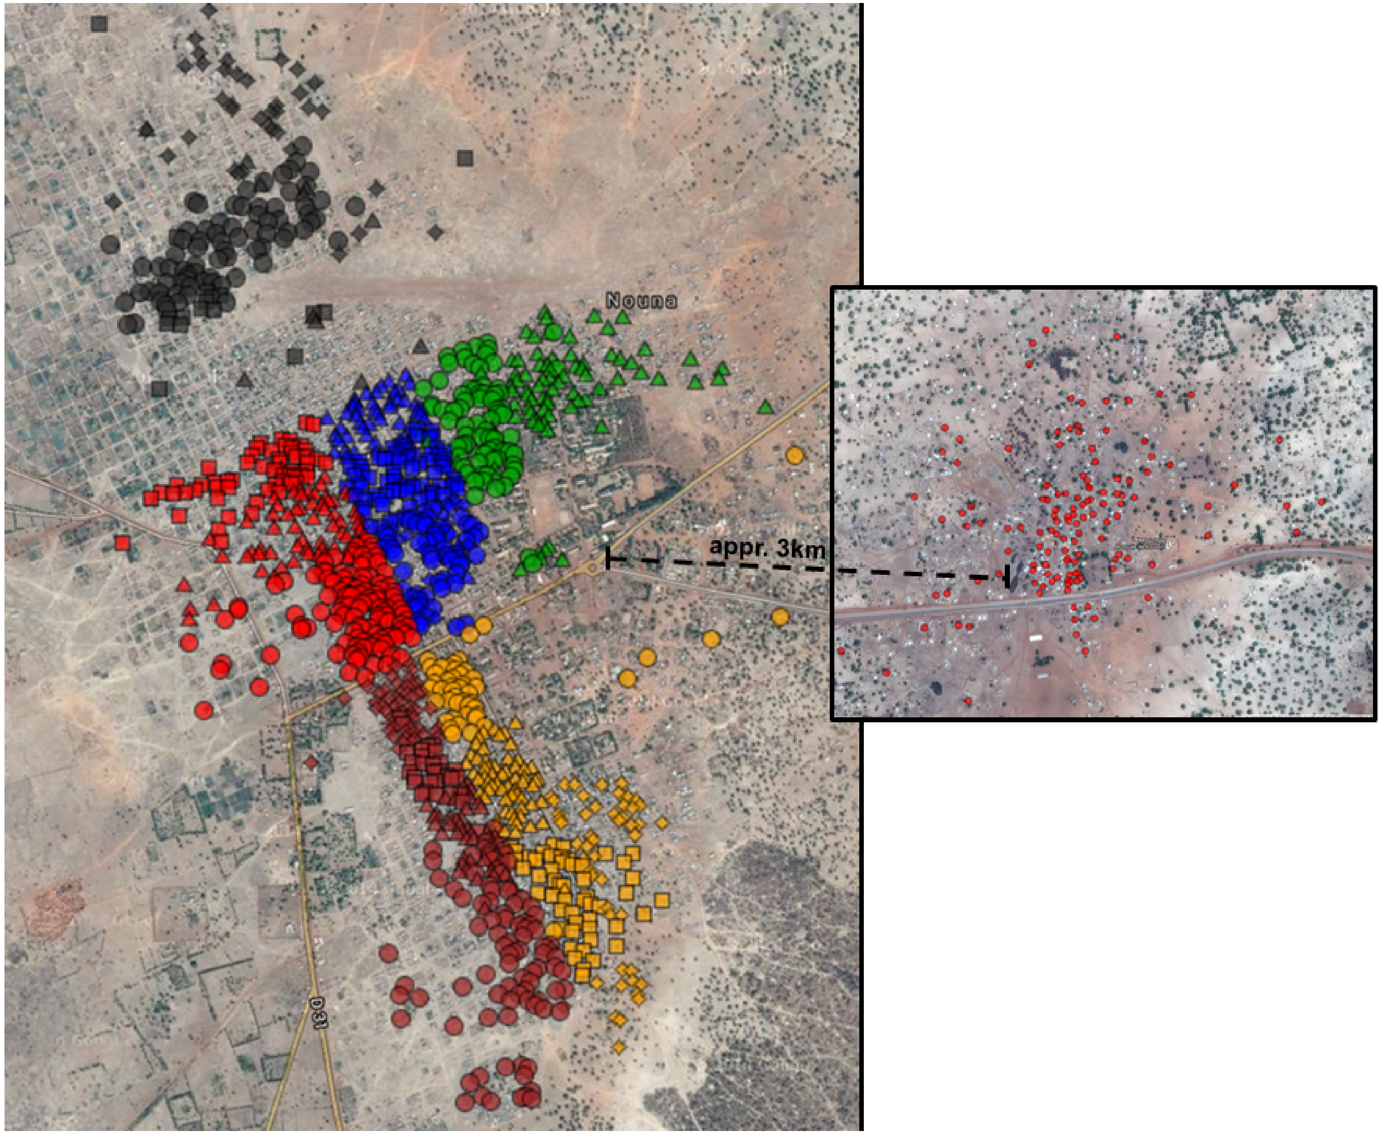

Figure A.1: Nouna town and its neighborhoods

Notes: Map depicts the distribution of households across the seven districts of Nouna town. Districts are distinguished by different colors. Within each sector there are two to four neighborhoods depicted by different symbols. District 7 is illustrated by a separate window as it is located approximately three kilometers in the West of the town's center. Created by the author with GPS Visualizer.

Figure A.2: Community-wealth ranking cards

Notes: Each staple of cards was designated for one out of three community representatives and comes with a different color. Each card contains the printed name of the household head and the household ID. After the wealth ranking exercise, facilitators manually tagged each card with its respective community rank and wealth category.

Figure A.3: Health insurance identification card

Notes: The front part on the left-hand side indicates the current insurance status and the insurance ID. The back contains personal information, namely village of residency, household ID, name of household head and insured individual, individual ID, birth date and sex. Confidential information was blackened.

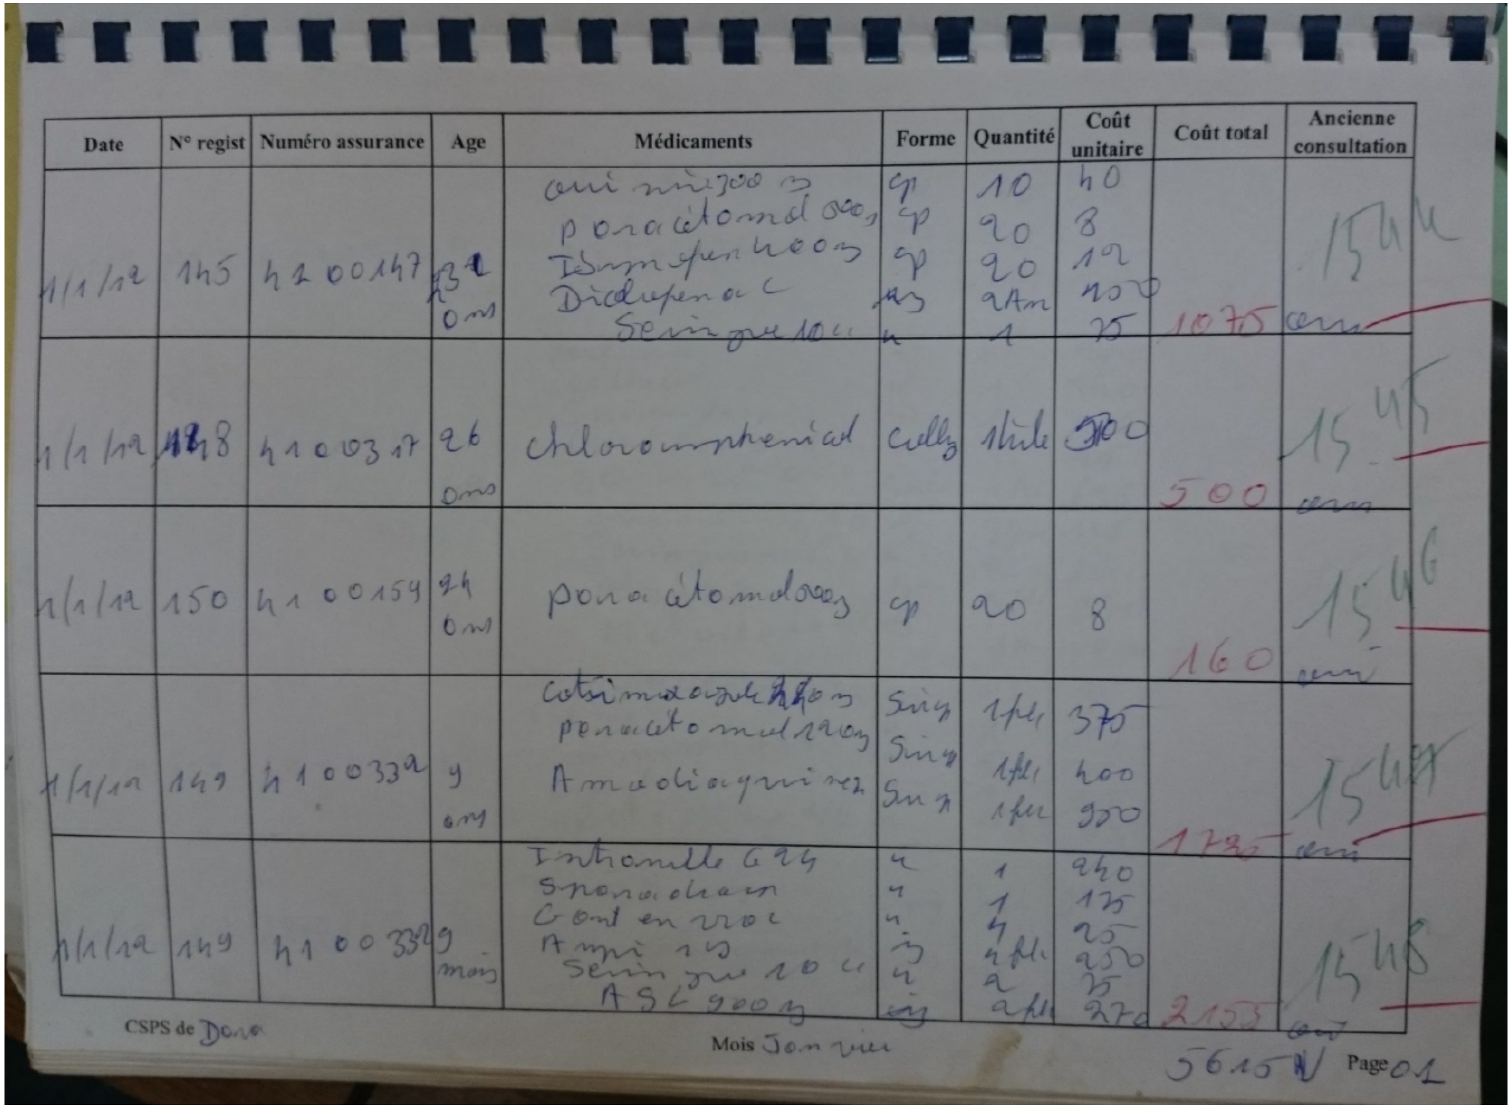

Figure A.4: Sample page of administrative health insurance records

Notes: Administrative health insurance records were kept by employees of the health care facility centers and contained ten pieces of information; date, consecutive number, insurance ID, age, medicine, unit, quantity, unit cost, total cost, and attendance of consultation.

A. Targeting Accuracy Analysis

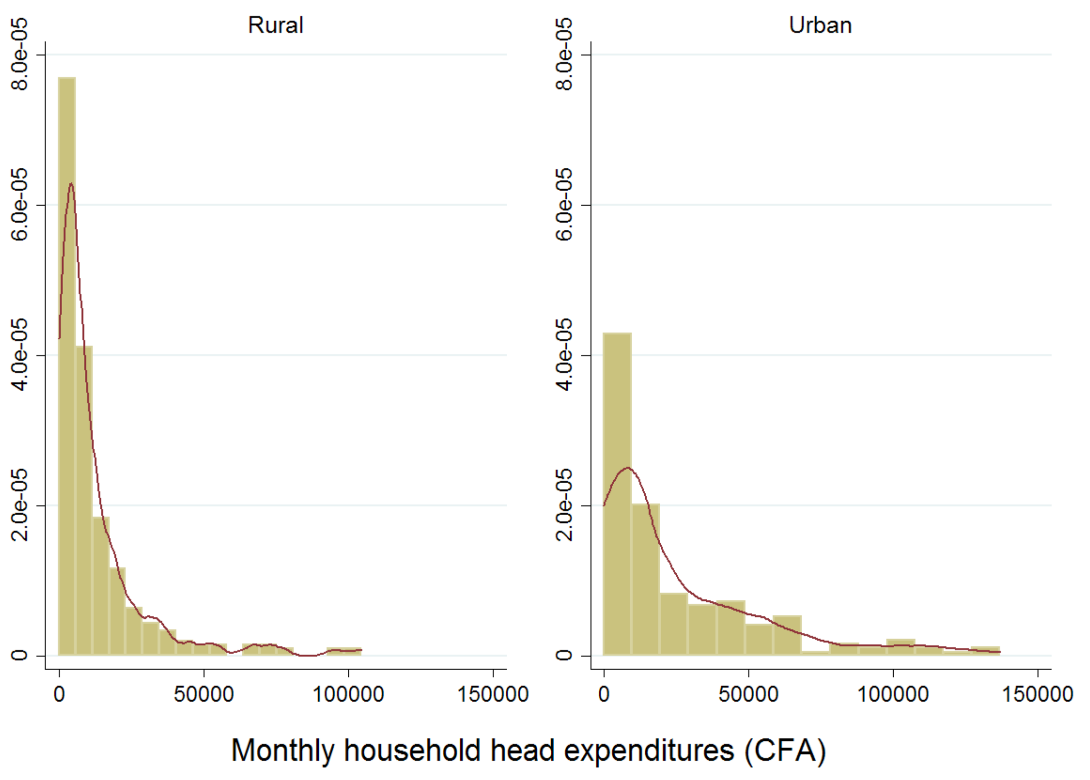

Figure B.1: Histogram and kernel densities of the consumption reference variable

Notes: Urban refers to the subsample of semi-urban households.

Table B.1: Five statistical targeting indices and their specifications

Indicators	**Number of categories used**	**Multidim. Poverty**	**Poverty Scorecard**	**Below the pov.line**	**Asset Index**	**Econometric PMT**
Total number of variables used		9	9	12	44	44
Demographic characteristics						
At least one male HH-Member at age 15	2				x	x
At least one female HH-Member at age 15	2				x	x
At least one male HH-Member at age 16 to 60	2				x	x
At least one female HH-Member at age16 to 60	2				x	x
At least one male HH-Member at age 60	2				x	x
At least one female HH-Member at age 60	2				x	x
How many household members are 15-years-old or younger?	7		x			
HHHead is married	2				x	x
HHHead is polyg married	2				x	x
HHHead is not widowed	2				x	x
HHHead is male	2				x	x
Occupational choice						
HHH can read or write	2		x		x	x
HHH response for literate	2			x	x	x
No one in the HH is literate	2	x				
Any HH member completed primary	2		x		x	x
No HH member has completed five years of schooling	2	x				
Any HH member completed secondary	2				x	x
HHH completed secondary	2				x	x
Any HH member completed tertiary	2				x	x
HHH is not employed in agriculture	2				x	x
Any HH member is not employed in agriculture	2		x		x	x
Share of employable household members	5			x		
Type of occupation (nothing, agric., non-agric.)	3			x		
Status of children (5-14 years)	3			x		
HH head is disabled	2				x	x
Dwelling characteristics						
HH uses running-water or good wells, any period	2	x			x	x
Drinking Water is changed at least every 2nd day	2				x	x
Wastewater by cesspool, gutters or septic tank	2				x	x
Water is not piped outside	2	x				
Source of drinking water	4		x			
Drinking Water arrangement	4			x		
Sanitation not at the open field	2				x	x
Toilet arrangement	3		x			
number of rooms	n.A.				x	x
Type of house	5			x		
Roof is made of concrete, metal sheets, or tile	2				x	x
Wall is not made of ordinary mud or straw	2				x	x
Floor is made of cement	2	x			x	x
Garbage evacuation through dustbin	2				x	x
No electricity or solar panel	2	x				
Main energy source of lighting	4		x			
Cooking fuel is wood	2	x				
Asset possession at household level						
At least one cart	2				x	x
At least one plow	2				x	x
At least one bike	2				x	x
At least one mbike	2		x		x	x
At least one car	2				x	x
At least one radio	2				x	x
At least one tv	2		x		x	x
At least one tel	2				x	x
At least one fridge	2				x	x
At least one kitchen	2				x	x
Ownership of 3 assets	4			x		
HH owns no assets at all	2	x				
Livestock possession at household level						
At least one horse_donkey	2				x	x
At least one goat_sheep	2				x	x
At least one chicken	2				x	x
At least one bullock	2				x	x
At least one pig	2				x	x
Number of bulocks owned by HH head	4		x	x		
Other						
HH experienced at least one severe illnes last month	2	x				
Type of risk coping	3			x		
At least one HH member emigrated last year	2			x		
Usage of transfers received	4			x		
HHs belongs to ethnic minority group	2				x	x
Relative size of primary agricultural output	5			x		

Notes: The second column *Categories* specifies the number of categories of the variable. The majority of variables consists of indicator variables which only take on two values.

Table B.2: Indicators and weights of the econometric PMT index

	Dependent variable: Eligible by consumption	
	Rural	Semi-urban
Any HH-member with primary education	-0.111	-0.155
Any HH-member with secondary education	-0.101	0.040
Any HH-member with tertiary education	0.000	0.023
Household head literate (incidence)	-0.222	-0.030
HHH response for level2	0.003	-0.086
HHH can read or write	0.147	-0.055
HH head occup. non-agric. (incidence)	0.061	0.035
Incidence of no_primagr at HH-level	-0.084	0.376
HHs belongs to ethnic minority group	0.196	-0.050
HHHead is disabled	0.033	0.134***
HH uses running-water or good wells, any period	0.015	0.078
Drinking Water is changed at least every 2nd day	-0.055	-0.053
Wastewater by cesspool, gutters or septic tank	0.021	-0.252
Garbage evacuation through dustbin	-0.064	0.171
Concrete, metal sheets, or tile	0.073	0.090
No ordinary mud or straw	-0.129	-0.143
Floor is made of cement	0.000	-0.089
Not at the open field	0.010	-0.011
Number of rooms	0.017	0.007
At least one Cart	-0.038	0.128
At least one Plow	0.017	-0.073
Bicycle	0.134	-0.107
Motorbike	0.031	-0.014
At least one Car	0.093	-0.295
At least one Radio	-0.111*	-0.080
At least one TV	-0.066	-0.194**
At least one Telephone	-0.064	0.069
At least one Fridge	-0.039	0.110
At least one Kitchen	0.000	-0.162
At least one Horse or donkey	0.067	-0.083
Goat or sheep	0.040	-0.073
At least one Chicken	-0.019	0.097
Bullock	-0.097	0.077
At least one Pig	0.061	-0.043
There is a male HH-Member at age15	-0.028	-0.102
There is a female HH-Member at age15	-0.099	0.061
There is a male HH-Member at age1660	-0.065	0.043
There is a female HH-Member at age1660	-0.104	-0.164
There is a male HH-Member at age60	0.157*	0.095
There is a female HH-Member at age60	0.118	0.018
HH head is married	0.175	-0.180
HH head is polyg married	-0.200	0.178
HH head is not widowed	-0.076	-0.288*
HH head is male	0.141	0.024
Observations	349	211
R^2	0.39	0.50
F-test (p-value)	0.00	0.00

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the household level. Regressions include community fixed effects.

Table B.3: Indicators and weights of the asset index

	1st principal component factor loadings	
	Rural	Semi-urban
Any HH-member with primary education	0.181	0.267
Any HH-member with secondary education	0.107	0.182
Any HH-member with tertiary education	omitted	0.030
Household head is literate	0.155	0.134
HHH response for level2	0.108	0.091
HHH can read or write	0.154	0.125
HH head employed non-agriculturally	-0.102	-0.050
No HH member applied in agriculture	0.165	0.122
HHs belongs to ethnic minority group	-0.029	-0.033
HHHead is disabled	0.002	-0.049
HH uses running-water or good wells, any period	-0.042	0.074
Drinking Water is changed at least every 2nd day	0.002	0.111
Wastewater by cesspool, gutters or septic tank	-0.015	-0.031
Garbage evacuation through dustbin	0.011	0.045
Concrete, metal sheets, or tile	0.137	0.142
No ordinary mud or straw	0.041	0.102
Floor is made of cement	omitted	0.012
Not at the open field	0.101	0.191
Number of rooms	0.228	0.231
At least one Cart	0.258	0.199
At least one Plow	0.271	0.166
At least one Bicycle	0.170	0.215
At least one Motorbike	0.217	0.208
At least one Car	0.014	0.053
At least one Radio	0.171	0.180
At least one TV	0.164	0.219
At least one Telephone	0.168	0.260
At least one Fridge	0.053	0.047
At least one Kitchen	omitted	0.067
At least one Horse or donkey	0.257	0.206
At least one Goat or sheep	0.239	0.183
At least one Chicken	0.169	0.150
At least one Bullock	0.227	0.145
At least one Pig	0.080	0.041
There is a male HH-Member at age15	0.200	0.173
There is a female HH-Member at age15	0.215	0.185
There is a male HH-Member at age1660	0.216	0.240
There is a female HH-Member at age1660	0.204	0.255
There is a male HH-Member at age60	-0.033	-0.017
There is a female HH-Member at age60	0.035	-0.032
HH head is married	0.164	0.128
HH head is polyg married	0.168	0.130
HH head is not widowed	0.141	0.150
HH head is male	-0.094	-0.114
Observations	349	211
Number of principal components	41	44

Notes: Weights are derived from a principal component analysis where all variables are first demeaned at the community level.

Table B.4: The Poverty Scorecard Index

Original scorecard for Burkina Faso		*Scorecard adjusted for our study*	
Indicator	**Score**	**Indicator**	**Score**
1. How many household members are 14-years-old or younger?		1. How many household members are 15-years-old or younger?	
A. Six or more	0	A. Six or more	0
B. Five	5	B. Five	5
C. Four	6	C. Four	6
D. Three	10	D. Three	10
E. Two	13	E. Two	13
F. One	19	F. One	19
G. None	29	G. None	29
2. In what languages can the male head/spouse read and write?		2. HH head can read and/or write	
A. None, or no male head/spous	0	A. No	0
B. French only	4	B. Yes	4
C. A non-French language (regardless of French literacy)	5		
3. Has the female head/spouse completed first grade?		3. First grade completed by HH head	
A. No	0	A. No	0
B. No female head/spouse	0		
C. Yes	9	B. Yes	9
4. What is the main source of energy for lighting?		4. What is the main source of energy for lighting?	
A. Firewood, or other	0	A. Firewood, or other	0
B. Candles, kerosene, or LPG	4	B. Candles or oil lamp	4
C. Flashlight, or batteries	5	C. Flashlight	5
D. Electricity, or solar energy	8	D. Electricity, solar panel or battery	8
5. What toilet arrangement does the household have?		5. What toilet arrangement does the household have?	
A. No toilet arrangement, or other	0	A. Open field	0
B. Non-ventilated pit latrine	4	B. Latrine	4
C. Ventilated pit latrine, or flush to a septic tank	15	C. Ventilated latrine and flush toilet	15
6. Does the household own a television?		6. Does the household own a television?	
A. No	0	A. No	0
B. Yes	10	B. Yes	10
7. Does the household own a bed or a mattress?		7. Omitted	
A. No	0		
B. Yes	3		
8. Does the household own a scooter or a motorcycle?		8. Does the household own a scooter or a motorcycle?	
A. No	0	A. No	0
B. Yes	6	B. Yes	6
9. Have any household members, in their main occupation in the last seven days, worked in agriculture, animal husbandry, fishing, or forestry?		9. Is the primary occupation of the HH head in agriculture?	
A. Yes	0	A. Yes	0
B. No	8	B. No	8
10. How many head of cattle or other large animals does the household now own?		10. How many head of bullocks does the household head now own?	
A. None, or one	0	A. None, or one	0
B. Two	2	B. Two	2
C. Three to five	3	C. Three to five	3
D. Six or more	7	D. Six or more	7

Notes: This was retrieved from the following link on September 10, 2016: `http://www.progressoutofpoverty.org/country/burkina-faso`

Table B.5: The Below the Poverty Line Scorecard

Original scorecard		Score				
	Indicator	**0**	**1**	**2**	**3**	**4**
1	Size group of operational holding of land	Nil	Less than 1 ha of un-irrigated land (or less than 0.5 ha of irrigated land)	1-2 ha of un-irrigated land (or 0.5-1 ha of irrigated land)	2-5 ha of un-irrigated land (or 1.0 -2.5 ha of irrigated land)	More than 5 ha of un-irrigated land (or 2.5 ha of irrigated land)
2	Type of house	Houseless	Kutcha	Semi-pucca	Pucca	Urban type
3	Average availability of normal wear clothing (per household in pieces)	Less than 2	2 or more, but less than 4	4 or more, but less than 6	6 or more, but less than 10	10 or more
4	Food security	Less than one square meal per day for a major part of the year	Normally, one square meal per day, but less than one square meal occasionally	One square meal per day throughout the year	Two square meals per day with occasional shortage	Enough food throughout the year
5	Sanitation	Open defecation	Group latrine with irregular water supply	Group latrine with regular water supply	Clean group latrine with regular water supply and regular sweeper	Private latrine
6	Ownership of consumer durables: Do you own TV, electric fan, radio, pressure cooker	Nil	Any one	Two items only	Any three or all items	All items and/or any one of a list of "luxury" items
7	Literacy status of the highest literate adult	Up to Primary (Class V)	Completed Secondary (passed class X)	Graduate/ Professional diploma	Post-graduate/ Professional graduate	Up to Primary (class V)
8	Status of the household labour force	Bonded labour	Female and child labour	Only adult females and no child labor	Adult males only	Others
9	Means of livelihood	Casual labour	Subsistence cultivation	Artisan	Salary	Others
10	Status of children (5-14 years) [any child]	Not going to school and working	Going to school and working			Going to school and not working
11	Type of indebtedness	For daily consumption purposes from informal sources	For production purpose from informal sources	For other purpose from informal sources	Borrowing only from institutional agencies	No indebtedness and possession of assets
12	Reason for migration from household	Casual work	Seasonal employment	Other forms of livelihood	Non-migrant	Other purposes
13	Preference of assistance	Wage Employment/ TPDS	Self-employment	Training and skill upgradation	Housing	Loan/subsidy of more than Rs. 1,00,000

Scorecard adjusted for our study		Score				
	Indicator	**0**	**1**	**2**	**3**	**4**
1	Size group of primary agricultural output	-	Lowest quartile	Second quartile	Third quartile	Fourth quartile
2	Type of house	Round hut (traditional)	House with outside facilities	House with inside facilities	Villa	Multi-storey building
3	Omitted	-	-	-	-	-
4	Drinking water source	Surface	Ordinary well	-	Good well	Running water
5	Sanitation	Open defecation	Regular latrine	-	Ventilated latrine	Flush toilet
6	Ownership of consumer durables: Do you own radio, telephone, bike	Nil	Any one	Two items only	All three items	All items and at least one of the following items: TV, fridge, motorbike, car
7	Literacy status of the highest literate adult	Not literate	Literate	Primary school	Secondary education	Tertiary education
8	Share of employable household members	0	0 < share < 0.25	0.25 < share < 0.5	0.5 < share < 0.75	0.75 < share < 1
9	Type of occupation	No occupation	Primarily agriculture	-	-	Not primarily agriculture
10	Status of children (5-14 years) [any child]	Not going to school and working	Not going to school and doing domestic work	Not going to school and not working	-	Going to school and not working
11	Type of risk coping	Nothing	Member of a savings group	-	-	Health insurance in 2006
12	At least one household member emigrated in the last year	-	Yes	-	No	-
13	Use of transfers received	Family support	Education	-	Celebrations	No transfer received

Notes: The original scorecard is from Alkire and Seth (2013)

Table B.6: The Multidimensional Poverty Index

	Original MPI		MPI adjusted for our study	
Indicator		**Weight**	**Indicator**	**Adjusted MPI Weight**
Education 0.333				
Years of Schooling	No household member has completed five years of schooling	0.167	No household member has completed five years of schooling	0.167
Child School Attendance	Any school-aged child is not attending school in years 1 to 8	0.167	No one is literate	0.167
Health 0.333				
Mortality	Any child has died in the family	0.167	Any severe-illness in the last month	0.333
Nutrition	Any adult or child for whom there is nutritional information is malnourished	0.167	-	
Standard of Living 0.333				
Electricity	The household has no electricity	0.056	The household has no electricity or solar panel	0.056
Sanitation	The households sanitation facility is not improved (according to the MDG guidelines), or it is improved but shared with other households	0.056	Water not piped outside	0.056
Water	The household does not have access to clean drinking water (according to the MDG guidelines) or clean water is more than 30 minutes walking from home.	0.056	Ordinary water source	0.056
Floor	The household has dirt, sand or dung floor	0.056	Floor is not cement	0.056
Cooking Fuel	The household cooks with dung, wood or charcoal.	0.056	Household cooks with wood	0.056
Assets	The household does not own more than one of: radio, TV, telephone, bike, motorbike or refrigerator, and does not own a car or truck.	0.056	HH owns no assets	0.056

Notes: The original MPI is from Alkire and Santos (2010)

Table B.7: Mean targeting error rates, alternative consumption definition

	Rural			Semi-urban		
	(1) Mean Targeting Error	(2) Mean Exclusion Error	(3) Mean Inclusion Error	(4) Mean Targeting Error	(5) Mean Exclusion Error	(6) Mean Inclusion Error
Econometric PMT	0.19	0.37	0.13	0.11	0.24	0.07
Asset index	0.25	0.48	0.17	0.25	0.54	0.17
Scorecards						
Below the poverty line	0.28	0.54	0.18	0.22	0.46	0.14
Poverty Scorecard Index	0.32	0.62	0.21	0.33	0.70	0.22
Multidimensional Poverty Index	0.33	0.65	0.22	0.30	0.64	0.20
Community-based targeting	0.28	0.55	0.19	0.25	0.52	0.16
Random targeting error	0.38	0.74	0.26	0.36	0.76	0.24
Number of households	349	89	260	212	50	162

Table B.8: Differences between pairs of MTEs, Rural Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		-4.0	-7.3*	-10.7**	2.3	-7.9*	-11.3***	-13.6***	-7.9*	-11.3***	-17.5***	-14.1***	-9.0**	-10.7***	-20.3***	-8.5*
MPI(#7)			-3.4	-6.8*	6.2*	-4.0	-7.3*	-9.6**	-4.0	-7.3*	-13.6***	-10.2***	-5.1*	-6.8*	-16.4***	-4.5
BPL(#12)				-3.4	9.6**	-0.6	-4.0	-6.2*	-0.6	-4.0	-10.2***	-6.8*	-1.7	-3.4	-13.0***	-1.1
PCA-based Weight																
PSI(#10)					13.0***	2.8	-0.6	-2.8	2.8	-0.6	-6.8**	-3.4	1.7	0.0	-9.6***	2.3
MPI(#7)						-10.2**	-13.6***	-15.8***	-10.2**	-13.6***	-19.8***	-16.4***	-11.3***	-13.0***	-22.6***	-10.7**
BPL(#12)							-3.4	-5.6*	0.0	-3.4	-9.6***	-6.2*	-1.1	-2.8	-12.4***	-0.6
All(#44)								-2.3	3.4	0.0	-6.2**	-2.8	2.3	0.6	-9.0***	2.8
Log(expd)-based Weight																
PSI(#10)									5.6*	2.3	-4.0*	-0.6	4.5	2.8	-6.8**	5.1
MPI(#7)										-3.4	-9.6***	-6.2*	-1.1	-2.8	-12.4***	-0.6
BPL(#12)											-6.2**	-2.8	2.3	0.6	-9.0***	2.8
All(#44)												3.4	8.5***	6.8**	-2.8	9.0***
Elig(expd)-based Weight																
PSI(#10)													5.1	3.4	-6.2**	5.6*
MPI(#7)														-1.7	-11.3***	0.6
BPL(#12)															-9.6***	2.3
All(#44)																11.9***
CBT																
Observations	354	354	354	354	354	354	354	354	354	354	354	354	354	354	354	354

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.9: Differences between pairs of MTEs, Semi-urban Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		-0.9	-10.4**	-4.7	6.6	-10.4**	-8.5*	-10.4**	-7.5	-11.3**	-17.9***	-15.1***	-9.4*	-15.1***	-24.5***	-11.3**
MPI(#7)			-9.4**	-3.8	7.5	-9.4**	-7.5*	-9.4*	-6.6	-10.4**	-17.0***	-14.2***	-8.5*	-14.2***	-23.6***	-10.4**
BPL(#12)				5.7	17.0***	0.0	1.9	-0.0	2.8	-0.9	-7.5**	-4.7	0.9	-4.7	-14.2***	-0.9
PCA-based Weight																
PSI(#10)					11.3*	-5.7	-3.8	-5.7	-2.8	-6.6*	-13.2***	-10.4**	-4.7	-10.4***	-19.8***	-6.6*
MPI(#7)						-17.0***	-15.1**	-17.0***	-14.2**	-17.9***	-24.5***	-21.7***	-16.0***	-21.7***	-31.1***	-17.9***
BPL(#12)							1.9	0.0	2.8	-0.9	-7.5**	-4.7	0.9	-4.7*	-14.2***	-0.9
All(#44)								-1.9	0.9	-2.8	-9.4***	-6.6*	-0.9	-6.6*	-16.0***	-2.8
Log(expd)-based Weight																
PSI(#10)									2.8	-0.9	-7.5*	-4.7*	0.9	-4.7	-14.2***	-0.9
MPI(#7)										-3.8	-10.4**	-7.5*	-1.9	-7.5**	-17.0***	-3.8
BPL(#12)											-6.6*	-3.8	1.9	-3.8	-13.2***	-0.0
All(#44)												2.8	8.5*	2.8	-6.6**	6.6*
Elig(expd)-based Weight																
PSI(#10)													5.7	-0.0	-9.4**	3.8
MPI(#7)														-5.7	-15.1***	-1.9
BPL(#12)															-9.4***	3.8
All(#44)																13.2***
CBT																
Observations	212	212	212	212	212	212	212	212	212	212	212	212	212	212	212	212

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.10: Differences between pairs of MTEs, Extremely poor households, Rural Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		-21.9*	-25.0*	-46.9***	3.1	-31.3*	-46.9***	-59.4***	-43.8***	-50.0***	-71.9***	-62.5***	-46.9***	-46.9***	-78.1***	-37.5**
MPI(#7)			-3.1	-25.0	25.0*	-9.4	-25.0	-37.5*	-21.9	-28.1*	-50.0***	-40.6**	-25.0*	-25.0*	-56.3***	-15.6
BPL(#12)				-21.9	28.1*	-6.3	-21.9*	-34.4**	-18.8	-25.0*	-46.9***	-37.5**	-21.9	-21.9*	-53.1***	-12.5
PCA-based Weight																
PSI(#10)					50.0***	15.6	0.0	-12.5	3.1	-3.1	-25.0**	-15.6	0.0	0.0	-31.3**	9.4
MPI(#7)						-34.4**	-50.0***	-62.5***	-46.9***	-53.1***	-75.0***	-65.6***	-50.0***	-50.0***	-81.3***	-40.6**
BPL(#12)							-15.6	-28.1*	-12.5	-18.8	-40.6***	-31.3*	-15.6	-15.6	-46.9***	-6.3
All(#44)								-12.5	3.1	-3.1	-25.0**	-15.6	0.0	0.0	-31.3**	9.4
Log(expd)-based Weight																
PSI(#10)									15.6	9.4	-12.5	-3.1	12.5	12.5	-18.8*	21.9
MPI(#7)										-6.3	-28.1**	-18.8	-3.1	-3.1	-34.4**	6.3
BPL(#12)											-21.9*	-12.5	3.1	3.1	-28.1**	12.5
All(#44)												9.4	25.0**	25.0**	-6.3	34.4**
Elig(expd)-based Weight																
PSI(#10)													15.6	15.6	-15.6	25.0*
MPI(#7)														0.0	-31.3**	9.4
BPL(#12)															-31.3**	9.4
All(#44)																40.6***
CBT																
Observations	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.11: Differences between pairs of MTEs, Moderately poor households, Rural Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		0.0	-9.6	-7.7	5.8	-7.7	-9.6	-9.6	0.0	-7.7	-15.4	-9.6	-1.9	-7.7	-21.2*	-5.8
MPI(#7)			-9.6	-7.7	5.8	-7.7	-9.6	-9.6	0.0	-7.7	-15.4	-9.6	-1.9	-7.7	-21.2*	-5.8
BPL(#12)				1.9	15.4	1.9	0.0	0.0	9.6	1.9	-5.8	0.0	7.7	1.9	-11.5	3.8
PCA-based Weight																
PSI(#10)					13.5	0.0	-1.9	-1.9	7.7	0.0	-7.7	-1.9	5.8	0.0	-13.5	1.9
MPI(#7)						-13.5	-15.4	-15.4	-5.8	-13.5	-21.2*	-15.4	-7.7	-13.5	-26.9**	-11.5
BPL(#12)							-1.9	-1.9	7.7	0.0	-7.7	-1.9	5.8	0.0	-13.5	1.9
All(#44)								0.0	9.6	1.9	-5.8	-0.0	7.7	1.9	-11.5	3.8
Log(expd)-based Weight																
PSI(#10)									9.6	1.9	-5.8	0.0	7.7	1.9	-11.5	3.8
MPI(#7)										-7.7	-15.4*	-9.6	-1.9	-7.7	-21.2*	-5.8
BPL(#12)											-7.7	-1.9	5.8	0.0	-13.5	1.9
All(#44)												5.8	13.5	7.7	-5.8	9.6
Elig(expd)-based Weight																
PSI(#10)													7.7	1.9	-11.5	3.8
MPI(#7)														-5.8	-19.2*	-3.8
BPL(#12)															-13.5	1.9
All(#44)																15.4
CBT																
Observations	52	52	52	52	52	52	52	52	52	52	52	52	52	52	52	52

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.12: Differences between pairs of MTEs, Around median households, Rural Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		-1.4	-3.5	-3.5	-4.9	-3.5	-3.5	-5.6	-4.2	-3.5	-8.3	-6.3	-5.6	-4.9	-11.8*	-1.4
MPI(#7)			-2.1	-2.1	-3.5	-2.1	-2.1	-4.2	-2.8	-2.1	-6.9	-4.9	-4.2	-3.5	-10.4*	-0.0
BPL(#12)				-0.0	-1.4	0.0	0.0	-2.1	-0.7	-0.0	-4.9	-2.8	-2.1	-1.4	-8.3*	2.1
PCA-based Weight																
PSI(#10)					-1.4	0.0	0.0	-2.1	-0.7	0.0	-4.9	-2.8	-2.1	-1.4	-8.3*	2.1
MPI(#7)						1.4	1.4	-0.7	0.7	1.4	-3.5	-1.4	-0.7	-0.0	-6.9	3.5
BPL(#12)							0.0	-2.1	-0.7	-0.0	-4.9	-2.8	-2.1	-1.4	-8.3	2.1
All(#44)								-2.1	-0.7	0.0	-4.9	-2.8	-2.1	-1.4	-8.3*	2.1
Log(expd)-based Weight																
PSI(#10)									1.4	2.1	-2.8	-0.7	0.0	0.7	-6.3	4.2
MPI(#7)										0.7	-4.2	-2.1	-1.4	-0.7	-7.6	2.8
BPL(#12)											-4.9	-2.8	-2.1	-1.4	-8.3*	2.1
All(#44)												2.1	2.8	3.5	-3.5	6.9
Elig(expd)-based Weight																
PSI(#10)													0.7	1.4	-5.6	4.9
MPI(#7)														0.7	-6.3	4.2
BPL(#12)															-6.9	3.5
All(#44)																10.4**
CBT																
Observations	144	144	144	144	144	144	144	144	144	144	144	144	144	144	144	144

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.13: Differences between pairs of MTEs, Affluent households, Rural Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		-4.0	-6.3	-11.1*	8.7	-7.1	-11.9*	-12.7**	-6.3	-11.9**	-15.1**	-12.7**	-6.3	-9.5*	-15.1***	-10.3*
MPI(#7)			-2.4	-7.1	12.7**	-3.2	-7.9*	-8.7*	-2.4	-7.9	-11.1**	-8.7*	-2.4	-5.6	-11.1*	-6.3
BPL(#12)				-4.8	15.1**	-0.8	-5.6	-6.3	-0.0	-5.6	-8.7*	-6.3	-0.0	-3.2	-8.7*	-4.0
PCA-based Weight																
PSI(#10)					19.8***	4.0	-0.8	-1.6	4.8	-0.8	-4.0	-1.6	4.8	1.6	-4.0	0.8
MPI(#7)						-15.9**	-20.6***	-21.4***	-15.1**	-20.6***	-23.8***	-21.4***	-15.1**	-18.3**	-23.8***	-19.0***
BPL(#12)							-4.8	-5.6	0.8	-4.8	-7.9*	-5.6	0.8	-2.4	-7.9*	-3.2
All(#44)								-0.8	5.6	0.0	-3.2	-0.8	5.6	2.4	-3.2	1.6
Log(expd)-based Weight																
PSI(#10)									6.3	0.8	-2.4	-0.0	6.3	3.2	-2.4	2.4
MPI(#7)										-5.6*	-8.7**	-6.3	0.0	-3.2	-8.7*	-4.0
BPL(#12)											-3.2	-0.8	5.6	2.4	-3.2	1.6
All(#44)												2.4	8.7*	5.6	0.0	4.8
Elig(expd)-based Weight																
PSI(#10)													6.3	3.2	-2.4	2.4
MPI(#7)														-3.2	-8.7*	-4.0
BPL(#12)															-5.6	-0.8
All(#44)																4.8
CBT																
Observations	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.14: Differences between pairs of MTEs, Extremely poor households, Semi-urban Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		-9.5	-42.9**	-9.5	14.3	-38.1*	-38.1*	-28.6	-28.6	-42.9*	-47.6**	-33.3*	-38.1*	-47.6**	-61.9***	-33.3
MPI(#7)			-33.3**	-0.0	23.8	-28.6*	-28.6	-19.0	-19.0	-33.3*	-38.1**	-23.8*	-28.6	-38.1**	-52.4***	-23.8
BPL(#12)				33.3*	57.1**	4.8	4.8	14.3	14.3	-0.0	-4.8	9.5	4.8	-4.8	-19.0	9.5
PCA-based Weight																
PSI(#10)					23.8	-28.6*	-28.6*	-19.0	-19.0	-33.3*	-38.1**	-23.8	-28.6	-38.1**	-52.4***	-23.8
MPI(#7)						-52.4**	-52.4**	-42.9**	-42.9*	-57.1**	-61.9***	-47.6**	-52.4**	-61.9***	-76.2***	-47.6**
BPL(#12)							0.0	9.5	9.5	-4.8	-9.5	4.8	0.0	-9.5	-23.8	4.8
All(#44)								9.5	9.5	-4.8	-9.5	4.8	-0.0	-9.5	-23.8	4.8
Log(expd)-based Weight																
PSI(#10)									0.0	-14.3	-19.0	-4.8	-9.5	-19.0	-33.3**	-4.8
MPI(#7)										-14.3	-19.0	-4.8	-9.5	-19.0*	-33.3*	-4.8
BPL(#12)											-4.8	9.5	4.8	-4.8	-19.0	9.5
All(#44)												14.3	9.5	0.0	-14.3	14.3
Elig(expd)-based Weight																
PSI(#10)													-4.8	-14.3	-28.6*	0.0
MPI(#7)														-9.5	-23.8*	4.8
BPL(#12)															-14.3	14.3
All(#44)																28.6*
CBT																
Observations	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.15: Differences between pairs of MTEs, Moderately poor households, Semi-urban Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		3.3	-6.7	-10.0	13.3	-10.0	-3.3	-16.7	-6.7	-10.0	-30.0*	-30.0*	-6.7	-20.0	-43.3**	-16.7
MPI(#7)			-10.0	-13.3	10.0	-13.3	-6.7	-20.0	-10.0	-13.3	-33.3*	-33.3*	-10.0	-23.3*	-46.7***	-20.0
BPL(#12)				-3.3	20.0	-3.3	3.3	-10.0	-0.0	-3.3	-23.3*	-23.3*	-0.0	-13.3	-36.7***	-10.0
PCA-based Weight																
PSI(#10)					23.3	0.0	6.7	-6.7	3.3	0.0	-20.0	-20.0	3.3	-10.0	-33.3**	-6.7
MPI(#7)						-23.3	-16.7	-30.0*	-20.0	-23.3	-43.3**	-43.3**	-20.0	-33.3*	-56.7***	-30.0*
BPL(#12)							6.7	-6.7	3.3	0.0	-20.0	-20.0	3.3	-10.0	-33.3**	-6.7
All(#44)								-13.3	-3.3	-6.7	-26.7*	-26.7**	-3.3	-16.7	-40.0***	-13.3
Log(expd)-based Weight																
PSI(#10)									10.0	6.7	-13.3	-13.3*	10.0	-3.3	-26.7*	0.0
MPI(#7)										-3.3	-23.3*	-23.3*	-0.0	-13.3	-36.7**	-10.0
BPL(#12)											-20.0	-20.0	3.3	-10.0	-33.3***	-6.7
All(#44)												0.0	23.3	10.0	-13.3	13.3
Elig(expd)-based Weight																
PSI(#10)													23.3*	10.0	-13.3	13.3
MPI(#7)														-13.3	-36.7**	-10.0
BPL(#12)															-23.3**	3.3
All(#44)																26.7*
CBT																
Observations	30	30	30	30	30	30	30	30	30	30	30	30	30	30	30	30

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.16: Differences between pairs of MTEs, Around median households, Semi-urban Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		-2.3	-7.0	-4.7	-16.3*	-8.1	-4.7	-9.3	-5.8	-9.3	-15.1*	-15.1*	-9.3	-12.8*	-24.4***	-8.1
MPI(#7)			-4.7	-2.3	-14.0*	-5.8	-2.3	-7.0	-3.5	-7.0	-12.8*	-12.8*	-7.0	-10.5	-22.1***	-5.8
BPL(#12)				2.3	-9.3	-1.2	2.3	-2.3	1.2	-2.3	-8.1	-8.1	-2.3	-5.8	-17.4***	-1.2
PCA-based Weight																
PSI(#10)					-11.6	-3.5	-0.0	-4.7	-1.2	-4.7	-10.5*	-10.5*	-4.7	-8.1	-19.8***	-3.5
MPI(#7)						8.1	11.6	7.0	10.5	7.0	1.2	1.2	7.0	3.5	-8.1	8.1
BPL(#12)							3.5	-1.2	2.3	-1.2	-7.0	-7.0	-1.2	-4.7	-16.3***	0.0
All(#44)								-4.7	-1.2	-4.7	-10.5*	-10.5*	-4.7	-8.1	-19.8***	-3.5
Log(expd)-based Weight																
PSI(#10)									3.5	-0.0	-5.8	-5.8	0.0	-3.5	-15.1**	1.2
MPI(#7)										-3.5	-9.3	-9.3	-3.5	-7.0	-18.6***	-2.3
BPL(#12)											-5.8	-5.8	-0.0	-3.5	-15.1**	1.2
All(#44)												0.0	5.8	2.3	-9.3*	7.0
Elig(expd)-based Weight																
PSI(#10)													5.8	2.3	-9.3*	7.0
MPI(#7)														-3.5	-15.1**	1.2
BPL(#12)															-11.6**	4.7
All(#44)																16.3**
CBT																
Observations	86	86	86	86	86	86	86	86	86	86	86	86	86	86	86	86

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

Table B.17: Differences between pairs of MTEs, Affluent households, Semi-urban Sector

	Fixed Weight			Weight based on PCA				Weight based on log(expd)-regression				Weight based on elig(expd)-regression				
	PSI (#10)	MPI (#10)	BPL (#12)	PSI (#10)	MPI (#10)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	PSI (#10)	MPI (#7)	BPL (#12)	All (#44)	CBT
Fixed Weight																
PSI(#10)		1.3	-6.7	-1.3	28.0***	-5.3	-6.7	-4.0	-4.0	-5.3	-8.0	-4.0	-2.7	-6.7	-6.7	-6.7
MPI(#7)			-8.0	-2.7	26.7***	-6.7	-8.0	-5.3	-5.3	-6.7	-9.3	-5.3	-4.0	-8.0	-8.0	-8.0
BPL(#12)				5.3	34.7***	1.3	-0.0	2.7	2.7	1.3	-1.3	2.7	4.0	0.0	0.0	-0.0
PCA-based Weight																
PSI(#10)					29.3***	-4.0	-5.3	-2.7	-2.7	-4.0	-6.7	-2.7	-1.3	-5.3	-5.3	-5.3
MPI(#7)						-33.3***	-34.7***	-32.0***	-32.0***	-33.3***	-36.0***	-32.0***	-30.7***	-34.7***	-34.7***	-34.7***
BPL(#12)							-1.3	1.3	1.3	0.0	-2.7	1.3	2.7	-1.3	-1.3	-1.3
All(#44)								2.7	2.7	1.3	-1.3	2.7	4.0	-0.0	-0.0	0.0
Log(expd)-based Weight																
PSI(#10)									-0.0	-1.3	-4.0	0.0	1.3	-2.7	-2.7	-2.7
MPI(#7)										-1.3	-4.0	0.0	1.3	-2.7	-2.7	-2.7
BPL(#12)											-2.7	1.3	2.7	-1.3	-1.3	-1.3
All(#44)												4.0	5.3	1.3	1.3	1.3
Elig(expd)-based Weight																
PSI(#10)													1.3	-2.7	-2.7	-2.7
MPI(#7)														-4.0	-4.0	-4.0
BPL(#12)															-0.0	0.0
All(#44)																-0.0
CBT																
Observations	75	75	75	75	75	75	75	75	75	75	75	75	75	75	75	75

Notes: Shaded cells contain targeting methods used in the main analysis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at household level.

B. Ethnic Favoritism Analysis

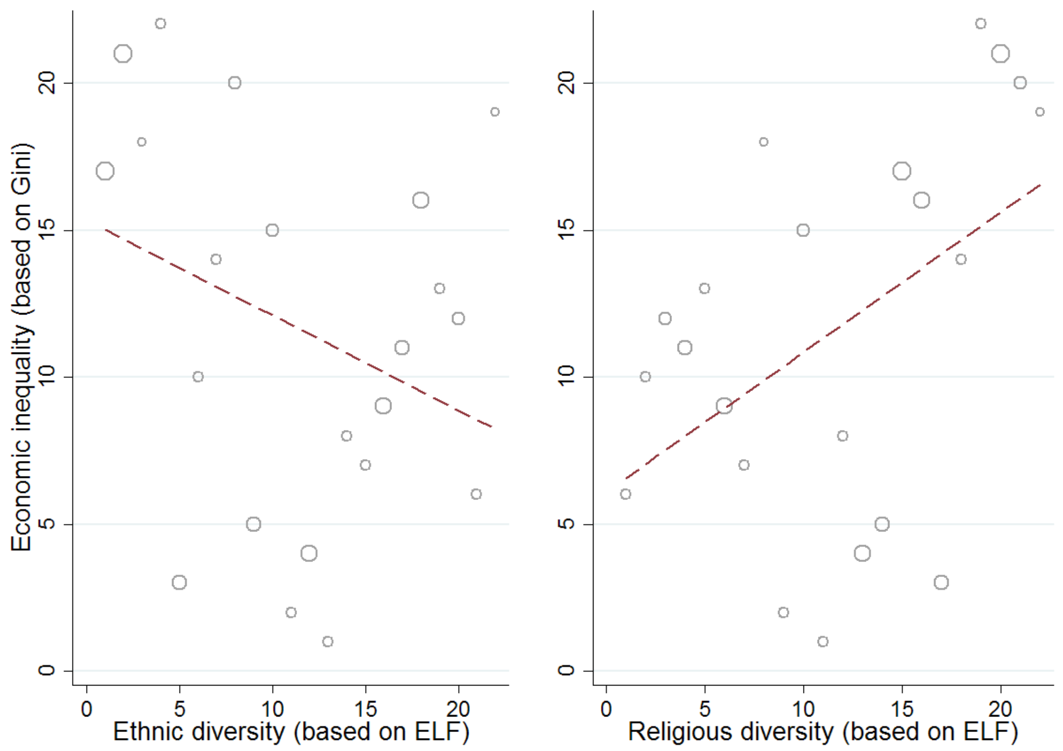

Figure C.1: Fractionalization and economic inequality for semi-urban communities

Notes: Grew circles indicate an community's relative population size. Values for ethnic diversity and economic inequality are the community ELF and Gini ranks, respectively. The Gini index is based on a PCA-based asset index comprising 28 variables.

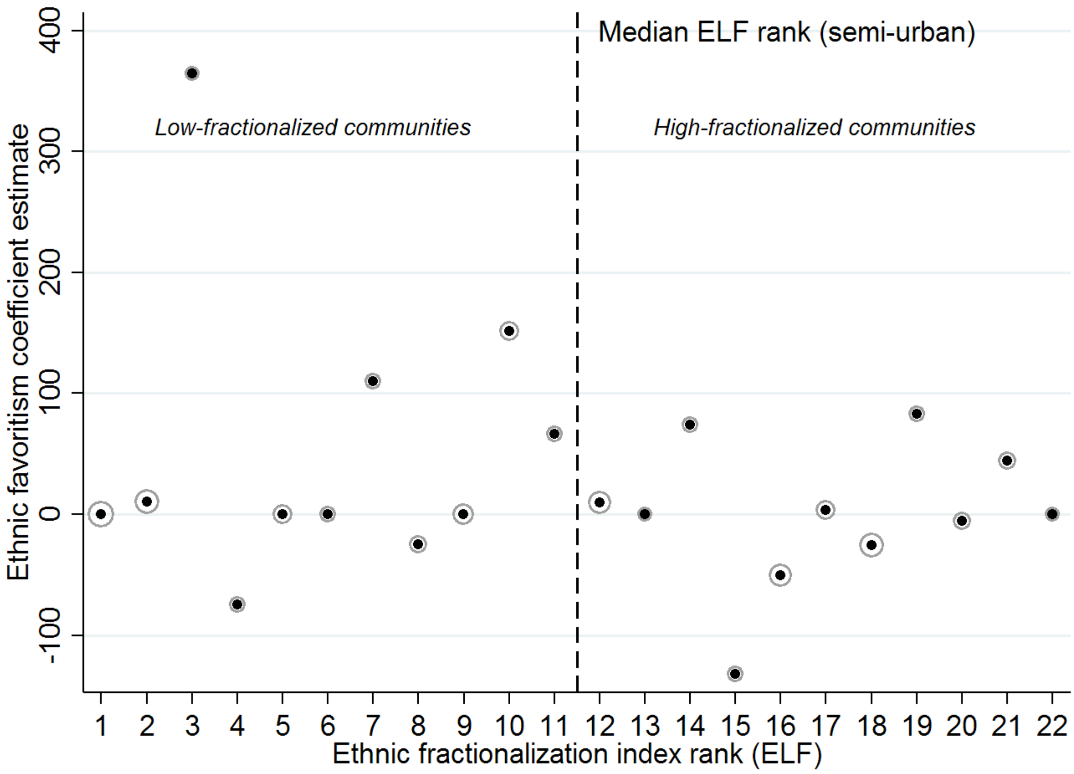

Figure C.2: Ethnic favoritism in the semi-urban committee decision by community

Notes: The graph depicts community-wise coefficient estimates of ethnic favoritism (particularly, $\hat{\beta}_1$ in equation 4.1) ordered by their rank in ethnic fractionalization. Fractionalization is based on the ethno-linguistic index of fractionalization (ELF). Grew circles around the dots indicate relative community size. For these community-wise coefficients I cannot reject the hypothesis of having zero value coefficients at conventional significance levels.

Table C.1: Distribution of six ethnic and four religious groups at the community-level

	Rural				Semi-urban		
	Average population share per community (1)	Relative frequency of being majority (2)	Relative frequency of being minority (3)	Rel. frequency of belonging to village leader group (4)	Average population share per community (5)	Relative frequency of being majority (6)	Relative frequency of being minority (7)
Ethnicity							
Dafin/Marka	0.48	0.54	0.11	0.54	0.33	0.27	0.00
Bwaba	0.20	0.21	0.25	0.29	0.12	0.14	0.27
Mossi	0.14	0.11	0.04	0.04	0.16	0.14	0.05
Peulh	0.13	0.11	0.18	0.11	0.09	0.14	0.27
Samo	0.04	0.04	0.25	0.04	0.25	0.32	0.00
Other	0.02	0.00	0.18	0.00	0.05	0.00	0.41
Religion							
Muslim	0.65	0.75	0.18	0.11	0.71	0.82	0.00
Catholic	0.23	0.21	0.21	0.61	0.25	0.18	0.18
Protestant	0.04	0.00	0.25	0.00	0.02	0.00	0.59
Animist	0.08	0.04	0.21	0.25	0.01	0.00	0.23
Other	0.00	0.00	0.14	0.04	0.00	0.00	0.00
Number of communities	28	28	28	28	22	22	22

Notes: The table depicts mean values that are calculated at the community level. Column 4 is based on additional information on the ethnic and religious affiliation of each community's traditional leader. There are no appointed traditional leaders for semi-urban neighborhoods. Majority and minority groups are defined as the group with the biggest and smallest population share, respectively.

Table C.2: Numeric values to calculate the distortion due to favoritism

	Rural communities			Semi-urban communities		
	All	**High-fract.**	**Low-fract.**	**All**	**High-fract.**	**Low-fract.**
Ethnicity	0.638	0.862	0.287	1.289	1.254	1.197
Religion	0.873	0.871	0.866	0.791	0.763	0.784

Notes: Each cell reflects the calculated value of the fraction $\frac{(1-p)p}{q}$ in equation 4.5. I obtain 12 different values by differentiating between ethnically and religiously represented sample households, as well as across all, high-fractionalized, and low-fractionalized communities.

Table C.3: Robustness checks

	Being finally targeted poor*100														
	Panel A: Main specification			*Panel B: Ethnic and religious favoritism estimated separately*						*Panel C: Probit estimation*			*Panel D: Alternative fractionalization measure*		
	Pooled (1)	Rural (2)	Urban (3)	Pooled (4)	Rural (5)	Urban (6)	Pooled (7)	Rural (8)	Urban (9)	Pooled (10)	Rural (11)	Urban (12)	Pooled (13)	Rural (14)	Urban (15)
Ethnicity															
HH ethnicity is reprsn. (in fract. commun.)	4.46***	1.38	6.40***	4.54***	1.72	6.34***				0.22***	0.08	0.33***	4.67**	1.49	8.61***
	(1.57)	(2.49)	(2.20)	(1.54)	(2.43)	(2.20)				(0.08)	(0.12)	(0.11)	(1.87)	(2.72)	(2.95)
Ethn. treatment * High majority-share													-0.95	-1.64	-4.91
													(2.83)	(4.45)	(4.12)
Religion															
HH religion is reprsn. (in fract. commun.)	-0.65	0.47	-1.87				0.18	0.70	-1.79	-0.01	0.02	-0.05	-1.24	0.03	-2.08
	(1.58)	(2.04)	(2.81)				(1.56)	(1.98)	(2.79)	(0.08)	(0.09)	(0.15)	(1.67)	(2.15)	(3.01)
Rel. treatment * High majority-share													5.22	10.15	0.62
													(4.07)	(7.35)	(5.07)
Controls (50)	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Community Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Ethnicity Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Religion Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Number of households	4550	2915	1635	4550	2915	1635	4550	2915	1635	4550	2889	1622	4550	2915	1635
Number of communities	50	28	22	50	28	22	50	28	22	50	28	22	50	28	22
Number of key informants	150	84	66	150	84	66	150	84	66	150	84	66	150	84	66

Notes: * p < 0.1, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses. Urban refers to the subsample of households from semi-urban communities. High Majority-Share: equals one if the share of a community's ethnic or religious majority group is below the median share of all communities.

Table C.4: Alternative specification for ethnic and religious favoritism across fractionalization

	Bein finally targeted poor$*$100		
	Pooled	Rural	Semi-urban
Ethnicity			
HH ethnicity is represented	3.14	-4.18	-1.87
	(3.18)	(6.90)	(4.52)
Ethnicity represented * ELF (ethnicity)	0.08	0.26	0.69**
	(0.17)	(0.30)	(0.34)
Religion			
HH religion is represented	0.61	2.24	-0.81
	(4.24)	(7.73)	(6.11)
Religion represented * ELF (religion)	-0.06	-0.07	-0.08
	(0.20)	(0.32)	(0.38)
Controls (50)	YES	YES	YES
Community Fixed Effects	YES	YES	YES
Ethnicity Fixed Effects	YES	YES	YES
Religion Fixed Effects	YES	YES	YES
Number of observations	4550	2915	1635
Number of households	4550	2915	1635
Number of communities	50	28	22
Number of key informants	150	84	66

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

C. Subsidy Evaluation

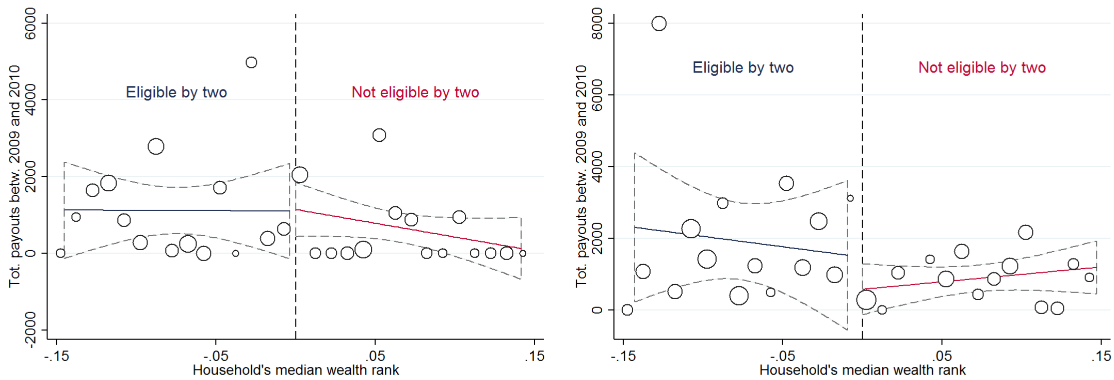

(a) Rural (b) Urban

Figure D.1: Adverse Selection: Health care utilization among insured households

Notes: Dots represent average outcome values by household median wealth percentile. Dot size indicates the relative number of observations per percentile. The forcing variable on the horizontal axis is based on three community wealth rankings, centered at the beneficiary cutoff, CDF-transformed and bounded between −0.2 and 0.8.

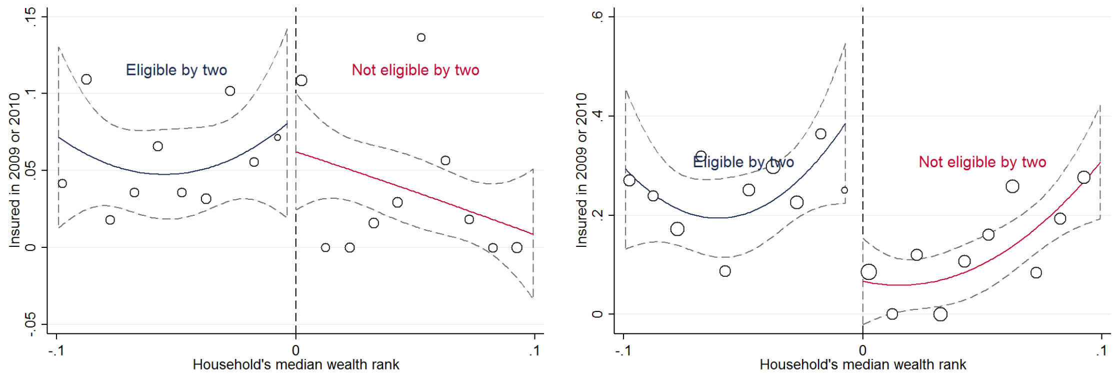

(a) Rural (b) Urban

Figure D.2: Local polynomial regression for health insurance enrollment

Notes: Dots represent average outcome values by household median wealth percentile. Dot size indicates the relative number of observations per percentile. The forcing variable on the horizontal axis is based on three community wealth rankings, centered at the beneficiary cutoff, CDF-transformed and bounded between −0.2 and 0.8. Fitted regression lines are based on the second-order polynomial regression equation 5.3.

Table D.1: First-stage regression results

	Eligible for subsidy in 2009					
	Rural (N=5692)			Semi-urban (N=2690)		
	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1
	(1)	(2)	(3)	(4)	(5)	(6)
Majority-eligible in 2009	0.839***	0.859***	0.885***	0.861***	0.856***	0.857***
	(0.031)	(0.036)	(0.044)	(0.042)	(0.049)	(0.059)
Median wealth rank	-0.264	0.192	1.070	-0.584	-0.675	-0.776
	(0.252)	(0.388)	(0.701)	(0.360)	(0.544)	(0.956)
Majority-eligible * Median wealth rank	0.181	-0.197	-0.990	0.601*	0.673	0.871
	(0.241)	(0.355)	(0.642)	(0.328)	(0.476)	(0.800)
Constant	0.146***	0.129***	0.106***	0.136***	0.140***	0.143***
	(0.027)	(0.030)	(0.034)	(0.039)	(0.044)	(0.052)
Mean of eligible	0.502	0.543	0.551	0.495	0.520	0.516
Community FEs	NO	NO	NO	NO	NO	NO
Observations (number of observations)	2088	1675	1143	1038	809	548
Observations left of cutoff	987	862	603	507	420	292
Observations right of cutoff	1159	871	598	570	428	295

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. The corresponding number of observations is given in the bottom panel, while the header contains the overall number of observations. Median wealth rank of household h is the median value of three community-based wealth ranks. A household is majority-eligible when its median wealth ranks is below the community-specific beneficiary cutoff.

Table D.2: Heterogeneous effects of the subsidy offer with respect to household wealth

	Insured in 2009 or 2010					
	Rural (N=5692)			Semi-urban (N=2690)		
	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1
	(1)	(2)	(3)	(4)	(5)	(6)
Majority-eligible in 2009 (at $\tilde{w}$)	0.000	-0.004	-0.002	0.230***	0.259***	0.318***
	(0.029)	(0.034)	(0.044)	(0.066)	(0.078)	(0.100)
Majority-eligible in 2009 * $\tilde{w}$	-0.203	-0.283	-0.362	0.092	0.220	0.573
	(0.166)	(0.197)	(0.247)	(0.342)	(0.385)	(0.453)
Median of community wealth at cutoff ($\tilde{w}$)	-0.013	0.024	0.089	0.449**	0.415**	0.286*
	(0.121)	(0.153)	(0.211)	(0.183)	(0.183)	(0.156)
Median wealth rank	-0.316*	-0.390	-0.499	0.312	0.827	1.270
	(0.174)	(0.276)	(0.541)	(0.312)	(0.525)	(0.928)
Majority-eligible * Median wealth rank	0.317	0.389	0.664	0.992	0.647	1.514
	(0.266)	(0.380)	(0.651)	(0.644)	(0.957)	(1.755)
Constant	0.055***	0.058***	0.063**	0.092***	0.072**	0.053
	(0.019)	(0.022)	(0.028)	(0.029)	(0.032)	(0.038)
Community Fixed Effects	NO	NO	NO	NO	NO	NO
Observations (number of households)	2088	1675	1143	1038	809	548
Observations left of cutoff	987	862	603	507	420	292
Observations right of cutoff	1159	871	598	570	428	295
Mean of 'being insured'	0.041	0.041	0.046	0.157	0.168	0.181
25th percentile of commun. wealth at cutoff (w)	0.162	0.162	0.162	0.216	0.216	0.216
75th percentile of commun. wealth at cutoff (w)	0.321	0.321	0.321	0.361	0.361	0.361

* p < 0.1, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses. w_c is the wealth percentile of a community c's 'cutoff household' and $\tilde{w}_c = w_c - median(w_c)$. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. The corresponding number of observations is given in the bottom panel, while the header contains the overall number of observations. Median wealth rank of household h is the median value of three community-based wealth ranks. A household is majority-eligible when its median wealth ranks is below the community-specific beneficiary cutoff.

Table D.3: Placebo Test: 2009 eligibility and enrollment in 2007/2008

	Insured in 2007 or 2008					
	Rural (N=5978)			Semi-urban (N=2620)		
	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1
	(1)	(2)	(3)	(4)	(5)	(6)
Majority-eligible in 09	-0.036	-0.032	-0.034	-0.035	-0.031	-0.025
	(0.023)	(0.028)	(0.035)	(0.049)	(0.056)	(0.068)
Median wealth rank	-0.227	-0.230	-0.532	0.070	0.141	0.108
	(0.166)	(0.277)	(0.504)	(0.359)	(0.576)	(1.060)
Majority-eligibility * Median wealth rank	0.046	0.134	0.637	-0.510	-0.552	-0.364
	(0.249)	(0.385)	(0.668)	(0.531)	(0.803)	(1.439)
Constant	0.068***	0.068***	0.076***	0.110***	0.107***	0.106**
	(0.016)	(0.019)	(0.025)	(0.033)	(0.038)	(0.047)
Mean insurance status	0.046	0.049	0.052	0.112	0.106	0.109
Community FEs	NO	NO	NO	NO	NO	NO
Observations (number of observations)	2280	1772	1224	1036	810	578
Observations left of cutoff	1008	818	586	516	426	312
Observations right of cutoff	1324	1006	690	558	422	304

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. The corresponding number of observations is given in the bottom panel, while the header contains the overall number of observations. Median wealth rank of household h is the median value of three community-based wealth ranks. A household is majority-eligible when its median wealth ranks is below the community-specific beneficiary cutoff.

Table D.4: Balancing test with six variables for semi-urban households in 2009

	HH size		Female HH share		Dafin HH		Muslim HH		Educ. HH member		Agric. HH share	
	h=0.15	h=0.1	h=0.15	h=0.1	h=0.15	h=0.1	h=0.15	h=0.1	h=0.15	h=0.1	h=0.15	h=0.1
Majority-eligible in 2009	-0.503	-0.319	-0.063	-0.069	0.031	-0.027	0.007	0.000	0.023	-0.058	-0.035	-0.073
	(0.823)	(0.990)	(0.056)	(0.070)	(0.096)	(0.113)	(0.061)	(0.074)	(0.083)	(0.106)	(0.047)	(0.056)
Median wealth rank	15.053*	17.641	-0.634	-1.134	0.546	-0.694	-0.320	-0.893	-0.265	-1.095	-0.216	-0.255
	(8.208)	(13.896)	(0.420)	(0.781)	(0.832)	(1.428)	(0.438)	(0.731)	(0.596)	(0.993)	(0.340)	(0.572)
Majority-eligible * Median wealth rank	-6.334	-4.579	0.574	1.342	-0.464	0.016	0.662	1.435	1.688*	0.860	-0.615	-1.674
	(10.574)	(17.180)	(0.688)	(1.232)	(1.169)	(2.058)	(0.727)	(1.356)	(0.989)	(1.666)	(0.578)	(1.039)
Constant	9.342***	5.988***	0.600***	0.684***	0.236*	0.771***	1.012***	0.996***	0.815***	0.615***	0.276***	0.086*
	(1.198)	(2.123)	(0.104)	(0.131)	(0.130)	(0.187)	(0.032)	(0.043)	(0.121)	(0.207)	(0.061)	(0.051)
Community FEs	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Mean of outcome with subsidy	5.983	6.269	0.477	0.464	0.426	0.403	0.836	0.845	0.749	0.806	0.235	0.245
Mean of outcome without subsidy	8.257	7.839	0.485	0.489	0.440	0.430	0.852	0.846	0.833	0.826	0.222	0.216
Observations	388	262	314	215	411	278	411	278	411	278	411	278

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. The corresponding number of observations is given in the bottom panel, while the header contains the overall number of observations. Median wealth rank of household h is the median value of three community-based wealth ranks. A household is majority-eligible when its median wealth ranks is below the community-specific beneficiary cutoff.

Table D.5: Intent-to-treat effect on enrollment with local polynomial regression

	Insured in 2009 or 2010					
	Rural (N=5692)			Semi-urban (N=2690)		
	h=0.2	h=0.15	h=0.1	h=0.2	h=0.15	h=0.1
	(1)	(2)	(3)	(4)	(5)	(6)
Majority-eligible in 2009	0.003	0.008	0.012	0.322***	0.346***	0.302**
	(0.043)	(0.050)	(0.059)	(0.091)	(0.104)	(0.129)
Median wealth rank	-0.498	-0.352	-0.688	2.302**	2.445	-0.106
	(0.570)	(0.849)	(1.339)	(1.152)	(1.622)	(2.674)
Median wealth rank^2	1.442	0.072	5.060	-12.320*	-11.385	27.239
	(3.003)	(5.819)	(13.137)	(6.622)	(12.401)	(32.101)
Majority-eligible * Median wealth rank	0.509	0.464	0.882	0.522	1.555	2.510
	(0.630)	(0.796)	(1.114)	(1.363)	(1.835)	(3.170)
Majority-eligible * Median wealth rank^2	-1.391	0.549	-5.980	21.643*	32.850	-24.365
	(6.186)	(11.280)	(27.025)	(12.252)	(23.186)	(60.563)
Constant w./o comm FEs	0.058	0.057	0.066	0.056	0.047	0.068
Mean of insured	0.041	0.041	0.046	0.157	0.168	0.181
Community FEs	YES	YES	YES	YES	YES	YES
Observations (number of households)	2088	1675	1143	1038	809	548
Observations left of cutoff	987	862	603	507	420	292
Observations right of cutoff	1159	871	598	570	428	295

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. h is the window bandwidth around the beneficiary cutoff: for instance, $h = 0.2$ indicates that the regression only includes households that are located 20 wealth percentiles above or below the beneficiary cutoff. The corresponding number of observations is given in the bottom panel, while the header contains the overall number of observations. Median wealth rank of household h is the median value of three community-based wealth ranks. A household is majority-eligible when its median wealth ranks is below the community-specific beneficiary cutoff.