

LOSING PROSOCIALITY IN THE QUEST FOR TALENT? SORTING, SELECTION, AND PRODUCTIVITY IN THE DELIVERY OF PUBLIC SERVICES *

Nava Ashraf
Oriana Bandiera
Edward Davenport
Scott S. Lee

April 20, 2019

Abstract

We embed a field experiment in a nationwide recruitment drive for a new healthcare position in Zambia to test whether career benefits attract talent at the expense of prosocial motivation. In line with common wisdom, offering career opportunities attracts less prosocial applicants. However, the trade-off only exists at low levels of talent; the marginal applicants in treatment are more talented and equally prosocial. These are hired, and perform better at every step of the causal chain: they provide more inputs, increase facility utilization, and improve health outcomes including a 25% decrease in child malnutrition.

JEL classification: J24, 015, M54, D82.

* Ashraf: Department of Economics, LSE, n.ashraf1@lse.ac.uk. Bandiera: Department of Economics and STICERD, LSE, o.bandiera@lse.ac.uk. Davenport: Department of Economics, MIT, edav@mit.edu. Lee: Department of Medicine and Institute for Global Health, Vanderbilt University, ssl@mail.harvard.edu. We thank the Ministry of Health of Zambia and especially Mrs. Mutinta Musonda for partnership on this project. We thank the IGC, JPAL Governance Initiative, USAID and HBS DFRD for financial support. We also thank Robert Akerlof, Charles Angelucci, Tim Besley, Robin Burgess, Paul Gertler, Edward Glaeser, Kelsey Jack, Giacomo Ponzetto, Imran Rasul, Jonah Rockoff and seminar participants at several institutions for useful comments. Adam Grant, Amy Wrzesniewski, and Patricia Satterstrom kindly provided guidance on psychological measurement. We thank Kristin Johnson, Conceptor Chilopa, Mardieh Dennis, Madeleen Husselman, Alister Kandyata, Allan Lalian, Mashekwa Maboshe, Elena Moroz, Shotaro Nakamura, Sara Lowes, and Sandy Tsai for the excellent research assistance and the Clinton Health Access Initiative in Zambia for their collaboration.

1 Introduction

Economic development entails the professionalization of public service delivery: qualified, career professionals replace informal providers with strong connection to beneficiaries, willing to work for little pay (Northcote and Trevelyan 1853; Weber 1922; North 1991). This shift fosters an identity based on a career in the civil service, but agents with altruistic preferences towards beneficiaries may be necessary for effective public service delivery (Akerlof and Kranton 2005; Besley and Ghatak 2005; Prendergast 2007; Brehm and Gates 1999; Wilson 1989). To what extent does professionalization crowd out those who care most deeply about their beneficiaries, and what impact does it have on welfare?¹

This paper provides the first experimental evidence on whether these identities attract different agents and whether this selection determines the effectiveness of service delivery. Does a career in the civil service attract talent at the expense of prosociality? Does it draw in agents who underperform relative to those attracted by “doing good”?² Or does it attract agents who are neither talented nor prosocial but attracted by the opportunity to extract rents? Understanding the conditions under which such trade-offs occur is critical to informing theory and to settling the policy debate on whether rewards for service delivery agents should be kept low so as to screen out individuals without altruistic preferences.

We design a nationwide recruitment experiment to identify causal impacts on the applicant pool, on the recruited agents, and on their performance. We collaborate with the Government of Zambia as they formalize primary health care in remote rural areas by creating a new health worker position in the civil service. This cadre is meant to replace informal service provision by religious and other charitable organizations, thereby following the typical evolution of the modern state. The stakes are high because, due to the shortage of medical staff, hiring effective agents can make a great difference for the quality of health services and, ultimately, health outcomes in these communities.

Our experiment varies the salience of a career in civil service at the recruitment stage, exploiting the fact that this position is new to potential applicants. In control districts, the recruitment ads reflect the status quo before the new position, when local health services were provided by individuals hired by non-governmental and charitable organizations. Helping the community is listed as the main benefit, and local agents are listed as peers. In treatment districts the ads are designed to highlight the civil servant identity: career advancement is listed as the main benefit, and doctors and nurses are listed as peers. Treatment and control posters differ only in the salience of career opportunities, while all other factors such as application requirements are kept equal.

To isolate the effect of selection on performance, we must sever the link between treatment and the marginal return to effort on the job. To this end, all hired agents are given the same information

¹This echoes the tension between intrinsic and extrinsic motivation on the job (Bénabou and Tirole 2003, 2006).

²Ambition towards a career in public service, with both its ability to attract the most able but also the most self-interested, has a long intellectual history; ambition was used by Romans, as *ambitio*, exclusively to refer to those in public life. In *De officiis*, Cicero referred to *ambitio* as a “malady” that can cause individuals to “lose sight of their claims to justice,” but one that seems to draw “the greatest souls” and “most brilliant geniuses”.

on career opportunities and social impact when they move to the same training school where they are trained together for one year, before deployment. A survey administered before and after the training program validates our design: before training, treatment and control agents differ in the perceived relevance of career benefits, but after training these perceptions converge.

The new health worker position effectively adds career opportunities to a job with social impact. A simple conceptual framework makes precise that, in line with prevailing policy concerns, this attracts applicants who are less prosocial conditional on a given level of talent. However, since the outside option is increasing in talent, adding career benefits will draw in more talented individuals, and the marginal, most talented applicant in both groups will have the highest prosociality. The treatment effect on recruited candidates will therefore depend on how candidates are chosen from the pool. If applicants are drawn randomly, there might be a trade-off between talent and prosociality. However, if only the most talented are hired, there will be no trade-off.

To evaluate the impact of treatment on the applicant pool and on hired candidates we collect information on the skills and prosociality of every applicant. This exercise reveals that, in line with common intuition, the average applicant in treatment is more talented and less prosocial. In line with the theoretical intuition, however, the most talented applicants have the same, high level of prosociality. We show that the selection panels in both treatment and control put a high weight on talent, leading them to recruit among the most talented in their pool; as a result, treatment recruits are more talented and equally prosocial.

To evaluate the impact of treatment on service delivery we combine three data sources: real-time data on service delivery in remote areas collected through a mobile platform, administrative data on health facility utilization, and our own survey of household health practices and outcomes, including immunization records and anthropometrics. This allows us to link the services delivered by the newly recruited health workers to the outcomes of the households who receive those services and, ultimately, their health impact.

We find that agents drawn by career opportunities are more effective at each step of the causal chain from the inputs they provide to the outcomes of the recipients. They provide more inputs (29% more household visits, twice as many community meetings) at the same cost. They increase facility utilization rates: the number of women giving birth at the health center is 30% higher, and the number of children undergoing health checks 24% higher, being weighed 22% higher, and receiving immunization against polio 20% higher. They improve a number of health practices among the households they serve: breastfeeding and proper stool disposal increase by 5 and 12 percentage points (pp), respectively, deworming treatments by 16%, and the share of children on track with their immunization schedule by 5pp (relative to a control mean of 6%). These changes are matched by changes in objective health outcomes: the share of children under age 5 who are underweight falls by 5pp.

Taken together, these results indicate that offering a civil service position with career opportunities attracts agents who deliver services with remarkable health impact in the communities.

The fact that we observe consistently positive impacts from three distinct and entirely independent data sources further strengthens our confidence in the findings.

In light of the evidence of poor bureaucratic performance in low-income countries (Collier 2009; Muralidharan et al. 2011) our findings suggest that this is not due to the fact that civil service careers attract poor performers when these jobs are first created.

The study of how individuals sort into jobs according to their preferences, skills, and the jobs' own attributes has a long tradition in economics (Roy 1951). More recently this has been enriched by the study of job missions as a selection and motivation mechanism (Besley and Ghatak 2005) and identity or self-image as components of preferences (Akerlof and Kranton 2005; Bénabou and Tirole 2011). Our findings provide empirical support to these contributions as we show that the identity associated with the job affects those drawn to it and that this selection affects performance.

The fact that career opportunities affect performance *through selection* complements the recent findings of Bertrand et al. (2016) that, on the intensive margin, better promotion prospects improve the effectiveness of Indian civil servants. Our findings also complement a large literature on the impact of financial incentives. On the selection margin, Dal Bó et al. (2013) and Deserranno (2019) study the effect of earnings levels on the traits of applicants for government and NGO jobs³ while several papers evaluate the effect of performance pay on the performance of agents after these have been hired either for the delivery of health services (Ashraf et al. 2014; Miller et al. 2012; Miller and Babiarz 2014; Celhay et al. 2019) or education (Muralidharan and Sundararaman 2011; Duflo et al. 2012; Glewwe et al. 2010; Fryer 2013; Rockoff et al. 2012; Staiger and Rockoff 2010). Our contribution is to provide the first experimental evidence that selection affects performance in public service delivery. In particular, we show that job design, of which incentives are a component, affects who sorts into these jobs in the first place, and that the effect of this selection on performance is of the same order of magnitude as the largest incentive effects estimates.^{4,5}

The rest of the paper is organized as follows. Section 2 describes the context and research design. Section 3 develops a conceptual framework to make precise the trade-off between talent and prosociality, and Section 4 tests for it in the applicant pool and among recruited candidates. Section 5 evaluates the treatment effect on performance in delivering health services. Section 6 evaluates the treatment effect on facility utilization, health behaviors, and health outcomes. Section 7 concludes with a discussion of external validity, welfare implications, and general equilibrium effects relevant for program scale-up.

³Dal Bó et al. (2013) find that higher salaries for civil service jobs attracts better qualified candidates with the same level of prosocial preferences. Deserranno (2019) finds that expectations of higher earnings discourage prosocial candidates from applying for an NGO job that encompasses both commercial and health promotion activities. While consistent with these selection effects, our experiment focuses on measuring the effect of selection on agents' performance and beneficiaries' outcomes, which encompasses the effect of all the attributes that determine effectiveness.

⁴There is a corresponding literature that studies the same issues in the private sector. This literature stresses the importance of the effect of incentives on selection, but empirical studies focus on incentives on the job (Lazear and Oyer, 2012; Oyer and Schaefer, 2011).

⁵Rothstein (2015) uses a model-based approach that simulates the selection effect of alternative teachers' contracts. He finds that bonus policies have small effects on selection while reductions in tenure rates accompanied by substantial salary increases and high firing rates can have larger effects.

2 Context and Research Design

2.1 Context: Health Services in Rural Communities

Delivering health services to remote rural areas is challenging at every level of development because trained medical staff are reluctant to be posted there and turnover rates are high (World Health Organization 2006).⁶ In Zambia, the average health post (the first-level government health facility) had 1.5 staff from the Ministry of Health, including those not permanently based there. The government community health assistant (CHA) position was created as a solution to this challenge. The position is meant to put formally trained government health staff in place of informal community health workers employed, often as volunteers, by religious and other non-profit organizations.⁷ In fact, the title of the newly created government position was changed to “Community Health Assistant” shortly after its inception precisely to distinguish it from the informal, untrained community health workers.

In 2010, the program’s first year, the Government sought to recruit, train, and deploy two health workers to each of 167 communities in 48 districts. The main task of CHAs is to visit households and refer them to health facilities as needed. The job requires both medical and social skills (World Health Organization 2006). Medical skills include weighing, taking vital signs, filling out patient registries, and determining whether a patient is pregnant. Social skills include counselling, supporting, advising, and educating patients and other laypeople.

Government-funded community health worker programs vary in the extent to which they integrate the health workers into the civil service. At one extreme there are programs that mimic the informal model with financing provided by the government and all other decisions including hiring, monitoring, and firing left to the non-governmental sector. At the other extreme is the model adopted in this program in Zambia where health workers become a cadre of civil servants and can advance to higher-ranked and better-paid cadres. The pay gradient is steep as the starting monthly wage is USD 290 for CHAs and USD 530 for entry-level nurses.⁸ Promotion into higher-ranked

⁶The US Health Resources and Services Administration estimates that 1 in 2 rural Americans lives in a medically underserved area with a shortage of primary care providers (PCPs), defined as a population-to-PCP ratio of greater than 3,500:1 (US Human Resources and Service Administration 2019).

⁷The history of community health work goes back at least to the early 17th century, when a shortage of doctors in Russia led to training community volunteers in providing basic medical care to military personnel. This role later became formalized with China’s “barefoot doctors,” laypeople who sometimes could not afford shoes but were trained to meet primary health needs in rural areas, and then became widespread in Latin America, in underserved areas in the United States, and, more recently, across Africa (La Familia Sana Program 1992; Perez and Martinez 2008). The original programs emphasized community self-reliance and participation. Like much of informal public service delivery, for example in the United Kingdom in the 18th and 19th centuries, these are provided by religious institutions, grassroots movements, and, more recently, non-governmental organizations. For this reason, however, they are often uncoordinated, lower-skilled efforts.

⁸At the time of the launch of the recruitment process in September 2010, the Government had not yet determined how much the health workers would be formally remunerated. Accordingly, the posters did not display any information about compensation. Although the health worker wage was unknown to applicants at the time of application (indeed, unknown even to the Government), applicants would likely have been able to infer an approximate wage, or at least an ordinal wage ranking, based on the “community health” job description and the relatively minimal educational qualifications required—both of which would intuitively place the job below facility-based positions in compensation. In Section 2.2, we present evidence against the hypothesis that wage perceptions may have differed by treatment.

cadres within the Ministry requires additional training (for example, nursing or medical school). Being part of the civil service, the health workers are eligible for in-service training, meaning that they attend school while continuing to receive salary and in some cases receive sponsorship of tuition as well. The official policy of the Ministry is to periodically ask the district medical officers to nominate a number of candidates on merit, but there is no mechanical link between quantitative measures of performance (say, the number of visits that a health worker makes) and nominations. Promotions to higher cadres are therefore not automatic, but the expected payoff is high even with low success rates, especially because job opportunities that allow for a career in central government are rare in the remote communities from which the health workers are recruited.⁹

The Government chose the latter model in the hope of attracting agents with strong technical skills to do community work. Nevertheless, they were fully aware that the focus on career advancement could backfire by crowding out applicants motivated to help the community.¹⁰ The possibility of this trade-off led to the experiment we describe below.

2.2 Experimental Design

Our experiment aims to assess whether a career in civil service attracts talent at the expense of prosociality, and whether this affects who is hired and their performance. This is relevant to evaluate the role of selection in public service delivery beyond health services in low-income countries, as the concern that material rewards attract the wrong types is pervasive. The key challenge is to separate the effect of selection from the effect of incentives on the job. We tackle this in two steps: the first opens the selection channel, and the second shuts down the incentive channel.

Experimental Design, Step I: Opening the Selection Channel

To open the selection channel we use the recruitment posters and the information materials distributed to health officers. In each community, paper advertisements for the job were posted in local public spaces, such as schools, churches, and the health post itself. District health officials were responsible for ensuring that the recruitment posters were posted.

To ensure that the recruitment process was carried out in a uniform manner across all the communities, the Government included detailed written instructions in the packets containing the recruitment materials (posters, applications, etc.) that were distributed to district health officials (see Appendix F).

⁹As of 2019, Ministry of Health records show that five of the original 307 CHAs have applied for higher-level professional training (two in registered nursing, two in public health nursing, and one in registered midwifery). Of these five, three were accepted by their training schools and granted study leave by the Ministry, with the understanding that, upon completing their programs, they would cease to work as CHAs and instead assume their new professional roles.

¹⁰Mr. Mwila, then Human Resources Director at the Ministry of Health, expressed this trade-off clearly when he asked us: “What is going to happen now that they (potential health workers) will see themselves as civil servants? Will they be connected to the community?”

The treatment poster stresses the civil service identity of the new position. It lists as the main benefit of the job the opportunity to ascend the civil-service career ladder to higher and better-paid positions such as environmental health technician, nurse, clinical officer, and doctor. This incentive is summarized in a bold caption stating, “Become a community health worker to gain skills and boost your career!” The poster also explicitly leverages a sense of belonging to the civil service by stating “become a highly trained member of Zambia’s health care system”. Finally it sets “experts in medical fields” as the peer group.

The control poster uses the standard approach of recruiting community health workers, stressing the social identity of the position by making salient community impact such as “[gaining] the skills you need to prevent illness and promote health for your family and neighbors”. The message is summarized in a caption stating, “Want to serve your community? Become a community health worker!” Finally, it lists local health post staff as the peer group candidates can expect to interact with.

Three points are of note. First, the social identity poster functions as control because the status quo community health worker jobs do not offer career opportunities. Second, treatment and control posters have exactly the same structure except the wording of the benefits. We chose this over a “neutral” control poster with no benefits whatsoever because in that case, the treatment effect would conflate the effect of interest with the effect of advertising benefits per se. While this might be of intrinsic interest, it would not allow us to answer the more general question of how agents who are attracted by a career in the civil service differ from those attracted by social impact and how this selection affects performance. Third, it is important to note that in these communities government jobs are scarce; therefore, a poster advertising a government job is likely to be highly visible.

Since recruitment was organized by district officials, we randomized treatment at the district level in order to maximize compliance with the experimental assignment, evenly splitting the 48 districts into two groups. This implies that each district official is only exposed to one treatment and is unaware of the other. As district officials are the main source of information for aspiring health workers, randomization at the district level minimizes the risk of contamination. Randomization at the district level also mitigates the risk of informational spillovers between communities, as the distance between health posts in different districts is large. Random assignment of the 48 districts is stratified by province and average district-level educational attainment.¹¹ To ensure compliance with the randomization protocol, we worked closely with the Government to standardize the information given to the district officials to organize the recruitment process.¹²

¹¹We stratify by the proportion of adults in the district who have a high school diploma, as reported in the most recent Living Conditions Monitoring Survey, conducted by the Central Statistical Office four years prior in 2006. We sort districts by province and, within each province, by high school graduation rate. Within each sorted, province-specific list of districts, we take each successive pair of districts and randomly assign one district in the pair to the career opportunities treatment and the other to the control group. For provinces with an odd number of districts, we pool the final unpaired districts across provinces, sort by educational attainment, and randomize these districts in the same pair-wise manner.

¹²District officials are given a packet containing 10 recruitment posters and 40 application forms for each health post and are asked to distribute each packet to the respective health center and, from there, to ensure that recruitment

Table A.1 reports balance tests on three sets of variables that can affect the supply of health workers, the demand for their services, and their working conditions. Overall, Table A.1 shows that the new health workers are recruited from similar areas and will work in similar areas. Besides showing balance between treatment and control, this exercise is useful to understand labor markets in rural Zambia. Two findings are of note. First, only 4.4% of the population have the necessary credentials (grade 12 education) to apply. Second, and more strikingly, just over half of the eligible were either out of work or in unpaid employment over the past twelve months.¹³ Among the 45% engaged in income-generating activities (either as employees or self-employed), fewer than one-third are employed in high-skill occupations (such as teachers, which account for 9% of the eligible population), and about half are employed in low-skill occupations, mostly in agriculture, which accounts for 18% of the eligible population. Given the scarcity of skilled jobs, the program might have the added benefit of creating job opportunities in these communities.

Experimental Design, Step II: Closing the Incentive Channel

To close down the incentive channel, all successful applicants are eligible for career opportunities once hired. After being recruited, all agents train together for one year, during which they receive the same information about the career opportunities they are entitled to as civil servants. As treatment and control health workers face the same incentives once hired, performance differences, if any, are attributable to selection.

The experiment aims to create differences in career opportunities at the application stage and then to eliminate these differences after candidates have been hired. To check whether it succeeded, we ask all agents about perceived benefits of the job when they first arrive at the training school in June 2011 and then again twenty-two months later in April-May 2013, which is after they have completed the one-year training. To elicit this information, we give each health worker a bag of 50 beans and ask them to allocate the beans to different cards describing potential benefits of the job. This method has two desirable features: (i) it forces respondents to take into account the trade-off between different benefits, namely that giving more weight to one benefit necessarily implies that other benefits will be given less weight, and (ii) it allows us to test whether the treatment affected other benefits besides career advancement and community service.

There are two sources of potential desirability bias, which might affect the magnitude of the treatment effects but not their sign. First, the fact that respondents say what they think the

posters are posted, application forms are made available, and so forth. We conduct a series of follow-up calls over several weeks to the district point-persons to ensure that the recruitment process is conducted as planned. To reinforce the treatment, we also include a basic written script that the district officials are invited to use to inform health centers and neighborhood health committees on the health worker program and recruitment process. In the career opportunities treatment, the script describes the new program as follows: “This is an opportunity for qualified Zambians to obtain employment and to advance their health careers. Opportunities for training to advance to positions such as Nurse and Clinical Officer may be available in the future.” In contrast, in the control group, the script states, “This is an opportunity for local community members to become trained and serve the health needs of their community.” (see Appendix F).

¹³The 28% who were out of work are either unemployed (13%), housewives (7.5%), or full time students (8.5%). Most (65%) of the unpaid jobs are in agriculture. These are balanced across treatments.

enumerators want to hear based on the information given on the posters does not invalidate this exercise; the aim of the exercise is precisely to test whether the information they have matches that given on the posters. Second, the fact that this is a community-based position, named “Community Health Worker,” might lead the health workers to overstate community benefits. This will bias the share put on community benefits upwards and the difference between treatments downwards, making it less likely for us to be able to detect a difference between treatment and control. This should be kept in mind when interpreting the magnitudes reported below.

The answers tabulated in Table A.2 show that differences in the perceived benefits reported by the health workers when they first arrive at the training school match those advertised in treatment and control posters and then disappear after the health workers are exposed to the training program. Table A.2, Panel A, shows that service to the community is listed as the main benefit in both groups. This might truly reflect preferences or be inflated by desirability bias as discussed above. Despite the fact that this biases treatment effects towards zero, we find that the treatment group places 38% more weight on career opportunities ($p=0.002$) and lower weight on both “allows me to serve the community” and “earn respect and status in the community” ($p=0.046$ and $p=0.031$, respectively). All other benefits are balanced across groups, suggesting that the poster did not convey different expectations about pay on the job or the nature of the job.

Table A.2, Panel B, shows that the answers converge after exposure to training and that there are no significant differences between the two groups. In line with the fact that control health workers receive information about career opportunities during training, the weight they give to career opportunities rises by 27%, while the weight they give to service to the community falls by 15%. In contrast, treatment health workers, who receive no new information during training, do not change their answers.

The experimental design allows us to identify the effect of career opportunities on performance through selection if the salience of career opportunities at the recruitment stage does not affect the agents’ behavior directly once the real career opportunities are known by both treatment and control health workers. This assumption fails if control agents react to the difference between advertised and actual benefits, rather than to the benefits themselves. If control agents value career benefits this will bias the treatment estimates downwards as they might respond to the positive surprise by working harder. Symmetrically, estimates will be biased upwards if control agents dislike career benefits or dislike finding out that the actual value of career opportunities is larger than the value advertised. Note that in this case, agents for whom the participation constraint is met ex-ante but not ex-post would drop out once hired. For instance, Deserranno (2019) finds that NGO health promoters who receive a negative surprise on earnings are 14pp more likely to drop out than those who do not over a two-year period. In contrast, the drop out rate of control CHAs was zero over the same period, and only 3% of those who started training did not complete it.

3 Framework

The experiment can be modelled as a three-stage game. In Stage 1, potential applicants choose whether to apply based on the information conveyed by the posters. In Stage 2, the panels select the CHAs from the applicant pool. In Stage 3, the selected applicants choose how much effort to exert on the job. Our treatment occurs at Stage 0, where we assign different posters to different districts so the information applicants have in Stage 1 depends on whether their district is in treatment or control.

Applicants differ in ability, $a \geq 0$, and social preferences towards the community, $s \in [0, 1]$, or prosociality for short. Below we make precise why both are desirable for successful applicants. General ability a comprises all cognitive (IQ) and non-cognitive (ambition, tenacity, work ethic) skills that make individuals productive in all occupations. Prosociality determines the utility individuals get from helping others, in this case by improving their health status. We assume that ability and prosociality are independently distributed in the population. If selected as a CHA, an individual exerts effort, $e \geq 0$, to produce community health according to $H(a, e)$ which is increasing in a and e , generating utility $sH(a, e)$ net of disutility of effort $d(e)$. We assume that output is concave in effort, so $H_{ee} \leq 0$, and also that ability and effort may be complementary, i.e. $H_{ae} \geq 0$, reflecting the idea that a given unit of effort for a highly talented CHA may generate more health output than a given unit of effort for a less talented CHA. To ensure an interior solution, we set $d' > 0$ and $d'' > 0$. All selected CHAs also receive material benefits equal to M , which can be thought of as reflecting the discounted sum of future wages; these accrue to all agents regardless of performance. We proceed via backwards induction.

3.1 Stage 3: Selected Applicants' Choice of Effort

If hired, the agent chooses e to maximize $U(a, s, e, M) = sH(a, e) - d(e) + M$. By the convexity of d , this yields an interior $e^*(a, s) \geq 0$ that satisfies $\frac{de^*}{ds} > 0$ and $\frac{de^*}{da} \geq 0$. This implies a realized health output function $H^*(a, s) \equiv H(a, e^*(a, s))$ with $H_a^* > 0$ and $H_s^* > 0$. This means that both ability and prosociality contribute to generating health output through different channels: prosociality increases effort, while ability increases output directly for any given level of effort and may also increase effort if $H_{ae} > 0$. This is the foundation of the trade-off: if career benefits lead to hiring CHAs with higher a but lower s , the net effect on health output is ambiguous. Finally, substituting realised health output into the utility function gives indirect utility $U^*(a, s, M) \equiv U(a, s, e^*(a, s), M) = sH^*(a, s) - d(e^*(a, s)) + M$, with $U_a^* > 0$, $U_s^* > 0$, and $U_M^* > 0$.

3.2 Stage 2: Panel's Choice of Applicants

In Stage 2, the panel observes applicants with different ability and prosociality. By conducting interviews, and using background information (such as O-level scores), the panel is able to observe these with some noise, so for individual i the panel observes $\tilde{a}_i = a_i + \epsilon_{ai}$ and $\tilde{s}_i = s_i + \epsilon_{si}$, where ϵ_i is individual-specific noise. For simplicity, we abstract from the exact optimization problem and,

as both effort and health output are increasing in a and s , simply assume that the panel follows some rule that picks candidates with high levels of \tilde{a} , \tilde{s} , or both. Even aside from the direct effect of s on effort and thus health output, we shall see below that ability and prosociality are positively correlated among the applicants with high ability and thus the panel may also take s into account as a secondary signal to improve prediction of a .

3.3 Stage 1: Applicant's Decision to Apply

In Stage 1, applicants decide whether or not to apply. If hired, an individual with traits (a_i, s_i) will receive utility in Stage 3 of

$$U^*(a_i, s_i, M)$$

To apply, an individual needs to pay cost c and whether she is hired depends on how her ability and prosociality compare to those of the other applicants. As the panel observes these with noise, an individual applicant i can have a non-zero probability of being selected even if they do not have the highest levels of a_i and s_i , and thus we can observe applicants with a range of a_i and s_i in equilibrium. In addition, the probability of selection for any applicant i will also depend on the total number of applicants (as the number of CHA positions is fixed), which is an equilibrium quantity that will depend on M . We thus write the equilibrium probability as $p(a_i, s_i, M)$, where the panel's decision rule means that p is increasing and concave in both a and s .¹⁴ We shall see later that p is decreasing in M . Taking into account the probability of selection in Stage 2 and the utility in Stage 3 if selected, individual i will apply if the expected utility of applying to be a CHA net of application costs exceeds the utility in their next best alternative occupation, which, in this setting, is mostly self-employment in agriculture or small trade where the agent is the residual claimant. We denote this by $V(a)$ and assume, as is standard, that the marginal return to ability is higher in the private sector $V_a > U_a^* > 0$ for every a . This is the empirically relevant case because, as is common in the public sector, CHAs' earnings are not linked to performance, while self-employed agents in the private sector are the residual claimants on the value they create.¹⁵ Thus, individual i applies if and only if $E(a_i, s_i, M) = p(a_i, s_i, M)U^*(a_i, s_i, M) - c > V(a_i)$. To capture the fact that in practice there are minimum qualification requirements, we assume that application costs are high enough that $E(0, s, M) < V(0)$ for any s and M , so that low-ability individuals who have little chance of being hired do not apply. Formally, there is a threshold of ability $\underline{a}_i(s_i, M)$ such that all i with $a_i < \underline{a}_i(s_i, M)$ do not apply. We shall suppress the subscript on the threshold, but it is important to note that it differs across individuals as it depends on their level of s_i . The Appendix shows that the structure of the solution depends on whether $E(a_i, s_i, M) > V(a_i)$ for

¹⁴We assume continuity of this and related equilibrium quantities (i.e. that there are sufficiently many applicants) for analytical simplicity, although this is clearly not literally true.

¹⁵Young (2012) runs Mincerian regressions for 24 countries (including Zambia), of which 14 are in sub-Saharan Africa, using a range of methods that yields estimates of returns to education of between 9.5% and 11.6%. His preferred estimate, used in the remainder of his paper, is 11.6%. The estimates are qualitatively unchanged when restricting the sample to Zambia only (results available upon request).

all $a_i > \underline{a}$. If so, everybody with $a_i > \underline{a}$ will apply. If not, there is a further threshold defined by $E(\bar{a}_i, s_i, M) = V(\bar{a}_i)$ such that only i with $\underline{a} < a_i < \bar{a}$ apply.

3.4 Stage 0: Treatment

The intervention takes place at Stage 0 when treatment districts receive the career posters and control districts do not. We model treatment as $M_T > M_C$ representing the fact that the career emphasis suggests a higher expected net present value of lifetime material benefits despite the starting wage being the same. Our goal is to make precise the conditions under which treatment creates a trade-off by attracting applicants with lower prosociality and higher ability.

Treatment increases U^* through several channels. Of these, the increase in the net present value of exogenous material benefits M is the one that makes the job relatively more attractive to individuals with low prosociality.¹⁶ This may occur through automatic salary progression (individuals are entitled to all increases negotiated collectively for government employees) or through promotion to higher-paying positions.

To create a trade-off between ability and prosociality, the increase in M must attract higher-ability applicants. This happens through the upper threshold of ability, $\bar{a}(s_i, M)$. In particular, as Result 1 below shows, the threshold $\bar{a}(s_i, M)$ is increasing in M . This also creates an additional effect in equilibrium, as the fact that more (high-ability) candidates apply reduces the probability of being selected for all applicants. This means that $\frac{dp}{dM} < 0$. This effect is weaker for candidates with high levels of a , that is $\frac{d^2p}{dMda} > 0$, regardless of whether they apply, because they are highly likely to be selected. The fact that M affects the probability of being selected as described above creates two offsetting effects on the lower threshold. On the one hand, higher material benefits increase U^* directly and thus lower $\underline{a}(s, M)$; on the other hand, the reduction in probability arising from more applicants discourages applying and thus raises $\underline{a}(s, M)$. Combining these effects yields:

Result 1: *Increasing material benefits M will attract higher-ability applicants who would not apply otherwise ($\frac{\partial \bar{a}}{\partial M} > 0$) and either (i) lower the ability of the lowest-ranked applicant ($\frac{\partial \underline{a}}{\partial M} < 0$) and increase the total number of applicants or (ii) discourage low-ability applicants ($\frac{\partial \underline{a}}{\partial M} > 0$) and have an ambiguous impact on the total number.*

To assess the effect on prosociality we note that U^* is increasing in both a and s . The threshold $\bar{a}(s_i, M)$ is increasing in s because, due to the fact that $V_a > U_a^* > 0$, higher-ability individuals need to have a high level of prosociality to meet their participation constraint. Symmetrically, the threshold $\underline{a}(s_i, M)$ is decreasing in s because more prosocial applicants have a higher expected payoff for the same probability of being chosen, thus the probability, and hence the level of ability, that makes them indifferent between applying and not is lower. This implies that talent and prosociality

¹⁶Treatment can affect U^* in several ways that do not create trade-offs, for instance by increasing the marginal product of ability. This is due to career benefits giving high-ability individuals the chance to be promoted to higher-ranked positions where they can benefit more people or have more influence on key decisions. This appeals to high a individuals who can benefit from it and high s individuals who care about it. In this respect, treatment improves the quality of the applicant pool in all dimensions without creating a trade-off.

are positively (negatively) correlated among the highest- (lowest-) ability applicants, even though they are not correlated in the population. Therefore we have:

Result 2: *Under any M , the most able applicant is also the most prosocial. An increase in M leaves the prosociality of the marginal applicant unchanged and has an ambiguous effect on the prosociality of the average applicant.*

Figure 2 illustrates treatment effects on ability and prosociality. There are many possible ways of modelling these functions. We make the starkest possible assumption and set $H_{ae} = 0$, meaning that prosociality increases the marginal return to effort, while ability increases output directly and does not affect effort. Thus, all else equal, more prosocial people will work hard, while more able people work less hard but produce more output. We consider a simple example where $H = 2\sqrt{\xi}e + a$, where $\xi > 0$ is a parameter, and $d(e) = e^2$, which implies $e^* = \sqrt{\xi}s$, $H^* = 2\xi s + a$, and $U^* = \xi s^2 + as + M$. The probability of selection must satisfy the following: (i) increasing in a and s , (ii) decreasing in M , and (iii) $\frac{\partial^2 p}{\partial a \partial M} > 0$. For the illustration we use the following functional form $p(a, s, M) = \frac{\gamma as^\beta + \mu(M_T - M)}{1 + \mu(M_T - M)}$, where $\gamma > 0$ is set to guarantee that $p \in [0, 1]$,¹⁷ $\beta > 0$ captures the weight that the panel puts on prosociality, and $\mu > 0$ scales the decrease in probability that arises from more (high-ability) people applying in equilibrium when M rises.¹⁸ Note that $\frac{dp}{dM} = -\frac{\mu(1-p)}{1 + \mu(M_T - M)} \rightarrow 0$ as $p(a, s, M) \rightarrow 1$, which says that the increased applicant effect is zero for candidates who are certain to be selected. We set $V = va^2 + a$, where we assume $v > \gamma$ which ensures that $V_a > U_a^* > 0$ everywhere and that $V_a > E_a > 0$ for all a sufficiently large. In Figure 2a, we see that, at low levels of a , the cost of applying is too high; for sufficiently high levels of a , the increased return to ability in the private sector begins to dominate, thus defining $\bar{a}(s, M)$. Figure 2b illustrates the application frontier, that is all the combinations of a and s such that an individual is indifferent between applying and not. The frontier is positively sloped, so that the ability threshold is higher for more prosocial applicants and all individuals below the frontier will apply. Note that the main results are robust to many functional form assumptions so long as they generate $\bar{a}(s, M)$ as upward sloping. The dashed line shows that an increase in M shifts the frontier upward. The effect on average prosociality depends on the balance of two forces, which can be seen in Figure 2c. First, the average prosociality of individuals whose a is high enough that they apply only with career benefits is higher than the average prosociality of those who apply without. In the Figure, these individuals must have $s > \hat{s}$. Intuitively, candidates with high ability that are attracted by treatment must have very high prosociality in order for them to prefer the CHA role to the private sector. Second, the average prosociality of individuals whose a is low enough that they apply without career benefits is lower when these are offered. This is because U^* is increasing in M ; thus for any a the level of s that makes individuals indifferent between applying or not falls. In the Figure, we see this through the fact that the dashed line is always to the left of the thick line. This is the standard substitution or crowding out effect.

¹⁷For ease of interpretation, note that $p = 1$ when $\gamma as^\beta = 1$.

¹⁸In principle, some $\bar{M} > M_T$ would be more suitable in the function for general usage, but for our purposes M_T is sufficient and helps to avoid unnecessary proliferation of parameters.

Given that the effect on the infra-marginal and marginal applicants differ, the treatment effect on hired CHAs depends on the selection mechanism. Mechanisms that pick the highest-ability candidates from the applicant pool will produce the largest possible positive difference in ability and no difference in prosociality because the most able applicants within each pool are also the most prosocial. Mechanisms that pick randomly from the pool will still produce a positive difference in ability and a negative difference in prosociality if the average applicant under career incentives is less prosocial. Thus, to understand how differences in the applicant pool translate into differences among hired CHAs we need to understand the selection mechanism. This is the aim of the next section.

4 Treatment Effect on the Applicant Pool and Selected Candidates

4.1 Treatment Effect on the Applicant Pool

The recruitment drive yielded 2,457 applications, an average of 7.4 applicants for each position. Overall, 1,804 (73.4%) applicants met the eligibility requirements and were invited for interviews;¹⁹ of these, 1,585 (87.9%), or 4.7 per position, reported on their interview day when we administered a questionnaire to collect information on skills, career ambition, and prosociality. These 1,585 form the applicant pool we analyze in this section.

To measure treatment effects on the composition of the applicant pool we collect measures of ability and prosociality at the application stage for the universe of applicants who were interviewed. To measure cognitive skills we use grade 12 scores in the final exam and the number of courses taken in biology and other natural sciences. These are the skills measures used as application requirements. For non-cognitive skills we focus on career ambition, measured by asking applicants the job they envisage doing in five years' time, and code as career-motivated those who aim to a higher-ranked position in the Ministry. To measure prosociality we combine the applicant's self reported willingness to stay in the community in the long term together with the "Inclusion of Others in Self (IOS)" scale that measures alignment of interests (Aron and Others 2004).²⁰ We do so to identify those who want to stay in the community because they care about community outcomes, as opposed to those who stay for other reasons.

Guided by the framework, we estimate the effect of treatment on skills and prosociality, both on the average applicant and as a function of skill rank. Table 1 shows that treatment attracts

¹⁹ All completed application forms were taken to the district Ministry of Health office where district health officials checked that requirements were met. No discretion was given at this stage; applicants who did not meet the objective criteria were rejected, and those who did were invited for interviews.

²⁰ IOS measures the extent to which individuals perceive community and self-interest as overlapping. Applicants are asked to choose between four pictures, each showing two circles (labeled "self" and "community") with varying degrees of overlap, from non-overlapping to almost completely overlapping. This variable equals 1 if the respondent chooses the almost completely overlapping picture, 0 otherwise. IOS has been validated across a wide variety of contexts, and adapted versions are found to be strongly correlated with environmental behavior (Schultz 2012) and connectedness to the community (Mashek et al. 2007).

individuals with higher cognitive skills and career motivation. The average effects are about one-fifth of a standard deviation in the control group and are all precisely estimated at the 5% confidence level or above. Both here, and throughout the paper, we also report p-values based on the effective degrees of freedom (EDF) correction procedure in Young (2016) and a randomization inference procedure (Young 2019).²¹ Average prosociality is lower, albeit not precisely estimated, while age, gender, and current occupation are very similar. The latter is due to the fact that, in line with the evidence from the Census in Table A.1, there is hardly any variation. Most applicants (70%) are farmers, a further 9% are housewives, 6% are traders, and 5% are teachers. Finally, the number of applicants per health post is not significantly different. Result 1 makes precise that this can happen if treatment increases both ability thresholds. To test this, Figure 3, Panels A and B report the kernel density estimate as well as the quantile treatment effects on total test scores. Both reveal a rightward shift, namely, all applicants in treatment, from the lowest- to the highest-ranked, have higher test scores. Panels C and D present the mean levels of ability and prosociality in treatment and control across different levels of the applicant’s skill rank in his or her health post, with confidence intervals based on a regression with standard errors clustered at the district level. Rank is based on final exam scores, within community. In line with Panels A and B, Panel C shows that the treatment group’s exam scores are higher at every rank. Most importantly, in line with the theoretical framework, the difference in prosociality is zero for top-ranked applicants and negative for lower-ranked applicants.

The results in Figure 3 are in line with the simple theoretical framework: ability increases throughout whereas the effect on prosociality is zero for top-ranked candidates and negative for lower-ranked candidates. The figure makes clear that the effect of treatment on CHAs themselves will depend on how these are chosen among the applicants. We analyze this next.

4.2 The Selection Mechanism and Treatment Effect on Selected Candidates

Selection panels are in charge of choosing the two candidates that will serve as CHAs in the health post. Panels have five members: the district health official, a representative from the health post’s associated health center, and three members of the local neighborhood health committee. Each panel was asked to nominate two top candidates and up to three reserves. The Government explicitly stated a preference for women and for those who had previously worked as community health workers, but the ultimate choice was left to the panels. Overall, selection panels nominated 334 applicants as “top 2” candidates and 413 as reserves.²²

To understand how differences in the applicant pool translate into differences in hired CHAs we analyze how panels select candidates. This analysis also sheds light on whether treatment affects

²¹Specifically, we use the randomized randomization-t p-value, computed using *randcmd* in STATA with 10,000 iterations.

²²The nominations were reviewed centrally by the Government of Zambia, and 334 final candidates were invited to join a yearlong training course. Of these, 314 applicants accepted the invitation and, in June 2011, moved to the training school in Ndola, Zambia’s second-largest city. Of the applicants who joined the program, 307 graduated and started working in August 2012. All the health workers were deployed back to their communities of origin.

panels' choices and on which traits panel members deem important for the job. Table 2 estimates the probability that candidate i in health post h is chosen as follows:

$$s_{ih} = \sum_{j \in J} \alpha_j^T C_h X_i^j + \sum_{j \in J} \alpha_j^C (1 - C_h) X_i^j + \gamma N_h + \zeta_{ih}$$

where $s_{ih} = 1$ if i is one of the two nominated candidates and 0 otherwise. C_h equals 1 if health post h is in treatment and 0 if it is in the control group. X_i^j are indicator variables that equal 1 if candidate i is in the top three of trait j , and the core set J includes skills, ambition, and prosociality. We also report regressions with an expanded set that includes social connections to local political leaders to test whether connections help with getting the job when material benefits are higher. To control for the strength of competition, we include the number of interviewed candidates in the same health post N_h . We control for the stratification variables and cluster standard errors at the district level, correcting for effective degrees of freedom using the procedure in Young (2016).

The coefficients of interest are α_j^T and α_j^C , which measure the weight given to trait j in the treatment and control groups, respectively. We test the null that panels use the same criteria in both groups, that is $\alpha_j^T = \alpha_j^C$. Panels are exposed to treatment as they see the posters, but in contrast to candidates, for whom the poster is the only source of information, panel members know the job attributes and who would be suitable for it. The two more senior panel members—the district health official and the health center representative—are employees of the Ministry of Health, and hence are familiar with career progression rules regardless of treatment. Thus this is likely not as powerful, or perhaps entirely moot.²³ Table 2 reports the estimates of α_j^T and α_j^C for all $j \in J$ and the p-value of the test of equality.

Column 1 in Table 2 shows that panels put a strong positive weight on skills and prosociality and do so equally in both treatment and control groups. The average probability of being nominated for an applicant who does not rank at the top of the skills and prosociality distributions and who has no career ambition is 0.09. This increases by 12pp for applicants at the top of the skill distribution, by 7 to 10pp for applicants with career ambitions, and by 6 to 9pp for applicants with high prosociality. The tests of equality of coefficients between treatment and control do not reject the null for any of these traits. Column 3 additionally shows that connections either to political leaders or to staff at the health facility do not affect the probability of selection in either treatment or control.

Taken together, the evidence suggests that career opportunities attract applicants who have different skills, career motivation, and prosociality and that all panels deem these traits to be valuable and are more likely to choose applicants who rank highly in all three.

Our conceptual framework makes clear that, compared to a random selection mechanism, this type of selection leads to higher skill differences and eliminates prosociality differences. To illustrate, we compare the traits of the CHAs selected by panels to 1000 random draws of two CHAs from each health post's applicant pool. Table 3 reports the average trait for panel-selected CHAs and the 10th, 50th, and 90th percentiles of the same traits when candidates are chosen randomly. The

²³Further analysis, available upon request, shows that treatment does not affect panel composition.

table shows that the distribution of skills and motivation in treatment is to the right of control, and that panels choose from the top in both groups. This implies that, compared to random selection, panels select CHAs who have higher exam scores and career motivation. Indeed, panel-selected CHAs score higher than the 90th percentile of randomly selected CHAs on all three measures. Despite the fact that panels put the same weight on talent and career motivation in treatment and control, the average skill and career motivation of selected CHAs is higher in the treatment group because treatment attracts candidates who do not apply in control.

In contrast, whereas randomly selected CHAs have lower prosociality in the treatment group, panel selection undoes this difference because, as shown in Figure 3, the most talented applicants in each pool have the same level of prosociality, and panels select these. The fact that treatment creates a trade-off between ability and prosociality for low-ability applicants is of no consequence because these are not hired.

5 Inputs in Service Delivery

5.1 Measuring Inputs in Service Delivery

The health workers’ main task, to which they are required to devote 80% of their time, or 4 out of 5 days per week, is to visit households. The input part of our analysis focuses on the number of visits completed over the course of 18 months, from August 2012 (when the health workers started work) until January 2014. The number of household visits is akin to an attendance measure for teachers or nurses: the health workers are supposed to work in people’s houses, and we measure how often they are there. Naturally, differences in the number of visits can be compensated for with differences in other inputs; we discuss this possibility in Section 5.3 after establishing the main results. Furthermore, differences in inputs ultimately are of interest only if they lead to better outcomes, which we will discuss in Section 6.

Our primary measure of household visits is built by aggregating information on each visit from individual receipts. All the health workers are required to carry receipt books and issue each household a receipt for each visit, which the households are asked to sign. The health workers are required to keep the book with the copies of the receipts to send to the Government when completed. They are also required to send all information on these receipts—consisting of the date, time, and duration of the visit, as well as the client’s phone number—via text message to the Ministry of Health. These text messages are collected in a central data-processing facility, which we manage.

Since visits are measured by aggregating text messages sent by the health workers themselves, identification can be compromised by the presence of measurement error that is correlated with treatment. For instance, health workers in the career treatment might put more effort in reporting visits via text messages or might report visits that never took place, leading to a positive bias in the estimated treatment effect.

We validate our visits measure by comparing it to administrative data and households’ own reports of health worker activity. The administrative data is drawn from the Health Management and Information System (HMIS), which is the Ministry of Health’s system for collecting routine health services data at government facilities. These are reported at the end of each month and sent electronically to the Ministry via a mobile platform, jointly by the two CHAs and the other staff working in each health post. As HMIS data are only available aggregated at the health post level (summed over the two workers in each health post) we regress these on our visit measure, also aggregated at the health post level. Columns 1 and 2 in Table A.3 show that the two measures are strongly correlated ($r=0.767$). The correlation in this measure is higher in the treatment group than the control group (0.836 vs. 0.644), although the difference is not significant, which contradicts the differential reporting hypothesis.

The households’ reports are collected via a survey that we administered to 16 randomly chosen households in each of 47 randomly selected communities chosen from the set of communities where the health workers operate, stratified by district. We ask respondents whether they know each of the health workers (97% do), whether they have ever been visited (44% of them have), and their level of satisfaction with each health worker. Columns 3-4 show a precisely estimated correlation between our visit measure and the probability that a household reports a visit. This correlation is slightly higher in the treatment group than the control group, although the difference is not significant. Columns 5-6 show a precisely estimated correlation between our visit measure and the household’s reported level of satisfaction with the health worker’s performance. The difference between treatment and control group is small and insignificant, casting doubt on the relevance of differential reporting.

Taken together, the findings in Table A.3 generally validate our visits measure. Ultimately, however, we will not be able to detect a treatment effect on households’ health outcomes in Section 6 if measured differences in visits capture differences in reporting rather than in actual visits.

5.2 Treatment Effect on Household Visits

Table 4 reports the reduced form effects of treatment on performance—that is, the estimates of:

$$v_{ihdp} = \alpha + \beta C_{id} + Z_h \gamma + \delta E_d + \rho_p + \epsilon_{ihdp} \quad (5.1)$$

where v_{ihdp} is the number of visits completed by health worker i in catchment area h , district d , and province p . C_{id} equals 1 if agent i is recruited and operates in a district assigned to the career opportunities treatment. Z_h is a vector of area characteristics, which includes the number of staff at the health post, cell network coverage, and the distribution of households between farms and villages described in Table A.1, although we note that the results are qualitatively unchanged if we remove these. We control for the stratification variables, district-level high school graduation rate E_d , and province indicators ρ_p throughout. Standard errors are clustered at the level of randomization, the

district, and, as mentioned before, we also report p-values from the effective degrees of freedom correction in Young (2016) and a randomization inference procedure (Young 2019).

The coefficient of interest is β , which measures the effect of making career opportunities salient at the recruitment stage on the number of visits completed over 18 months. Considering that all the health workers are given the same information on career opportunities during the year-long training, β captures the effect of career opportunities on performance through selection. Recall that selection can affect performance by increasing productivity for a given level of effort or by increasing the marginal return to effort. An example of the former is talent for logistics: for the same amount of effort, a more talented health worker plans better and reaches more households in the same amount of time. An example of the latter is the utility weight put on career advancement: health workers who value career more draw a higher marginal benefit from a given unit of effort and therefore exert more effort.

The causal effect of career opportunities on performance can be identified under the assumptions that (i) C_{id} is orthogonal to ϵ_{ihdp} , and (ii) there are no spillovers between the two groups. Orthogonality is obtained via random assignment. Spillovers via movements of health workers between treatment and control areas are ruled out by the program requirement that health workers must have been residing in the community they want to work in prior to applying. This implies that career opportunities cannot draw in talent from control areas. Spillovers of information, caused for example by potential applicants in control seeing the treatment poster, would introduce a downward bias because they would reduce the information differences between treatment and control. Information spillovers are minimized by design, as recruitment messages were randomized at the district level—which, given the travel distance between rural communities in different districts, makes it very unlikely that applicants in one group might have seen the poster assigned to the other group. Importantly, information cannot accidentally spillover through the district officials that implement the program or through the recruitment panels, as these are only exposed to one treatment.

Column 1 of Table 4 reveals a large and precisely estimated effect of career opportunities on household visits: health workers recruited by making career opportunities salient do 94 more visits (29% more than control) over the course of 18 months. The median treatment effect is 104.4 (bootstrapped s.e. 43.9), which allays the concern that the average effect is driven by outliers. The magnitude of the difference is economically meaningful: if each of the 147 health workers in control had done as many visits as their counterparts in the career treatment, 13,818 more households would have been visited over the 18-month period. Given that for most of these households, health workers are the only providers of health services, the difference between treatments is likely to have implications for health outputs in these communities. We return to this issue in Section 6.

Columns 2-4 divide the 18-month period into three and show that the estimated treatment effect is identical in the three semesters. This casts doubt on the alternative hypothesis that agents in the two groups have the same traits, but agents in the treatment group perceive stronger career incentives because they have known about them for longer (about 2 years vs. 1 year for the control

group). Such a difference should wane with time, while the difference due to stable traits should be stable.²⁴

To shed light on what treatment health workers do differently, we administer a time use survey to all health workers after they have started working. The findings, reported in detail in the Appendix, indicate that treatment and control health workers work similar hours and allocate their time similarly across similar activities. This indicates that treatment health workers are more efficient at their jobs. Household visits take place in remote, low-density areas: the median 78 square km area has 200 households, with an interquartile range of 130 to 360. It is thus rather time-consuming to go from house to house, and this is compounded by the fact that roads are bad. In this setting, the ability to plan—e.g., by making appointments with specific households or collecting information as to whether members are likely to be home before setting out to visit them—is an important determinant of completing visits successfully. Our model was agnostic as to whether effort and ability are complementary, i.e., as to whether higher ability, for a given level of prosociality, increases output directly only or also through increased effort. Given that these two groups have similar levels of prosociality but different levels of ability, these findings, suggestive of similar effort levels, appear more consistent with the version in which they are not complementary.

To conclude we establish the extent to which differences in performance are due to selection on observables. We search for the vector of observables that explains the largest possible share of variation of performance in the control group and use the estimated coefficients to predict performance in the treatment group.²⁵ This yields the predicted difference between treatment and control on the basis of the observables that best predict performance. The best predictors explain 31% of the observed variation in control and the predicted difference between treatment and control is 44 visits. Given that the actual, unconditional performance gap is 101, differences in observables explain 44% of it. The remaining 56% is due to traits we do not measure.

The finding that observables have limited power in explaining performance differences echoes the well-established finding that differences in teachers’ effectiveness are large and only weakly correlated with observable traits. It is also consistent with other settings where agents self-select, such as in applying for welfare programs (Alatas et al. n.d.) or purchasing health products (Ashraf et al. 2010). In those settings, like in ours, self-selection cannot be mimicked by targeting on observable traits.

²⁴The fact that the treatment effect is stable also casts doubt on the event that it is driven by a negative “surprise” for agents in the control group (i.e., their effort response to finding out about career opportunities is negative and larger—in absolute value—than what it would have been had they known the career opportunities at the outset). In addition, if there really were a substantial intrinsic crowd-out of this kind, we would likely expect at least some greater attrition in the control group as participation constraints that were met *ex ante* would not be met *ex post* - yet, as discussed later, we find that attrition is neither economically nor significantly different between treatment and control. Nonetheless, it remains possible that intrinsic crowd-out may be contributing to the results if the effects do not fade with time and do not lead to dropping out.

²⁵Specifically, we select the five best predictors, in addition to our stratification variables, using the Furnival-Wilson branch and bound algorithm (implemented using *vselect* in STATA) in the control group. We then use these to predict the mean number of visits in both treatment and control, and calculate the difference in these predicted means.

5.3 Beyond Number of Visits: Compensation Mechanisms and Other Activities

Table 5 investigates the hypothesis that health workers in the control group take other actions that compensate for the lower number of visits. Column 1 tests whether control health workers are more likely to be retained while career health workers leave with their newly acquired skills as soon as it is feasible to do so. Since the health workers are bonded to their position for one year,²⁶ we measure retention by the number of health workers who make at least one visit after the one-year commitment has elapsed. We find that, by this measure, 18% of health workers drop out, though some of this may be due to a combination of malfunctioning phones and the rainy season (falling between months 15-18 in our analysis window), making travel to cell network-accessible areas difficult. This attrition rate is balanced across treatments. It is important to note that according to the Ministry’s rule, health workers have to wait two years before applying for higher-ranked positions, such that none of those who left their positions did so for career progression. It is possible that career opportunities will affect retention rates after the two-year mark. Whether this entails a welfare cost depends on whether the workers can be easily replaced and whether the Government can use their skills in other jobs. In our context, replacement is straightforward; the number of applicants per post was above seven, and the government faces scarcity of health staff at all levels, such that promoting high-performing health workers to nursing and other higher-level cadres is likely to be welfare-improving.

The number of visits can hide heterogeneity on a variety of dimensions that can make the health workers less effective in generating health outcomes, such as doing shorter visits, targeting the head of household rather than women and children, or targeting easier-to-reach households. We provide evidence that career health workers do not do worse on any of these dimensions. They devote the same time to a single visit (column 2), and are equally likely to target their primary clients—women and children (column 3). They also reach more households (column 4) and make more follow-up visits (column 5). The point estimates indicate that just over one-third (36/94) of the total treatment effect is due to career health workers visiting more households, and two-thirds to them visiting the same household more than once. This is consistent with the two groups of health workers having a similar number of households in their catchment area and visiting them at least once, but treatment health workers doing more follow-up visits. Note that follow-ups are considered an integral part of the health worker job, in view of which Ministry of Health guidelines state health workers should attempt to visit each household on a quarterly basis. Finally, Table A.5 shows that treatment health workers allocate their time in a similar way to control health workers during household visits. This allays the concern that health workers who see themselves as health professionals neglect “soft” tasks like counseling.

Besides household visits, the health workers are expected to assist staff at the health post by seeing patients, assisting with antenatal care, and maintaining the facility. They are also supposed

²⁶The health workers were told that if they quit before one year of service, they would be required to pay monthly wages for any months not worked (rather than simply relinquishing pay) to compensate the Government for the free one-year training that they received.

to organize community meetings such as health education talks at the health post and in schools. Columns 6-7 investigate whether differences in household visits are compensated by differences in secondary tasks using HMIS data on the number of community meetings health workers organize and the number of patients they attend to at the health post. The latter should be seen as a proxy of the quantity of services delivered by the health workers at the health post, as seeing patients is mostly a nurse’s job. We find that health workers recruited by making career opportunities salient organize twice as many meetings over 18 months, and the difference is precisely estimated. The effect of career opportunities on the number of patients the health workers see at the health post is also positive, but small and not precisely estimated.

6 Facility Utilization, Health Practices, and Health Outcomes

The program leads to a substantial increase in the number of health staff operating in the communities where the health workers are deployed: the number of staff associated with the community health post increases on average from 1.5 to 3.5. Given the size of the increase and the magnitude of the treatment effect on household visits and community mobilization meetings, it is reasonable to expect treatment to affect health outcomes in these communities. The health workers can directly affect facility utilization and health practices by increasing both demand, e.g., by providing information and promoting behavioral changes, and supply, e.g., by helping cover staff shortages at the health post or delivering medical treatments to households. In turn, improved facility utilization and practices should lead to better outcomes.

Besides their intrinsic importance for the welfare of these communities, treatment effects on facility utilization and household outcomes allow us to shed light on whether health workers in the control group perform better on dimensions we cannot observe enough to improve outcomes. For instance, treatment health workers could target households that are more interested in health services and would use facilities when necessary anyway, while control health workers could target households that they need to persuade to change behavior, and that require more work, leading to fewer visits overall. If this were true, treatment would be uncorrelated (or even negatively correlated) with facility utilization and health outcomes.

To provide evidence on whether treatment affected facility utilization, we use data from the Ministry’s HMIS administrative records; to measure effects on health practices and outcomes, we survey households residing in the communities where the health workers operate. As the main remit of the health worker job is maternal and child health, we focus on this throughout.

6.1 Treatment Effect on Facility Utilization

The Ministry’s HMIS administrative records are compiled by facilities’ senior staff and transmitted to the Ministry of Health via an electronic platform. Two levels of facilities serve these communities:

health centers and health posts.²⁷ The health workers are supposed to encourage women to give birth at the closest health center and to bring in children for regular visits and immunizations at the closest facility (health center or health post). The importance of institutional deliveries in this context cannot be understated: Zambia’s maternal mortality rates are very high and health centers have the equipment and medical supplies that can prevent these deaths. Regular children’s visits ensure that conditions such as diarrhea are treated before they become dangerous. Immunizations protect children from potentially fatal illnesses.

To test whether the treatment affected facility utilization, we obtain information on institutional deliveries, children’s visits, and immunizations for the period January 2011-June 2014 and estimate the following specification:

$$y_{hdpt} = \alpha + \beta C_{hd} + \gamma A_t + \delta C_{hd} * A_t + Z_h \theta + E_d \phi + \rho_p + \xi_{hdpt}$$

where y_{hdpt} is the outcome in health facility h in district d and province p at quarter t .²⁸ h represents the lowest level of government facility to which the health workers can refer their patients. This is the health post if operational or the closest health center otherwise. The only exception is childbirths, which are always measured at the health center level, as that is where they are supposed to take place. $C_{hd}=1$ if facility h is located in a district randomly assigned to the career treatment. We have data for 14 quarters, equally divided before and after the health workers’ arrival, and $A_t=1$ after the health workers’ arrival (4th quarter of 2012). To minimize composition bias and to test for robustness to facility fixed effect models, we restrict the sample to the facilities for which we have at least three observations before and after the health workers’ arrival.²⁹ Z_h is a vector of area characteristics, which includes the number of staff at the health post, cell network coverage, and the distribution of households between farms and villages described in Table A.1. We control for the stratification variables, district-level high school graduation rate E_d , and province indicators ρ_p throughout. Standard errors are clustered at the level of randomization, the district, and, as mentioned before, we also report p-values from the effective degrees of freedom correction in Young (2016) and a randomization inference procedure (Young 2019).

The parameter of interest is δ , the difference-in-differences between facilities in treatment and control districts before and after the health workers’ arrival. Under the parallel trend assumption, δ captures the effect of career opportunities for health workers on these outputs.

Table 6 shows that, indeed, career opportunities improve clinic utilization outputs. In particular, the number of women giving birth at a health center increases by 30% relative to the mean in

²⁷Health facilities in Zambia are structured according to a population-based hierarchy. Health posts are the first-level health facility for most rural communities and provide basic medical care (no inpatient or surgical services). Health centers, which typically serve a population encompassing four to five health posts, provide both outpatient and inpatient services, including labor and delivery and minor surgical procedures. District hospitals in turn encompass several health center catchment areas and are primarily focused on inpatient care.

²⁸HMIS data should be transmitted to MoH monthly, but in practice (due to poor connectivity), reports are missing for some months and the information added to the following month. We aggregate the data at the quarterly level to smooth out monthly fluctuations due to this.

²⁹This restriction keeps an average of 76% of the health posts and 71% of the health centers in the sample.

control areas at baseline. The effect on institutional deliveries is thus the same order of magnitude as the effect of performance pay for clinics as evaluated in Rwanda (23%; Basinga et al. 2011) and Cambodia (25%; de Poel et al. 2014). Selection and incentive effects of similar magnitudes (22% each) are also found in the only firm study that identifies the two separately (Lazear 2000).

Table 6 also shows that the number of children under age five visited increases by 24%, the number of children under 5 weighed increases by 22%, and the number of children under 12 months of age receiving polio vaccination increases by 20%. The effects on postnatal visits for women and BCG and measles vaccinations are also positive and in the 8-22% magnitude range, but are not precisely estimated. The average standardized treatment effect (Kling et al. 2007) over all outcomes is 0.278, significantly different from zero at the 1% level. Reassuringly, there are no significant differences between treatment and control in any of these outcomes before the health workers' arrival: all the estimated β coefficients are small and not significantly different from zero.

To provide support to our identifying assumption, in Table A.6 (Panel A) we run a placebo test where we split the pre-health worker period in two halves and test whether outcomes improve in treatment areas over time even in the absence of the health workers. Reassuringly, they do not. Finally, Table A.6 (Panel B) estimates (2) with facility fixed effects; the fact that all estimated δ coefficients remain stable provides evidence that they are not biased by time-invariant facility unobservables correlated with treatment.

6.2 Treatment Effect on Health Practices and Outcomes

To provide evidence on the effect of treatment on health practices and outcomes, we survey households in 47 randomly chosen communities located in each of the 47 districts where the health workers operate.³⁰ We randomly choose 16 households in each community, surveying 738 in total.³¹ These surveys are administered by a team of enumerators who are trained by us and unconnected to the health workers or the Ministry of Health. As the main focus of the health worker job is maternal and child health, we only survey households that contain at least one child under five. The survey contains modules on health and sanitation knowledge, health practices, incidence of illnesses, and anthropometrics for the youngest child. Knowledge, practices, and illnesses are self-reported; deworming and immunization data are drawn from the child health card, and anthropometrics are measured by trained enumerators. We interview the main carer of the child, which is their mother in 90% of the cases and either a grandparent or a sibling in the remaining 10%. All questions are drawn from the Demographic and Health Survey (DHS) Zambia questionnaire, with the exception

³⁰Although 48 districts were randomized, one district did not participate in the recruitment process or submit nominations for CHA candidates, and hence was excluded from the survey sample.

³¹The sample frame had 752 households but we interviewed 738. The missing households are evenly spread across communities as the number of households surveyed in a community varies between 13 and 16. The difference is due to several factors. In some communities, safety concerns related to local political tensions forced the survey team to leave the community before completing surveying. In other communities, especially low-density communities where travel times between households could exceed one hour, the survey team was unable to find a sufficient number of eligible households within the allotted survey time. One household interview was lost due to malfunction of the mobile device on which the interview was recorded.

of the health knowledge module, which we designed based on the health worker curriculum, and mid-upper arm circumference (MUAC), which the DHS does not measure.

Table 7 reports the estimates of:

$$y_{idp} = \alpha + \beta C_{id} + D_i \gamma + \delta E_d + \rho_p + \epsilon_{idp} \quad (6.1)$$

where y_{idp} is the outcome of child (or respondent) i in district d and province p . C_{id} equals 1 if child (or respondent) i lives in a district that is assigned to the career opportunities treatment. D_i is a vector of child, respondent, and household characteristics that includes child age and gender, household size and number of assets, and the education level of the respondent. As above, we control for the stratification variables, district-level high school graduation rate E_d and province indicators ρ_p throughout, and cluster standard errors at the district level, with additional p-values reported based on the effective degrees of freedom correction procedure in Young (2016) and a randomization inference procedure (Young 2019).

Column 1 shows that the average respondent answers 74% of the knowledge questions correctly and that this does not differ by treatment status. In contrast, treatment affects all the health practices we collect information on. In particular, Columns 2 and 3 show that children under 2 living in treatment areas are 5pp more likely to be breastfed,³² and their stools are 12pp more likely to be safely disposed; these effects represent an 8% and 20% increase from the control group mean, respectively. Columns 4 and 5 show that treatment also increases the number of deworming treatments by 16% and the likelihood that the child is on track with the immunization schedule by 4.7pp, which is 81% of the control group mean (5.8%).³³ Importantly, the treatment affects the incidence of immunizations for children who are young enough to have been exposed to the health workers when their immunization period started (as shown in Column 5), but not for those who were too old to start the cycle when the health workers started working (coefficient -0.014, standard error 0.022). This echoes the findings in Table 6 that show no difference in immunization rates between treatment and control areas before the health workers started working.

Columns 6-8 measure treatment effects on the incidence of three main illness symptoms: fever, diarrhea, and cough. These are fairly common, as 47%, 26%, and 45% of children in control areas had experienced them, respectively, in the past two weeks. As it is widely acknowledged, self-reported symptoms can actually worsen as knowledge improves and individuals learn how to recognize them, so these effects are lower bounds. We find that treatment reduces the incidence of cough symptoms by 7pp while leaving the others unchanged. Finally, Columns 9-13 show treatment effects on anthropometric measurements. We report weight-for-age z-scores and mid-upper arm circumference (MUAC). The combination of these two allows us to measure both chronic and acute

³²WHO recommends breastfeeding until the age of two years.

³³A child is defined to be on track if she has completed all immunizations required for her age. At age 3 months, this includes BCG, OPV 0-2, PCV 1-2, DPT-HepB-Hib 1-2, and rotavirus 1-2. At 4 months, this includes, additionally, OPV 3, PCV 3, and DPT-HepB-Hib 3. At 9 months, this includes OPV 4 if OPV 0 was not given, and measles 1. The immunization series is complete at age 18 months with measles 2. Finally, we consider a child to be on track for vitamin A supplementation if she has ever been supplemented.

malnutrition.³⁴ Following WHO’s guidelines, we use the -2SD and -3SD thresholds for weight-for-age z-scores to measure moderate and severe underweight, respectively, and 12.5 cm and 11.5 cm for MUAC to measure moderate and severe wasting, respectively (Food and Nutrition Technical Assistance Project 2011). According to these measures, 21% of the children in control areas are underweight, and 5% severely so. The incidence of wasting is much lower, with 3.6% of the children exhibiting some wasting and 1.5% severe wasting. These data, which match the corresponding DHS figures for rural Zambia (Government of Zambia 2014), suggest that these areas are characterized by high rates of chronic malnutrition but low rates of acute malnutrition.

The findings in columns 9-10 show that children in treatment areas are 5pp less likely to be underweight (25% of the control group mean) and 3pp less likely to be severely underweight (55% of the control group mean). In line with this, columns 11 and 12 show a large percentage reduction in wasting, but given the limited occurrence of this in our sample, the effects are not precisely estimated. The average standardized effect of the two measures is precisely estimated with p-value 0.019 for the less severe measures and 0.045 for the more severe, with the latter shown in the Table in column 13.

The average standardized treatment effect across all variables (coded so that higher values correspond to better outcomes) is 0.108, significantly different from zero at the 1% level.

Taken together, the findings in this and the previous section show that differences in the inputs provided by treatment and control health workers are matched by differences in facility utilization and household health practices. The selection effect of career opportunities is strong enough to generate discernible differences in household behaviors and child health outcomes.

7 Conclusion

Attracting effective employees is a core objective for all organizations. This can be a particularly challenging objective to achieve for public organizations because both effective performance (in, for example, generating health impact) and desirable employee attributes are difficult to measure. But the stakes to getting this right are high. Our paper has shown that offering a civil service position with career opportunities for community-based work attracts agents who deliver health services with substantial impact. This significant effect on the health and well-being of communities is driven entirely by a selection effect of the types of agents drawn into the position.

The civil service job we study is one sometimes referred to as a “street-level bureaucrat” (Lipsky 1980), a job where internalizing the utility of beneficiaries could be particularly helpful. Yet it was in just such a job that offering a career in the civil service, in posters that clearly attracted ambitious types, provided large impacts. Of course, the career opportunities which attracted ambitious

³⁴We did not measure weight-for-height, an alternative to MUAC for assessing acute malnutrition, for three reasons. First, compared to weight and MUAC, height measurement is more invasive, requiring, for children under two, laying the child down on a height board and having two enumerators hold the child while collecting the measurement. During survey piloting, many respondents (and the children themselves) balked at this procedure. Second, accurate height measurement is made difficult by high measurement error relative to standard effect sizes (Mwangome et al. 2012). Finally, MUAC is a more accurate predictor of mortality (Myatt et al. 2006).

types—a career in the Ministry of Health—entail some social benefit, and the community-oriented nature of the job attracted a basic level of altruism across the board. But it is in precisely these types of jobs where it has been argued that adding individualistic benefits, such as career opportunities, might attract the “wrong” type of individual. Our experiment reveals that this is indeed the case, as the lower-ability applicants in the treatment group have lower prosociality. Thus if candidates were picked by a random draw we would expect fewer prosocial recruited candidates in treatment.

In practice, however, selection mechanisms, in Zambia and elsewhere, do not choose applicants randomly. To the extent that the mechanism picks from the top of the ability distribution, the sorting equilibrium guarantees that these are the most prosocial. This allays the concern, often expressed by policy makers, that offering material rewards will crowd out prosocial applicants in education and health (World Health Organization 2006; Lehmann and Sanders 2007; Muralidharan and Sundararaman 2013). The findings also stress the importance of giving the right incentives to selectors. The two components of recruitment—sorting and selection—are equally important because good candidates cannot be hired if they do not apply and improving the applicant pool is useless unless the best candidates are selected.

Tailoring job design to attract star performers has three advantages over the more common alternative of setting higher requirements on observable characteristics. First, these characteristics are difficult to identify; in our case observables only explain 44% of the performance gap. Second, requirements can only screen out those who do not meet them; they do not necessarily draw in those who do. Most importantly, requirements create barriers to entry and rents, which can draw in applicants attracted by rent extraction rather than public service delivery.³⁵

The findings measure the productivity gains that come from effective selection via recruitment: treatment health workers provide more inputs at the same cost, since wages are the same across both treatments.³⁶ The fact that the health workers are recruited locally from the communities where they are meant to serve implies that there is no competition for talent across communities: career opportunities can thus be offered in each community without losing effectiveness, as each community can only hire from their own pool, and most communities in these areas have access to a pool of skilled individuals who are either unemployed or in low-skill jobs.

³⁵In this case, pay and qualifications might end up being negatively correlated with equilibrium performance as illustrated by the evidence on effort and performance among civil-service versus contract teachers, as in, for example, Muralidharan and Sundararaman (2013), Duflo et al. (2015), and Kremer et al. (2013). Muralidharan and Sundararaman (2013) find that, in rural Andhra Pradesh, contract teachers (who have less education and qualifications) are at least as effective as regular civil-service teachers, suggesting that the substantial wage differential (of over a factor of 5) “is unlikely to reflect differences in productivity and mostly represents rents accruing to unionized civil-service teachers” (p. 3). We owe this suggestion to a constructive referee.


³⁶Due to political constraints, all agents had to be paid the same amount. This implies that we cannot judge whether agents attracted by career opportunities have a higher reservation wage, such that their higher performance comes at a price; in other words, the government could get the agents in the control group to work for a lower wage. A priori, the difference in reservation wages between applicants in the two treatments is difficult to sign: that applicants to the career opportunities treatment are more skilled suggests that it might be positive, whereas the fact that they expect to move on to better-paid positions suggests that it might be negative (in the manner that interns are typically willing to forego compensation for the sake of career opportunities).

While retention rates after 18 months are the same in the two groups, agents in the career incentives treatment might leave their posts for higher-ranked positions sooner than those in the control group. Whether this entails a welfare cost depends on whether they can be easily replaced and whether the government can use their skills in other jobs. In our context, replacement is straightforward; the number of applicants per post was above seven, and the government faces scarcity of health staff at all levels, such that promoting strong performers to nursing and other higher-level cadres is likely to be welfare-improving. In contexts where retention in the original post is more important, the welfare cost of attracting agents who expect to move on will be higher.

More generally, we cannot quantify the opportunity cost of the health workers' time, namely the value of the activities they give up to become full-time health workers, and the size of this difference between treatment and control. If productivity in these alternative occupations is increasing in the same qualities that make a health worker productive, the findings imply that the opportunity cost is higher in the treatment group. By revealed preference, we know that the private value of the health worker job must be at least equal to the private value of these activities. Otherwise these individuals would have not switched occupations. To the extent that health work generates more social value than the outside option, the private and social optima coincide, and a career-oriented position for community-based public service delivery improves social welfare.

Figure 1: Treatment and control posters

**REPUBLIC OF ZAMBIA
MINISTRY OF HEALTH**



DESIGNATED HEALTH CENTRE: _____	FOR POSTING AT: _____
---------------------------------	-----------------------

TRAINING OPPORTUNITY

ONE-YEAR COURSE IN COMMUNITY HEALTH

The Ministry of Health of the Republic of Zambia is launching a new national Community Health Worker (CHW) strategy and invites applicants to participate in the inaugural training of community health workers.

The training will begin on **30th August 2010** and will be held at the Provincial level for selected applicants. All participation costs, including transportation, meals and accommodation will be covered by the Ministry of Health.

BENEFITS:

- Become a highly trained member of Zambia's health care system
- Interact with experts in medical fields
- Access future career opportunities including:
 - Clinical Officer
 - Nurse
 - Environmental Health Technologist

QUALIFICATIONS:

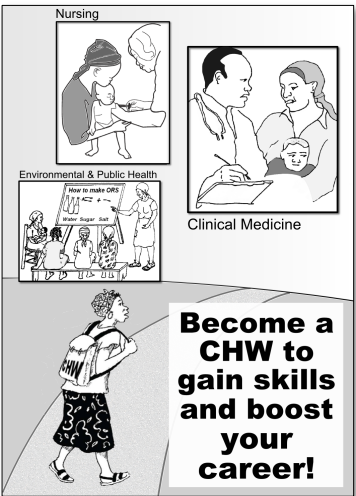
- Zambian National
- Grade 12 completed with two "O" levels
- Age 18-45 years
- Endorsed by Neighborhood Health Committee within place of residence
- Preference will be given to women and those with previous experience as a CHW

APPLICATION METHOD:

Submit to the DESIGNATED HEALTH CENTRE indicated above:

- Completed application form with necessary endorsements. If no blank forms are attached to this notice, kindly obtain a blank one at the nearest health centre.
- Photocopy of school certificate documenting completion of Grade 12 and two "O" levels.
- Photocopy of Zambian national registration card.


For more information: Contact the designated health centre indicated above.



CLOSING DATE: 30th JULY 2010.

Only shortlisted candidates will be contacted for interview.

**REPUBLIC OF ZAMBIA
MINISTRY OF HEALTH**



DESIGNATED HEALTH CENTRE: _____	FOR POSTING AT: _____
---------------------------------	-----------------------

TRAINING OPPORTUNITY

ONE-YEAR COURSE IN COMMUNITY HEALTH

The Ministry of Health of the Republic of Zambia is launching a new national Community Health Worker (CHW) strategy and invites applicants to participate in the inaugural training of community health workers.

The training will begin on **30th August 2010** and will be held at the Provincial level for selected applicants. All participation costs, including transportation, meals and accommodation will be covered by the Ministry of Health.

BENEFITS:

- Learn about the most important health issues in your community
- Gain the skills you need to prevent illness and promote health for your family and neighbors
- Work closely with your local health post and health centre
- Be a respected leader in your community

QUALIFICATIONS:


- Zambian National
- Grade 12 completed with two "O" levels
- Age 18-45 years
- Endorsed by Neighborhood Health Committee within place of residence
- Preference will be given to women and those with previous experience as a CHW

APPLICATION METHOD:

Submit to the DESIGNATED HEALTH CENTRE indicated above:

- Completed application form with necessary endorsements. If no blank forms are attached to this notice, kindly obtain a blank one at the nearest health centre.
- Photocopy of school certificate documenting completion of Grade 12 and two "O" levels.
- Photocopy of Zambian national registration card.

For more information: Contact the designated health centre indicated above.



CLOSING DATE: 30th JULY 2010.

Only shortlisted candidates will be contacted for interview.

Figure 2: Treatment effects on the applicant pool, in theory

Figure 2a




Figure 2b



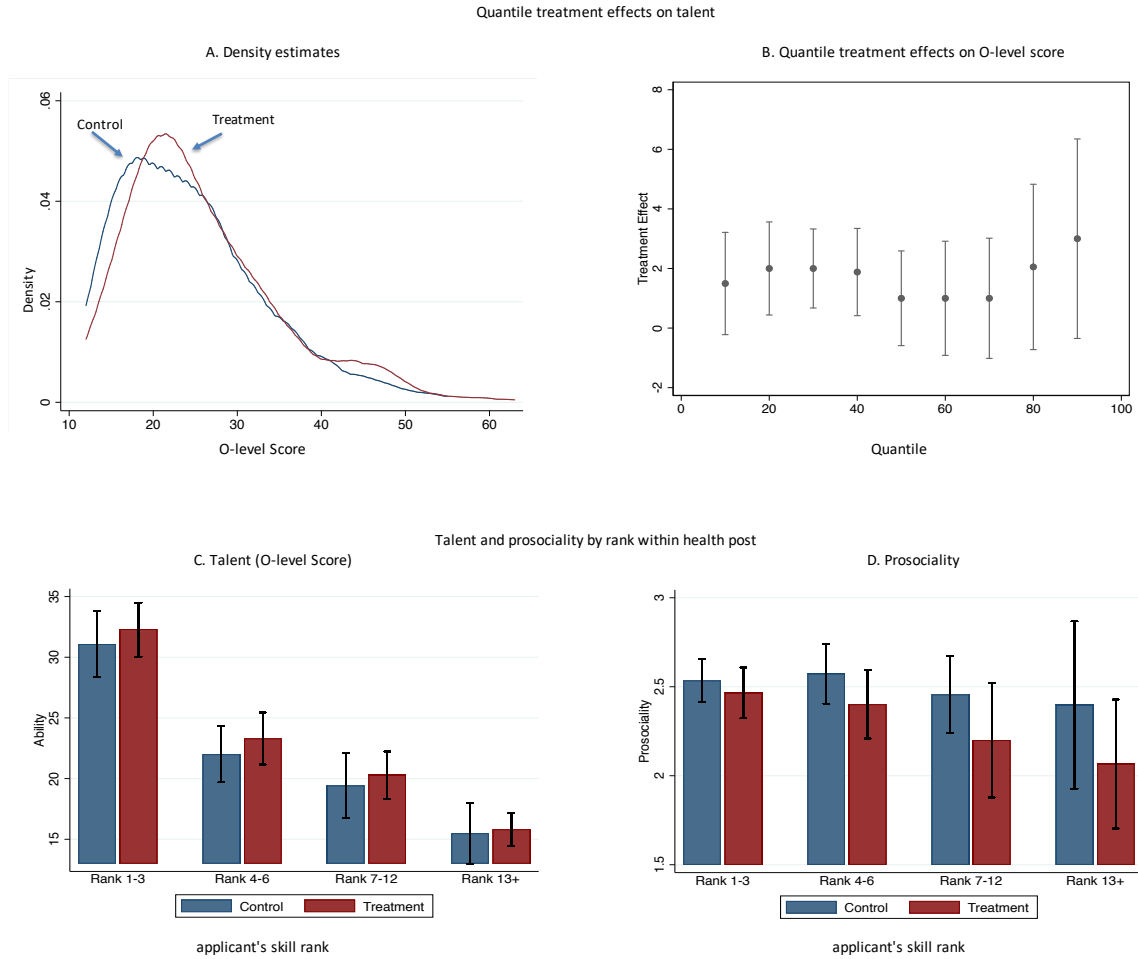


Figure 2c



The first figure presents the expected utility to applying against the payoff of the outside option. The intersections represent the thresholds of ability at which agents are indifferent between applying and not. The upper threshold is increasing in s while the lower threshold is decreasing in s . The second figure presents the upper threshold of ability, at which agents are indifferent between applying and not, as a function of s . Higher monetary rewards shift the frontier, with a larger positive effect at higher levels of s as the negative equilibrium probability effect decreases with s . The third figure illustrates the change in composition of applicants. The applicants with ability levels not attracted in control, but attracted in treatment, all have higher s . The functional forms and equations are provided in the Appendix.

Figure 3: Treatment effects on the applicant pool



Notes: Panel A reports kernel density estimates of exam scores in treatment and control. Panel B reports quantile treatment effect estimates on exam scores and 95% bootstrapped confidence intervals. Panels C and D present the mean levels of ability and prosociality in treatment and control across different levels of the applicant's skill rank in his or her health post. The confidence intervals are at the 95% level, based on standard errors clustered at the district level. "Applicants" are the 1,585 candidates who were interviewed for the position. Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. Ordinary levels or O-levels are administered by the Examinations Council of Zambia (ECZ) to 12th-grade students, the highest grade in the Zambian secondary education system. O-levels total exam score is constructed as the sum of inverted O-levels scores (1=9, 2=8, and so on) from all subjects in which the applicant wrote the exam, so that larger values correspond to better performance. O-levels passed in biology and other natural sciences equals the number of O-levels passed in biology, chemistry, physics, science, and agricultural science. Prosociality is the average of "Do you see yourself in the community in 5-10 years" (yes/no) and the Inclusion of Others in Self scale (Aron et al. 2004). Applicants are asked to choose between sets of pictures, each showing two circles (labeled "self" and "community") with varying degrees of overlap, from non-overlapping to almost completely overlapping.

Table 1: The effect of career opportunities on the applicant pool

	treatment	control	p-value (clustered)	p-value (EDF)	p-value (RI)
Applicants per health post	8.9 (5.6)	10.0 (6.5)	0.611	0.641	0.666
Cognitive skills (O-levels total exam score)	24.8 (9.8)	23.3 (9.3)	0.019	0.040	0.002
Cognitive skills (number of science O-levels)	1.44 (0.86)	1.24 (0.89)	0.006	0.017	0.024
Career motivation	0.25 (0.43)	0.19 (0.39)	0.026	0.050	0.050
Prosociality	2.34 (0.79)	2.51 (0.64)	0.237	0.295	0.416
Farmer	0.715 (0.45)	0.684 (0.47)	0.415	0.470	0.349
Age	25.7 (5.5)	26.0 (5.8)	0.433	0.487	0.487
Female	0.292 (0.45)	0.304 (0.46)	0.800	0.822	0.819

Notes: Columns 1 and 2 show means and standard deviations in parentheses. Sample includes the 1,585 candidates who were interviewed for the position. Treatment=1 if the candidate is interviewed in a district where career opportunities were made salient. Column 3 reports the p-values of the null hypothesis that the career treatment effect equals zero conditional on stratification variables and with standard errors clustered at the district level. Column 4 reports the p-value using the Young (2016) effective degrees of freedom (EDF) correction, clustered at the district level. Column 5 reports the p-value using a Randomization Inference (RI) procedure, clustered at the district level (specifically, the randomization-t p-value from Young (2018)). Ordinary levels or O-levels are administered by the Examinations Council of Zambia (ECZ) to 12th-grade students, the highest grade in the Zambian secondary education system. O-levels total exam score is constructed as the sum of inverted O-levels scores (1=9, 2=8, and so on) from all subjects in which the applicant wrote the exam, so that larger values correspond to better performance. O-levels passed in biology and other natural sciences equals the number of O-levels passed in biology, chemistry, physics, science and agricultural science. Career motivation=1 if the candidate chooses any combination of being an "environmental health technician," "clinical officer," or "doctor" in response to the question, "When you envision yourself in 5-10 years' time, what do you envision yourself doing?". Prosociality is the average of "Do you see yourself in the community in 5-10 years" (yes/no) and the Inclusion of Others in Self scale (Aron et al. 2004). Applicants are asked to choose between sets of pictures, each showing two circles (labeled "self" and "community") with varying degrees of overlap, from non-overlapping to almost completely overlapping. Farmer=1 if the applicant's main occupation is self-employment or work in the family farm in agriculture. Age is in years. Female=1 if the applicant is female.

Table 2: The effect of career opportunities on candidate selection by panels

	=1 if selected	p-value of difference	=1 if selected	p-value of difference
=1 if top 3 in skills X treatment	0.121 (0.029)	0.98	0.159 (0.036)	0.54
=1 if top 3 in skills X control	0.122 (0.038)		0.128 (0.040)	
=1 if top 3 prosociality X treatment	0.095 (0.039)	0.40	0.081 (0.042)	0.63
=1 if top 3 prosociality X control	0.058 (0.030)		0.056 (0.035)	
=1 if aims to higher rank X treatment	0.097 (0.042)	0.57	0.093 (0.039)	0.59
=1 if aims to higher rank X control	0.070 (0.032)		0.065 (0.034)	
=1 if connected to village leader X treatment			0.010 (0.040)	0.68
=1 if connected to village leader X control			0.028 (0.027)	
=1 if connected to health center staff X treatment			-0.038 (0.070)	0.62
=1 if connected to health center staff X control			-0.001 (0.041)	
Adjusted R-squared	0.096		0.096	
N	1518		1279	

Notes: OLS estimates. Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. All regressions include the stratification variables (province dummies and share of high school graduates in the district) and standard errors clustered at the district level, using the Young (2016) effective degrees of freedom correction. Independent variables are interacted with the treatment and control dummies. Top 3 skills=1 if the applicant's exam score is one of the 3 highest among applicants to the same health post. "Aims to be a higher-rank health professional in 5-10 years"=1 if the candidate chooses any combination of being an "environmental health technician," "clinical officer," or "doctor" in response to the question, "When you envision yourself in 5-10 years' time, what do you envision yourself doing?". Prosociality is the average of "Do you see yourself in the community in 5-10 years" (yes/no) and the Inclusion of Others in Self scale (Aron et al. 2004). Applicants are asked to choose between sets of pictures, each showing two circles (labeled "self" and "community") with varying degrees of overlap, from non-overlapping to almost completely overlapping. "Connected to village leader (health center staff)"=1 if the candidate chooses "political leader" or "village committee member" ("formally trained health worker") in response to the question, "Are any of your relatives or members of your household in the following positions?"

Table 3: Panel selection vs. random selection

	treatment	control
<i>Cognitive skills (O-levels total exam score)</i>		
panel selection	27.2	25.6
random draw: median	24.9	23
random draw: 10th and 90th pctl	[24.0; 25.8]	[22.2; 23.9]
standard deviation (main sample)	(9.8)	(9.3)
<i>Cognitive skills (number of science O-levels)</i>		
panel selection	1.55	1.46
random draw: median	1.42	1.25
random draw: 10th and 90th pctl	[1.35; 1.50]	[1.17; 1.34]
standard deviation (main sample)	(0.86)	(0.89)
<i>Career motivation (=1 if aims to higher rank in 5-10 yrs)</i>		
panel selection	0.36	0.25
random draw: median	0.27	0.19
random draw: 10th and 90th pctl	[0.24; 0.31]	[0.15; 0.22]
standard deviation (main sample)	(0.43)	(0.39)
<i>Prosociality</i>		
panel selection	2.55	2.55
random draw: median	2.46	2.54
random draw: 10th and 90th pctl	[2.40; 2.51]	[2.50; 2.59]
standard deviation (main sample)	(0.79)	(0.64)

Notes: Sample includes the 1,585 candidates who were interviewed for the position. Treatment=1 if the candidate is interviewed in a district where career opportunities were made salient. Panel selection reports the average trait of the two CHAs chosen by the panels in each health post. Random selection reports the average trait of two CHAs chosen randomly over 1000 draws. Ordinary levels or O-levels are administered by the Examinations Council of Zambia (ECZ) to 12th-grade students, the highest grade in the Zambian secondary education system. O-levels total exam score is constructed as the sum of inverted O-levels scores (1=9, 2=8, and so on) from all subjects in which the applicant wrote the exam, so that larger values correspond to better performance. O-levels passed in biology and other natural sciences equals the number of O-levels passed in biology, chemistry, physics, science and agricultural science. Career motivation=1 if the candidate chooses any combination of being an "environmental health technician," "clinical officer," or "doctor" in response to the question, "When you envision yourself in 5-10 years' time, what do you envision yourself doing?". Prosociality is the average of "Do you see yourself in the community in 5-10 years" (yes/no) and the Inclusion of Others in Self scale (Aron et al. 2004). Applicants are asked to choose between sets of pictures, each showing two circles (labeled "self" and "community") with varying degrees of overlap, from non-overlapping to almost completely overlapping.

Table 4: The effect of career opportunities on the number of visits

dependent variable	Household visits					
	source time horizon unit of observation	SMS receipts months 1-18	SMS receipts months 1-6	SMS receipts months 7-12	SMS receipts months 13-18	
		CHA	CHA	CHA	CHA	
		(1)	(2)	(3)	(4)	
		Treatment	93.86 (37.11)	33.86 (15.93)	29.57 (13.47)	30.42 (12.90)
Area characteristics	Yes	Yes	Yes	Yes		
Mean of dependent variable in control	319.0	167.1	92.1	59.8		
Adjusted R-squared	0.113	0.116	0.064	0.106		
N	307	307	307	307		
Median Treatment Effect (se)	104.4 (43.9)	51.3 (21.0)	46.6 (21.0)	31.8 (15.0)		
EDF p-value	0.030	0.066	0.058	0.042		
RI p-value	0.030	0.087	0.024	0.051		

Notes: OLS Estimates, standard errors clustered at the district level. EDF p-value refers to the p-value from a null hypothesis that the treatment effect is zero (in the same regression), using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)). Standard errors for the median treatment effect are bootstrapped and clustered at the district level. The dependent variable is total number of households visited over the relevant time horizon. SMS receipts are sent by individual CHAs to MOH for each visit. Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. All regressions include the stratification variables (province dummies and share of high school graduates in the district). Area characteristics include: number of staff in the health post, geographical distribution of households in the catchment area, and an indicator variable that equals 1 if the CHA reports to have good cell network coverage most of the time or all the time.

Table 5: Compensation mechanisms

dependent variable	source	women and children visited per HH				community mobilization meetings		patients seen at health post		emergency calls
		retention		visit duration		unique HHs visited		visits per HH		
		SMS receipts	SMS receipts	CHA	CHA	HMIS records	HMIS records	SMS receipts	SMS receipts	
		CHA	CHA	health post	health post	CHA	CHA	health post	health post	
unit of observation		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
Treatment		0.0025 (0.0477)	0.260 (1.849)	0.044 (0.0946)	36.34 (15.47)	0.487 (0.246)	17.08 (5.226)	32.65 (260.2)		0.047 (0.0581)
Area characteristics		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dep var in control		0.820	33.7	2.07	187.2	1.760	20.32	1126.6		0.460
Adjusted R-squared		0.041	0.012	0.072	0.121	0.125	0.027	0.030		0.006
N		307	304	142	307	307	146	146		298
EDF p-value		0.963	0.900	0.672	0.043	0.085	0.006	0.910		0.473
RI p-value		0.966	0.889	0.658	0.046	0.029	0.012	0.919		0.517

Notes: OLS Estimates, standard errors clustered at the district level. EDF p-value refers to the p-value from a null hypothesis that the treatment effect is zero (in the same regression), using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)). Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. Retention=1 if the CHA still reports visits after 1 year. Visit duration is computed as end time minus start time in minutes. Emergency calls=1 if the CHA takes at least 1 out of hours call in a typical week. SMS receipts are sent by individual CHAs to MOH for each visit. The Health Management and Information System (HMIS) is the Zambian Ministry of Health's system for reporting health services data at government facilities. The two CHAs are required to submit monthly reports that summarize their activities at the health post/community level. The number of observations varies because some health posts do not submit the reports; these are equally distributed between treatments. The time use survey was administered in May 2013 during a refresher training program. All regressions include the stratification variables (province dummies and share of high school graduates in the district). Area characteristics include: number of staff in the health post, geographical distribution of households in the catchment area, and an indicator variable that equals 1 if the CHA reports to have good cell network coverage most of the time or all the time.

Table 6: The effect of career opportunities on facility utilization

Dependent variable: total over each quarter 2011:1-2014:2	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	institutional deliveries	postnatal (0-6 weeks) visits	children under 5 visited	children under 5 weighed	children under 1 receiving BCG vaccinations	children under 1 receiving polio vaccinations	children under 1 receiving measles vaccinations	average standardized effect
Treatment	-2.17 (11.27)	-12.88 (9.314)	-65.80 (142.7)	-72.88 (133.3)	11.01 (11.96)	-0.286 (9.118)	1.730 (10.00)	-0.005 (0.156)
After	4.25 (4.27)	15.50 (5.11)	61.78 (63.40)	108.9 (64.07)	-1.270 (4.618)	-1.124 (3.711)	-1.168 (3.604)	0.042 (0.058)
Treatment*After	13.65 (6.21)	7.79 (9.380)	312.1 (97.67)	277.8 (109.7)	7.147 (8.881)	14.60 (4.809)	11.19 (7.259)	0.278 (0.092)
Area characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dependent variable in control in year 1	45.3	55.4	1285.7	1236.1	83.0	72.4	70.4	na
Adjusted R-squared	0.348	0.215	0.254	0.253	0.151	0.151	0.118	na
Number of facilities	89	119	123	123	120	121	120	na
Number of observations	1301	1543	1618	1610	1518	1531	1535	1102
EDF p-value	0.038	0.418	0.003	0.017	0.433	0.005	0.138	
RI p-value	0.014	0.455	0.010	0.033	0.463	0.003	0.134	

Notes: OLS Estimates, standard errors clustered at the district level. EDF p-value refers to the p-value from a null hypothesis that Treatment*After is zero (in the same regression), using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)). Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. Data source is the Health Management and Information System (HMIS) available monthly from January 2011 until June 2014. Health center and health post staff are required to submit monthly reports that summarize their activities at the health post/community level. These are aggregated at the quarter level in the regressions. The variable in Column 1 is defined at the health center level because health centers are equipped for child births and health posts are not. The variables in Columns 2-7 are defined at the health post level if this reports data, at the health center otherwise. The average standardized treatment effect is computed using the methodology in Kling et al. (2001). After=1 after September 2012 (from 2012:4 onwards), when CHAs started working. All regressions include the stratification variables (province dummies and share of high school graduates in the district). Area characteristics include: number of staff in the health post, geographical distribution of households in the catchment area, and an indicator variable that equals 1 if the CHA reports to have good cell network coverage most of the time or all the time.

Table 7: The effect of career opportunities on health practices and outcomes

Dependent variable	Information		Health practices			Incidence of illness			Anthropometrics				All	
	% of correct answers in medical knowledge test		=1 if child is breastfed up to 2 years	=1 if stool are safely disposed	number of deworming treatments	=1 if child's CHA is on track with immunization schedule	=1 if child has experienced fever in the last two weeks	=1 if child has experienced diarrhea in the last two weeks	=1 if weight for age z score <2 SD (moderately or severely undernourished)	=1 if weight for age z score <3 SD (severely undernourished)	=1 if MUAC<12.5 (moderately or severely wasted)	=1 if MUAC<11.5 (severely wasted)	average standardized effect (anthro only)	average standardized effect
Treatment	0.002 (0.010)		0.051 (0.023)	0.121 (0.039)	0.225 (0.129)	0.047 (0.020)	-0.003 (0.037)	0.037 (0.027)	-0.070 (0.033)	-0.053 (0.030)	-0.028 (0.015)	-0.014 (0.014)	0.124 (0.062)	0.108 (0.037)
Household controls	yes		yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Child controls	no		yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Mean of dep var in control	0.740		0.641	0.597	1.45	0.058	0.468	0.256	0.448	0.210	0.051	0.036	0.015	
Adjusted R-squared	0.057		0.561	0.161	0.263	0.024	0.077	0.017	0.021	-0.006	0.003	0.018	0.017	
N	738		613	736	659	462	731	731	582	582	581	581	579	376
EDF p-value	0.827		0.054	0.007	0.119	0.038	0.946	0.216	0.060	0.114	0.084	0.179	0.342	
RI p-value	0.784		0.072	0.005	0.176	0.045	0.951	0.262	0.107	0.058	0.099	0.205	0.359	

Notes: OLS estimates, standard errors clustered at the district level. EDF p-value refers to the p-value from a null hypothesis that the treatment effect is zero (in the same regression), using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)). Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. The medical knowledge test contains 14 questions on topics that CHAs are supposed to cover; these questions were drafted by the researchers in consultation with CHA program officials and the CHA curriculum. Breastfeeding and stool disposal are self-reported. In line with UNICEF (2014), we define stools as safely disposed if flushed in toilet/laraine. Deworming and immunization data are as reported in the child health card. A child is defined as on track if they have completed all immunizations required for their age in months. The immunization sample is restricted to children who were 3 months or younger (including unborn) when the CHAs started working. Thresholds for weight-for-age and MUAC are taken from WHO guidelines; following these, data are restricted to children between 6-59 months. Household controls include size, education level of the respondent, and number of assets. Child controls include age and gender. All regressions include the stratification variables. The average standardized treatment effect is computed using the methodology in Kling et al. (2001) after recoding all variables so that higher values indicate better outcomes. For weight-for-age z-score and MUAC we use the lowest thresholds.

Table A.1: Eligible population by treatment (randomization balance)

	treatment	control	p-value (clustered)	p-value (EDF)	p-value (RI)
A. Characteristics of the eligible population					
Share of eligibles in the district (18-45 year olds with grade 12 or above)	.044 (.205)	.043 (.203)	.917	.919	.961
Share of women among the eligibles	.371 (.483)	.391 (.488)	.241	.258	.142
Main activity of eligible candidates during the past 12 months:					
<i>not working</i>	.296 (.456)	.280 (.449)	.480	.494	.613
<i>unpaid work</i>	.201 (.401)	.229 (.420)	.344	.361	.522
<i>paid work</i>	.457 (.498)	.437 (.496)	.353	.371	.408
<i>of which: mid skill</i>	.240 (.427)	.230 (.421)	.705	.714	.724
<i>of which: low skill</i>	.484 (.500)	.453 (.498)	.173	.189	.149
B. Catchment area characteristics					
Number of staff in health post*	1.49 (1.10)	1.39 (1.17)	.635	.640	.604
Geographical distribution of households in catchment area:*					
<i>Most people live in their farms, none in villages</i>	.082 (.277)	.091 (.289)	.846	.848	.825
<i>Some people live in farms, some in small villages (5-10hh)</i>	0.518 (.503)	.532 (.502)	.878	.879	.879
<i>Most people live in medium/large villages (more than 10hh), a few on their farms</i>	.388 (.490)	.364 (.484)	.816	.819	.804
Poor cell network coverage*	.082 (.277)	.065 (.248)	.681	.685	.743
C. Target population characteristics					
Share of district population under 5	.187 (.390)	.187 (.390)	.915	.916	.877
Main type of toilet: Pit latrine or better**	.718 (.450)	.676 (.468)	.554	.561	.449
Household water supply: Protected borehole or better**	.361 (.480)	.400 (.490)	.416	.425	.300

Notes: Columns 1 and 2 show means and standard deviations in parentheses. Column 3 reports the p-value of the test of equality of means based on standard errors clustered at the district level. Column 4 reports the p-value using the Young (2016) effective degrees of freedom (EDF) correction, clustered at the district level. Column 5 reports the p-value using a Randomization Inference (RI) procedure, clustered at the district level (specifically, the randomization-t p-value from Young (2018)). Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. Variables are drawn from the 2010 Census (10% PUMS sample) except those indicated by *, which are drawn from our surveys, and those indicated by **, which are drawn from the 2010 Living Conditions Monitoring Survey (LCMS), which covers 20,000 HHs and is representative at the district level. Activities codes follow the ILO ISCO88 convention. Mid-skill includes ISCO codes between 300 and 599, namely technicians, clerical workers and services and sales workers. Low-skill includes ISCO codes above 600, namely agriculture, crafts, basic manufacturing and elementary occupations. Number of staff in health post is the total number of nurses, environmental health technicians, and clinical officers assigned to the health post as reported by district officials surveyed by phone. Information on the geographical distribution of HHs was obtained from a survey of the deployed CHAs before deployment. CHAs were shown stylized maps accompanied by a description and asked to choose the one that most closely resembled the catchment area of their health post. Questions were asked to each CHA individually so that two CHAs from the same health post could give different answers. For the 5 out of 161 cases in which the two CHAs gave different answers, we use the information provided by supervisors to break the tie. To measure cell network coverage we attempt to call all CHAs after deployment. We make daily calls for 118 consecutive days. The health post is classified as having poor coverage if we do not manage to reach either of its two CHAs during this period. Main type of toilet: Pit latrine or better=1 if the surveyed household uses a pit latrine, ventilated improved pit (VIP), or flush toilet, and 0 if bucket, other, or no toilet. Household water supply: Protected borehole or better=1 if the water supply comes from a protected borehole or well, communal tap, or other piped water system, and 0 if it comes from an unprotected well or borehole, river/dam/stream, rain water tank, or other.

Table A.2: Experimental checks

Expected job benefits	at entry (June 2011)				on the job (May 2013)					
	treatment	control	p-value (clustered)	p-value (EDF)	p-value (RI)	treatment	control	p-value (clustered)	p-value (EDF)	p-value (RI)
Good future career	.164 (.157)	.119 (.113)	.002	.006	.007	.160 (.122)	.151 (.108)	.650	.682	.644
Allows me to serve the community	.397 (.226)	.436 (.243)	.046	.072	.078	.364 (.184)	.371 (.178)	.914	.922	.923
Earns respect and status in the community	.037 (.094)	.059 (.110)	.031	.053	.074	.039 (.069)	.038 (.061)	.882	.893	.909
Interesting job	.149 (.163)	.148 (.138)	.987	.988	.991	.132 (.103)	.138 (.106)	.354	.403	.438
Allows me to acquire useful skills	.181 (.169)	.161 (.137)	.254	.304	.365	.216 (.132)	.219 (.111)	.888	.899	.894
Offers stable income	.026 (.056)	.024 (.055)	.645	.677	.683	.038 (.069)	.039 (.060)	.989	.990	.993
Pays well	.032 (.093)	.025 (.057)	.414	.461	.532	.051 (.089)	.043 (.067)	.664	.695	.723

Notes: Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. CHAs were given 50 beans and asked to allocate them on cards, listing different reasons in proportion to the importance of each benefit for working as a CHA. The cards were scattered on a table in no particular order. "At entry" variables are drawn from a survey administered at the beginning of the training program. "On the job" variables are drawn from a survey administered eight months after the CHAs started working. We show means with standard deviations in parentheses and the p-value of the test of equality of means based on standard errors clustered at the district level (conditional on stratification variables). Column 4 reports the p-value using the Young (2016) effective degrees of freedom (EDF) correction, clustered at the district level. Column 5 reports the p-value using a Randomization Inference (RI) procedure, clustered at the district level (specifically, the randomization-t p-value from Young (2018)).

Table A.3: Validation of household visit measures

dependent variable	source	number of visits from		= 1 if HH reports a visit by		HH satisfaction: overall	
		HMIS records		CHA		CHA's services	
		unit of observation		HH survey		HH survey	
		(1)	(2)	(3)	(4)	(5)	(6)
Number of visits (in 00s) reported by CHA via SMS receipts		0.767 (0.083)	0.644 (0.161)	0.0208 (0.0083)	0.0154 (0.0173)	0.0393 (0.0194)	0.0444 (0.0162)
Number of visits (in 00s) reported by CHA via SMS receipts* <i>Treatment</i>			0.192 (0.187)		0.00991 (0.0192)		-0.00261 (0.0356)
Mean of dependent variable		643.6		0.438		4.332	
Adjusted R-squared		0.473	0.473	0.014	0.013	0.013	0.018
N		145	145	1284	1284	1253	1253
EDF p-value of treatment interaction		-	0.320	-	0.641	-	0.947
RI p-value of treatment interaction		-	0.358	-	0.617	-	0.950

Notes: OLS estimates, standard errors clustered at the health post level in Columns 3-6. The independent variable is visits reported by SMS between September 2012 and January 2014. The dependent variable in Columns 1 and 2 is the total number of visits done by the two CHAs in the health post drawn from HMIS administrative data over the period between September 2012 and January 2014. The dependent variables in Columns 3-6 are drawn from a HH survey administered to 16 HHs in each of 47 communities where CHAs are active. Satisfaction measures range from 1 (very dissatisfied) to 5 (very satisfied). EDF p-value refers to the p-value from a null hypothesis that the treatment interaction is zero (in the same regression), using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)).

References

- Akerlof, George A and Rachel E Kranton, “Identity and the Economics of Organizations,” *The Journal of Economic Perspectives*, 2005, 19 (1), 9–32.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin Olken, Ririn Purnamasari, and Matthew Wai-Poi, “Self-Targeting: Evidence from a Field Experiment in Indonesia,” *Journal of Political Economy*, 124 (2), 371–427.
- Aron, Arthur and Others, “Including Others in the Self,” *European Review of Social Psychology*, 2004, 15 (1), 101–132.
- , Elaine N Aron, and Danny Smollan, “Inclusion of Other in the Self Scale and the Structure of Interpersonal Closeness,” *Journal of Personality and Social Psychology*, 1992, 63 (4), 596.
- Ashraf, Nava, James Berry, and Jesse M Shapiro, “Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia,” *American Economic Review*, 2010, 100 (5), 2383–2413.
- , Oriana Bandiera, and B Kelsey Jack, “No Margin, No Mission? A Field Experiment on Incentives for Public Service Delivery,” *Journal of Public Economics*, 2014, 120, 1–17.
- Basinga, Paulin, Paul J Gertler, Agnes Binagwaho, Agnes L B Soucat, Jennifer Sturdy, and Christel M J Vermeersch, “Effect on Maternal and Child Health Services in Rwanda of Payment to Primary Health-care Providers for Performance: An Impact Evaluation,” *The Lancet*, 2011, 377 (9775), 1421–1428.
- Bénabou, Roland and Jean Tirole, “Intrinsic and Extrinsic Motivation,” *Review of Economic Studies*, 2003, 70, 489–520.
- and –, “Incentives and Prosocial Behavior,” *The American Economic Review*, 2006, 96 (5), 1652–1678.
- and –, “Identity, Morals, and Taboos: Beliefs as Assets,” *The Quarterly Journal of Economics*, 2011, 126 (2), 805–855.
- Bertrand, Marianne, Robin Burgess, Arunish Chawla, and Guo Xu, “The Costs of Bureaucratic Rigidity: Evidence from the Indian Administrative Service,” 2016.
- Besley, Timothy and Maitreesh Ghatak, “Competition and Incentives with Motivated Agents,” *American Economic Review*, 2005, 95 (3), 616–636.
- Brehm, John and Scott Gates, *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public*, University of Michigan Press, 1999.
- Celhay, Pablo, Paul Gertler, Paula Giovagnoli, and Christel Vermeersch, “Long-Run Effects of Temporary Incentives on Productivity in Medical Care Clinics,” *American Economic Journal: Applied Economics*, 2019, *forthcomin*.
- Collier, Paul, “Rethinking the Provision of Public Services in Post-Conflict States,” *Contracting Out Government Functions and Services: Emerging Lessons from Post-Conflict and Fragile Situations*, 2009, pp. 115–122.
- Dal Bó, Ernesto, Frederico Finan, and Martin A Rossi, “Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service,” *The Quarterly Journal of*

- Economics*, 2013, 128 (3), 1169–1218.
- de Poel, Ellen, Gabriela Flores, Por Ir, Eddy Van Doorslaer, and Others**, “Can Vouchers Deliver? An Evaluation of Subsidies for Maternal Health Care in Cambodia,” *Bulletin of the World Health Organization*, 2014, 92 (5), 331–339.
- Deserranno, Erika**, “Financial Incentives as Signals: Experimental Evidence from the Recruitment of Health Workers,” 2019.
- Duflo, Esther, Pascaline Dupas, and Michael Kremer**, “School Governance, Teacher Incentives, and Pupil–Teacher Ratios: Experimental Evidence from Kenyan Primary Schools,” *Journal of Public Economics*, 2015, 123, 92–110.
- , **Rema Hanna, and Stephen P Ryan**, “Incentives Work: Getting Teachers to Come to School,” *American Economic Review*, jun 2012, 102 (4), 1241–1278.
- Food and Nutrition Technical Assistance Project**, “Anthropometry: Assessing Children Under 5 Pocket Reference,” 2011.
- Fryer, Roland**, “Teacher Incentives and Student Achievement: Evidence from New York City Public Schools,” *Journal of Labor Economics*, 2013, 31 (2), 373–427.
- Glewwe, Paul, Nauman Ilias, and Michael Kremer**, “Teacher Incentives,” *American Economic Journal: Applied Economics*, 2010, 2 (July), 205–227.
- Government of Zambia**, “Zambia Demographic and Health Survey 2013-14: Preliminary Report,” Technical Report 2014.
- Kling, Jeffrey R, Jeffrey B Liebman, and Lawrence F Katz**, “Experimental Analysis of Neighborhood Effects,” *Econometrica*, 2007, 75 (1), 83–119.
- Kremer, Michael, Conner Brannen, and Rachel Glennerster**, “The Challenge of Education and Learning in the Developing World,” *Science*, 2013, 340 (6130), 297–300.
- La Familia Sana Program**, *A Historical Overview of Lay Health Worker Programs*, La Familia Sana, Inc., 1992.
- Lazear, Edward and Paul Oyer**, “Personnel Economics,” in Robert Gibbons and John Roberts, eds., *The Handbook of Organizational Economics*, Princeton, NJ: Princeton University Press, 2012, chapter 12, pp. 479–519.
- Lazear, Edward P**, “Performance Pay And Productivity,” *American Economic Review*, 2000, 90 (2), 1346–1361.
- Lehmann, Uta and David Sanders**, “Community Health Workers: What Do We Know About Them? The State of the Evidence on Programmes, Activities, Costs and Impact on Health Outcomes of Using Community Health Workers,” *World Health Organization*, 2007, 2, 1–42.
- Lipsky, Michael**, *Street-Level Bureaucracy: Dilemmas of the Individual in Public Service*, Russell Sage Foundation, 1980.
- Mashek, Debra, Lisa Cannaday, and June Tangney**, “Inclusion of Community in Self Scale: A Single-item Pictorial Measure of Community Connectedness,” *Journal of Community Psychology*, 2007, 35, 257–275.
- Miller, Grant and Kim Babiarz**, “Pay-for-Performance Incentives in Low- and Middle-Income

- Country Health Programs,” in Anthony J. Culyer, ed., *Encyclopedia of Health Economics*, Elsevier, 2014, pp. 457–466.
- , **R. Luo, L. Zhang, S. Sylvia, Y. Shi, P. Foo, Q. Zhao, R. Martorell, A. Medina, and S. Rozelle**, “Effectiveness of Provider Incentives for Anaemia Reduction in Rural China: A Cluster Randomised Trial,” *BMJ*, jul 2012, *345* (jul27 1), e4809–e4809.
- Muralidharan, Karthik and Venkatesh Sundararaman**, “Teacher Performance Pay: Experimental Evidence from India,” *Journal of Political Economy*, feb 2011, *119* (1), 39–77.
- and —, “Contract Teachers: Experimental Evidence from India,” sep 2013.
- , **Nazmul Chaudhury, Jeffrey Hammer, Michael Kremer, and Halsey Rogers**, “Is There a Doctor in the House? Medical Worker Absence in India,” 2011.
- Mwangome, Martha K, Greg Fegan, Ronald Mbunya, Andrew M Prentice, and James A Berkley**, “Reliability and Accuracy of Anthropometry Performed by Community Health Workers Among Infants Under 6 Months in Rural Kenya,” *Tropical Medicine and International Health*, 2012, *17* (5), 622–629.
- Myatt, Mark, Tanya Khara, and Steve Collins**, “A Review of Methods to Detect Cases of Severely Malnourished Children in the Community for their Admission into Community-based Therapeutic Care Programs,” *Food and Nutrition Bulletin*, 2006, *27* (3 Supplement), S7—S23.
- North, Douglass C**, “Institutions,” *The Journal of Economic Perspectives*, 1991, *5* (1), 97–112.
- Northcote, Stafford and C E Trevelyan**, “Report on the Organisation of the Permanent Civil Service,” *Report to the House of Commons, London*, 1853.
- Oyer, Paul and Scott Schaefer**, “Personnel Economics: Hiring and Incentives,” in David Card and Orley Ashenfelter, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, 2011, chapter 20, pp. 1769–1823.
- Perez, Leda M and Jacqueline Martinez**, “Community Health Workers: Social Justice and Policy Advocates for Community Health and Well-Being,” *American Journal of Public Health*, jan 2008, *98* (1), 11–14.
- Prendergast, Canice**, “The Motivation and Bias of Bureaucrats,” *American Economic Review*, feb 2007, *97* (1), 180–196.
- Rockoff, Jonah E, Douglas O Staiger, Thomas J Kane, and Eric S Taylor**, “Information and Employee Evaluation: Evidence from a Randomized Intervention in Public Schools,” *The American Economic Review*, 2012, pp. 3184–3213.
- Rothstein, Jesse**, “Teacher Quality Policy When Supply Matters,” *American Economic Review*, jan 2015, *105* (1), 100–130.
- Roy, Andrew D**, “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 1951, *3* (2), 136–146.
- Schultz, P Wesley**, “Inclusion with Nature: The Psychology of Human-Nature Relations,” in Peter Schmuck and P Wesley Schultz, eds., *Psychology of Sustainable Development*, Dordrecht: Kluwer Academic Publishers, 2012, pp. 61–78.
- Staiger, Douglas O and Jonah E Rockoff**, “Searching for Effective Teachers with Imperfect

- Information,” *The Journal of Economic Perspectives*, 2010, *24* (3), 97–117.
- US Human Resources and Service Administration**, “Shortage Areas,” 2019.
- Weber, Max**, *Economy and Society*, Tübingen, 1922.
- Wilson, James Q**, *Bureaucracy: What Government Agencies Do and Why They Do It*, New York: Basic Books, 1989.
- World Health Organization**, “The World Health Report 2006: Working Together for Health,” 2006.
- Young, Alwyn**, “The African Growth Miracle,” *Journal of Political Economy*, 2012, *120* (4), 696–739.
- , “Improved, Nearly Exact, Statistical Inference with Robust and Clustered Covariance Matrices using Effective Degrees of Freedom Corrections,” 2016.
- , “Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results,” *The Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

WEB APPENDIX (for online publication only)

A Randomization Balance

Table A.1 describes three sets of variables that can affect the supply of health workers, the demand for their services, and their working conditions. For each variable, the table reports the means and standard deviations in treatment and control, as well as the p-value of the test of means equality, with standard errors clustered at the level of randomization, the district. Table A.1 shows that the randomization yielded a balanced sample, as all p-values of the tests of equality are greater than 0.05. As treatment and control means are very close throughout, we comment on treatment group values in the rest of this section.

Panel A reports statistics on the eligible population drawn from the 2010 Census. This shows that the eligibles—namely, 18-45 year-old Zambian citizens with at least grade 12 education—account for 4.4% of the district population, and that among them, 37% are female. Just over half of the eligible were either out of work or in unpaid employment over the past twelve months.³⁷ Among the 45% engaged in income-generating activities (either as employees or self-employed), fewer than one-third are employed in high-skill occupations (such as teachers, who account for 9% of the eligible population), and about half are employed in low-skill occupations, mostly in agriculture which accounts for 18% of the eligible population. Taken together, the evidence suggests that, despite their educational achievements, the majority of the eligible population is either out of work or employed in occupations below their skill level.

Panel B illustrates the characteristics of the catchment areas. These variables are drawn from surveys administered to district officials and the health workers themselves. Three points are of note. First, health posts are poorly staffed in both the treatment and control groups; the average number of staff (not including the new health workers) is 1.5. Given that the aim is to assign two health workers to each health post, the program more than doubles the number of health staff in these communities. Second, the areas vary in the extent to which households live on their farms or in villages, but the frequency of either type is similar in the treatment and control groups. This is relevant as travel times between households depend on population density and are higher when households are scattered over a large area, as opposed to being concentrated in a village. Third, over 90% of the catchment areas in both groups have at least some cell network coverage, which is relevant for our analysis, as some performance measures are collected via SMS messages.

Panel C illustrates the characteristics of the target population that are relevant for the demand for health worker services. First, children under 5, who (together with pregnant women) are the main targets of the health workers, account for 19% of the population. Second, Panel C shows that access to latrines and—most noticeably—protected water supply is limited in these areas. Lack of latrines and protected water supply favors the spread of waterborne infections, to which pregnant

³⁷The 28% who were out of work are either unemployed (13%), housewives (7.5%), or full-time students (8.5%). Most (65%) of the unpaid jobs are in agriculture. These are balanced across treatments.

women and children are particularly vulnerable and, through this, affects the demand for health workers' services.

B Model

B.1 Solution

In Stage 3, we have an interior $e^*(a, s)$ defined by $sH_e(a, e^*) - d'(e^*) = 0$ which is a global maximum as $U_{ee} = sH_{ee}(a, e) - d''(e) < 0$ for all $e > 0$. For the comparative statics of $e^*(a, s)$ with respect to a and s , we have $\frac{de^*}{da} = \frac{-sH_{ea}}{sH_{ee} - d''(e^*)} \geq 0$ and $\frac{de^*}{ds} = \frac{-sH(a, e^*)}{sH_{ee} - d''(e^*)} > 0$ as both numerator and denominator are negative in both cases. Defining $H^*(a, s) \equiv H(a, e^*(a, s))$, we have $\frac{dH^*}{da} = \frac{\partial H}{\partial a} + \frac{\partial H}{\partial e} \cdot \frac{de^*}{da} > 0$ and $\frac{dH^*}{ds} = \frac{\partial H}{\partial e} \cdot \frac{de^*}{ds} > 0$. A straightforward application of the Envelope Theorem shows that $U_a^* > 0$ and $U_s^* > 0$, and, finally, it is clear that $U_M^* = 1 > 0$.

In Stage 1, with $E(0, s, M) < V(0)$ an interior solution requires $V_a < E_a$ for some a . If so, there is a threshold of ability \underline{a} such that $E(\underline{a}, s, M) = V(\underline{a})$ and $E_a(\underline{a}, s, M) > V_a(\underline{a})$ so that all i with $a_i < \underline{a}$ do not apply. If $E(a, s, M) > V(a)$ for all $a_i > \underline{a}$ everybody with $a_i > \underline{a}$ will apply. If however there is a value of a such that $E(\bar{a}, s, M) = V(\bar{a})$, it must be that $E_a(\bar{a}, s, M) < V_a(\bar{a})$ and such that only i with $\underline{a} < a_i < \bar{a}$ apply.

B.2 Comparative statics with respect to M

Result 1: Increasing material benefits M will attract higher-ability applicants who would not apply otherwise ($\frac{\partial \bar{a}}{\partial M} > 0$) and either (i) lower the ability of the lowest-ranked applicant ($\frac{\partial \underline{a}}{\partial M} < 0$) and increase the total number of applicants or (ii) discourage low-ability applicants ($\frac{\partial \underline{a}}{\partial M} > 0$) and have an ambiguous impact on the total number.

To prove the first statement, note that the total differential of $E(\bar{a}, s, M) = V(\bar{a})$ implies $\frac{d\bar{a}}{dM} = \frac{E_M}{(V_a - E_a)}$. The denominator is positive since $E_a(\bar{a}, s, M) < V_a(\bar{a})$; the numerator is equal to $E_M = p_M(\bar{a}, s, M)U(.) + pU_M$. The first term captures the effect of M on the probability of being selected, which depends on the relative ability of the applicant vs. the other applicants; from $\frac{d^2 p}{dM da} > 0$, we assume for simplicity that the marginal applicant's value of a is sufficiently large that this effect is negligible. The second term is positive as shown above. Thus $E_M > 0$ and $\frac{d\bar{a}}{dM} = \frac{E_M}{(V_a - E_a)} > 0$. Likewise, $\frac{d\underline{a}}{dM} = \frac{E_M}{(V_a - E_a)}$. The denominator is negative since $E_a(\underline{a}, s, M) > V_a(\underline{a})$, and again the numerator is equal to $E_M = p_M(\underline{a}, s, M)U(.) + pU_M$. The first term is negative because an applicant with skill a is less likely to be selected under high M since, as seen above, this attracts higher-ability applicants. The second term is positive as discussed above. Thus, if the increase in payoff U^* is larger than the discouragement due to lower probability of being selected, then $E_M = p_M U(.) + pU_M > 0$ which implies that the lower threshold decreases and overall more people apply. In contrast, if $E_M = p_M U(.) + pU_M < 0$, then the lower threshold increases, and the effect on the number of applicants depends on the distribution of a in the population that, in turn, determines whether the number of low-ability applicants who no longer apply is larger than the number of high-ability applicants who only apply with high M .

Result 2: *Under any M , the most able applicant is also the most prosocial. An increase in M leaves the prosociality of the marginal applicant unchanged and has an ambiguous effect on the prosociality of the average applicant.*

Taking the total differential of $E(\bar{a}, s, M) - V(\bar{a})$ with respect to a and s gives $E_s ds + E_a da = V_a da$. Hence $ds/da = \frac{V_a - E_a}{E_s}$. Given that $E_a(\bar{a}, s, M) < V_a(\bar{a})$ we have $ds/da > 0$, and the applicant with the highest a , i.e. with the highest \bar{a} , has $s = 1$. This shows that *for any M , the most able applicant is also the most prosocial*. As M increases, and we have $\frac{d\bar{a}}{dM} = \frac{E_M}{(V_a - E_a)} > 0$, the marginal applicant has higher \bar{a} and the same $s = 1$, *which proves that an increase in M leaves the prosociality of the marginal applicant unchanged*. Now consider two levels of M , $M_T > M_C$. The marginal candidate when $M = M_T$ has ability $\bar{a}(1, M_T)$ while when $M = M_C$ the marginal candidate has ability $\bar{a}(1, M_C) < \bar{a}(1, M_T)$. Both candidates have $s = 1$. Define \hat{s} the level of s such that $E(\bar{a}(1, M_C), \hat{s}, M_T) = V(\bar{a}(1, M_C))$ that is the applicant with ability $\bar{a}(1, M_C)$ is indifferent between applying at the higher level of M or not. Then all the candidates who do not apply when M is low and apply when M is high must have $s > \hat{s}$, whilst those who apply when M is low have $0 \leq s \leq 1$. Thus the new applicants increase the mean prosociality. However all applicants with $a < \bar{a}(s, M_C)$, whose prosociality was such that they were indifferent between applying and not when $M = M_C$ will be strictly better off when $M = M_T$. This implies that for any $a < \bar{a}(s, M_C)$, raising M will attract lower prosociality applicants. The net effect depends on the relative strength of these two channels and is therefore ambiguous.

In the graphical illustration, we use $H = 2\sqrt{\xi}e + a$, where $\xi > 0$ is a parameter, and $c(e) = e^2$. A selected applicant thus chooses e to maximise $s(2\sqrt{\xi}e + a) - e^2 + M$, which yields $e^* = \sqrt{\xi}s$, $H^* = 2\xi s + a$, and $U^* = \xi s^2 + as + M$. The probability of selection must satisfy the following: (i) increasing and concave in a and s , (ii) decreasing in M , (iii) $\frac{\partial^2 p}{\partial a \partial M} > 0$. As explained in the main body, we use the functional form $p(a, s, M) = \frac{\gamma as^\beta + \mu(M_T - M)}{1 + \mu(M_T - M)}$, where $\gamma > 0$ is set to guarantee that $p \in [0, 1]$, $\beta > 0$ captures the weight that the panel puts on prosociality, and $\mu > 0$ scales the decrease in probability that arises from more (high-ability) people applying in equilibrium when M rises. For interpretation purposes, it is useful to note that the probability is 1 when $\gamma as^\beta = 1$. This function satisfies all three criteria: (i) is straightforward from inspection, (ii) can be seen in $\frac{dp}{dM} = -\frac{\mu(1-p)}{1 + \mu(M_T - M)} < 0$, (iii) can be seen in $\frac{d^2 p}{dM da} = \frac{\mu}{1 + \mu(M_T - M)} \cdot \frac{dp}{da} > 0$. Moreover, for (iii), we see that $\frac{dp}{dM}$ goes to zero as the probability goes to one. We set $V = va^2 + a$, where we assume $v > \gamma$ which ensures that $V_a > U_a^* > 0$ everywhere as $V_a \geq 1 \geq s^2 = U_a^*$. This also ensures that $V_a > E_a > 0$ for all a sufficiently large, as this ensures that the quadratic term in the polynomial below is positive for any s and M . Combining these, an applicant in Stage 1 thus applies if $\frac{\gamma as^\beta + \mu(M_T - M)}{1 + \mu(M_T - M)}(\xi s^2 + as + M) - c > va^2 + a$. This defines a polynomial of degree 2 given by $\phi(a) \equiv [v(1 + \mu(M_T - M)) - \gamma s^{\beta+1}]a^2 + [1 + \mu(1 - s)(M_T - M) - \gamma s^\beta(M + \xi s^2)]a + [c + \mu(M_T - M)(c - \xi s^2 - M)]$. The two thresholds are thus given by $\phi(a) = 0$. The upper root is plotted for a range of s in Figure 2b and Figure 2c. It can be shown that $\frac{d\bar{a}}{ds} > 0$, which reflects the fact that the curve is upward sloping, while $\frac{d\bar{a}}{dM}$ is ambiguous. In the Figure, we have $\frac{d\bar{a}}{dM} > 0$ for all s , although this is not necessary for the main results: the important aspect for the panel's selection mechanism is the set

of candidates with high a and high s (i.e., in the north-east section of the graph) who will have high probability of being selected and thus will have $\frac{da}{dM} > 0$.

C Time Use

We surveyed the health workers in May 2013, nine months after they started working.³⁸ The survey asked the health workers to report the frequency of emergency visits typically done outside of working hours. The median health worker does one emergency call per week, and Column 8 of Table 5 shows that this holds true for health workers in both groups.

The time use survey is designed to collect information on hours worked and the time allocated to different activities. This allows us to assess whether the differences in performance documented above are due to differences in time allocation across tasks—namely, whether treatment health workers do more visits because they devote more time to that task. To collect information on the latter, health workers were given 50 beans and asked to allocate the beans in proportion to the time devoted to each activity within each task. Besides household visits, community meetings, and time at the health post, we allow for two further activities: traveling and meeting with supervisors. For each activity, we calculate the share of time devoted to each activity by dividing the number of beans allocated to that activity by the total number of beans allocated to all activities. The share of time allocated to these five activities is 0.32, 0.22, 0.16, 0.22 and 0.09, respectively. We then estimate a system of equations for hours worked and share of time devoted to each task, omitting traveling. Table A.4 reports our findings.

Column 1 shows that the average health worker reports working 43 hours per week in the typical week and that there is no difference in reported working hours by treatment. This suggests that health workers in the control group do not compensate for visiting fewer households by devoting more hours to other, possibly informal, tasks. It also provides further assurance that health workers in the career treatment do not have differential incentives to overstate their contribution, as self-reported hours are unverifiable and hence easy to “game.”

Columns 2-5 show that health workers in the two groups allocate their time in a similar manner; thus, observed performance differences are not driven by differences in time allocation. Two potentially complementary explanations are possible. First, treatment health workers might work more effective hours—e.g., by taking shorter breaks over the 43 weekly hours. Second, treatment health workers might be more efficient at their jobs. These effects might be strengthened by peer externalities because each health worker works alongside another health worker hired through the same treatment. Thus, health workers in the treatment group are more likely to have a highly productive peer than health workers in the control group. Peer effects might be driven by imitation, social comparison, or a perception that the other health worker competes for the same promotion.

³⁸To implement this survey, we took advantage of a refresher course organized by the Government in the health worker school in Ndola. Of the 307 health workers, 298 (97%, equally split by treatment group) came to training and took part in the survey.

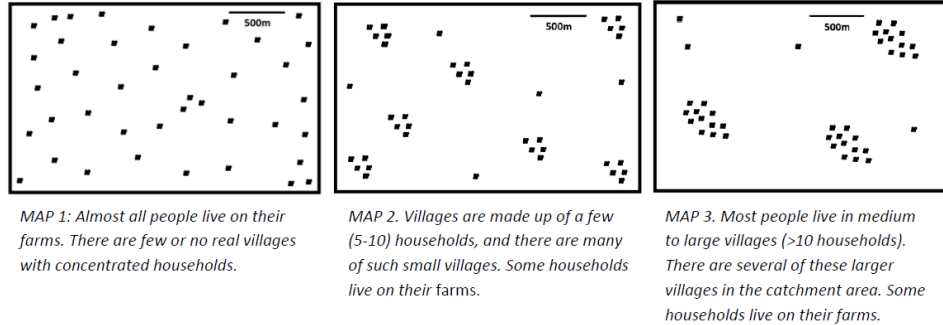
Finally, Table A.5 tests whether health workers in the two groups allocate their time differently within each activity, namely whether they have different work “styles.” Panel A shows that health workers in the treatment group devote more time to counseling, inspections, and visiting sick household members, but, taken one-by-one, these differences are small and not precisely estimated. Health workers in the treatment group devote 1.6% less time to filling in forms and receipts and submitting SMSs, but the difference is not precisely estimated at conventional levels. Because the quality of reports is the same, this implies that career health workers are more productive at this task. Panel B shows a similar pattern for time allocation during work at the health post: collecting data and filling in reports is an important component of the job, which takes 23% of the health workers’ time in the control group, but only 18% in the career treatment. As with household visits, there is no evidence that health workers in the career treatment collect less data at the health post level or that these data are of worse quality. Health workers in the two groups are equally likely to submit HMIS reports in a given month, and these are equally accurate. Thus, the evidence suggests that health workers in the career treatment are more productive, and this frees time for other tasks.

D Data Appendix

In this section, we describe each of the variables used in our analysis, including its source and unit of measurement. We collect data at each stage of the program: application, selection, training, and performance in the field. A description of each source, including the sample, can be found in Section E.

Eligible population and catchment area characteristics

- *Number of staff in health post* (source: district health officials survey, by phone) - Total number of nurses, environmental health technicians, and clinical officers assigned to the health post, as reported by district health officials we surveyed by phone.
- *Geographical distribution of households in catchment area* (source: health worker survey, in person, at refresher training) - Health workers were shown stylized maps and asked to choose the one that most closely resembled the catchment area of their health post. Questions were asked to each health worker individually so that two health workers from the same health post could give different answers. For the 5 out of 161 cases in which the two health workers gave different answers, we used information provided by supervisors to break the tie.



- *Poor cell network coverage* (source: attempted phone calls) - We attempted to call all health workers after deployment. We made daily calls for 118 consecutive days. The health post was classified as having poor coverage if we did not manage to reach either of its two health workers during this period.

Experiment Validation

- *Relative weight variables* (source: health worker survey, in person, at training) - These were derived from survey questions that asked the trainees to allocate 50 beans between different potential reasons for applying to the health worker position: “good future career,” “allows me to serve the community,” “earns respect and high status in the community,” “pays well,” “interesting job,” “allows me to acquire useful skills,” and “offers stable income.”
- *Expects to be employed in Ministry of Health in 5-10 years* (source: interviewee questionnaire, in person, at interview) - Circled any combination of being a “community health worker,” “nurse,” “environmental health technician,” “clinical officer,” or “doctor” in response to the question, “When you envision yourself in 5-10 years’ time, what do you envision yourself doing?”

Performance in Service Delivery

Household Visits

Source: SMS Receipts

Each variable is constructed from SMS messages sent by CHAs in which each message corresponds to a single household visit. The household visit receipt is reproduced below.

- *Unique households visited*
- *Number of visits per household*
- *Average visit duration, in minutes*

MINISTRY OF HEALTH
HOUSEHOLD VISIT RECEIPT

CHA ID:

1 START TIME: : :

END TIME: : :

DATE: / / 20

2 Client's Name

Client's Village Household ID

Client's Phone Number (if available)

3 I, the Client, certify that this receipt is truthful and accurate. CLIENT'S SIGNATURE

Source: HMIS (monthly reports)

Each reported variable is the sum of each indicator's monthly values from September 2012 to January 2014.

- *Number of households visited*
- *Number of women and children visited per household visit*
- *Number of patients seen at health post*
- *Number of community mobilization meetings*

Time Use

Source: health worker survey, in person, at refresher training

- *Number of hours worked in a typical week* - Health workers were asked "In a typical week, how many total hours do you spend doing health worker work? Please count work that you do at the health post and in the village, including moving from household to household."
- *Frequency of out-of-hours calls in a typical week* - Health workers were asked "In a typical week, how often do you have to leave your house at night and do CHA work due to emergencies like pregnancies or accidents?" Possible responses were "5-7 days per week," "3-4 days per week," "1-2 days per week," "2-3 times per month," "Once per month," "Sometimes, but less than once per month," and "Never."
- *Share of time allocated to* - To obtain time allocations, health workers were asked to allocate 50 beans between different activities. The instructions were as follows:

Please use the beans to show how much time you spend doing each activity. If you spend more time in an activity, you should place more beans on the card. If you never do an activity, you

should place no beans on the card. Place the beans any way you would like. For instance, you can place all beans on one card, or 0 beans on any card.

Household visits - Now I would like you to think about household visits specifically. Here are some cards that list different activities you may do during household visits.

- greeting household members*
- assessing and referring sick household members*
- reviewing and discussing the household's health profile and goals*
- asking questions about household health behaviors and knowledge*
- providing health counseling*
- doing household inspections (waste disposal, latrines, etc.)*
- documentation (filling registers/books and sending visit receipts via SMS)*

Health Post - Now here are some cards that list different activities you may do at the HEALTH POST OR RURAL HEALTH center.

- seeing sick patients at the OPD*
- dispensing medications from the pharmacy*
- helping with ANC visits*
- cleaning and maintaining the facility*
- assisting with deliveries and other procedures when needed*
- documentation (filling registers/books and sending monthly reports through HMIS)*

In the Community - Now here are some cards that list different activities you may do as a health worker.

- campaigns for polio, measles, child health, and other health issues*
- health talks and other community mobilization activities*
- school health talks and other school activities*
- meeting with NHC and volunteer CHWs for planning*

Health workers' observable traits


Skills

- *Average test score at training [0-100]* - Average score in 11 tests on basic medical practices taken during the training program.
- *O-levels total exam score* (source: MOH application files) - This variable is constructed as the sum of inverted O-levels scores (1=9, 2=8, and so on) from all subjects in which the applicant wrote the exam, so that larger values correspond to better performance.

- *O-levels passed in biology and other natural sciences* (source: MOH application files) - Includes biology, chemistry, physics, science, and agricultural science.

Preferences

- *Envisions remaining in the community in 5-10 years* (source: interviewee questionnaire, written, at interview) - Answered affirmatively in response to the question, “Do you see yourself in the community in 5-10 years?”
- *Perceives community interests and self-interest as overlapping* (source: interviewee questionnaire, written, at interview) - Based on the “Adapted Inclusion of Others in Self (IOS) scale” (Aron and Others 2004), which measures the extent to which individuals perceive community- and self-interest as overlapping. The Inclusion of Other in the Self scale was originally designed by Dr. Art Aron and colleagues (Aron et al. 1992) as a measure of self-other inclusion and relationship closeness. The Continuous IOS makes use of the basic design of the original IOS,³⁹ but allows for (a) the measure to be embedded within a web-based questionnaire, (b) the output values to be continuously scaled, and (c) modifications in the appearance and behavior of the measure. IOS has been validated across a wide variety of contexts, and adapted versions are found to be strongly correlated with environmental behavior (Schultz 2012) and connectedness to the community (Mashek et al. 2007). The measure is coded as 0-1, where 1 implies highest overlap. Applicants are asked to choose between sets of pictures, each showing two circles (labeled “self” and “community”) with varying degrees of overlap, from non-overlapping to almost completely overlapping. This variable equals 1 if the respondent chooses the almost completely overlapping picture (D), 0 otherwise.



- *Aims to be a higher-rank health professional in 5-10 years* (source: interviewee questionnaire, written, at interview) - Circled any combination of being an “environmental health technician,” “clinical officer,” or “doctor” in response to the question, “When you envision yourself in 5-10 years’ time, what do you envision yourself doing?”

Other traits

³⁹http://www.haverford.edu/psych/ble/continuous_ios/originalios.html

- *Social connections* (source: interviewee questionnaire, written, at interview) - For connectedness to a village leader, circled “political leader” or “village committee member,” and for connectedness to health center staff, circled “formally trained health worker (e.g. nurse aid, nurse, clinical officer)”, in response to the question, “Are any of your relatives or members of your household in the following positions?”

E Data Sources

- **Source: Application** (sample: all applicants) - Applications were submitted in August-September 2010. The initial application stage was comprised of the initial application form, which includes fields for gender, date of birth, village of residence, and educational qualifications. The application form also included a question asking through what means the applicant first learned of the health worker job opportunity: recruitment poster, facility health worker, community health worker, government official, word-of-mouth, or “other.”
- **Source: Interviewee Questionnaire** (sample: subset of applicants called for an interview) - Written questionnaires were completed by interview candidates at the time of their interviews in September-October 2010. The questionnaire (written in English) included a series of questions about the interviewee’s demographic background, community health experience, social capital, and work preferences and motivations. Notably, we included a measure employed by social psychologists, “Inclusion of Others in Self” (Aron and Others 2004) to measure connection with the community. The questionnaire stated that the answers would not be used for selection purposes but rather as part of a research project, although we cannot rule out that panelists could have seen the questionnaire or referred to it when making their decisions.
- **Source: Ranking Sheet** (sample: members of interview panels) - Ranking sheets were filled and collected in September-October 2010. Each panel consisted of five members: the district health officer, a representative from the health center, and three neighborhood health committee members. Once all interviews were completed, every member of the selection panel completed a private and individual ranking sheet by ranking their top ten candidates. This ranking exercise occurred *before* panel members formally deliberated and discussed the candidates. After interviewing all candidates and deliberating, interview panels were requested to complete and submit a consensus-based “Selection Panel Report” that included fields for the two nominated candidates as well as three alternates.
- **Source: Baseline Survey** (sample: all trainees) - The baseline survey was conducted in June 2011 and consisted of two main components:
 1. Questionnaire - Conducted one-on-one by an enumerator and collected information on the trainees’ socioeconomic background and livelihood, motivations and preferences, and expectations of the program.

2. Self-assessment - A three-hour exam with multiple choice questions to determine the baseline knowledge of health matters that each student had prior to the training.

- **Source: Catchment Area Survey** (sample: all deployed CHWs and supervisors) - Just prior to graduation in July 2012, all CHWs and supervisors were given a short survey that asked about characteristics of their health posts, including population density, rainy-season information, and general community health measures.
- **Source: Time Use Survey** (sample: all deployed CHWs) - This survey was conducted in April-May 2013 in Ndola, Zambia. The respondents were pilot health workers who reported to Ndola for a supplemental in-service training to introduce new tasks as part of a revised health worker scope of work. The survey was administered by Innovations for Poverty Action, in partnership with the Ministry of Health, the Health Worker Training School, and the Clinton Health Access Initiative.
- **Source: SMSs** (sample: all deployed health workers) - All health workers carry with them receipt books for each visit, which require the signature of the client visited. The information on these receipts—consisting of the data, time, and duration of the visit, as well as the client’s phone number—is then SMS’ed in real time to the Ministry of Health and our central data-processing facility.

F District Instruction Appendix

The health worker program was introduced differently to health centers depending on the treatment group. In each district, the district health official was given a package that contained a script, a memo from the Permanent Secretary, and detailed instructions about the health worker recruitment process. In addition, district health officials received “health center packages” for each participating health center in the district, which contained a set of posters and application forms and instructions for the health center representative on how to post posters and collect applications. The district health officials were to visit each health center and meet with the staff and neighborhood health committee members to introduce the program and distribute the health center packages, using the script provided to them in their packages. The script was only provided to the district health officials, and was addressed directly to them. It is unlikely that the applicants or health center staff were able to read this script themselves.

The following script was given to district health officials in the treatment group:

To Health center and Neighborhood Health Committee: I would like to let you know about a new government program to strengthen the country’s health workforce. Applications are currently being accepted for a new Community Health Worker position. This is an opportunity for qualified Zambians to obtain employment and to advance their health careers. Opportunities for training to advance to positions such as Nurse and

Clinical Officer may be available in the future. Successful applicants will receive 1 year of training, both theoretical and practical. All training costs, including transportation, meals and accommodation during the one-year training program, will be covered by the Ministry of Health. Please encourage all qualified persons to apply so that they can benefit from this promising career opportunity.

The district health officials in the control group received the following script:

To Health center and Neighborhood Health Committee: I would like to let you know about a new government program to improve health care services in your community. Applications are currently being accepted for a new Community Health Worker position. This is an opportunity for local community members to become trained and serve the health needs of their community. The new CHWs will work at the Health Post and community level in coordination with an affiliated Health center. Successful applicants will receive 1 year of training, both theoretical and practical. All training costs, including transportation, meals and accommodation during the one-year training program, will be covered by the Ministry of Health. Please encourage all qualified persons to apply so that they can benefit from this promising community service opportunity.

Table A.4: The effect of career opportunities on time use

dependent variable	Share of time spent in:				
	Hours worked	HH visits	Health Post	Community meetings	Meeting with supervisor
	(1)	(2)	(3)	(4)	(5)
Treatment		.007 (.014)	-.021 (.013)	.011 (.011)	-.001 (.008)
Area characteristics		yes	yes	yes	yes
Mean of dependent variable in control	42.8	.312	.171	.213	.085
Adjusted R-squared	.031	.055	.081	.031	.063
N	298	298	298	298	298
EDF p-value	.609	.560	.085	.297	.892
RI p-value	.583	.622	.140	.259	.887

Notes: Column 1: OLS Estimates, standard errors clustered at the district level. Columns 2-5: SURE Estimates, standard errors clustered at the district level bootstrapped with 1500 replications. EDF p-value refers to the p-value from a null hypothesis that the treatment effect is zero with OLS estimates, using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)). Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. Data source is the Time Use Survey that was administered in May 2013 during a refresher training program. Hours worked is defined as the number of hours worked in a typical week as reported by the CHAs. To measure the "Share of time spent in," CHAs were given 50 beans and asked to allocate them on cards listing the different activities listed above plus travel. The cards were scattered on a table in no particular order. All regressions include the stratification variables (province dummies and share of high school graduates in the district). Area characteristics include: number of staff in the health post, geographical distribution of households in the catchment area, and an indicator variable that equals 1 if the CHA reports to have good cell network coverage most of the time or all the time.

Table A.5: The effect of career opportunities on time allocation

Panel A: Time allocation during household visits

share of time allocated to:	counseling	inspections	filling in receipts and forms	asking questions about health behaviors and knowledge	discussing health profile and goals	visiting sick household members
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	.006 (.012)	.007 (.015)	-.016 (.010)	-.011 (.009)	-.003 (.012)	.010 (.009)
Mean of dependent variable in control	0.207	0.196	0.146	0.137	0.122	0.100
Area characteristics	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	.030	.041	.049	.026	.014	.027
N	292	292	292	292	292	292
EDF p-value	.580	.627	.108	.197	.759	.268
RI p-value	.594	.633	.068	.198	.795	.314

Panel B: Time allocation during work at the health post

share of time allocated to:	seeing sick patients	filling in forms	dispensing medications	helping with antenatal care visits	cleaning and maintaining the health post
	(1)	(2)	(3)	(4)	(5)
Treatment	-.002 (.011)	-.050 (.020)	.006 (.013)	.019 (.019)	.019 (.013)
Mean of dependent variable in control	0.262	0.228	0.207	0.160	0.104
Area characteristics	Yes	Yes	Yes	Yes	Yes
R-squared	.051	.104	.091	.094	.133
N	271	271	271	271	271
EDF p-value	.827	.009	.634	.293	.112
RI p-value	.783	.008	.636	.231	.133

Notes: System estimates (SURE), bootstrapped standard errors clustered at the district level in parentheses. All regressions include the stratification variables (province dummies and share of high school graduates in the district). EDF p-value refers to the p-value from a null hypothesis that the treatment effect is zero with OLS estimates, using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)). Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. All 298 participants in the refresher training program were given 50 beans and asked to allocate the beans to show how much time they spent doing each activity within each task. They were instructed to place more beans on a card if they spent more time on an activity, to place no beans if they never do an activity, and to place the beans any way they would like, including placing all beans on one card, or 0 beans on any card. Panel A activities are: greeting household members, assessing and referring sick household members, reviewing and discussing the household's health profile and goals, asking questions about health behaviors and knowledge, providing health education and counseling, doing household inspections (waste disposal, latrines, etc.), and documentation (filling registers/books and sending SMS visits). The omitted category in Panel A is "greetings." The sample in Panel A covers the 292 out of 298 CHAs who reported spending time doing visits. Panel B activities are: seeing sick patients in the health post, dispensing medications from the pharmacy, helping with ANC visits, cleaning and maintaining the facility, assisting with deliveries and other procedures when needed, and documentation (filling registers/books and sending monthly reports through DHIS2). The omitted category in Panel B is "assisting with deliveries." The sample in Panel B covers the 271 out of 298 CHAs who reported spending time at the health post. Area characteristics include: number of staff in the health post, geographical distribution of households in the catchment area, and an indicator variable that equals 1 if the CHA reports to have good cell network coverage most of the time or all the time.

Table A.6: The effect of career opportunities on facility utilization – robustness checks

Panel A. Placebo test						
dependent variable: total over each quarter 2011:1-2014:2						
	women giving birth at the health center	postnatal (0-6 weeks) visits	children under 5 visited	children under 5 weighed	children under 1 receiving BCG vaccinations	children under 1 receiving polio vaccinations
	(1)	(2)	(3)	(4)	(5)	(6)
	Yes	Yes	Yes	Yes	Yes	Yes
Treatment	-3.750 (12.18)	-12.17 (9.392)	-6.500 (179.5)	-6.454 (161.1)	12.03 (13.54)	-3.404 (10.59)
After	0.864 (3.887)	15.43 (4.433)	92.22 (78.05)	153.5 (74.10)	2.659 (5.177)	3.843 (3.941)
Treatment*After	12.29 (5.349)	8.409 (9.553)	363.8 (116.9)	335.3 (137.4)	7.949 (10.06)	11.77 (5.232)
Placebo After	7.267 (2.347)	0.138 (4.528)	-64.76 (81.76)	-94.40 (66.74)	-8.341 (4.548)	-10.45 (5.173)
Treatment*Placebo After	2.944 (5.834)	-1.329 (5.206)	-111.0 (141.6)	-123.8 (137.9)	-1.828 (7.704)	5.992 (7.841)
Area characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dependent variable in control in year 1	45.3	55.4	1285.7	1236.1	83.0	72.4
Adjusted R-squared	0.350	0.214	0.254	0.255	0.152	0.117
Number of facilities	89	119	123	123	120	120
Number of observations	1301	1543	1618	1610	1518	1535
EDF p-value of Treatment*Placebo After	0.622	0.803	0.444	0.381	0.816	0.456
RI p-value of Treatment*Placebo After	0.599	0.808	0.418	0.381	0.810	0.474
						0.728
Panel B. Health post fixed effects.						
dependent variable: total over each quarter 2011:1-2014:2						
	women giving birth at the health center	postnatal (0-6 weeks) visits	children under 5 visited	children under 5 weighed	children under 1 receiving BCG vaccinations	children under 1 receiving polio vaccinations
	(1)	(2)	(3)	(4)	(5)	(6)
	Yes	Yes	Yes	Yes	Yes	Yes
After	4.975 (4.167)	15.46 (5.217)	63.77 (66.27)	106.9 (65.97)	-1.094 (4.790)	-1.273 (3.887)
Treatment*After	13.02 (6.411)	8.665 (9.840)	306.3 (107.4)	278.4 (119.1)	8.409 (8.842)	15.38 (5.321)
Area characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dependent variable in control in year 1	45.3	55.4	1285.7	1236.1	83.0	72.4
Adjusted R-squared	0.822	0.664	0.618	0.591	0.496	0.565
Number of facilities	89	119	123	123	120	121
Number of observations	1301	1543	1618	1610	1518	1535
EDF p-value of Treatment*After	0.047	0.374	0.006	0.022	0.337	0.006
RI p-value of Treatment*After	0.024	0.412	0.017	0.043	0.352	0.003
						0.140

Notes: OLS Estimates, standard errors clustered at the district level. Treatment=1 if the health worker is recruited in a district where career opportunities were made salient. Data source is the Health Management and Information System (HMIS) available monthly from January 2011 until June 2014. Health center and health post staff are required to submit monthly reports that summarize their activities at the health post/community level. These are aggregated at the quarter level in the regressions. The variable in Column 1 is defined at the health center level because health centers are equipped for child births and health posts are not. The variables in columns 2-7 are defined at the health post level if this reports data, at the health center otherwise. After=1 after September 2012 (from 2012:4 onwards), when CHAs started working. Placebo After=1 after September 2011, halfway through the period before the CHAs started working. All regressions include the stratification variables (province dummies and share of high school graduates in the district). Area characteristics include: number of staff in the health post, geographical distribution of households in the catchment area, and an indicator variable that equals 1 if the CHA reports to have good cell network coverage most of the time or all the time. EDF p-value refers to the p-value from a null hypothesis that the treatment interaction, as specified in the tables, is zero (in the same regression), using the Young (2016) effective degrees of freedom correction. RI p-value refers to the equivalent p-value using a Randomization Inference procedure (specifically, the randomization-t p-value from Young (2018)).