

Governance and the effectiveness of public health subsidies: Evidence from Ghana, Kenya and Uganda*

Rebecca Dizon-Ross Pascaline Dupas Jonathan Robinson[†]

January 16, 2017

Abstract

Distributing subsidized health products through existing health infrastructure could substantially and cost-effectively improve health in sub-Saharan Africa. There is, however, widespread concern that poor governance – in particular, limited health worker accountability – seriously undermines the effectiveness of subsidy programs. We audit targeted bednet distribution programs to quantify the extent of agency problems. We find that around 80% of the eligible receive the subsidy as intended, and up to 15% of subsidies are leaked to ineligible people. Supplementing the program with simple financial or monitoring incentives for health workers does not improve performance further and is thus not cost-effective in this context.

JEL codes: D73, H11, I15, I38

Keywords: leakage, extortion, shirking, motivation

*This study was funded by the NICHD (grant P01HD061315-01), the Hellman Fellows Fund, UCLA and Stanford University. Dupas gratefully acknowledges support from the National Science Foundation (grant 1254167) and the Sloan Foundation. The study protocols received approval from the Ghana Ethical Research Committee, the Ugandan National Council for Science and Technology, the Ugandan Joint Clinical Research Centre, Kenya's Ministry of Science and Technology (research permit office), the IRBs of MIT, UCLA, UCSC, Stanford, and IPA, and received foreign clearance from the NIH for the work in Ghana. The Ghana experiment was registered on the AEA RCT Registry (AEARCTR-0000331) and a pre-analysis plan for it was submitted to the JPAL Hypothesis Registry in December 2011; the data is publicly archived at the Harvard Dataverse (DOI: 10.7910/DVN/P04NKK). We thank Kwame Abrokwa, Jonathan Addie, Liz Schultz, Sarah Green, Catlan Reardon and Ishita Ahmed for outstanding field research assistance. We are grateful to Francisca Antman, Jessica Cohen, and Ben Olken for detailed comments and suggestions, as well as to numerous seminar and conference participants for insightful discussions. All errors are our own.

[†]Dizon-Ross: Booth School of Business, University of Chicago, rdr@uchicago.edu. Dupas: Department of Economics, Stanford University, and NBER, pdupas@stanford.edu. Robinson: Department of Economics, University of California at Santa Cruz, and NBER, jmrtwo@ucsc.edu.

1 Introduction

Free or highly-subsidized distribution is often advocated as a necessary strategy for rapidly increasing coverage of essential health products in rural areas of poor countries. There is widespread concern in the policy community, however, that agency issues and corruption in the health sector will compromise the success of these programs (UNDP, 2011). To implement subsidy programs, government agencies and other major players must be able to identify eligible beneficiaries and effectively deliver products to them, typically using existing health infrastructure.¹ Yet many public service providers are paid a fixed wage and are hard to fire, and thus may have little direct incentive to refrain from corruption or to exert the effort needed to effectively deliver services or products (World Bank, 2004).

These concerns are bolstered by existing evidence regarding agency issues in other types of public programs in developing countries. Anecdotal and survey evidence suggests that corruption is endemic in developing countries (see Svensson 2005 for a review of the survey evidence). Several quantitative studies document alarming levels of health worker absence (among others, see Chaudhury et al. 2006; Chaudhury and Hammer 2004; and Muralidharan et al. 2011).² Others show very high levels of mistargeting for non-health-related subsidy programs such as the Indian public distribution system (PDS) food subsidy.³ The policy community recommends that interventions to improve accountability in the health sector are needed (UNDP, 2011). Indeed, interventions that improve governance, such as performance incentives or increased monitoring, have been shown to have large effects on health care quality in Rwanda (Gertler and Vermeersch 2013) and Uganda (Björkman and Svensson 2009).

Despite the high levels of policy concern, however, there are still relatively few reliable estimates of the magnitude of the agency costs, and there is considerable heterogeneity in the existing estimates (Olken and Pande 2012). Moreover, much of the existing literature focuses on corruption in procurement, fiscal transfers, and food subsidies. There has been little prior research evaluating whether local government workers effectively implement public

¹The existing health infrastructure is normally used since it is viewed as having the potential to be the most cost-effective distribution method. This is particularly true for products for which medical diagnosis is necessary to identify eligibility (for example, ARVs or anti-malarials), but substantial cost-savings have also been documented for other products, like bed nets, where diagnosis is not required (De Allegri et al., 2010).

²These examples focus on nurses and other health workers, since they are the ones commonly doing distribution of the type we study. However, other work has documented poor performances of doctors as well (see Das, Hammer and Leonard 2008 and Das and Hammer 2014 for reviews).

³For example, Niehaus et al. (2013) show that 48% of households are misclassified for the PDS food subsidy in Karnataka, India, with 70% of ineligibles tagged as eligible and 13% of eligibles tagged as ineligible. Small bribes are also common, with 75% of households reporting paying prices above the statutory fee, and the mean bribe amount roughly USD 0.20.

distribution systems for health products, where the estimated returns to the subsidized products are often high enough that, as long as a sufficient level of coverage can be achieved, zero leakage is not required in order to justify the cost of implementing the subsidy. In this paper, we examine the effectiveness of public distribution for targeted subsidies for one such preventative health product in three countries in sub-Saharan Africa. We document that, while agency issues do exist, they are not as extensive as might be thought given some of the findings from studies of government programs in other settings. We then evaluate whether supplementing the program with simple governance interventions can mitigate the remaining agency problems and further improve performance.

Our setting is a subsidy program recommended by the World Health Organization (WHO) that provides free bed nets to those most vulnerable to malaria – pregnant women and their unborn children – through antenatal care clinics. There are three main ways that health workers may undermine this type of subsidy program. First, they may demand under-the-counter payments from eligible clients – *extortion*. Second, they may provide the product to ineligible people – *leakage*, which may increase the total cost of a subsidy program and may crowd out eligible recipients, who presumably have higher health returns from the subsidy. Third, health workers may slack off, for example by failing to attend work or monitoring inventories to prevent stockouts – *shirking*, the key potential downside of which is that it compromises coverage among the eligibles. Our goal is to assess the extent of leakage and extortion relative to the coverage achieved by the program.

A key measurement challenge in these types of settings is that agents who engage in petty corruption typically do not readily report doing so. To overcome this, we devised a suite of measures that include audits on health center registers, back-check surveys with prenatal clients, and decoy visits to communities and health centers. Together, these measures generate a comprehensive picture of the performance of health workers. Our study takes place in three countries, Ghana, Kenya, and Uganda, which vary in corruption. Out of 178 countries in the 2012 Transparency International Corruption Index (where the least corrupt country is Denmark at rank 1, and Somalia, Afghanistan and North Korea tie for last place), Ghana ranks 64th, Uganda 130th, and Kenya 139th. At the time of data collection, both the Kenya and Uganda governments were implementing free bed net distribution schemes for pregnant women enrolling for prenatal care. In Ghana, there was no such government program, but we set one up, and so could randomize several features of program governance.

Our primary finding is that health worker performance, while not perfect, is higher than might be expected ex ante. 70-90% of eligible subsidy recipients received the subsidy at the clinic. Extortion appears to be rare, as only 1.4% of eligible subsidy recipients were asked to pay bribes. Leakage does exist, however: comparing administrative records of bed net

deliveries with our coverage estimates, we estimate an upper bound on the leakage rate at clinics of 15% in Ghana. Leakage to outsiders, however, appears to be low: “mystery clients” (ineligible men from outside the village making decoy visits) were able to obtain a subsidized net from the health facility in only 4.7% of their visits. While our study does not directly measure leakage higher in the distribution chain (e.g before subsidized products reach health facilities), we perform back of the envelope calculations for the public distribution programs in Kenya and Uganda by comparing observed population coverage rates in our study areas to publicly available data on the total number of nets distributed by these programs, and estimate that total leakage in these two countries is smaller than the 15% upper bound estimated for Ghana.

While program performance exceeds that of many previously-studied settings, and does not appear to jeopardize the cost-effectiveness of the program, 70-90% coverage and 15% leakage still leave considerable room for improvement. Indeed, these levels would be considered unacceptable in many OECD countries. The secondary contribution of our paper is thus to experimentally test strategies to improve efficiency further. Specifically, we randomized four aspects of the distribution scheme in Ghana: (1) whether health centers were threatened with audits, (2) whether health workers received bonus pay for implementing the program, (3) the size of the initial stock of bed nets received by clinics, and (4) whether health workers were asked to distribute nets directly or, instead, distribute vouchers redeemable for free nets at local shops.

If poor governance were responsible for the gap between current and “OECD-level” performance, one might expect these reforms to close much of the gap. We find no effect, however, for the bonus compensation, the threat of audits or the stock size. The voucher scheme, which was intended to improve outcomes (by reducing the effort cost of implementing the program for health workers and also reducing their ability to charge a bribe – since people may be less willing to pay for a voucher than for an actual net), if anything had the opposite effect: it mildly worsened some aspects of performance. The fact that none of these simple tweaks improved provider behavior suggests that the remaining barriers to perfect implementation may be prohibitively costly to lift, and as such, bare-bone free distribution through existing systems may be the most cost-effective approach.

Given that many existing studies of other types of programs suggest agency issues are potentially more problematic than in our setting, it is natural to question whether our findings are generalizable to other similar programs. Originally, Ghana was intended to be the sole site for our study. Having seen the outcomes from Ghana, we chose to expand to two other countries with higher corruption indices with this exact question in mind. The fact that we found similar results in the other countries as well suggests that our findings

are more general to this type of program.

Another important consideration in generalizing these results is that we study compliance with a targeting rule that is easy to verify, making it harder to hide leakage. In practice, however, many subsidized products are targeted at easily identified subgroups such as pregnant women or young children; thus, assessing leakage levels for “easy” targeting rules is important as a standalone research question even if the external validity to other targeting rules is unknown.

A final consideration is that the willingness to pay for ITNs in rural Africa is modest: in 2007, only 25% of households in rural Kenya were willing to pay more than \$2 for a \$7-bednet (Dupas, 2014). Even though there is little scope to demand large bribes in this setting, there remains ample room for health workers to extort small amounts of money from the ineligible. In other contexts, public employees demand bribes even at a low level – for example, in a study conducted in the same time period as ours, Foltz and Opoku-Agyemang (2014) find that policemen in Ghana allow only 19% of trucks to pass roadblocks without taking a bribe, and the most frequent bribe amount paid is 1 Ghanaian Cedi (around 60 US cents) – an amount that 75% of rural Kenyan households are willing to pay for an ITN (Dupas, 2014). In fact, based on the demand curve estimated in Dupas (2014), the profit-maximizing price a monopolist health worker would want to charge ineligibles in Kenya is \$1.6 (in 2009 USD), a price 34% of households are willing to pay. They could also extort the eligible: based on the demand among pregnant women observed in Cohen and Dupas (2010), a health worker would maximize profit from the eligible by charging them 30 US cents (in 2009 USD).

Evidently, these sums of money are not enough to tempt health workers to accept bribes. And indeed, even when our mystery clients had a high willingness to pay, health workers still turned down the vast majority of their net requests. This suggests that, despite the lack of monitoring, health workers face some costs of deviating from distribution protocols. In the final section of the paper, we present survey evidence that health workers in all three countries appear positively selected in terms of other-regarding preferences, have a higher level of intrinsic job motivation, and feel more accountable than workers in other professions, three factors that each raise the cost of corrupt behavior. We also present data suggesting that the health workers in our sample performed well on other tasks unrelated to the bed net distribution program, which could also reflect this positive selection.

Would performance be worse on more highly valued products? While answering this question comprehensively is outside the scope of this study, for two of the countries, Kenya and Uganda, we have household survey data on a curative product, for which willingness-to-pay tends to be higher than prevention – specifically, we asked whether children who presented with malaria received the recommended Artemisinin-based Combination Therapy

(ACT) drugs as they should. We find a coverage rate of 92%, even higher than for the nets. However, we find some evidence that extortion for this product is more frequent: though the drugs should be free at public facilities, 14% of mothers report having had to pay for the drugs - a number much higher than our 2% for nets, although still substantially lower than the 75% found by Niehaus et al. (2013) for food subsidies or 81% found by Foltz and Opoku-Agyemang (2014) for trucks. While this figure may also be overstated because women conflated payment for the drugs with consultation fees, these potential bribe requests did not threaten the near-universal coverage rate (perhaps because health workers could price discriminate), suggesting that even with highly-valued products, public distribution can be effective.

The remainder of the paper is organized as follows. Section 2 presents the theoretical framework that guided our study design. Section 3 presents the study settings and Section 4 describes the data collection protocols used. Section 5 presents performance results from the three audit studies and discusses their implication for the cost-effectiveness of free distribution schemes. Section 6 presents the experimental results from Ghana. Section 7 provides evidence on motivation levels of health workers, and Section 8 concludes.

2 Theoretical Framework

The theoretical framework underlying our research questions and study design is based on Becker and Stigler (1974), Banerjee (1997) and Niehaus et al. (2013). We consider three types of agents: the government, a health worker, and the clients of a health clinic. The government has a set of publicly provided private health products that the clients need (in our case, insecticide-treated bed nets). The goal of the government is to allocate these health products in a way that maximizes social welfare. We assume that the health products are scarce in the sense that the number of people exceeds the number of products the government can afford. Individuals vary in their need for the products: we assume that there is a group of people that value the products more highly, and thus have a higher willingness to pay for the product (for example, the private health returns to using a bed net are higher for pregnant women than for the rest of the population).

In a frictionless environment, the market would deliver the efficient outcome that those who value the product most would receive it.⁴ However, this is not true if there exist credit constraints, which could create a wedge between willingness and ability to pay for the credit-constrained (Cohen and Dupas 2010, Tarozzi et al. 2014). Since targeting based

⁴The products could also have positive externalities in which case the market outcome would not be efficient; we ignore these here for simplicity but they do not affect our predictions.

on credit constraints is typically infeasible and/or costly, the government could adopt an alternative rule of targeting the population with high returns on average (pregnant women) for a specified subsidy level.

The allocation of the health products is the responsibility of the health worker at a health facility. The health worker has to expend effort to deliver the product, and effort is costly. Assume that the health worker can observe the eligibility of the person who comes to the facility (e.g. the health worker can observe if the person is a pregnant woman or not). However, the government cannot monitor the health worker's behavior: in particular, the government observes neither the actual price charged by the health worker for the product, nor the eligibility of the people to whom the health worker offers (or does not offer) the product.

The utility function of the health worker has four components: the cost of effort, a monetary payoff (the sum of the worker's salary and any side payments from corruption), and two types of non-monetary payoffs. If the monetary payoff is higher than in alternative jobs available, then the threat of losing the job generates "extrinsic motivation" to perform well. The first non-monetary payoff represents the health worker's satisfaction from clients receiving health benefits from the product. The *weight* the health worker puts on this non-monetary payoff is a function of her level of pro-social motivation, and the *magnitude* depends on her beliefs about the health impact of the tasks she has to perform. The second non-monetary payoff is a psychic payoff from doing one's job well, i.e. "intrinsic motivation." This psychic payoff goes to zero if the health worker breaks any rule on the job (i.e., gives the subsidy to an ineligible person or fails to give the subsidy to an eligible person for the specified price) in order to reduce her effort or increase her monetary payoff.

In this set-up, the health worker can engage in three forms of corruption, defined broadly as "incidents where a public employee breaks a rule for private gain" (Olken and Pande, 2013) (note that we take "private gain" here to exclude the non-monetary utility payoffs). In each case, she will choose to be corrupt if the increase in monetary payoff (or decrease in effort) exceeds the reduction in non-monetary payoffs:

- *Extortion—extracting rents from the eligible*: Since the health worker has local market power (since other health facilities distributing the product are located some distance away), the health worker can charge eligible people a price higher than the official subsidized price. Extortion reduces social welfare in two ways: (1) if demand is price-sensitive and health workers are unable to price discriminate, extortion will lower the number of recipients; and (2) part of the government subsidy is captured by the health worker, for whom the marginal utility of cash is likely lower than for the eligible individual. In the contexts we consider, willingness to pay absent access to credit is

likely low (see Dupas and Miguel 2016 for a review), and so health workers will not be able to demand large bribes for ITNs. The potential for small bribes still exists, though: in Kenya, over 75% of households are willing to pay 60 cents (in 2009 USD) for an ITN (Dupas, 2014) and 80% of pregnant women are willing to pay 30 cents (Cohen and Dupas, 2010), meaning that both bribe taking from the ineligible and extortion from the eligible is possible. Based on this Kenya data, the profit-maximizing price a monopolist health worker would want to charge is \$1.6 (in 2009 USD) to the ineligible and 30 cents to the eligible. If the monopolist could price-discriminate, the prices charged would vary across the population, but the general point still holds: although willingness to pay is low, it is not zero, and so a health worker only concerned with maximizing profits would charge positive prices for many eligible and ineligible individuals.

- *Leakage—diverting the subsidized product to the ineligible*: Some ineligible people may have a higher ability to pay than some eligible people (due for example to credit constraints). It could then be profit-maximizing for the health worker to allocate some of the products to ineligible individuals, which could crowd out receipt by the eligible if the total number of subsidies is limited. Leakage reduces social welfare and leads to inefficiency (from a public health perspective) if there is a wedge between willingness and ability to pay such that those ineligible individuals who purchase the product have lower health returns to the product than those eligible individuals who decide not to. Note that high demand for nets at low price points could translate into frequent requests for free or cheap ITNs from the ineligible.
- *Shirking—not implementing the program*: the health worker may not pay the effort cost associated with delivering the subsidy to eligible clients, even though it is part of her job description. Shirking can reduce social welfare by reducing coverage rates among the eligible, assuming the social cost (health worker effort) of delivering the product is outweighed by the social benefit, as it should be for products like bednets with large benefits.

In this set-up, higher levels of extrinsic, pro-social and intrinsic motivation unambiguously lead to lower levels of extortion and shirking. However, pro-socially motivated health workers could break the targeting rule *for public gain* – namely, if they have private information on how the health returns of receiving the product will vary across clients, the health workers may be able to implement a more efficient targeting rule than the blunt government rule. For example, health workers may deny the free bed net to pregnant women that are wealthy

enough to already have one; or grant a free bed net to poor ineligible men who have an unprotected child.

3 Study design

3.1 Study Sites

Our sample consists of 168 rural health facilities (72 in Ghana, 48 in Kenya and 48 in Uganda). We chose these three countries as follows. First, we picked Ghana as our experimental site since it did not have a bed net distribution program and therefore it was possible for the research team to implement one with randomized variation in program features. Second, and *after* observing relatively high performance in Ghana, we picked Kenya and Uganda because they both had government-led bed net distribution programs through prenatal clinics, are among the countries perceived as most corrupt according to Transparency International (TI), and because, as we shared our initial Ghana results, we received anecdotal reports that leakage was high in both these countries. We thus chose two additional countries where the *prior* was that leakage would likely be on the high side, hence likely providing estimates from the top end of the distribution for the costs of corrupt behavior on the part of health workers. Ghana, ranking much better on the TI corruption index, was perceived as providing an estimate from the bottom end of the distribution.⁵

In each country, health facilities were selected for inclusion based on a census conducted of all of the public and private health facilities in a given region/province.⁶ Primary inclusion criteria for health centers in the study were: (1) having an antenatal care clinic (ANC); and (2) being rural or semi-rural.⁷ Our final sample spans 21 districts in Ghana, 10 districts in Kenya, and 6 districts in Uganda.⁸

Table 1, Panel A presents some statistics from Demographic and Health Surveys on

⁵As noted by an anonymous reviewer, Uganda and Kenya may not be representative since they already had a distribution program at the time while other countries may lack implementation capacity to do this. However, according to the 2013 World Malaria Report, 34 of 44 countries with ongoing malaria transmission in sub-Saharan Africa had an ITN distribution program through antenatal clinics in place as of 2012, so our results remain relevant and interesting for a large part of the continent. .

⁶We sampled one of 10 regions in Ghana, one of 8 provinces in Kenya, and one of 4 regions in Uganda.

⁷For the Ghana sample, we had the following additional criteria: (3) having no other healthcare facilities within a 2 km radius, no hospitals within a 5 km radius, and not more than one other ANC within a 5 km radius; (4) having no free bed net distribution program currently in place (very few clinics had this); (5) having at least two stores within a 2 km radius willing to participate in the voucher scheme (only 6% of clinics were excluded by this criterion); and, (6) being accessible for net deliveries (less than 2% were inaccessible).

⁸At the time of sampling, the three countries had 170, 208, and 112 districts each, respectively. Ghana and Kenya have had administrative changes since then.

coverage with insecticide-treated bed nets (ITNs) in each of the three countries. While these surveys were not concurrent with our data collection (the DHS in Ghana was conducted a few months before our study, but the DHS in Kenya and Uganda were separated from data collection by several years), they are nevertheless likely indicative. These statistics suggest that large numbers of people do not have nets: in the latest DHS, between 17 and 47% of households had no net, and a much higher fraction (58 to 78%) did not have enough to cover all household members.

Table 1, Panel B presents average baseline characteristics for the 168 health centers in the study. The average health clinic has been operating for 16 years and 85% of clinics are public. Health centers in the study sample enroll around 28 new ANC patients every month on average, and receive 63 revisits by existing ANC patients. The average clinic has 2.9 health workers (trained nurses and/or midwives) in charge of ANC patients and roughly 50% of health centers conduct outreach visits (i.e., they go to remote communities and provide “mobile” ANC services there). Only 13% are located within 10 kilometers of a store selling bed nets, suggesting bed nets are not widely available outside subsidized distribution schemes.

3.2 Timeline

The data collection in Ghana took place between October 2011 and April 2012. At the time, there was no bed net distribution program through antenatal care clinics, nor any other distribution scheme, although the Ministry of Health in Ghana had done some limited distributions of bed nets through antenatal care clinics in the past so health workers were somewhat accustomed to this type of scheme.

In Kenya and Uganda, data collection took place between May and September 2013. Since 2009, national policy in Kenya is that all pregnant women are provided a free long-lasting insecticide treated net (LLIN) at their first antenatal care visit.⁹ In Uganda, such distribution is not yet a national policy, but funding from the President’s Malaria Initiative and USAID enabled such distribution in parts of the country, including our study area, starting in 2012 and ending in October 2013. Both public and private facilities were included in the distribution program in Uganda.

⁹Funding for this distribution program comes primarily from the Global Fund to Fight AIDS, Tuberculosis and Malaria, and from the President’s Malaria Initiative.

3.3 Ghana set-up

The 72 health facilities in the Ghana sample were invited to participate in an NGO program called SALI (“Saving Lives”). The program mimicked those ongoing in Kenya and Uganda, and consisted of distributing free Long Lasting Insecticide-Treated Nets (LLINs) to pregnant women during routine antenatal care visits. The program was approved by the Ghana Health Service and was implemented by antenatal care clinic staff (most of them midwives or nurses). We hired and trained SALI staff, whose job was to roll out the distribution program – namely, to visit health facilities to introduce the program, deliver bales of nets and train health workers on the eligibility criteria for the free net and on record-keeping. The SALI staff was completely unaware that an evaluation of the SALI program would be implemented, as were the health workers.¹⁰

The program was rolled out into the 72 clinics over a 7-week period, from mid-October to early December 2011. The distribution program was announced as a continuous scheme, with health centers (or shops when applicable) given instructions on how to get a new delivery of LLINs before their stock ran out. In practice, the program stopped abruptly in all study health centers in mid-March 2012 when the Ghana Health Service rolled out a separate (unannounced) free distribution scheme for LLINs. Given this, health centers in our sample were exposed to the SALI program for up to 150 days, with an average of 109 days. The effects we identify in Ghana are thus specific to a relatively new program, in contrast to our results in Uganda and Kenya which are from longer-lasting programs.

Prior to rolling out the program, we grouped the 72 clinics into 6 strata with comparable average characteristics (size, remoteness, and proximity to district borders). Within each stratum, we randomly assigned clinics into either of two distribution mechanisms: 48 health centers were assigned to “direct distribution” (as in the Kenya and Uganda government programs), and 24 health centers to a “voucher scheme” wherein the health workers would distribute vouchers that could then be redeemed for a free LLIN at a local store. Since the great majority of health centers did not have any store selling bed nets in their vicinity, we stocked one or two local stores (located within two kilometers of health center) with LLINs and instructed shopkeepers to give one free LLIN to anyone who came in with a voucher from the local clinic and a corresponding ANC registration card.¹¹ We overlaid randomized governance features of the program onto this basic split between direct distribution and

¹⁰As such, the protocol involved deception of research subjects (health workers). As per IRB requirements, health workers were “debriefed” on the true intent of the research study in the Fall 2013, after the study had been completed.

¹¹Shops in our program were not permitted to mark up the nets, so testing whether vouchers decrease stockouts was not a motivation for this treatment, although that could in principle be another advantage of vouchers.

vouchers. Those are described in section 6.

4 Data

In all three countries, we collected data through two completely independent teams unaware of each other’s existence. The first is a team of “mystery clients” (undercover enumerators) asked to do decoy visits to health centers and their surrounding communities. The second is a team of regular surveyors, who administered surveys to ANC clients, health workers, as well as other professionals. In Ghana, we also have administrative data from the SALI staff (itself completely independent from the two data collection teams), which kept program implementation records and also asked health centers to keep a log of program beneficiaries.

4.1 Decoy visits

We sent undercover enumerators to local communities, who were trained to perform two types of decoy visits:

(1) *“Mystery Client” visits*: To measure how well health workers respect the targeting rule, we arranged to have undercover enumerators visit clinics (and stores, in the case of the voucher treatment) to try to obtain a subsidized net. To make this measure clearly interpretable as an estimate of the percent of requests by ineligibles yielding nets, we chose to employ only men as mystery clients (since they are clearly ineligible for the program). After concluding the interaction and once out of sight, mystery clients recorded the details of their encounter with the health workers, including whether they were asked to pay a bribe. To minimize possible suspicion among health workers, mystery clients dressed casually and never visited the same health center twice. They were not asked to follow a specific script. Both the order in which mystery clients visited a health center and the timing (across and within days) were randomized.

In Ghana, we paid the mystery clients 5 GHC (about \$3, roughly half the retail price) for any bed net they were able to bring back from such decoy visits (in addition to a salary of \$15 per day). We scheduled 10 mystery visits per health center spread over the study period. This led to an average of 0.6 mystery visits per health center per week. In Kenya and Uganda, because the bed nets distributed were from the government programs rather than our own NGO program, we did not incentivize enumerators to pay bribes for bed nets. We scheduled 3 mystery visits per health center over the course of two months.

(2) *Informal community interviews*: In all three countries, enumerators spoke with a convenience sample of community members about whether bed nets were available in the

area and, if so, where, at what price, what the eligibility criteria were, whether they thought an ineligible person would be able to obtain a net, and whether they themselves had received a net at the health center.¹² To elicit truthful answers, surveyors posed as visitors and did not introduce themselves as enumerators in Ghana and Kenya. In Uganda, we were unable to obtain an IRB waiver to conduct these visits undercover so they were conducted by the regular survey team described below in Section 4.2. Again, no enumerator performed this activity in the same area twice. We polled around 18 community members per health center in each of the three countries.

An important point to note when interpreting this set of results is that all mystery clients were from areas outside the study area (so that local communities would not realize the true intent of the decoy visits). If leakage is lower to men than to ineligible women, if health workers behave differently with strangers than with local community members, or if local community members under-report leakage to strangers, the estimates will be a lower bound on the percent of requests by ineligibles that yield nets.

4.2 Regular survey data

We use three types of regular survey data.

(1) *ANC Client Back-checking surveys*: An important outcome is whether eligible clients received nets. To verify this, we hired regular teams of enumerators to survey ANC clients at their homes. We sampled the ANC clients to be surveyed as follows. In all three countries, health workers keep registers of the women who come in for antenatal visits. These ANC registers include a record of each visit, as well as some rudimentary contact information (typically just the woman’s name and the area she is from). In each facility, this register was used to randomly sample 20 pregnant women who had visited the facility for antenatal care in the previous 4 months.¹³ The survey team attempted to visit these women at their

¹²During each visit, the mystery clients were instructed to speak with three individuals at the local market, as well as three households in their homes. The questions about whether the community members themselves had obtained nets were only asked in Uganda and Kenya.

¹³These data are used to estimate coverage and extortion (but not leakage, which is estimated from the program’s administrative records). From the point of view of our estimation, the threats associated with this sampling strategy are of two sorts: (a) some women who truly visited the ANC may not have been listed on the register; (b) some “fake” ANC clients were listed to boost the numbers and justify a smaller number of nets left in inventory. If those who are not listed are those who are not given a net or are asked for bribes, then (a) would lead us to overestimate coverage and/or underestimate extortion. Our visits to ANC clients (and, in particular, the percentage of clients not found) allow us to estimate an upper bound on (b). That said, we consider it very unlikely that health workers in any of the three countries modified how they fill ANC registers in response to the bed net distribution programs, especially since (1) they likely did not expect anyone to use the ANC registers for audit purposes since they were asked to keep alternate records of bed net recipients, (2) the registers are formatted with one registrant per row (so all revisits are recorded on the initial row) which makes them very difficult to use for monitoring how many nets the clinic

homes and administer a short survey about their experience at the ANC clinic, their bed net ownership and usage, including whether they received a bed net from the local health center and at what price. A subset of them were also asked to play a dictator game (so that we have a basis of comparison for the health workers, see below). The survey team tracked 92% of women in Ghana, but only 71% in Kenya and 66% in Uganda. The rest could not be traced, typically because ANC registers contain almost no information on clients' addresses and therefore tracing women from outer villages was particularly difficult. The higher attrition rate in Kenya and Uganda likely reflects the fact that the samples in these two countries included semi-urban health centers, which have wider catchment areas. In the analysis, in one specification we assume those that could not be tracked did not receive a net, providing a lower bound on coverage.¹⁴

(2) *Surveys of Health Workers:* Health workers involved with prenatal clients (on average, 3.1 per clinic) were sampled for a survey that measured basic demographics, other-regarding preferences (including a dictator game) intrinsic motivation, and other personality traits. This survey was administered by our regular surveyor teams and took place after all other data collection exercises had been completed, since being surveyed on one's level of intrinsic motivation and altruism could temporarily affect on-the-job performance. We were not able to survey all health workers involved with prenatal clients due to official leaves, or because some health workers were too busy for a survey or simply absent. We successfully interviewed 89% of health workers in Ghana, 74% in Kenya and 70% in Uganda. In Kenya and Uganda, we also recorded attendance of ANC health workers, by conducting one (unannounced) spot check per clinic.

(3) *Surveys of Other Professionals:* To get some sense of how pro-social motivation and intrinsic and extrinsic job motivation differ between health workers and other people, our regular survey teams also conducted the same survey modules used with health workers with non-health workers, in particular, teachers, as well as (in Ghana and Uganda only) shopkeepers and microfinance agents.

4.3 Ghana administrative data

Since we set-up the distribution scheme ourselves in Ghana, we have administrative data on bed net deliveries by our SALI teams at each facility (timing and quantities delivered). We also have logs kept by health workers, in which they were asked to record some basic

should have distributed (i.e., the number of eligible clients who visited in a given timeframe). As such, our prenatal client survey sample is likely a representative sample of the population of prenatal clients.

¹⁴In Uganda and Ghana, 7% of the interviews were conducted with proxies because the sampled respondent was unavailable at the time of interview; in Kenya, 1% of interviews were conducted with proxies.

information (name, prenatal card number and address) for each person to whom they gave a bed net or voucher. The survey team attempted to visit these individuals at their homes and administered a short survey identical to the one administered to ANC clients. Surveyors were successful at reaching 94% of the individuals sampled for this survey. We use these records to estimate the share of nets unaccounted for, which provides an upper bound on the leakage rate. (It is an upper bound since health workers could forget to record eligible clients who received nets, and we may have failed to track truly eligible clients).

5 Performance Results From Three Countries

We first examine the performance of the standard program (bed net distribution at clinics). For this we exclude from the analysis the 24 health centers in Ghana with a voucher scheme, keeping 48 health centers from each of the three countries studied, and we present the overall mean as well as country-specific means for each outcome, and the p-values for tests of equality for each country-pair.

5.1 Do the Eligible Receive the Full Subsidy?

We first consider the effectiveness of the bed net subsidy programs in reaching intended beneficiaries. The results are presented in Table 2. Our interviews with women randomly sampled from the ANC registers reveal that 76% of intended beneficiaries received a net at their first prenatal visit as per program protocols (Table 2, column 1), and 80% received one at some point (column 3).¹⁵ When we exclude clinics with reported stockouts over the sample period, the share of the eligible population receiving the subsidized net at their first prenatal care visit reaches 81% (column 2). Only 1% of them were asked to pay something in exchange for the net.

There is some meaningful heterogeneity across countries. Coverage is significantly higher in Kenya (90%, column 2) than in the other two countries (Ghana - 77%; Uganda - 69%). When asked why they did not receive a bed net, around 9% of the Ugandan women (20% of those who ventured an explanation) mentioned that it was because they already had a net or were rich enough to afford one, suggesting that health workers may have been, to some

¹⁵These estimates exclude any ANC registrants who were sampled from the registers but we were unable to track. A very conservative lower bound approach would be to assume all of the registrants who we were unable to track and interview did not receive nets – note that we believe this is too conservative since the likelihood of tracking was based on factors that were plausibly orthogonal to likelihood of net receipt (in particular, we do not think that ANC clients we could not trace are “fake patients,” since facilities have no incentive to over-report ANC clients in their registers, as discussed in footnote 13). The conservative lower bound approach would yield a 64% coverage rate for Kenya, 71% for Ghana, and 46% for Uganda.

extent, targeting the subsidized nets to those needing them the most. Request for payment is also significantly higher in Uganda, although the absolute level (3%) is still low.

Errors of exclusion or efficient targeting? In Table 3, we present some descriptive evidence on whether incomplete coverage is likely to represent errors of exclusion (truly poor women who should have received a net) or the outcome of health workers exercising discretion over how to allocate the subsidy based on local information. We investigate whether, *within clinic*, prenatal clients which have a higher socio-economic status (proxied by education), and thus are more likely to be able to afford a net on their own, are less likely to receive a free net. We find suggestive but weak evidence that more educated women were less likely to get nets – an additional year of education is associated with a 0.7 percentage point reduction in the likelihood of receiving a net at the first visit (Table 3, panel A, column 1). This implies that a woman at the 75th percentile in the distribution of years of education (7.5 years) is 3.1 percentage points (13%) less likely to receive a free net at her first visit than a woman at the 25th percentile (3 years). Targeting appears more pronounced in Ghana (col. 2) and Uganda (col. 4) than in Kenya (col. 3), where the coverage rate was closer to 100%.¹⁶

Altruistic targeting towards needier women should only occur if health workers are concerned about the stock of nets running out, and should be more pronounced when the net stock is low. We shed light on this in Column 5, in which we limit the sample to ANC visits that took place within one month of a stockout. The magnitude of the effect increases, though the standard errors are too large to reject that the effect for this subsample is the same as in Column 1.

5.2 Leakage to the Ineligible

Success rates of Mystery Client requests Table 4 shows the results of the decoy health center visits by our “mystery clients” trying to obtain nets for which they were ineligible, and for the community interviews. We first note that on 20.5% of the visits in Kenya and Uganda, the clinic was actually out of stock (we know this from independent visits made by the survey team – see column 1). Stockouts were not measured independently in Ghana, but mystery client reports suggest it was rarer (4.6%, column 2). We thus focus on non-stockout visits to make meaningful comparisons across countries, but we note that the stockout number used in Ghana is likely an upper bound, since health workers could have told ineligibles asking

¹⁶Outcomes from the first visit are likely more reflective of health worker targeting than outcomes across all visits (for example, more-educated people who did not get nets might return to the clinic specifically for a net). The results across all visits, shown in Panel B of Table 3, are consistent but somewhat weaker, indicating that there is some “correction” at visit 2.

for nets that the clinic was out of stock as a way to gently deny their requests.

Only 4.7% of the 685 non-stockout attempts made across all three countries were successful (Table 4, column 3). This varied from 11% in Uganda to 8.7% in Kenya to only 2.2% in Ghana. All of the nets leaked to mystery clients in Kenya and Uganda were given out for free (column 4), which is not surprising since the mystery clients in these two countries were not incentivized to pay for the (government sponsored) nets. What is more surprising is that while mystery clients in Ghana had a higher reservation price – recall that we paid them \$3 per (NGO sponsored) net they successfully acquired from health centers– they were less successful at obtaining nets than their counterparts in Kenya and Uganda. Only 1.3% obtained a net by bribing; and only 0.9% got a net for free. Requests for payments, shown in column 5, were rarer in Kenya (3.6%) and Uganda (1.1%) than Ghana (5.1%). However, the average amount requested in Ghana was very high (more than the full price of the net, even after bargaining had happened – column 6), and actual transactions occurred after only a quarter of the payment requests, possibly suggesting that some health workers may have been asking for a high price as a way to get rid of the visitor, or that they intended to replace the leaked net with a purchased net.

The informal interviews mystery clients conducted with randomly sampled community members (columns 8-10 of Table 4) also suggest modest levels of leakage. In all three countries, about 10% of community members thought a male (and so obviously ineligible) client could get a net from the local facility. When asked if they themselves had acquired a net, less than 4 percent of men said they had (note that some of these may have been received legitimately while taking a pregnant wife or child under five to the facility).

Errors of inclusion or efficient targeting? In Table 5, we provide some suggestive evidence that much of this leakage is consistent with health workers bending the targeting rule for the public good. In particular, we test whether health workers are more likely to leak nets to mystery clients whose “narrative” made them appear like they had a higher return to bed net usage, for example because they had a pregnant wife or vulnerable child. This analysis is limited by several factors (which is why we take them as suggestive), notably that the mystery clients’ narratives were not randomly assigned and so could be endogenous to the interaction or correlated with other mystery client characteristics, such as their power of persuasion, and also that the narratives were only stories and so obviously not truly tied to need. In addition, there were systematic cross-country differences in the strategies used, so that the regressors in Table 5 vary across countries.¹⁷

¹⁷In Ghana, mystery clients almost always said that the net would be for personal use (a few of them mentioned a pregnant wife), while in Uganda and Kenya, some mystery clients mentioned a sick child at home, though in Kenya they only did so if asked by health workers

That said, the results in Table 5 are consistent with benevolent leakage to children: mystery clients who mentioned a child were significantly more likely to obtain a free net. Mentioning a pregnant wife was not effective, often because health workers responded by asking the mystery client to bring his wife to the next visit.¹⁸

Total Local Leakage Rate (Ghana only) To get a sense of the likely magnitude for the leakage rate, we estimate an upper bound using the administrative records of bed net deliveries to health centers, as well as the records on beneficiaries kept by health centers, which we audited. We can do this exercise only in Ghana, the only country for which we have information on the total number of nets supplied in each clinic (we have this data since we implemented the program ourselves through the SALI team; despite our best efforts, it was not possible for us to obtain detailed governmental records of bed net deliveries to individual health centers in Kenya and Uganda). Combining all of the available information for Ghana, we estimate an upper bound for the percent of bed nets unaccounted for as follows: we compare N , the number of nets delivered by the program to the facility to N_e , the estimated number of nets that reached an eligible person.¹⁹ To estimate N_e , we use the log of beneficiaries kept by health workers at the request of the SALI program, and subtract invalid entries. First, all duplicate entries were considered invalid. Second, among non-duplicate entries, we estimate the share of invalid entries using the data from our random audits of the logs. Audited entries were considered invalid if the person named in the entry could not be found (6.4% of entries), was found and not eligible (1.5% of entries), or had not received a net (2.7% of entries). This estimate is an upper bound since some of those not found or not recorded in the administrative ledger may have been legitimate program beneficiaries.²⁰

Figure 1A shows the distribution across the clinics of the estimated (upper bound) number of bed nets unaccounted for, while Figure 1B shows the distribution of the estimated share. The average estimated leakage rate is 14.7%. The median number of bed nets unaccounted for is 20 and the 75th percentile is 37 nets.²¹

¹⁸Of course, since these narratives are only cheap talk, these results are not necessarily indicative of improved targeting.

¹⁹Data is only available for 47 of the 48 direct distribution clinics because, for one clinic, the survey team lost the surveys for respondents sampled from the SALI ledger, and so we cannot compute the “% of valid entries” for that clinic.

²⁰Note that we exclude audits where the person was found, eligible, and surveyed, but where the data about whether the person received a net is missing (most missing data are in surveys conducted with proxies, so these missing entries are often because the proxy did not know whether the respondent had received a net or not). If we assume instead that none of the respondents with missing data received nets (a very conservative assumption), the total leakage rate would increase to 16.8%.

²¹Note that the number of unaccounted-for nets can be negative either as a result of sampling error in the estimate of the percent valid entries (so we may underestimate the number of valid entries), or errors in the

Leakage Across the Entire Delivery Chain Our results so far concern leakage at the point of distribution (once the nets have reached health facilities). What about leakage higher up in the chain? We do a simple accounting exercise to estimate the total rate of leakage in the Uganda and Kenya programs, for which aggregate supply data is available online. The nets procured for distribution must be equal to the number of nets received by eligibles plus the number of nets leaked (unaccounted for).

In Kenya, 2,800,000 free ITNs were procured for the year 2013, to be distributed to an estimated eligible population of 2,837,475 eligible pregnant women and children under the age of one.²² This implies a maximum potential coverage rate of 98%. Our observed coverage rate of 91% among pregnant women therefore suggests minimal leakage (the rate would be 7% if our study area were representative of Kenya as a whole, though this is unlikely to be exactly correct). This is lower than the 15% upper bound for local leakage in Ghana.

In Uganda, the program we audited reports distributing a total of 268,804 bed nets to 34 districts for the October 2012 to September 2013 period. The total estimated ANC population over that period is 493,631.²³ Thus the program delivered enough bed nets to cover 55% of eligible pregnant women. Our estimated coverage rate is higher than that, at 66%. This could be because the districts we studied received more bednets per pregnant woman than average (perhaps because our study area is centrally located), or because women who registered for ANC for the first time during the program were more likely to receive nets than women who visited the ANC but had already registered.

The number of nets procured in both countries is thus very similar to our estimate of the number of nets received by eligibles. While this accounting exercise is not definitive (since our coverage estimates come from a single region in each country, while the supply figures are national), we took this evidence as strongly suggestive that leakage higher up is limited.²⁴

records kept by the SALI program on the number of nets they delivered (though we think this is unlikely). In any case it is reassuring that only three clinics have negative values, and the value is never less than -15 nets. The one clinic with -15 nets is likely due to sampling error, since 100% of the respondents sampled from the SALI ledger in that clinic were found, were eligible, and had received nets.

²²See http://www.pmi.gov/docs/default-source/default-document-library/malaria-operational-plans/fy12/kenya_mop_fy12.pdf?sfvrsn=6, Table 4 on p. 18, accessed August 26 2014.

²³See http://www.usaid.gov/sites/default/files/documents/1860/SMP_Year_5_Annual_Report-Final_Oct2012-Sep2013.pdf, p. 9, accessed August 26 2014.

²⁴If the region we study was favored by the government and received more nets per capita than others, our calculations would be underestimating leakage, as seems to be the case in Uganda. This is less of a concern in Kenya where the study took place in Western Province, which overwhelmingly favored the loser in the presidential elections of both 2007 and 2012. A second caveat is that our 48 facilities per country fall under the jurisdiction of only 10 districts in Kenya and 6 districts in Uganda; therefore our delivery chain leakage estimates reflect the performance of only 16 agencies.

5.3 Cost-Effectiveness Implications

Our administrative records from Ghana suggest a conservative upper bound on leakage at the facility level of around 15%. Leakage higher up in the chain is minimal, according to the estimates from Kenya and Uganda. Overall our upper bound estimate is comparable to the lower bound of 18% found by Olken (2005) for leakage of food in a food relief program in Indonesia, and on par with recent evidence on the allocation of government transfer benefits in Indonesia (Alatas et al. 2013).

Is 15% leakage high or low? Assuming no positive health impacts from leakage (so counting leakage as a pure deadweight loss – a strong assumption), 15% leakage implies an increase in the price per bed net delivered to an eligible person by $0.15/(1-0.15)=18\%$ with no change in total health benefits. Thus, except for products that are only marginally cost-effective to begin with, such leakage is unlikely to substantially affect cost-effectiveness. In the case of bednets, Cohen and Dupas (2010, Table IX) estimate costs between \$200 and \$662 per child life saved (depending on assumptions), which would increase to \$236-\$781 with 15% leakage. These remain orders of magnitude below the cost-effectiveness threshold of approximately \$20,000 per life saved (\$241 per disability-adjusted life year saved times 80 years) suggested by the 1993 World Development Report (World Bank, 1993).

6 The Ghana Experiment

Our results show that health worker performance does not threaten the viability of public distribution systems. Nevertheless, the imperfect coverage and the level of leakage we observe would likely be considered unacceptable in OECD countries – are there ways of further improving performance? We explore this possibility in this section, using our experimental results in Ghana.

6.1 Experimental design

In Ghana, where we implemented the program, we were able to randomly vary four aspects of the program to test specific hypotheses about what could improve performance. The hypotheses were outlined in our pre-analysis plan and built upon the theoretical framework presented in Section 2. Figure A1 shows the experimental design.

First, we randomized *direct vs. voucher distribution*. In the direct treatment, health workers had the nets on site to distribute; in the voucher treatment, health workers gave out vouchers which were redeemable at a local shop. The hypothesis here is (H1): the voucher could reduce extortion and leakage because people are less willing to pay a bribe for a voucher

than for the product itself, and could reduce shirking because the effort cost of handing out a voucher should be lower than a net (since vouchers are small and can be kept conveniently in the visitation room, and so take less effort to hand over than larger nets which must be retrieved from a storage room).

The other three randomized features were the following:

Staff monitoring (Audit vs. No Audit): Half of the clinics were randomly selected for an “audit treatment.” The audit treatment was rolled out in those clinics in January 2012 (recall that the program began in October 2011). Health workers were informed that the NGO implementing the program would perform audits, starting within the next month. How the audits would be performed was not disclosed, but health centers were warned that the program would be shut down if the audits revealed that either leakage or extortion had occurred after the audit announcement. The hypothesis here is (H2): performing a top-down audit (and shutting down the program if it reveals foul play) can reduce corruption if health workers are sufficiently pro-socially motivated to not want to deprive the community of the program, if health workers are concerned that their jobs (and corresponding monetary payoff) would be threatened if they are caught breaking program rules (extrinsic motivation), or if it is profit-maximizing for health workers to decrease corruption when being caught being corrupt threatens the continuation of the program.

Pay (Compensation vs. No Compensation): In clinics with direct distribution, where health workers had the responsibility to manage the LLIN inventory and coordinate with the SALI project staff to avoid stockouts, we randomly varied whether health workers received compensation for implementing the program. The compensation was a fixed monthly fee of 100 Ghana cedis (US\$60, corresponding to approximately 17% of the median monthly salary of a nurse or midwife or 25% of the median monthly salary of any healthworker) paid via direct deposit into health workers’ bank accounts. The hypothesis here is (H3): Increasing existing health workers’ compensation can reduce corruption if pro-social motivation, intrinsic job motivation, or effort have positive income elasticities, or if payments increase health workers’ pro-social and/or intrinsic job motivation as a form of gift exchange. Note that there is no selection margin here since we are increasing compensation *after* health workers select into the profession.

Scarcity (Small vs. Large Delivery): Finally, within direct distribution clinics, we randomly varied whether the stock of LLINs delivered to the health center at the onset of the program was high or low. Since clinics were instructed to call the SALI program officer to restock the nets whenever they would run out, this variation in the level of the initial stock should have no effect (besides increasing the potential for early stockouts) unless it affected the salience of the “budget constraint.” The hypothesis is (H4): Increasing the (perceived)

tightness of the budget constraint by giving out fewer nets to the clinic can reduce leakage if health workers believe that eligibles have higher returns (as they should under most targeting rules) and are sufficiently pro-socially motivated to care whether those people receive the subsidy.

To the best of our knowledge, this work represents the first experimental tests of hypotheses 1, 3, and 4.²⁵ Regarding hypothesis H2, Olken (2007) finds that top-down audits reduce corruption among local officials granting road-construction contracts, but bottom-up monitoring does not. In our context, misconduct by frontline providers is likely observable by the local community since the targeting rule is based on pregnancy status, so the scope for bottom-up monitoring is higher and the need for top-down audits might be lower.

6.2 Experimental results

We present the experimental results in Table 6. Because the audit threat was rolled out midway through the program, we use a difference-in-differences specification to determine the effects of audits, comparing the results before and after the time when audit threats were rolled out at clinics that were vs. were not sampled for the audit threat. Thus, we estimate regressions of the following form:

$$y_{ict} = \alpha + \beta_1 Voucher_c + \beta_2 Audit_c \times Post_t + \beta_3 Pay_c + \beta_4 SmallDelivery_c + \beta_5 Post_t + \delta Audit_c + \eta' X_c + u_{ict} \quad (1)$$

where y_{ict} is the outcome of individual i at health center c in period t , X_c is a vector of health center level controls (specifically: randomization strata fixed effects, baseline ANC attendance, and ANC staff size), and $Audit_c$, $Voucher_c$, Pay_c and $SmallDelivery_c$ are treatment dummies. $Post_t$ is a dummy equal to 1 if t is the post-audit period. We cluster the

²⁵To the best of our knowledge, there is no evidence on the first hypothesis in the literature. However, despite the lack of evidence, hypothesis H1 is a standard argument given by voucher proponents (Sexton, 2011). Regarding hypothesis H3, there is some evidence that higher public wages are correlated with lower corruption, but most of the evidence is cross-sectional (Van Rijckeghem and Weder, 2001; Rauch and Evans, 2000). Dal Bo, Finan and Rossi (2013) provide experimental evidence that higher wages improve selection into public service, but our experiment differs because it involves increasing the pay of existing health workers when asking them to take on a new task, thus shutting down the selection margin. Di Tella and Schargrodsky (2003) provide some evidence that higher-paid procurement officers respond more to corruption crackdowns in Argentina than lower-paid officers; however, their results could reflect factors other than wages and could run through channels other than corruption. Regarding hypothesis H4, we know from other contexts that local officials have and use local information to target scarce publicly-provided goods (i.e. Bardhan and Mookherjee 2000; Niehaus et al. 2013; Alatas et al. 2013; and Basurto, Dupas and Robinson 2016), but we do not know whether the quality of targeting increases with (perceived) scarcity.

standard errors at the health center level.^{26, 27}

We see little effect of most of the experimental treatments, except for the voucher scheme. While the voucher did not affect extortion among eligibles, it actually lowered eligible coverage (in contrast to our hypothesis that it would improve coverage): the first two columns of Table 6 (row 1) show that health workers were 19.5 percentage points less likely to deliver the subsidy to prenatal patients (17.8 percentage points less likely for the first visit).²⁸ We also find that mystery clients were 3.1% more likely to obtain a program net in voucher clinics than in direct clinics, a large increase relative to the base level of 2% in direct clinics (although only significant at the 10% level). However, when we look at the total local leakage rate estimated using administrative records, it was over 50% lower in voucher clinics than direct clinics (column 11), perhaps suggesting that the decrease in awareness/demand outweighed the increase in leakage conditional on solicitation.

There are two potential explanations for lower performance in the voucher clinics. The first is awareness: the percentage of community members who were aware of the nets program was 8.5 percentage points lower (off of a base of 72 percentage points) in the areas surrounding voucher clinics relative to direct clinics (col. 6). That voucher schemes might have lower awareness is intuitive: vouchers, which can fit in a pocket, have lower visibility than bulky bed nets.²⁹ Coverage might be lower because women did not know about the program, and leakage may have been higher because there was less community monitoring of health worker.

The second potential explanation for the voucher scheme's lower performance is its impact on the intrinsic job motivation of health workers. Qualitative evidence provided by the SALI program staff suggests that the voucher scheme lowered health worker autonomy and morale, because health workers felt as though they were not trusted. In fact, 2 out of 24 voucher clinics refused to implement the program at all, while 0 of the 48 direct distribution clinics

²⁶The treatments were assigned in cross-cutting fashion. We are powered to look at main effects (e.g., payment vs no-payment) but not interactions (e.g., payment in audit clinics vs no-payment in audit clinics). This of course means that when we look at the main effects of the various subtreatments (e.g., payment vs no-payment), we will be pooling across groups that received different other subtreatments (e.g., audit and non-audit) and cannot distinguish whether interaction effects drive the pooled effects we see.

²⁷Columns (2)-(5) of Table A1 test for balance across our experimental groups in Ghana, both in terms of baseline characteristics (Panel A) and program implementation details (Panel B). We regress each dependent variable on a dummy variable for being assigned to the voucher treatment, audit treatment, pay treatment, and large stock treatment. Columns (2)-(5) present coefficients and standard errors from these regressions. None of the differences are significant at the 5% level.

²⁸Because vouchers add an extra step to the process (pregnant women have to go to the store to redeem the voucher), incomplete redemption could also increase the coverage gap between voucher and direct clinics. In our sample, however, redemption rates are very high (over 95%). Thus, the gap in overall net receipts (19.9 percentage points) is only a little larger than the gap in product receipt at the clinic.

²⁹Of course, bed nets were also being distributed by the stores in the voucher scheme, but community awareness of what happens in a few stores may be lower than awareness of what happens in clinics.

refused.³⁰

Turning to the other randomized program features, the audit threat did not affect overall performance. One caveat to the audit results is that, in our audits, health workers' jobs were not on the line, and so potentially audits that threaten jobs could have had larger effects. However, audits that do not threaten employment could still have large effects for all the reasons listed above, and, indeed, such audits have been shown to have large effects in other contexts; for example, for the highly-effective government audits studied by Olken (2007), the primary mechanism for effects was that "the results of the audits were read publicly... and so could result in substantial social sanctions." Also notable is the fact that the payment treatment had no effect on any outcomes measured. The reason does not seem to be low power, since we can often rule out even modest effects (particularly because the sign of the point estimates is often opposite of the hypotheses laid out in Section 2). For example, our 95% confidence intervals allow us to reject effects of payments on coverage of eligible as small as 6 percentage points. Our results therefore provide no evidence that corruption declines with provider income.

Finally, we find inconclusive evidence on whether the small delivery treatment decreased leakage. The effect on leakage to mystery clients is large and negative (a 1.6 percentage point decrease off of a base rate of 2%),³¹ but not statistically significant. In addition, there is a large and significant positive effect on the likelihood that community members refer mystery clients to the health facility for a bed net.³²

³⁰We thank Kathleen Beegle for pointing out that this morale dampening effect may have been particularly pronounced in our experiment for the following reason. Health workers knew that we delivered bed nets to local shops for safekeeping, instead of delivering them to the health center. Therefore health workers likely perceived the voucher program as an anti-corruption scheme. If health workers had instead been asked to distribute vouchers for bed nets readily available in the retail sector, morale might have been higher. In practice, this was not an option we could pursue since 90% of health centers did not have a shop selling bed nets within 10 kilometers (see Table 1).

³¹This could be the reason why we observed more *free* nets being given to ineligible mystery clients in Kenya and Uganda than in Ghana (what we call "benevolent leakage"). The Ghana program was implemented by a small NGO at a small scale with unknown quantities of nets to offer, whereas in Kenya and Uganda where regular deliveries from district headquarters mean that leakage may not increase the risk of exclusion errors (pregnant women missing out on nets).

³²The impact of the small delivery treatment on the initial stock was substantial: direct distribution clinics that got a "large" delivery received 270 nets on average, compared to only 140 or so for small delivery clinics. However, 140 nets is still enough for over 1 month of projected net distribution based on the baseline registration and visit data (see Table A1: the average clinic saw 116 patients, including both new registrants and follow-up visits, each month). Thus, it is possible that the small delivery treatment was not small enough to seriously increase the salience of the budget constraint. Since delivering enough nets for only one week worth of distribution at a time would have large financial implications in terms of delivery costs, estimating whether it could improve targeting would not be policy-relevant, however.

7 Discussion

In this section, we explore possible explanations for the relatively high performance levels observed, and for why the monitoring and financial incentives provided in Ghana failed to further increase compliance with program rules. As discussed in section 2, the likelihood that health workers engage in what we call “corrupt” behavior may decrease with their level of motivation, either intrinsic, pro-social or extrinsic. Using survey data, we document that health workers appear positively selected in terms of pro-social motivation compared to other professionals, and that they also exhibit higher levels of job-specific intrinsic and extrinsic motivation. Consistent with high intrinsic and extrinsic motivation, health workers perform relatively well on a number of tasks beyond the bed net distribution scheme.

7.1 What Motivates Health Workers?

Tables 7 and 8 present results from surveys administered to health workers and other professionals. In Panel A, we present country-level averages of various motivation measures among health workers. In Panel B, we run the following regressions:

$$y_{i,c} = \beta_1 + \beta_2 \times Teacher_i + \beta_3 \times Shopkeeper_i + \beta_4 \times MFI_i + X_i + v_c + \varepsilon_{ic}$$

where y_{ic} is the outcome of worker i in country c , and $Teacher_i$, $Shopkeeper_i$ and MFI_i are dummy variables equal to 1 if the individual is a teacher, shopkeeper or microfinance agent, respectively; and v_c is a vector of country fixed effects. We present the coefficient estimates for β_2 , β_3 and β_4 , each representing the difference between the respective profession and health workers. X_i is a vector of demographic controls (age and gender).

Pro-social Motivation. The majority of health workers say that they receive personal satisfaction from helping people and do not expect anything in return (Table 7, Panel A). Within country, health workers appear substantially more pro-socially motivated than the other professionals we surveyed (Table 7, Panel B).

Besides survey responses, which can be subject to social desirability bias, we have an incentivized measure of other-regarding preferences from a dictator game. We gave players an envelope with ten bills and told them that the money they left in the envelope would be delivered to a randomly-selected community member living in their community.³³ Interestingly, the share left in the envelope by health workers is identical across all three countries

³³The total amount of the 10 bills was about \$6 in Ghana and Kenya, and about \$4 in Uganda. To avoid social pressure effects, the respondent was told that the surveyors would not open the envelope themselves.

(26-27%). More detailed results in Figure A2 show that health workers in Kenya and Uganda tend to leave more than other professionals (with the one exception of MFI employees in Uganda, who are more generous than Uganda health workers): pooling the other professionals together, health workers in Kenya (Uganda) gave 60% (22%) more than workers in other professions, with both differences significant at the 5% level. In Ghana, however, health workers give less.³⁴

Intrinsic Job Motivation. Columns 1-5 of Table 8 indicate that the majority of health workers see their jobs as benefiting society, believe their work is appreciated, and have high levels of job satisfaction. On all these measures the differences with other professions are significant. The experimental result that providing financial compensation to health workers for implementing the Ghana program had no effect on performance is consistent with health workers being sufficiently motivated without it. The relatively poor performance we observed under the voucher scheme in Ghana, which undermined health workers autonomy, is also consistent with an important role for intrinsic motivation.

Extrinsic Motivation. Table 8, columns 6 to 8, show that health workers are also highly extrinsically motivated. They generally report high levels of job insecurity, with the majority of health workers “strongly disagreeing” with the statement that “health worker jobs are very secure” (column 6).³⁵ Health workers’ perceived level of job security is significantly lower than that of teachers. This may be because performance is easily observable: the tasks health workers do are fairly standard, and in the specific context of the bed net distribution schemes we audited, leakage is easily observable since the eligibility criteria is a commonly observable one (pregnancy). Finally, we find that health workers feel very closely monitored by the Ministry of Health, though relatively less so in Uganda, where they also performed relatively worse overall (Table 8, column 7). Health workers also report higher levels of monitoring than other professions (Panel B). The fact that health workers believe they are accountable could explain the experimental finding that the threat of top-down audits was ineffective at increasing performance in Ghana: health workers may have expected there to be an audit even before the threats were made.

³⁴In Ghana, the dictator games were played in 2012 with health workers and ANC clients, and in 2014 with the other professionals. Given an inflation rate of approximately 9% over this time period, the two rounds of dictator games are not directly comparable because the stakes differed (though the evidence on this is mixed, see Andersen et al. (2011) for evidence that stakes may matter in the ultimatum game).

³⁵This is consistent with surveys we conducted with head nurses in Kenya and Uganda, which revealed that around 10% of them knew at least one health worker who had lost her job due to poor performance. In Ghana, no head nurses reported knowing someone who had lost their job, but they argued it did not mean misconduct would go unpunished. In equilibrium, the threat of job loss may be enough to deter misconduct.

Correlations with clinic performance Do these characteristics map into higher performance? Because we do not have health worker specific performance outcomes, we can only test for correlations at the clinic level (the level at which outcomes are measured). We do so in Table A2, in which we regress clinic performance on standardized indices of pro-social motivation, intrinsic motivation, and extrinsic motivation. We find that other-regarding preferences and extrinsic motivation are positively correlated with some measures of performance, but that intrinsic motivation is not. In particular, we find that other-regarding preferences is positively correlated with the coverage rate among the eligible and negatively correlated with leakage to the ineligible,³⁶ and that coverage among the eligible was higher in health facilities whose health workers had a higher extrinsic motivation score. Overall the correlations are weak, however, possibly owing to the fact that we only have clinic-level rather than individual-level performance.

7.2 Performance on other tasks

Table 9 presents data on health worker performance on other tasks. Data from the ANC patient interviews show that nurses spent an average of 18 minutes with each patient. Nurses conducted palpation (the key prenatal check that nurses are supposed to perform) for 96% of the clients, and this is comparable across countries (col 2). Average wait time for a checkup is about one hour on average. ANC staff members also seem to engage in good record-keeping, with 96% of the key identifier fields filled in the registers (col 8).³⁷

In Kenya and Uganda, we also asked about additional services that should be performed during ANC visits: 84% of women received intermittent preventive treatment for malaria (IPT, col 4) and 73% received iron tablets (col 5). These figures are broadly similar to the coverage rates from bed net distribution. If women brought a child into a government clinic for malaria treatment, 92% received the recommended Artemisinin-based Combination Therapy (ACT), an even higher level of coverage than for the nets.³⁸ However, extortion appears higher here: while 78% of women report receiving the drugs for free, as they should have (col 6), 14% of women reported that they received the drugs but had to pay (col 7). One caveat is that it is highly possible that people confused payments for ACTs with other ancillary fees (e.g., lab tests) associated with malaria visits, an often-cited issue complicating

³⁶These results are consistent with Brock, Lange and Leonard (2013), who find a positive correlation between altruism as measured through the dictator game and effort provided by doctors in Tanzania, and with Callen et al. (2013), who find that doctors with higher public sector motivation are less likely to be absent and less likely to falsify reports.

³⁷The “key fields” are name, registration date, ANC card #, address, # children, and gestational age at registration. We digitized data from the ANC registers for sampled patients only.

³⁸We exclude non-government clinics from these statistics since ACTs are not supposed to be free at these facilities.

the study of informal payments for drugs and other health services (Lewis, 2007; Stepurko et al., 2010). But, another potential explanation is that health workers are more likely to request payments for drugs than they are with bed nets. This could be because willingness to pay for treatment is typically higher than for prevention (Dupas, 2011; Cohen, Dupas and Schaner 2015), and therefore extorting for drugs may be more effective than for bed nets. In this case, higher willingness to pay may have the negative effect of increasing bribe requests, but at the same time may have the positive effect of increasing the coverage rate because women are more likely to demand treatment. In this setting, near-universal coverage does not seem much dampened by bribe requests, perhaps because health workers can price discriminate. Interestingly, we observe a positive correlation between payment requests for ACTs and bribe requests from mystery clients (correlation of 0.3, significant at the 5% level) at the health facility level.

Figure A3 shows the results of unannounced attendance spot checks conducted at each clinic in the Kenya and Uganda samples. Of the health workers who were officially supposed to be on duty, 8% ($=.05/ (.05+.61)$) were absent in Kenya and 20% ($=.13/ (.13+.53)$) in Uganda.^{39,40} Perhaps more important than the individual-level attendance results are the results on clinic closure, since clinics in this setting are large enough that health workers can load-share and so absenteeism may not compromise performance: clinics were closed for fewer than 1% of mystery client visits (Table 9, col 9). Also suggestive that attendance may not directly impact service quality is the fact that we do not find a correlation between attendance and any of our other performance measures (see Table A3).

8 Conclusion

Increasing coverage of life-saving health products in rural sub-Saharan Africa requires distribution at heavily subsidized prices, and the potentially most cost-effective way to do this is through existing health systems. Whether government health workers can do this effectively

³⁹Note that the “away on official duty” and “official day off” categories that we exclude may include absenteeism disguised as official days off. Chaudhury et al. (2006) excludes workers who are not on duty (which would correspond most closely to our “official day off” category) from their absence calculations, but count as absent those away on official duty (less than 4% of health workers on duty in their case, much lower than the 19% (16/82) we observe); if we tabulate the results that way, it would mean absence rates of 26% in Kenya ($=.21/ (1-.18)$) and 36% in Uganda, relatively similar to the 35% rate found across countries and 38% rate found in Uganda by Chaudhury et al. (2006).

⁴⁰These attendance rates differ from the survey compliance rates for the healthworker survey for several reasons: first, 33% of the healthworkers who were present during the surveyors’ first visit were not interviewed because they were too busy or because surveyors ran out of time; second, some health workers were off duty or away on official leave; and third, surveyors returned to most clinics on multiple dates to increase completion rates. The first two would cause compliance rates to underestimate attendance rates; the third would cause them to overestimate.

is an open question. Will they respect the eligibility rule? Will they demand bribes? Will they even bother to implement the program? How can we design distribution schemes to avoid these problems? We shed some light on these questions by auditing government distribution schemes in Kenya and Uganda, and implementing and auditing a program in Ghana in which we randomly varied several features including bonus pay, the threat of audits, and whether the distribution is direct or indirect (through vouchers). A key contribution is measuring performance, extortion and leakage in various unobtrusive ways (e.g., “mystery client” visits).

We find that distribution programs administered through existing health centers perform better than conventionally believed. Across the three countries, 80% of eligible women received the subsidy, only 1% of eligible women were asked to pay bribes, and at most 15% of local subsidies leaked to ineligible people, most often for free, and more often when ineligibles mentioned having a needy child. While the limited level of leakage could come from a lack of demand from the ineligible, in Ghana we experimentally “tempted” health workers by sending ineligible men with a high willingness to pay for the subsidized product, and we found that very few of them were successful at obtaining it from healthworkers.

Our experimental results suggest further improving performance beyond its current level is not easy: neither the threat of top-down audits nor bonus pay, two commonly proposed interventions to improve worker performance, made a difference in Ghana. There thus appears to be no low hanging-fruit to be picked on the governance side of things, possibly because health workers’ motivation levels are fairly high. On the other hand, voucher schemes, which are popular as a way to reduce health workers’ discretion and which have been implemented in a number of countries (most notably in Tanzania), appear to actually modestly reduce program effectiveness. This could be because health workers’ intrinsic job motivation might have been undermined by a lack of “trust” (i.e., the fact that they are not given responsibility for the subsidized bed nets themselves), or that people in the community may have been less aware of the (less visible) voucher scheme. The first mechanism is similar to the one proposed by Rasul and Rogger (2016), who find that bureaucrats in Nigeria are less likely to complete projects if they are given less autonomy.

While our data comes from only three countries, the programs we consider were implemented by different institutions (the government in Kenya and Uganda, vs. an NGO in Ghana) and audited at very different points in their implementation (after more than a year in Kenya and Uganda, vs. in the first few months in Ghana), suggesting that our results may be representative of similar programs in other contexts. An interesting question is whether the results carry through to distribution schemes for curative products, which may be more valued by patients and thus may be more easy to charge payment for. While this was not

the focus of our study, in the countries where antimalarial treatment is supposed to be distributed free of charge at public facilities we find a mixed picture in which both coverage rates and requests for payments are higher for anti-malarial drugs than for bed nets. Though side payment requests are a clear cause for concern, the fact that near-universal coverage is attainable despite them suggests that health workers are at least able to effectively price discriminate, and further suggests that public distribution can be effective.

The relatively high pass-through we observe of the subsidy to its intended beneficiaries goes against the growing conventional wisdom that service provision in developing countries is universally poor. In fact, permissions to conduct this study were difficult to obtain, because of what was perceived as its highly sensitive nature. Anyone we discussed our study design with *ex ante* expected us to observe poor performance levels, in particular much higher non-benevolent leakage. Why are our results different from expectations?

One consideration is that we consider a targeting rule that is very easy to verify, and so makes it harder to hide leakage. While we fully acknowledge that leakage rates could be different under a different targeting mechanism, assessing leakage levels for “easy” targeting rules, as we have done, is important since many products/subsidies are in practice targeted at women or other easily identified subgroups such as young children.

A second consideration is that service quality varies from country to country and so far service provision has only been studied rigorously in a small subset of countries, mostly outside of sub-Saharan Africa, and with a strong focus on South Asia. What’s more, absenteeism is typically considered to be a good proxy for performance (the Chaudhury et al. 2006 study exclusively measured absenteeism), and while this may be an important metric in contexts where health centers are small enough that absenteeism of one worker means the health center is closed (as in Banerjee, Duflo and Glennerster 2008) or for services that require scarce expertise (as in Goldstein et al. 2013), absenteeism may not have direct consequences on service in more common contexts.

References

- [1] Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari and Matthew Wai-Poi (2013). Does Elite Capture Matter? Local Elites and Targeted Welfare Programs in Indonesia. NBER Working paper 18798.
- [2] Andersen, Steffen, Seda Ertac, Uri Gneezy, Moshe Hoffman, and John A. List (2011). Stakes Matter in Ultimatum Games. *American Economic Review* 101 (7): 3427-39.
- [3] Banerjee, Abhijit. V. (1997). A Theory of Misgovernance. *Quarterly Journal of Economics* 112(4), 1289-332.
- [4] Banerjee, Abhijit. V., Esther Duflo and Rachel Glennerster (2008). Putting A Band-Aid On A Corpse: Incentives For Nurses In The Indian Public Health Care System. *J Eur Econ Assoc*; 6(2-3): 487–500.
- [5] Bardhan, Pranab and Dilip Mookherjee (2000). Capture and Governance at Local and National Levels. *American Economic Review* 90 (2): 135–139.
- [6] Basurto, Pia, Pascaline Dupas and Jonathan Robinson (2016). Decentralization and Efficiency of Subsidy Targeting: Evidence from Chiefs in Rural Malawi. Working paper.
- [7] Becker, Gary and George Stigler (1974). Law Enforcement, Malfeasance and the Compensation of Enforcers. *Journal of Legal Studies* 3(1), 1-19.
- [8] Björkman, Martina and Jakob Svensson (2009). Power to the People: Evidence from a Randomized Field Experiment on Community-Based Monitoring in Uganda. *Quarterly Journal of Economics* 124(2): 735–769.
- [9] Brock, Michelle J., Andreas Lange and Kenneth L. Leonard (2013). Generosity Norms and Intrinsic Motivation in Health Care Provision: Evidence from the Laboratory and Field. EBRD Working Paper #147.
- [10] Callen, Michael, Saad Gulzar, Ali Hasinain, Yasir Khan, and Arman Rezaee (2013). Personalities and Public Sector Performance: Experimental Evidence from Pakistan. Working paper.
- [11] Chaudhury, Nazmul and Jeffrey Hammer (2004). Ghost Doctors: Absenteeism in Rural Bangladeshi Health Facilities. *World Bank Economic Review* 18 (3): 423-441.
- [12] Chaudhury, Nazmul, Jeffrey Hammer, Michael Kremer, Karthik Muralidharan, and F. Halsey Rogers (2006). Missing in Action: Teacher and Health Worker Absence in Developing Countries. *Journal of Economic Perspectives* 20 (1): 91-116.
- [13] Cohen, Jessica and Pascaline Dupas (2010). Free Distribution or Cost-Sharing? Evidence from a randomized malaria experiment. *Quarterly Journal of Economics* 125(1): 1-34.
- [14] Das, Jishnu, and Jeffrey Hammer (2014). Quality of Primary Care in Low-Income Countries: Facts and Economics. *Annual Review of Economics* 6: 525-553.
- [15] Das, Jishnu, Jeffrey Hammer and Kenneth Leonard (2008). The Quality of Medical Advice in Low-Income Countries. *Journal of Economic Perspectives* 22(2): 93-114.

- [16] Das, Jishnu, Alaka Holla, Michael Kremer and Karthik Muralidharan (2015). Quality and Accountability in Health: Audit Evidence from Primary Care Clinics in India. Working paper.
- [17] De Allegri, M., P. Marschall, S. Flessa, J. Tiendrebéogo, B. Kouyaté, A. Jahn, and O. Müller (2010). Comparative cost analysis of insecticide-treated net delivery strategies: sales supported by social marketing and free distribution through antenatal care. *Health policy and planning* 25(1): 28-38.
- [18] Dal Bo, Ernesto, Frederico Finan and Martin A. Rossi (2013). Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service. *Quarterly Journal of Economics* 128(3): 1169-1218.
- [19] Di Tella, Rafael and Ernesto Schargrodsky (2003). The Role of Wages and Auditing during a Crackdown on Corruption in the City of Buenos Aires. *Journal of Law and Economics* 46: 269-92.
- [20] Dupas, Pascaline (2014). Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment. *Econometrica* 82(1): 197-28.
- [21] Dupas, Pascaline, and Edward Miguel (2016). Impacts and Determinants of Health Levels in Low-Income Countries. Forthcoming, *Handbook of Field Experiments*, Banerjee and Duflo, Eds.
- [22] Foltz, Jeremy B. and Kweku A. Opoku-Agyemang (2014). Do higher salaries lower petty corruption? A policy experiment on West Africa's highways". Working paper.
- [23] Gertler, Paul, and Christel Vermeersch (2013). Using Performance Incentives to Improve Health Outcomes. NBER Working Paper No. 19046.
- [24] Markus Goldstein, Josh Graff Zivin, James Habyarimana, Christian Pop-Eleches and Harsha Thirumurthy (2013). The Effect of Absenteeism and Clinic Protocol on Health Outcomes: The Case of Mother-to-Child Transmission of HIV in Kenya, *American Economic Journal: Applied Economics* 5(2): 58-85.
- [25] Lewis, Maureen (2007). Informal payments and the financing of health care in developing and transition countries. *Health Affairs*, 26(4), 984-997.
- [26] Muralidharan, Karthik, Nazmul Chaudhury, Jeffrey Hammer Michael Kremer, and F. Halsey Rogers (2011). Is there a Doctor in the House? Absent Medical Providers in India. Working paper.
- [27] Niehaus, Paul, Antonia Attanassova, Mariane Bertrand and Sendhil Mullainathan (2013). Targeting with Agents. *American Economic Journal: Economic Policy* 5 (1): 206-238.
- [28] Olken, Benjamin (2005). Corruption and the Costs of Redistribution: Micro Evidence from Indonesia. *Journal of Public Economics* 90 (4-5): 853-870.
- [29] Olken, Benjamin (2007). Monitoring Corruption: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy* 115(2): 200-249.
- [30] Olken, Benjamin, and Rohini Pande (2013). Governance Review Paper, J-PAL Governance Initiative.

- [30] Olken, Benjamin, and Rohini Pande (2012). Corruption in Developing Countries. *Annual Review of Economics* 4: 479-505.
- [31] Rasul, Imran and Daniel Rogger (2016). Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service. Forthcoming, *Economic Journal*.
- [32] Rauch, James E and Peter B. Evans (2000). Bureaucratic Structure and Bureaucratic Performance in Less Developed Countries. *Journal of Public Economics* 75(1): 49-71.
- [33] Sexton, Alexis R. (2011). Best practices for an insecticide-treated bed net distribution programme in sub-Saharan eastern Africa. *Malaria journal* 10(1): 157.
- [34] Svensson, Jakob (2005). Eight questions about corruption. *The Journal of Economic Perspectives* 19(3): 19-42.
- [35] Stepurko, Tetiana, Milena Pavlova, Irena Gryga, and Wim Groot (2010). Empirical studies on informal patient payments for health care services: a systematic and critical review of research methods and instruments. *BMC Health Services Research*, 10(1), 1.
- [36] Tarozzi, Alessandro, Aprajit Mahajan, Brian Blackburn, Dan Kopf, Lakshmi Krishnan, and Joanne Yoong (2014). Micro-Loans, Insecticide-Treated Bednets, and Malaria: Evidence from a Randomized Controlled Trial in Orissa, India. *American Economic Review* 104(7): 1909-41.
- [37] United Nations Development Programme (2011). Kohler, Jilian. "Fighting corruption in the health sector: methods, tools and good practices." New York: United Nations Development Programme.
- [38] Van Rijckeghem, Caroline and Beatrice Weder (2001). Bureaucratic corruption and the rate of temptation: do wages in the civil service affect corruption, and by how much? *Journal of Development Economics* 65(2): 207-331.
- [39] World Bank (1993). *World Development Report 1993: Investing in Health*. New York: Oxford University Press.
- [40] World Bank (2004). *World Development Report 2004: Making Services Work for Poor People*. Washington DC: The World Bank and Oxford University Press.

Table 1. Baseline bed net coverage in study areas, and Characteristics of health facilities in the sample

<u>Panel A. Bed net coverage^a</u>		Ghana	Kenya	Uganda
DHS Survey year		2011	2010/2014	2011
Average household size		4.3	4.4/n.a.	4.9
Share of households with at least one insecticide-treated net (ITN)		0.53	0.48/0.83	0.60
Average number ITNs per household		0.80	0.80/1.80	1.30
Percentage of households with at least one ITN for every two persons who stayed in the household last night ^b		0.22	n.a./0.417	0.33
<u>Panel B. Study Facility characteristics</u>		<i>Sample Mean [Std. Dev.] for:</i>		
	Full Sample	Ghana	Kenya	Uganda
Survey year		2011	2013	2013
Years since facility is operating	16.01 [15.86]	17.25 [13.74]	13.96 [16.86]	16.30 [17.67]
Public facility	0.85 [0.36]	0.90 [0.30]	1.00 [0.00]	0.60 [0.49]
# of monthly ANC new registrants ^c	27.99 [20.06]	25.73 [20.31]	29.64 [19.31]	29.75 [20.50]
# of monthly ANC follow-up visits	63.06 [61.25]	90.22 [77.59]	49.40 [36.44]	35.98 [29.23]
# of midwives and nurses for ANC	2.93 [2.28]	2.01 [1.20]	4.15 [3.31]	3.08 [1.58]
Facility conducts outreach ANC activities	0.49 [0.50]	0.23 [0.42]	0.92 [0.28]	0.46 [0.50]
Nets available for sale within 10km	0.13 [0.34]	0.10 [0.30]	0.10 [0.31]	0.21 [0.41]
Has a maternity ward	0.82 [0.39]	0.85 [0.36]	0.96 [0.20]	0.63 [0.49]
Accessible during the rainy season	0.89 [0.32]	0.81 [0.40]	0.94 [0.24]	0.96 [0.20]
Health worker privately sells nets at facility	0.02 [0.13]	0.04 [0.20]	0.00 [0.00]	0.00 [0.00]
Number of Health Facilities	168	72	48	48

Notes: For Ghana sample, includes all health facilities/communities, whether sampled for direct or indirect (voucher) distribution.

^a Source: Demographic and Health Surveys. For each country, we show the average for the region included in our study.

^b This is the DHS definition for universal coverage

^c ANC stands for antenatal care

Table 2. Coverage and extortion among eligibles

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Data Source:	Backcheck surveys with random subset of ANC Clients						
		<i>If no stockouts^a :</i>		<i>If offered net at some visit:</i>	<i>If not offered net:</i>		
	Received net at first ANC visit	Received net at first ANC visit	Received net at some ANC visit	Was requested to pay	Thinks reason is stockout	Does not know why	Thinks reason is already has one
Overall Mean	0.76	0.81	0.80	0.01	0.25	0.50	0.07
Ghana	0.74	0.77	0.79	0.01	0.11	0.63	0.04
Kenya	0.90	0.90	0.91	0.01	0.66	0.17	0.06
Uganda	0.63	0.69	0.66	0.03	0.25	0.49	0.09
Observations	2,028	1,495	2,028	1,605	473	473	473
P-value for equality of means:							
Ghana = Kenya	0.00***	0.00***	0.00***	0.72	0.00***	0.00***	0.69
Ghana = Uganda	0.03**	0.16	0.01***	0.04**	0.01***	0.03**	0.08*
Kenya = Uganda	0.00***	0.00***	0.00***	0.02**	0.00***	0.00***	0.54

Notes: There are 144 health facilities in the sample (48 per country). Ghana sample: Only includes facilities sampled for direct distribution.

^a Clinics with no stockouts during the sampled time period identified using data from attendance checks in Uganda and Kenya, and from mystery client visits from Ghana.

Table 3. Non-coverage among eligibles: Errors of Exclusion or Efficient Targeting?

	(1)	(2)	(3)	(4)	(5)
Data Source:	Backcheck surveys with random subset of ANC Clients				
	All	Ghana	Kenya	Uganda	All countries - ANC visits in month preceding stockout only
<u>Panel A. Dependent Variable: Received net - visit 1</u>					
Years of education	-0.0068** [0.0030]	-0.0083* [0.0044]	-0.0026 [0.0034]	-0.0065 [0.0065]	-0.0155 [0.0178]
Health facility Fixed effects	X	X	X	X	X
Observations	2,028	771	671	586	131
R-Squared	0.357	0.274	0.369	0.349	0.686
Dep. Var. Mean	0.762	0.744	0.903	0.625	0.542
<u>Panel B. Dependent Variable: Received net - any visit</u>					
Years of education	-0.0046* [0.0028]	-0.0039 [0.0040]	-0.0022 [0.0037]	-0.0066 [0.0063]	-0.0099 [0.0151]
Health facility Fixed effects	X	X	X	X	X
Observations	2,028	771	671	586	131
R-Squared	0.361	0.285	0.390	0.364	0.711
Dep. Var. Mean	0.796	0.794	0.914	0.664	0.588

Notes: Standard errors in brackets, clustered at the level of the health facility. Ghana sample: Only includes facilities sampled for direct distribution. Regressions have month fixed effects to control for potential stockouts, and include controls for registrant age and parity. The mean number of years of education among clients surveyed is 5 (4.2 in Ghana, 4.6 in Kenya and 6.9 in Uganda, which adopted free primary education much earlier). Across all countries (column 1), the gap between the 25th and the 75th percentile in terms of years of education is 4.5 years.

Table 4. Leakage to ineligible: Evidence from mystery client visits and community interviews

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Data Source:	"Mystery client" (MC) visits					Community Interviews				
	Clinic out of stock during visit (as measured through other data sources) ^a	MC told that the clinic is out of stock	If clinic had nets in stock ^b		Payment requested for net	If Payment requested:		Respondent thinks that (male) surveyor could get a net at a health center	Ineligible young man received net (lower bound) ^d	Ineligible young man received net (upper bound) ^e
			Acquired program net	Acquired program net for free		Initial amount requested as percent of full price ^c	Final amount requested as percent of full price ^c			
Overall Mean	0.205	0.134	0.0467	0.0378	0.0424	0.902	0.745	0.0993	0.00699	0.0381
Ghana		0.046	0.022	0.009	0.051	1.070	0.882	0.092		
Kenya	0.042	0.090	0.087	0.087	0.036	0.121	0.097	0.093	0.025	0.069
Uganda	0.368	0.472	0.110	0.110	0.011	0.323	0.323	0.113	0.000	0.007
Observations	288	766	685	687	684	27	27	2,559	143	289
P-value for equality of means:										
Ghana = Kenya		0.19	0.03**	0.01***	0.51	0.00***	0.00***	0.98		
Ghana = Uganda		0.00***	0.01***	0.00***	0.01**	0.02**	0.17	0.31		
Kenya = Uganda	0.15	0.00***	0.59	0.59	0.24	0.52	0.60	0.35	0.31	0.00***

Notes: Missing coefficients mean that outcome not collected in that country. There are 144 health facilities in the sample (48 per country). Ghana sample: Only includes facilities sampled for direct distribution.

^a Stockout data collected during attendance spot checks done by the research team in Kenya and Uganda.

^b Conditional on there not being a stockout (calculated from attendance spot check data for Kenya and Uganda, and from MC reports in Ghana). Upper bound in Ghana since health workers could have told mystery clients they were stocked out to get rid of them.

^c Full price calculated as the wholesale price converted to local currency using exchange rate during the middle of the ANC registration times.

^d Lower bound because set to missing when he received net but someone in the household was pregnant or under 5

^e Upper bound because includes people who got a net but were in a household that may have been eligible (was reported to have a pregnant woman or child under 5). Only upper bound if assume percent of eligibles that get nets weakly higher than percent of ineligible.

Table 5. Leakage: Errors of Inclusion or Efficient targeting?

	(1)	(2)	(3)	(4)
Data Source:	Mystery Client visits			
Dependent Variable:	Received free net			
	All	Ghana	Kenya	Uganda
Requested for pregnant woman	0.012 [0.023]	0.011 [0.014]	0.023 [0.14]	
Requested for child	0.18*** [0.045]			0.19** [0.089]
If asked, said that had child	0.11** [0.047]		0.11 [0.075]	
MC signaled that educated	0.0077 [0.016]	0.0077 [0.0090]		
Healthworker female	-0.012 [0.022]	-0.012 [0.013]		
Health facility Fixed effects	X	X	X	X
Observations	683	455	137	91
R-Squared	0.402	0.114	0.466	0.375
Mean of the Dependent Variable	0.0378	0.009	0.087	0.110

Notes: The first three regressors are mutually exclusive indicator variables reflecting the strategy used by the mystery client when seeking a bed net from the healthworker. The choice of strategy was left to the mystery clients, and, possibly owing to differences in training, there were systematic cross-country differences in the strategies they used, hence the regressors vary across countries in columns 2-4. Standard errors in brackets, clustered at the level of the health facility. There are 144 facilities in the sample (for Ghana sample: Direct distribution clinics only).

Table 6. Ghana: experimental results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Data Source:	Backcheck surveys with random subset of ANC Clients			Surveys with Program Log beneficiaries	Community Interviews		"Mystery Client" visits				Administrative data estimate
Dependent Variable:	Received net/voucher at first ANC visit	Received net/voucher at some ANC visit	<i>If offered net at some visit:</i> Was requested to pay	Listed beneficiary ineligible	Respondent (male) surveyor thinks that could get a net at a health center	Respondent aware of net distribution program	Payment requested for net	Acquired program net	Acquired program net for free	Clinic stocked out	Total local leakage rate
<u>Effect of:</u>											
Voucher	-0.178** [0.072]	-0.195*** [0.069]	-0.002 [0.009]	0.013 [0.012]	-0.067*** [0.025]	-0.085** [0.034]	0.035 [0.026]	0.031* [0.018]	0.007 [0.008]	0.073** [0.030]	-0.084* [0.043]
Audit	0.033 [0.056]	0.069 [0.054]	0.017 [0.013]	-0.032** [0.014]	-0.060 [0.042]	0.008 [0.066]	0.003 [0.031]	0.010 [0.025]	-0.012 [0.013]	0.010 [0.040]	-0.027 [0.029]
Small delivery	0.007 [0.059]	0.006 [0.055]	-0.009 [0.007]	0.001 [0.012]	0.061*** [0.021]	0.038 [0.035]	-0.015 [0.022]	-0.015 [0.012]	-0.011 [0.007]	0.020 [0.021]	0.015 [0.034]
Health worker payment	-0.061 [0.061]	-0.056 [0.058]	-0.003 [0.007]	0.010 [0.012]	-0.009 [0.021]	-0.002 [0.035]	0.010 [0.021]	0.002 [0.012]	0.002 [0.007]	0.005 [0.020]	-0.002 [0.035]
Later program period	-0.079* [0.041]	-0.120*** [0.041]	-0.021** [0.010]	0.003 [0.008]	0.058* [0.029]	-0.024 [0.052]	-0.046* [0.025]	-0.036* [0.019]	0.001 [0.011]	-0.001 [0.031]	
Observations	1,158	1,158	815	1,229	1,322	1,322	717	717	720	718	71
R-Squared	0.094	0.108	0.023	0.027	0.030	0.025	0.028	0.029	0.029	0.060	0.241
Dep. Var. Mean in direct distribution facilities	0.744	0.794	0.00873	0.0144	0.0916	0.716	0.0503	0.0209	0.00833	0.0460	0.147

Notes: There are 72 health facilities in the sample, all in Ghana. Each column corresponds to an OLS regression with randomization strata fixed effects and two health-center level controls: baseline ANC attendance and ANC staff size. Standard errors clustered at the health facility level. ***, **, * indicates significance at 1, 5 and 10 percent. "Effect of Audit" is the difference in differences coefficient, so the coefficient on AuditXLater program period, where "later program period" is the period after the initial audit threat in audit clinics, or after the audit threats were conducted at nearby clinics for non-audit clinics; all regressions also control for the

Table 7. Pro-social Motivation and Selection into Health Work

	(1)	(2)	(3)	(4)	(5)	(6)
	Pro-Social Motivation				Job Characteristics	
	Survey response on a scale from 1 (don't agree) to 5 (strongly agree) to the statement:					
	Helping people brings personal satisfaction	If you help someone, they should do you a favor in return	It is important to do good things for my community	My family comes first, my work second	Monthly pay ^a	Years of education
<u>Panel A. Country-level averages across health workers</u>						
Ghana	4.848 [0.0481]	1.463 [0.109]	4.939 [0.0288]	3.331 [0.105]	235.8 [13.99]	14.02 [0.228]
Kenya	4.898 [0.0501]	1.321 [0.105]	5 [0]	3.265 [0.129]	387.7 [16.28]	15.87 [0.0630]
Uganda	4.631 [0.0482]	2.403 [0.135]	4.785 [0.0316]	3.295 [0.136]	141.2 [7.149]	15.67 [0.127]
Observations	450	450	449	444	412	449
P-value for equality of means:						
Ghana = Kenya	0.47	0.35	0.04**	0.69	0.00***	0.00***
Ghana = Uganda	0.00***	0.00***	0.00***	0.83	0.00***	0.00***
Kenya = Uganda	0.00***	0.00***	0.00***	0.87	0.00***	0.17
Overall Mean	4.791	1.731	4.906	3.300	250.7	15.13
<u>Panel B. Comparisons with other professions (omitted category: health workers)</u>						
Teacher	-0.282*** [0.0563]	0.473*** [0.103]	-0.114*** [0.0330]	0.238** [0.102]	-33.10** [12.95]	0.309*** [0.116]
Shop owner	-0.347*** [0.0801]	1.261*** [0.145]	-1.608*** [0.136]	0.963*** [0.130]	-106.1*** [17.90]	-5.985*** [0.457]
Microfinance agent	-0.324*** [0.0769]	0.961*** [0.146]	-0.163*** [0.0525]	0.708*** [0.123]	-155.4*** [12.71]	-1.908*** [0.209]
Observations	1,136	1,138	1,134	1,127	1,094	1,172

Notes: Panel A: Data from surveys with health workers only. Ghana sample only includes workers from facilities sampled for direct distribution. Panel B: Data from surveys with health workers, shopkeepers, teachers and MFI agents. Regressions control for country fixed effects. Cols (5) and (6) of panel B also control for worker gender and age. *** p<0.01, ** p<0.05, * p<0.1.

^a Winsorized at 99th percentile. Converted to USD using the exchange rate at the time the survey was taken, except for the other profession salaries (i.e., non-healthworker salaries) in Ghana: since healthworkers in Ghana were surveyed two years before the other professions (2012 vs 2014), the other profession salaries were first converted to 2012 Ghana cedis using inflation rates, then converted to USD using exchange rates at the time the healthworker surveys were conducted.

Table 8. Intrinsic and Extrinsic Job Motivation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Survey response on a scale from 1 (don't agree) to 5 (strongly agree) to the statement:							
	Intrinsic Job Motivation				Extrinsic Motivation			
	As a [<i>profession</i>], I have the responsibility to be a role model in my community	This [<i>work place</i>] plays a very important role in this community	I am very satisfied with my job	People in remote areas do not have enough appreciation for [<i>profession</i>]	[<i>Profession</i>] should be paid more	[<i>Profession</i>] jobs are very secure -- even if a worker does a bad job they won't be fired	This [facility/school] is closely monitored by the Ministry of [Health/Education]	This [facility/school] is closely monitored by local NGOs
<u>Panel A. Country-level averages across health workers</u>								
Ghana	4.909 [0.0334]	4.896 [0.0378]	4.341 [0.125]	3.135 [0.133]	4.710 [0.0734]	1.288 [0.0726]	4.834 [0.0609]	2.276 [0.154]
Kenya	4.912 [0.0643]	4.971 [0.0291]	4.081 [0.133]	2.378 [0.163]	4.657 [0.0930]	1.241 [0.0760]	4.912 [0.0646]	3.504 [0.165]
Uganda	4.879 [0.0304]	4.866 [0.0303]	3.831 [0.0979]	3.087 [0.113]	4.691 [0.0491]	1.523 [0.0652]	4.365 [0.0960]	4.007 [0.0932]
Observations	450	449	448	447	448	449	448	446
P-value for equality of means:								
Ghana = Kenya	0.96	0.12	0.15	0.00***	0.66	0.65	0.38	0.00***
Ghana = Uganda	0.52	0.54	0.00***	0.78	0.83	0.02**	0.00***	0.00***
Kenya = Uganda	0.64	0.01**	0.13	0.00***	0.74	0.01***	0.00***	0.01***
Overall Mean	4.900	4.909	4.094	2.890	4.688	1.352	4.703	3.220
<u>Panel B. Comparisons with other professions (omitted category: health workers)</u>								
Teacher	-0.0120 [0.0383]	-0.152*** [0.0346]	-0.447*** [0.106]	0.499*** [0.121]	-0.0756 [0.0646]	0.443*** [0.0763]	-0.382*** [0.0642]	-0.152 [0.122]
Shop owner	-0.524*** [0.0797]	-1.008*** [0.111]	-0.139 [0.114]	0.458*** [0.150]	-0.782*** [0.105]	N/A	N/A	N/A
Microfinance	-0.524*** [0.0719]	-0.330*** [0.0571]	-1.026*** [0.146]	0.0481 [0.144]	-0.168** [0.0759]	0.0823 [0.110]	N/A	N/A
Observations	1,138	1,133	1,134	1,126	1,131	997	870	867

Notes: Panel A: Data comes from surveys with health workers at 144 health facilities. Ghana sample only includes workers from facilities sampled for direct distribution. Panel B: Data from surveys with health workers, shopkeepers, teachers and MFI agents. Each column is a separate regression, with controls for country fixed effects. Standard errors in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table 9. Performance beyond net distribution

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Data Source:	Backcheck surveys with random subset of ANC Clients							ANC Register Audit	Mystery Client visits
Dependent Variable:	Avg. minutes spent with nurse during ANC checkup	Palpated by nurse during ANC visit	Average wait time for checkup	Received IPT during ANC visit	Received iron tablets during ANC visit	Got free ACT drugs last time brought child for malaria treatment ^a	Was charged fee for ACT last time brought child for malaria treatment ^a	% of key fields filled in ^b	Clinic closed during first visit attempt
Overall Mean	17.91	0.960	63.70	0.835	0.730	0.780	0.144	0.964	0.00595
Ghana	12.294	0.937	57.900	N/A	N/A	N/A	N/A	0.963	0.006
Kenya	25.249	0.985	72.374	0.932	0.660	0.777	0.158	0.965	0.000
Uganda	20.956	0.977	65.741	0.732	0.803	0.789	0.100	0.964	0.014
Observations	2,283	2,310	2,246	1,180	1,183	382	382	2,617	1,008
P-value for equality of means:									
Ghana = Kenya	0.00***	0.00***	0.03**	N/A	N/A	N/A	N/A	0.84	0.04**
Ghana = Uganda	0.00***	0.00***	0.32	N/A	N/A	N/A	N/A	0.94	0.41
Kenya = Uganda	0.01***	0.31	0.41	0.00***	0.00***	0.82	0.16	0.87	0.15

Notes: Standard errors for regressions clustered at clinic level. There are respondents from 168 health facilities in the sample (Ghana sample includes all facilities, including those sampled for indirect distribution via vouchers). Results are identical when the voucher facilities are excluded. Columns 4-7: these questions were only asked in Kenya and Uganda.

^a Excludes non-government clinics because ACT drugs only supposed to be free in government clinics.

^b Key fields are name, reg date, ANC card #, address, # children, and gestational age at registration.

Figure 1. Clinic-level Estimate of the Number of Program Bed Nets Unaccounted for (Ghana only)

Figure 1A. Number of nets unaccounted for

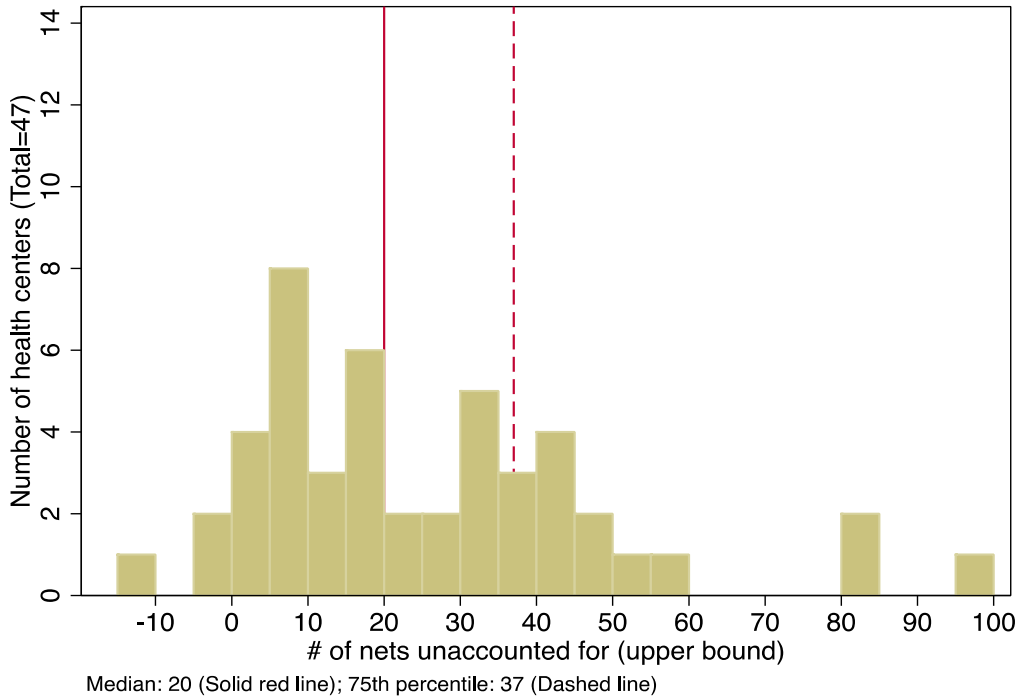
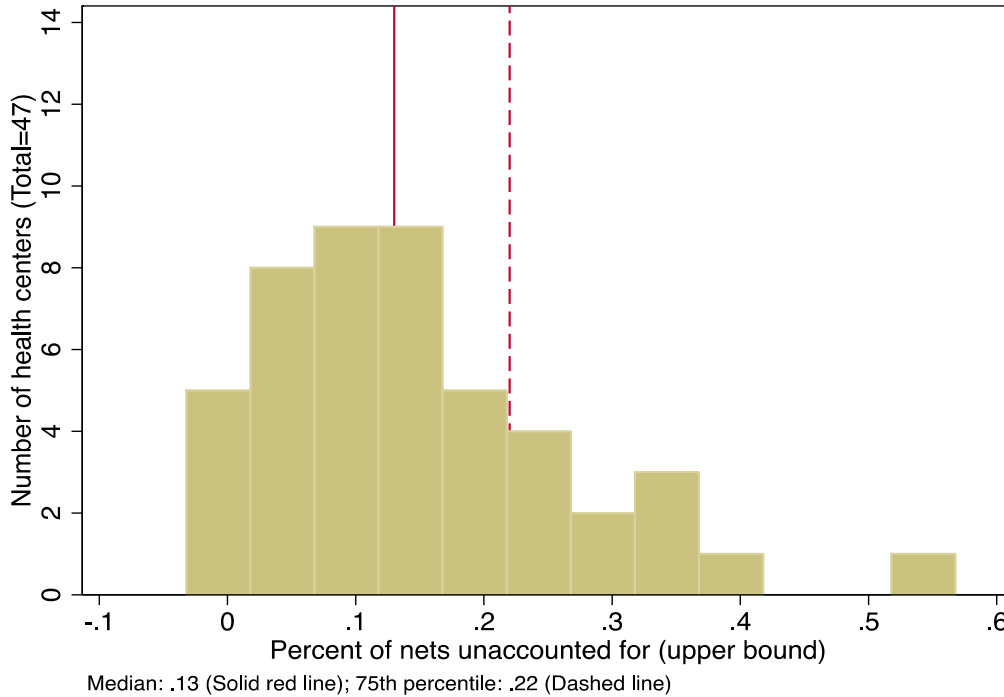


Figure 1B. Percent of nets unaccounted for



Note: Based on 47 health facilities with direct distribution (all in Ghana). For each facility, the estimated number of program nets unaccounted for is the difference between the total number of nets delivered to the facility and the estimated number of "valid" (eligible) beneficiaries listed in the program ledgers. The prevalence of invalid ledger entries was established through audits of randomly selected listed beneficiaries. These audit surveys were misplaced by the survey team for one facility that is why we have only 47 and not 48 observations.

Table A1. Ghana experimental sample: Summary statistics on participating health centers and balance check

	<i>Coeff. Estimate (s.e) on Treatment Dummy:</i>				
	Sample Mean (Std. Dev.)	Voucher	Audit Threat	Small Delivery	Health worker Payment
<u>Panel A: Baseline Characteristics of Health Centers</u>					
# of monthly ANC new registrants ^a	25.73 (20.24)	-0.45 (6.33)	2.56 (4.79)	-1.49 (6.26)	-1.79 (6.26)
# of monthly ANC follow-up visits	90.22 (77.32)	11.67 (24.51)	11.01 (19.41)	1.56 (22.05)	10.40 (22.05)
# of midwives and nurses for ANC	2.01 (1.20)	-0.06 (0.43)	-0.42 (0.29)	-0.29 (0.35)	-0.37 (0.35)
Facility conducts outreach ANC activities	0.23 (0.42)	-0.05 (0.14)	0.00 (0.11)	-0.05 (0.13)	-0.12 (0.13)
Years since facility is operating	17.25 (13.69)	-3.81 (4.49)	-0.43 (3.43)	-1.12 (4.32)	-2.69 (4.35)
Facility is a CHPS compound ^b	0.25 (0.43)	0.02 (0.12)	0.06 (0.11)	-0.12 (0.13)	0.04 (0.13)
Public facility	0.06 (0.23)	-0.02 (0.06)	0.00 (0.05)	0.13 (0.070)*	0.13 (0.070)*
Has a maternity ward	0.85 (0.36)	-0.02 (0.12)	-0.08 (0.09)	-0.04 (0.10)	-0.04 (0.10)
# of other ANC facilities within 10 km radius	2.08 (2.91)	-0.27 (1.10)	-0.50 (0.70)	0.96 (0.90)	-0.21 (0.90)
Has distributed nets in the past	0.07 (0.26)	-0.08 (0.08)	0.03 (0.06)	0.00 (0.08)	-0.08 (0.08)
Accessible during the rainy season	0.81 (0.40)	-0.08 (0.13)	0.11 (0.09)	0.00 (0.11)	0.00 (0.11)
Distance (in km) from region capital	86.29 (49.43)	-1.54 (16.72)	8.49 (12.77)	-10.62 (14.34)	-6.80 (14.34)
Nets available for sale within 10km	0.10 (0.30)	-0.02 (0.11)	-0.03 (0.07)	0.04 (0.09)	0.04 (0.09)
Health worker privately sells nets at facility	0.04 (0.20)	0.06 (0.06)	0.03 (0.05)	0.04 (0.04)	0.04 (0.04)
ANC client Dictator Game: Amount given (out of 10 GHC)	1.91 (1.48)	-0.24 (0.49)	0.05 (0.37)	0.34 (0.42)	-0.16 (0.42)
<u>Panel B: Program Implementation Details</u>					
Phase-in Rank (1 to 6)	3.50 (1.71)	0.40 (0.54)	-0.06 (0.41)	-0.62 (0.51)	0.54 (0.51)
Initial stock of nets delivered	184.03 (146.24)	2.08 (34.13)	12.50 (32.48)	-129.17 (43.800)***	8.33 (43.80)
Total # of staff who attended training	4.59 (2.16)	-0.58 (0.81)	-0.73 (0.52)	0.21 (0.63)	0.20 (0.63)
Share of ANC staff trained on SALI program	0.81 (0.28)	-0.16 (0.080)**	-0.05 (0.07)	0.17 (0.080)**	-0.02 (0.08)
In-Charge present for training	0.68 (0.47)	-0.10 (0.13)	0.08 (0.11)	0.13 (0.14)	-0.29 (0.140)**
Duration of program (days)	109.06 (22.39)	-8.73 (7.80)	-0.89 (5.74)	2.96 (4.23)	3.04 (4.23)
Number of Health Facilities with Treatment	Total N=72	24	24	24	24

Notes: Each row corresponds to one OLS regression. Standard errors in brackets. ***, **, * indicate significance at 1, 5 and 10% levels.

^a ANC stands for Antenatal Care. ^b CHPS stands for Community and Health Planning Services, these are community-based services in remote areas.

Table A2. Correlations between clinic performance and healthworkers characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
Data Source:	Backcheck surveys with random subset of ANC Clients		Community Interviews	"Mystery Client" visits		
Dependent Variable:	Got net/voucher at first visit to clinic	Got net/voucher at any visit to clinic	Respondent thinks that (male) surveyor could get a net at a health center	Acquired program net	High potential opportunity to acquire net	Clinic stocked out
<i>Healthworker characteristics</i>						
Pro-social motivation (standardized index) ^a	0.037* [0.022]	0.038* [0.020]	0.006 [0.008]	-0.030*** [0.011]	-0.045** [0.021]	-0.029 [0.024]
Intrinsic job motivation (standardized index) ^a	0.022 [0.017]	0.015 [0.016]	0.01** [0.006]	0.011 [0.0078]	0.023 [0.021]	-0.013 [0.018]
Extrinsic job motivation (standardized index) ^a	0.038* [0.023]	0.03 [0.021]	-0.0006 [0.009]	0.014 [0.012]	0.014 [0.022]	0.016 [0.026]
Similarity to local population (standardized index) ^b	-0.0021 [0.019]	0.0029 [0.017]	-0.008 [0.009]	0.012 [0.013]	0.03 [0.020]	0.059*** [0.021]
<i>Clinic characteristics</i>						
Above-median ANC registrants	-0.046 [0.040]	-0.038 [0.037]	0.03* [0.02]	-0.038** [0.018]	-0.0019 [0.031]	0.069* [0.036]
Clinic accessible in the rainy season	0.026 [0.077]	0.019 [0.074]	0.01 [0.03]	0.014 [0.021]	0.074** [0.037]	0.018 [0.049]
Total number of staff working in ANC	0.0089* [0.0052]	0.011** [0.0045]	-0.003 [0.003]	-0.00087 [0.0017]	-0.0057 [0.0042]	-0.0063 [0.0055]
Private facility	-0.43*** [0.12]	-0.49*** [0.12]	0.05 [0.05]	-0.063** [0.024]	-0.11* [0.057]	0.078 [0.057]
NGO facility	-0.15** [0.071]	-0.14** [0.068]	0.06* [0.03]	-0.026 [0.042]	0.029 [0.074]	0.16 [0.11]
Observations	2028	2028	2559	766	762	766
R-squared	0.065	0.073	0.014	0.043	0.082	0.133
Mean of the dependent variable	0.76	0.8	0.1	0.043	0.14	0.13

Notes: Each column is a separate regression, with controls for country fixed effects, the experimental treatments in Ghana, and healthworker age and gender. Columns 1 and 2 also control for whether there were any stockouts in those clinics, as measured by the attendance checking team in Uganda and Kenya (there is no control in Ghana where stockouts were determined by healthworker effort). Standard errors clustered at the level of the clinic in brackets. ***, **, * indicates significance at 1, 5 and 10 percent. Ghana sample: Direct distribution clinics only.

^a Indexes are averages across standardized variables as grouped in Tables 7 and 8, which are then re-standardized to have a mean of 0 and std dev of 1. The pro-social motivation index also contains the share left in the dictator game (shown in Fig A2).

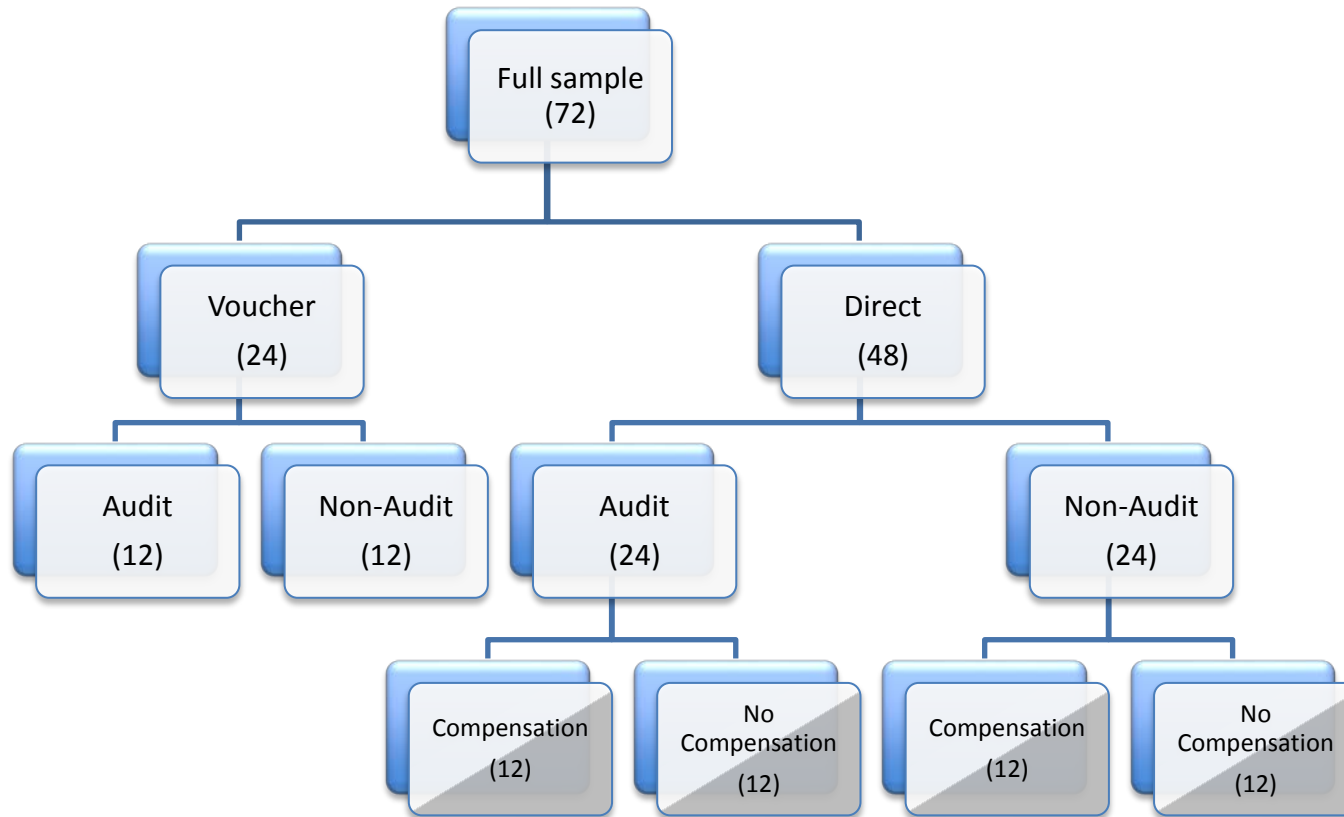
^b The similarity index is the average across standardized versions of three variables: (1) percentage of ANC registrants who have the same ethnicity as the health worker; (2) indicator for being born in same region as clinic; (3) indicator for being born in same district as clinic.

Table A3. Correlations between performance measures and absenteeism

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Data Source:	Backcheck surveys with random subset of ANC Clients			Community Interviews	"Mystery Client" visits			
Dependent Variable:	Got net/voucher at first visit to clinic	Got net/voucher at any visit to clinic	Payment requested for net (conditional on offer)	Respondent thinks that (male) surveyor could get a net at a health center	Payment requested for net	Acquired program net	High potential opportunity to acquire net	Clinic stocked out
<i>Healthworker characteristics</i>								
Share absent	0.0015 [0.091]	0.00028 [0.088]	-0.006 [0.008]	-0.03 [0.03]	-0.025 [0.023]	0.0041 [0.063]	0.025 [0.12]	-0.023 [0.16]
<i>Clinic characteristics</i>								
Above-median ANC registrants	-0.033 [0.039]	-0.026 [0.035]	0.008 [0.008]	0.03 [0.02]	-0.0087 [0.014]	-0.046** [0.018]	-0.026 [0.028]	0.026 [0.032]
Clinic accessible in the rainy season	0.019 [0.079]	0.012 [0.078]	0.0008 [0.006]	0.01 [0.03]	0.015 [0.024]	0.0064 [0.021]	0.054* [0.028]	-0.02 [0.045]
Total number of staff working in ANC	0.0048 [0.0052]	0.0071 [0.0043]	0.0002 [0.0006]	-0.002 [0.003]	-0.0052** [0.0024]	-0.002 [0.0019]	-0.0066 [0.0042]	-0.0053 [0.0046]
Private facility	-0.46*** [0.16]	-0.51*** [0.16]	0.004 [0.008]	0.04 [0.05]	-0.04 [0.037]	-0.054*** [0.019]	-0.066 [0.042]	0.11*** [0.036]
NGO facility	-0.19** [0.079]	-0.17** [0.076]	0.07** [0.03]	0.06* [0.03]	-0.018 [0.021]	-0.036 [0.042]	-0.014 [0.079]	0.12 [0.11]
Observations	2028	2028	1549	2559	765	766	762	766
R-squared	0.078	0.082	0.034	0.01	0.016	0.035	0.086	0.124
Dep. Var. Mean	0.76	0.8	0.01	0.1	0.041	0.043	0.14	0.13

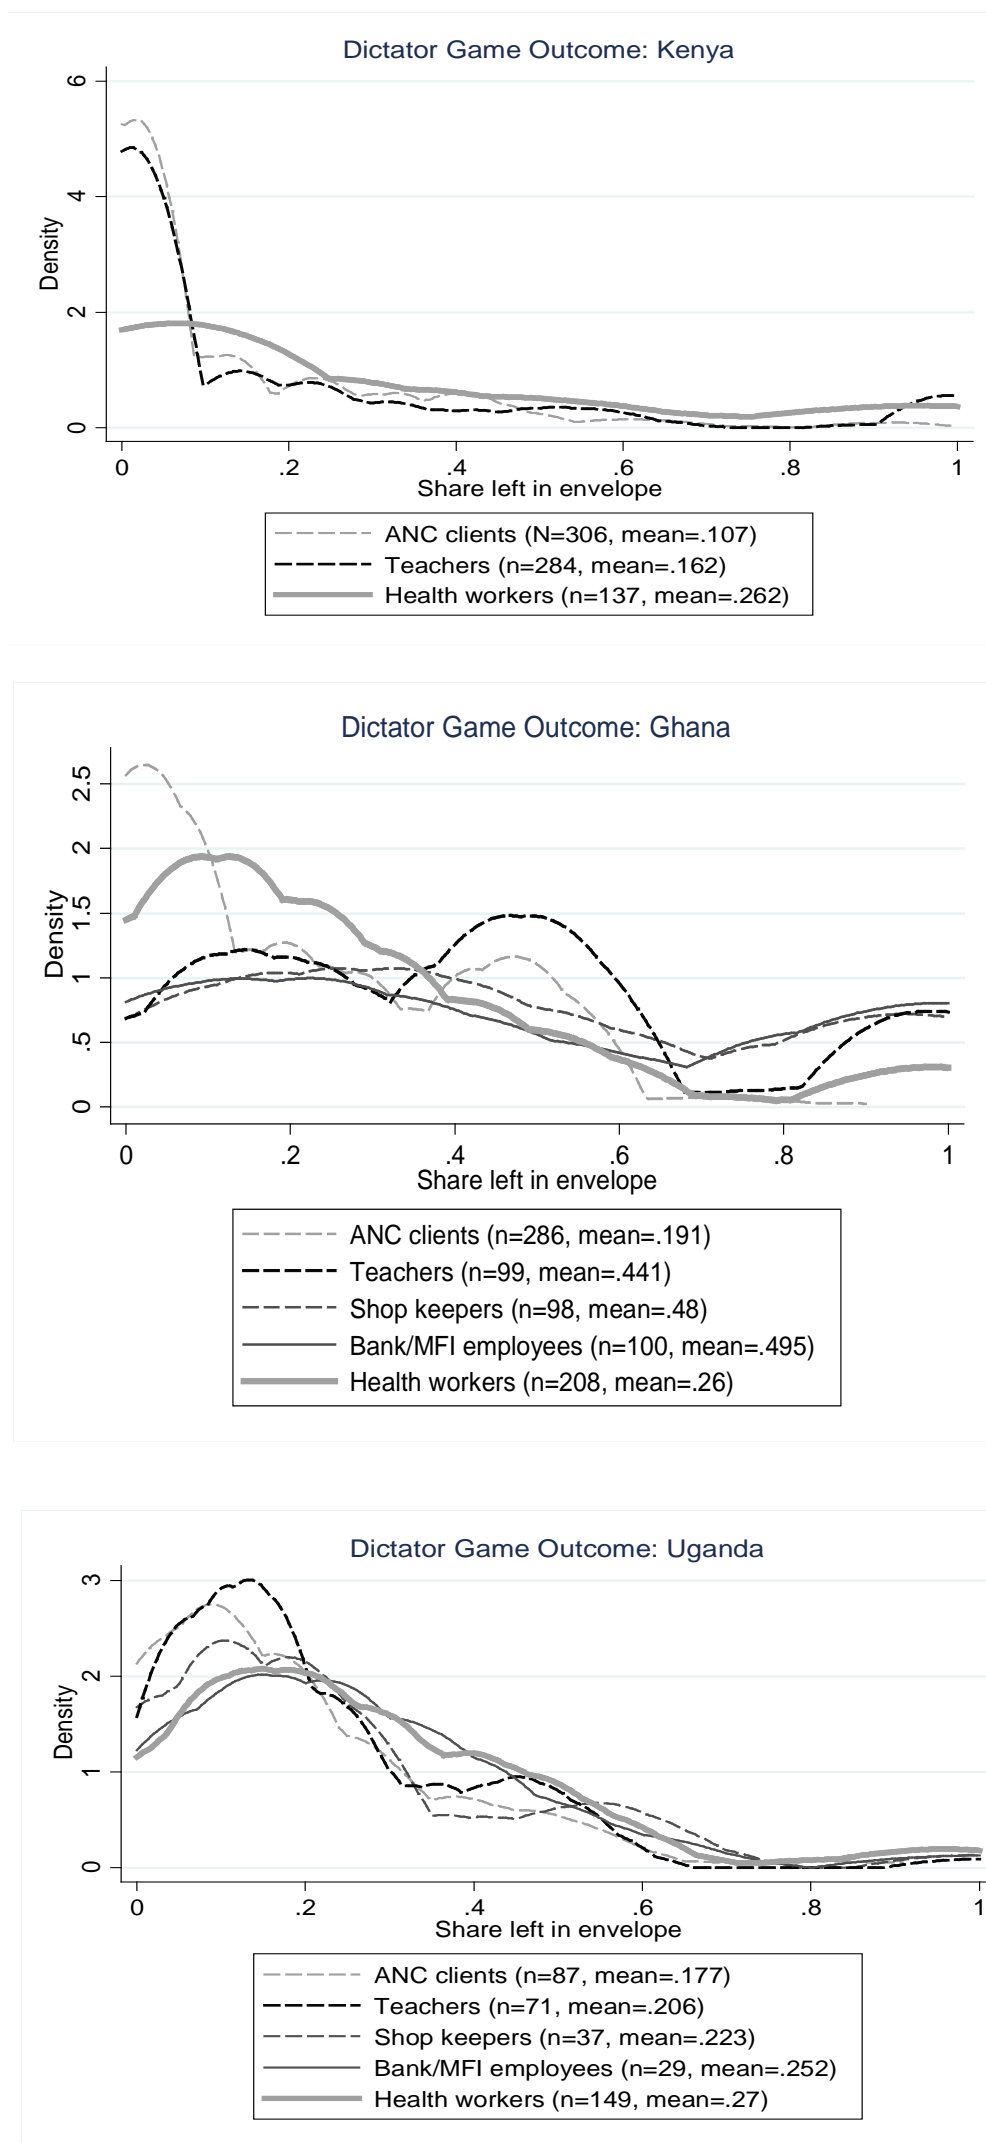
Each column is a separate regression, with controls for country fixed effects, the experimental treatments, and whether the healthworkers had the "honesty test" as a part of their dictator game. Columns (1) and (2) also control for whether there were any stockouts in those clinics, as measured by the attendance checking team in Uganda and Kenya (there is no control in Ghana where stockouts were determined by healthworker effort). Standard errors clustered at the level of the health facility in brackets. ***, **, * indicates significance at 1, 5 and 10 percent. Ghana sample includes only facilities sampled for direct distribution.

Figure A1. Ghana experimental design



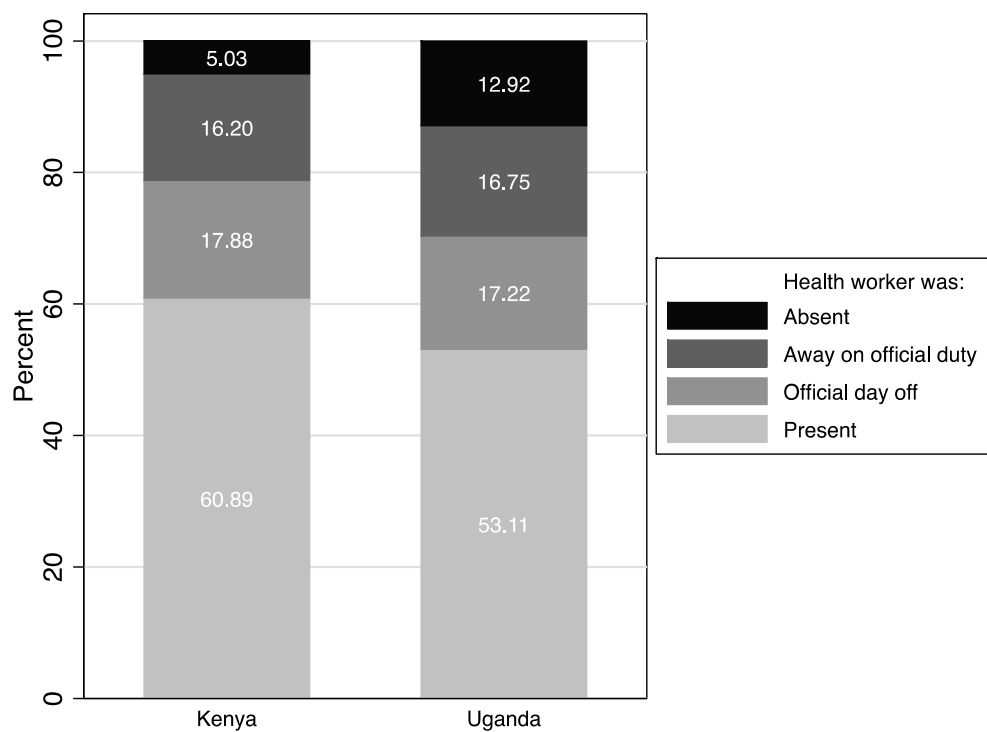
Notes: Sample size numbers are given in terms of number of clinics. Gray shading means split 50/50 between small and large delivery.

Figure A2. Other-regarding preferences of health workers compared to others



Notes: The dictator game allowed the participant to leave as much money as they wanted, anonymously, for a community member, out of an envelope with 10 bills.

Figure A3. Health worker attendance data



Notes: Individual-level data from one unannounced spot check. Kenya: 188 health workers from 48 facilities. Uganda: 214 health workers from 48 facilities. No unannounced spot check were performed in Ghana.