Universitat Pompeu Fabra Barcelona

ESSAYS IN DEVELOPMENT ECONOMICS

Philipp Kastrau

DIRECTOR DE LA TESI:

Gianmarco León-Ciliotta
Department of Economics and Business

TESI DOCTORAL UPF / 2022

ACKNOWLEDGMENTS

I am deeply grateful to my advisor Gianmarco León and my co-author Erika Deserranno for providing me the opportunity to conduct field research in Sierra Leone and all their guidance, encouragement, patience, and understanding throughout my PhD studies. I am also thankful to Niccoló Meriggi for his amazing support and advice in Sierra Leone and to Luz Azlor for having been the best imaginable field coordinator. I am indebted to the whole team at NestBuilders, especially Alieu Mansaray, for being both so welcoming and hard-working - without your work in the field none of the chapters in this dissertation would have been possible to write and you made my time in the field a lot of fun. I would also like to thank my co-author Stefano Caria for his invaluable contributions and Andre Cazor Katz, Margaux Jutant, Mustapha Adedolapo Kokumo, Raquel Lorenzo Vidal, Theo Durandard, and Malachy Gavan for their outstanding research assistance.

I was fortunate to learn from the amazing faculty and students at Universitat Pompeu Fabra and UC Berkeley. I am especially grateful for the helpful suggestions and comments from Alessandro Tarozzi and Giacomo Ponzetto and from the participants at the Applied Lunch Seminar. I would also like to thank Fred Finan for hosting me at UCB. Marta Araque and Laura Agusti were incredibly helpful, friendly, and supportive during all my time at UPF. I am very grateful for the people who made the PhD experience much more fun, especially Sampreet, Laurenz, Willy, and Viraj. None of my research projects would have been possible without our collaborators at the Ministry of Health and Sanitation and the IRC in Sierra Leone and without the generous financial support by UK aid from the UK government (through the Economic Development and Institutions initiative), the Rockefeller Foundation (through the International Growth Center), J-PAL Governance Initiative, and the Weiss Family Fund.

Finally, I want to thank the most important people: Thank you to my parents

Ursula and Harald for whom my education was always a priority and without whom I would have never been able to make it to grad school. Dad, may you rest in peace. Thank you to those who supported me so much over the course of my PhD, especially BAK, Benji, Jonny, Katrina, Malte, Mel, Mgone, Moe, and Sophia.

ABSTRACT

This thesis consists of three independent chapters which all utilize primary data from Sierra Leone. The first chapter uses a field experiment to show that financial incentives maximize output in an hierarchical organization when they are equally shared between frontline health workers and their supervisors. The second chapter studies promotion incentives in the public sector using a field experiment and finds that meritocratic promotions lead to higher productivity for workers who expect a steep pay increase and those who are highly ranked in terms of performance. When promotions are not meritocratic, increasing the pay gradient instead reduces worker productivity through negative morale effects. The third chapter uses an instrumental variables approach and finds that exposure to the 2014 Ebola outbreak in Sierra Leonean villages leads to significant increases in risk perceptions regarding Covid-19 and in the likelihood of public distributions of face masks, but a decrease in health access.

RESUMEN

Esta tesis consta de tres capítulos independientes que utilizan datos primarios de Sierra Leona. El primer capítulo usa un experimento de campo para mostrar que los incentivos financieros maximizan la producción en una organización jerárquica cuando se comparten por igual entre los trabajadores de salud de primera línea y sus supervisores. El segundo capítulo estudia los incentivos de promoción en el sector público utilizando un experimento de campo y encuentra que las promociones meritocráticas conducen a una mayor productividad para los trabajadores que esperan un fuerte aumento salarial y aquellos que están altamente clasificados en términos de desempeño. Cuando las promociones no son meritocráticas, aumentar el gradiente salarial reduce la productividad del trabajador a través de efectos morales negativos. El tercer capítulo utiliza un enfoque de variables instrumentales y encuentra que la exposición al brote de ébola de 2014 en las aldeas de Sierra Leona conduce a aumentos significativos en las percepciones de riesgo con respecto a la Covid-19 y en la probabilidad de distribución pública de máscaras faciales, pero a una disminución en el acceso a la salud.

PREFACE

This thesis consists of three self-contained chapters. While all three chapters make independent scientific contributions, they all have in common that they use primary data from health workers and households in Sierra Leone and address questions broadly in the field of development economics.

The first chapter (joint with Erika Deserranno and Gianmarco León-Ciliotta) tackles a classic problem faced by organizations, namely how to distribute incentives among their different layers. By means of a field experiment with a large public-health organization in Sierra Leone, we show that financial incentives maximize output when they are equally shared between frontline health workers and their supervisors. The impact of this intervention on completed health visits is 61% larger than the impact of incentive schemes that target exclusively the worker or the supervisor. Also, the shared incentives uniquely improve overall health service provision and health outcomes. We use these experimental results to structurally estimate a model of service provision and find that shared incentives are effective because worker and supervisor effort are strong strategic complements, and because side payments across layers are limited. Through the use of counterfactual model experiments, we highlight the importance of effort complementarities across the different layers of an organization for optimal policy design.

The second chapter (joint with Stefano Caria, Erika Deserranno, and Gianmarco León-Ciliotta) studies promotion incentives in the public sector by means of a field experiment with the Ministry of Health in Sierra Leone. The experiment creates exogenous variation in meritocracy by linking promotions to performance for the lowest tier of health workers and in perceived pay progression by revealing to them the salary of higher-tier workers. We find that meritocratic promotions lead to higher productivity for workers who expect a steep pay increase and those who are highly ranked in terms of performance. When promotions are not meritocratic, increasing the pay gradient instead reduces worker productivity through negative morale effects. The findings highlight the importance of taking into account the interactions between different tools of personnel policy.

The third chapter provides micro evidence for one mechanism behind the dramatically different political responses to the Covid-19 pandemic, namely how an increase in the perceived risk of Covid-19 among individuals stemming from past exposure to similar health crises generates citizen demand for containment measures. Exploiting exogenous variation in exposure to the 2014 Ebola outbreak across villages in Sierra Leone, I find that past exposure leads to significant increases in risk perception regarding Covid-19 and trust in health professionals among households. I then show that this also translates into Ebola-affected villages being significantly more likely to have organized the public distribution of face masks. However, the increased caution comes at the cost of reduced health access as households in Ebola-affected villages are more likely to avoid health clinics during the pandemic out of fear of contracting Covid-19.

Contents

List of Figures

List of Tables

Chapter 1

The Allocation of Incentives in Multi-Layered Organizations

Joint with Stefano Caria, Erika Deserranno, and Gianmarco León-Ciliotta

1.1 Introduction

Reaping the benefits of worker effort complementarities is a key reason for the existence of organizations (Alchian and Demsetz, 1972). These complementarities are particularly salient in vertical organizations, where workers and managers jointly contribute to production (Wilson, 1989; Garicano, 2000). Without good management, workers are often ineffective, and similarly, the efforts of managers can only pay off if workers are motivated to do their job. How to allocate incentives to maximize organizational performance in the presence of these complementaries is, however, not fully understood. At one extreme, if agents redistribute financial rewards among themselves, the initial allocation of incentives is inconsequential and the organization should only worry about the total level of incentives provided. In contrast, if transfers are constrained, the precise allocation of incentives across the different layers of the organization becomes crucial. Surprisingly, empirical evidence on the optimal allocation of incentives in vertical organizations and on its structural determinants remains limited.

In this paper, we show that the allocation of financial incentives within a large public-health organization has substantial impacts on the provision of health care services in poor communities across Sierra Leone. In particular, we document experimentally that equally sharing a piece-rate incentive between a worker and a supervisor generates an increase in health visits - the main output of the organization - that is 61% larger than the gain in visits achieved when the entire piece rate is offered either to the worker or to the supervisor. Through a structural model and novel empirical evidence, we then shed light on the key factors that underpin these results - (i) the strong complementarity in worker and supervisor effort, and (ii) the limited redistribution of the incentive - and explore their quantitative implications for optimal policy design.

The program we study is a large community-based health initiative designed to improve health-service provision in Sierra Leone, with a focus on pre- and post-natal care. Community-health services play a crucial role in reducing the burden of common diseases and child mortality (Nyqvist et al., 2019; Deserranno et al., 2020). Yet, access to primary health care is still a major issue in rural areas of developing countries and the expansion of community health worker programs is an important part of the global strategy to ensure universal health-care access (Campbell et al., 2013). Finding ways of optimizing the performance of community health workers is hence a first-order policy priority.

We introduce a new piece-rate scheme that pays 2,000 Sierra Leone Leones (SLL) per completed health visit, and create random variation in its recipient in a sample of 372 health units across the country. Each unit is composed of an average of 8 health workers, who directly carry out health visits, and one supervisor, who provides training, support and advice. The incentive is either paid (i) only to the health worker who carried out the visit, (ii) only to the supervisor of this worker, or (iii) is shared equally between the worker and the supervisor. In all these treatments, the organization relies on workers' reports to determine the amount of incentive to be paid. Importantly, we also collect an independent measure of completed health visits by interviewing a random sample of households in each village. This independent measure does not suffer from self-reporting bias and is the main outcome variable in the study.

To guide our empirical analysis, we propose a simple model of service provision that illustrates the trade offs involved in the choice of how to allocate the incentive. In the model, a supervisor and a worker interact over two time periods. In the first period, the supervisor chooses how much effort to invest in training and advising the worker, and offers her a side payment conditional on the amount of services delivered at the end of the game. In the second period, the worker chooses how much effort to exert to provide services. A key intuition is that the optimal share of the incentive to be offered to each agent depends on: (i) the strategic complementarity of worker and supervisor effort, and (ii) the extent to which side payments offset the initial allocation of the incentive. In our setting, strategic complementarities are likely to be strong as supervisors play a key "enabling" role: they raise the health workers' ability to conduct household visits by training and advising them, providing the necessary skills to perform their tasks, and helping them build trust in the community. Further, workers and supervisors' ability to offer side payments to each other is constrained by different contractual frictions.[1]

[1] Contractual frictions can derive from the limited observability and predictability of worker effort (Duflo et al., 2012), the difficulty of making binding commitments (Casaburi and Macchiavello, 2019), social norms on the appropriateness of side payments or institutional rules that limit managerial autonomy (Banerjee et al., 2020; Bandiera et al., 2021), or flypaper effects whereby payments are expected to stay in the layer of the organization to which they are originally allocated (Hines and Thaler, 1995).

Our model highlights how, in a setting where these features are present, sharing the piece rate is an optimal policy.

In the first part of the paper, we present the causal effects of our treatments on the number of visits carried out by the health workers, as reported by the households. Our central empirical finding is that the shared incentives treatment maximizes the number of completed health visits. Workers in the control group without performance-based incentives (status quo) carried out 5.3 visits per household in the six months prior to our endline survey. This number significantly increases to 7.4 visits (a 40% increase over the control condition) when the incentive is offered either only to the worker or only to the supervisor, and to 8.7 visits (a 63% increase over the control condition) when the incentive is shared between the worker and supervisor. Overall, the shared incentives generate an increase in health visits that is 61% larger than the increase caused by either of the one-sided incentives treatments. We rule out concerns related to quantity-quality trade-offs. The observed increase in the quantity of household visits provided in the shared incentives treatment is not compensated by a reduction in visit length, nor by changes in the targeting of poor and deserving households. Moreover, the share of households who report trusting the health worker is the highest in the shared incentives treatment. This is important because trust in health service providers is known to be one of the main determinants of the demand for health services (Alsan, 2015; Lowes and Montero, 2021; Martinez-Bravo and Stegmann, 2022; León-Ciliotta et al., 2022). We also find that the health worker's knowledge about how to adequately provide health services to the community is the highest in the shared incentives treatment.

The large positive impact of the shared incentives treatment on household visits translates into better access to pre- and post-natal care and lower disease incidence. Pregnant or expecting women are more likely to report having received at least four pre-natal visits from any provider and having delivered in a health facility (rather than at home) in the shared incentives treatment than in the one-sided incentives treatments or the control. Households also report fewer instances of fever among children below the age of five, and are more aware of how to prevent diseases.

Importantly, shared incentives outperform both one-sided incentives also in terms of cost-effectiveness. The incentive is only paid when a visit is *reported* by the health worker. Thanks to a system of extensive back-checks, we find that over-reporting is minimal. Instead, health visits are often under-reported, plausibly due to high reporting costs, which we discuss in Section 1.2. Crucially, under-reporting decreases with the share of the incentive offered to the worker. This makes shared incentives particularly cost effective: we find that each 2,000 SLL spent on the program generate 16.1 extra visits in the shared incentives treatment, 9.6 extra visits in the supervisor incentives treatment, and 6.5 extra visits in the worker incentives treatment.

In the second part of the paper, we study the mechanisms explaining the large boost in output generated by shared incentives. In line with our model, we show that both effort complementarity and limited side payments play an important role. Three key results point to the presence of large effort complementarities. First, shared incentives generate the same increase in supervisor effort as supervisor incentives. This could seem surprising, since the *direct* incentive offered to the supervisor is lower in the shared incentives treatment. However, as predicted by our model, shared incentives compensate for this by providing a strong boost to worker effort, which raises the return to supervisor effort and hence *indirectly* incentivizes the supervisor to raise effort. Second, shared incentives generate a larger increase in visits and supervisor effort when effort complementarity is plausibly higher due to the low level of experience of the worker. Third, we carry out a formal mediation analysis which shows that the boost in visits due to worker effort increases with the level of supervisor effort.

Next, we turn to the role of side payments. We leverage data on inter-personal transfers to show that, on average, net transfers from the supervisor to the worker are positive, but very small: less than 10% of the overall incentive payment of the average supervisor. Why are transfers limited? One possibility is that the poor observability of worker effort makes contracting hard. In line with this, we show that supervisors who plausibly cannot observe worker output accurately make lower transfers. Additionally, as we argue below, in many cases the worker may have a higher stake in the production of output than the supervisor. Transfers from workers to supervisors, however, are almost never observed in the data, suggesting that frictions may also prevent bottom-up transfers.

We present several pieces of evidence which are inconsistent with two alternative explanations of our results. First, one-sided incentives treatments could be ineffective due to a negative morale effect arising from pay inequality (Breza et al., 2018). Our experimental design minimizes this concern, as workers are not informed of the presence of supervisor incentives (if any) and only few seem to learn about it from the supervisors. Moreover, we find no evidence suggesting that workers in the supervisor incentives treatment are less satisfied with their payment or their job compared to the control group. Second, we consider the possibility of strong non-linearities in the utility, cost or production functions. Shared incentives could be highly effective in the absence of effort complementarities if, for both agents, the marginal utility generated by the incentive declines rapidly after 1,000 SLL (the size of the incentives paid in the shared incentives scheme) or the marginal cost (product) of effort increases (decreases) steeply after the level of effort generated by a 1,000 SLL incentive. However, when we analyze non-parametrically the relationship between treatment effects and proxies of utility (wealth) and costs (distance between the worker and her patients, or between the supervisor and the worker), we do not observe any sharp non-linearities. Similarly, we do not find evidence of sharp non-linearities in the relationship between supervisor effort and visits completed.

In the final part of the paper, we leverage the experimental variation to structurally estimate our model of service provision and perform different counterfactual simulations. For the estimation, we use moments capturing household visits and supervisor effort in the three treatment conditions and in the control group. The model is able to match these moments with precision. The estimated model parameters confirm that our results are driven by a strong complementarity of effort. In particular, we estimate that the marginal return to worker effort is up to 116% higher due to the complementarity with supervisor effort. Second, our calibrated contractual friction parameter implies that side transfers are 45% more expensive due to difficulties in contracting. Third, we find that, in the absence of the intervention, supervisors have weaker incentives to provide effort than workers. This underscores the importance of incentive schemes that ensure supervisors are adequately incentivized.

We derive three lessons on optimal policy based on the structural model. First, given the estimated parameters, we calculate that the optimal policy would offer 59% of the value of the incentive to the worker, and 41% to the supervisor. Second, we study how the optimal policy changes for different levels of effort complementarity. We find that the optimal allocation of the incentive is sensitive to the exact value of this parameter, which emphasizes the importance of re-calibrating the policy in new contexts. Third, the strong complementarity determines a large positive external effect of individual effort, which the agents fail to internalize. This makes interventions that tie incentives to joint output more effective than interventions that incentivize effort directly, even in settings where effort is perfectly observable. This result has broad implications for optimal pay structure in organizations where workers at different layers complement each other in the production of output.

This paper contributes to four strands of the literature. First, we show that the allocation of incentives in an organization with multiple tiers is highly consequential due to a combination of strong effort complementarities and a limited redistribution of incentives. The existing empirical literature has largely been unable to shed light on this important point since most studies to date have explored the effects of raising incentives in one layer of the organization, while holding incentives in the other layer fixed. These include papers focusing on the bottom layer - e.g., frontline workers (Glewwe et al., 2010; Muralidharan and Sundararaman, 2011; Duflo et al., 2012; Ashraf et al., 2014), sales associates (e.g., Lazear, 2000) - and papers focusing on the top layer — e.g., high-level public sector officials (Rasul and Rogger, 2018; Luo et al., 2019), private sector CEOs/managers (Bandiera et al., 2007; Bertrand, 2009; Frydman and Jenter, 2010) - with Behrman et al. (2015) as an exception.[2] Our results have two implications for organizations working in contexts similar to ours. First, agents engage in very limited fine-tuning of the allocation of incentives through transfers. Thus, there are large returns from picking the optimal allocation from the start. Second, when the interests of principals and agents are not aligned, the ability of one layer of the organization to distort behavior in other layers through transfers is likely to be limited.

Second, we provide evidence on the *productive* role of middle managers in hierarchical organizations. A long-standing literature focuses instead on the *monitoring*

[2]Behrman et al. (2015) evaluate the effectiveness of three alternative performance incentive schemes on mathematics tests scores in Mexican schools: (1) individual incentives for students only, (2) individual incentives for teachers only, and (3) individual and group incentives for students, teachers, and administrators. Program impact estimates reveal the largest average effects for (3). The paper cannot assess whether this is because of complementarities across layers or because of the different incentives structure (e.g., individual vs. group).

role of managers. This literature - which spans seminal theoretical contributions (e.g., Tirole, 1986, 1992) and a number of recent empirical papers (Cilliers et al., 2018; Bandiera et al., 2021; Dal Bó et al., 2021; Dodge et al., 2021) - studies how to optimally delegate authority and how to avoid harmful collusion between workers and supervisors, but it typically ignores the enabling (and thus productive) role of supervisors, and therefore remains silent on how effort complementarities can be best leveraged. In our experiment, we explicitly minimize the scope for collusion through frequent back-checks of worker reports. This enables us to shed light on how the top layer of the hierarchy enables the frontline layer to be productive, and on the implications of this complementarity for the design of incentives.

Third, we advance the literature on effort complementarities in organizations. Seminal theoretical work by Alchian and Demsetz (1972); Itoh (1991); Ray et al. (2007) has reflected on the implications of complementarities for incentive design. Empirically, a number of papers have demonstrated that in "horizontal" teams - composed of workers from the *same* layer of the organization - group incentives that reward joint (rather than individual) output are effective even if at the potential cost of increasing free-riding (Muralidharan and Sundararaman, 2011; Babcock et al., 2015; Friebel et al., 2017). Unlike our paper, this literature does not shed light on the optimal allocation of such incentives. This is partly because, in horizontal teams, offering anything but symmetric incentive schemes is often not a policy option due to, e.g., fairness concerns (Breza et al., 2018; Card et al., 2012) or other rigid contractual arrangements that prevent from offering different incentives to workers performing comparable tasks. In "vertical" teams, on the other hand, asymmetric incentives are more acceptable, since workers in the different layers of the organization have different responsibilities and different levels of experience. Further, vertical teams are likely to display different levels of strategic complementarities and contractual frictions compared to horizontal teams, since task heterogeneity is likely to be higher and team members are often not peers, but managers and subordinates.

Finally, this paper documents the presence of contractual frictions *within an organization* (Adhvaryu et al., 2020). Most of the literature has instead focused on contractual frictions *across organizations or firms* (Coase, 1937; Gibbons, 2005; Lafontaine and Slade, 2007; Lee et al., 2021). Our results highlight that frictions such as the limited observability of subordinates' effort limit the scope for Coasian bargaining within the firm, and make the allocation of incentives central to organizational performance.

The paper is structured as follows. Section 1.2 discusses the context and research design. Section 1.3 presents a simple model of service delivery with effort complementarity across layers and contractual frictions. Section 1.4 studies the effect of our incentives treatments on output and reporting, and their cost-effectiveness. Section 1.5 explores the mechanisms underlying our main output results. Section 1.6 presents the structure model and performs a number of relevant counterfactual policies. Section 1.7 concludes. The Appendix presents further results and discusses key aspects of research ethics (e.g., the AEA pre-registration and IRB).

1.2 Context and Research Design

1.2.1 The Community Health Program

Sierra Leone is one of the poorest countries in the world, with the third-highest maternal mortality rate and the fourth-highest child mortality rate in 2017 (World Health Organization, 2017). Such elevated mortality rates have been attributed to the slow post-civil war recovery, the 2014 Ebola outbreak, and a critical shortage of health workers, together with limited access to health facilities throughout the country (World Health Organization, 2016). In order to strengthen the provision of primary health care, Sierra Leone's Ministry ofHealth and Sanitation (MoHS) created a national Community Health Program in 2017. The program is organized around Peripheral Health Units (PHUs), small health facilities staffed with doctors, nurses, and midwives. Each PHU has typically a catchment area of seven to ten villages with one community health worker per village and one supervisor per PHU, for a total of approximately 15,000 health workers and 1,500 supervisors nationwide.

Health workers (bottom layer) The role of the health workers is to provide a package of basic healthcare services in their community. They do so by making home visits to expecting mothers or mothers who recently gave birth, during which they provide: (i) health education (e.g., about the benefits of a hospital delivery); (ii) timely pre- and post-natal check-ups, and (iii) accompany women for birth to the health facility. They also conduct visits to households with young children in which they: (i) educate them on how to prevent and recognize symptoms of malaria, diarrhea, and pneumonia, (ii) treat non-severe cases of malaria and diarrhea, (iii) screen for danger signs and refer for further treatment at an health facility when necessary. To ensure a high visit quality, workers are asked to follow a checklist each time they provide a service. We describe the checklists in Appendix 1.B.1.

Health workers are hired locally and typically have no experience in the health sector prior to joining the program. They work part-time and are paid a fixed monthly allowance of 150,000 SLL ($17.5) by the MoHS.[3] In Appendix 1.B.2, we provide additional information about hours worked and earnings from secondary activities.

[3]Throughout the draft, we use the January 2019 exchange rate: 1 USD = 8,550 SLL (Sierra Leonean Leones).

Supervisors (top layer) The role of the supervisors is to train and advise health workers located in their PHU (typically, seven to ten health workers per supervisor). They do so in three ways. First, they organize monthly one-day "general trainings" at the local health facility which cover key health topics, such as diagnosing, treating and recognizing danger signs for referral to health facilities. Second, supervisors organize "one-to-one" trainings with health workers on a monthly basis in their respective villages. Third, supervisors provide "in-the-field supervision" by accompanying health workers on household visits. During these visits, supervisors are not tasked to provide services themselves to the households, but rather to provide health workers with concrete feedback on how to improve service delivery and continuous on-site training. Supervisors' presence during these household visits also helps build trust towards the health worker in the community and reinforces the demand for her services. This is particularly important, since community members may initially have doubts about the expertise of the health worker - who is typically known by the community as a farmer or shopkeeper - and the supervisor can play a key role in legitimizing their position in the eyes of the community. Thus, overall, a substantial share of the support offered to the worker is personalized, which limits the potential for economies of scale in supervisor effort. Supervisors do not take any personnel decisions (hiring, firing, promotions, etc.). These are taken by the head of the PHU.

Most supervisors are former health workers. They also work part-time and are paid a fixed monthly allowance of 250,000 SLL ($29.2) by the MoHS. We provide more details on the supervisors' tasks and earnings in Appendix 1.B.3.

Complementarities across layers In the setting we study, supervisors are mostly engaged in supporting frontline workers. This is a common arrangement in many organizations, and sets our paper apart from a recent literature that focuses on the monitoring role played by middle managers (Callen et al., 2020; Muralidharan et al., 2021; Bandiera et al., 2021; Dal Bó et al., 2021; Dodge et al., 2021). In our context, supervisors generate demand for the workers' services by training the workers and by building trust towards them in the community. This can create a strategic complementarity between worker and supervisor effort. When a supervisor increases her effort, the worker is able to generate more visits for the same amount of time spent in the community. Similarly, the effort of the supervisor has a larger return when the worker is motivated and makes the most of the stronger demand for their services created by the supervisor.

1.2.2 Intervention and Research Design

We study the introduction of a new incentive scheme that pays a piece-rate of 2,000 SLL ($0.23) for each reported household visit. We have four experimental conditions. In the *worker incentives treatment* (T_{worker}), the incentive of 2,000 SLL is paid entirely to the health worker who provides the visit.[4] In the *supervisor incentives treatment* (T_{supv}), the incentive of 2,000 SLL is paid entirely to the supervisor of the health worker who provides the visit. In the *shared incentives treatment* (T_{shared}), the incentive is equally shared between the health worker and the supervisor (1,000 SLL each). In the *control group* (status quo), the incentive is paid neither to the health worker nor to the supervisor.[5]

Our experiment takes place in 372 PHUs throughout Sierra Leone, with the intervention running from May 2018 to August 2019. Appendix 1.B.5 provides details on the location of the 372 PHUs, and Appendix 1.C discusses research ethics and pre-registration of our study.

[4] The size of the piece rate is substantial: a health worker can earn up to 14% of her monthly fixed allowance if she provides one visit every other day.

[5] We chose the three most natural ways of splitting the incentives in our context, as further discussed in Appendix 1.B.4.

The 372 PHUs were randomly assigned to one of the four experimental groups, in equal proportions. The randomization was performed at the PHU level to limit spillovers across treatments. The randomization was stratified by district, average distance between the residence of the supervisor and the health workers in the PHU, and by the number of health workers in the PHU. Importantly, a subsample of the health workers in our study experienced a change in the promotion process six months after the start of the new incentive scheme, which we study in Deserranno et al., 2022b. In Appendix 1.B.7, we describe the change in the promotion system and show that the results of this paper are orthogonal to this variation.[6]

Description of the intervention The incentive scheme has three important features. First, the incentives were disbursed by a reputable external organization independent from the government. Subjects were paid on a monthly basis through mobile money and without any delay. This enabled us to establish the credibility of the new incentive scheme in the eyes of all experimental participants.

Second, incentives were paid based on worker self-reports. This is a common arrangement for incentives schemes with decentralized delivery agents, as directly monitoring output is typically expensive and impractical (e.g., Soeters and Griffiths, 2003; Shapira et al., 2018). To report a visit, the worker has to send an SMS from their main phone number to a toll free number. To trigger a payment, the SMS needs to indicate the date of the service and the contact number of the patient, and needs to be sent from the worker's registered phone number. The latter implies that supervisors or households are unable to report services for the workers. All health workers of our study (including those in the control group)

[6] Specifically, we show that: (a) the results hold if we restrict the analysis to the sample of health workers who did not experience any change in the promotion system, (b) the treatment effects are orthogonal to whether the health worker experienced a change in the promotion system or not.

were asked to report their visits, but only those in the treatments were incentivized based on the SMSs. We present more information about the reporting system in Appendix 1.B.

Our set-up discourages over-reporting through extensive back-checks and strong penalties. A random 25% of reports are verified by contacting the household mentioned in the report, and a worker caught reporting a visit that did not take place would not be eligible for any further incentive payment and would be reported to the MoHS. Back-checks and penalties were discussed extensively during the training on the reporting system that workers received prior to the start of the intervention.[7] We will later show that the threat was credible and nearly eliminates over-reporting.

Our design, however, does not prevent under-reporting. Even though the SMS reporting tool is free to use, reporting is inherently costly. First, reporting takes time and requires gathering information on the patients' name and phone number, which patients may not always be willing to share. Second, mobile phone coverage is unreliable and unpredictable in rural areas of Sierra Leone, thus limiting health workers' ability to send the SMS on the spot. This can lead to under-reporting if the worker subsequently forgets to send the SMS or sends an incomplete SMS with missing information. In Section 1.4.2, we will show that under-reporting is frequent in our setting. Similarly low reporting rates have been documented in other low-income countries. Karing (2021), for example, shows that local health facilities in Sierra Leone under-report vaccination entries, despite the presence of financial incentives, likely due to hassle costs.

Third, the incentive scheme rewards (reported) output, rather than (reported) effort. Output incentives are widespread both in the private and public sector.[8]

[7] To keep things as comparable as possible across experimental groups, all workers received the same training and the same number of back-checks, including workers in the control and supervisor incentive groups.

[8] In the financial sector, for example, a large fraction of the pay of financial analysts is variable

They have the advantage of rewarding workers based on a measure (output) which is more verifiable than effort. As we will show later, they also have the advantage of incentivizing both the worker and the supervisor to internalize some of the positive spillovers of their effort on the productivity of other subjects.

Transparency of the incentive scheme To mirror most workplace environments where supervisors have information about the pay structure of the subordinates but subordinates are not informed about their superior's compensation (Cullen and Perez-Truglia, 2019, 2021; Deserranno et al., 2022b), we informed all supervisors in the study about the worker incentives but did not inform the workers about the supervisor incentives. As we discuss in Section 1.5.3, this limits the presence of negative morale concerns resulting from pay inequality. Workers could only learn about the presence of supervisor incentives from the supervisors themselves, and few supervisors seem to have shared this information with their workers.

A second important point is that supervisors were not given information about the number of visits reported by each worker, nor about worker earnings from the incentive scheme. Since worker reporting is not constant across experimental groups, disclosing this information would have introduced differential observability of worker effort across treatments, and hence would have confounded the interpretation of our results. Additionally, the fact that supervisors are not aware of worker earnings further minimizes the possibility that the supervisor and the worker collude to report visits that have not actually been carried out.

---

and proportional to the amount of capital they raise, while the head of the unit is typically paid a bonus proportional to the amount of capital raised in the entire unit. In the retail sector, the commissions earned by both managers and frontline salesmen are a function of total revenues. In most micro-finance or agriculture extension programs, frontline workers are rewarded for the number of clients who take up the financial/agriculture product in their village, while their supervisors are rewarded for the total number of clients in the district.

Side payments We made clear to all supervisors that they could share all or part of their incentive with workers at their discretion. These transfers could potentially be used to incentivize worker effort. However, whether supervisors will choose to provide such payments is an open question. While recent studies have highlighted the use and sharing of bribes in organizations to access jobs or rents (Weaver, 2021; Lameke et al., 2020), there is still very little evidence on whether individuals are able to use transfers to incentivize each other to exert more effort. A number of features of our context, which are common across organizations, may make such informal incentives hard to implement. First, supervisors have a limited ability to precisely observe the worker's level of effort and reporting behavior, since production is decentralized (also, as explained above, we did not inform supervisors of the number of reports filed by each worker). This makes it hard for the supervisor to assess whether workers exert the level of effort that was requested from them in exchange for a side payment. Second, making binding commitments may be difficult because side contracting is inherently informal and the worker would have limited means to punish the supervisor for defaulting on a side payment (e.g., the worker's threat to reduce future effort would not be credible, since the organization may punish the worker for such low effort). Given this difficulty, the supervisor may need to compensate the worker for the perceived risk of default (Bubb et al., 2018). Third, there may be social norms or psychological factors that limit redistributions within the boundaries of the same organization (Hines and Thaler, 1995). The second and third factor are also likely to inhibit transfers from workers to supervisors.

1.2.3 Data and Balance Checks

Data Sources

Our study leverages three main sources of data.

Staff surveys. All 372 supervisors and 2,970 health workers in the 372 PHUs were surveyed at baseline in April-May 2018 and at endline in June-September 2019 (fifteen to sixteen months after the implementation of the treatments). They were surveyed on their demographic background, their health knowledge, and their job. We also have access to village-level information (e.g., distance to the health facility, mobile network coverage) collected from a leaflet that is given to each health worker by the PHU.

Household surveys. A random sample of three eligible households per village (∼7% of the households) were surveyed at endline in June-September 2019. The respondent of the survey was the female household head, who is the most knowledgeable about health topics. Each respondent was asked questions on the number of visits received by the health worker and the quality of these visits, trust in the health worker, disease incidence among young children, access to pre- and post-natal care. We will later use these data as our main measures of health worker performance.

Administrative data. Throughout the duration of our experiment, we have access to two sources of administrative data. First, we observe the number of valid SMS reports sent by each health worker, along with the incentive payments. Second, the MoHS provided us with information on the number of health services/patients treated by each local health facility at the monthly level (e.g., number of institutional births at the facility, number of children fully immunized at the facility, number of fever/malaria/diarrhea cases treated at the facility).

Summary Statistics and Balance Checks

Table 1.1 reports summary statistics and balance checks for the characteristics of the supervisors (Panel A), health workers (Panel B), households (Panel C), and villages (Panel D).[9] Panel E reports statistics on the number of health services provided by the local health facility (one per PHU) in the month before the start of the experiment.

Panel B shows that 71% of the health workers are male, 70% have completed primary education and 8% have completed secondary school. On average, health workers are 37 years old, are responsible for 55 households each, and live 3.4 km away from the supervisor. Panel A shows that supervisors are more likely to be men than the health workers (92%) and are more likely to have completed secondary school (25%). They are responsible for an average of 8 health workers each. Panel C shows that household respondents are less educated than health workers and supervisors, with only 25% having completed primary school. Household members are also less wealthy, as measured by a wealth score from 0 to 8 that counts the number of items owned on a list of household items (e.g., clothes, pair of shoes, cooking pots). On average, a household owns 1 out of the 8 items while workers and supervisors own 2.5 and 3 items respectively. Households live on average 1.4 km away from the health worker.

Panel D shows that 77% of the villages have an accessible road to the health facility. Phone network is available in 84% of the villages but is mostly unreliable. We will later show that the lack of reliable phone availability substantially increases the cost of SMS reporting. Finally, Panel E shows that health facilities record 47 pregnant women visits per month, 13 institutional births, 11 infants immunized, 66 cases of malaria and diarrhea among children under five.

[9] Given the absence of a baseline household survey, we asked households in our endline survey a set of retrospective questions that are unlikely to vary over time (i.e., age, education, location) and report those in Panel C.

To perform the balance checks, we regress each baseline characteristic on a dummy variable for each of the 3 treatments, controlling for stratification variables and clustering standard errors at the PHU level in worker/village level regressions. Column (11) reports the p-value from a joint F-test of the equality of all treatment groups. The baseline characteristics are balanced across treatments except for the age of the health worker (p-value of 0.062). In Table 1.A.1, we report the p-value for each pairwise treatment comparison. Out of 156 pairwise comparisons, 16 are statistically significant with a p-value below 0.1.

1.3 Model

We propose a simple model of service provision that features both contractual frictions and a positive complementarity between worker and supervisor effort. The model illustrates how the combination of effort complementarities and contractual frictions makes one-sided incentive schemes sub-optimal.

For simplicity, we consider the case of a single frontline worker (player 1) and a single supervisor (player 2).[10] The worker's task is to visit households and offer them health services. The supervisor's task is to make it easier for the worker to deliver this service, as explained in Section 1.2 (e.g., by training and advising the worker). The players interact over two periods. In the first period, the supervisor chooses a level of effort e_2, and offers to pay the worker a side payment of $s \in [0, \infty)$ for every visit that the worker completes. In the second period, the worker observes the effort choice of the supervisor and the side payment she offers, and then chooses effort e_1. This sequential structure reflects the hierarchical nature of the relationship as well as the fact that much of the supervisor's support

[10]This departs from our empirical setting, in which supervisors are responsible for multiple workers. As explained at the end of the section, this simplification does not affect the main results of the model.

offered to the worker (e.g., training) is given in advance of the worker's choice of effort.

Offering side payments is costly. We model this by assuming that a side payment of s costs to the supervisor zs, with $z \geq 1$. z is a reduced form parameter that captures any barrier to the offer of a side payment (e.g., the poor observability of worker effort, social norms, stickiness of payments), or the difficulty of making binding commitments (e.g., the supervisor may need to compensate the worker for the perceived risk of default). These contractual frictions limit the scope for Coasian bargaining.

Household visits y are produced as a result of both worker and supervisor effort. We capture this with the following output function:

$$y = \alpha e_1 + \gamma e_1 e_2 \tag{1.1}$$

where α is weakly positive. Importantly, when $\gamma > 0$, efforts are strategic complements: the higher the effort of one player, the larger the return to the effort of the other player. Also, this functional form captures the intuition that, when $e_1 = 0$, the supervisor cannot generate any visit no matter how much effort she spends training and advising the worker.

Both players maximize a private payoff that is given by the benefit that the player gets from the visits completed by the worker minus the cost of effort. We assume that each player i gets a benefit of b_i for every completed visit. This captures the combination of intrinsic and extrinsic motives that players may have to exert effort in the absence of performance-based incentives (e.g., there may be a threat of losing the job or social status that decreases in y).[11] Additionally, the worker gets a monetary payment of pm per visit in the three treatments, where $p \in [0, 1]$

[11] In the empirical setting, agents also receive a fixed wage. Given the linear utility specification, the introduction of this additional term will not affect our conclusions.

is the share of the output incentive assigned to the worker, i.e., in the *worker incentives treatment*, $p = 1$; in the *shared incentives treatment*, $p = 0.5$; and in the *supervisor incentives treatment*, $p = 0$. The supervisor, on the other hand, is paid an incentive of $(1-p)m$ per visit completed by the worker.[12] Further, the worker also receives a transfer from the supervisor of s per visit, and the supervisor pays an amount zs per visit in order to make this transfer.[13] Finally, both agents bear a convex cost of effort: $c(e_i) = c_i e_i^2$.

In sum, the payoffs of the worker and of the supervisor are given by:

$$\pi_1 = (b_1 + pm + s) * y(e_1, e_2) - c(e_1) \tag{1.2}$$

$$\pi_2 = (b_2 + (1-p)m - zs) * y(e_1, e_2) - c(e_2). \tag{1.3}$$

We solve the model using backward induction. To obtain our main analytical results, we simplify the problem and assume that $b_1 = b_2 = 0$, $c_1 = c_2 = c$, $m = 1$ and $\alpha = 1$. This enables us to illustrate the core features of the model, which are determined by the production function, the possibility of side payments, and the sequential interaction, while setting aside additional considerations that emerge when costs or benefits are asymmetric. We will relax these assumptions when we take the model to the data in Section 1.6.

In this simplified setting, the optimal side payment is given by:

[12] In the empirical setting, agents are paid uniquely for the visits they performed *and* subsequently reported. We abstract from modeling worker reporting behavior because it complicates the model without affecting its main results. See the discussion at the end of the section.

[13] In practice, transfers from supervisors to workers could be fixed (not proportional to visits) or based on the number of visits reported by the worker. Again, such extensions do not affect the main intuition of the results.

$$s^* = \begin{cases} \frac{1-p(1+z)}{2z} & p \leq \frac{1}{1+z} \\ 0 & p > \frac{1}{1+z} \end{cases} \tag{1.4}$$

This formula shows that the optimal side payment decreases with the contractual frictions (z) and the incentive offered to the worker (p). If contractual frictions are large and the worker receives a large share of the incentive ($p > \frac{1}{1+z}$), the supervisor will not make any side payment. We derive optimal efforts for these two cases - positive side payments ($p \leq \frac{1}{1+z}$) and zero side payments ($p > \frac{1}{1+z}$) - and present the complete mathematical analysis of the model in Appendix 1.D. As expected, the efforts of both players increase in the strength of the complementarity. Further, due to the complementarity, agents' efforts do not necessarily increase monotonically in the share of the incentive that is offered to them.

We can use the model to illustrate how the optimal incentive scheme depends on contractual frictions and complementarities in effort. In particular, we consider a policy maker that aims to find the level of p that maximizes visits. In what follows, we will call incentive schemes that only incentivize one player ($p = 1$ or $p = 0$) *"one-sided,"* and schemes that incentivize both players ($0 < p < 1$) *"two-sided."* Also, we will refer to incentive schemes that weakly maximize visits as "optimal." Finally, we restrict attention to values of γ and c such that $z\gamma^2 < 8c^2$. This condition limits the relative size of the complementarity, guaranteeing positive optimal efforts (as we show in Appendix 1.D.2). We can prove the following result.

Result 1 *When effort complementarity is lower than a threshold level t, there is a unique optimal incentive scheme, which is one-sided: $p^* = 1$. When effort complementarity is equal or larger than t, there is always a two-sided scheme which is optimal: $p^* \epsilon (0, 1)$. If there are contractual frictions, this optimal two-sided scheme is the unique optimal scheme. If there are no contractual frictions, $p = 0$ may also be optimal.*

This result is established in two steps, which are discussed in detail in Appendix 1.D and summarized here.

When complementarities are low ($\gamma < t$), supervisor effort has only a limited effect on the worker's ability to carry out household visits. In this case, it is straightforward to show that household visits are maximized by offering the entire incentive to the worker.

When complementarities are large ($\gamma \geq t$), supervisor effort becomes central to the optimal incentive decision. If contracting is costly ($z > 1$), incentive schemes that concentrate most of the rewards on one subject are not effective, since the drop in productivity that comes from the low effort of one subject more than offsets the monetary incentive offered to the other subject. Instead, efforts are maximized by intermediate values of p. Thus, the optimal incentive scheme is two-sided, as we show in Figure 1.1.[14]

If complementarities are large ($\gamma \geq t$) and there are no contractual frictions ($z = 1$), the supervisor is able to perfectly match any changes in incentive in the interval $\left[0, \frac{1}{1+z}\right]$ with a commensurate change in side payments. All values of p in that interval result in the same number of visits. If this is the highest possible number of visits (as shown, for instance, in the example analyzed in Figure 1.A.1b), then all $p \, \epsilon \left[0, \frac{1}{1+z}\right]$ are optimal.

In sum, the model clarifies that, when efforts are strong strategic complements, it is optimal to offer a two-sided incentive scheme that rewards both players. Furthermore, in this case, we may observe that subjects' own efforts do not increase monotonically with the incentive that is offered to them. One final implication

[14]More precisely, the optimal incentive is either $p^* = \frac{1}{1+z}$ (which is the optimal incentive in the interval $[0, \frac{1}{1+z}]$) or $p^* = \frac{1}{6} + \frac{\sqrt{\gamma^2+24c^2}}{6\gamma}$ (which is the optimal incentive in the interval $(\frac{1}{1+z}, 1]$). In Figure 1.A.2, we show how optimal efforts and side payments change as p changes.

Figure 1.1: Optimal Incentives ($\gamma \geq t$ and $z > 1$)

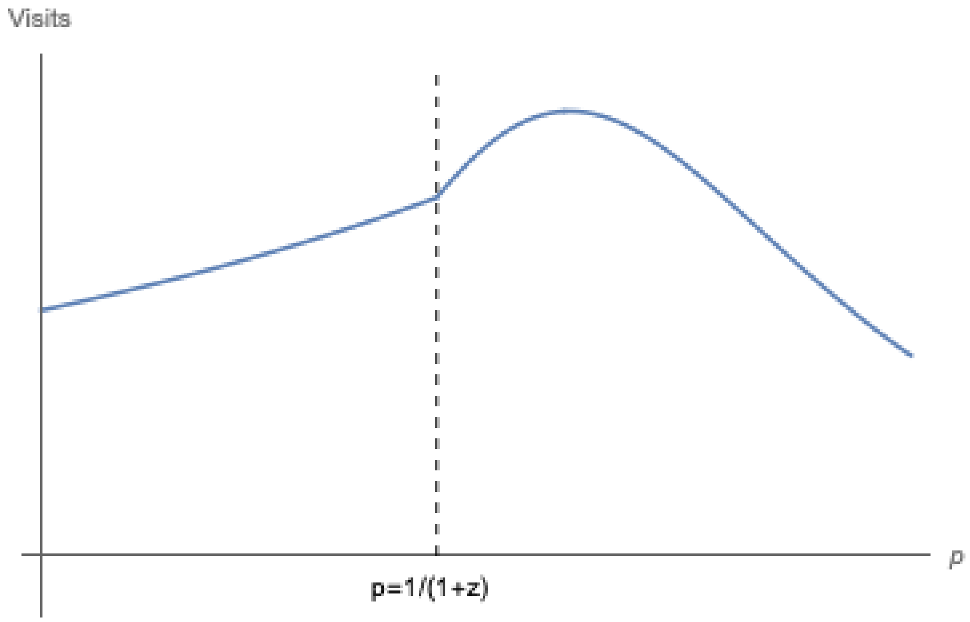

of the model, which we explore in Appendix 1.D, is that the difference in output between the optimal two-sided incentive scheme and the one-sided scheme $p = 1$ increases in the complementarity γ. Thus, if in the experiment we find that a two-sided incentive scheme is optimal, we would also expect that the difference in output between this scheme and the worker incentive scheme is larger for supervisor-worker pairs that have a high γ. We will explore these predictions empirically in Section 1.5.

The model also sheds light on the important role played by side payments. In particular, two predictions will help us interpret our experimental results. First, the model shows that, when there are no contractual frictions, all incentive schemes that motivate positive side payments produce the same number of visits. In contrast, when there are contractual frictions, changes in the allocation of incentives always affect output. In other words, if we observe positive side payments and differential treatment effects on output, this indicates that the supervisor and the worker cannot contract costlessly. Second, the model shows that there is an additional factor that can limit side payments. In Appendix 1.D.7, we present an extension of the model that allows for heterogeneity in benefits and costs. This extended model shows that the supervisor will not offer any side payment when the benefit b_2 that she receives from household visits absent our intervention is low compared to the benefit b_1 that is received by the worker.[15] In these cases, it would be optimal for the worker to pay the supervisor to exert effort - an action which we do not allow in the model and do not observe in the data, presumably because the frictions preventing transfers from the bottom to the top of the hierarchy are even larger than the frictions that impede transfer from the top to the bottom. In sum, the lack of side payments is theoretically consistent either with high contractual frictions preventing the supervisor from offering side payments, or with an asymmetry in how much workers and supervisors value output. However, in the latter case, these limited side payments are sufficient to equalize output.

Finally, we note that, to keep the model tractable, we depart from our empirical setting in two main ways. First, in the model, we abstract from the fact that each supervisor has multiple workers. This prevents us from exploring the optimal targeting of supervisor effort across heterogeneous workers, but does not affect the model's main predictions. Second, in the model, the incentive is paid on the basis of the number of actual visits completed, rather than the number of visits reported. In the structural estimation Section 1.6, we present a version of the model in which incentives are based on the number of visits reported. To model under-reporting, we posit that the reporting process suffers from random shocks (e.g., bad network), which prevent some visits from being reported. We allow the

[15]Both b_1 and b_2 may stem from the agents' worry that low output will result in their dismissal from the organization. As supervisors tend to be more experienced and established in the organization, it is likely that they are less concerned by the possibility of being fired, compared to workers. Alternatively, supervisors may have better outside options and would hence suffer a smaller utility loss if they lose their position.

reporting rate to differ by treatment, since presumably workers can take costly actions to over-ride the shock, but their willingness to take these actions is a function of the incentive they get paid for each report. Indeed, empirically, the reporting rate increases in p. This raises the relative attractiveness of the worker incentive scheme compared to the other schemes. However, as long as the elasticity of reporting with respect to p does not exceed a threshold, all model results remain qualitatively unchanged.

1.4 Main Results

1.4.1 Output

Quantity of visits We start by assessing the causal effect of our incentives treatments on the number of visits provided by health workers. We estimate the following regression equation:

$$Y_{ij} = \alpha + \beta_1 T_{worker,j} + \beta_2 T_{supv,j} + \beta_3 T_{shared,j} + Z_j + \varepsilon_{ij} \quad (1.5)$$

where Y_{ij} represents the number of household visits provided by health worker i in PHU j. $T_{worker,j}$, $T_{supv,j}$, and $T_{shared,j}$ are indicators for whether incentives in PHU j were assigned to health workers only, supervisor only, or were shared between the two. In our model's notation, these correspond to $p = 1$, $p = 0$ and $p = 1/2$, respectively. Z_j are the stratification variables, discussed in Section 1.2.2. ε_{ij} is an error term clustered at the level of the treatment assignment, the PHU.

To measure the number of household visits provided by the health worker, we do not rely on the number of visits reported by the worker because this often differs from the true number of visits due to under-reporting, as discussed in Section 1.4.2. Instead, we interviewed each sampled household on the total number of natal- and disease-related visits received from the health worker in the six months preceding the endline survey.[16] For each worker, we then calculate the mean number of visits received by a household (mean of 7.3). We also study the coverage and range of services provided by the health worker, which we proxy with the share of households who were visited at least once (mean of 71%) and the number of different visit types received by a household (mean of 1.7).

Our main results are reported in Table 1.2 column (1) and the corresponding Figure 1.2. They show that introducing performance-based incentives significantly boosts the number of household visits provided by the health worker, regardless of whether the incentives are one- or two-sided. The mean number of visits per household in the control group is 5.334. This number increases by 2.090 (39%) in the worker incentives treatment, by 2.145 (40%) in the supervisor incentives treatment, and by 3.356 (63%) in the shared incentives treatment. Interestingly, offering the entire incentive to the health workers is equally effective than offering the entire incentive to the supervisor. Both interventions, however, are outperformed by the two-sided incentive scheme, which achieves 17% more visits overall. Relative to the control group, the boost in visits generated by the two-sided incentive scheme is 61% larger than the boost that results from either one of the one-sided schemes.[17]

When we break down household visits by their type, we find that, compared to the one-sided treatments, shared incentives generate significant gains over both natal-related visits and disease-related visits (Table 1.A.3). Health workers in the

[16]To minimize recall bias, households were asked about visits received "since the start of the year," which roughly corresponds to the past six months.

[17]These results estimate the treatment effects on the average number of visits provided by the health worker to a single sampled household in the six months preceding the endline survey. For completeness, in Table 1.A.2, we also report the corresponding treatment effects on the *total* number of visits provided to sampled households per month (column 1) and on the *total* number of visits provided in the community per month (column 2). The latter outcome variable is measured as the number of visits per month in our sample divided by the share of households included in our sample. We estimate that health workers provide a total of 41 monthly household visits in the community in the control group. This number goes up to 59 in T_{worker} and T_{supv}, and to 67 in T_{shared}.

Figure 1.2: Effect of Incentives on the Number of Visits

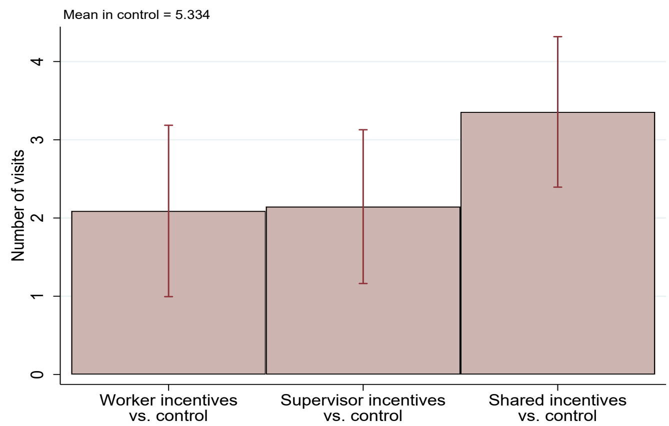

Notes: The figure plots the difference in the number of visits provided by the health worker between each treatment group and the control group. The coefficients are estimated from a regression of the number of visits on the treatment dummies, controlling for stratification variables with standard errors clustered at the PHU level. Bars are 95% confidence intervals.

shared incentives treatment also achieve higher coverage (Table 1.2 column 2) and provide a higher variety of services (column 3).

Quality of the visits and targeting The higher number of visits provided by workers in the shared incentives treatment may potentially come at the expense of visit length (which is not incentivized), so that the aggregate amount of time dedicated to the job remains unchanged. This would be problematic: as discussed earlier, workers are expected to follow a checklist when they visit a household and short visits may be an indication that such checklist is not properly followed, and thus that the service provided may be of lower quality.

We do not find a quantity-quality trade-off. Table 1.2 (column 5) shows that, conditional on having received at least one visit, the average visit length reported by a household (23 minutes) did not decrease in the shared incentives treatment.[18] The

[18]Obviously, this result has the caveat that we are conditioning the sample on a potentially endogenous variable. When we assign an average visit length of zero to the 29% of households who were never visited by the worker, we obtain that the shared incentives increase visit length (see column 4). This captures both the intensive and the extensive margin of effort.

average number of health topics the household discussed with the health worker during a visit increases by 26% in the shared incentives treatment (column 6). This is consistent with these workers receiving more training from the supervisor, as further discussed in Section 1.5.1. Importantly, the share of households who report trusting the health worker in the shared incentives treatment also increases: it is 7.1 percentage points (10%) higher in the shared incentives treatment than in the control, and 3.5 percentage points (5%) higher than in both one-sided incentives treatments than in the control (column 7).

We also explore whether the higher number of visits in the shared incentives treatment comes at the expense of worse targeting of households: health workers may switch from visiting the most deserving households (i.e., poor households) to households who can be visited at a lower cost (i.e., households who are located close by or who are friends or family members). An increase in mis-targeting would be concerning as it would offset some of the gains coming from a larger number of visits.

To analyze targeting, we run a household-level regression of the number of visits received by the household on the treatments dummies interacted with whether the household is poor (wealth score below median), lives within 30 minutes of the health worker's home and is a family member or a friend of the health worker.[19] Table 1.A.4 shows that households who are socially or geographically close to the health worker are less likely to be targeted in the control group (columns 2 and 3), while the household's wealth score does not predict visits (column 1). Such targeting remains nearly unchanged when workers are paid a higher incentive.

[19] We run the following household-level regression: $Y_{hij} = \alpha + \beta_1 T_{worker,j} + \beta_2 T_{supv,j} + \beta_3 T_{shared,j} + \beta_4 X_h + \beta_5 T_{worker,j} * X_h + \beta_6 T_{supv,j} * X_h + \beta_7 T_{shared,j} * X_h + Z_j + \varepsilon_{hij}$, where Y_{hij} represents the number of visits that the household h received from health worker i in PHU j and X_h is a household characteristic (e.g., poor, social/geographical distance to health worker). All the other variables are defined as in equation (1.5). ε_{hij} is an error term clustered at the PHU level.

Overall, these results alleviate concerns related to quantity-targeting trade-offs. The fact that households who are friends or family members of the health worker are equally likely to report visits than households who are socially distant also indicates that visit misreporting is limited in our context.

Access to natal-care services and disease incidence We now test whether the increase in the number of natal- and disease-related services provided by the health worker in the shared incentives treatment translates into better access to health services and better health outcomes.

We start by analyzing households' access to pre- and post-natal care. We measure access with an equally-weighted average of z-scores of the key indicators used for the assessment of pre- and post-natal care quality under the World Health Organization framework (four pre-natal visits, institutional birth, post-natal care within two days of birth, up-to-date vaccination, breastfeeding).[20] Table 1.3 (column 1) shows that the shared incentives treatment leads to better access to pre- and post-natal care. More precisely, the pre- and post-natal care index is 0.092 standard deviations higher in the shared treatment relative to the control (significant at the 1% level). Columns (2) to (6) present the results for each each single component of the index.

Next, we analyze diseases incidence among children under the age of five, which we proxy with an equally-weighted average of z-scores of three variables: the share of households who report that at least one child under five years of age had fever,

[20] Questions on pre-natal an post-natal care were asked to households composed of a woman who gave birth in the year preceding the endline survey. Table 1.A.5 (column 1) shows that this fertility measure is not affected by our treatments. Pre-natal care is measured by asking women who gave birth in the past year whether they received at least four pre-natal visits from any provider, and post-natal care is measured by asking them whether they gave birth in a health facility (vs. at home), whether they received at least one post-natal visit within two days of birth, whether they breastfed their infant for at least six months, and whether their infants are up-to-date on the vaccination schedule.

diarrhea or cough in the past month.[21] Table 1.3 (column 7) shows that disease incidence index is 0.053 standard deviations lower in the shared incentives treatment than in the control group (significant at the 5% level). This is driven by households in the shared incentives treatment reporting fewer fever instances, while we see no effect for diarrhea and cough (columns 8-10). These households also have better knowledge about how to prevent malaria (i.e., sleep under a treated bednet) and diarrhea (i.e., wash hands with soaps, drink clean water): Table 1.A.5 (column 2). We find no significant effects on under-five mortality rates (Table 1.A.5, column 3), presumably due to the relatively short timeframe of the experiment.

We corroborate these results on health outcomes using administrative records from the local health facility (PHU-level data), which do not suffer from any recall or response bias. The results are presented in Table 1.A.6, columns (1)-(7). In line with the household survey data, we find that the number of recorded pregnant women services, institutional births and fully immunized infants at the health facility is higher in the shared incentives treatment than in the other groups, albeit the results are less precisely estimated. All three incentives treatments appear to increase the number of malaria and diarrhea cases treated at the health facility relative to the control group. Given the lower disease incidence rate reported by our sampled households, these positive coefficients are consistent with health workers referring sick children to the health facility more frequently in the treatment groups than in the control.

[21] The three most common diseases among children in Sierra Leone are malaria, pneumonia and diarrhea. Because households may not be aware of which disease a child suffered from, we asked them to report whether any child had common symptoms associated with each disease (fever, cough and diarrhea).

1.4.2 Reporting and Cost-Effectiveness of the Intervention

This section assesses the relative cost-effectivness of the three incentive schemes. All schemes pay 2,000 SLL per visit *reported* by the health worker. Cost-effectiveness is thus a function of both the actual number of visits carried out and the number of visits reported. We start by evaluating, in Table 1.4, whether the number of visits reported is impacted by our treatments. Column (1) shows that reported visits are highest in the worker incentives treatment, even though we have shown that actual visits are maximized by shared incentives. More precisely, we find that, in the six months preceding the endline survey, workers send an average of 8.7 SMS reports per month in T_{worker}, 6.3 in T_{shared}, and 3.7 in T_{supv}. The reporting differences across treatments are relatively stable over time (see Figure 1.A.3). These results imply that the most expensive incentive scheme for the organization is T_{worker}. More precisely, Table 1.4 (column 8) shows that the new incentive scheme costs the organization an average of 131,593 SLL in T_{worker}, 93,953 SLL in T_{shared}, and 54,108 SLL in T_{supv}.

In Table 1.4 (column 2), we present results on the reporting rate, i.e., the ratio between the number of SMS reports per month (column 1) and the *actual* number of visits per month.[22] We also present results on dummy variables capturing whether a worker under-reports or over-reports actual visits. This analysis shows that health workers generally under-report the number of visits provided, especially when they are not incentivized to do so: they report 30.3% of the actual visits in T_{worker}, 17.1% in T_{shared}, 13.8% in T_{supv}. Moreover, the share of workers who

[22]The actual number of visits per month is calculated as the number of actual visits among the random sample of households we interviewed scaled up for the number of households in the community, as in Table 1.A.2 column (2). While the reporting rate we obtain from this calculation may be over- or under-estimated for a single health worker, average differences across treatments are meaningful and accurate. Note that households have no strategic incentive to misreport the number of visits received by the health worker and that the survey was not announced beforehand so that the health worker could not have influenced households to give favorable answers during the survey.

under-report is 12 times larger than the share of workers who over-report (Table 1.4, columns 3-4).[23]

These results confirm that reporting is costly, so that workers under-report even in T_{shared} and T_{worker} when they receive monetary incentive for reporting. As explained in Section 1.2, we hypothesize that a key driver of reporting costs is the poor quality of the mobile phone network, which makes it hard to send SMS messages from some villages. Consistently with this hypothesis, we find that the reporting rate is close to zero in the 16% of villages where network connectivity is virtually absent, regardless of the level of the incentives (Table 1.A.7, Panel A). In villages with some network connectivity (even if often unreliable), the reporting rate increases with the level of the incentives (Table 1.A.7, Panel B). In these villages, we estimate an elasticity of reporting to incentives of 0.75, and hence estimate that an incentive of 3,800 SLL would lead health workers to report 100% of their visits.[24] The results are robust to controlling for correlates of network connectivity (e.g., distance to urban area) interacted with the treatment dummies (column 2).

Policy choice What policy should the organization adopt on the basis of these results? Suppose that the organization wants to maximize household visits, conditional on the payment per *actual* visit not exceeding 2,000 SLL. In this case, the shared incentive intervention is unambiguously optimal for the organization. On the one hand, over-reporting is minimal in all treatments, so the cost per actual visit never exceeds 2,000 SLL. On the other hand, shared incentives maximize

---

[23] Given that under-reporting is widespread and differential across groups, in the structural estimation section we extend our basic model to explicitly take under-reporting into account.

[24] In villages with network connectivity, the reporting rate is 13.7 percentage points (75%) higher in T_{worker} relative to T_{shared}, while the incentives paid to the worker for reporting is twice as high. The estimated elasticity of reporting to incentives is thus 0.75 (75%/100%).

actual visits, and so satisfy the organization's objective.[25]

The shared incentive intervention is also optimal if the organization's objective is to maximize cost-effectiveness, i.e., to maximize the amount of actual visits generated per dollar spent. In Table 1.4 (column 9), we show that in the worker incentives treatment, the organization obtains an additional 6.5 visits per worker for each 2,000 SLL spent on incentives. This figures goes up to 9.6 visits for each 2,000 SLL spent in the supervisor incentives treatment, and to 16.1 visits for each 2,000 SLL spent in the shared incentives treatment (a significant difference of 9.6 visits compared to worker incentives).[26]

Shared incentives, however, impose a larger total cost compared to supervisor incentives. If this cost breaks the organization's budget constraint, the organization could either opt for supervisor incentives, which offer a similar increase in visits as worker incentives, for a lower cost; or it may decrease the amount of the incentive paid in the shared incentive scheme.

1.5 Mechanisms

The previous section showed that health workers provide significantly more household visits under shared incentives than under the one-sided incentives schemes,

[25] Note that over-reporting is minimal in our setting thanks, in part, to the presence of a fairly sophisticated and independent monitoring system. In the absence of such monitoring system, over-reporting may occur more frequently, especially in the worker incentives treatment, and this would presumably make the worker incentives even less attractive relative to the other treatments.

[26] A key caveat is that these results are partly driven by the differential rate of under-reporting. Due to under-reporting, the organization saves on incentive payouts that should instead accrue to workers and supervisors. If under-reporting was reduced, e.g., by a lowering of reporting costs, differences across treatments in the number of additional visits produced for each 2,000 SLL spent would be smaller. Further, the organization may want to design a scheme to reimburse agents for the incentives that have not been claimed, for equity reasons or to comply with labor laws, which would also reduce differences in cost-effectiveness. Finally, differences in cost-effectiveness would be smaller in settings where output is observable and hence incentives can be tied to actual output.

with no concomitant reduction in visit quality. In this section, we explore the mechanisms underlying this result. Guided by the theoretical framework developed in Section 1.3, we provide evidence consistent with the presence of *both* complementarities in the effort exerted by the supervisor and the health worker and limited side payments. We then present evidence against two alternative mechanisms: inequality aversion and a kink in the utility, cost or production functions.

1.5.1 Effort Complementarities

Three pieces of evidence point to the presence of strong effort complementarities in our setting. We discuss each in turn.

Supervisor effort First, we estimate the effects of our three incentive schemes on the levels of effort exerted by the supervisor. If effort complementarities were weak ($\gamma < t$), the effort of the supervisor should monotonically increase with the level of the supervisor's incentives, i.e., be higher in the supervisor incentives treatment relative to the other groups. We show next that this is not the case.

Recall from Section 1.2.1 that supervisors have three main tasks: (i) they provide in-the-field training and advising by accompanying health workers on household visits (henceforth, an "accompanied visit"), (ii) organize one-to-one meetings with each health worker, and (iii) organize monthly one-day general trainings. We measure (i) with the fraction of households who report having received an accompanied visit in the six months preceding the endline survey (mean of 20%).[27] We measure (ii) and (iii) by asking health workers the number of times the supervisor provided them one-to-one meetings in the six months preceding the endline survey (mean

[27]No household reports having received a visit from the supervisor without the presence of the health worker. Among households who received at least one accompanied visits, 97% received one accompanied visit, and 3% received two accompanied visits. This implies that roughly one fifth of the households have seen the supervisor once in the past six months, and the vast majority of the remaining households have never seen the supervisor.

of 1.4) and whether the supervisor organized a general training in the last month (mean of 99.4%).

Table 1.5 column (1) shows that the share of households who report having received an accompanied visit is 5.7 percentage points (35%) and 6.2 percentage points (38%) higher in T_{supv} and T_{shared} respectively, relative to the control group, while there is no difference between T_{worker} and the control group. Importantly, the coefficients for T_{supv} and T_{shared} are nearly identical, and this is despite the fact that the supervisor is paid an incentive which is twice as high in the former than in the latter. This suggests that the overall returns to supervisor effort are similar in the supervisor and shared incentive schemes, which is consistent with the existence of effort complementarities that indirectly compensate the supervisor in the shared incentive scheme for the lower monetary payment. Also, note that the treatment effects on the "accompanied visits" is much smaller in magnitude than the effects on any visit (accompanied or not) we documented in the previous section (Table 1.2, column 1). This implies that the increase in visits in the shared incentives treatment is mostly due to an increase in "unaccompanied" visits in which the health worker was not accompanied by the supervisor.

Columns (2) and (3) of Table 1.5 show that our treatments neither affect the number of times the supervisor provided one-to-one meetings to the health worker, nor do they affect the likelihood that the supervisor organized a general training. The latter is not surprising as supervisors are required to organize such trainings on a monthly basis, and 99% of them do so.

Two pieces of evidence provide direct support to the fact that supervisors play an "enabling" role in our context, rather than only a "monitoring" role. First, Table 1.A.8 column (1) shows that health workers improve their health knowledge the most in the shared incentives treatment.[28] Specifically, health workers in the

[28] Health knowledge is measured with a quiz on health knowledge which was administered to health workers at baseline and endline.

group incentives treatment seem to have improved their knowledge of when and whether to refer a child under 5 if she has fever or loose stools, and the threshold of breaths per minute above which a baby needs to be referred for fast-breathing. Second, Table 1.A.4 (column 5) shows that the boost in visits in the shared incentives treatment is similar for households that received an accompanied visit and those that did not, with the obvious caveat that this variable is endogenous. If the role of the supervisor was limited to monitoring, we would expect health workers to target their visits towards households that were in direct contact with the supervisor in the past, since presumably the supervisor would find it easier to contact these households again and to monitor whether the worker has visited them. That shared incentives boost visits for households who were never in direct contact with the supervisor suggests instead that health workers in this treatment have received better training and are able to raise demand for their services even when unaccompanied.

Finally, note that only 16% of the health workers report that their supervisors ever helped them with the SMS reporting (Table 1.A.8, column 2). This is not surprising as all health workers received extensive training on how to report at the start of the experiment (see Section 1.2and Appendix 1.B.6). Interestingly, the share of supervisors who helped health workers with reporting is comparable in the two one-sided treatments relative to the control group and is slightly lower in the shared incentives treatment. This indicates that the introduction of supervisor incentives did not divert supervisor's time away from productive tasks (e.g., training workers on health issues) towards helping with reporting.

Heterogeneity by health worker's experience Next, we present heterogeneous treatment effects by an empirical proxy of effort complementarity: limited health worker's experience. Health workers with little experience are less well-trained about health issues and less-known in the community, and they thus plausibly benefit more from the training and advice of the supervisor. We thus expect the shared incentives treatment to be more effective in boosting output and supervisor effort for these health workers, compared to their more experienced counterparts.

Table 1.A.10 estimates a fully interacted model and presents the treatment effects for workers with experience below the median (i.e., below 4 years) in Panel A and for workers with experience above the median in Panel B. For inexperienced workers, the shared incentives treatment increases the number of household visits provided by the health worker by 4 (85%) relative to the control group (column 1), and increases supervisor effort (measured with the share of households in which the health worker was accompanied by the supervisor) by 9.2 percentage points (70%; column 3). For experienced workers, these effects are significantly lower: they are about half the magnitude for visits and one third of the magnitude for supervisor effort. The results are robust to controlling for all worker characteristics correlated with health worker experience (listed in Table 1.A.9 column 1), and their interaction with the three treatment indicators (columns 2 and 4). Overall, the results confirm that the shared incentives treatment is particularly effective in boosting output and supervisor effort when effort complementarity between the layers of the organization is likely high.[29]

Mediation analysis As a final evidence in favor of effort complementarities, we perform a mediation analysis to test whether the boost in visits attributable to worker effort increases when supervisors exert more effort. Following Acharya et al.

[29] In contrast with Bandiera et al. (2007), Table 1.A.13 shows that supervisors are not more likely to target their effort towards health workers who they perceived as highly ranked in terms of performance at baseline, and are also equally likely to target their friends/family members. We also find no heterogeneity in supervisor effort and household visits with respect to the supervisor's span of control (the number of workers per supervisor). This might be explained by limited variation in span of control in our setting.

(2016), we estimate the Controlled Direct Effect (CDE) of the worker incentives treatment on visits *net of a mediator* - here, supervisor's effort. This quantity captures the treatment effect that would be observed if supervisor effort was fixed at an exogenous level, while worker's effort (which is not directly observable in our setting) was allowed to respond to the incentives.[30] We then present this "de-mediated" effect for different levels of supervisor's effort. In the presence of effort complementarities we would expect the increase in visits generated by the worker to grow in supervisor effort (when supervisor effort increases, the worker exerts more effort and the return to worker effort increases).

In line with this, Panel A of Figure 1.A.4 shows that the effect of worker effort on output increases substantially with supervisor effort, as measured with the fraction of household visits in which the health worker was accompanied by the supervisor. Indeed, the CDE of the worker incentives treatment on visits is close to zero when 0% of the household visits were accompanied by the supervisor and goes up to more than 2 at the opposite extreme when 100% of the household visits were accompanied.[31] This is consistent with a strategic complementary between worker effort and the in-the-field-training offered by the supervisor. We also find evidence of a strong complementarity between worker effort and the general training provided by the supervisor, while we see no complementarity with respect to the one-to-one meetings (Figure 1.A.4, Panels B and C).

[30]We focus on the comparison between the worker incentives treatment and the control group since a mediation analysis performed on the other treatments would be confounded by the fact that in those treatments the supervisor is directly incentivized to exert effort.

[31]We produce Figure 1.A.4 by following the steps outlined in Acharya et al. (2016). First, we regress the number of visits provided by a health worker on the worker incentives treatment, the mediator (supervisor's effort), and their interaction. From this, we obtain a de-mediated outcome, defined as the difference between actual visits and predicted visits based on the coefficients of all covariates (except the treatment) estimated at different levels of the mediator. Finally, we run a regression of the de-mediated outcome on the treatment and report the coefficients for different levels of the mediator.

1.5.2 Limited Side Payments

In this section, we document that side payments are limited in our context, and present suggestive evidence that this is because of the presence of contractual frictions.

To measure side payments, we collected detailed data on monetary and in-kind transfers from both the supervisors and the health workers. At endline, all supervisors were asked whether they transferred a portion of their incentive to health workers since baseline. If they did, we then asked each health worker to assess the value (in-cash or in-kind) of this side payment.[32]

Side payments are generally small and infrequent. In Table 1.6 column (1), we show that the share of supervisors who make positive side payments increases with the level of the supervisor incentive (1.1% in the control group, 1.6% in T_{worker}, 11.3% in T_{shared}, and 19.4% in T_{supv}), but that the large majority of supervisors do not make any transfer across all treatment groups. In column (3), we document that the average amount that a supervisor transfers to a worker *over an entire month* is 702 SLL (resp., 431 SLL) in T_{supv} (resp., T_{shared}). These amounts are very small if one considers that the supervisor earns an incentive of 2,000 SLL (resp., 1,000 SLL) per visit reported in T_{supv} (resp., T_{shared}), and that the average supervisor earns 55,280 SLL per month in T_{supv} (resp., 47,097 SLL in T_{shared}) from the incentive payment, as shown in Table 1.4 (column 7). Workers also occasionally make side payments to their supervisor when they are paid an incentive, but the amount of such transfers is negligible (average of 151 SLL in T_{worker}; see Table 1.6 column 4). Overall, this evidence shows that side payments do happen in our context, but their frequency and magnitude is minimal.

[32]This was asked to health workers rather than supervisors to limit recall bias. To make sure supervisors did not under-report transfers, they were made aware from the very start of the experiment that they were free to share incentives with their workers. See Section 1.2.2 for details.

Why are side payments limited? In Section 1.3, we discussed two possible explanations. First, the supervisor may find it optimal to offer a sizable side payment to the worker, but contractual frictions partially limit her ability to offer these payments. Second, the optimal side payment may be small or even zero if the value that the supervisor attaches to household visits is small compared to that attached by the health worker.[33] We also pointed out that, as long as we observe positive net payments from the supervisor to the worker, we can disentangle these two potential explanations by looking at the impacts on visits of the different treatments. If the level of side payments is low due to contractual frictions, we expect that changes in the share of the incentives allocated to the worker can generate large differences in visits. In contrast, if there are no frictions and we observe a low level of side payments to the worker due to the relative low value that the supervisor attaches to output compared to the worker, we should observe the same number of visits for all incentive schemes that generate positive side payments. Our results in the previous section, which show that visits are far from being equalized across treatments (see Figure 1.2) despite side payments being positive, point to the likely presence of contractual frictions in our setting. In what follows, we present two additional pieces of evidence pointing to the presence of these frictions.

Heterogeneity by supervisor's observability of worker output First, we study the sensitivity of side payments to proxies for "top-to-bottom" contractual frictions. These frictions are more likely to be present when worker effort or output is not observable to the supervisor, since this makes contracts hard to enforce. To measure the observability of output, we leverage the fact that, at endline, we asked each supervisor to rank the workers she supervises from the best to the worse in terms of their "overall work as a health worker." We correlate this

[33]If the asymmetry is large enough, the worker may actually find it desirable to offer a payment to their supervisor, but bottom-to-top contractual frictions may prevent them from doing so.

perceived rank with the *actual* rank of health workers obtained on the basis of the number of households visits completed at endline. The correlation is positive for most supervisors, except for 10% of them for whom the correlation is negative and who thus have poor observability. Table 1.A.9 shows that these poorly-informed supervisors tend to live further away from the health workers, while they have the same education, age, and wealth score.

Table 1.A.11 (column 1, Panel A) shows that side payments in both T_{supv} and T_{shared} are inexistent for the supervisors who observe worker output poorly. In contrast, side payment are positive (even though limited) for the remaining supervisors, who can better observe worker output (column 1, Panel B). These heterogeneous effects are robust to controlling for correlates of the observability of output interacted with the treatment dummies (column 2).

Overall, these results are consistent with side payments being larger when worker output is more observable and hence when contractual frictions are likely weaker. This result provides evidence on the likely importance of contractual frictions in preventing transfers from supervisors to workers. Importantly, output observability seems to be limited for most supervisors, which could make contracting difficult even for supervisors who are in the upper part of our proxy measure of observability.[34]

Results for workers with better outside options than their supervisor

A second piece of evidence pointing to the presence of contractual frictions comes from the analysis of worker-supervisor pairs where the worker has a better outside option than her supervisor, as proxied by the worker having a higher hourly wage

[34]We do not observe any heterogeneity in side payments with respect to whether the worker is a friend or family member of the supervisor (Table 1.A.13). This suggests that relational contracts have limited ability to attenuate contractual frictions within our organization. This is in contrast with a number of papers showing that relational contracts attenuate frictions *across* organizations (McMillan and Woodruff, 1999; Macchiavello and Morjaria, 2015; Adhvaryu et al., 2020; Macchiavello and Morjaria, 2021).

from their second job than her supervisor. In these pairs, the worker is likely to exert less effort than the supervisor would find optimal, and we thus expect the supervisor to have strong reasons to offer a sizable side payment to the worker.[35] Yet, Table 1.A.12 shows that, even within that sample, side payments are limited and visits are not equalized across treatments. This points to the presence of additional constraints to side payments, such as contractual frictions.

1.5.3 Alternative Mechanisms

The previous section provides empirical support for our theoretical framework, in which two-sided incentives outperform one-sided incentives due to the presence of *both* effort complementarities and limited side payments. This section provides evidence against two alternative mechanisms that are not considered in our model, but could explain why two-sided incentives outperform one-sided incentives: inequality aversion and a kink in agents' utility or cost functions.

Inequality aversion The ineffectiveness of the one-sided incentives treatments could be explained by aversion to pay inequality. For example, in the supervisor incentives treatment, the health workers may think that it is unfair that the supervisor earns money for services provided by the worker, while the worker herself does not earn anything. Similarly, the supervisor may think that it is unfair that she is not paid any incentive in the worker incentives treatment. If this was the case, then one-sided incentives may *reduce* the effort of the non-incentivized person, while raising the effort of the incentivized one. This could, in turn, explain why one-sided incentives are outperformed by two-sided incentives.

[35] In our theoretical framework, it is natural to think of outside options as a key driver of parameters b_1 and b_2, since outside options change the extent to which agents are concerned about losing their job due to under-performance. Agents with strong outside options will have a low value of parameter b and will, all else equal, exert less effort.

We provide three pieces of evidence against this mechanism. First, recall from Section 1.2.2 that health workers were not told about the introduction of supervisor incentives and few seem to have learned it from the supervisor: in T_{supv} (resp., T_{shared}), only 15% (resp., 20%) of workers reported knowing that their supervisor receives an incentive. Of these, only 2% (resp., 10%) were aware of the exact amount earned by the supervisor while the rest underestimated this amount. Second, Table 1.A.14 shows that there is no evidence for health workers in T_{supv} and supervisors in T_{worker} to be less satisfied with their payment, the organization, or their job in general compared to the control group. If inequality aversion or fairness concerns were the main mechanism driving our results we would instead expect the non-incentivized health workers in T_{supv} and the non-incentivized supervisors in T_{worker} to be less satisfied than workers and supervisors in the control group.[36] Third, we observe that the supervisor's effort is higher (and not lower) in T_{worker} relative to the control group, which cannot be reconciled with supervisors being demotivated by workers receiving incentives. All in all, these three pieces of evidence make it unlikely that inequality aversion alone drives our results.

Finally, we note that the absence of changes in satisfaction with the job and the organization is also inconsistent with the hypothesis that agents increase effort in the shared incentives treatment due to positive reciprocity. Under this story, any incentive payment would elicit a positive effort response that does not depend on the amount of the incentive paid. However, it is unlikely that reciprocal agents would increase effort, but not report higher satisfaction with the organization.

Sharp non-linearities in utility, cost or production functions We would also expect two-sided incentives to be more effective than one-sided incentives

[36]Workers in all treatments are equally likely to find the work environment competitive or to self-identify with their job: columns (7) and (8) in Table 1.A.14. We also find no differential treatment effects on visits depending on the workers' level of inequality aversion, estimated with a set hypothetical questions (columns 9 to 11).

if the returns to offering a piece rate above 1,000 SSL were low due to strong non-linearities in agents' utility, cost or production functions. In this section, we provide evidence against these strong non-linearities explaining our results.

Intuitively, in the absence of effort complementarities, shared incentives could outperform the one-sided incentives if there was of a sharp discontinuity in agents' utility function for incentive payments above 1,000 SLL. In this case, one-sided incentives would fail to motivate either of the two agents substantially more than the shared incentives treatment. Moreover, such a discontinuity in the utility function of both supervisors and workers would explain why the supervisors provide the same amount of effort in T_{supv} and T_{shared} (as shown in Table 1.5), and why the shared incentives treatment leads to more visits relative to the other treatments.

To investigate this, Panel A of Figure 1.A.5 displays non-parametric plots of the treatment effects on output and supervisor effort by worker and supervisor wealth score (a proxy for background utility). In the presence of strong non-linearities in utility, treatment effects would decline steeply in wealth, at least for some range of the wealth distribution. The figure shows instead that the treatment effects are fairly stable over the whole wealth score distribution (if anything supervisor effort appears to slightly increase with supervisor wealth). This is not surprising since even the wealthiest workers and supervisors in our sample are fairly poor, and doubling the incentive from 1,000 to 2,000 SLL can boost their income substantially.[37]

Alternatively, there may be a similar discontinuity in the cost function. Here, the marginal cost of effort would need to rise sharply at the level of efforts agents provide for a 1,000 SLL incentive. This would be the case if, for example, the distance of households from the health worker had a bimodal distribution, with

[37]Health workers earn an average of 1,443 SLL per hour in their alternative occupation. This is low relative to the amount they can earn for providing one household visit (1,000 or 2,000 SLL for a 15 minutes visit).

a group of near-by households that can be reached at low cost, and a group of far-away households that can only be reached at a cost above 2,000 SLL per visit. Again, this is a very unlikely scenario as it requires a sharp convexity in the workers' and supervisor's cost of effort around the 1,000 SLL cutoff.

Panel B of Figure 1.A.5 presents non-parametric plots of impacts on output and supervisor effort over the distribution of household-worker distance (a proxy for the worker's cost of visiting a household) and worker-supervisor distance (a proxy for the supervisor's cost of training/monitoring a health worker). Again, we do not find evidence of strong non-linearities.

A last possibility is that the results are explained by a discontinuity in the production function, such that the return to the worker effort increases with the effort of the supervisor only up to a threshold which coincides with the effort level exerted by the supervisor in the shared incentives treatment. In contrast with this story, Figure 1.A.6 shows that the non-parametric relationship between realized visits and supervisor effort is positive and close to linear.

1.6 Structural Model

In this section, we use the exogenous variation generated by the interventions to structurally estimate the model presented in Section 1.3, allowing for worker and supervisor-specific costs and benefits (see Appendix 1.D.7). First, we present our identification and estimation strategy. We then discuss the fit of the empirical and simulated moments. Finally, we present parameter estimates, and conclude with a set of counterfactual policy exercises.

1.6.1 Identification and Estimation

Our main objective is to estimate the following parameters of the model: complementarity γ, contractual friction z, the two cost of effort c_1 and c_2, the baseline incentives b_1 and b_2, and the production function parameter α. We calibrate z with a regression exercise that is described below. We jointly identify the the remaining six parameters using eight empirical moments, i.e., the mean of output (household visits) and the mean of supervisor effort in the four experimental conditions.[38] Intuitively, the moments capturing supervisor effort are informative about the cost and benefit parameters of the supervisor. Conditional on those parameters, the moments capturing output are informative about the cost and benefit of the worker, the complementarity of effort, and the parameter α.

We calibrate contractual frictions by using data on side payments. In particular, our model shows that $s = k - \frac{z+1}{2z}mp$. This suggests that the slope of a regression line of side payments s on mp - the product of the piece rate times the share of the piece rate offered to the worker - is informative of the size of contractual frictions z. When there are no frictions ($z = 1$), the slope of the regression line is 1. As frictions grow, the slope drops below 1 and approaches 0.5 from above. This result is intuitive: the stronger the frictions, the less responsive to p the side payment.[39]

[38] In our model there is no individual heterogeneity and so we only rely on empirical moments capturing mean outcomes. The specific measure of visits we use for the structural analysis is total visits per month. We obtain this by multiplying the number of visits per month per surveyed household by the number of households served by the health worker (as reported in Table 1.A.2, column 2). The measure of supervisor effort we use is the fraction of household visits in which the health worker was accompanied by the supervisor (as reported in Table 1.5, column 1). Also, note that we do not have good data on worker effort, since it is hard to obtain a clean measure of worker effort that is empirically distinct from output (household visits). We thus do not use any moment describing worker effort.

[39] We note two features of this calibration exercise. First, this exercise does not rely on the information on the absolute level of side payments which is contained in the intercept of the regression line, as this is likely to be observed with noise due to misreporting and poor memory. This is also a key reason why we calibrate the friction before the main structural estimation procedure. Second, we proxy s, the side payment offered, which we do not observe in the data, with the side payment paid, which we observe in the data.

To make the model more realistic, we introduce under-reporting by assuming that, for each completed visit, a shock that prevents the worker from reporting the visit occurs with probability $(1-q)$. This shock occurs after efforts have been exerted, and so its realization is not factored into effort decisions. However, agents know that a shock may occur and hence expect the value of the piece rate to be $m * q$. Except for this change in the expected value of the piece rate, the model remains unchanged. In our headline results, we assume conservatively that supervisors form an expectation about q using the average reporting rate across the experimental conditions. We then show robustness to assuming instead that supervisors realize that q varies with p, and thus that they form separate expectations about the reporting rate in each treatment.

To estimate the model we use a classical minimum distance estimator (Wooldridge, 2010). We save the empirical moments in a vector $\boldsymbol{m}$. For a parameter vector $\boldsymbol{\theta}$, we solve the model and calculate the simulated moments $\boldsymbol{m}_S(\boldsymbol{\theta})$. We update $\boldsymbol{\theta}$ in order to solve:

$$\hat{\boldsymbol{\theta}} = \min_{\boldsymbol{\theta}} \left[\boldsymbol{m}_S(\boldsymbol{\theta}) - \boldsymbol{m}\right]' \cdot J(\boldsymbol{m})^{-1} \cdot \left[\boldsymbol{m}_S(\boldsymbol{\theta}) - \boldsymbol{m}\right]. \tag{1.6}$$

$J(\boldsymbol{m})$ is a diagonal matrix that contains the variance of each moment, ensuring that more precisely estimated moments get a greater weight in estimation. We calculate $J(\boldsymbol{m})$ using a bootstrap with 1,000 replications. Table 1.7 presents our main structural results and Table 1.8 describes the empirical fit of the simulated moments.

The estimated model fits the empirical moments tightly: it matches both the moments related to supervisor effort and those related to household visits. Crucially, the estimated model is able to reproduce the key result that visits are maximized by the shared incentives treatment.

1.6.2 Parameter Estimates

Our structural estimates show that worker and supervisor effort are strongly complementary, and that contracting through side payments is very costly (Table 1.7). The estimated complementarity parameter γ determines a substantial increase in the marginal product of worker effort. Compared to a setting where $\gamma = 0$, the number of household visits generated by a unit of worker effort is 82% larger when the supervisor exerts the control level of effort, and 116% larger when the supervisor exerts the shared incentives level of effort. Supervisor effort thus plays a key role in enabling the worker to carry out household visits, and this results in a strong strategic complementarity between the efforts of the two agents.

The calibrated value of parameter z implies that side payments are 45 percent more costly due to contractual frictions. This constitutes a strong disincentive to offering side transfers, though we are not aware of other estimates of contractual frictions that we can use as a benchmark. A further disincentive against side transfers comes from the fact that the baseline incentive of the supervisor to exert effort (b_2) is lower than that of the worker (b_1). This is not surprising, since her role is probably harder to monitor and incentivize. Low supervisor motivation also suggests that reforms that target contractual frictions without also addressing supervisor motivation risk to backfire, as the supervisor may not necessarily use the greater ability to influence the worker in a way that is consistent with the objectives of the organization.

We also find that the supervisor has a high unit cost of effort (c_2). As a result, interventions that fail to incentivize the supervisor may be ineffective: the contribution of the supervisor is key to ensure the worker can be productive, but, absent additional incentive, the supervisor will underprovide her key support to the worker.

These results are robust to changing assumptions about the expected reporting rate, as shown in Tables 1.A.15 and 1.A.16. When we assume that supervisors have correct expectations about the reporting rate in each treatment group, we estimate very similar levels of effort complementarity (the worker's marginal product increases to 83% in the control group and to 117% in the group incentive condition) and an extremely high contractual friction ($z = 11.74$). This confirms that our core results on the importance of effort complementarity and contractual frictions do not depend on the specific assumption we make on reporting rate expectations.

1.6.3 Counterfactual Policies

We conduct three counterfactual policy experiments that explore, in turn, how to optimally share the incentive between the two agents, how the optimal incentive changes as key structural parameters vary, and the impact of an alternative policy that directly incentivizes effort.

We find that offering an equal share of the incentive to the worker and the supervisor is almost optimal. In Figure 1.3, we show that, in order to maximize household visits, the worker should be offered 59% of the overall incentive, which is very close to the equal share that we offered in the shared incentives treatment. In other words, given the strong complementarity and large contractual frictions we have estimated, the optimal incentive scheme is one that rewards both agents with a similar payment.[40]

This result, however, depends strongly on the strength of the complementarity between worker and supervisor effort. We illustrate this point with our second counterfactual experiment in Figure 1.4. Here, we plot the optimal share of the

[40]This is a similar exercise than the one done in the simulations shown in Figure 1.1, but here we are using the estimated parameters from the model to simulate the optimal incentive split between the layers.

Figure 1.3: Optimal Incentive p^*

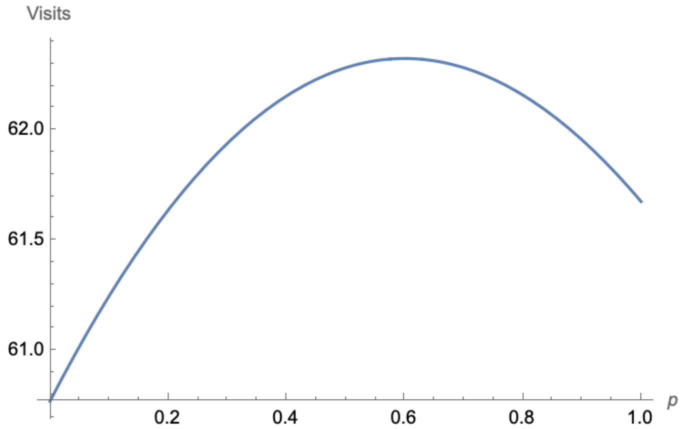

incentive offered to the worker (p^*) for different levels of complementarity. A key result that emerges from this analysis is that, as the complementarity parameter shrinks, the optimal incentive offered to the worker increases substantially. Quantitatively, if the complementarity parameter was 10 percent *lower* than what we estimate, the optimal incentive would give 80 percent of the piece rate to the worker. If the complementarity parameter was instead 10 percent *higher* than what we estimate, the optimal incentive would give 60 percent of the piece rate to the supervisor.

Figure 1.4: Optimal Incentive p^* by Complementarity γ

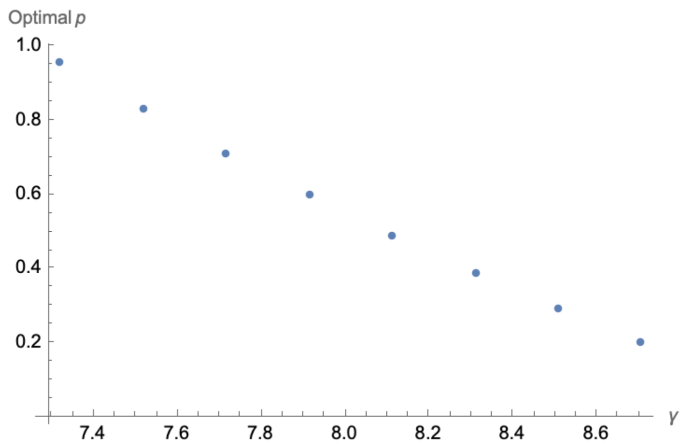

Thus, these results suggest that in organizations in which effort complementary is weaker than in our settings — e.g., settings in which the role of the supervisor is limited to monitoring, distributing tasks or making personnel decisions, but not to train and advise workers - the optimal split is one that allocates significantly more to the worker. And in organizations in which effort complementarity is stronger - e.g. organizations where supervisors are closely involved in production - the optimal incentive scheme allocates the largest share of the piece rate to the supervisor.

Our final key result highlights that tying incentives to joint output is more effective than directly incentivizing effort (e.g., incentivizing supervisors on the amount of supervision and training, and incentivizing health workers on the number of times they attempt to approach a household, regardless of whether this results in a visit or not). In Figure 1.5, we compare the maximum number of visits that are generated through (i) a scheme that equally shares a payment of 2,000 SLL per visit between the worker and the supervisor, and (ii) a scheme of the same cost that optimally offers incentives directly tied to individual effort.[41]

What emerges is that, at the current level of complementarity, incentivizing output through an equally-shared piece rate generates 18% more visits that optimally incentivizing effort, for the same cost. This is because, when efforts are highly complementary, output incentives implicitly help agents internalize the positive external effect that their effort has on the other player. This makes output incentives particularly effective.

[41]In this comparison, we assume that effort can be observed and is perfectly predictive of output. Hence, we abstract from issues related to asymmetric information, which may decrease the effectiveness of both incentive schemes. In the effort incentive case, since effort can be observed, the payoff to the worker becomes $\pi_1 = e_1 * (b_1 + mp)$ and the payoff to the supervisor $\pi_2 = e_2 * (b_2 + (1 - m) * p)$. In this model, the supervisor always offers zero side transfer since her reward only depends on her own effort.

Figure 1.5: An Alternative Policy that Targets Effort

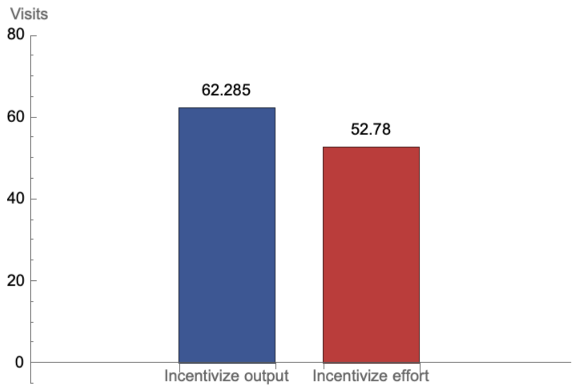

1.7 Conclusion

In this paper, we provide novel evidence on the optimal structure of performance incentives in a hierarchical organization. We do so by designing a field experiment in collaboration with the national Community Health Program in Sierra Leone, which is structured in two layers: frontline health workers and supervisors. The experiment creates random variation in the recipient of a new incentive scheme that rewards household health visits, while holding the total payout per visit constant. The visit piece-rate is offered either entirely to the worker, entirely to the supervisor, or is split equally between the two agents. We find that all treatments increase household visits relative to the control condition, and that the shared incentives treatment generates the largest increase in visits.

Before releasing the results of the field experiment, we invited social scientists to forecast them on the online Social Science Prediction Platform. The majority (52%) of survey participants forecasted that one-sided worker incentives would maximize health visits.[42] This is not surprising, since worker incentives have received much attention in the existing empirical literature, as we discussed in the

[42] 28% forecasted shared incentives to maximize visits, 4% forecasted supervisor incentives to be the most effective, and 18% forecasted that all three treatments would have the same effect. See Appendix 1.E for more details about the prediction survey.

Introduction. In line with most of this literature, this paper shows that incentivizing frontline health workers increases output relative to a control group without incentives. However, our results also indicate that sharing the same piece rate equally between the worker and her superior generates a treatment effect that is almost two thirds larger.

The key finding that output is highly responsive to the allocation of the incentive has a number of important implications. First, when incentives are not redistributed through side transfers, organizations need to fine-tune the allocation of payments to maximize output. Our empirical results document a context where the complementarity between worker and supervisor effort is strong, and hence where sharing a piece-rate across the layers of the organization has a large impact on output. We would expect weaker complementarities, and hence a less pronounced impact of shared incentives, in contexts in which the role of the supervisor is more focused on monitoring. This emphasizes the importance of calibrating incentives in each organization based on effort complementarities.

Second, organizations may consider introducing reforms that facilitate or limit side-transfers across layers. Facilitating side-transfers could be attractive in contexts where the organization has an imperfect understanding of the production function, and may prefer to rely on supervisors (who have better local knowledge) to redistribute incentives to bottom-tier workers at their discretion. Such delegation is however optimal only if the interests of the supervisors are aligned with those of the organization. If supervisors are poorly motivated to generate output - as the structural estimates suggests for our context - expanding the scope for transfers from supervisors to workers may instead be detrimental to the performance of the organization. Understanding how to facilitate side payments across layers of an organization and assessing its effect on organizational performance is a great avenue for future research.

Finally, to introduce an incentive scheme such as the one considered in this paper, organizations need to be able to reliably measure output. In our setting, we pay the incentive on the basis of workers' self reports, while performing extensive checks to prevent over-reporting. As digital technologies improve, the costs of monitoring worker self-reports will likely decrease, enabling more organizations to set up incentive schemes like ours (Kelley et al., 2021; Dodge et al., 2021; Adhvaryu et al., 2022). An interesting additional implication is that when managers have access to new monitoring technologies, contracting frictions with front-line workers may decrease, as our results suggest that these frictions stem at least in part from the managers' difficulty of observing workers' output. As argued above, increasing the scope of side transfers may or may not be in the organization's best interest. More work is needed studying how to allocate access to information in organizations fairly and efficiently.

Table 1.1: Summary Statistics and Balance Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Sample of villages:	All		Control		Worker Incentives Treatment		Supervisor Incentives Treatment		Shared Incentives Treatment		F-test of joint significance (p-value)
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	
A. Characteristics of the supervisors (N=372)											
Male = {0, 1}	0.919	0.273	0.925	0.265	0.925	0.265	0.914	0.282	0.914	0.282	0.988
Age (in years)	37.84	8.856	37.91	9.329	37.46	7.869	36.85	8.690	39.13	9.433	0.316
Completed primary education = {0, 1}	0.739	0.440	0.763	0.427	0.731	0.446	0.785	0.413	0.677	0.470	0.405
Completed secondary education = {0, 1}	0.253	0.435	0.226	0.420	0.269	0.446	0.215	0.413	0.301	0.461	0.533
Wealth score (0 to 8)	3.013	1.227	3.097	1.269	2.914	1.222	2.914	1.239	3.129	1.182	0.507
Number of health workers responsible for	7.984	2.861	7.559	2.799	8.355	2.831	8.011	2.902	8.011	2.899	0.289
B. Characteristics of the health workers (N=2,970)											
Male = {0, 1}	0.708	0.455	0.727	0.446	0.721	0.449	0.710	0.454	0.675	0.469	0.407
Age (in years)	37.12	11.47	35.95	11.14	37.79	11.72	37.48	11.72	37.17	11.21	0.062
Completed primary education = {0, 1}	0.697	0.460	0.727	0.446	0.694	0.461	0.703	0.457	0.666	0.472	0.301
Completed secondary education = {0, 1}	0.077	0.267	0.070	0.255	0.076	0.265	0.078	0.268	0.085	0.278	0.867
Wealth score (0 to 8)	2.454	1.167	2.430	1.231	2.400	1.116	2.438	1.120	2.550	1.199	0.273
Number of households responsible for	55.19	78.59	62.72	120.2	54.08	62.92	53.16	56.37	51.26	60.24	0.375
Distance to supervisor (in km)	3.415	2.945	3.267	2.887	3.815	3.610	3.107	2.141	3.447	2.895	0.190
C. Characteristics of the female household respondent, aggregated to village level (N=2,970)											
Age (in years)	27.79	4.576	28.13	4.741	27.69	4.410	27.56	4.572	27.84	4.586	0.266
Completed primary education = {0, 1}	0.248	0.269	0.259	0.268	0.225	0.261	0.261	0.272	0.247	0.273	0.203
Completed secondary education = {0, 1}	0.035	0.119	0.039	0.126	0.033	0.118	0.036	0.118	0.033	0.116	0.912
Wealth score (0 to 8)	1.103	0.872	1.199	1.021	1.044	0.822	1.117	0.843	1.062	0.790	0.111
Distance to health worker (in km)	1.433	2.630	1.189	2.124	1.591	2.575	1.438	2.894	1.483	2.785	0.506
D. Characteristics of the village (N=372)											
Accessible road to health facility= {0, 1}	0.766	0.424	0.778	0.416	0.762	0.426	0.775	0.418	0.747	0.435	0.797
Phone network available	0.838	0.368	0.824	0.381	0.845	0.362	0.862	0.345	0.821	0.384	0.490
E. Services provided by the health facility per month (N=372)											
Number of pregnant women services	47.71	45.80	43.76	42.98	51.22	57.17	48.43	37.10	47.34	44.10	0.769
Number of institutional births	13.44	8.266	12.58	6.087	13.67	7.814	13.38	6.845	14.13	11.33	0.509
Number of fully immunized infants	11.41	10.75	10.88	10.10	11.92	12.58	10.67	7.060	12.14	12.39	0.684
Number of malaria cases treated	45.89	32.03	42.29	26.88	51.44	38.88	46.63	31.52	43.23	29.38	0.286
Number of diarrhea cases treated	20.45	17.03	19.62	13.25	21.81	19.47	20.58	21.58	19.76	12.02	0.809

Notes: Each row states the sample mean and standard deviation of a variable, and by treatment group. The last column reports the p-value from the F-test of joint significance of the treatment dummies. This is calculated from a regression of each variable on the 3 treatment dummies, controlling for the stratification variables and using standard errors clustered at the PHU level in worker/village level regressions or robust standard errors in PHU/supervisor level regressions.

Table 1.2: Household Visits

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Household visits provided by the health worker in the past 6 months						
Dep. Var.	Number of visits	% households visited	Number of visit types	Average visit length	Average visit length (conditional on at least one visit)	Number of health topics discussed per visit	% households who trust the health worker as a health provider
Worker incentives	2.090***	0.072***	0.250***	2.024**	0.603	0.164	0.037
	(0.558)	(0.025)	(0.094)	(0.941)	(1.423)	(0.127)	(0.023)
Supervisor incentives	2.145***	0.082***	0.325***	1.933**	-0.070	0.173	0.031
	(0.501)	(0.025)	(0.100)	(0.926)	(1.409)	(0.130)	(0.024)
Shared incentives	3.356***	0.127***	0.565***	4.134***	1.590	0.528***	0.071***
	(0.492)	(0.023)	(0.092)	(0.927)	(1.290)	(0.134)	(0.024)
Unit	Worker	Worker	Worker	Worker	Worker	Worker	Worker
Observations	2,926	2,926	2,926	2,926	1,803	2,926	2,926
Mean dep. var.	7.296	0.709	1.745	14.388	23.404	2.248	0.745
Mean dep. var. in Control	5.334	0.637	1.448	12.324	22.736	2.015	0.707
p-value Worker = Supervisor	0.932	0.710	0.492	0.927	0.635	0.946	0.808
p-value Supervisor = Shared	0.038	0.060	0.026	0.024	0.196	0.017	0.102
p-value Worker = Shared	0.046	0.026	0.002	0.033	0.451	0.013	0.147

Notes: All regressions include stratification variables. Standard errors are clustered at the PHU level. In col. (4), we assign an average visit length of zero to households who were never visited by the worker. *** p<0.01, ** p<0.05, * p<0.1

Table 1.3: Health Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Pre- and post-natal care in the past 2 years						Diseases incidence			
Dep. Var.	Index (cols. 2-6)	% women who received at least 4 ante-natal visits before birth	% women with institutional birth	% women who received post-natal visit within 2 days of birth	% women with at least 6 months of breastfeeding	% households with up-to-date infant vaccination	Index (cols. 8-10)	% children under-5 who had fever	% children under-5 who had cough	% children under-5 who had diarrhea
Worker incentives	0.029	0.017	0.022	0.007	-0.002	-0.008	0.010	-0.028	0.016	0.005
	(0.028)	(0.024)	(0.021)	(0.026)	(0.025)	(0.018)	(0.032)	(0.022)	(0.012)	(0.007)
Supervisor incentives	0.042	0.032	0.035*	-0.022	0.016	0.015	-0.028	-0.014	-0.005	-0.005
	(0.029)	(0.026)	(0.019)	(0.024)	(0.025)	(0.019)	(0.032)	(0.028)	(0.012)	(0.005)
Shared incentives	0.092***	0.058**	0.036*	0.017	0.040	0.025	-0.053**	-0.058***	-0.007	-0.001
	(0.028)	(0.025)	(0.019)	(0.027)	(0.025)	(0.019)	(0.026)	(0.022)	(0.011)	(0.005)
Unit	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker
Observations	2,499	2,499	2,499	2,499	2,499	2,499	2,826	2,823	2,825	2,826
Mean dep. var.	-0.006	0.778	0.868	0.305	0.666	0.230	-0.009	0.183	0.072	0.016
Mean dep. var. in Control	-0.048	0.750	0.845	0.303	0.652	0.222	0.009	0.208	0.071	0.017
p-value Worker = Supervisor	0.656	0.509	0.491	0.258	0.439	0.223	0.273	0.580	0.088	0.129
p-value Supervisor = Shared	0.077	0.243	0.963	0.138	0.330	0.630	0.397	0.086	0.868	0.343
p-value Worker = Shared	0.024	0.052	0.469	0.703	0.077	0.091	0.025	0.105	0.056	0.360

Notes: The index in col. (1) [resp., col. (7)] estimates an equally weighted average of the z-scores of variables in cols. (2)-(6) [resp., cols. (8)-(10)]. The sample in cols. (1)-(6) is restricted to households with a woman who gave birth in the past year. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.4: Reporting and Incentives Payouts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Reporting				Cost of the intervention (in 1,000 SLL)				
Dep. Var.	Number of reports per month	Reporting rate = number of reports/ number of visits	Over-reporting = {0, 1} [number of reports > number of visits]	Under-reporting = {0, 1} [number of reports < number of visits]	Incentives payment per health worker	Total incentives payment across all health workers	Incentives payment per supervisor	Total incentives payments in the PHU (cols. 6+7)	Number of visits per worker for each 2,000 SLL spent
Worker incentives	7.182***	0.222***	0.046**	-0.002	16.054***	132.903***	-1.310	131.593***	-
	(0.759)	(0.043)	(0.018)	(0.024)	(1.352)	(12.000)	(1.832)	(11.962)	-
Supervisor incentives	2.181***	0.058*	-0.006	0.041*	-0.047	-1.172	55.280***	54.108***	3.126
	(0.586)	(0.034)	(0.015)	(0.021)	(0.185)	(2.631)	(8.095)	(8.377)	(4.151)
Shared incentives	4.757***	0.091**	0.009	0.046**	5.584***	46.857***	47.097***	93.953***	9.574**
	(0.829)	(0.036)	(0.015)	(0.021)	(0.691)	(6.680)	(6.534)	(13.164)	(4.826)
Unit	Worker	Worker	Worker	Worker	PHU	PHU	PHU	PHU	PHU
Observations	2,970	2,624	2,926	2,926	372	372	372	372	279
Mean dep. var.	5.213	0.177	0.069	0.862	5.431	45.535	25.858	71.392	10.787
Mean dep. var. in Control	1.525	0.078	0.055	0.840	0.000	0.000	0.000	0.000	
Mean dep. var. in Worker T.									6.508
p-value Worker = Supervisor	<0.001	<0.001	0.004	0.059	<0.001	<0.001	<0.001	<0.001	-
p-value Supervisor = Shared	0.008	0.398	0.333	0.802	<0.001	<0.001	0.428	0.009	0.223
p-value Worker = Shared	0.026	0.005	0.050	0.035	<0.001	<0.001	<0.001	0.032	-

Notes: In col. (2), the reporting rate is measured as the number of reports (col. 1) divided by the number of household visits. In cols. (3)-(4), under (over) reporting is an indicator for whether the number of reports is lower (higher) than the number of household visits. Col. (9) excludes the control group and the omitted group is the worker incentives treatment. The outcome variable in col. (9) calculates the number of visits performed by the average worker in the PHU minus the mean number of visits in the control, divided by the total incentive payout in the PHU (col. 8). We winsorize the top and bottom 1% of the outcome variable due to the presence of outliers and inpute the maximum value of the outcome variable for the few PHUs in which the total incentives payout is zero. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.5: Supervisor Effort

	(1)	(2)	(3)
	In-the-field training	*One-to-one meetings*	*General training*
Dep. Var.	% accompanied household visits	% health workers visited in the past 6 months	Supervisor organized training in the past month = {0, 1}
Worker incentives	0.030	-0.050	0.004
	(0.022)	(0.137)	(0.005)
Supervisor incentives	0.057**	-0.041	0.006
	(0.023)	(0.137)	(0.005)
Shared incentives	0.062***	-0.043	0.003
	(0.021)	(0.139)	(0.005)
Unit	Worker	Worker	Worker
Observations	2,919	2,833	2,864
Mean dep. var.	0.204	1.375	0.994
Mean dep. var. in Control	0.164	1.417	0.991
p-value Worker = Supervisor	0.293	0.950	0.443
p-value Supervisor = Shared	0.846	0.987	0.463
p-value Worker = Shared	0.181	0.963	0.916

Notes: Variable in col. (1) is reported by the household and aggregated to worker level. Variables in cols. (2)-(3) are reported by the health worker. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.6: Side Payments

	(1)	(2)	(3)	(4)	(5)
			Monthly transfer amount (in 1,000 SLL)		
Dep. Var.	Supervisor shared incentive with health worker = {0, 1}	Health worker shared incentive with supervisor = {0, 1}	… from supervisor to health worker	… from health worker to supervisor	Net transfer (col. 3-4)
Worker incentives	0.005	0.073***	0.110	0.151***	-0.042
	(0.016)	(0.014)	(0.090)	(0.056)	(0.077)
Supervisor incentives	0.183***	-0.001	0.702***	0.104**	0.598***
	(0.047)	(0.008)	(0.190)	(0.043)	(0.190)
Shared incentives	0.102***	0.041***	0.432***	0.084*	0.348**
	(0.039)	(0.015)	(0.158)	(0.043)	(0.164)
Unit	Worker	Worker	Worker	Worker	Worker
Observations	2,915	2,909	2,488	2,488	2,488
Mean dep. var.	0.084	0.049	0.308	0.101	0.207
Mean dep. var. in Control	0.011	0.019	0.000	0.016	-0.016
p-value Worker = Supervisor	<0.001	<0.001	0.004	0.484	0.001
p-value Supervisor = Shared	0.171	0.005	0.273	0.725	0.318
p-value Worker = Shared	0.013	0.100	0.068	0.325	0.026

Notes: All regressions include stratification variables. Standard errors clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.7: Parameter Estimates

	(1)
Complementarity γ	7.9
	(3.3)
Worker unit cost of effort c_1	2.1
	(1.2)
Supervisor unit cost of effort c_2	12542.0
	(15071.5)
Worker baseline incentive b_1	23.5
	(13.7)
Supervisor baseline incentive b_2	16.5
	(12.6)
α	1.5
	(0.7)
Calibrated friction z	1.45
Δ in marginal product of worker effort (shared incentive)	116%
Δ in marginal product of worker effort (control)	82%
Total worker cost of effort (control)	31.6
Total supervisor cost of effort (control)	1961.1

Notes: The first panel of the table shows parameter estimates obtained using minimum distance estimation. We use eight empirical moments: supervisor effort in each one of the four treatments, and number of visits per month in each one of the four experimental groups. Supervisor effort is proxied by the proportion of households that receive a visit where the worker is accompanied by the supervisor. Costs are expressed in thousand SLL. Boostrapped standard errors are reported in parenthesis (we bootstrap the estimation 500 times and truncate the estimated coefficients at the 99th percentile of the distribution). The second panel first shows the calibrated value of contractual frictions. Second, it shows some quantities implied by the parameter estimates. To calculate the change in the marginal product of worker effort we take the derivative of the production function with respect to worker effort (i) with $\gamma = 7.9$ and supervisor effort fixed at the level indicated in parenthesis, and (ii) with $\gamma = 0$. To calculate the total cost of an agent effort we multiply the unit cost of effort by the average effort exerted by the agent in the control group.

Table 1.8: Moment Fit

Moments	Targeted Real	Simulated
Supervisor effort in worker incentives group	0.198	0.205
Supervisor effort in supervisor incentives group	0.225	0.231
Supervisor effort in shared incentives group	0.228	0.221
Supervisor effort in control group	0.164	0.156
Output in worker incentives group	59.679	61.679
Output in supervisor incentives group	58.896	60.773
Output in shared incentives group	66.895	62.285
Output in control group	41.040	41.156
Value loss function	6.6	

Notes: The table shows the targeted empirical moments used for minimum distance estimation as well as the simulated moments.

Appendix

Appendix 1.A Appendix Figures and Tables

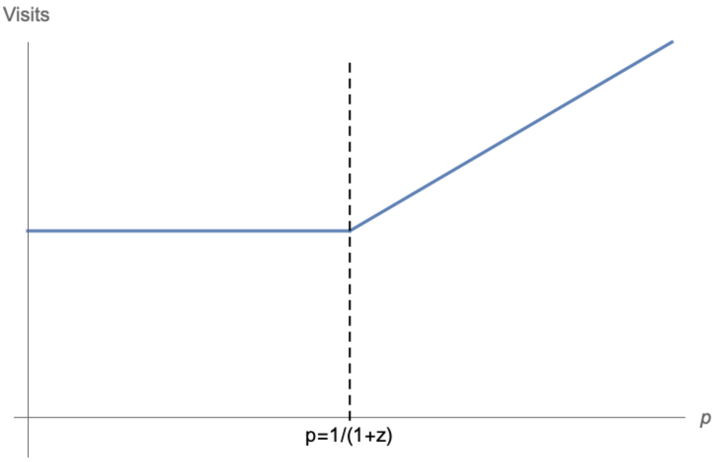

(a) Weak Effort Complementarities and No Contractual Frictions ($\gamma < t$, $z = 1$)

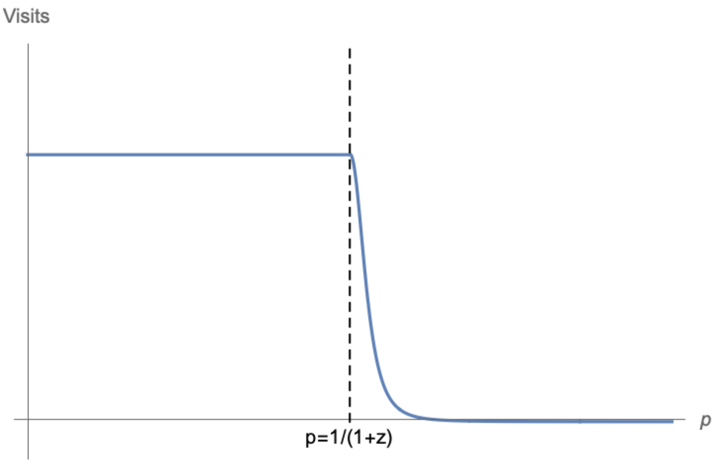

(b) Strong Effort Complementarities and No Contractual Frictions ($\gamma \geq t$, $z = 1$)

Figure 1.A.1: Optimal Incentives (Continued)

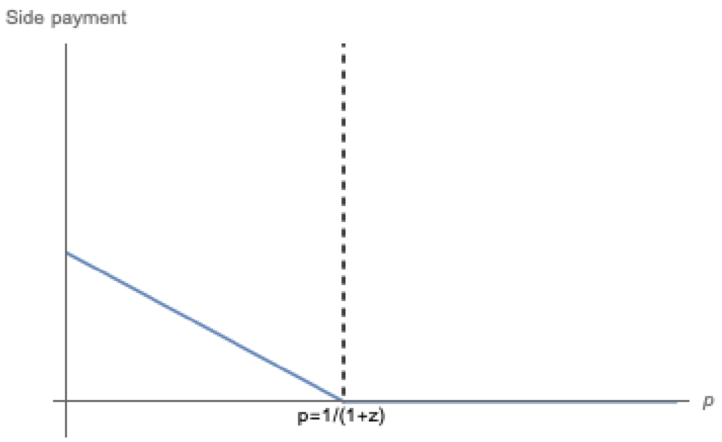

(a) Side Payment

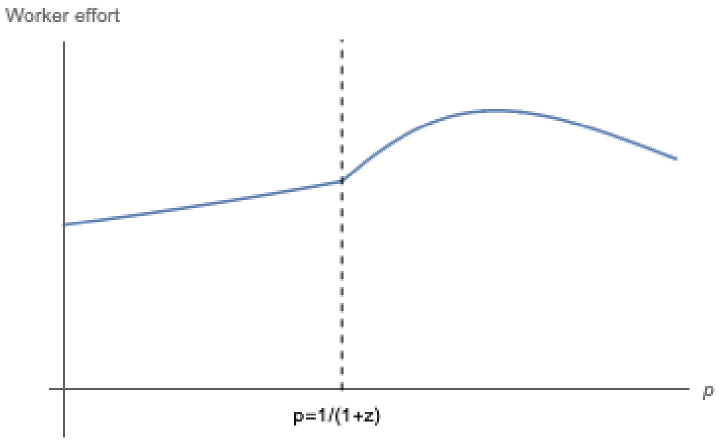

(b) Worker Effort

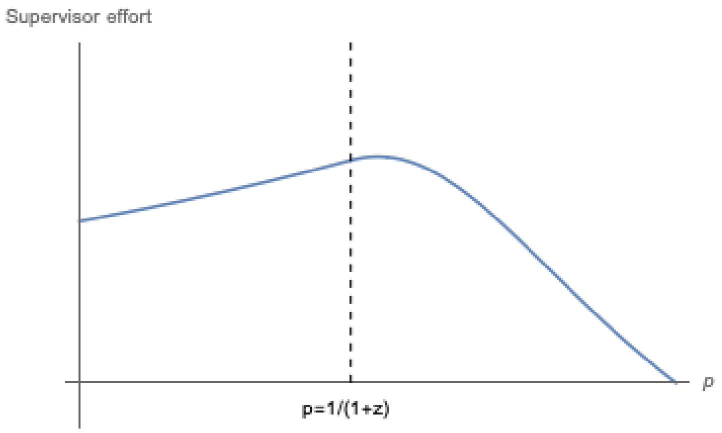

(c) Supervisor Effort

Figure 1.A.2: Side Payment and Efforts as a Function of the Share of the Incentive Offered to the Worker ($\gamma \geq t$, $z > 1$)

Figure 1.A.3: Time Evolution of SMS Reporting

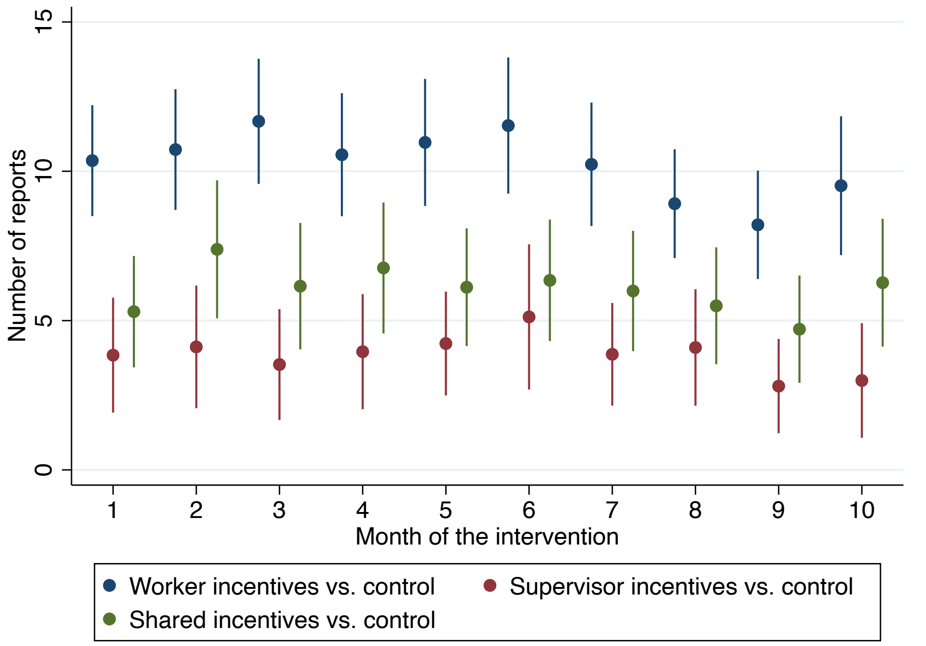

Notes: The figure plots the difference in the number of SMS reports between each treatment group and the control group. The coefficients are estimated from a regression of the number of SMS reports in each single month on the treatment dummies, controlling for the stratification variables and with standard errors clustered at the PHU level. Bars are 95% confidence intervals.

Figure 1.A.4: Mediation Analysis

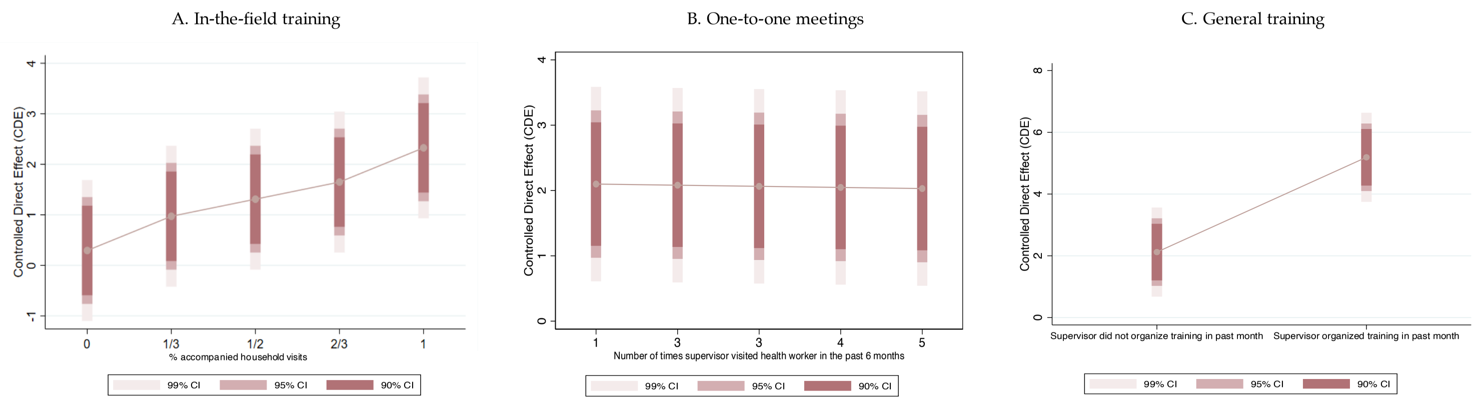

Notes: This figure plots the controlled direct effect (CDE) of the worker incentives treatment on the number of visits provided by a health worker for different values of supervisor's effort, as measured with in-the-field training (Panel A), one-to-one meetings (Panel B), general training (Panel C). Each figure is produced by following the steps outlined in Acharya et al. (2016). First, we regress the number of visits provided by a health worker on the worker incentives treatment, the mediator (supervisor's effort), and their interaction. From this, we obtain a de-mediated outcome, defined as the difference between actual visits and predicted visits based on the coefficients of all covariates (except the treatment) estimated at different levels of the mediator. Finally, we run a regression of the de-mediated outcome on the treatment and report the coefficients for different levels of the mediator.

Figure 1.A.5: Non-Parametric Estimates of Visits and Supervisor Effort by Wealth and Distance

A. Treatment Effects on Visits and Supervisor Effort by Wealth

B. Treatment Effects on Visits and Supervisor Effort by Distance

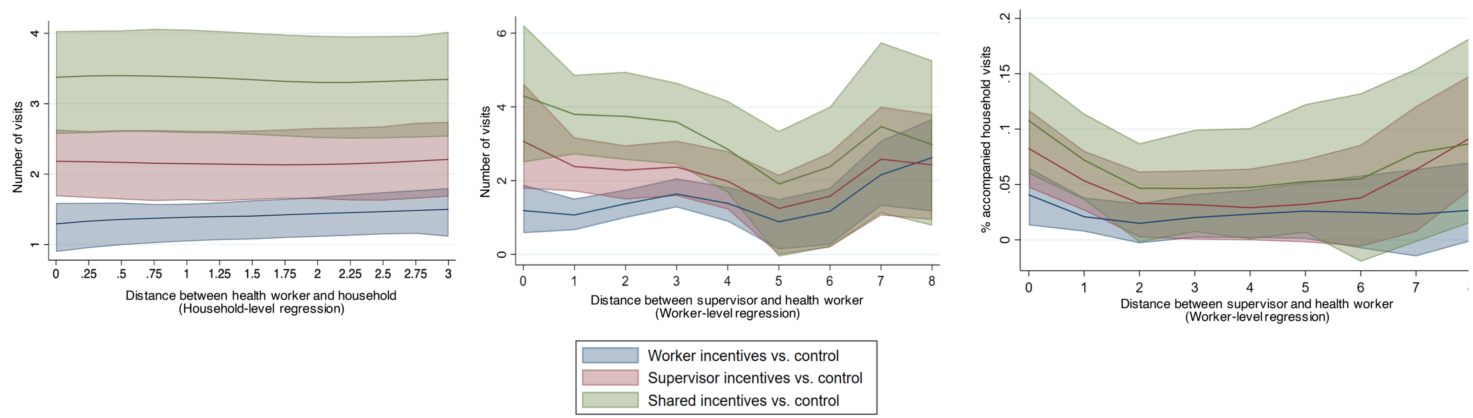

Notes: This figure plots non-parametric estimates of the treatment effects on the number of visits (left and middle figures) and on the supervisor effort index (right figures) by wealth score and by distance. Standard errors are bootstrapped for each value of the x-axis, with 100 repetitions and the re-sampling is with replacement. 95% confidence intervals presented in the figures. In Panel B, the "distance between health worker and household" is measured at the household level and the analysis is performed at the household level. In all the other figures, the analysis is performed at the health worker level.

Figure 1.A.6: Non-Parametric Estimates of Visits by Supervisor Effort

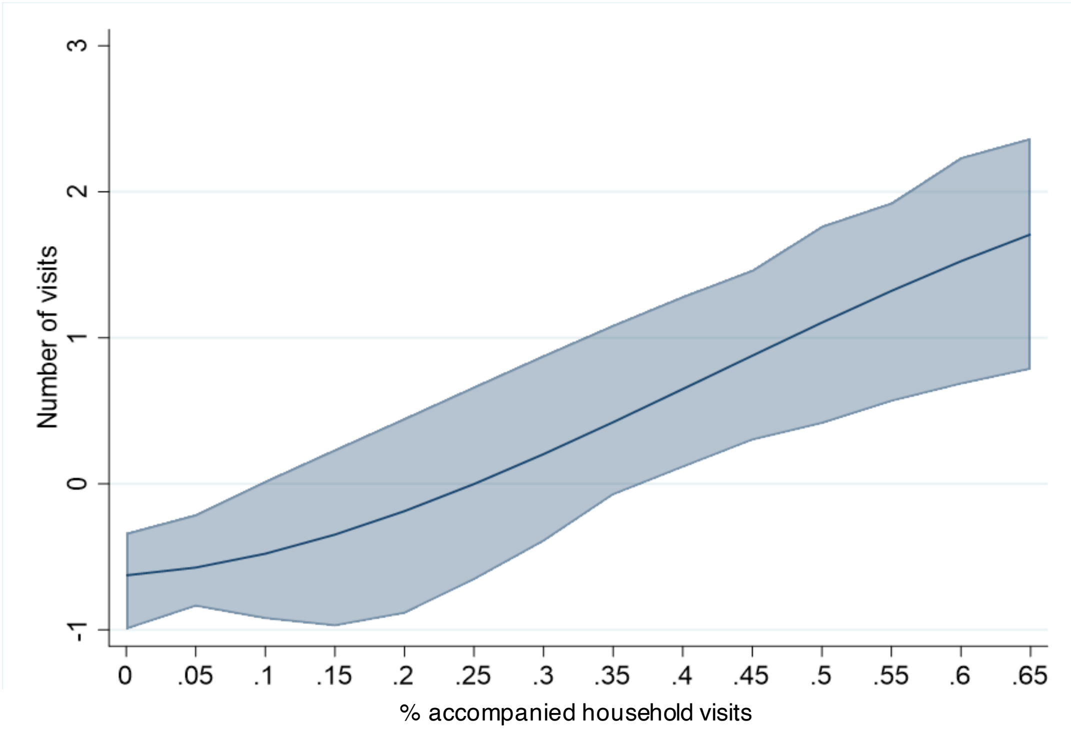

Notes: This figure plots non-parametric estimates of the number of visits provided by the health worker on the fraction of accompanied household visits (supervisor effort). Standard errors are bootstrapped for each value of the x-axis, with 100 repetitions and the re-sampling is with replacement. 95% confidence intervals presented in the figure.

Table 1.A.1: Balance Checks (Pairwise Treatment Comparisons)

	(1)	(2)	(3)	(4)	(5)	(6)
	P-values from Pairwise Treatment Comparisons					
	Worker Incentives = Supervisor Incentives	Worker Incentives = Shared Incentives	Supervisor Incentives = Shared Incentives	Worker Incentives = Control	Supervisor Incentives = Control	Shared Incentives = Control
A. Characteristics of the supervisors (N=372)						
Male = {0, 1}	0.823	0.869	0.957	0.923	0.750	0.796
Age (in years)	0.664	0.151	0.074	0.615	0.370	0.371
Completed primary education = {0, 1}	0.399	0.454	0.109	0.592	0.748	0.195
Completed secondary education = {0, 1}	0.395	0.671	0.199	0.473	0.883	0.249
Wealth score (0 to 8)	0.901	0.285	0.215	0.371	0.295	0.888
Number of health workers responsible for	0.375	0.450	0.904	0.054	0.304	0.253
B. Characteristics of the health workers (N= 2,970)						
Male = {0, 1}	0.912	0.218	0.170	0.678	0.749	0.102
Age (in years)	0.472	0.338	0.838	0.009	0.067	0.088
Completed primary education = {0, 1}	0.812	0.329	0.201	0.405	0.528	0.059
Completed secondary education = {0, 1}	0.944	0.708	0.738	0.666	0.590	0.397
Wealth score (0 to 8)	0.915	0.138	0.112	0.835	0.736	0.094
Number of households responsible for	0.532	0.353	0.711	0.291	0.138	0.096
Distance to supervisor (in km)	0.043	0.443	0.228	0.204	0.443	0.636
C. Characteristics of the female household respondent, aggregated to village level (N=2,970)						
Age (in years)	0.851	0.477	0.388	0.099	0.080	0.347
Completed primary education = {0, 1}	0.072	0.257	0.440	0.065	0.923	0.469
Completed secondary education = {0, 1}	0.924	0.776	0.712	0.669	0.755	0.470
Wealth score (0 to 8)	0.785	0.581	0.324	0.122	0.141	0.015
Distance to health worker (in km)	0.727	0.907	0.818	0.184	0.327	0.225
D. Characteristics of the village (N=372)						
Accessible road to health facility= {0, 1}	0.784	0.511	0.361	0.809	0.991	0.400
Phone network available	0.715	0.341	0.210	0.361	0.222	0.955
E. Services provided by local health facilities per month (N=372)						
Number of pregnant women services	0.669	0.539	0.811	0.311	0.467	0.637
Number of institutional births	0.749	0.740	0.565	0.229	0.432	0.206
Number of fully immunized infants	0.358	0.983	0.319	0.520	0.817	0.488
Number of malaria cases treated	0.345	0.102	0.458	0.076	0.368	0.860
Number of diarrhea cases treated	0.674	0.383	0.746	0.379	0.727	0.973

Notes: Each row presents p-values from pairwise treatment comparisons. These are calculated from a regression, where the variable is regressed on an indicator for worker, supervisor and shared incentives treatment. All regressions include stratification variables. Standard errors are clustered at the PHU level in worker/village level regressions and we use robust standard errors in PHU/supervisor level regressions.

Table 1.A.2: Household Visits (Aggregated to Entire Village)

	(1)	(2)	(3)
Dep. Var.	Total number of visits provided by the health worker to sampled households per month	Total number of visits provided by the health worker in the entire village per month	Total number of visits provided by all health workers in the PHU per month
Worker incentives	0.982***	17.725***	173.143***
	(0.245)	(4.429)	(40.189)
Supervisor incentives	0.998***	18.007***	154.653***
	(0.229)	(4.139)	(38.853)
Shared incentives	1.414***	25.505***	215.047***
	(0.214)	(3.862)	(35.299)
Unit	Worker	Worker	Supervisor
Observations	2,926	2,926	372
Mean dep. var.	3.154	56.906	448.434
Mean dep. var. in Control	2.275	41.040	303.529
p-value Worker = Supervisor	0.957	0.957	0.709
p-value Supervisor = Shared	0.110	0.110	0.184
p-value Worker = Shared	0.116	0.116	0.369

Notes: Col. (2) divides the total number of household visits provided by the health worker to sampled households per month (col. 1) by the share of households in the village sampled. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.A.3: Household Visits by Type

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. Var.	% households who received [visit type] from the health worker in the past 6 months							
	Natal-related visits				Disease-related visits			
Visit Type	Pregancy visit	Accompanied woman for birth to the health facility	Pre and post-natal visit	Index (cols. 1-3)	Routine visit	Treatment/ referrals of under-5 children	Follow-up visit of under-5 children	Index (cols. 5-7)
Worker incentives	0.037**	-0.005	0.027	0.069	0.068**	0.053**	0.042*	0.155***
	(0.016)	(0.007)	(0.017)	(0.049)	(0.033)	(0.025)	(0.022)	(0.060)
Supervisor incentives	0.027*	0.004	0.037*	0.092*	0.089***	0.071**	0.031	0.178***
	(0.016)	(0.008)	(0.019)	(0.051)	(0.030)	(0.028)	(0.024)	(0.061)
Shared incentives	0.064***	0.008	0.051***	0.168***	0.151***	0.111***	0.079***	0.324***
	(0.017)	(0.008)	(0.017)	(0.051)	(0.029)	(0.024)	(0.020)	(0.056)
Unit	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker
Observations	2,926	2,926	2,926	2,926	2,926	2,926	2,926	2,926
Mean dep. var.	0.179	0.041	0.132	0.000	0.517	0.504	0.201	0.000
Mean dep. var. in Control	0.145	0.038	0.103	-0.087	0.437	0.443	0.162	-0.171
p-value Worker = Supervisor	0.543	0.193	0.604	0.650	0.540	0.527	0.669	0.733
p-value Supervisor = Shared	0.037	0.660	0.455	0.154	0.046	0.143	0.046	0.024
p-value Worker = Shared	0.132	0.069	0.151	0.056	0.014	0.020	0.094	0.008

Notes: The index in col. (4) [resp., col. (8)] estimates an equally weighted average of the z-scores of variables in cols. (1)-(3) [resp., cols. (5)-(7)]. All regressions include stratification variables. Standard errors clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.4: Household Targeting

	(1)	(2)	(3)	(4)	(5)
Dep. Var.	Number of visits				
Definition of covariate X:	Household's wealth score (0 to 8)	Household's distance to health worker (in km)	Household's respondent is a family member of the health worker = {0, 1}	Household's respondent is a friend of the health worker = {0, 1}	Household received no visit accompanied by the supervisor = {0, 1}
Worker incentives	2.118***	2.002***	2.113***	2.080***	1.661*
	(0.555)	(0.673)	(0.568)	(0.587)	(0.912)
Supervisor incentives	2.129***	1.972***	1.889***	1.932***	0.670
	(0.509)	(0.584)	(0.505)	(0.505)	(0.806)
Shared incentives	3.385***	3.632***	3.515***	3.461***	2.227***
	(0.492)	(0.695)	(0.540)	(0.523)	(0.859)
X	0.046	-0.333***	2.583***	0.261	-3.285***
	(0.095)	(0.097)	(0.595)	(0.581)	(0.460)
Worker incentives * X	-0.208	0.206	-0.353	0.244	0.368
	(0.158)	(0.198)	(0.828)	(1.010)	(0.822)
Supervisor incentives * X	0.042	0.199	0.474	1.629	1.594*
	(0.134)	(0.152)	(0.772)	(1.073)	(0.867)
Shared incentives * X	-0.000	0.072	-0.836	-1.066	1.118
	(0.138)	(0.142)	(0.853)	(0.977)	(0.829)
Unit	Household	Household	Household	Household	Household
Observations	8,559	5,538	8,601	8,601	8,459
Mean Dep. Var.	7.314	7.314	7.314	7.314	7.314
Mean Dep. Var. in Control	5.360	5.360	5.360	5.360	5.360
Mean X	0.000	1.465	0.308	0.112	0.793
p-value Worker*X = Supervisor*X	0.151	0.974	0.275	0.261	0.221
p-value Supervisor*X = Shared*X	0.783	0.441	0.096	0.028	0.634
p-value Worker*X = Shared*X	0.232	0.519	0.566	0.246	0.447

Notes: All regressions include stratification variables. Standard errors clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.A.5: More Health Outcomes

	(1)	(2)	(3)
Dep. Var.	Presence of a woman who gave birth in the past year in the household = {0, 1}	% household respondents who know how to prevent malaria and diarrhea	Under-5 mortality rate
Worker incentives	0.036	0.062*	-0.318
	(0.023)	(0.032)	(2.220)
Supervisor incentives	0.035	0.053*	3.083
	(0.023)	(0.032)	(3.838)
Shared incentives	0.010	0.086***	-1.485
	(0.024)	(0.030)	(2.093)
Unit	Worker	Worker	Worker
Observations	2,970	2,970	2,824
Mean dep. var.	0.841	0.563	4.135
Mean dep. var. in Control	0.819	0.511	3.822
p-value Worker = Supervisor	0.976	0.796	0.380
p-value Supervisor = Shared	0.273	0.312	0.218
p-value Worker = Shared	0.263	0.460	0.571

Notes: A household respondent "knows how to prevent malaria and diarrhea" if she reports that children should sleep under mosquito nets that is spayed with mosquito repellent/insecticide and use soap, use toilet facility to defecate, drink clean water. Mortality under-5 is measured as child mortality per 1000 years of exposure to the risk of death (self reported by households and aggregated to village level). The measure follows the method used Nyqvist et al. (2019). All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.6: Health Facility Services

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Pre- and post-natal care at the health facility in the past month				Disease treatments at the health facility in the past month		
Dep. Var.	Index (cols. 2-4)	Number of pregnant women services	Number of institutional births	Number of fully immunized children	Index (cols. 6-7)	Number of malaria cases treated	Number of diarrhea cases treated
Worker incentives	0.108	5.001	0.895	2.302	0.175*	9.857*	4.309**
	(0.094)	(5.511)	(1.020)	(1.553)	(0.099)	(5.839)	(2.102)
Supervisor incentives	0.097	4.799	1.731	0.852	0.187**	10.104*	2.420
	(0.083)	(4.220)	(1.090)	(1.352)	(0.094)	(5.767)	(2.231)
Shared incentives	0.244*	13.918*	2.552*	3.042*	0.223*	9.455	6.309*
	(0.126)	(7.283)	(1.389)	(1.625)	(0.132)	(6.621)	(3.269)
Unit	PHU	PHU	PHU	PHU	PHU	PHU	PHU
Observations	371	371	371	371	371	371	371
Mean dep. var.	0.000	41.889	13.776	12.406	0.000	57.464	18.936
Mean dep. var. in Control	-0.110	36.063	12.513	10.892	-0.155	49.595	15.582
p-value Worker = Supervisor	0.916	0.971	0.483	0.386	0.908	0.967	0.494
p-value Supervisor = Shared	0.269	0.209	0.584	0.203	0.800	0.923	0.291
p-value Worker = Shared	0.323	0.263	0.246	0.690	0.733	0.953	0.569

Notes: This table present results using admin data from the local health facilities (one per PHU). The number of observations is 371 instead of 372 because of one missing variable. The index in col. (1) [resp., col. (5)] estimates an equally weighted average of the z-scores of variables in cols. (2)-(4) [resp., cols. (6)-(7)]. All regressions include stratification variables. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.7: Reporting by Network Availability

	(1)	(2)
Dep. Var.	Reporting rate = number of reports/number of visits	
A. Treatment effects for villages <u>without</u> phone network		
No network * Worker incentives	0.080	0.070
	(0.059)	(0.068)
No network * Supervisor incentives	0.011	0.002
	(0.045)	(0.054)
No network * Shared incentives	0.022	0.022
	(0.044)	(0.050)
B. Treatment effects for villages <u>with</u> phone network		
Network * Worker incentives	0.238***	0.240***
	(0.046)	(0.051)
Network * Supervisor incentives	0.063*	0.068*
	(0.036)	(0.040)
Network * Shared incentives	0.101**	0.089**
	(0.039)	(0.040)
Network	0.014	0.015
	(0.027)	(0.033)
Unit	Worker	Worker
Extra Controls	No	Yes
Observations	2,532	2,227
Mean Dep. Var.	0.177	0.177
Mean Dep. Var. in Control & No Network	0.080	0.078
Treatment comparisons in Panel A (No network)		
p-value Worker=Supv	0.186	0.232
p-value Worker=Shared	0.257	0.378
p-value Supv=Shared	0.748	0.587
Treatment comparisons in Panel B (Network)		
p-value Worker=Supv	0.001	0.002
p-value Worker=Shared	0.010	0.007
p-value Supv=Shared	0.389	0.641
Treatment comparisons across Panels (No network vs. network)		
p-value for Worker incentives	0.016	0.025
p-value for Supervisor incentives	0.239	0.209
p-value for Shared incentives	0.078	0.138

Notes: The table reports coefficients from a fully interacted model in which the treatment dummies are interacted with a dummy for whether the network is available in the village. Col. (2) controls for the correlates of network availability (p<.1) -- i.e., age and wealth of the health worker, number of households the health worker is responsible for, distance to supervisor -- interacted with the treatment dummies. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.8: Supervisor Effort - More Outcomes

	(1)	(2)
Dep. Var.	Difference in health workers' knowledge about health between baseline and endline	Supervisor ever helped health worker with SMS reporting = {0, 1}
Worker incentives	0.158	-0.018
	(0.116)	(0.032)
Supervisor incentives	0.063	-0.007
	(0.113)	(0.035)
Shared incentives	0.266**	-0.059**
	(0.121)	(0.029)
Unit	Worker	Worker
Observations	2,927	2813
Mean dep. var.	-0.313	0.158
Mean dep. var. in Control	-0.433	0.181
p-value Worker = Supervisor	0.372	0.717
p-value Supervisor = Shared	0.074	0.075
p-value Worker = Shared	0.350	0.120

Notes: The outcome variable in col. (1) reports the difference in health worker knowledge between baseline and endline. The health knowledge is measured through a test that counts the number of correct or almost correct answers out of 5 questions asked to the health worker about when and whether to refer a child under-5. The outcome variable in col. (2) is reported by the health worker. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 1.A.9: Correlations

	(1)	(2)	(3)	(4)	(5)	(6)
	Number of years of experience of the health worker	Health worker has high experience (above median of 4 years) = {0, 1}	Supervisor's observability of output [Correlation between actual and perceived ranking]	Supervisor has low observability of output = {0, 1} [Correlation between actual and perceived ranking is in bottom decile]	Difference between supervisor and health worker hourly pay in other job	Supervisor has higher hourly pay than health worker in other job = {0,1}
Unit of Observations	Worker	Worker	Supervisor	Supervisor	Worker	Worker
Mean	5.039	0.473	0.067	0.144	-0.958	0.514
S.D.	4.481	-	0.425	-	18.358	-
A. Characteristics of the health workers						
Male = {0, 1}	0.129***	0.119***			-0.028	0.006
Age (in years)	0.166***	0.178***			0.026	0.045
Completed primary education = {0, 1}	-0.024	-0.051***			-0.002	-0.038
Completed secondary education = {0, 1}	0.022	-0.006			0.027	-0.064
Wealth score (0 to 8)	-0.052***	-0.073***			-0.070*	-0.048
Number of households responsible for	0.017	0.007			0.020	-0.002
Distance to supervisor (in km)	0.069***	0.084***			-0.018	0.013
B. Characteristics of the supervisors						
Male = {0, 1}			-0.013	0.066		
Age (in years)			0.044	-0.061		
Completed primary education = {0, 1}			0.086	-0.008		
Completed secondary education = {0, 1}			-0.098*	0.016		
Wealth score (0 to 8)			0.032	-0.067		
Number of health workers responsible for			0.086	-0.245***		
Average distance to supervisor (in km)			-0.020	0.092*		
% health workers > 1 km away			-0.038	0.133**		

Notes: This table presents summary statistics for each column variable, and pairwise correlations of the column variable with the raw variable. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.10: Heterogeneity by Worker Experience

	(1)	(2)	(3)	(4)
Dep. Var.	Number of visits		Supervisor effort: % accompanied household visits	
A. Treatment effects for workers with experience below the median:				
Low experience * Worker incentives	2.054***	2.395***	0.030	0.030
	(0.628)	(0.722)	(0.025)	(0.027)
Low experience * Supervisor incentives	2.576***	2.661***	0.067**	0.067**
	(0.598)	(0.646)	(0.029)	(0.031)
Low experience * Shared incentives	4.022***	4.335***	0.092***	0.104***
	(0.684)	(0.751)	(0.026)	(0.027)
B. Treatment effects for workers with experience above the median:				
High experience * Worker incentives	2.246***	2.056***	0.030	0.031
	(0.780)	(0.756)	(0.031)	(0.032)
High experience * Supervisor incentives	1.720**	1.657**	0.045	0.045
	(0.669)	(0.643)	(0.030)	(0.031)
High experience * Shared incentives	2.583***	2.638***	0.030	0.022
	(0.608)	(0.670)	(0.030)	(0.032)
High experience	1.057**	1.141*	0.017	0.033
	(0.532)	(0.594)	(0.025)	(0.028)
Unit	Worker	Worker	Worker	Worker
Extra Controls	No	Yes	No	Yes
Observations	2,909	2,552	2,902	2,547
Mean Dep. Var.	7.296	7.296	0.204	0.204
Mean Dep. Var. in Control & Low experience	4.749	4.749	0.131	0.131
Treatment comparisons in Panel A (Low experience)				
p-value Worker=Supv	0.455	0.733	0.226	0.236
p-value Worker=Shared	0.011	0.026	0.029	0.010
p-value Supv=Shared	0.057	0.038	0.431	0.248
Treatment comparisons in Panel B (High experience)				
p-value Worker=Supv	0.551	0.630	0.643	0.676
p-value Worker=Shared	0.684	0.492	0.990	0.781
p-value Supv=Shared	0.234	0.192	0.632	0.482
Treatment comparisons across Panels (Low vs. High experience)				
p-value for Worker incentives	0.824	0.716	0.994	0.973
p-value for Supervisor incentives	0.270	0.218	0.535	0.572
p-value for Shared incentives	0.094	0.077	0.086	0.039

Notes: The table reports the coefficients from a fully interacted model in which the treatment dummies are interacted with a dummy for whether the worker's experience is high or low. "Low experience" is an indicator that takes value one if the health worker has less than the median number of of experience (i.e., less than 4 years of experience) as a health worker at baseline. Cols. (2) and (4) control for the health worker characteristics that are significantly correlated (p<.1) with experience -- i.e., gender, age, wealth score, distance to supervisor -- interacted with the treatment dummies. All regressions include stratification variables. Standard errors clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.11: Heterogeneity by Output Observability

	(1)	(2)
Dep. Var.	Side-payment: Supervisor shared incentive with health worker = {0, 1}	
A. Treatment effects for supervisors with low observability of output:		
Low observability * Worker incentives	-0.028	-0.027
	(0.032)	(0.031)
Low observability * Supervisor incentives	0.032	0.044
	(0.065)	(0.064)
Low observability * Shared incentives	0.067	0.062
	(0.089)	(0.087)
B. Treatment effects for supervsiors with high observability of output:		
High observability * Worker incentives	0.009	0.007
	(0.018)	(0.019)
High observability * Supervisor incentives	0.205***	0.200***
	(0.052)	(0.052)
High observability * Shared incentives	0.107**	0.110**
	(0.042)	(0.043)
High observability	-0.008	-0.006
	(0.023)	(0.025)
Unit	Worker	Worker
Extra Controls	No	Yes
Observations	2,915	2,915
Mean Dep. Var.	0.084	0.084
Mean Dep. Var. in Control & Low observability	0.000	0.000
Treatment comparisons in Panel A (Low observability)		
p-value Worker=Supv	0.370	0.293
p-value Worker=Shared	0.305	0.335
p-value Supv=Shared	0.750	0.871
Treatment comparisons in Panel B (High observability)		
p-value Worker=Supv	<0.001	<0.001
p-value Worker=Shared	0.021	0.017
p-value Supv=Shared	0.135	0.160
Treatment comparisons across Panels (Low vs. High observability)		
p-value for Worker incentives	0.315	0.380
p-value for Supervisor incentives	0.040	0.061
p-value for Shared incentives	0.680	0.619

Notes: The table reports the coefficients from a fully interacted model in which the treatment dummies are interacted with a dummy for whether the supervisor has high/low observability of output. "Low observability" is an indicator that takes value one if the correlation between the actual worker ranking (based on endline household data on visit) and the supervisor's perceived worker ranking at endline is in the bottom decile (i.e. is negative). Col. (2) also controls for correlates of observability (i.e., supervisor completed secondary school), interacted with the treatment dummies. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.12: Sample of Workers with Higher Outside Option than Supervisors

	(1)	(2)
Dep. Var.	Number of visits	Side-payment: Supervisor shared incentive with health worker = {0, 1}
Sample:	Workers with higher outside option than their supervisor	
Worker incentives	0.453	-0.033
	(1.344)	(0.065)
Supervisor incentives	2.303	0.248**
	(1.396)	(0.103)
Shared incentives	3.286**	0.029
	(1.322)	(0.080)
Unit	Worker	Worker
Observations	291	293
Mean dep. var.	7.641	0.116
Mean dep. var. in Control	5.848	0.044
p-value Worker = Supervisor	0.184	0.020
p-value Supervisor = Shared	0.498	0.088
p-value Worker = Shared	0.033	0.323

Notes: Sample restricted to workers with higher outside option than supervisor. These are workers with an average hourly earnings from any outside (secondary) job which is higher than the one of their supervisor, conditional on both the worker and the supervisor being engaged in an outside job with a positive income. All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.13: Heterogeneity by Performance Ranking, Social Distance and Span of Control

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dep. Var.	Number of visits	Supervisor effort: % accompanied household visits	Side-payment: Supervisor shared incentive with health worker = {0, 1}	Number of visits	Supervisor effort: % accompanied household visits	Side-payment: Supervisor shared incentive with health worker = {0, 1}	Number of visits	Supervisor effort: % accompanied household visits	Side-payment: Supervisor shared incentive with health worker = {0, 1}
Panel A	*Effects for "Low Rank" workers: ranked in the bottom half by the supervisor at baseline*			*Effects for "Low Social Distance" workers: friends/family of the supervisor at baseline*			*Effects for "Low Span of Control" PHU: supervisor responsible for less than 7 health workers*		
Low X * Worker incentives	2.366***	0.038	0.006	1.726***	0.010	0.004	1.828**	-0.038	0.042
	(0.692)	(0.026)	(0.016)	(0.612)	(0.025)	(0.016)	(0.850)	(0.034)	(0.040)
Low X * Supervisor incentives	2.356***	0.050*	0.182***	2.322***	0.042	0.170***	1.730*	0.013	0.077
	(0.583)	(0.028)	(0.047)	(0.616)	(0.030)	(0.044)	(1.003)	(0.041)	(0.049)
Low X * Shared incentives	3.752***	0.078***	0.097**	3.165***	0.052**	0.108**	2.714***	0.004	0.107**
	(0.592)	(0.026)	(0.038)	(0.580)	(0.024)	(0.045)	(0.711)	(0.032)	(0.053)
Panel B	*Effects for "High Rank" workers: ranked in the top half by the supervisor at baseline*			*Effects for "High Social Distance" workers: not friends/family of the supervisor at baseline*			*Effects for "High Span of Control" PHU: supervisor responsible for 8 health workers or more*		
High X * Worker incentives	1.939***	0.020	0.002	2.559***	0.064**	0.007	2.350***	0.055**	-0.007
	(0.626)	(0.026)	(0.016)	(0.788)	(0.032)	(0.019)	(0.700)	(0.027)	(0.019)
High X * Supervisor incentives	1.589**	0.033	0.173***	2.013***	0.086***	0.193***	2.448***	0.076***	0.218***
	(0.722)	(0.027)	(0.048)	(0.646)	(0.030)	(0.054)	(0.589)	(0.028)	(0.062)
High X * Shared incentives	3.301***	0.044*	0.110**	3.592***	0.076**	0.095**	3.722***	0.090***	0.099*
	(0.668)	(0.026)	(0.043)	(0.752)	(0.032)	(0.043)	(0.634)	(0.027)	(0.052)
High X	0.515	0.009	0.004*	-0.244	-0.065***	0.004	-1.120*	-0.016	0.012
	(0.481)	(0.019)	(0.002)	(0.519)	(0.024)	(0.010)	(0.578)	(0.031)	(0.036)
Meaning of X	**Performance Ranking**			**Social Distance**			**Span of Control**		
Unit	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker	Worker
Observations	2,696	2,689	2,685	2,915	2,908	2,904	2,926	2,919	2,915
Mean Dep. Var.	7.296	0.204	0.084	7.296	0.204	0.084	7.296	0.204	0.084
Mean Dep. Var. in Control & Low X	5.118	0.161	0.012	5.410	0.145	0.014	5.619	0.181	0.000
Treatment comparisons in Panel A (Low X)									
p-value Worker=Supv	0.989	0.696	0.000	0.420	0.312	0.000	0.935	0.224	0.574
p-value Worker=Shared	0.072	0.145	0.017	0.045	0.107	0.022	0.364	0.201	0.311
p-value Supv=Shared	0.038	0.331	0.145	0.237	0.746	0.303	0.376	0.837	0.665
Treatment comparisons in Panel B (High X)									
p-value Worker=Supv	0.647	0.648	0.000	0.494	0.455	0.001	0.894	0.505	0.000
p-value Worker=Shared	0.059	0.405	0.012	0.242	0.711	0.040	0.081	0.238	0.028
p-value Supv=Shared	0.033	0.717	0.313	0.040	0.736	0.138	0.066	0.643	0.120
Treatment comparisons across Panels (Low vs. High X)									
p-value for Worker incentives	0.523	0.505	0.214	0.312	0.121	0.807	0.637	0.033	0.268
p-value for Shared incentives	0.538	0.246	0.161	0.627	0.490	0.757	0.288	0.041	0.921
p-value for Supervisor incentives	0.341	0.576	0.281	0.681	0.210	0.473	0.536	0.211	0.082

Notes: The table reports the coefficients from a fully interacted model in which the treatment dummies are interacted with a dummy for whether the worker is ranked in the bottom/top half by the supervisor in cols (1)-(3), whether the worker is a friend/family member of the supervisor in cols (4)-(6), whether the span of control is high or low in cols. (7)-(9). All regressions include stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.A.14: Job Satisfaction

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Health worker is satisfied with [...] = {0, 1}			Supervisor is satisfied with [...] = {0, 1}			Health worker perceives the work environment as non-competitive	Health worker self-identifies herself through her job = {0, 1}	Number of visits		
Dep. Var.	Incentives payments	Organization paying the incentives	Job	Incentives payments	Organization paying the incentives	Job			Sample: Workers with "low" inequality aversion	Sample: Workers with "medium" inequality aversion	Sample: Workers with "high" inequality aversion
Worker incentives	0.344***	0.074***	-0.039	0.086	0.012	-0.067	-0.046	0.006	1.982*	2.359***	0.922
	(0.032)	(0.024)	(0.036)	(0.056)	(0.049)	(0.061)	(0.044)	(0.026)	(1.063)	(0.571)	(1.257)
Supervisor incentives	-0.003	0.014	-0.040	0.305***	0.082*	-0.073	-0.023	-0.019	1.823***	2.419***	1.757
	(0.030)	(0.028)	(0.038)	(0.062)	(0.042)	(0.062)	(0.043)	(0.030)	(0.679)	(0.598)	(2.064)
Shared incentives	0.198***	0.048**	-0.054	0.329***	0.092**	-0.040	-0.024	-0.018	2.803***	3.675***	4.369**
	(0.036)	(0.024)	(0.036)	(0.063)	(0.041)	(0.059)	(0.043)	(0.029)	(0.835)	(0.572)	(1.831)
Unit	Worker	Worker	Worker	Supervisor	Supervisor	Supervisor	Worker	Worker	Worker	Worker	Worker
Observations	2,709	2,825	2,876	360	364	359	2,923	2,923	772	1,913	236
Mean dep. var.	0.357	0.870	0.793	0.311	0.909	0.738	0.727	0.829	7.150	7.442	6.532
Mean dep. var. in Control	0.219	0.837	0.828	0.132	0.860	0.787	0.746	0.838	5.795	5.017	5.478
p-value Worker = Supervisor	<0.001	0.023	0.986	0.001	0.098	0.929	0.598	0.379	0.884	0.929	0.662
p-value Supervisor = Shared	<0.001	0.186	0.709	0.738	0.757	0.605	0.987	0.963	0.252	0.062	0.266
p-value Worker = Shared	<0.001	0.254	0.679	<0.001	0.048	0.664	0.609	0.394	0.487	0.040	0.074

Notes: A worker/supervisor is defined as unsatisfied with the incentive payment if she reports that the incentive she is paid per valid SMS report is "not fair" (too little). A worker/supervisor is defined as unsatisfied with the environment if she reports that the environment is competitive rather than cooperative. ``Health worker self-identifies herself through her job" is a dummy variable that takes value one if the health worker answers "my job as a community health worker" to the following question: "We have spoken with many people in Sierra Leone and they identify themselves to different groups. Some people self identify themselves as belong to an ethnic group, a language, a religion, etc. Others identify themselves describe themselves in terms of their job. Besides being a citizen of Sierra Leone, which specific group do you feel you belong to first and foremost?". Questions in cols. (7) and (8) were not asked to the supervisor. The sample size changes across columns because a number of health workers and supervisors answered "don't know" to the questions. Cols. (9) to (11) are restricted to workers indicated in the panel headings. Inequality aversion is measured by asking each health worker the following hypothetical questions: "There is a local farm that hires workers to help with the potato harvest. Sheka accepts a contract to work at the farm for 20,000 SLL per day. He arrives at work the next morning. The farm is very big and there is one supervisor for the 20 workers helping with the harvest. He learns that his supervisor gets paid [amount] SLL per day. Do you think Sheka will show up to work the next day?," and amount = {20,000; 30,000; 120,000}. Our measure of inequality aversion takes value 0 ("low") if the worker answers that Sheka would always show up to work, regardless of the amount; value 1 ("medium") if worker reports that Sheka would not show up only if amount=120,000 and would show up otherwise; value 2 ("high") if the worker reports that Sheka would not show up if amount > 30,000. All regressions include stratification variables. Standard errors clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.A.15: Parameter Estimates (Alternative Assumption on Expected Reporting Rate)

	(1)
Complementarity γ	7.4
Worker unit cost of effort c_1	10.6
Supervisor unit cost of effort c_2	70914.0
Worker baseline incentive b_1	133.1
Supervisor baseline incentive b_2	93.2
α	1.4
Calibrated friction z	11.74
Δ in marginal product of CHW effort (shared incentive)	117 %
Δ in marginal product of CHW effort (no incentive)	83 %
Total CHW cost of effort (no incentive)	170.6
Total PS cost of effort (no incentive)	11088.1

Notes: The first panel of the table shows parameter estimates obtained using minimum distance estimation for the version of the model where the supervisor correctly expects the reporting rate to differ by treatment. The second panel first shows the calibrated value of contractual frictions. Second, it shows some quantities implied by the parameter estimates.

Table 1.A.16: Moment Fit (Alternative Assumption on Expected Reporting Rate)

Moments	Targeted Real	Simulated
Supervisor effort in worker incentive group	0.198	0.205
Supervisor effort in supervisor incentive group	0.225	0.231
Supervisor effort in shared incentive group	0.228	0.221
Supervisor effort in control group	0.164	0.156
Output in worker incentive group	59.679	61.679
Output in supervisor incentive group	58.896	60.773
Output in shared incentive group	66.895	62.285
Output in control group	41.040	41.157
Value loss function	6.6	

Notes: The table shows the targeted empirical moments used for minimum distance estimation as well as the simulated moments. In this version of the model the supervisor correctly expects the reporting rate to differ by treatment.

Table 1.A.17: Heterogeneity by Promotion Incentives

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)
	Household visits provided by the health worker in the past 6 months					% households who trust the health worker as a health provider
	Number of visits	% households visited	Number of visit types	Average visit length	Number of health topics discussed per visit	
Worker incentives	1.635	0.094*	0.305	1.221	0.006	0.005
	(1.125)	(0.048)	(0.189)	(2.128)	(0.203)	(0.045)
Supervisor incentives	1.664*	0.063	0.414*	2.116	0.386	0.064
	(0.992)	(0.051)	(0.237)	(2.157)	(0.312)	(0.045)
Shared incentives	3.335***	0.139***	0.611***	4.432**	0.521**	0.125***
	(1.186)	(0.047)	(0.190)	(2.041)	(0.238)	(0.044)
Meritocratic promotions	0.651	0.072*	0.264	2.369	0.224	0.070*
	(0.766)	(0.042)	(0.163)	(1.730)	(0.190)	(0.039)
Pay progression	-0.895	0.004	0.011	-1.980	0.026	0.020
	(0.844)	(0.048)	(0.182)	(1.905)	(0.265)	(0.043)
Meritocratic promotions + Info about supv. fixed salary	0.272	-0.031	0.065	-0.914	0.080	-0.017
	(0.848)	(0.044)	(0.163)	(1.555)	(0.203)	(0.048)
Worker incentives * Meritocratic promotions	-0.784	-0.140**	-0.485*	-3.099	-0.216	-0.020
	(1.700)	(0.068)	(0.263)	(2.765)	(0.309)	(0.061)
Supervisor incentives * Meritocratic promotions	2.352	0.037	0.128	0.271	-0.194	-0.084
	(1.429)	(0.066)	(0.307)	(2.761)	(0.393)	(0.062)
Shared incentives * Meritocratic promotions	0.064	-0.068	-0.172	-2.104	-0.114	-0.158**
	(1.533)	(0.064)	(0.270)	(2.672)	(0.389)	(0.065)
Worker incentives * Info about supv. fixed salary	0.491	-0.010	-0.033	3.265	0.322	0.045
	(1.427)	(0.073)	(0.263)	(2.829)	(0.356)	(0.065)
Supervisor incentives * Info about supv. fixed salary	-0.046	-0.018	-0.261	-1.068	-0.315	-0.068
	(1.248)	(0.071)	(0.293)	(2.744)	(0.412)	(0.067)
Shared incentives * Info about supv. fixed salary	0.217	-0.045	-0.121	-0.481	-0.200	-0.082
	(1.376)	(0.067)	(0.259)	(2.795)	(0.356)	(0.062)
Worker incentives * Merit. + Info about supv. fixed salary	2.157	0.059	0.292	2.954	0.521	0.102
	(1.569)	(0.065)	(0.251)	(2.657)	(0.316)	(0.065)
Supervisor incentives * Merit. + Info about supv. fixed salary	-0.416	0.057	-0.233	-0.011	-0.354	0.017
	(1.303)	(0.070)	(0.279)	(2.559)	(0.372)	(0.067)
Shared incentives * Merit. + Info about supv. fixed salary	-0.290	0.058	0.080	1.039	0.289	0.016
	(1.510)	(0.064)	(0.253)	(2.475)	(0.337)	(0.064)
Unit	Worker	Worker	Worker	Worker	Worker	Worker
Observations	2,926	2,926	2,926	2,926	2,926	2,926
Mean dep. var.	7.296	0.709	1.745	14.39	2.248	0.745
Mean dep. var. in Control	5.334	0.637	1.448	12.32	2.015	0.707

Notes: All regressions include stratification variables. Standard errors clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 1.A.18: Sample of Health Workers without Promotion Incentives

	(1)	(2)	(3)	(4)	(5)	(6)
	Household visits provided by the health worker in the past 6 months					
Dep. Var.	Number of visits	% households visited	Number of visit types	Average visit length	Number of health topics discussed per visit	% households who trust the health worker as a health provider
Worker incentives	2.357***	0.119***	0.345***	3.915***	0.330**	0.051
	(0.815)	(0.040)	(0.132)	(1.359)	(0.165)	(0.033)
Supervisor incentives	2.350***	0.138***	0.452***	1.988	0.362**	0.048
	(0.686)	(0.038)	(0.129)	(1.233)	(0.171)	(0.038)
Shared incentives	3.122***	0.154***	0.636***	4.267***	0.603***	0.060*
	(0.655)	(0.037)	(0.134)	(1.251)	(0.165)	(0.035)
Unit	Worker	Worker	Worker	Worker	Worker	Worker
Observations	960	960	960	960	960	960
Mean dep. var.	6.753	0.682	1.660	13.702	2.168	0.730
Mean dep. var. in Control	4.777	0.571	1.268	11.060	1.811	0.680
p-value Worker = Supervisor	0.993	0.600	0.425	0.147	0.868	0.928
p-value Supervisor = Shared	0.258	0.661	0.178	0.062	0.209	0.746
p-value Worker = Shared	0.344	0.330	0.036	0.794	0.146	0.780

Notes: Sample restricted to health workers who did not experience any change in promotion incentives. All regressions include stratification variables. Standard errors clustered at the PHU level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Appendix 1.B Additional Material on the Context and Intervention

1.B.1 Context: Checklists

Workers are expected to follow a checklist when they visit a household. The checklist differs depending on the type of visit the health worker conducts:
i) Prenatal visits to a pregnant woman: Health workers are asked to visit expecting mothers at least four times over the course of a pregnancy. During these visits, health workers should first make sure not only the pregnant woman but also her husband or other decision-makers in the family are present. Second, they assess the pregnant woman for danger signs (e.g., convulsion or fever) that would require an immediate referral to the PHU. Third, they use the Mother, Newborn, and Child Health Card to assess previously agreed actions and current health practices related to the pregnancy with the family. Fourth, health workers present new visit-specific information to the family (e.g., helping with planning for the birth including arranging transportation so the woman can give birth at the PHU). Fifth, health workers and families identify barriers together and agree on an action plan until the next visit. Finally, health workers must fill a register that documents what they have done during the visit.
(ii) Accompanying a pregnant woman to the PHU for child birth: The health workers should accompany pregnant women to the PHU for giving birth. At the PHU, the health worker should help the family to obtain all necessary drugs and other supplies. In case a woman delivers at her home rather than the PHU, the health worker should assist during the birth, communicate the birth to the head of the PHU, and accompany the woman for a post-natal visit at the PHU as soon as possible after the birth.

(iii) Postnatal visits within one month of birth: Health workers are asked to visit mothers with newborn babies at least four times during the first month after birth. During these visits, health workers first assess the mother and baby for the presence of any danger signs (e.g., fever or convulsions) that would require a referral to the PHU. Second, they discuss with the family how well they were able to implement health practices agreed upon with the health worker during the previous visit. Third, health workers present new visit-specific information about health behaviors relevant to the mother and baby (e.g., telling the mother to keep the baby warm and only breastfeed the baby). Fourth, they go over a checklist of recommended health behaviors and check whether or not the family knows about and follows them. Fifth, for the items on the checklist which the family does not follow yet, health workers discuss barriers and possible solutions with the family and make a new action plan to be discussed during the next visit. Finally, health workers fill a register that documents what they have done during the visit.
(iv) Child health checkup visits: Health workers are asked to visit mothers and their young children five times between the age of 1 - 15 months. During these visits, health workers first assess the child for danger signs (e.g., convulsions or being unable to breast feed) that would require an immediate referral to the PHU. Second, they use the Mother, Newborn, and Child Health Card to assess previously agreed actions and current health practices related to the pregnancy with the family. Third, health workers present visit-specific information to the mother (e.g., advising the mother how to transition from exclusive breast feeding to other foods after the age of 6 months or reminding the mother of scheduled vaccinations for the child). Fourth, health workers and families identify barriers together and agree on an action plan until the next visit. Finally, health workers must fill a register that documents what they have done during the visit.
(v) Visits in which a disease is diagnosed and the patient is either treated or

referred to the health facility: The main focus of health workers is on children who are younger than 5 years. They are trained to identify whether a child has diarrhea, malaria, or pneumonia and to decide whether or not the child can be treated by the health worker or whether it needs to be referred to the PHU. First, health workers assess the child for general dangers signs (e.g., convulsions or the child being unable to breastfeed or drink) which would require an immediate referral to the PHU. Second, they assess the child for the three conditions above (e.g., they count the breaths per minute and compare this to age-specific threshold values in order to assess a child for pneumonia) and decide whether or not the child requires treatment and whether or not the child needs to be referred to the PHU. Health workers also should always assess children for malnutrition.

(vi) Follow-up visits of sick patients: For sick children that were not referred to the PHU, health workers are supposed to do at least two follow-up visits at the child's home on the third and sixth days after the start of the treatment. During these follow-up visits, health workers re-assess the sick child following the same steps as during the initial visit. They also should discuss the condition of the child with the caregiver and counsel the caregiver on disease-specific steps they need to undertake as well as general recommended health behaviors (e.g., hand washing or bed net use).

(vii) Routine household visits: First, health workers introduce themselves and the purpose of the visit. Second, they use the Family Health Card and assess previously agreed upon actions as well as current household health practices with the family. Third, health workers present new health information (e.g., on topics like hand-washing and sanitation, bed net use, or family planning) to the family. Finally, health workers and families identify barriers together and agree on an action plan until the next routine household visit by the health worker.

1.B.2 Context: Pay of Health Workers and Supervisors

In our sample, health workers report dedicating 18 hours per week to their health worker job and 22 hours to other jobs from which they earn another 127,000 SLL ($14.85) per month. The hourly rate from the health worker job is thus comparable to their outside option. On the other hand, supervisors report working 11 hours per week on the health program. They dedicate 21 hours to other jobs from which they earn another 156,000 SLL ($28.1) per month. When a supervisor's position becomes available, one of the health workers in that PHU is promoted to take over the position.

1.B.3 Context: Supervision

Supervisors have three main tools to train and advise health workers:
(i) Monthly trainings: Supervisors host a monthly meeting at the PHU which all health workers under their supervision are supposed to attend. During these trainings, supervisors provide information on health knowledge (how to prevent diseases, recognize dangerous signs). Central to these monthly meetings is the facilitation of mutual learning among health workers. They are asked to share both successes and barriers they experienced during their work in the previous month. Depending on the number of affected health workers, supervisors help them individually or collectively find solutions for the barriers that have been identified. This often involves re-training health workers on the checklists mentioned above or advising them on effective communication strategies health workers can use with households.
(ii) One-to-one trainings: Supervisors are asked to visit each health worker under their supervision in their village once per month. During these field visits, supervisors go through the records of health workers and randomly select three recent

households the health worker provided a service to. For each of these three cases, supervisors ask the health worker about the detailed actions the health worker took and validate whether the steps on the checklists mentioned above have been followed. Supervisors then provide detailed feedback in which they identify gaps in the health worker's knowledge and explain again in detail how to provide the health services correctly.
(iii) In-the-field supervision / direct observation: Supervisors are asked to accompany the health worker to household visits and directly observe how the health worker conducts the visit. During these household visits, supervisors identify both the strengths and weaknesses of the health worker and raise awareness about the importance of her work with the family. After the household visit, supervisors provide personal feedback to the health worker in private.

1.B.4 Intervention: Choice of the Treatments

Theoretically, the set of possible splits an organization can select from is larger than the three splits in our design (100%-0%, 50%-50% or 0%-100%). An organization could for instance decide to give 25% of the incentive to the worker and 75% to the supervisor (or vice-versa). Due to the limited sample size of the experiment, we could not test the effect of a wider set of possible splits. We chose the 50%-50% split because informal discussions we had with supervisors (outside of our experimental areas) and government officials indicated that this split was the most natural in our setting. More precisely, we asked these informants how they would split an incentive of 2,000 SLL between supervisors and workers such that the number of visits provided in the PHU is maximized. 63% of the respondents answered that the supervisor should be assigned half of the incentive (1,000 SLL), 8% answered that they should be assigned 60% of the incentive (1,200 SLL), 21% answered that they should be assigned 75% of the incentive (1,500 SLL), and the remaining

8% chose another split. In line with this, our structural model confirms that the optimal split is indeed very close to the 50%-50% one: see Section 1.6.

1.B.5 Intervention: Location of the Experiment

Our experiment takes place in 372 PHUs across six districts of Sierra Leone. One district is located in the south (Bo), one in the east (Kenema), three in the north (Bombali, Tonkolili and Kambia) and one in the west (Western Area Rural). Out of the existing 823 PHUs across the six districts, we excluded half because no up-to-date and verified list of health workers was available, and selected 372 PHUs from the remaining eligible PHUs to be part of the experiment.

1.B.6 Intervention: The Reporting System

The reporting system works in three steps:

(i) Each time a household visit is provided, the health worker is asked to send an SMS to a toll-free number indicating the date of the service, the name and phone number of the patient, and a one-letter code corresponding to the service type. If the SMS does not include all the required information, the system returns an error message.[43] All health workers in our study (including those in the control group) are asked to report their visits. The incentive was only paid for household visits that fall in one of these categories: (i) prenatal visits to a pregnant woman, (ii) accompanying a pregnant woman to the PHU for child birth, (iii) postnatal visits within 1 month of birth, (iv) child health checkup visits (for children 1-15 months), (v) visits in which a disease is diagnosed and the patient is either treated or referring to the health facility, (vi) follow-up visits of sick patients, (vii) routine

[43] When the patient is a child, the health worker reports the name and phone number of the primary care giver. When the household does not have a phone, the health worker reports the phone number of a neighbor.

household visits (e.g., providing health education on how to prevent diseases).

(ii) The SMS information is automatically uploaded to a server from which the performance incentives are calculated on a monthly basis and are paid without delay.

(iii) The SMS information is continuously back-checked by a team of monitors who contact a random 25% of households each week either by phone or in-person (unannounced visits), and ask them to confirm the date and the type of the household visit.

All health workers were promised a fixed bonus of SLL 10,000 conditional on truthful reporting at the end of the experiment. Despite this, we show in the paper that the reporting rate is low in all treatments.

1.B.7 Intervention: Promotion Incentives

A random sample of 2,081 health workers out of the 2,970 health workers in this study experienced a change in the promotion system. More specifically, six months after the start of the experiment which is the focus of this paper, the promotion system became meritocratic in a half of the 372 PHUs while the rest of the PHUs kept the status-quo system (in which the promotion decision is at the discretion of the PHU in-charge). See Deserranno et al. (2022b) for more details.

Table 1.A.17 shows that our main treatment effects on visits are orthogonal to the random variation in the promotion system and orthogonal to providing information about the supervisor's fixed wage. This is not surprising as the short-run incentives analyzed in this paper are paid by an external organization and have no role in the government promotion decision, nor do they influence the supervisor's fixed wage. Table 1.A.18 moreover shows that the effects of our incentives treatments persist if we restrict the analysis to the sub-sample of health workers that did not take part in this separate study.

Appendix 1.C Research Ethics

Following Asiedu et al. (2021), we detail key aspects of research ethics.

IRB: The project received IRB from the Universitat Pompeu Fabra (Parc de Salut MAR: 2018/7834/I), Northwestern University (ID: STU00207110) and from the Sierra Leone Ethics and Scientific Review Committee (no IRB number assigned by this local institution).

We obtained informed consent from all participants prior to the study. The consent form described the participants' risks and rights, confidentiality, and contact information. Research staff and enumerator teams were not subject to additional risks in the data collection process. None of the researchers have financial or reputation conflicts of interest with regard to the research results. No contractual restrictions were imposed on the researchers limiting their ability to report the study findings.

The interventions under study did not pose any potential harm to participants and non-participants. The intervention rollout took place according to the evaluation protocol. Our data collection and research procedures adhered to protocols around privacy, confidentiality, risk-management, and informed consent. Participants were not considered particularly vulnerable (beyond some households residing in poverty). Besides individual consent from study participants, consultations were conducted with local representatives at the district levels. All the enumerators involved in data collection were aware about implicit social norms in these communities.

We plan to present the findings from the project to district and national level authorities in Sierra Leone in 2022. However, no activity for sharing results to participants in each study village is planned due to resource constraints. We do not foresee risks of the misuse of research findings.

AEA RCT Registry: The study was pre-registered on the AEA RCT Registry with the number AEARCTR 0003345. We follow the pre-analysis plan closely. The outcomes variables we use in the paper were pre-registered.

In the pre-analysis plan, we specified that we would study treatment effects on the number of hours that the workers self-report to dedicate to the health worker job per week, as a measure of their effort. We ended up not using this self-reported variable because it suffers from a reporting bias. Indeed, it does not correlate with the average number of hours households report having been visited (number of visits $\times$ average visit length): the correlation is -0.019 and is not statistically significant. Treatment effects on self-reported hours are null.

Appendix 1.D Model Appendix

1.D.1 Set Up

This section solves the model under the assumption that $b_1 = b_2 = 0$, $c_1 = c_2 = c$, $m = 1$ and $\alpha = 1$. We will later relax these assumptions.

We first quickly summarize the simplified set-up. A supervisor (player 2) and a worker (player 1) exert efforts e_1 and e_2 to produce output y, where $y = e_1 + \gamma e_1 e_2$. Thus, output depends on the efforts of players 1 and 2 and on the level of effort complementarity (γ). Effort is costly to both the worker and the supervisor, and we assume that the cost of effort is quadratic: ce_i^2 (with $c > 0$). Before the start of the game, a principal offers to pay p to the worker and $1 - p$ to the supervisor for every unit of output produced, where $p \in [0, 1]$. There are two time periods. In period 1, the supervisor chooses effort e_1 and offers a side transfer s to the worker for every unit of output produced. Contractual frictions increase the cost of the side transfer to the principal by a factor of z ($z > 1$). Transfers can only go from the supervisor to the worker: $s \geq 0$. In period 2, the worker observes e_1 and s, and chooses e_2.

The payoff of the worker is as follows:

$$\pi_1 = (e_1 + \gamma e_1 e_2)(s + p) - ce_1^2$$

And the payoff of the supervisor:

$$\pi_2 = (e_1 + \gamma e_1 e_2)(1 - p - sz) - ce_2^2$$

1.D.2 A Key Assumption

In what follows, we will make the following assumption about the strength of the effort complementarity:

Assumption 1: $\frac{8c^2}{z} > \gamma^2$; $c, \gamma \in \mathbb{R}^+$.

As it will become clear in the next section, this assumption guarantees that both agents exert positive efforts. We can show that the following claim is true.

Claim 0: If assumption 1 ($\frac{8c^2}{z} > \gamma^2$) holds; then, it is also true that:

a) $2c^2 - \gamma^2 p(1-p) > 0$

b) $8zc^2 - \gamma^2(1 + p(z-1))^2 > 0$

Proof:

The proof will be divided in two parts. First, we show that assumption 1 implies a). Then, we show that it also implies b).

Part 1: Consider the following maximization problem:

$$\max_{p \in [0,1]} p(1-p)$$

The solution is $p = \frac{1}{2}$, such that, at its maximum, the objective function attains the value of $\frac{1}{4}$. By the definition of maximum, we have that:

$$\frac{\gamma^2}{4} \geq \gamma^2 p(1-p) \;\; \forall p \in [0,1]$$

By our assumption 1, we have that: $\frac{2c^2}{z} > \frac{\gamma^2}{4}$. Thus, by the above and the transitivity of the inequality this also implies that $\frac{2c^2}{z} > \gamma^2 p(1-p)$, and by $2c^2 \geq \frac{2c^2}{z}$ implies $2c^2 > \gamma^2 p(1-p)$ (what we wanted to show).

Part 2: First note that:

$$8zc^2 - \gamma^2(1+p(z-1))^2 > 0 \iff \frac{8zc^2}{(1+p(z-1))^2} > \gamma^2$$

Therefore, we want to show $\frac{8zc^2}{(1+p(z-1))^2} \geq \frac{8c^2}{z}$ since it is sufficient to show that Assumption 1 implies b):

$$\frac{8zc^2}{(1+p(z-1))^2} \geq \frac{8c^2}{z} \iff z^2 \geq 1 + 2p(z-1) + p^2(z-1)^2$$

$$\iff z^2(1-p)(1+p) \geq 2zp(1-p) + (1-p)^2 \iff z^2(1+p) - 2zp - (1-p) \geq 0$$

The quadratic function $z^2(1+p) - 2zp - (1-p)$ has roots $z_1 = 1$ and $z_2 = \frac{p-1}{p+1} < 0$, taking negative values between the two (in $(\frac{p-1}{p+1}, 1)$) and weakly positive elsewhere. Since $z \geq 1$, this means that for all values of z, $z^2(1+p) - 2zp - (1-p) \geq 0$ and so $\frac{8zc^2}{(1+p(z-1))^2} \geq \frac{8c^2}{z}$.

1.D.3 The Model: Main Analysis

We solve the model by backward induction:

Period 2:

The maximization problem of the worker in the second period is:

$$\max_{e_1}(e_1 + \gamma e_1 e_2)(s + p) - ce_1^2$$

Thus, her optimal level of effort is:

$$e_1^* = \frac{(s + p)(1 + \gamma e_2)}{2c}$$

Period 1:

Player 2 anticipates the optimal action of player 1 in period 2. Thus the maximization problem of player 2 is:

$$\max_{e_2,s} \frac{(s + p)(1 - p - sz)(1 + \gamma e_2)^2}{2c} - ce_2^2$$

Thus, the optimal effort and side transfer are:

$$e_2^* = \frac{\gamma(s + p)(1 - p - sz)}{2c^2 - \gamma^2(s + p)(1 - p - sz)}$$

$$s^* = \begin{cases} \frac{1-p(1+z)}{2z}, & p \leq \frac{1}{1+z} \\ 0, & p > \frac{1}{1+z} \end{cases}$$

Let us first focus on the case where $p \leq \frac{1}{1+z}$. In this case, the side transfer is strictly positive and the optimal effort of the supervisor is given by::

$$e_2^* = \frac{\gamma(1 + p(z - 1))^2}{8zc^2 - \gamma^2(1 + p(z - 1))^2}$$

Plugging e_2^* into e_1^*, we get:

$$e_1^* = \frac{2c(1 + p(z - 1))}{8zc^2 - \gamma^2(1 + p(z - 1))^2}$$

In this case, the output y is given by: is:

$$y = \frac{16zc^3(1+p(z-1))}{(8zc^2 - \gamma^2(1+p(z-1))^2)^2}$$

We then consider the case where $p > \frac{1}{1+z}$. Now the side transfer is censored at zero. Optimal efforts are given by: $s = 0$ and:

$$e_2^* = \frac{\gamma p(1-p)}{2c^2 - \gamma^2 p(1-p)}$$

$$e_1^* = \frac{pc}{2c^2 - \gamma^2 p(1-p)}$$

And so output is given by:

$$y = \frac{2pc^3}{(2c^2 - \gamma^2 p(1-p))^2}$$

What level of p maximises output?

Suppose the principal wants to set p to the level that maximizes output. This maximization problem is divided in two parts: first, we maximize y assuming that $p \leq \frac{1}{1+z}$; then, we maximize y assuming that $p > \frac{1}{1+z}$.. We will refer to the first part of the problem as the "left-hand side" problem (or LHS problem for brevity), and to the second part of the problem as the "right-hand side" problem (or RHS problem for brevity). Also, we will use $p_a^* = p^*(p \leq \frac{1}{1+z})$ to denote the level of p that maximizes output in the LHS problem and $y(p_a^*)$ as the level of output when $p = p_a^*$. We will use p_b^* and $y(p_b^*)$ symmetrically to denote the level of p that maximizes output in the RHS problem, and the corresponding level of output. After solving the two problems, we compare $y(p_a^*)$ to $y(p_b^*)$. If $y_a^* > y_b^*$ ($y_a^* < y_b^*$), the solution to the overall problem is given by p_a^* (p_b^*).

We now solve the LHS problem:

$$\max_{p\leq\frac{1}{1+z}} \frac{16zc^3(1+p(z-1))}{(8zc^2-\gamma^2(1+p(z-1))^2)^2}$$

The derivative of the objective function with respect to p is given by:

$$\frac{dy}{dp}=\frac{16zc^3(z-1)(8zc^2+\gamma^2(1+p(z-1))^2)}{(8zc^2-\gamma^2(1+p(z-1))^2)^3}.$$

Assumption 1 implies that this derivate is positive for any value of p. To see this, note that (i) $c>0$ and $z>1$ (which guarantee that the numerator is positive), and (ii) the second part of Claim 0 shows that Assumption 1 implies that $8zc^2-\gamma^2(1+p(z-1))^2>0$ for any p, such that the denominator is positive for any level of p.

This shows that, as long as $p\leq\frac{1}{1+z}$, output grows in p. Thus, the LHS problem is solved by choosing the largest possible value for p: $p_a^*=\frac{1}{1+z}$.

To find the solution to the RHS problem, we solve:

$$\max_{p>\frac{1}{1+z}} \frac{2pc^3}{(2c^2-\gamma^2p(1-p))^2}$$

In this case, the optimal p is given by the solution to:

$$\frac{dy}{dp}=\frac{2c^3(2c^2+\gamma^2p(1-3p))}{(2c^2-\gamma^2p(1-p))^3}=0$$

Claim 0 shows that Assumption 1 implies that $2c^2-\gamma^2p(1-p)>0$,. Thus, in the RHS problem, the optimal p is given by the solution to:

$$3\gamma^2p^2-\gamma^2p-2c^2=0$$

The unique positive middle solution for the optimal p is then:

$$p_b = \frac{1}{6} + \frac{\sqrt{\gamma^2 + 24c^2}}{6\gamma}$$

Interestingly, p_b decreases with γ, as can be seen from the derivative of p_b^* with respect to γ:

$$\frac{dp_b}{d\gamma} = \frac{-4c^2}{\sqrt{\gamma^2 + 24c^2}} < 0$$

In order for p_b to be the global maximum of the RHS problem, we need to ensure that (i) $\frac{d^2y}{d^2p} < 0$ (the second derivative is negative), (ii) that the objective function $(\frac{2pc^2}{(2c^2-\gamma^2p(1-p))^2})$ is continuous on $p \in [\frac{1}{1+z}, 1]$ and (iii) that $\frac{1}{6} + \frac{\sqrt{\gamma^2+24c^2}}{6\gamma} \leq 1$. We tackle each one of these requirements in turn:

- A negative second derivative at $p = \frac{1}{6} + \frac{\sqrt{\gamma^2+24c^2}}{6\gamma}$:

  $$\frac{d^2y}{d^2p} = \frac{2c^3\gamma^2((2c^2 - \gamma^2p(1-p))(1-6p) - 3(2c^2 + \gamma^2p(1-3p))(2p-1)))}{(2c^2 - \gamma^2p(1-p))^4} < 0$$

  $$\iff (2c^2 - \gamma^2p(1-p))(1-6p) - 3(2c^2 + \gamma^2p(1-3p))(2p-1) < 0$$

  Note that $p = \frac{1}{6} + \frac{\sqrt{\gamma^2+24c^2}}{6\gamma} > \frac{1}{2}$. Now take the minimum of $(2c^2 - \gamma^2p(1-p))(1-6p) - 3(2c^2 + \gamma^2p(1-3p))(2p-1)$ with respect to $p \in [\frac{1}{2}, 1]$.

  As the first derivative of $(2c^2-\gamma^2p(1-p))(1-6p)-3(2c^2+\gamma^2p(1-3p))(2p-1)$ is negative, its minimum is achieved at $p = \frac{1}{2}$. At this point: $(2c^2 - \gamma^2p(1-p))(1-6p)-3(2c^2+\gamma^2p(1-3p))(2p-1) = -\frac{1}{2}(8c^2-\gamma^2) < 0$ since $8c^2-\gamma^2 > 0$ by Assumption 1 and $8c^2 \geq \frac{8c^2}{z}$.

- The objective function is continuous:

  $$2c^2 - \gamma^2p(1-p) \neq 0 \iff p \neq \frac{1}{2} \pm \frac{\sqrt{\gamma^2 - 8c^2}}{\gamma^2}$$

A sufficient condition for this is to assume $\gamma^2 < 8c^2$ (again, implied by Assumption 1 and $8c^2 \geq \frac{8c^2}{z}$).

- The condition $\frac{1}{6} + \frac{\sqrt{\gamma^2+24c^2}}{6\gamma} \leq 1$ is equivalent to $c^2 \leq \gamma^2$. That is, the complementarity has to be high enough for a two-sided incentive to be generate higher output compared to a one-sided incentive paid to to the worker.

To sum up, the possible candidates for the optimal p^* when $c^2 \leq \gamma^2$ are:

$$p_a^* = \frac{1}{1+z}$$

$$p_b^* = \frac{1}{6} + \frac{\sqrt{\gamma^2 + 24c^2}}{6\gamma}$$

And the corresponding levels of output are:

$$y(p_a^*) = \frac{2c^3(1+z)^3}{(2c^2(1+z)^2 - \gamma^2 z)^2}$$

$$y(p_b^*) = \frac{27c^3(\gamma + \sqrt{\gamma^2 + 24c^2})}{(24c^2 - \gamma(\gamma + \sqrt{\gamma^2 + 24c^2}))^2}$$

. The optimal p is found by comparing $y(p_a^*)$ to $y(p_b^*)$.

1.D.4 Comparative Statics on the Advantage of Each Optimal Incentive Candidate

Let $\mathcal{A}_{p,q}$ be the advantage of choosing the incentive that gives p to the worker and $1-p$ to the supervisor compared to choosing the incentive that pays q to the worker and $1-q$ to the supervisor. Using this tool we can compare different incentive schemes and analyze how certain parameters affect the advantage of one versus the other.

Comparing $p = p_a^*$ and $p = 1$:

$$\mathcal{A}_{p_a^*,1} = y(p_a^*) - y(1) = \frac{2c^3(1+z)^3}{(2c^2(1+z)^2 - \gamma^2 z)^2} - \frac{1}{2c}$$

We have that:

$$\frac{d\mathcal{A}_{p_a^*,1}}{d\gamma} = \frac{8\gamma c^3 z(1+z)^3}{(2(1+z)^2 c^2 - \gamma^2 z)^3} > 0$$

since $2(1+z)^2 c^2 - \gamma^2 z > 0$ by our previous assumption: $2c^2 - \gamma^2 p(1-p) > 0$.

In a similar fashion, comparing $p = p_b^*$ and $p = 1$:

$$\mathcal{A}_{p_b^*,1} = y(p_b^*) - y(1) = \frac{27c^3(\gamma + \sqrt{\gamma^2 + 24c^2})}{(24c^2 - \gamma(\gamma + \sqrt{\gamma^2 + 24c^2}))^2} - \frac{1}{2c}$$

$$\frac{d\mathcal{A}_{p_b^*,1}}{d\gamma} = \frac{27c^3(\gamma + \sqrt{\gamma^2 + 24c^2})(24c^2 - \gamma(\gamma + \sqrt{\gamma^2 + 24c^2}) + 2\gamma + 2\gamma^2\sqrt{\gamma^2 + 24c^2})}{(8c^2 - \gamma(\gamma + \sqrt{\gamma^2 + 24c^2}))^3\sqrt{\gamma^2 + 24c^2}} > 0$$

again using $8c^2 - \gamma(\gamma + \sqrt{\gamma^2 + 24c^2}) > 0$ by our previous assumption: $8zc^2 - \gamma^2(1 + p(z-1))^2 > 0$.

This means that the advantage of choosing the optimal $p^* \in (0,1)$ compared to $p^* = 1$ is increasing in γ: the larger γ is, the more harming it is (in terms of final output), to pay all the incentive to the worker.

Let us now try the analogous comparison between $p = p_a^*$, $p = p_b^*$ and $p = 0$.
For $p = p_a^*$ versus $p = 0$:

$$\mathcal{A}_{p_a^*,0} = y(p_a^*) - y(0) = \frac{2c^3(1+z)^3}{(2c^2(1+z)^2 - \gamma^2 z)^2} - \frac{16zc^3}{(8zc^2 - \gamma^2)^2}$$

We have that:

$$\frac{d\mathcal{A}_{p_a^*,0}}{d\gamma} = \frac{8\gamma c^3 z(1+z)^3}{(2(1+z)^2c^2 - \gamma^2 z)^3} - \frac{\gamma 64 c^3 z}{(8zc^2 - \gamma^2)^3}$$

And comparing $p = p_b^*$ with $p = 0$:

$$\mathcal{A}_{p_b^*,0} = y(p_b^*) - y(0) = \frac{27c^3(\gamma + \sqrt{\gamma^2 + 24c^2})}{(24c^2 - \gamma(\gamma + \sqrt{\gamma^2 + 24c^2}))^2} - \frac{16zc^3}{(8zc^2 - \gamma^2)^2}$$

$$\frac{d\mathcal{A}_{p_b^*,0}}{d\gamma} = \frac{2c^3(\gamma + \sqrt{\gamma^2 + 24c^2})(56c^2 + \gamma^2 + 2\gamma + 3\gamma\sqrt{\gamma^2 + 24c^2})}{(8c^2 - \gamma(\gamma + \sqrt{\gamma^2 + 24c^2}))^3\sqrt{\gamma^2 + 24c^2}} - \frac{\gamma 64 c^3 z}{(8zc^2 - \gamma^2)^3}$$

As one can see from the derivatives, the effect of γ on the advantage of $p = p^*$ with respect to $p = 0$ is unclear and will depend on the specific value of γ, but also on the cost of effort of the players c and the contracting cost of the supervisor z. Intuitively, when z is small it is more likely that γ has a positive effect on the advantage of $p = p^*$ with respect to $p = 0$; while a large z makes $p = 0$ more attractive and the increase in the advantage of $p = p^*$ with respect to $p = 0$ less responsive to γ.

1.D.5 Special Cases

$\gamma = 0$, $z = 1$:

In this case, the supervisor has no incentive to exert effort, since his effort is not leading to any rise in productivity $\gamma = 0$. Therefore, his optimal level of effort is $e_2^* = 0$. And, as in the general case, he chooses to pay a positive side payment $(s \geq 0)$ as long as $p \leq \frac{1}{1+z}$. As $z = 1$, this condition simplifies to $p \leq \frac{1}{2}$.

On the other hand, the worker exerts effort:

$$e_1^* = \frac{s+p}{2c}$$

Let us then analyze the maximization problem of the principal:

- If $p \leq \frac{1}{2}$ and so $s = \frac{1-2p}{2}$, then $y = \frac{1}{4c}$. This is independent of p; that is, any $p \leq \frac{1}{2}$ would lead to the same output level y.
- If $p > \frac{1}{2}$ and $s = 0$, the principal's problem becomes:

$$\max_p \frac{p}{2c}$$

  The solution is $p^* = 1$ since the objective function is increasing in p. Note that, in this case, as $c > 0$, we have that $\gamma < c$ (unlike before).

Finally, the principal compares the two possible optimal p^*:

$$y(p^* \leq \frac{1}{2}) = \frac{1}{4c}$$

$$y(p^* = 1) = \frac{1}{2c}$$

And, as $y(p^* = 1) > y(p^* \leq \frac{1}{2})$, he chooses $p^* = 1$. This is intuitive given that the supervisor does not contribute directly to production.

$\gamma = 0,\ z > 1$:

Again here, the supervisor chooses to exert no effort $e_2^* = 0$ and offers a side payment of $s = \frac{1-p(1+z)}{2z}$ if $p \leq \frac{1}{1+z}$, while the worker exerts effort $e_1^* = \frac{s+p}{2c}$.
The two-step maximization problem of the principal is now:

- When $s > 0$ and $p \leq \frac{1}{1+z}$:

$$\max \frac{1 - p(1-z)}{4zc}$$

  solved by $p^* = \frac{1}{1+z}$ as the objective function increases in p.

- When $s = 0$ and $p > \frac{1}{1+z}$:

$$\max \frac{p}{2c}$$

  just like in the previous case, maximized at $p^* = 1$.

Now, the principal would compare the output levels under the 2 candidate:

$$y\left(p^* = \frac{1}{1+z}\right) = \frac{1}{2c(1+z)}$$

$$y(p^* = 1) = \frac{1}{2c}$$

Again, $p^* = 1$ turns out to be the optimal incentive from the point of view of the principal, since $y(p^* = 1) > y(p^* = \frac{1}{1+z})$. Indeed, the result above is nested in this example.

<u>$\gamma > 0$, $z = 1$:</u>

Using the results above and plugging in for $z = 1$ one can obtain:

- When $p \leq \frac{1}{2}$ and so $s > 0$:

$$e_2^* = \frac{\gamma}{8c^2 - \gamma^2}$$

$$e_1^* = \frac{2c}{8c^2 - \gamma^2}$$

$$y = \frac{16c^3}{(8c^2 - \gamma^2)^2}$$

- When $p > \frac{1}{2}$ and $s = 0$:

$$e_2^* = \frac{\gamma p(1-p)}{2c^2 - \gamma^2 p(1-p)}$$

$$e_1^* = \frac{pc}{2c^2 - \gamma^2 p(1-p)}$$

$$y = \frac{2pc^3}{(2c^2 - \gamma^2 p(1-p))^2}$$

The solution to the two-step principal's problem is given by one of the following p^*:

- When $p \leq \frac{1}{2}$, any $p^* \in [0, \frac{1}{2}]$ would work.
- When $p > \frac{1}{2}$, $p^* = \frac{1}{6} + \frac{\sqrt{\gamma^2+24c^2}}{6\gamma}$, as long as $\gamma > c$

Finally, the optimum will be determined by comparing:

$$y\left(p^* = \frac{1}{6} + \frac{\sqrt{\gamma^2+24c^2}}{6\gamma}\right) = \frac{27c^3(\gamma + \sqrt{\gamma^2+24c^2})}{(24c^2 - \gamma(\gamma + \sqrt{\gamma^2+24c^2}))^2}$$

$$y\left(p^* \leq \frac{1}{2}\right) = \frac{16c^3}{(8c^2 - \gamma^2)^2}$$

The p^* generating the largest level of output y will be chosen and this will depend on the specific values of γ and c.

1.D.6 Proof of Result 1

As before, we assume that Assumption 1 ($\frac{8c^2}{z} > \gamma^2;\ c, \gamma \in \mathbb{R}^+$) holds.

Result 1: When effort complementarity is lower than a threshold t, there is a unique optimal incentive scheme, which is one sided ($p^* = 1$). When effort complementarity is larger than t, there is always a two-sided scheme which is optimal ($p^* \in (0,1)$). If there are contractual frictions, this optimal two-sided scheme is the unique optimal scheme. If there are no contractual frictions $p^* = 0$ may also be optimal.

Proof: To prove this statement we will first separately prove the following claims (given assumption 1):

Claim 1. The interior solution to the left-hand side problem ($\max_{p \le \frac{1}{1+z}} y$) is strictly optimal when there are contractual frictions ($z > 1$). Otherwise, any $p \le \frac{1}{1+z}$ leads to the same level of output.

Claim 2. When $\gamma^2 > c^2$, the principal's maximization problem always has an interior solution.

Claim 3. There exists a point $t = \frac{2c^2((1+z)^2 - (1+z)^{\frac{3}{2}})}{z}$ such that for all γ such that $c^2 > \gamma^2 > 0$, $y(1) < y(\frac{1}{1+z})$ i.f.f. $\gamma^2 > t$; while $y(1) > y(\frac{1}{1+z})$ i.f.f. $\gamma^2 < t$.

Proof of Claim 1: When solving the model, we showed that the solution to the principal's left-hand side (LHS) problem, that is, $\max_{p \le \frac{1}{1+z}} y$ has a unique global solution $p^* = \frac{1}{1+z}$ when $z > 1$ and multiple solutions, namely any $p \le \frac{1}{1+z}$ when $z = 1$. This follows from the derivative of the objective function (y) with respect to p, which is increasing in p whenever $z > 1$ and is flat and equal to 0 whenever $z = 1$:

$$\frac{dy}{dp} = \frac{16zc^3(z-1)(8zc^2 + \gamma^2(1+p(z-1))^2)}{(8zc^2 - \gamma^2(1+p(z-1))^2)^3}$$

Proof of Claim 2: As explained above, $p^* = \frac{1}{1+z}$ is a global (not necessarily strict) solution to the principal's LHS maximization problem regardless the value of z. For the right-hand side (RHS) problem ($\max_{p > \frac{1}{1+z}} y$) we found that there is an interior solution (which is also the global solution to the RHS problem) whenever $\gamma^2 > c^2$. Therefore, there will always be an interior value $p^* \in (0,1)$ that solves the principal's problem (since the overall solution follows from the comparison of the value of output achieved under the solution to the LHS and RHS maximization problems).

Proof of Claim 3: First, note: $y(p=1) = \frac{1}{2c}$ and $y\left(p = \frac{1}{1+z}\right) = \frac{2c^3(1+z)^3}{(2c^2(1+z)^2-\gamma^2 z)^2}$. Now we want to analyze when the following inequality is true:

$$y(p=1) > y\left(p = \frac{1}{1+z}\right) \iff \frac{1}{2c} > \frac{2c^3(1+z)^3}{(2c^2(1+z)^2-\gamma^2 z)^2}$$

$$\iff (2c^2(1+z)^2-\gamma^2 z)^2 > 4c^4(1+z)^3 \iff 4c^4(1+z)^3 - 4c^2(1+z)^2\gamma^2 + \gamma^4 z > 0$$

The LHS of the above inequality is a quadratic function in γ^2. Therefore, we solve for its roots to understand when it takes positive or negative values (that is, when the inequality holds) and we find the following two roots:

$$\gamma_1^2 = \frac{2c^2}{z}((1+z)^2 - (1+z)^{\frac{3}{2}})$$

$$\gamma_2^2 = \frac{2c^2}{z}((1+z)^2 + (1+z)^{\frac{3}{2}})$$

Then plugging in for some value of γ^2 in the middle of the two roots, e.g. $\frac{2c^2(1+z)^2}{z}$, we see that the quadratic function takes negative values:

$$4c^4(1+z)^3 - 4c^2(1+z)^2\frac{2c^2(1+z)^2}{z} + \left(\frac{2c^2(1+z)^2}{z}\right)^2 z = -\frac{4c^4(1+z)^3}{z} < 0$$

This means that $4c^4(1+z)^3 - 4c^2(1+z)^2\gamma^2 + \gamma^4 z > 0$ i.f.f. $\gamma^2 \in (-\infty, \gamma_1^2) \cup (\gamma_2^2, \infty)$ and, conversely, $4c^4(1+z)^3 - 4c^2(1+z)^2\gamma^2 + \gamma^4 z < 0$ i.f.f. $\gamma^2 \in (\gamma_1^2, \gamma_2^2)$.

Finally, note that $c^2 \leq \gamma_2^2$, which is equivalent to $1 \leq \frac{2}{z}((1+z)^2 + (1+z)^{\frac{3}{2}})$, that is true for all $z \geq 1$ since $1 < \frac{(1+z)^2+(1+z)^{\frac{3}{2}}}{z}$. This implies that $\forall \gamma^2 < c^2$ it is true that $4c^4(1+z)^3 - 4c^2(1+z)^2\gamma^2 + \gamma^4 z > 0$ (and so $y(1) > y\left(\frac{1}{1+z}\right)$) i.f.f. $\gamma^2 \in (-\infty, \gamma_1^2)$. And by analogy, $4c^4(1+z)^3 - 4c^2(1+z)^2\gamma^2 + \gamma^4 z < 0$ (and so $y(1) < y\left(\frac{1}{1+z}\right)$) i.f.f. $\gamma^2 \in (\gamma_1^2, c^2)$. Noting that $\gamma_1^2 = t$ completes the proof of Claim 3.

We showed that if $c^2 > t > \gamma^2$, then $y(1) > y(\frac{1}{1+z})$. Since the only two candidates for being the global optimum of y with respect to p when $c^2 > \gamma^2$ and $z > 1$ are precisely $p = 1$ and $p = \frac{1}{1+z}$, under contractual frictions ($z > 1$) the global optimum is attained when $p = 1$. In addition, since under $z = 1$ $y(\frac{1}{1+z}) = y(0)$, as shown in the special case in Section 1.D.5; $y(1) > y(\frac{1}{1+z})$ also implies that $y(1) > y(0)$, such that when $c^2 > t > \gamma^2$ and $z = 1$, $p = 1$ is still the global maximum. This shows: "When effort complementarity is lower than a threshold t, there is a unique optimal incentive scheme, which is one sided ($p^* = 1$)."
"When effort complementarity is larger than t, there is always a two-sided scheme which is optimal ($p^* \in (0, 1)$)." follows from Claim 2 when $\gamma^2 > c^2 > t$ and from Claim 3 when $c^2 > \gamma^2 > t$. On the other side, "If there are contractual frictions, this optimal two-sided scheme is the unique optimal scheme." follows from the previous discussion together with Claim 1.
Finally, the last statement: "If there are no contractual frictions $p^* = 0$ may also be optimal." is directly proved in the special case in Section 1.D.5 where $z = 1$ and $\gamma > 0$.

1.D.7 The Model with Heterogeneity

In this final section we extend the model to allow workers and supervisors to have different costs and benefits. Output is now given by: $\alpha e_1 + \gamma e_1 e_2$. Further, we assume that the cost of effort is given by: $c(e_1) = c_1 e_1^2$, $c(e_2) = c_2 e_2^2$. Moreover, both players get a different benefit (b_1 and b_2) for each unit of production. Finally, the payment per unit of output is given by m.
The payoff of the worker will look as follows:

$$\pi_1 = (\alpha e_1 + \gamma e_1 e_2)(b_1 + mp + s) - c_1 e_1^2$$

And the payoff of the supervisor:

$$\pi_2 = (\alpha e_1 + \gamma e_1 e_2)(b_2 + m(1-p) - sz) - c_2 e_2^2$$

Let us solve the model by backward induction:

<u>Period 2:</u>

The maximization problem of the worker in the second period is:

$$\max_{e_1} (\alpha e_1 + \gamma e_1 e_2)(b_1 + mp + s) - c_1 e_1^2$$

Thus, the worker's optimal level of effort is:

$$e_1^* = \frac{(b_1 + s + mp)(\alpha + \gamma e_2)}{2c_1}$$

<u>Period 1:</u>

Anticipating the optimal effort of player 1, the maximization problem of player 2 becomes:

$$\max_{e_2, s} \frac{(b_1 + s + mp)(b_2 + m(1-p) - sz)(\alpha + \gamma e_2)^2}{2c_1} - c_2 e_2^2$$

Thus, the optimal effort of player 2 and the optimal side transfer are:

$$e_2^* = \frac{\gamma\alpha(b_1 + s + mp)(b_2 + m(1-p) - sz)}{2c_1c_2 - \gamma^2(b_1 + s + mp)(b_2 + m(1-p) - sz)}$$

$$s^* = \begin{cases} \frac{(b_2+m) - zb_1 - mp(z+1)}{2z}, & p \leq \frac{b_2+m-zb_1}{m(z+1)} \\ 0, & p > \frac{b_2+m-zb_1}{m(z+1)} \end{cases}$$

Let us first focus in the case where $p \leq \frac{b_2+m-zb_1}{m(z+1)}$. In this situation:

$$e_2^* = \frac{\gamma\alpha\eta^2}{8zc_1c_2 - \gamma^2\eta^2}$$

where $\eta = b_1 z + b_2 + m(1 + p(z-1))$.

And plugging e_2 into e_1:

$$e_1^* = \frac{2\alpha\eta c_2}{8zc_1c_2 - \gamma^2\eta^2}$$

In this case, the output y as a function of p is:

$$y = \frac{16\alpha^2 c_1 c_2^2 z\eta}{(8zc_1c_2 - \gamma^2\eta^2)^2}$$

In the case in which $p > \frac{b_2+m-zb_1}{m(z+1)}$, we will assume that $s = 0$:

$$e_2^* = \frac{\gamma\alpha(b_1+mp)(b_2+m(1-p))}{2c_1c_2 - \gamma^2(b_1+mp)(b_2+m(1-p))}$$

$$e_1^* = \frac{\alpha(b_1+mp)c_2}{2c_1c_2 - \gamma^2(b_1+mp)(b_2+m(1-p))}$$

And so the output is:

$$y = \frac{2\alpha^2 c_1 c_2^2(b_1+mp)}{(2c_1c_2 - \gamma^2(b_1+mp)(b_2+m(1-p)))^2}$$

Implications There are at least two implications of this model's extension. First, the condition for positive side payments is now $p \leq \frac{b_2+m-zb_1}{m(z+1)}$. This condition becomes harder to satisfy as z grows and as $b_2 - b_1$ shrinks. Second, as long as side payments are positive, output is $y = \frac{16\alpha^2 c_1 c_2^2 z\eta}{(8zc_1c_2-\gamma^2\eta^2)^2}$. When $z = 1$, output is not a function of p: all levels of p result in the same level of output. On the other hand, when $z > 1$, output is a function of p.

Appendix 1.E Prediction Survey Appendix

In collaboration with the Social Science Prediction Platform,[44] we invited social scientists to forecast how our treatments affect household visits compared to the control group. The participants made their forecasts before the results of this study were made public. Participants were paid to participate in the survey. 90% of the participants are economists; 41% of whom are faculty members and 45% are graduate students.

Participants were asked to forecast the average number of household visits health workers conduct in T_{worker}, T_{supv}, and T_{shared} after giving them a 700-word description of the study and informing them about the average number of household visits and its standard deviation for control group workers:

"We are interested to hear your predictions about the effects of the different incentive schemes on the main outcome variable, the number of household visits conducted by the community health worker in the previous 6 months as reported by the household's female primary caregiver during the endline household survey. Control Group Reference: As a reference point, community health workers in the control group conducted on average 5.3 visits per household in the 6 months preceding the endline survey, with a standard deviation of 5.6. We would like you to predict the number of visits that the health workers conducted in the other three experimental conditions: How many visits do you think the health workers carried out when the 2,000 incentive was paid in full to the community health worker? How many visits do you think the health workers carried out when the 2,000 incentive was paid in full to the supervisor? How many visits do you think the health workers carried out when the 2,000 incentive was shared equally between the community health worker and the supervisor?"

[44]See Social Science Prediction Platform. This prediction platform enables the systematic collection and assessment of expert forecasts of the effects of untested social programs.

The average forecasts for the number of household visits by survey participants are 7.73 in T_{worker} (compared to 7.42 we find in the data), 6.28 inT_{supv} (7.48), and 7.41 in T_{shared} (8.7). 52% of participants forecasted T_{worker} to be the most effective treatment in our paper, 4% chose T_{supv}, 28% chose T_{shared}, and 18% forecasted either two or all three treatments to have the same effect.

Chapter 2

Promotions and Productivity: The Role of Meritocracy and Pay Progression in the Public Sector

Joint with Erika Deserranno and Gianmarco León-Ciliotta

2.1 Introduction

Many organizations face constraints on their ability to dismiss workers or to offer them performance pay, especially in the public sector. As such, they often rely on promotion incentives to motivate their employees (Cullen and Perez-Truglia, 2021; Finan et al., 2017). But to what extent are workers motivated by the opportunity to climb the organization's ladder? Despite the long-standing theoretical literature on the effects of promotion incentives on worker productivity (e.g., Lazear and Rosen, 1981; Waldman, 1984; Gibbons and Waldman, 1999b), credible empirical evidence has remained elusive.

The design of promotion incentives involves two distinct but interrelated components. To motivate lower-tier workers to exert extra effort, promotion rules should be predominantly performance-based (*high meritocracy*) and the prize associated with a promotion should be large enough (*steep pay progression*). In this paper, we provide causal estimates of the isolated and combined effect of both of these components by means of a field experiment with a large public sector organization in Sierra Leone.

We show that meritocracy and pay progression complement each other. Raising the extent to which promotions are meritocratic increases the productivity of lower-tier workers, but this is only the case when combined with sufficiently steep pay progression. Similarly, higher pay progression boosts worker productivity, but this result holds only when promotions are meritocratic. Meanwhile, when promotions are non-meritocratic, a higher pay progression demotivates workers, causing a reduction in their productivity. These findings highlight the importance of taking into account the interactions between different tools of personnel policy.

The public-sector organization we focus on is the Community Health Worker Program implemented by the Ministry of Health and Sanitation in Sierra Leone. The experiment takes place in 372 health units, each located in a different geographical area and composed of an average of eight Community Health Workers (CHWs), who provide basic health services to households in their community, and one Peer Supervisor (PS), who monitors and trains the CHWs. CHWs receive a fixed pay that equals 60% of the PS salary, and they have the opportunity of being promoted to PS whenever a position becomes vacant in their own health unit.

Before our experiment, promotion decisions were entirely left to the discretion of the local health authority (i.e., the person in charge of the health unit) and were perceived by CHWs as being non-meritocratic: half of the CHWs in our sample expressed the belief that the best-performing CHW was unlikely to be promoted unless she had a connection with the local health authority. As part of our experiment, we collaborated with the Ministry of Health and Sanitation to transition a random half of the 372 health units to a new meritocratic promotion system that promotes the best-performing CHW based on the quantity and the quality of the health services provided (as measured by the research team). This creates random variation in the *actual* promotion criteria, which we cross-randomize with variation in the *perceived* pay gap between the PS and the CHWs. Leveraging the low initial awareness of pay disparities, we provided CHWs in a random half of the 372 health units with information about the true PS pay, thus affecting their perception of the pay progression. Our 2×2 research design allows us to assess the effect of a more meritocratic promotion regime, steeper (perceived) pay progression and the interplay between the two on CHW productivity.

To guide the empirical analysis, we develop a simple theoretical framework in which we model the promotion mechanism as a single prize contest where workers (CHWs) compete for a promotion by exerting effort. Meritocratic contests, in which promotions are based uniquely on worker performance, are predicted to boost worker effort relative to less-meritocratic contests if the pay gap between lower- and upper-tier workers is large enough. Similarly, raising the pay progression is predicted to motivate workers to climb the organization's ladder and to prompt an increase in their effort, but this is true only if the system is meritocratic enough. In a non-meritocratic system, a steeper pay progression can instead *reduce* workers' effort if they perceive promotions as being awarded in an unfair or unequal manner (i.e., a negative morale effect), or if they divert time away from providing health services into "lobbying" their superiors (de Janvry et al., 2021).

Our empirical analysis proceeds in two steps. We first study the direct causal effect of a more meritocratic promotion regime on CHW performance while holding perceptions about pay progression fixed. In line with the theoretical framework, we find that the introduction of a more meritocratic promotion rule increases the performance of workers who believe that the pay progression is steep enough at baseline: the number of visits they provide goes up by 27% with no concomitant decrease in the average visit length.[1] The effect of meritocracy on worker performance is positive also for workers who are likely to see the PS turn over within five years, while we find no effect for workers whose supervisor is unlikely to turn over soon. Finally, we document a 30% increase in the performance of workers who are ranked among the top-three in the health unit, while we find no effect on other lower-ranked workers. Overall, our findings are consistent with promotion incentives being effective at motivating two types of workers: (1) those for whom the prize associated with the promotion has a high present value, and who are presumably more interested in the promotion, and (2) those who are highly ranked in terms of performance, and who have higher chance of being promoted in a meritocratic regime. The rest of the workforce does not respond to promotion incentives.

In the second part of the empirical analysis, we study the causal effect of pay progression on CHW performance in the meritocratic promotion regime vis-a-vis the old regime. Increasing perceived pay progression - by revealing the true PS pay to workers who initially underestimated pay progression - has two contrasting effects depending on the prevailing promotion rule. In the new meritocratic promotion regime, higher (perceived) pay progression raises the number of visits provided by 24%, with an even larger effect among high-ranked workers. This indicates that even for public sector workers - who have been argued to be "intrinsically motivated" (Besley and Ghatak, 2005; Bénabou and Tirole, 2006) - extrinsic incentives in the form of a potential future higher pay play an important role, especially for high ability workers.

In the old (non-meritocratic) regime, higher (perceived) pay progression instead

[1] Higher meritocracy also increases the retention of these workers. Through a bounding exercise, we show that worker retention is not the main driver of the productivity gains.

decreases the number of visits by 26%. Two potential mechanisms can explain such a reduction in productivity: one possibility is that workers may perceive the large pay gap between the different layers of the organization as being unfair or unequal if the system does not reward highly productive workers, leading to a negative morale effect that decreases their motivation. Alternatively, the larger perceived pay gap may increase workers' interest in a promotion, incentivizing them to substitute productive activities (household visits) for non-productive ones (lobbying). We provide suggestive evidence that our results are consistent with a morale effect rather than a lobbying effect. First, the drop in the number of visits provided is not compensated by workers being more likely to interact with the local health authority nor with workers dedicating a larger fraction of their time to non-patient-oriented activities, which we would expect if they were diverting time into lobbying-related activities. Second, the reduction in the number of visits is concentrated among high-ranked workers and workers who are unsatisfied with the work of the PS, both of whom are expected to view a non-meritocratic regime with a high pay progression as the most unfair.

From a policy perspective, the results of this paper show that organizations seeking to increase the productivity of lower-tier workers should simultaneously enforce promotion rules that reward performance *and* ensure that the prize associated with promotions is large enough. This is particularly important as a large number of organizations, both in the public and private sector, adopt only one of the two above components rather than both. In large public organizations in developing countries, for example, pay progression is often steep while promotions are non-meritocratic, largely due to patronage, nepotism, or strict seniority-based rules (Wade, 1985; Shepherd, 2003; World Bank, 2016; Sahling et al., 2018; Besley et al., 2021). This is illustrated in Figures 2.A.1 and 2.A.2 which show, respectively, that many bureaucracies of low-income countries combine high pay progression with low meritocracy and that this combination negatively correlates with government performance.[2] Similarly, in the private sector, promotion rates have been shown to be significantly lower for women and minorities across all ranks of firm hierarchies, even after controlling for their performance and especially in firms with steep pay gradients (e.g., Castilla, 2008; Kunze and Miller, 2017; Cullen and Perez-Truglia, 2019; Macchiavello et al., 2020; Benson et al., 2021). While raising the pay progression in these "non-meritocratic" organizations may potentially improve the selection of high-tier workers (a mechanism we do not capture in our experiment),[3] our findings indicate a consequent demotivation of the "unfavored" low-tier workers which may hinder organizational performance.

This paper contributes to different strands of the literature. First, it adds to the literature studying the effects of promotion incentives, which has been predominantly theoretical in scope (Lazear and Rosen, 1981; Harris and Holmstrom, 1982; Waldman, 1984; Rosen, 1986; Gibbons and Murphy, 1992; Gibbons and Waldman, 1999a,b; Bose and Lang, 2017; Ke et al., 2018). A few recent empirical papers have documented the positive effects of increasing upward mobility on the performance of workers for whom a new senior position becomes "attainable", while holding the promotion rule fixed (Karachiwalla and Park, 2017; Nieddu and Pandolfi, 2018; Bertrand et al., 2020; Li, 2020).[4] There is also recent empirical work exploring

[2]Pay progression and meritocracy are measured using the Worldwide Bureaucracy Indicators, and government performance is measured using the Gothenburg's Quality of Government Indicators. Refer to the figure notes for more details. In a regression with country and time fixed effects, Figure 2.A.2 shows that government performance is negatively correlated with pay progression in non-meritocratic regimes and positively correlated with meritocracy when combined with high pay progression.

[3]The experiment allows us to assess the effect of pay progression and meritocracy on the productivity of low-tier workers (CHWs), holding the productivity of high-tier workers (PSs) fixed. However, it does not capture the effect on the productivity of high-tier workers (PSs) and how this, in turn, affects CHW performance. Indeed, we did not change the actual pay progression, and promotions are infrequent in our context.

[4]Using retrospective panel data on teachers in China, (Karachiwalla and Park, 2017) show that promotions are associated with better performance in the years leading up to promotion eligibility but reduce performance if workers are repeatedly passed over for promotion. (Nieddu

whether managerial discretion improves or deteriorates the extent to which the promotion system is performance-based (Xu, 2018; Aman-Rana, 2020; Voth and Xu, 2021).[5] In contrast with our paper, these studies do not assess the causal effect of a more meritocratic promotion rule on worker productivity, nor its interaction with pay progression.

Our paper differs from the large literature on non-tournament-based incentives, such as pay-for-performance schemes that do not involve competition across workers (e.g., Lazear, 2000; Muralidharan and Sundararaman, 2011; Khan et al., 2016, among many others). The tournament structure of promotion incentives implies that only the winner is rewarded. As a result, the types of workers who respond to promotion incentives and the magnitude of their response may sharply differ from non-tournament-based incentives - e.g., only the subset of workers who have a chance of being promoted may respond and their response may be particularly strong. Promotion incentives also differ in that their effectiveness is a function of pay progression. Whether promotion incentives are more cost-effective than non-tournament-based schemes is ultimately an empirical question. We discuss this in more detail in the concluding Section 2.7.

---

and Pandolfi, 2018) show that promotion incentives in academia prompt higher productivity, but this is only the case when the goals set are attainable. (Bertrand et al., 2020) show that strict seniority-based rules in the Indian public sector prompt an increase in effort among workers for whom the promotion is attainable while demotivating workers who are too young to be promoted in the foreseeable future. (Li, 2020) shows that exposure to unfair promotions in Chinese high schools adversely affects the productivity of non-favored teachers, a result that echoes our negative morale effects. Unlike (Li, 2020), we show that such morale effects materialize only when pay progression is large enough.

[5] In the Pakistani public sector, (Aman-Rana, 2020) shows that discretionary promotions - which are not based on any strict promotion rule - improve meritocracy if the incentives of mid-level bureaucrats (who decide on promotions) are aligned with the organization's objectives. (Voth and Xu, 2021) show that discretion in promotions in the Royal British Navy improved the selection of captains whenever the admirals had superior information about candidates; while (Xu, 2018) shows that discretion in promotions in the British Empire promoted governors connected to their superiors (patronage) who subsequently underperformed. (Weaver, 2021) studies managerial discretion in hiring (rather than in promoting) workers, and shows that letting managers select new hires based on whether they receive a bribe leads to the selection of high-quality workers.

The second strand of the literature we contribute to is the one on the effects of pay inequality within organizations on worker performance. Most of the existing empirical evidence has focused on *horizontal* pay inequalities (i.e., between workers in the same layer of an organization) while shutting down dynamic incentives, and documents negative morale effects (Card et al., 2012; Cohn et al., 2014; Mas, 2017; Breza et al., 2018). In contrast, we center our attention on *vertical* pay inequalities between supervisors and their subordinates for which the theoretical predictions are less clear. On the one hand, a steeper pay progression can demotivate workers who are averse to vertical pay inequalities. On the other hand, it can prompt an increase in effort through career incentives. Understanding which of the two effects prevails is of obvious policy relevance given the recent rapid growth of the manager-worker pay ratio (Ashraf and Bandiera, 2018). The only paper we are aware of that studies vertical pay inequalities is (Cullen and Perez-Truglia, 2021). In the context of a private-sector firm with a relatively meritocratic promotion regime, their study shows that lower-tier workers exert more effort when their perceptions of their supervisor's salary are revised upward. We complement (Cullen and Perez-Truglia, 2021) by focusing on a large public-sector organization in which promotions have only recently started to become more meritocratic and by studying how the effects of vertical pay inequalities vary with the level of meritocracy. This focus allows us to bridge the literature on pay inequalities with that on promotions.

Finally, our study contributes to investigations that explore how to build effective state capacity in developing countries (see Finan et al., 2017 for a literature review). While the low productivity of frontline public-sector workers has often been attributed to low-powered incentives, low monitoring, or inadequate selection, we argue that the lack of meritocratic promotions combined with steep pay progression – commonly seen in large bureaucracies of developing countries (as shown in Figure 2.A.1) - may also constrain the state's ability to provide high-quality public services. Our study is also related to a few recent papers which study the effect of meritocracy in personnel decisions other than promotions, i.e., transfers and hiring (Khan et al., 2019; Xu and Adhvaryu, 2020).[6] To the best of our knowledge, this is the first paper exploring the effect of performance-based promotions in the public sector, and its interaction with pay progression.

The paper is structured as follows. Section 2.2 discusses the context and research design. Section 2.3 shows how our treatments affect worker perceptions about meritocracy and pay progression. Section 2.4 introduces a theoretical framework that models worker effort responses to an increase in meritocracy and pay progression. Sections 2.5 and 2.6 present the effects of higher meritocracy and pay progression, respectively, on worker productivity. Section 2.7 concludes. In the Appendix, we discuss further results and key aspects of research ethics.

2.2 Context and Research Design

2.2.1 The Community Health Worker Program

Sierra Leone is one of the poorest countries in the world, with the third-highest maternal mortality rate and the fourth-highest child mortality rate in 2017 (World Health Organization, 2017). Such elevated mortality rates have been attributed to the slow post-civil war recovery, the 2014 Ebola outbreak, and the critical shortage of health workers together with limited access to health facilities throughout the country (World Health Organization, 2016). In order to strengthen the provision of primary health care, Sierra Leone's Ministry of Health and Sanitation (MoHS) created a national Community Health Worker program in 2017. The program

[6]In the context of property tax inspectors in Pakistan, (Khan et al., 2019) show that allowing workers to choose their location based on their performance improves their productivity. (Xu and Adhvaryu, 2020) show that more meritocracy in the recruitment system of bureaucrats in Taiwan incentivizes future potential job applicants to invest in human capital in order to increase their chance of admission, and this may improve the selection of these bureaucrats.

is organized around Peripheral Health Units (PHUs), small health posts staffed with doctors (when available), nurses, and midwives. Each PHU has typically a catchment area of seven to 10 villages with one Community Health Worker (CHW) per village and one Peer Supervisor (PS) per PHU, for a total of approximately 15,000 CHWs and 1,500 PSs nationwide.

The role of the CHWs is to provide a basic and polyvalent package of healthcare services at the community level. They do so by making home visits to households with expecting mothers or young children, during which they provide the following services: (i) health education (e.g., about the benefits of a hospital delivery), (ii) pre- and post-natal check-ups, and (iii) basic medical care and referrals to health clinics. This model of local preventive health service provision has been shown to increase the use of maternal and child health services, improve child health, and reduce child mortality in other similarly poor contexts (e.g., Darmstadt et al., 2010; Nyqvist et al., 2019; Deserranno et al., 2020).

CHWs are hired locally and typically have no experience in the health sector prior to joining the program. The role of the PS is to ensure that each CHW acquires the skills and knowledge necessary to provide primary care services. To do so, the PS organizes a monthly one-day training that CHWs are asked to attend, and subsequently advises, trains and monitors CHWs through in-person visits and by accompanying them on household visits. The PS thus has the responsibility of enabling health workers to perform their tasks (Deserranno et al., 2022a). Almost all PSs have previous experience as a CHW, and have thus already acquired health knowledge.

Both CHWs and PSs are part-time employees who typically have a secondary occupation such as farming, petty trading, or small shopkeeping. In our sample, CHWs and PSs report dedicating an average of 18 and 11 hours per week to their CHW/PS job, respectively. CHWs are paid a fixed monthly allowance of 150,000

SLL (17.5 USD) and PSs are paid 250,000 SLL (29.2 USD).[7] The pay gap between PSs and CHWs is thus large: CHWs earn 40% less than the PSs even though they report working more hours on average. Using the self-reported number of hours as a reference, the hourly wage of PSs is 2.7 times higher than that of CHWs.

As with most public-sector employees, CHWs and PSs are almost never fired and new vacancies open up when CHWs or PSs voluntarily decide to quit. PSs usually leave their jobs at the time of retirement (55 years old), and are not pushed out by "upstart" high-performing CHWs.[8] In our study, the age distribution of PSs at baseline implies that at least 10% of the positions are expected to become vacant in the following five years. Consistent with this observation, we see nine of the 372 PS positions in our sample becoming vacant during the ten months of our study, which amounts to a 15% chance of having an opening in a five years span at any given PHU.

When a PS position becomes available, one of the CHWs in that PHU is promoted to take over the position. The competition for a promotion thus happens within the PHU because CHWs are never promoted in PHUs other than their own. The District Health Management Teams (DHMTs), which oversee the implementation of the CHW program at the district level, are in charge of the promotions. Historically, the DHMTs have always delegated the promotion decision to the head of the PHU (the "PHU in-charge"), who is responsible for all personnel and administrative matters in the PHU. While delegating the promotion decision to a specific person may be optimal if that person has private information on which CHW is

[7] We use the January 2019 exchange rate: 1 USD = 8,550 SLL (Sierra Leonean Leones). This payment is formally split between their wage and a transportation and communication allowance. In practice, this distinction only serves as a way to earmark the money. These salaries are in line with earnings from other non-CHW activities: CHWs and PSs report earning 200,000 and 240,000 SLL from other non-CHW activities, to which they dedicate 18 and 19 hours per week respectively.

[8] After they retire at 55 years old, PSs are paid 10% of their wage. See data from "Social Security Programs Throughout the World: Africa" for Sierra Leone.

best fitted to serve as PS, the system is also subject to patronage and nepotism. As we describe later, our data show that there is a wide perception among CHWs that this system is not meritocratic, and that connections to the PHU in-charge, rather than productivity, is the key predictor of promotions.

The set of skills required for the PS and CHW jobs do not perfectly overlap - e.g., the PS position requires managerial skills that the CHW position does not. As a result, promoting CHWs based on their current performance (as we do in the new meritocratic system that we discuss below) is not necessarily the best possible system to select high-performing PSs.[9] Yet, such a system is likely more effective than the status-quo system that puts more weight on connections. The PS work is indeed mostly independent of the PHU in-charge and having a connection to PHU in-charge has limited added value in our context, as shown in Table 2.A.1. In contrast, promoting a high performing CHW presumably implies selecting someone who is highly motivated and with good health knowledge, both of which predict PS performance in our sample of workers (see Table 2.A.1).[10]

2.2.2 Research Design

Our experiment took place in 372 PHUs in six of the 14 districts of Sierra Leone and covers 372 PSs and 2,009 CHWs.[11] These PHUs were cross-randomized into

[9]E.g., see the "Peter Principle" (Peter et al., 1969; Benson et al., 2019). It might be more effective, for example, to promote CHWs based on their "potential" as a good manager. Such systems are however more subjective and have been shown to lead to more discrimination (Benson et al., 2021). Understanding which promotion system leads to selecting the best supervisor is outside the scope of this paper and a good avenue for future research.

[10]Table 2.A.1 shows that the high-performing PSs in our sample - i.e., those who supervise and motivate their CHWs by regularly visiting them or by frequently accompanying them on household visits - tend to have greater health knowledge and are predicted to have provided more visits when they themselves were CHWs (columns 1-4). In contrast, connections to the PHU in-charge, proxied with the number of years the PS has known the PHU in-charge before joining the program, do not predict PS performance (columns 5-6).

[11]One district is located in the south (Bo), one in the east (Kenema), three in the north (Bombali, Tonkolili and Kambia) and one in the west (Western Area Rural). In the 372 PHUs, we were able to reach 372 PSs and 2,081 CHWs by phone. Out of the 2,081 CHWs, 72 refused to

two treatment arms: (1) the "meritocratic promotion treatment", which introduced a new meritocratic promotion regime (henceforth, T_{merit}), and (2) the "pay progression treatment" which created variation in the *perceived pay progression* (henceforth, T_{pay}). We discuss these two sources of variation in turn.

Meritocratic Promotion Treatment In November 2018, we collaborated with the MoHS and the DHMTs to transition a random 186 PHUs to a new meritocratic promotion system ($T_{merit} = 1$), while the status quo was left unaltered in the remaining 186 PHUs ($T_{merit} = 0$). In the new promotion regime, the DHMTs promoted CHWs based on objective measures of CHW performance collected by the research team. Performance data were collected in $T_{merit} = 1$ and in $T_{merit} = 0$ by measuring the number of visits and the average visit length of those visits through a household survey and unannounced random spot checks with potential patients.[12] Every time a vacancy became available in a treated PHU ($T_{merit} = 1$), we provided the DHMTs with information on the number and average length of the visits provided by each CHW in the PHU, which is then used to decide on whom to promote. No information on performance was shared with DHMTs in the control PHUs ($T_{merit} = 0$).[13]

Two weeks after the new promotion system was introduced, we provided information on the new promotion system to CHWs in the 186 PHUs in which the change was implemented ($T_{merit} = 1$). The information was provided by phone by operators trained to read the following script:

be interviewed at baseline and are excluded from the sample. All the staff members interviewed at baseline were then re-interviewed at endline. See Section 2.2.3 for more details on the data.

[12] Refer to Section 2.2.3 for details on the data shared with the DHMTs and a discussion of the accuracy of these performance measures.

[13] Our data confirm that in $T_{merit} = 1$ the DHMTs used the information we provided to them: all four health workers promoted in $T_{merit} = 1$ during our experiment ranked among the top 10% in terms of number of visits, while none of the five health workers promoted in $T_{merit} = 0$ ranked that high.

> *"I would like to tell you about a new policy of how promotions from CHW to PS will be done. From now on, the number of services and the quality of services a CHW provides every month will be the key criteria for promotion decisions. The next time a new PS vacancy comes up at a PHU, the best-performing CHW at the PHU will be recommended to the DHMT for promotion to PS."*

To keep the saliency of promotions constant between the treatment and control group, we also reminded CHWs in the 186 control PHUs about the status quo promotion system ($T_{merit} = 0$). The same operator who called workers in the meritocratic promotion group read the following script to workers in the control group:

> *"I would like to tell you about the official policy of how promotions from CHW to PS should be done. The PHU in-charge or the PHU CHW Focal can nominate one of the CHWs as the new PS to the DHMT. This means that the decision whether a CHW gets promoted depends mainly on whether the PHU in-charge thinks highly of the CHW."*

Before reading the script in $T_{merit} = 1$ and $T_{merit} = 0$, the phone operators introduced themselves as belonging to a reputable survey firm, and explicitly mentioned that the information they were conveying was officially approved by the DHMT and the MoHS. In Section 2.3.1, we will demonstrate that CHWs in $T_{merit} = 1$ updated their perception of meritocracy upward after receiving the information above, indicating that they trusted and understood the information. In contrast, CHWs in $T_{merit} = 0$ did not change their perception and were thus presumably aware of the status quo system.

This variation in perceived meritocracy across treatments will allow us to quantify the effect of meritocracy on CHW performance in *anticipation* of future promotions, without the need for promotions *actually* occurring during the study period. We will thus assess whether CHWs work harder when they perceive future promotions as being more meritocratic, but will not estimate the effects of more meritocratic promotions on PS performance and on how this, in turn, affects CHW performance. Because the new meritocratic system likely improves the quality of the PS selected relative to the status-quo (as discussed in the previous section), our results are likely an underestimate of the long-run effect of meritocratic promotions on CHW performance. We discuss this in more detail in the concluding Section 2.7.

Pay Progression Treatment As explained above, PSs and CHWs are paid 250,000 SLL and 150,000 SLL per month, respectively. Importantly, this pay gap was unknown to most CHWs at baseline: only 30% of the CHWs reported knowing the exact PS pay. We took advantage of this lack of information to create random variation in *perceived* pay progression. Cross-randomizing by the meritocratic promotion treatment, we informed CHWs in a random selection of 186 PHUs of the true pay differential between their own salary and their supervisor's ($T_{pay} = 1$). The information was provided by phone, immediately after informing them about the promotion system:

> *"CHWs are entitled to 150,000 SLL per month. PSs are entitled to 250,000 SLL per month, which is 100,000 SLL more per month than CHWs."*

To keep the saliency of pay constant across all treatment groups, we reminded CHWs in the remaining 186 PHUs ($T_{pay} = 0$) about their own pay:

> *"CHWs are entitled to 150,000 SLL per month."*

As we will show in Section 2.3.2, CHWs in $T_{pay} = 1$ shifted their perception of the pay gap in different directions depending on their priors: workers who underestimated PS pay at baseline revised their perceptions upward, while those who overestimated PS pay revised downward. This variation in perceived pay progression will allow us to quantify the effect of a steeper or flatter pay progression on CHW productivity due to shifting perceptions of the pay progression rather than by changing *it per se*. Importantly, we will estimate the effects of steeper or flatter pay progression on CHW productivity, holding PS productivity fixed. Estimating the effects of *actually* changing the PS pay on the selection and the performance of the PS and how this, in turn, affects CHW performance is beyond the scope of this paper.

In sum, the 372 PHUs of this study were randomly divided into four groups of equal size varying in T_{merit} and T_{pay}. The randomization was performed at the PHU level because promotions are done at this level, as well as to limit information spillover between different treatment arms.[14] We stratified the randomization by district and by the presence of temporary performance-based incentives, which were introduced by an external organization in a sub-sample of the PHUs and which are the focus of (Deserranno et al., 2022a). In Appendix 2.B, we show that the temporary incentives did not interact with our treatments. Finally, note that all the CHWs in this study were on the job when the experiment started. As a result, our treatment effects do not capture any response on the recruitment margin.

[14]While CHWs and PSs frequently interact within a PHU, these interactions are minimal across PHUs. As a result, CHWs in $T_{pay} = 0$ are unlikely to learn about the PS pay from CHWs in $T_{pay} = 1$. We provide evidence of this later in the paper.

2.2.3 Data and Balance Checks

Data Sources

We leverage survey data collected from CHWs, PSs, and households.

CHW and PS surveys 372 PSs and 2,009 CHWs in the 372 PHUs were surveyed at baseline (in April-May 2018) and at endline (ten months after the implementation of the treatments, in July-September 2019). CHWs were surveyed on their demographic background (age, gender, education, wealth), their knowledge about health, and their CHW job (number of years of experience as a CHW, number of hours dedicated to the CHW job). The PS interviews contained similar questions, though PSs were also asked to rank the CHWs from 1 to N in terms of performance, where N is the total number of CHWs in that PHU. We will later use this as a baseline measure of relative CHW rankings and show that it correlates with other predictors of CHW performance, like CHW health knowledge and education level.

Two weeks before the implementation of the treatments (November 2018) and two weeks after (December 2018), we surveyed each CHW to assess their perceptions about how meritocratic the promotion system is and about pay progression in the organization. We discuss these measures in detail in the next section.

We also have access to baseline village-level information (i.e., accessible road to government hospital, primary school in the village, number of water sources in the village, and mobile network availability) collected from a leaflet that is given to each CHW by the PHU.

Household surveys A random sample of three eligible households per village were surveyed ten months after the implementation of the treatments (in July-

September 2019).[15] This represents roughly 7% of the total number of health workers' potential patients. The respondent was the main female household head. She was asked about the number of visits received by the CHW and the average length of those visits. Given the absence of a baseline household survey, we also asked retrospective questions (e.g., connection with the CHW a year ago, household composition) as well as questions that were unlikely to vary over time (e.g., distance from the CHW house or the PHU, education), which we use in the household balance checks.

All CHWs (both in $T_{merit} = 1$ and $T_{merit} = 0$) were made aware at baseline that we would measure their performance by interviewing households on the visits they received. As explained in the previous section, the CHWs in $T_{merit} = 1$ were also aware that this information would then be used by the DHMTs to decide on promotions. To avoid collusion with the households on misreporting visits, CHWs were not informed about how many households we would interview, which ones and when. In line with the absence of collusion, we show in Section 2.5 that the share of friends and family members of the CHW who report having received a visit is comparable to the share of non-friends who report having received a visit. While interviewing a sub-sample of households increases the noisiness of the performance data (relative to interviewing the entire village), we will later show that the measure of performance is accurate enough to affect CHW effort in $T_{merit} = 1$.

[15] In the absence of a full listing of households in each village, the sampling was done through a random walk starting from the house of the CHW and with pre-specified sampling intervals between households. To cover a random sample of households across the *entire* village (and not only households who live near the CHW), the intervals were calculated based on the total number of households in the community. In order to be eligible for the household survey, the respondent had to be female, be one of the primary caregivers, be between 18 and 49 years old, and have lived in the household for at least 6 months during the study period. We set these eligibility criteria so that sampled households would belong to the group targeted to receive the services of the CHW.

Summary Statistics and Balance Checks

Table 2.1 (Panel A) reports summary statistics and balance checks for the CHW characteristics. 73% of the CHWs are male, 71% have completed primary education and 8% have completed secondary school. On average, CHWs are 37 years old, have worked as a CHW for 2.2 years, are responsible for 57 households each, and report working 18 hours per week as a CHW. On a health knowledge test with 7 questions, they answered an average of 2.9 questions correctly, indicating low health knowledge. To perform the balance checks, we regress each baseline CHW characteristic on a dummy for the meritocratic promotion treatment, the pay progression treatment and the interaction of both, controlling for stratification variables and clustering standard errors at the PHU level. Columns (3) to (8) show that CHW characteristics are well balanced across treatments.

Panel B reports summary statistics on PS characteristics. PSs are 38 years old on average, with 10% being above 50 years old and expecting to retire within five years. Relative to the CHWs, PSs are more likely to be men (92%) and are more likely to have completed secondary school (25%). They are also more knowledgeable about health services and dedicate fewer hours per week to the program (11 hours per week). They are responsible for an average of eight CHWs each, and have worked an average of 3.5 years as a PS and an average of 1.8 years as a CHW prior to becoming a PS. PS characteristics are balanced across treatments.

Panel C presents summary statistics on CHW perceptions about meritocracy and pay progression before the implementation of the treatments. We discuss these in the next section.

Table 2.A.2 presents summary statistics at the village level (Panel A) and at the household level, aggregated to the village level (Panel B). Household respondents are less educated than both CHWs and PSs, with only 28% having completed

primary school; household members are also less wealthy. Nearly all (97%) of the households knew the CHW at baseline. Most (87%) live within 30 minutes of the CHW's house and 39% live within 30 minutes of a government hospital. The village and household characteristics are balanced across treatments.

Importantly, the data show that there is a wide perception among CHWs that the status-quo promotion system is not meritocratic. Indeed, only 45% of the CHWs reported that the PS was the best-performing CHW at the time of their promotion (last variable of Table 2.1, Panel A) and 50% reported perceiving the system as non-meritocratic at baseline, a finding that we revisit in Section 2.3.1. Moreover, we calculate that, at the time they were promoted, more than 60% of the PSs in our sample were more connected to the PHU in-charge (in term of number of years they had known each other) than any other potential PS candidate, while less than 25% of them ranked highest in terms of (predicted) performance as a CHW (see Figure 2.A.3 for details). Table 2.A.3 presents a horse race between the different CHW characteristics in predicting promotion, and shows that connections matter twice as much as performance and education, and more than 10 times as much as tenure.[16] We interpret this as evidence that social connections are the key determinant of promotions when these are decided by the PHU in-charge. Interestingly, the correlation between social connections and CHW performance is only 0.018 within the pool of CHWs we interviewed and is not statistically significant. Thus, promoting CHWs based uniquely on connections rather than based on performance presumably leads to substantially different candidate selection.

[16]We follow a two-steps procedure to predict PS past performance when they were CHWs. Refer to the notes of Figure 2.A.3 or Table 2.A.3 for details on the procedure. For each PS in our dataset, we identify the CHWs who competed for the PS position as those who were on-the-job at the time of the promotion and which we interviewed at baseline. In a dataset composed of all competing CHWs and the PS, we regress an indicator for "being promoted" (1 for the PS and 0 for the CHWs) on individual characteristics at the time of the promotion.

2.3 Beliefs Updating

In this section, we show that our treatments create exogenous variation in workers' perceptions about how meritocratic the promotion system is and about pay progression.

2.3.1 Beliefs about Meritocratic Promotions

To measure how workers updated their beliefs about meritocracy in the promotion system, we analyze CHWs' perceptions about meritocracy before and after we announced the introduction of the new promotion regime. We measure perceived meritocracy using a set of hypothetical survey questions. We asked each CHW which of the following workers she perceived as having a higher chance of being promoted: a CHW who ranks *first* out of 10 in terms of performance but who does not know the PHU in-charge outside of work vs. another CHW who ranks X out of 10 and who knows the PHU in-charge outside of work, where $X = \{2, 5, 10\}$.[17] Our measure of perceived meritocracy takes a value of -1, 0 or 1. It is coded as 1 if the CHW perceives the system as meritocratic, that is if she believes that the best-performing worker is always more likely to be promoted than the well-connected worker, regardless of whether the connected worker is ranked second, fifth or tenth. It is coded as -1 if the CHW perceives the system as non-meritocratic, that is if she believes that the best-performing worker is never promoted, even when the connected worker is the worst performer (ranked tenth). It is coded as 0 for intermediary situations in which the CHW believes that the best-performing worker is more likely to be promoted only when the well-connected worker has a low enough performance (ranked either fifth or tenth).

[17] The exact wording of the questions is: "*A PHU needs a new PS. Whom of the following two CHWs is most likely promoted to PS? (1) Alpha is the best-performing CHW (out of 10). Alpha does not know the PHU in-charge outside of work. (2) Foday is the second-best/ fifth-best/worst-performing CHW (out of 10). Foday is a very good friend of the PHU in-charge.*"

Figure 2.1 presents the distribution of meritocracy perceptions before and after treatment among CHWs in the meritocratic promotion treatment ($T_{merit} = 1$) and the rest ($T_{merit} = 0$). Consistent with randomization, perceptions are comparable in $T_{merit} = 1$ and $T_{merit} = 0$ before treatment (Panels A vs. C) with roughly 50% of CHWs perceiving the promotion system as meritocratic (prior of 1). A formal balance check of the perception variable is presented in Table 2.1 (Panel C).

After the introduction of the new promotion system, CHWs updated their beliefs upward in $T_{merit} = 1$, with an extra 28.4% of CHWs perceiving the system as meritocratic (Figure 2.1 Panels A vs. B). Interestingly, the CHWs who updated perception of meritocracy upward are those who had a prior of 0, while the 2.3% of workers with a more extreme prior of -1 did not update upward. In $T_{merit} = 0$, CHWs did not significantly update their perceptions (Panels C vs. D). The corresponding regression results are presented in Table 2.2 where we estimate the effect of the meritocratic promotion treatment on post-treatment perceptions, controlling for the stratification variables and clustering standard errors at the PHU level. Column (1) shows that the average perception of meritocracy in $T_{merit} = 1$ is 63% higher than in $T_{merit} = 0$ following treatment (statistically significant at the 1% level). Consistent with Bayesian models, CHWs whose prior of meritocracy is the highest in $T_{merit} = 1$ updated their beliefs less strongly (Table 2.A.4, column 1).

Interestingly, T_{merit} did not affect the expected time until the next promotion in the PHU, which is equal to 47 months in both treatment groups (Table 2.2, column 2).[18] It also did not affect perceptions about PS pay, PS workload (number of working hours), or PS work-related expenses (transportation and communication): see columns (3) to (5). In sum, the meritocratic promotions treatment appears

[18] These results should be taken as suggestive because 30% of the CHWs said they were not sure when the next promotion will take place. While this is not surprising - it is often hard to precisely predict a superior's future exiting behavior - this forces us to code the answer of these CHWs as missing, and to effectively run the regression on a potentially endogenous sample of CHWs.

to have changed perceptions about the promotion criteria (which is perceived as more performance-based), without affecting the perceived prize associated with the promotion and the perceived duration until the next promotion.

2.3.2 Beliefs about Pay Progression

Figure 2.2 plots the difference between perceived and true PS pay for CHWs in the pay progression treatment ($T_{pay} = 1$) and those not assigned to that treatment ($T_{pay} = 0$). To measure perceived PS pay, we asked each CHW: "*How much does your PS earn from the government each month?*" and offered a reward conditional on giving the right answer to elicit truthful responses.[19] We did not ask CHWs about perceptions of their own pay as this information was revealed to everyone at baseline, as explained in Section 2.2.2.

Consistent with the randomization, perceptions of PS pay are comparable in $T_{pay} = 1$ and $T_{pay} = 0$ before the treatment (Panels A vs. C). In both groups, roughly 30% of the CHWs knew that PSs earn 250,000 SLL per month. 37% of the CHWs underestimated PS pay and 33% overestimated it.[20] Table 2.3 (columns 3-4) shows that the size of the misperception about PS pay is unrelated with most CHW characteristics, except with the number of years of experience and age. Interestingly, the size of the misperception is comparable for CHWs who are connected to the PS or connected to the PHU in-charge relative to unconnected CHWs.

After receiving information about PS pay, almost all CHWs's beliefs in $T_{pay} = 1$ converged to the true PS pay (250,000 SLL). In contrast, few CHWs updated

[19] We offered a reward of 2,000 SLL if the answer is correct. In order to avoid revealing the true pay to CHWs who are not in the pay progression treatment, we disbursed the reward only at the end of the study period.

[20] Large misperceptions about supervisors' pay are common. In (Cullen and Perez-Truglia, 2021), for example, only 12% of respondents knew their manager's salary. In our context, large misperceptions about PS pay exist because this information is not publicized to CHWs. Additionally, discussions between colleagues about each other's pay is not the norm.

their beliefs in $T_{pay} = 0$, in which only 38% of the CHWs correctly guessed PS pay in our post-treatment survey. The absence of significant belief updating in $T_{pay} = 0$ corroborates the lack of information spillover across treatment groups. The corresponding regression results in Table 2.4 (column 1) show that the mean absolute difference between perceived PS pay and the truth is 482 SLL in $T_{pay} = 1$ vs. 35,320 SLL in $T_{pay} = 0$. Table 2.A.5 column (1) shows that, consistent with Bayesian models, CHWs update their beliefs more strongly the further their baseline perception about PS pay was from the truth. Column (2) shows that belief updating about PS pay is orthogonal to T_{merit}.

Throughout the paper, we will study the effect of T_{pay} in three separate groups of workers: (i) CHWs who underestimated PS pay at baseline, (ii) those who overestimated PS pay, and (iii) those with accurate beliefs. This is because these workers revised their beliefs in different directions in $T_{pay} = 1$, and are thus expected to respond differently to the treatment: the first group revised their perceptions of PS pay *upward* by 29,043 SLL (+13%), the second group revised them *downward* by 59,685 SLL (-19%), and the third group did not update their views significantly (Table 2.4, column 6). The magnitude of the update is smaller for the first than the second group because the level of CHW pay (150,000 SLL) provides a lower bound for perceptions.

In columns (7) and (8) of Table 2.4, we explore whether changes in CHWs' perceptions of PS pay were associated with changes in perceived PS workload (number of working hours) and PS work-related expenses (transportation and communication). Workers who revised their perception of PS pay downward did not change their perceptions in either area, while those who revised their perception of PS pay upward increased their estimates of PS work-related expenses slightly, but did not change their perceptions of the PS workload. Overall, this indicates that the pay progression treatment affected *perceptions of gross* PS pay as well as *net* PS pay (i.e., the PS pay accounting for total working hours and work expenses). Finally, columns (9) and (10) show that CHWs who update their beliefs of PS pay upward or downward did not change their perceptions about meritocracy in the promotion system or about the duration until the next promotion.

2.4 Theoretical Framework

Having established that our treatments had significant effects on CHWs' beliefs about meritocracy and pay progression, we now set up a simple model of promotion tournaments. The model provides a set of theoretical predictions on how workers respond to meritocratic promotions and pay progression that will guide our empirical analysis. These predictions are distinct from those of models studying non-tournament-based incentives because workers are rewarded based on their relative (rather than absolute) performance.

2.4.1 The Setup

Players Several Community Health Workers (CHWs) compete to be promoted to the position of Peer Supervisor (PS). They are risk-neutral and value the promotion in proportion to the pay progression from CHW to PS. The promotion mechanism is modeled as a single-prize contest, in which CHWs compete by exerting effort. In what follows, we study the case of two CHWs competing for the promotion. The case of N CHWs leads to similar predictions under additional mild assumptions.

The Promotion Tournament We are interested in a promotion tournament in which a principal can observe the effort of both workers, $(e_1, e_2) \in \mathbb{R}^2_+$, and can commit to a promotion rule that maps any effort pair to a promotion decision.

Since the promotion contest is characterized by this promotion rule, we start by specifying it.

We denote a *meritocratic promotion rule* by $P = (P_1, P_2)$ where $P_i : \mathbb{R}^2_+ \to [0,1]$ such that

$$(e_1, e_2) \to P_i(e_1, e_2) = \begin{cases} 0 & \text{if } e_i < e_{-i} \\ p & \text{if } e_i = e_{-i} \\ 1 & \text{if } e_i > e_{-i} \end{cases}$$

where $p \in (0,1)$ and $\sum_{i=1,2} P_i(e_1, e_2) = 1$. This promotion rule is the standard winner-take-all-allocation rule which has been extensively used in the promotion tournament literature (e.g., Lazear and Rosen, 1981; Siegel, 2010, 2014).

We are also interested in *non-meritocratic promotion rules*. Let $b = (b_1, b_2) \in \mathbb{R}^2$ denote the extent to which a promotion tournament is non-meritocratic. The b-biased contest is a promotion tournament characterized by $P^b = (P^b_1, P^b_2)$, where $P^b_i(e_1, e_2) = P(b_1 e_1, b_2 e_2)$.[21] Therefore, a promotion tournament is meritocratic if $b_1 = b_2$. If $b_1 \neq b_2$, the promotion rule favors one of the workers, and we will say that it is non-meritocratic.

Note that any b-biased contest is strategically equivalent to the $b' = (\frac{b_1}{b_2}, 1)$-biased contest. In what follows, we will use b to refer to contest $(b, 1)$. In this setting, the meritocratic contest is then simply the 1-biased contest. Implicitly, we also assume that any non-meritocratic contest favors player 1, i.e., $b \geq 1$. The case in which the contest favors player 2 ($b < 1$) is similar.

Payoffs The CHWs decide how much effort to exert. Effort is costly and each worker is characterized by a cost function of effort $c_i : \mathbb{R}_+ \to \mathbb{R}_+$. Workers exert

[21] All model's results hold if the bias is instead assumed to be additive, i.e., if $\tilde{P}^b_i(e_1, e_2) = P(e_1 + b_1, e_2 + b_2)$.

effort in the hope of being promoted, which increases their wage from $\underline{w}$ to $\bar{w}$. We refer to $\bar{w} - \underline{w} > 0$ as the *pay progression* associated with the promotion.

Given a promotion rule P^b and an effort pair (e_1, e_2), player i's payoff is

$$u_i(e_1, e_2) = \underline{w} + P_i^b(e_1, e_2)\,[\bar{w} - \underline{w}] - c_i e_i. \tag{2.1}$$

The payoff is a function of how meritocratic the promotion rule is (P^b), the pay progression ($\bar{w} - \underline{w}$), and the cost of effort $c_i > 0$ which is assumed to be linear.[22] We define worker i to have higher ability than worker i' if $c_i \leq c_{i'}$.

The model is divided into two parts. We first consider the cost function, c_i as independent of pay progression $\bar{w} - \underline{w}$ and meritocracy b (Section 2.4.2). We then extend the model by assuming that workers display *morale concerns* and that their costs instead depend on pay progression $\bar{w} - \underline{w}$ and meritocracy b (Section 2.4.3). This assumption is motivated by recent empirical evidence showing that morale concerns about pay differences and unfair promotions negatively affect effort within the workplace (Card et al., 2012; Cohn et al., 2014; Mas, 2017; Breza et al., 2018; Li, 2020). As such, we hypothesize that workers perceive a high pay progression (high $\bar{w} - \underline{w}$) in a non-meritocratic regime (high b) as unfair, leading to higher perceived costs. This is modeled by adding an extra morale cost-shift function $g_i : \mathbb{R}_+^2 \to \mathbb{R}_{++}, (b, \bar{w} - \underline{w}) \mapsto g_i(b, \bar{w} - \underline{w})$ in player i's payoff:

$$u_i(e_1, e_2) = \underline{w} + P_i^b(e_1, e_2)[\bar{w} - \underline{w}] - c_i g_i(b, \bar{w} - \underline{w}) e_i \tag{2.2}$$

The addition of the morale cost-shift function will only be consequential for a subset of the results, while other results will hold regardless. This will be made clear later in the model.

[22] The assumption of cost linearity is common in the literature on promotion rules (e.g., Nti, 2004; Franke, 2012; Franke et al., 2013) and can be relaxed in the model. Most of the results indeed hold if we assume convex costs and make minimal assumptions on the cost elasticities.

Throughout, we assume that the participation constraints of both players are satisfied. We are interested in Nash equilibria in which no players play a weakly dominated action with positive probability. See Appendix 2.D for a more formal and detailed exposition of the model.

2.4.2 Predictions without Morale Concerns

This section studies the b-biased contest ($b \geq 1$) with pay progression $\bar{w} - \underline{w} > 0$ when there are no morale concerns for any player. The morale cost-shift function is thus normalized to 1 for both players i.e., $g_i(b, \bar{w} - \underline{w}) = 1$ for all $b, \bar{w} - \underline{w}$, and i.

Following (Siegel, 2010), the b-biased promotion tournament with effort costs (c_1, c_2) has a unique equilibrium in mixed strategies. From Propositions 1 - 7 presented in Appendix 2.D.1, we obtain the following predictions for all players:

Prediction 1 *All else equal, more meritocratic promotions (lower b) increase worker's effort.*[23]

Prediction 2 *All else equal, higher pay progression (higher $\bar{w} - \underline{w}$) increases worker's effort.*

Prediction 3 *The effect of higher meritocracy (resp., pay progression) on worker's effort increases as pay progression (resp., meritocracy) increases.*

Prediction 4 *The effort response in Predictions 1 - 3 is stronger for higher-ability workers.*

[23]The increase in effort for the average worker is larger in a model with 2 players (like ours) than in a model with many players. This is because the increase in effort is stronger for high ability (high ranked) workers (see Prediction 4) and the average effect thus decreases with the number of workers who are not "high ranked." In Section 2.5, we show that in teams of 8 health workers, the effect of meritocracy on the average worker are positive but not significant. Refer to (Boudreau et al., 2016) for empirical evidence in a lab setting that high-ability workers respond more strongly to promotion incentives.

See Appendix 2.D.1 for details on the propositions and Appendix 2.D.2 for their proofs.

Note that the intensity of the effort response described in the Predictions 1-3 is comparable for players 1 and 2 as long as their costs are symmetric. See Appendix 2.D.1 for more details.

2.4.3 Predictions with Morale Concerns

This section derives the model's results under the assumption that workers display *morale concerns*, which we model by adding an extra morale cost-shift function $g_i : \mathbb{R}^2_+ \to \mathbb{R}_{++}, (b, \bar{w} - \underline{w}) \to g(b, \bar{w} - \underline{w})$ in workers' payoffs.

We make three assumptions about g_i. Each of these are explained intuitively below and formally presented in Appendix 2.D. The first assumption is that the only player who faces morale concerns is the "unfavored" player (2), i.e., $g_1(b, \bar{w} - \underline{w}) = 1$ for all $(b, \bar{w} - \underline{w}) \in \mathbb{R}^2_+$. This assumption is made for simplicity and the results that follow hold if g_1 was instead decreasing in both of its arguments. The second assumption is that a more-biased contest, or a contest with higher pay progression, increases the morale cost-shift function for player 2, and does so in a log-supermodular way.[24] Finally, we assume that for a higher pay progression $\bar{\bar{w}} - \underline{\underline{w}} > \bar{w} - \underline{w}$, $g_2(b, \bar{\bar{w}} - \underline{\underline{w}})$ dominates $g_2(b, \bar{w} - \underline{w})$, and therefore that the morale cost-shifts increase faster in the bias when the pay progression is higher.

Given these assumptions, we can rewrite the players' payoffs as:

$$u_1(e_1, e_2) = \underline{w} + P_1^b(e_1, e_2)[\bar{w} - \underline{w}] - c_1 e_1$$

$$u_2(e_1, e_2) = \underline{w} + P_2^b(e_1, e_2)[\bar{w} - \underline{w}] - c_2 g_2(b, \bar{w} - \underline{w}) e_2$$

[24]Log supermodularity implies that the morale cost-shift function becomes less elastic in b as the pay progression increases.

From Propositions 8 - 13 presented in Appendix 2.D.1, we obtain the following predictions for all players:

Prediction 5 *All else equal, more meritocratic promotions (lower b) increase worker effort.*

Prediction 6 *All else equal, higher pay progression (higher $\bar{w} - \underline{w}$) increases worker effort if the promotion rule is meritocratic enough ($b \leq \bar{b}$), while it reduces effort if the promotion rule is non-meritocratic enough ($b \geq \bar{\bar{b}}$).*

Prediction 7 *The effect of higher meritocracy (resp., pay progression) on worker's effort increases as pay progression (resp., meritocracy) increases if $b \leq \bar{b}$.*

Prediction 8 *The effort response in Predictions 5 - 7 is stronger for higher-ability workers.*

See Appendix 2.D.1 for a formal definition of $\bar{b}$ and $\bar{\bar{b}}$ and for details on the propositions, and Appendix 2.D.2 for the proofs.[25]

The theoretical framework makes clear that the addition of morale concerns does *not* affect the direction of workers' reactions to meritocracy: higher meritocracy in the promotion rule always increases worker effort, regardless of the presence of morale concerns (Predictions 1 and 5). The addition of morale concerns, however, does affect the direction in which workers respond to pay progression. Without morale costs (g_i), greater pay progression always boosts workers' effort regardless of how meritocratic the promotion rule is (Prediction 2). With morale costs (g_i), greater pay progression boosts workers' effort only if the promotion rule is meritocratic enough, while it*reduces* worker effort if the rule is not meritocratic

[25]The intensity of the effort response described in Prediction 5 is comparable for players 1 and 2 as long as their costs are symmetric. For Predictions 6 and 7, the relative intensity of the effort response is theoretically ambiguous, and therefore not explored empirically. See Appendix 2.D.1 for more details.

(Prediction 6).[26] We will later show that, empirically, the effect of pay progression is consistent with Prediction 6 rather than Prediction 2, and thus consistent with the presence of morale concerns.

Finally, note that Prediction 6 can be obtained in an alternative multitasking model (without morale concerns) in which workers not only choose how much effort to exert on productive tasks $e_i \in \mathbb{R}_+$ but also choose whether and how much to lobby their principal for the promotion (unproductive task): $l_i \in \mathbb{R}_+$.[27] If productive effort (e_i) and lobbying (l_i) are substitutes, such a model predicts that if the promotion rule is not meritocratic enough, greater pay progression reduces productive effort while increasing lobbying effort. We do not focus on this alternative model since it is proven to be inconsistent with the empirical results in Section 2.6.2.

2.5 The Effect of Meritocratic Promotions on Worker Productivity

The main results of this paper are divided in two sections. In this section, we study the effect of greater meritocracy in the promotion system on worker productivity, while shutting down any effect of pay progression. To do so, we compare the

[26]Intuitively, morale concerns introduce a tension when assessing the effect of pay progression on productivity. Steeper pay progression raises the effective prize for any given level of effort, which prompts player 2 to exert more effort. At the same time, it leads player 2 to perceive the promotion tournament as more unfair, which increases the effective costs and reduces her effort. Morale concerns instead unambiguously amplify the effect of meritocracy on productivity. A more biased tournament decreases the likelihood that player 2 wins the contest (and therefore reduces the effective prize for any given level of effort), and it increases morale concerns (and therefore increases the cost of effort).

[27]Imagine that the principal promotes the worker who obtains the highest score $s_i^\alpha = \alpha e_i + (1-\alpha) l_i$, where $\alpha \in \mathbb{R}$ captures how efficient lobbying is in getting the promotion, then the CHWs compete by simultaneously and independently choosing a score $s_i^\alpha \in \mathbb{R}_+$. Given the scores (s_1^α, s_2^α), CHW i's payoff becomes $u_i(s_1^\alpha, s_2^\alpha) = \underline{w} + P_i(s_1^\alpha, s_2^\alpha)\,[\bar{w} - \underline{w}] - \min_{e_i, l_i | \alpha e_i + (1-\alpha) l_i = s_i^\alpha} c_i(e_i, l_i)$.

productivity of workers in $T_{merit} = 1$ vs. $T_{merit} = 0$ restricting the sample to the 186 PHUs where CHWs received no information on the pay gap ($T_{pay} = 0$). In the next section, we study the interactions of meritocracy and pay progression by leveraging the 2×2 design in the full sample of workers.

From Predictions 3-4 and 7-8 of our theoretical framework, we expect the effect of our meritocratic promotion treatment to be concentrated among two types of workers: (1) workers who perceive the prize associated with the promotion to be large enough to be interested in the promotion, and (2) workers who are highly ranked in terms of performance (i.e., high ability), as they have a higher chance of being promoted in a meritocratic regime. To test this, we estimate the following equation:

$$Y_{ij} = \alpha + \beta_1 T_{merit,j} \times X_{ij} + \beta_2 T_{merit,j} \times (1 - X_{ij}) + \gamma X_{ij} + Z_j \gamma + \varepsilon_{ij}, \quad (2.3)$$

where Y_{ij} represents the performance of CHW i in PHU j, $T_{merit,j}$ is a dummy for whether the PHU j is assigned to the meritocratic promotion treatment, X_{ij} is a dummy for whether workers have a high perceived pay progression or a high ranking at baseline, Z_j are the stratification variables. We estimate standard errors clustered by PHU, and report p-value corrections for multiple hypothesis testing.

Our main measure of worker performance is the total number of visits that households report having received from the CHW in the six months prior to the endline survey (mean of 7.9).[28] To obtain this measure, we take the total number of times a household has received a routine visit, ante- or post-natal visit, or has been treated/referred for sickness, and then average these data at the CHW level. We will later also present results on the length of the visits (mean of 15 minutes) -

[28]To minimize recall bias, households were asked about visits received "since the start of the year" which corresponds to the past 6 months.

which we will use as a proxy of work quality - and on retention.

The results are presented in Table 2.5 and Figure 2.3. For completeness, we start by assessing the effect of T_{merit} on the average worker's productivity by estimating the non-interacted version of equation (2.3). We find that making the promotion system more performance-based raises the number of visits provided by the average CHW by 0.932 (12.5%), but this effect is not statistically significant (column 1 of Table 2.5 and first bar of Figure 2.3). Table 2.A.6 breaks down the result by type of visit and shows that CHWs treat significantly more patients in $T_{merit} = 1$, while other types of visits increase, but not significantly.[29]

The remainder of the section presents the heterogeneous effects of meritocracy by perceived pay progression and performance ranking using equation (2.3). The analysis of these heterogenous effects was pre-registered (see Section 2.C for more details). Because we study multiple heterogeneous effects, we will correct our p-values for multiple hypothesis testing.

Heterogeneous Effect by Perceived Pay Progression In columns (2)-(3) of Table 2.5, we estimate equation (2.3) with X_{ij} defined as a dummy variable for whether the worker's perceived pay progression is above the median, that is above the actual rate of 250,000 SLL. Consistent with the model, the effect of meritocracy on worker productivity increases with perceived pay progression. The effect of meritocracy on the number of visits is strong and significant for the CHWs with a high (above-median) perceived pay progression ($\hat{\beta}_1 = 2.014$, a 27% increase), while no effect is detected among workers with a low (below-median) perceived pay progression ($\hat{\beta}_2 = 0.323$, not statistically significant). The difference between $\hat{\beta}_1$ and $\hat{\beta}_2$ is statistically significant at the 10% level (p-value reported at the

[29]Table 2.A.7 presents the elasticity of CHW performance (number of visits) with respect to meritocracy by instrumenting CHW post-treatment perceived meritocracy with the meritocratic promotion treatment. We find that a one-unit increase in perceived meritocracy (on a scale of -1 to 1) raises the number of visits by 3.235.

bottom of the table). These results remain statistically significant with p-values corrections for multiple hypothesis testing.[30]

Importantly, the variation in perceived pay progression we leverage in equation (2.3) is not random. As discussed in Section 2.3.2, misperceptions about the PS pay are correlated at baseline with age and experience. In column (3), we show that our results are robust - and even become slightly more precise - when we further control for these correlates and their interaction with T_{merit} in equation (2.3). The heterogeneity in the effect of T_{merit} we attribute to perceived pay progression is thus unlikely to be explained by variation in age and experience.[31]

In the next section, we study the causal effect of pay progression by leveraging *random variation* in perceived pay progression.

So far, we have proxied for the perceived prize associated with a promotion with CHWs' prior about pay progression. An alternative strategy is to assess how likely the PS is to leave her position in the near future. Holding perceived pay progression fixed, CHWs who expect a PS to leave her position soon should have a higher present value of the prize associated with the promotion and therefore respond more strongly to the meritocracy treatment.

We explore this heterogeneity in Table 2.5 (column 4), where we proxy the likelihood that the PS will leave her position soon by an indicator for whether the supervisor is within five years of the standard retirement age (that is, above 50 years old). Using this definition, 10% of the CHWs in our sample have a supervisor who is likely to retire soon. For these workers, making promotions more performance-based increases the number of visits by 4.894 (a 66% increase, statis-

[30] Figure 2.A.4 (Panel A) presents the effect of the meritocratic promotion treatment on the number of visits by quintiles of prior PS pay. The difference in productivity between $T_{merit} = 1$ and $T_{merit} = 0$ is positive and statistically significant only among workers in the top quintile.

[31] The magnitude of the results is unaffected if we control for the entire list of CHW-level variables presented in Table 2.1 and their interaction with T_{merit}, but we lose precision due to the addition of 30 covariates in the regression.

tically significant at the 1% level). In contrast, higher meritocracy has no effect on workers who are unlikely to experience a promotion in the next five years. The difference in the effect of meritocracy for these two types of workers is statistically significant at the 1% level and is robust to controlling for correlates of the supervisor's age (column 5). Table 2.A.8 (columns 1-6) shows that, as expected, the results attenuate as the PS is expected to retire further in the future: the effect shrinks by half but remains positive when the PS is within 10 years of retirement, while it disappears when the PS is within 15 years of retirement. Column 7 tests for heterogeneous effects based on whether the CHW's perception of the duration until the next promotion is above or below the median, and shows that the productivity boost is 10 times larger for the latter, but this result is imprecisely estimated.[32]

Table 2.A.8 (columns 9-10) expands the heterogeneous effects to four types of workers, depending on whether their priors of PS pay are high *and* whether the promotion is expected within five years. The effect of meritocratic promotions on worker performance is small and not significant for workers for whom a promotion is unlikely to occur within the next five years, regardless of the prior of PS pay. Among workers for whom a promotion is more likely to occur within the next five years, those with a high prior respond very strongly (they double the number of visits provided), while those with a low prior respond more moderately (the number of visits increases by 36%).

[32] As explained in footnote 18, a nontrivial fraction of CHWs is unable to precisely predict when the PS position will become vacant. To avoid excluding this potentially endogenous sample of workers from the regression, we make the assumption that their expected time until the next promotion is above the median. Within the rest of the workers (who gave us an answer), an extra year until supervisor's retirement age is associated with an increase in their perceived duration until the next promotion of 1.1 years. This indicates that workers who report knowing when the PS will leave her position are probably implicitly assuming that the PS will exit at retirement age.

Heterogeneous Effect by Performance Ranking As explained above, we expect the effect of meritocracy to be stronger among high-ranked workers, as they have a higher chance of being promoted in a meritocratic regime. Our preferred measure for the ranking of each CHW within the PHU is the one provided by the PS at baseline. The PS has indeed frequent interactions with all CHWs and is in the best position to compare and rank her subordinates. The PS also has no incentive to misreport the ranking because she does not decide on promotions (the PHU in-charge does). Table 2.3 (columns 9-10) shows that the ranking - as reported by the PS - is correlated with variables that we expect to predict performance: health knowledge, education, years of experience, and number of household visits reported by the CHW. It is also correlated with the number of years the CHW has known the PS, a variable we will later control for, while it does not correlate with connections to the PHU in-charge (the number of years the CHW has known the PHU in-charge) or with the CHWs' perceived PS pay at baseline.

Table 2.5 (column 6) reports the coefficients $\hat{\beta}_1$ and $\hat{\beta}_2$ estimated from equation (2.3) with X_{ij} defined as a dummy variable for whether the worker is ranked among the top three of her PHU (henceforth, "high rank" workers). Increasing the meritocracy of the promotion system significantly boosts the number of visits provided by high-ranked workers ($\hat{\beta}_1 = 2.251$, a 30% increase), but does not affect the productivity of lower-ranked workers ($\hat{\beta}_2 = 0.066$, not statistically significant). The difference between $\hat{\beta}_1$ and $\hat{\beta}_2$ is statistically significant at the 5% level.

Figure 2.A.4 (Panel B) breaks down the results for workers ranked 1-3, 4-6, 7-9. The effect of meritocracy is positive for workers ranked 1-3, and zero for workers ranked above 4. Note that the effect of meritocracy for low-ranked workers is not negative. This is presumably because these workers had only weak incentives to exert effort in the old non-meritocratic system and have equally weak incentives

in the new meritocratic system (as they have limited chance of promotion).

Table 2.5 column (7) shows that the results are robust to further controlling in equation (2.3) for the variables that are significantly correlated with a worker's ranking in Table 2.3 and their interaction with T_{merit}. This ensures that the observed heterogeneous effects are driven by the performance ranking, rather than other observable characteristics. The results are also robust, though less precise, if we measure the ranking of each CHW as reported by other CHWs in the PHU rather than as reported by the PS (Table 2.A.8, columns 11-12).[33]

Overall, the results of this section show that meritocratic promotions boost the productivity of "top" workers - who have a chance of being promoted under the new meritocratic regime - while there is no effect on the rest of the workforce. The fact that only a selected sample of workers react to meritocratic promotions is consistent with the tournament structure of promotion incentives, which exclusively rewards the top worker.[34]

Other Outcome Variables: Visit Length, Targeting and Worker Retention We have shown that the effect of our meritocratic promotion treatment raises the number of visits for workers who perceive the prize associated with the promotion to be large enough and those who are highly ranked. In Table 2.6 (columns 1-7), we test for the possibility that these CHWs compensate for the higher number of visits by providing shorter visits, i.e., by skipping some of the checklist items they are supposed to follow and thus presumably reducing visit quality. We find that visit length of the average worker increases by 12.3% (sta-

[33]The ranking as reported by other CHWs is positively and significantly correlated with the PS ranking. While CHWs may not be as good as the PS in ranking their colleagues, this indicates that CHWs do have an idea of what the ranking looks like, even in the old promotion regime where effort is not incentivized as much as in the new system. This is not surprising as CHWs know each other and regularly attend trainings together.

[34]The incentive for the very best workers to exert effort may be weaker in contexts in which the incentives do *not* have a tournament structure, as those analyzed in e.g., (Lazear, 2000; Muralidharan and Sundararaman, 2011; Khan et al., 2016).

tistically significant at the 10% level), and that this is mostly driven by workers with high perceived pay progression, a promotion expected soon or a high ranking. This is consistent with workers being aware that the quality of the visits matters for promotions in $T_{merit} = 1$, as explained in Section 2.2.2.

The higher number of visits may also potentially be compensated by CHWs targeting only households who live nearby or those who are friends or family members (and who are thus presumably less costly to reach), at the expense of other more deserving households. Table 2.A.9 shows that this is not the case: targeting by physical or social distance does not change with T_{merit}.

Table 2.6 (columns 8-14) presents the effect of meritocracy on worker retention, as measured by whether the CHW self-reports not having dropped out and provided at least one visit to surveyed households in the six months before the endline survey. According to this definition, the retention rate in our sample is 89%. Column (9) shows that higher meritocracy increases the retention of workers with high perceived pay progression by 7.9 percentage points (from 88% in $T_{merit} = 0$ to 96% in $T_{merit} = 1$). In contrast, it does not affect retention for workers with low perceived pay progression. Similarly, column (13) shows that our meritocracy treatment increases the retention of high-ranked workers by 5.4 percentage points, while it does not affect the retention of low-ranked workers.[35]

The positive effect of meritocracy on the retention of workers who have high perceived pay progression or who are highly ranked raises the question of whether the increase in visits provided by these workers is driven by selection (i.e., meritocracy increasing the retention of the most productive of these workers) or by higher effort of those retained. To separate the two, we perform a bounding exercise. Assuming that the increase in retention in the meritocratic regime comes from

[35]This might be the case because high-ranked workers have better outside options and become frustrated if they do not see opportunities for career progression in absence of a fully meritocratic promotion system. We further explore this "demotivation effect" in the next section.

workers belonging to the top or bottom decile of the productivity (visits) distribution, and using the estimates identified earlier, we calculate that the direct effect of meritocracy on the number of visits provided by workers with high perceived pay progression - net of selection - is between 1.28 and 2.52 (which correspond to a 17% and 34% increase, respectively).[36] For high ranked workers, the direct effect is between 1.39 and 2.35 (which correspond to a 19% and 32% increase, respectively). This indicates that the "on-the-job" effort response of these workers are non-trivial, even in the lower bound scenario.

Alternative Mechanisms In our model, the increase in the performance of workers with high perceived pay progression and high ranking in the meritocratic promotion treatment is explained by these workers exerting more effort in anticipation of a future promotion, due to a greater interest in the promotion (for the former) or a higher chance of being promoted (for the latter).

An alternative story is that these workers become more productive because supervisors start monitoring them more than other workers. Table 2.A.10 rejects this possibility by showing that the PSs did not adjust their effort in the meritocratic system relative to the old system: the likelihood that they visited a CHW or accompanied them on a household visit is unchanged across all workers types.

Another story consistent with our results is that the boost in the productivity of workers with high perceived pay progression or high ranking is due to these workers

[36] Assuming that productivity (Y) is a function of both meritocracy (M) and retention (R), which itself is a function of M, the elasticity of worker productivity with respect to meritocracy can be written as: $\frac{dY}{dM} = \frac{\delta Y}{\delta M} + \frac{\delta Y}{\delta R} * \frac{dR}{dM}$, where $\frac{dY}{dM} = 2.073$ and $\frac{dR}{dM} = 0.077$ for workers with high perceived pay progression (Table 2.5 column 3 and Table 2.6 column 10, respectively). $\frac{\delta Y}{\delta M}$ is the behavioral response of interest, namely the direct effect of meritocracy due to changes in effort; and $\frac{\delta Y}{\delta R}$ is the change in productivity of the marginal retained worker. We obtain the bounds for $\frac{\delta Y}{\delta M}$ by assuming that the productivity gain from the marginal retained worker corresponds to the difference between the 90^{th} or 10^{th} percentile of the productivity distribution - which correspond to 17.67 or 1.67 visits, respectively - and the average productivity in the control group (7.46 visits).

revising their perceptions of meritocracy more strongly. Table 2.A.4 (columns 2-4) shows that this is not the case.

2.6 The Effect of Pay Progression on Worker Productivity

Having established that a meritocratic promotion system boosts productivity of CHWs who believe that pay progression is large at baseline, we now assess the causal effect of a change in perceived pay progression on CHW productivity in the status-quo (non-meritocratic) promotion regime and in the new (meritocratic) promotion regime.

Estimating the effect of the pay progression treatment on the productivity of the average worker is uninformative in our setting because it pools together workers who over- or underestimate PS pay at baseline, who revise their beliefs in opposite directions, and who have opposite reactions to the treatment.[37] In our main specification, we estimate the treatment effects in three separate samples of workers: (i) those with priors of PS pay below the actual pay level at baseline (who revise their beliefs upward), (ii) those with priors above the actual pay level (who revise their beliefs downward), and (iii) those with accurate priors (no revision):

$$Y_{ij} = \alpha + \beta_1 T_{pay,j} \times T_{merit,j} + \beta_2 T_{pay,j} \times (1 - T_{merit,j}) + \gamma T_{merit,j} + Z_j \delta + \varepsilon_{ij}. \quad (2.4)$$

For workers with perceived PS pay below (above) the truth at baseline, β_1 and β_2 capture the causal effect of increasing (decreasing) perceived pay progression on productivity in a high meritocracy regime ($T_{merit} = 1$) and a low meritocracy

[37] For completeness, Table 2.A.11 reports the results pooling all workers together, regardless of their baseline priors, but these are hard to interpret.

regime ($T_{merit} = 0$), respectively. γ captures the effect of a more meritocratic system when $T_{pay} = 0$, which was the focus of the previous section and which we do not discuss again.[38]

Instead of estimating (2.4) in different sub-samples of workers, one can alternatively estimate a fully interacted equation with triple interactions $T_{pay} \times T_{merit} \times \mathbf{1}(Perceived\ PS\ pay \lesseqqgtr Truth)$. We do not use this model as our main specification because comparisons across worker types (for example, between workers who underestimate or overestimate PS pay at baseline) are not necessarily causal in our empirical design. Table 2.A.12 (columns 1-2) shows for example that, relative to workers who underestimate PS pay (Panel A), those who overestimate it (Panel B) have half a year of experience more and are one year older, and this may affect their effort response. We focus instead on assessing the effect of raising pay progression in meritocratic and non-meritocratic regimes *within a worker type*, for which we can confidently claim that our estimates are causal.[39]

2.6.1 Pay Progression in Meritocratic Regimes

In this section, we assess the effect of pay progression on worker productivity in the new meritocratic system ($T_{merit} = 1$). The next section presents the corresponding effects in the old non-meritocratic system ($T_{merit} = 0$).

Predictions 2 and 6 of our theoretical framework say that when the promotion system is meritocratic enough ($b < \bar{b}$), raising (reducing) pay progression $\bar{w} - \underline{w}$ should *boost* (*reduce*) worker productivity. In line with this, Figure 2.4 (first and third bars) and the corresponding Table 2.7 (row [i]) show that, within the sample of workers who revise their perception of pay progression upward, the number of

[38] The estimates for γ are reported in Table 2.A.11, row [iii] (columns 1, 4, 7 and 10).

[39] Table 2.A.12 (columns 3-8) shows that CHWs' characteristics are balanced across treatments within a worker type. For completeness, we report the results of the fully interacted model in Table 2.A.11, in which we control for all CHW characteristics interacted with the treatments, but we will not discuss the results of this table in the main text.

visits provided goes up by 1.871 (24%). Within the sample of workers who revise their perception downward, the number of visits instead goes down by 2.062 (26%). Both of these results are consistent with standard theory of promotion/career incentives, i.e., worker effort moves in the same direction as the perceived pay gap. Within the sample of workers whose priors were equal to the truth at baseline (and who did not update their beliefs about the pay gap), the number of visits did not change. This is reassuring as it indicates that providing information about true PS pay unlikely affects workers' behavior through channels unrelated to a reassessment of their prior beliefs.[40]

Consistent with Prediction 4 of our theoretical framework, the effect of pay progression on worker productivity is found to be more pronounced for higher-ranked workers, who have greater chances of being promoted in a meritocratic regime, while the effect is muted for lower-ranked workers (Table 2.8, columns 3-6, rows [i] and [ii]).

Finally, Table 2.A.15 (column 1) computes the elasticity of CHW performance with respect to PS pay. To do so, we use the entire sample of workers and instrument the updating of CHWs' beliefs about PS pay with $T_{pay} \times \mathbf{1}(Perceived PS pay < Truth)$ and $T_{pay} \times \mathbf{1}(Perceived\ PS\ pay > Truth)$.[41] Revising PS pay upward by 10% (25,518 SLL) increases the number of visits provided by the average CHW by

[40] Table 2.A.13 shows that pay progression does not significantly impact visit length (columns 1 and 4, row [i]) but it does affect retention. Higher perceived pay progression increases retention by 8.7 percentage points (column 2, row [i]). Lower perceived pay progression instead reduces retention by 4.8 percentage points, albeit not significantly (column 5, row [i]). As before, PS behavior is unaffected by changes in CHW perceived pay progression (columns 3 and 6, row [i]) and pay progression does not affect household targeting by physical or social distance (Table 2.A.14).

[41] Using this approach, the Cragg-Donald F-statistic is around 180. If we only use T_{pay} as an instrument, we predictably obtain a low first stage, as workers update in opposite directions depending on whether they over- or underestimate PS pay at baseline. Alternatively, we could split the sample by whether the CHW over- or underestimates PS pay at baseline, and use T_{pay} as an instrument for the perceived PS pay following the treatment (rather than using the extent to which they updated perceptions). The results are shown in Table 2.A.15 (columns 2-3) and are discussed later.

9.4% (0.028*25.518/7.560), giving us a cross-wage elasticity of 0.94.[42]

Overall, the results indicate that extrinsic incentives in the form of a potential future higher pay play an important role even for public sector workers who have been argued to be more "intrinsically motivated" (Besley and Ghatak, 2005; Bénabou and Tirole, 2006).

2.6.2 Pay Progression in Non-Meritocratic Regimes

We now turn our attention to the effects of pay progression in a non-meritocratic regime ($T_{merit} = 0$). Figure 2.4 (second bar) and the corresponding Table 2.7 (columns 1, row [ii]) show that higher pay progression reduces the number of visits provided by CHWs by 1.982 (26%). This suggests that the combination of a steep pay progression and a promotion regime with low meritocracy, commonly seen in the public and private sectors,[43] can be detrimental to the productivity of workers at the bottom of the organization.

Two potential channels can explain the observed reduction in worker productivity. The first is the negative morale effect proposed in Section 2.4.3 of our theoretical framework: workers may become less motivated and provide fewer visits if they perceive a non-meritocratic organization as being unfair or unequal when increasing its pay progression (Prediction 6). The second channel is one of multitasking and lobbying: when pay progression increases, workers may become more interested in a promotion and may start devoting more time to lobbying (e.g.,

[42]This is not a trivial elasticity in comparison to the own-wage labor supply elasticity of 1.12-1.25 identified in the experimental literature (Fehr and Goette, 2007). The only other estimate of vertical cross-wage elasticity in the literature is provided by (Cullen and Perez-Truglia, 2021). They document that raising the perceived salary of a manager by 10% increases the number of hours worked by lower-tier employees by 4.31% when these employees are told that the manager position is attainable. Their elasticity might be lower than ours because they use different metrics for performance and (perhaps more importantly) because their promotion system may not be as meritocratic as the system in our meritocratic promotion treatment.

[43]Refer to the discussion in the Introduction.

talking with the PHU in-charge) so as to increase their chances of promotion in a non-meritocratic regime. If lobbying and productive effort are substitutes, this behavior would reduce the number of visits because the extra time spent on lobbying would crowd out time spent on productive tasks (visits).[44]

Two pieces of evidence provide suggestive evidence that the reduction in worker productivity we find in the data is more likely driven by a demotivational effect caused by morale concerns than by workers spending more time lobbying. First, we find limited evidence of increased lobbying when pay progression increases. Lobbying is inherently hard to measure, as it can take different forms, but should presumably entail CHWs being more likely to interact with the PHU in-charge. At endline, we asked CHWs whether they had talked to the PHU in-charge in the past year. While an average of 54% had done so, this variable did not increase with pay progression (Table 2.8, column 1). Moreover, we asked CHWs what fraction of their time as a CHW was dedicated to non-patient-related activities, which include communications with the PHU in-charge (mean of 21%). Once again, we document no effect of the pay progression treatment on this variable (Table 2.8, column 2).

Second, we find that the negative effect of pay progression on worker productivity is stronger among the two types of workers who presumably perceive the combination of pay progression and non-meritocracy as the most unfair: high-ranked workers, who would be the first to benefit from the steeper pay progression under a meritocratic regime, and workers who are unsatisfied with the work of the PS, who should find a steep vertical pay gap as more unjustified. Table 2.8 shows that high-

[44] de Janvry et al., 2021 defines this type of lobbying as an "un-productive influence activity." Another type of un-productive influence activity would consist in CHWs bribing the PHU in-charge to get the promotion. This could reduce the number of visits if bribing forces the CHW to devote more time to another secondary job in order to raise the money. This is unlikely in our context because bribes and side-payments across the different layers of the organization are minimal (Deserranno et al., 2022a).

ranked workers and those unsatisfied with the PS react to the increase in perceived pay progression by providing 2.511 and 3.231 fewer visits respectively (columns 3 and 5, row [iii]). These demotivational effects are instead much smaller (and often not statistically significant anymore) for lower-ranked workers and workers who are satisfied with the work of their PS (row [iv]). These heterogeneous results are robust to controlling for all observed CHW characteristics and their interaction with the treatment dummies (Table 2.8, columns 4 and 6). This ensures that the heterogeneity in the treatment effects we are attributing to ranking and satisfaction with the PS is likely not due to variation in other observables. Table 2.A.5 (columns 6 and 7) shows that the larger reduction in effort observed among CHWs who are high ranked or unsatisfied with their PS is not explained by these workers updating their beliefs about pay progression more strongly than other workers.

Table 2.A.15 presents IV results in which CHWs' post-treatment perception of PS pay is instrumented by T_{pay}, separately for the subsample of workers who overestimated PS pay at baseline and those who underestimated it. Column 2 (row [ii]) shows that, in the non-meritocratic regime, workers who perceive the level of PS pay as being 10% higher (23,571 SLL higher) provide 19% fewer visits (-0.061*23.571/7.560), leading to an elasticity of -1.9. This level of elasticity of vertical pay inequalities in non-meritocratic regimes is large relative to what the literature has identified as the demotivational effect created by horizontal pay inequality across peers (Breza et al., 2018; Cullen and Perez-Truglia, 2021).[45] It is however smaller than the demotivational effect created by mass layoffs or pay cuts (Akerlof et al., 2020; Coviello et al., 2021).

Finally, the last bar of Figure 2.4 and Table 2.7 (column 2, row [ii]) show that a downward update of beliefs about pay progression has a precisely estimated zero

[45] Cullen and Perez-Truglia, 2021 find that a 10% increase in employees' perception of their peers' salaries decrease the number of hours they work by 9.4%, leading to an elasticity of -0.94. (Breza et al., 2018) show that when coworkers' productivity is difficult to observe, horizontal pay inequality reduces output by 0.45 standard deviations and attendance by 18 percentage points.

effect on worker productivity. This may indicate that a reduction in perceived pay progression in a system that is non-meritocratic does not make workers more likely to perceive the system as fair, or at least does not increase it by enough to raise worker productivity.

2.7 Conclusion

Despite the popular definition of organizations as "pyramids of opportunities" (Alfred P. Sloan) and the wide attention that promotions have received both in the theoretical literature (e.g., Lazear and Rosen, 1981; Waldman, 1984; Gibbons and Waldman, 1999b) and in public policy (e.g., McKinsey, 2015; World Bank, 2018), empirical evidence on promotion incentives is scarce. This paper fills this gap by providing the first experimental evidence on the causal effect of meritocratic promotions and pay progression on worker productivity.

We collaborated with the Ministry of Health and Sanitation in Sierra Leone to introduce exogenous variation in (i) the extent to which the promotion process from frontline workers (lower-tier) to supervisor (upper-tier) is meritocratic or not, and (ii) the perceived gap between these two positions. Our findings show that promotion systems should have two components to maximize the productivity of frontline workers: promotions based on performance (meritocratic) *and* a large enough pay progression associated with promotions. Crucially, raising the extent to which promotions are meritocratic causes an increase in worker productivity only if combined with a high enough pay progression, otherwise the effect is muted. A higher pay progression can have contrasting effects depending on whether promotions are decided solely based on performance or not. In meritocratic regimes, a steeper pay progression motivates frontline workers to climb the organization's ladder and prompts an increase in their effort. In non-meritocratic regimes, in contrast, a steeper pay progression reduces worker productivity. We provide suggestive evidence that this latter effect is consistent with a negative morale effect.

Our findings have several important policy implications. In recent years, the manager-worker pay ratio has rapidly grown around the world. In the United States, it has increased more than tenfold over the past 50 years, from approximately 20 in the 1960s to over 300 in 2015 (Ashraf and Bandiera, 2018; Mishel and Wolfe, 2019). The salaries of high-level officials in public-sector agencies in developing countries have also substantially increased in recent years, partly motivated by recommendations from the World Bank and other international organizations (Shepherd, 2003; World Bank, 2014). While raising pay at the top of the organization may improve the quality of managerial staff, the results of this paper show that this can come at the expense of demotivating workers at the bottom of the organization if the promotion system is not meritocratic enough. When, however, the promotion system is meritocratic, higher pay progression instead unambiguously increases the productivity of bottom-tier workers.

There are also several additional implications that are less straightforward and require further research. First, the positive effect of promotion incentives identified in this paper may amplify in the longer-run. During the timeframe of our experiment, few promotions took place, and thus most workers reacted to what they believe the future promotion rule will look like. In the longer run, the number of workers up-for-promotion will mechanically increase, and our results indicate that this may intensify their effort response in the years leading up to promotion eligibility. Moreover, the quality of higher-level staff may change as the number of promotions increases. Shifting the promotion system from one that is mostly based on connections to one that rewards performance more prominently may improve the quality of the supervisors selected, and in turn further boost the effort of lower-tier workers.

Second, the effectiveness of performance-based promotions (or any other type of performance-based incentives) depends on the organization's ability to accurately measure worker performance. The noisier is the measure of performance, the lower is the worker incentive to exert effort. While our measure of worker performance is not entirely accurate, as it relies on the visits received by a random sample of the potential patients rather than the full population, it is likely more accurate than in the many settings in which it is measured by governments that lack resources to monitor workers closely. The fact that worker performance was measured by outside researchers may also have helped maintain fidelity to the design (Banerjee et al., 2008; De Ree et al., 2018).

Finally, many organizations face the trade-off of whether to incentivize workers through performance-based promotions or, alternatively, through performance-based incentives without a tournament structure. In our context, promotion incentives are shown to be very cost-effective: they prompt 37% of the workers to raise their effort (by 66% on average) at the cost of increasing the wage only for the promoted worker (by 50% or 11.7 dollars per month). Only a small share of the productivity gains is thus being passed on to workers in the form of higher wages. Promotion incentives may be even more cost-effective in contexts in which workers have greater opportunities to rise in the organization, or with a steeper pay progression. Even if cost-effective, we have shown that promotion incentives tend to concentrate the increase in productivity among a subset of the workers: those with a high perceived pay progression and with a high performance ranking. An organization that aims to achieve a more uniform distribution of effort across workers may thus prefer incentives that do not have a tournament structure. Further research is needed to get a better grasp of these trade-offs.

Figure 2.1: Beliefs Updating about Meritocracy

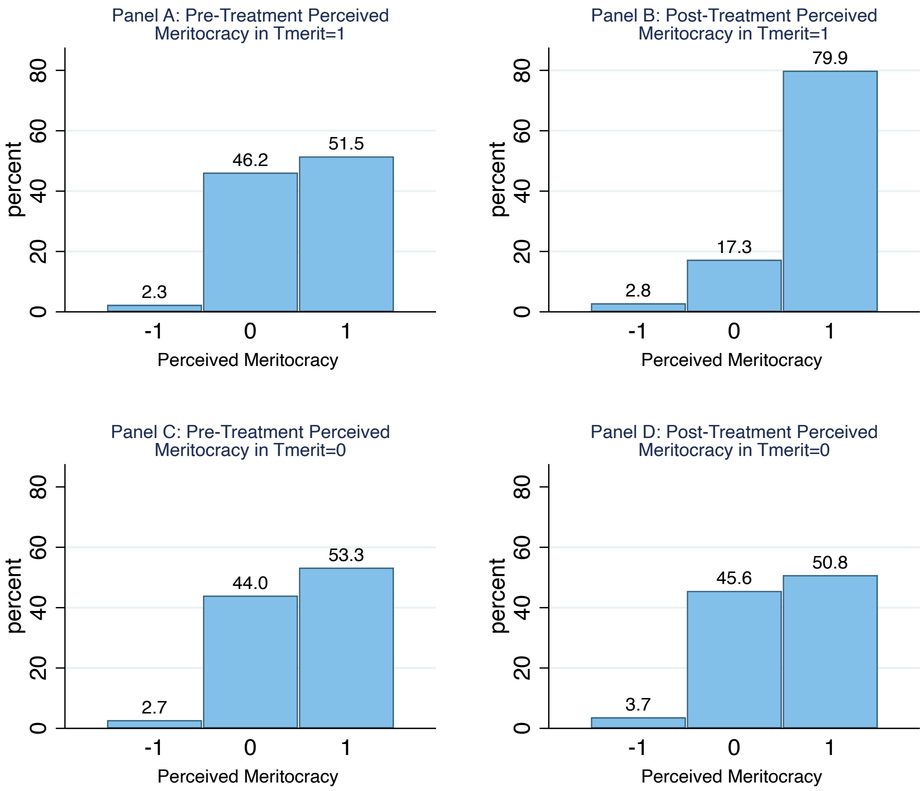

Notes: This figure plots the distribution of perceived meritocracy in the promotion system, which ranges from -1 to 1. Refer to the text for an exact definition. Panels A and B are restricted to Tmerit=1 and Panels C and D to Tmerit=0. Panels A and C (B and D) plot perceptions before (after) the information on meritocracy was provided to the CHWs.

Figure 2.2: Beliefs Updating about Pay Progression

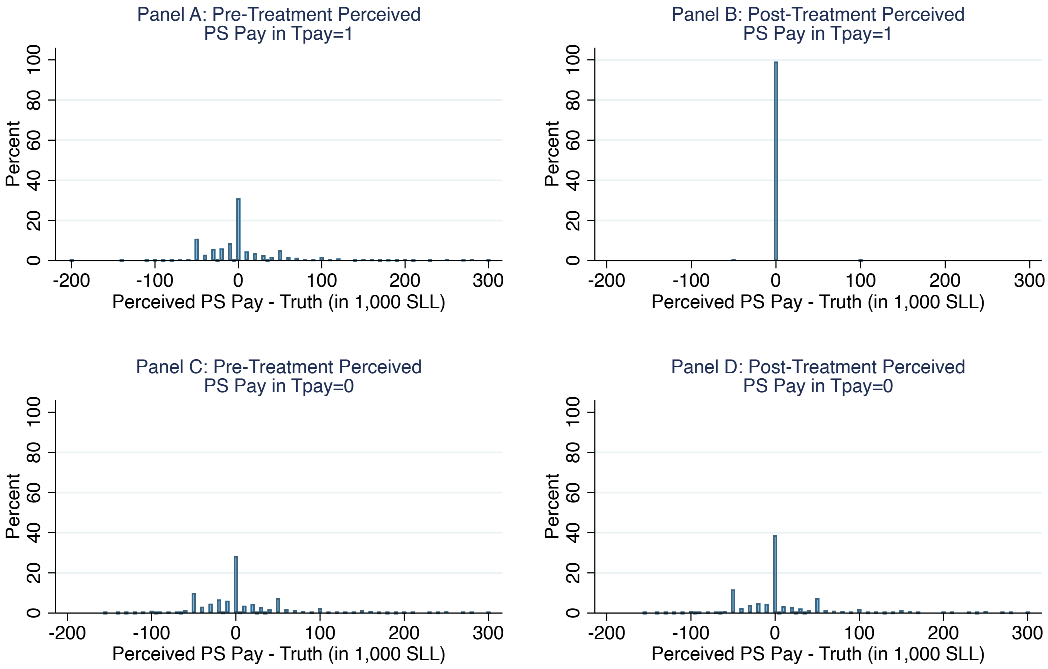

Notes: This figure plots the difference between perceived PS pay and the truth (250,000 SLL). Panels A and B are restricted to Tpay=1 and Panels C and D to Tpay=0. Panels A and C (B and D) plot perceptions before (after) the information on PS pay was provided to the CHWs.

Figure 2.3: Effect of Meritocracy on the Number of Visits, by Worker Type

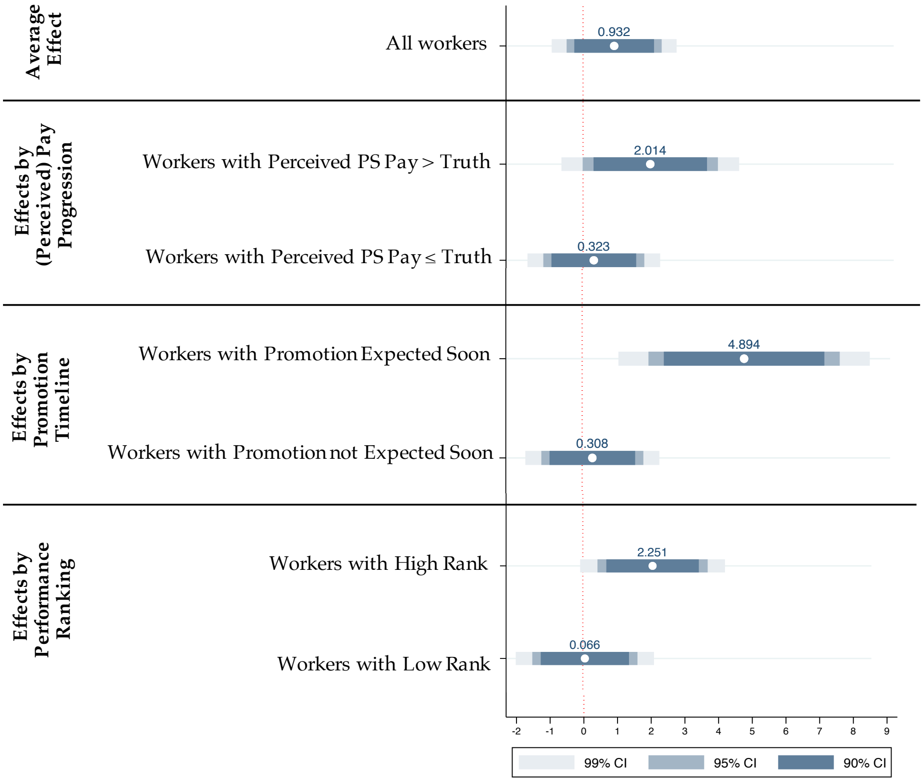

Notes: The first coefficient plots the effect of Tmerit on the number of visits for the average worker. The other coefficients plot the effect of Tmerit for different samples of workers. "Perceived PS Pay > Truth" equals 1 if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). All regression coefficients correspond to those shown in Table 4, in which we control for the stratification variables and cluster standard errors at the PHU level. The sample is restricted to CHWs in Tpay=0.

Figure 2.4: Effect of Pay Progression on the Number of Visits, by Meritocracy

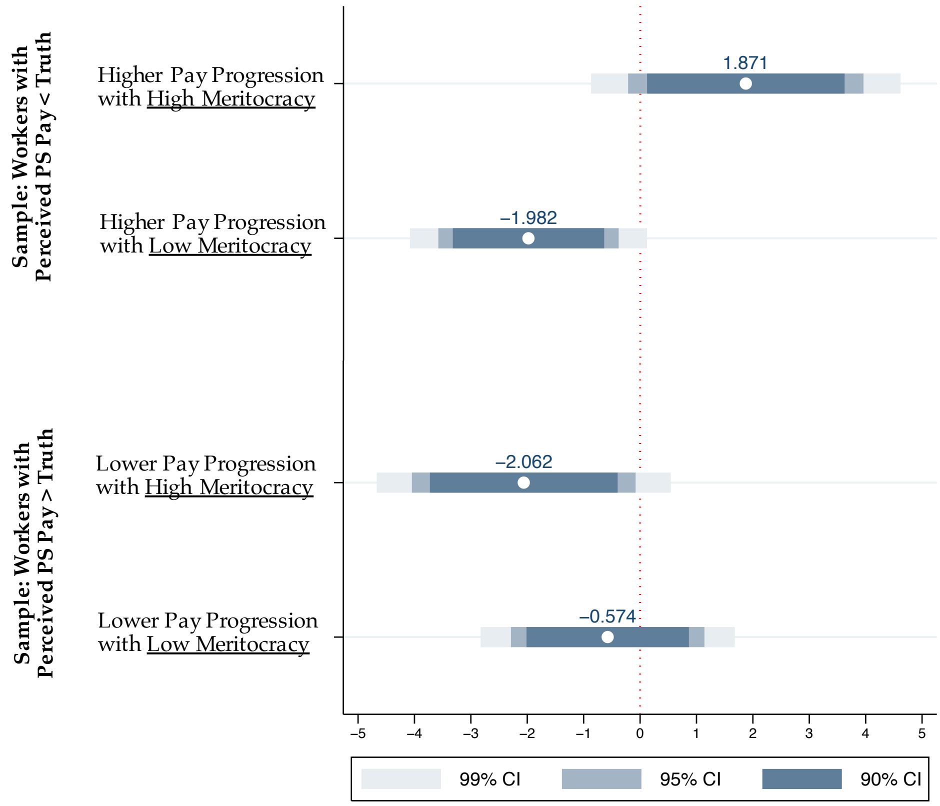

Notes: This figure plots the effects of Tpay on the number of visits for High Meritocracy (Tmerit=1) vs. Low Meritocracy (Tmerit=0) using a single regression with an interaction term. The sample is restricted to workers with baseline "Perceived PS Pay < Truth" in the top half of the figure and on the sample of workers with baseline "Perceived PS Pay > Truth" in the bottom half of the figure. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). All regression coefficients correspond to those shown in Table 6 (columns 1 and 2), in which we include stratification variables and cluster standard errors at the PHU level.

Table 2.1: Summary Statistics and Balance Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			Tmerit		Tpay		Tmerit × Tpay	
	Mean	S.D.	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.
A. CHW characteristics (N=2,009)								
Male = {0, 1}	0.726	0.446	-0.017	(0.034)	-0.023	(0.030)	-0.001	(0.048)
Age (in years)	37.03	11.22	0.111	(0.848)	-0.731	(0.780)	1.255	(1.117)
Completed primary education = {0, 1}	0.713	0.453	-0.024	(0.036)	0.018	(0.035)	0.009	(0.050)
Completed secondary education or above = {0, 1}	0.083	0.275	0.019	(0.020)	-0.018	(0.019)	-0.001	(0.027)
Wealth score (0 to 8)	2.496	1.157	0.084	(0.083)	0.008	(0.068)	0.025	(0.116)
Health knowledge score (0 to 7)	2.895	1.425	-0.065	(0.115)	-0.039	(0.110)	0.111	(0.155)
Number of years as CHW	2.212	2.828	0.346	(0.218)	0.083	(0.180)	-0.164	(0.280)
Number of households CHW is responsible for	56.90	73.98	0.944	(6.278)	-1.014	(5.520)	2.109	(8.457)
Number of hours worked as CHW per week	17.78	34.71	-0.070	(3.010)	-2.410	(2.979)	2.824	(3.832)
Number of household visits provided per week	21.47	19.93	0.350	(1.753)	0.775	(1.606)	-1.488	(2.198)
Satisfied with the PS = {0, 1}	0.762	0.426	0.073**	(0.034)	0.058	(0.036)	-0.040	(0.046)
Number of years CHW has known PS for	7.774	8.430	0.038	(0.706)	-0.283	(0.632)	0.843	(0.949)
Ever talked to the PHU in-charge = {0, 1}	0.530	0.499	-0.022	(0.048)	-0.032	(0.048)	-0.040	(0.067)
Number of years CHW has known PHU in-charge for	2.926	4.645	-0.652	(0.479)	-0.825*	(0.491)	0.613	(0.599)
PS was the best-performing CHW when promoted = {0, 1}	0.451	0.498	-0.054	(0.080)	-0.021	(0.081)	0.116	(0.113)
B. PS characteristics (N=372)								
Male = {0, 1}	0.919	0.273	0.043	(0.031)	-0.000	(0.037)	-0.105*	(0.054)
Age (in years)	37.84	8.856	0.433	(1.336)	-1.449	(1.281)	0.715	(1.785)
Completed primary education = {0, 1}	0.739	0.440	-0.001	(0.066)	0.031	(0.065)	0.015	(0.091)
Completed secondary education or above = {0, 1}	0.253	0.435	0.022	(0.065)	-0.010	(0.065)	-0.047	(0.091)
Wealth score (0 to 8)	3.013	1.227	0.128	(0.169)	-0.092	(0.175)	0.117	(0.240)
Health knowledge score (0 to 7)	3.481	1.371	0.045	(0.198)	0.100	(0.202)	-0.119	(0.282)
Number of years as PS	3.529	2.734	-0.139	(0.377)	-0.072	(0.386)	0.122	(0.521)
Number of CHWs PS is responsible for	7.984	2.861	-0.381	(0.405)	-0.441	(0.407)	0.743	(0.575)
Number of hours worked as PS per week	11.16	33.97	-0.420	(5.636)	-5.758	(4.217)	9.114	(7.459)
Number of years as CHW before promotion	1.823	1.978	-0.007	(0.345)	-0.243	(0.338)	-0.284	(0.458)
Ever talked to the PHU in-charge = {0, 1}	1.000	0.000	-	-	-	-	-	-
Number of years PS has known PHU in-charge for	4.073	6.521	1.890	(1.247)	1.038	(1.570)	-1.961	(2.000)
C. CHW pre-treatment perceptions (N=2,009)								
Perceived Meritocracy = {-1, 0, 1}	0.498	0.548	-0.032	(0.030)	-0.041	(0.034)	0.030	(0.044)
Perceived PS Pay (in 1,000 SLL)	261.7	64.23	0.352	(3.634)	-4.474	(3.731)	0.744	(5.029)

Notes: This table presents summary statistics and balance checks for baseline CHW and PS characteristics in Panel A and B, and for pre-treatment CHW perceptions about meritocracy and PS pay in Panel C. Each row states the sample mean and standard deviation of a variable, as well as the estimates from a regression, where the variable is regressed on an indicator for Tmerit, Tpay and Tmerit × Tpay. All regressions control for stratification variables and cluster standard errors at the PHU level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 2.2: Meritocracy and Beliefs Updating

	(1)	(2)	(3)	(4)	(5)
	Post-Treatment Perceptions About Promotions		Post-Treatment Perceptions About PS Pay		
Dep. Var.:	Perceived Meritocracy = {-1, 0, 1}	Number of Months until Next Promotion	PS Pay (in 1,000 SLL)	PS Number of Hours Worked	PS Work-Related Expenses (in 1,000 SLL)
Tmerit	0.296***	0.653	2.848	0.104	1.840
	(0.025)	(5.049)	(1.880)	(0.594)	(3.015)
Observations	1,982	1,387	2,009	1,940	1,932
Mean Dep. Var. if Tmerit=0	0.471	46.35	253.8	14.15	95.43

Notes: All regressions control for stratification variables. "Work-related expenses" include communication and transportation costs. The sample size varies across columns because of CHWs answering "don't know" and their answer being coded as missing. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.3: Correlations Between Worker Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Perceived Meritocracy = {-1, 0, 1}		|Perceived PS Pay - Truth| (in 10,000 SLL)		1(Perceived PS Pay > Truth)		Promotion Expected Soon		Performance Ranking [Low Ranking = High Performance]		High Rank	
	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.
Male = {0, 1}	0.003	(0.018)	0.000	(0.002)	-0.018	(0.021)	0.067	(0.049)	-0.018***	(0.003)	0.104***	(0.021)
Age (in years)	-0.715	(0.438)	0.102**	(0.050)	1.468***	(0.544)	3.783***	(1.039)	-0.066	(0.075)	0.529	(0.484)
Completed primary education = {0, 1}	0.027	(0.019)	-0.003	(0.002)	-0.024	(0.023)	0.021	(0.055)	-0.015***	(0.004)	0.069***	(0.023)
Completed secondary education or above = {0, 1}	-0.003	(0.012)	-0.000	(0.001)	-0.014	(0.013)	-0.026	(0.022)	-0.008***	(0.002)	0.046***	(0.014)
Wealth score (0 to 8)	0.091**	(0.042)	0.004	(0.004)	-0.018	(0.050)	-0.012	(0.102)	-0.029***	(0.007)	0.176***	(0.049)
Health knowledge score (0 to 7)	0.025	(0.057)	-0.004	(0.006)	0.092	(0.063)	-0.261**	(0.129)	-0.035***	(0.011)	0.227***	(0.065)
Number of years as CHW	-0.039	(0.107)	0.025*	(0.013)	0.291**	(0.143)	0.272	(0.256)	-0.061***	(0.021)	0.324***	(0.124)
Number of households CHW is responsible for	1.856	(2.769)	0.349	(0.343)	2.119	(3.788)	4.574	(9.149)	-1.512***	(0.521)	9.788***	(3.465)
Number of hours worked as CHW per week	0.573	(1.029)	0.015	(0.097)	-1.682	(1.150)	8.542	(6.703)	-0.214	(0.235)	0.764	(1.571)
Number of household visits provided per week	0.915	(0.729)	0.021	(0.076)	0.165	(1.031)	0.792	(1.543)	-0.381***	(0.131)	3.508***	(0.868)
Satisfied with the PS = {0, 1}	0.045**	(0.018)	-0.001	(0.002)	0.014	(0.021)	-0.010	(0.036)	-0.004	(0.003)	0.022	(0.019)
Number of years CHW has known PS for	-0.575*	(0.341)	0.042	(0.038)	0.534	(0.403)	2.705**	(1.046)	-0.124*	(0.067)	1.134***	(0.381)
Ever talked to the PHU in-charge = {0, 1}	-0.007	(0.020)	-0.003	(0.002)	-0.015	(0.024)	0.010	(0.052)	-0.006*	(0.004)	-0.004	(0.023)
Number of years CHW has known PHU in-charge for	-0.171	(0.182)	-0.017	(0.018)	-0.301	(0.214)	0.158	(0.632)	-0.010	(0.039)	-0.179	(0.226)
PS was the best-performing CHW when promoted = {0, 1}	-0.008	(0.017)	-0.001	(0.002)	-0.019	(0.022)	-0.109	(0.078)	0.010**	(0.004)	-0.037*	(0.019)

Notes: All variables reported in this table measure a CHW characteristic at baseline. Each row states the estimates from six regressions, where the CHW characteristic in each row is regressed on the variable in each column. "Perceived PS Pay > Truth" equals one if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. All regressions control for stratification variables and cluster standard errors at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.4: Pay Progression and Beliefs Updating

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Post-Treatment Perceptions About PS Pay			Post-Treatment Perceptions About Promotions		Post-Treatment Perceptions About PS Pay			Post-Treatment Perceptions About Promotions	
Dep. Var.:	|PS Pay - Truth| (in 1,000 SLL)	PS Number of Hours Worked	PS Work-Related Expenses (in 1,000 SLL)	Perceived Meritocracy = {-1, 0, 1}	Number of Months until Next Promotion	PS Pay (in 1,000 SLL)	PS Number of Hours Worked	PS Work-Related Expenses (in 1,000 SLL)	Perceived Meritocracy = {-1, 0, 1}	Number of Months until Next Promotion
Tpay	-34.838***	0.832	4.499	-0.035	-4.081					
	(1.480)	(0.600)	(2.999)	(0.030)	(5.039)					
Tpay × 𝟙(Perceived PS Pay < Truth)						29.043***	0.134	8.052*	0.014	-8.138
						(1.823)	(0.771)	(4.318)	(0.044)	(6.837)
Tpay × 𝟙(Perceived PS Pay > Truth)						-59.685***	0.687	-1.083	-0.078	4.160
						(3.427)	(0.789)	(4.287)	(0.048)	-7.198
Tpay × 𝟙(Perceived PS Pay = Truth)						0.848	1.864**	6.087	-0.050	-7.174
						(0.929)	(0.872)	(4.905)	(0.044)	(6.820)
Observations	2,009	1,940	1,932	1,982	1,387	2,009	1,940	1,932	1,982	1,387
Mean Dep. Var. if Tpay=0	35.32	13.79	94.30	0.643	49.46	260.7	13.79	94.30	0.643	49.46
… & 𝟙(Perc. PS Pay < Truth)	32.71	14.05	92.75	0.598	50.50	220.7	14.05	92.75	0.598	50.50
… & 𝟙(Perc. PS Pay > Truth)	63.44	13.95	95.60	0.648	46.56	309.7	13.95	95.60	0.648	46.56

Notes: All regressions control for the stratification variables. Columns (6) to (10) also control for two dummy variables: 𝟙(Perceived PS Pay < Truth) and 𝟙(Perceived PS Pay > Truth). 𝟙(Perceived PS Pay < Truth) [resp., 𝟙(Perceived PS Pay > Truth)] equals one if the PS salary pre-treatment perception of the CHW is below (resp., above) the actual salary of SLL 250,000 and 0 otherwise. "Work-related expenses" include communication and transportation costs. The sample size varies across columns because of CHWs answering "don't know" and their answer being coded as missing. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.5: Meritocracy and Worker Performance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	Number of Visits						
Tmerit	0.932						
	(0.726)						
Tmerit × 𝟙(**Perceived PS Pay > Truth**) [i]		2.014*	2.073**				
		(1.033)	(1.038)				
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth**) [ii]		0.323	0.306				
		(0.772)	(0.786)				
Tmerit × **Promotion Expected Soon** [i]				4.894***	4.818***		
				(1.475)	(1.534)		
Tmerit × **Promotion not Expected Soon** [ii]				0.308	0.367		
				(0.786)	(0.784)		
Tmerit × **High Rank** [i]						2.251**	2.185**
						(0.907)	(0.853)
Tmerit × **Low Rank** [ii]						0.066	0.191
						(0.866)	(0.860)
Observations	995	995	986	995	989	932	921
Mean Dep. Var. if Tmerit=0	7.455	7.455	7.455	7.455	7.455	7.455	7.455
p-value H_0: [i] - [ii] = 0		0.099	0.090	0.007	0.012	0.026	0.038
p-value MHT correction for [i]		0.016	0.016	0.004	0.004	0.008	0.008
p-value MHT correction for [ii]		0.888	0.896	0.896	0.853	0.912	0.896
Extra Controls: Tmerit × Correlates		No	Yes	No	Yes	No	Yes

Notes: The sample is restricted to CHWs in Tpay=0. All regressions control for stratification variables and for the uninteracted x-variable indicated in bold in the table. Columns with odd numbers additionally control for CHW characteristics that are correlated with the uninteracted x-variable (see Table A.5) and their interaction with Tmerit. At the foot of the table, we present p-values adjusted for multiple hypothesis testing across all columns computed using Romano and Wolf [2016] step-down procedure. "Perceived PS Pay > Truth" equals one if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.6: Meritocracy and Worker Performance (Other Measures)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Dep. Var.:	Visit Length (in Minutes)							Retention = {0, 1}						
Tmerit	1.797*							0.032*						
	(1.083)							(0.019)						
Tmerit × 𝟙(**Perceived PS Pay > Truth**) [i]		1.998	2.091						0.079**	0.077**				
		(1.522)	(1.538)						(0.031)	(0.032)				
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth**) [ii]		1.685	1.613						0.007	0.007				
		(1.297)	(1.331)						(0.024)	(0.024)				
Tmerit × **Promotion Expected Soon** [i]				5.676*	6.476**						0.089	0.080		
				(3.013)	(3.003)						(0.054)	(0.058)		
Tmerit × **Promotion not Expected Soon** [ii]				1.199	1.114						0.023	0.024		
				(1.149)	(1.151)						(0.020)	(0.020)		
Tmerit × **High Rank** [i]						2.230	2.315						0.054*	0.060*
						(1.552)	(1.611)						(0.032)	(0.032)
Tmerit × **Low Rank** [ii]						1.457	1.557						0.009	0.008
						(1.409)	(1.406)						(0.024)	(0.024)
Observations	995	995	986	995	989	932	921	1,004	1,004	995	1,004	998	940	929
Mean Dep. Var. if Tmerit=0	14.602	14.602	14.602	14.602	14.602	14.602	14.602	0.878	0.878	0.878	0.878	0.878	0.878	0.878
p-value H_0: [i] - [ii] = 0		0.861	0.793	0.165	0.095	0.686	0.702		0.079	0.086	0.257	0.364	0.269	0.231
p-value MHT correction for [i]		0.171	0.167	0.024	0.012	0.159	0.159		0.016	0.024	0.112	0.195	0.100	0.052
p-value MHT correction for [ii]		0.171	0.171	0.183	0.183	0.183	0.179		0.857	0.857	0.287	0.267	0.857	0.857
Extra Controls: Tmerit × Correlates		No	Yes	No	Yes	No	Yes		No	Yes	No	Yes	No	Yes

Notes: The sample is restricted to CHWs in Tpay=0. All regressions control for the stratification variables and for the uninteracted x-variable indicated in bold in the table. Columns with odd numbers additionally control for CHW characteristics that are correlated with the uninteracted x-variable (see Table A.5) and their interaction with Tmerit. At the foot of the table, we p-values adjusted for multiple hypothesis testing across all columns computed using Romano and Wolf [2016] step-down procedure. "Perceived PS Pay > Truth" equals 1 if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. "Visit Length" is the average visit length as reported by households. A visit length of zero is inputed to households that are never visited by the CHW. "Retention" equals 1 if CHW self-reported not having dropped out and visited at least one household, and 0 otherwise. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.7: Pay Progression and Worker Performance

	(1)	(2)	(3)
Dep. Var.:		Number of Visits	
Sample:	*Perceived PS Pay < Truth* *[Higher Perceived Pay Progression with Tpay=1]*	*Perceived PS Pay > Truth* *[Lower Perceived Pay Progression with Tpay=1]*	*Perceived PS Pay = Truth* *[Same Perceived Pay Progression with Tpay=1]*
Tpay × High Meritocracy (Tmerit=1) [i]	1.871*	-2.062**	-0.251
	(1.065)	(1.012)	(1.016)
Tpay × Low Meritocracy (Tmerit=0) [ii]	-1.982**	-0.574	-1.010
	(0.816)	(0.875)	(0.827)
Observations	701	668	597
Mean Dep. Var.	7.577	7.763	7.313
Mean Dep. Var. if Tpay=0	7.702	8.473	7.656
p-value H_0: [i] - [ii] = 0	0.005	0.269	0.555

Notes: All regressions control for a dummy variable for "High Meritocracy (Tmerit=1)" and for the stratifications variables. Sample described in column headings. The sample of CHWs with Perceived PS Pay < / > / = Truth consists of those with PS salary pre-treatment perception below/above/equal to the actual salary of SLL 250,000. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.8: Pay Progression and Worker Performance – Morale Concerns vs. Lobbying

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Talked to PHU In-Charge = {0,1}	Fraction of Time on Non-Patient-Related Activities	Number of Visits			
Definition of **Z**:	-		High Rank		Unsatisfied with the PS	
Tpay × High Meritocracy (Tmerit=1) [i]	-0.043	-0.000				
	(0.063)	(0.016)				
Tpay × Low Meritocracy (Tmerit=0) [ii]	-0.038	0.020				
	(0.056)	(0.018)				
Tpay × High Meritocracy (Tmerit=1) × **Z** [i]			3.434***	3.781***	4.842***	4.655***
			(1.292)	(1.433)	(1.630)	(1.670)
Tpay × High Meritocracy (Tmerit=1) × **1-Z** [ii]			-1.915	-1.509	1.108	1.212
			(1.829)	(1.938)	(1.191)	(1.287)
Tpay × Low Meritocracy (Tmerit=0) × **Z** [iii]			-2.511**	-2.112**	-3.231***	-3.289***
			(1.000)	(0.997)	(1.160)	(1.244)
Tpay × Low Meritocracy (Tmerit=0) × **1-Z** [iv]			-0.997	-1.160	-1.486*	-1.227
			(1.007)	(0.985)	(0.889)	(0.829)
Observations	738	715	660	652	701	691
Mean Dep. Var.	0.543	0.212	7.577	7.577	7.577	7.577
Mean Dep. Var. if Tpay=0	0.556	0.210	7.702	7.702	7.702	7.702
p-value H_0: [i] - [ii] = 0	0.954	0.391	0.016	0.019	0.040	0.072
p-value H_0: [iii] - [iv] = 0			0.241	0.474	0.140	0.113
p-value H_0: [i] - [iii] = 0			<0.001	0.001	<0.001	<0.001
p-value H_0: [ii] - [iv] = 0			0.660	0.872	0.082	0.113
Extra Controls	No	No	No	Yes	No	Yes

Notes: Sample restricted to workers with "Perceived PS Pay < Truth". All regressions control for stratification variables and for a dummy variable for "High Meritocracy (Tmerit=1)". Columns (3)-(6) additionally control for "Z" and columns (4) and (6) also control for all CHW characteristics that are correlated with "Z" and their interactions with Tpay, Tmerit and Tpay × Tmerit. Refer to the paper for details on the list of controls. "Talked to PHU In-Charge" is self-reported by the CHW at endline. "Non-Patient Related Activities" include administrative tasks and liaising with PHU staff. The time spent on different tasks is self-reported by the CHW at endline. "High Rank" equals 1 if the CHW is ranked first, second or third in terms of performance by the PS at baseline, and 0 otherwise. "Unsatisfied with the PS" equals 1 if the CHW was not happy with the PS at baseline and 0 otherwise. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix

Appendix 2.A Appendix Tables and Figures

Figure 2.A.1: Meritocracy, Pay Progression and Government Performance by GDP Level: Country-Level Analysis

Panel A: Correlation between Meritocracy and GDP per Capita

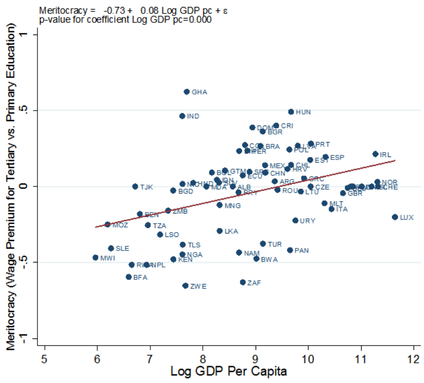

Panel B: Correlation between Pay Progression and GDP per Capita

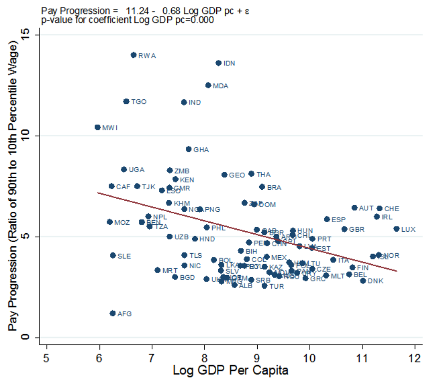

Panel C: Correlation between Government Performance and GDP per Capita

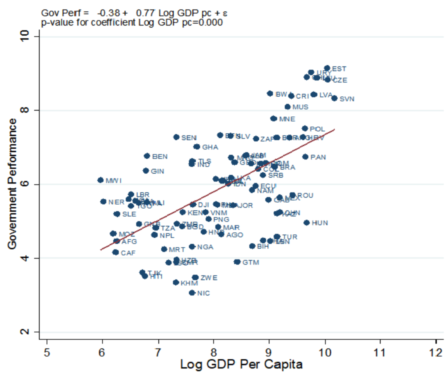

Notes: One observation per country. The red solid line represents the linear regression of meritocracy (Panel A), pay progression (Panel B) and government performance (Panel C) on log GDP per capita, with robust standard errors and no controls. For each country, we use data for the most recent year for which we have information on meritocracy, pay progression, government performance and GDP per capita (2018 or 2017 in most countries). Pay progression is measured by the World Bank's Worldwide Bureaucracy Indicators as the ratio of the 90th percentile wage to the 10th percentile wage in the public sector. Meritocracy is measured by the World Bank's Worldwide Bureaucracy Indicators as the average wage premium for workers with a tertiary education vs. a primary education in the public sector relative to the private sector. (Differences between the public and private sectors are used to hold fixed country-level characteristics such as the fraction of workers with a tertiary or primary education.) Government performance is measured by the Gothenburg's Quality of Government Indicators as an index of 4 government scores (1-10): steering capability, resource efficiency, consensus building, and international cooperation. Log GDP per capita is measured by the World Development Indicators.

Figure 2.A.2: Association between Meritocracy, Pay Progression and Government Performance: Country-Level Analysis

Notes: One observation per country-year. The red solid line represents the linear regression of government performance on pay progression (Panels A-B) or meritocracy (Panels C-D), with country and year fixed effects and with standard errors clustered at the country level. Panels A and B focus on the sample of countries with average meritocracy below and above the sample median, respectively. Panels C and D focus on the sample of countries with average pay progression below and above the sample median, respectively. "Residuals Meritocracy" ("Residuals Pay Progression") are measured as the residuals from a regression of meritocracy (pay progression) on country and year fixed effects. Pay progression is measured by the World Bank's Worldwide Bureaucracy Indicators as the ratio of the 90th percentile wage to the 10th percentile wage in the public sector. Meritocracy is measured by the World Bank's Worldwide Bureaucracy Indicators as the average wage premium for workers with a tertiary education vs. a primary education in the public sector relative to the private sector. (Differences between the public and private sectors are used to hold fixed country-level characteristics such as the fraction of workers with a tertiary or primary education.) Government performance is measured by the Gothenburg's Quality of Government Indicators as an index of 4 government scores (1-10): steering capability, resource efficiency, consensus building, and international cooperation. All variables vary across countries but also within countries over time.

Figure 2.A.3: Comparison of Supervisor's vs. Worker's Connections and Performance in the Status-Quo Promotion System

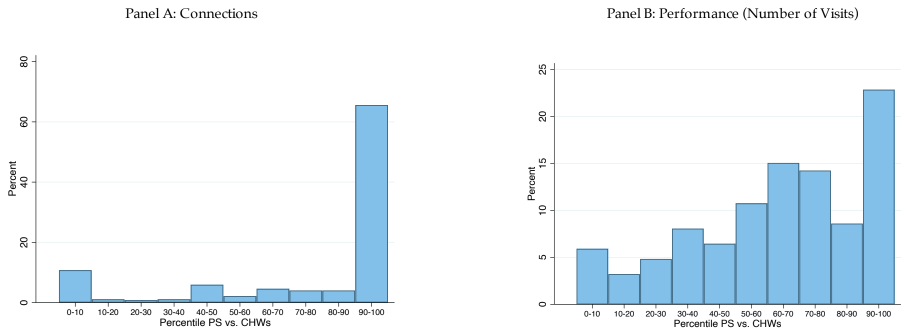

Notes: Panel A plots the distribution of the number of years the PS had known the PHU in-charge before joining the health program relative to the number of years other CHWs in the PHU (i.e., other candidates for the PS position) had known the PHU in-charge. PS connections is the x^{th} percentile if she had known the PHU in-charge for more years than x% of the CHWs in her PHU. Panel B plots the distribution of PS performance as a CHW relative to the performance of other CHWs in the PHU. PS performance is the x^{th} percentile if she performed better as a CHW than x% of the CHWs in her PHU. Because PS past performance when they were CHWs is not observed, we predict it in two steps. In the sample of all CHWs, we first regress the number of endline visits provided by a CHW within a given PHU on CHW characteristics: gender, age, primary/secondary education, tenure as a CHW. The R-squared of the first-stage is 38%. We then calculate the PS predicted number of visits by multiplying the obtained coefficients from the first step by the actual PS characteristics at the moment in which she was promoted. We do not include health knowledge and the wealth score in our two-step procedure because we do not know their values at the time of the promotion.

Figure 2.A.4: Meritocracy and Worker Performance by Prior PS Pay and Performance Ranking

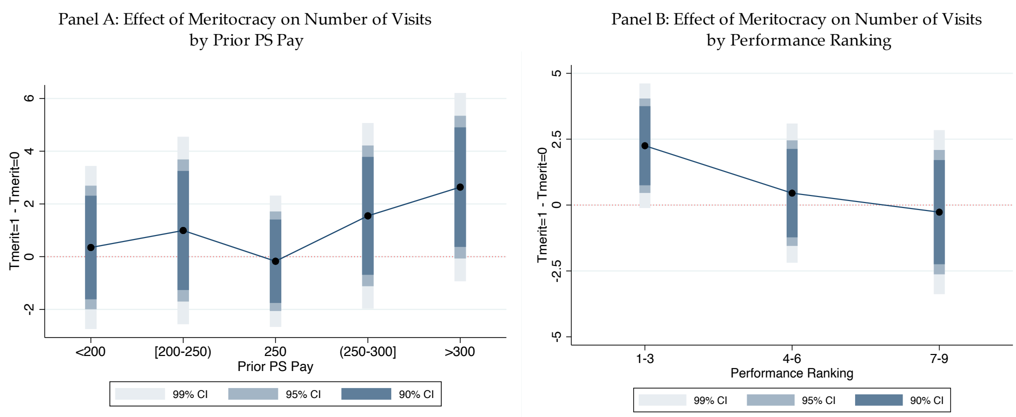

Notes: This figure plots the effect of Tmerit by perceived PS pay (Panel A) and by performance ranking as reported by the PS (Panel B). It plots the coefficients from regressing the number of visits on Tmerit, a dummy for the category reported on the x-axis and the interaction of Tmerit with each dummy, controlling for the stratification variables and with standard errors clustered at the PHU level. The sample is restricted to CHWs in Tpay=0. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households).

Table 2.A.1: Correlates of Supervisor Performance

	(1)	(2)	(3)	(4)	(5)	(6)
	Health knowledge score (0 to 7)		Predicted number of visits as a CHW		Number of years PS has known the PHU in-charge for	
	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.
Number of times the PS visited or called a CHW	0.122**	(0.056)	0.174	(0.217)	0.009	(0.018)
Number of times the PS accompanied a CHW to HH visit	0.010**	(0.005)	0.030**	(0.015)	-0.003*	(0.001)
Total number of HH visits provided by all CHWs supervised by the PS	0.600	(1.393)	9.383**	(4.130)	-0.104	(0.255)

Notes: Each row states the estimates from three regressions, where the variable in each row is regressed on the column variable. The regressions are at the PS level (sample of all 372 PSs). All regressions control for stratification variables, and for the two treatment indicators: Tmerit and Tpay. "Number of times PS visited or called a CHW" is reported by each CHW and aggregated to PS level. "Number of times PS accompanied a CHW to a HH visit" is reported by each household and aggregated to PS level. "Total number of household visits provided by all CHWs supervised by the PS" is reported by each household and aggregated to the PS level. "Predicted number of visits as a CHW" (columns 3-4) is measured in two steps. In the sample of all CHWs, we first regress the number of endline visits provided by a CHW within a given PHU on CHW characteristics: gender, age, primary/secondary education, tenure as a CHW. The R-squared of the first-stage is 38%. We then calculate the PS predicted number of visits by multiplying the obtained coefficients from the first step by the actual PS characteristics at the moment in which she was promoted. Robust standard errors presented in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 2.A.2: Summary Statistics and Balance Checks at Village and Household Level

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			Tmerit		Tpay		Tmerit × Tpay	
	Mean	S.D.	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.
A. Village characteristics (N=2,009)								
Accessible road to government hospital= {0, 1}	0.788	0.409	0.009	(0.039)	0.014	(0.044)	-0.022	(0.058)
Primary school in the village = {0, 1}	0.477	0.500	-0.003	(0.040)	0.024	(0.039)	0.027	(0.056)
Number of water sources in the village	2.742	26.24	2.456	(2.193)	0.980	(0.870)	-2.718	(2.497)
Mobile network available = {0, 1}	0.868	0.338	-0.009	(0.028)	-0.031	(0.030)	0.012	(0.042)
B. Household respondents, aggregated to village level (N=2,009)								
Age (in years)	29.15	4.990	0.115	(0.396)	0.288	(0.364)	-0.829	(0.527)
Completed primary education = {0, 1}	0.284	0.292	0.041*	(0.021)	0.024	(0.023)	-0.028	(0.032)
Number of children under 5	0.731	0.280	0.015	(0.022)	-0.020	(0.023)	-0.017	(0.033)
Wealth score (0 to 8)	-0.220	2.175	0.280	(0.194)	0.225	(0.189)	-0.268	(0.259)
Main occupation is farming = {0, 1}	0.605	0.369	-0.017	(0.027)	-0.045	(0.028)	0.011	(0.041)
Knew the CHW at baseline = {0, 1}	0.971	0.121	-0.005	(0.007)	-0.003	(0.007)	0.001	(0.012)
CHW is localed <30 min = {0, 1}	0.870	0.273	-0.002	(0.021)	0.002	(0.022)	0.000	(0.028)
Government hospital is localed <30 min = {0, 1}	0.389	0.409	0.046	(0.037)	0.031	(0.031)	-0.060	(0.047)

Notes: This table presents summary statistics and balance checks for baseline village characteristics in Panel A and for household characteristics (aggregated to the village level) in Panel B. Each row states the sample mean and standard deviation of a variable, as well as the estimates from a regression where the variable is regressed on an indicator for Tmerit, Tpay and Tmerit × Tpay. All regressions control for stratification variables and cluster standard errors at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.A.3: Worker Characteristics that Predict Promotions

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Promoted to PS = {0, 1}				
Connected to the PHU in-charge = {0, 1}	0.663***			0.555***	0.607***
	(0.047)			(0.045)	(0.050)
High performance (predicted number of visits > median) = {0, 1}		0.460***		0.272***	
		(0.029)		(0.029)	
Male = {0, 1}			0.116***		0.103***
			(0.036)		(0.033)
Age (in years)			0.002		0.002
			(0.002)		(0.001)
Completed primary education = {0, 1}			0.193***		0.078**
			(0.033)		(0.038)
Completed secondary education or above = {0, 1}			0.476***		0.264***
			(0.085)		(0.068)
High tenure (tenure > median) = {0, 1}			0.221***		0.033
			(0.032)		(0.027)
Observations	746	746	743	746	743
Mean Dep. Var.	0.217	0.217	0.217	0.217	0.217
R-squared	0.553	0.381	0.304	0.620	0.586

Notes: The dependent variable "Promoted to PS" equals one for the PSs in our sample and zero for the CHWs in our sample who were present in the PHU at the time of the promotion. The dependent variable is regressed on characteristics of the potential candidates for the PS position at the time of the promotion, who are assumed to be the current PS and the CHWs who were present in the PHU at the time of the promotion. "Connected to the PHU in-charge" equals one if the number of years the candidate has known the PHU in-charge before joining the program is in the top quartile. "High performance" equals one if the number of visits performed by the candidate is above the median. Because PS past performance when they were CHWs is not observed, we predict it in two steps. In the sample of all CHWs, we first regress the number of endline visits provided by a CHW within a given PHU on CHW characteristics: gender, age, primary/secondary education, tenure as a CHW. We then calculate the PS predicted number of visits by multiplying the obtained coefficients from the first step by the actual PS characteristics at the moment in which she was promoted. "High tenure" equals one if the number of years as a CHW is above the median. The sample is restricted to Tmerit=0. All regressions control for PHU fixed effects and cluster standard errors at the PHU level *** p<0.01, ** p<0.05, * p<0.1

Table 2.A.4: Meritocracy and Beliefs Updating – Heterogeneous Effects

	(1)	(2)	(3)	(4)
Dep. Var.:	Post-Treatment Perceived Meritocracy = {-1, 0, 1}			
Definition of (pre-treatment) **Z** variable:	Perceived Meritocracy = {-1, 0, 1}	𝟙(Perceived PS Pay > Truth)	Promotion Expected Soon	High Rank
Z	0.739***	-0.033	0.010	0.018
	(0.028)	(0.040)	(0.050)	(0.036)
Tmerit	0.574***	0.283***	0.297***	0.322***
	(0.030)	(0.029)	(0.026)	(0.032)
Tmerit × **Z**	-0.543***	0.039	-0.016	-0.053
	(0.039)	(0.050)	(0.074)	(0.047)
Observations	1,982	1,982	1,982	1,842
Mean Dep. Var.	0.626	0.626	0.626	0.626

Notes: The variable Z is defined in the column headings and measured before the treatment. "𝟙(Perceived PS Pay > Truth)" equals 1 if the PS salary pre-treatment perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals 1 if the supervisor of the CHW is within 5 years of retirement age at baseline. "High Rank" equals 1 if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. All regressions control for stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.A.5: Pay Progression and Beliefs Updating – Heterogeneous Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	|Post-Treatment Perceived PS Pay - Truth| (in 1,000 SLL)						
Definition of (pre-treatment) **Z** variable:	| Perceived PS Pay - Truth|	Tmerit	High Rank	Satisfied with the PS	Tmerit	High Rank	Satisfied with the PS
Z	0.774***	-1.180	-2.582	-1.180	-0.317	0.931	-2.282
	(0.048)	(3.679)	(3.430)	(3.679)	(2.471)	(2.996)	(3.372)
Tpay	-3.449**	-36.063***	-35.549***	-36.063***			
	(1.524)	(3.114)	(1.961)	(3.114)			
Tpay × **Z**	-0.769***	1.625	2.524	1.625			
	(0.049)	(3.696)	(3.421)	(3.696)			
Tpay × 𝟙(Perceived PS Pay < Truth)					-32.552***	-31.862***	-34.274***
					(2.060)	(1.766)	(2.859)
Tpay × 𝟙(Perceived PS Pay > Truth)					-62.084***	-62.991***	-65.066***
					(3.678)	(3.689)	(4.519)
Tpay × 𝟙(Perceived PS Pay = Truth)					-2.274	-1.474	-3.624
					(1.611)	(1.729)	(2.697)
Tpay × 𝟙(Perceived PS Pay < Truth) × **Z**					-0.287	-0.268	2.136
					(2.637)	(3.013)	(3.379)
Tpay × 𝟙(Perceived PS Pay > Truth) × **Z**					-1.283	-1.039	3.079
					(2.673)	(3.169)	(3.428)
Tpay× 𝟙(Perceived PS Pay = Truth) × **Z**					-0.225	-1.844	1.700
					(2.632)	(3.032)	(3.491)
Observations	2,009	2,009	1,867	2,009	2,009	1,867	2,009
Mean Dep. Var.	17.90	17.90	17.90	17.90	17.90	17.90	17.90

Notes: All regressions control for stratification variables. Columns (5) and (7) also control for two dummy variables: 𝟙(Perceived PS Pay < Truth) and 𝟙(Perceived PS Pay > Truth). 𝟙(Perceived PS Pay < Truth) [resp., 𝟙(Perceived PS Pay > Truth)] equals one if the PS salary pre-treatment perception of the CHW is below (resp., above) the actual salary of SLL 250,000 and 0 otherwise. "High Rank" equals 1 if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. "Satisfied with the PS" equals 1 if the CHW was "very happy" with the PS at baseline and 0 otherwise. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.A.6: Meritocracy and Worker Performance, by Type of Visit

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
Dep. Var.:	Number of Routine Visits				Number of Cases Treated				Number of Cases Referred				Number of Ante-natal Visits				Number of Post-natal Visits			
Tmerit	1.325				1.019*				0.213				-0.011				-0.014			
	(0.909)				(0.574)				(0.160)				(0.177)				(0.024)			
Tmerit × 𝟙**(Perceived PS Pay > Truth)** [i]		3.198*				2.192**				0.247				0.334				-0.025		
		(1.650)				(0.957)				(0.262)				(0.244)				(0.057)		
Tmerit × 𝟙**(Perceived PS Pay ≤ Truth)** [ii]		0.265				0.353				0.195				-0.204				-0.008		
		(0.805)				(0.539)				(0.187)				(0.239)				(0.024)		
Tmerit × **Promotion Expected Soon** [i]			3.091				4.127**				0.966**				0.780				0.114*	
			(1.999)				(1.596)				(0.464)				(0.487)				(0.060)	
Tmerit × **Promotion not Expected Soon** [ii]			1.066				0.536				0.091				-0.137				-0.034	
			(0.988)				(0.599)				(0.159)				(0.180)				(0.025)	
Tmerit × **High Rank** [i]				1.573				1.800*				0.522**				-0.169				-0.047
				(1.768)				(1.067)				(0.248)				(0.458)				(0.039)
Tmerit × **Low Rank** [ii]				1.174				0.530				0.028				0.031				0.007
				(1.006)				(0.626)				(0.208)				(0.082)				(0.037)
Observations	995	995	995	932	995	995	995	932	995	995	995	932	995	995	995	932	995	995	995	932
Mean Dep. Var. if Tmerit=0	4.038	4.038	4.038	4.038	2.846	2.846	2.846	2.846	0.805	0.805	0.805	0.805	0.312	0.312	0.312	0.312	0.064	0.064	0.064	0.064
p-value H_0: [i] - [ii] = 0		0.064	0.357	0.834		0.042	0.035	0.291		0.866	0.076	0.111		0.118	0.075	0.663		0.788	0.019	0.337

Notes: The sample is restricted to CHWs in Tpay=0. All regressions control for stratification variables and for the uninteracted x-variable indicated in bold in the table. "Perceived PS Pay > Truth" equals 1 if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. The dependent variable is reported by the households. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.7: Meritocracy and Worker Performance – IV Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	Number of Visits						
Perceived Meritocracy	3.235						
	(2.746)						
Perceived Meritocracy × 𝟙(**Perceived PS Pay > Truth**) [i]		6.767*	7.240*				
		(3.923)	(4.141)				
Perceived Meritocracy × 𝟙(**Perceived PS Pay ≤ Truth**) [ii]		1.051	0.894				
		(2.983)	(2.957)				
Perceived Meritocracy × **Promotion Expected Soon** [i]				18.548**	19.201**		
				(8.376)	(9.523)		
Perceived Meritocracy × **Promotion not Expected Soon** [ii]				0.936	1.102		
				(2.916)	(2.871)		
Perceived Meritocracy × **High Rank** [i]						10.942*	11.898*
						(5.622)	(6.123)
Perceived Meritocracy × **Low Rank** [ii]						0.082	0.116
						(2.847)	(2.864)
Observations	981	981	972	981	975	919	908
Mean Dep. Var.	7.965	7.965	7.965	7.965	7.965	7.965	7.965
F-stat 1st Stage (Cragg Donald Test)	64.94	29.554	31.026	30.328	26.175	16.689	12.498
p-value H_0: [i] - [ii] = 0		0.151	0.123	0.045	0.066	0.045	0.053
Extra Controls: Tmerit × Correlates		No	Yes	No	Yes	No	Yes

Notes: The sample is restricted to CHWs in Tpay=0. In odd columns, we present IV regression with Tmerit as an IV. In even columns, we present IV regressions with two IVs: Tmerit × 𝟙(Perceived PS Pay > Truth) and Tmerit × 𝟙(Perceived PS Pay ≤ Truth). All regressions control for the dummy variable "𝟙(Perceived PS Pay > Truth)" and the stratification variables. Columns with odd numbers additionally control for CHW characteristics that are correlated with the uninteracted x-variable (see Table A.5) and their interaction with Tmerit. Perceived PS Pay > Truth" equals 1 if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.8: Meritocracy and Worker Performance – Additional Heterogeneous Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Dep. Var.:						Number of Visits						
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion Expected Soon (≤ 2 years for PS to retire)** [i]	4.666***	4.264**										
	(1.781)	(1.878)										
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion not Expected Soon (> 2 years for PS to retire)** [ii]	0.645	0.713										
	(0.771)	(0.771)										
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion Expected Soon (≤ 10 years for PS to retire)** [i]			2.600**	2.750**								
			(1.301)	(1.306)								
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion not Expected Soon (> 10 years for PS to retire)** [ii]			0.215	0.212								
			(0.856)	(0.861)								
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion Expected Soon (≤ 15 years for PS to retire)** [i]					0.734	0.804						
					(1.034)	(1.015)						
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion not Expected Soon (> 15 years for PS to retire)** [ii]					0.936	0.924						
					(1.013)	(1.015)						
Tmerit × **Promotion Expected Soon (self-reported)** [i]							1.378	1.321				
							(0.859)	(0.852)				
Tmerit × **Promotion not Expected Soon (self-reported)** [ii]							0.132	0.295				
							(0.914)	(0.905)				
Tmerit × 𝟙(**Perceived PS Pay > Truth) & Promotion Expected Soon (≤ 5 years for PS to retire)** [i]									7.467***	7.396***		
									(2.020)	(2.118)		
Tmerit × 𝟙(**Perceived PS Pay > Truth) & Promotion not Expected Soon (> 5 years for PS to retire)** [ii]									0.863	0.901		
									(1.108)	(1.115)		
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion Expected Soon (≤ 5 years for PS to retire)** [iii]									2.697*	2.504		
									(1.617)	(1.721)		
Tmerit × 𝟙(**Perceived PS Pay ≤ Truth) & Promotion not Expected Soon (> 5 years for PS to retire)** [iv]									0.017	0.080		
									(0.855)	(0.862)		
Tmerit × **High Rank (as reported by other CHWs)** [i]											1.483*	1.617*
											(0.814)	(0.850)
Tmerit × **Low Rank (as reported by other CHWs)** [ii]											0.511	0.661
											(1.044)	(1.028)
Observations	995	989	995	989	995	989	995	987	995	985	899	883
Mean Dep. Var. if Tmerit=0	7.455	7.455	7.455	7.455	7.455	7.455	7.455	7.455	7.455	7.455	7.455	7.455
p-value H_0: [i] - [ii] = 0	0.041	0.086	0.132	0.114	0.889	0.934	0.229	0.337	0.004	0.007	0.389	0.431
p-value H_0: [i] - [iii] = 0									0.035	0.035		
p-value H_0: [i] - [iv] = 0									0.001	0.002		
Extra Controls: Tmerit × Correlates	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

Notes: The sample is restricted to CHWs in Tpay=0. All regressions control for stratification variables and for the uninteracted x-variable(s) indicated in bold in the table. Columns with even numbers additionally control for CHW characteristics that are correlated with the uninteracted x-variable(s) and their interaction with Tmerit. Refer to the paper for details on the list of controls. "Promotion Expected Soon (≤ X years for PS to retire)" equals one if the supervisor of the CHW is within X years of retirement age at baseline and 0 otherwise, where X can be 2, 10 or 15 depending on the variables. "Promotion Expected Soon (self-reported)" equals one if the perceived duration until the next promotion as reported by CHWs at baseline is below the median. CHWs who answered that they don't know when the next promotion will take place are assumed to have a perception above the median. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by other CHWs at baseline and 0 otherwise. At baseline, each CHW was asked to assess the rank of other CHWs in the PHU. We define a CHW to be "High Rank" if she is ranked in the top three by pooling together answers from all other CHWs in the PHU. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.9: Meritocracy and Household Targeting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)
Dep. Var.:	% Visits to Households Living Within 30 Minutes Walk of the CHW							Median Distance Between the Visited Households and the CHW							% Visits to Friends/Family of the CHW						
Tmerit	0.011							0.251							0.030						
	(0.023)							(0.515)							(0.027)						
Tmerit × 𝟙**(Perceived PS Pay > Truth)** [i]		0.047	0.042						-0.037	-0.205						0.026	0.025				
		(0.030)	(0.029)						(0.613)	(0.660)						(0.039)	(0.039)				
Tmerit × 𝟙**(Perceived PS Pay ≤ Truth)** [ii]		-0.010	-0.011						0.430	0.453						0.031	0.026				
		(0.031)	(0.030)						(0.746)	(0.783)						(0.033)	(0.032)				
Tmerit × **Promotion Expected Soon** [i]				0.068	0.045						0.030	-0.122						0.059	0.017		
				(0.066)	(0.066)						(1.329)	(1.422)						(0.052)	(0.056)		
Tmerit × **Promotion not Expected Soon** [ii]				0.002	0.004						0.284	0.267						0.025	0.031		
				(0.024)	(0.024)						(0.574)	(0.590)						(0.030)	(0.029)		
Tmerit × **High Rank** [i]						-0.014	-0.008						1.690	1.734						0.005	-0.001
						(0.038)	(0.038)						(1.258)	(1.348)						(0.037)	(0.038)
Tmerit × **Low Rank** [ii]						0.019	0.009						-0.679	-0.891						0.046	0.037
						(0.029)	(0.028)						(0.523)	(0.583)						(0.033)	(0.033)
Observations	940	940	933	940	935	880	871	741	741	733	741	737	692	682	970	970	961	970	964	909	898
Mean Dep. Var. if Tmerit=0	0.872	0.872	0.872	0.872	0.872	0.872	0.872	1.767	1.767	1.767	1.767	1.767	1.767	1.767	0.451	0.451	0.451	0.451	0.451	0.451	0.451
p-value H_0: [i] - [ii] = 0		0.181	0.210	0.349	0.561	0.477	0.733		0.634	0.552	0.865	0.809	0.100	0.119		0.916	0.984	0.577	0.822	0.365	0.428
Extra Controls: Tmerit × Correlates		No	Yes	No	Yes	No	Yes		No	Yes	No	Yes	No	Yes		No	Yes	No	Yes	No	Yes

Notes: The sample is restricted to CHWs in Tpay=0. All regressions control for the stratification variables and for the uninteracted x-variable indicated in bold in the table. Columns with odd numbers additionally control for CHW characteristics that are correlated with the uninteracted x-variable (see Table A.5) and their interaction with Tmerit. "Perceived PS Pay > Truth" equals 1 if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. The sample size varies from one column to another because of missing values in the dependent variable. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.10: Meritocracy and Supervisor Performance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dep. Var.:	PS Visited CHW or Accompanied Her to HH Visit						
Tmerit	0.003						
	(0.034)						
Tmerit × $\mathbb{1}$(**Perceived PS Pay > Truth**) [i]		0.011	0.005				
		(0.044)	(0.044)				
Tmerit × $\mathbb{1}$(**Perceived PS Pay ≤ Truth**) [ii]		-0.002	0.005				
		(0.040)	(0.040)				
Tmerit × **Promotion Expected Soon** [i]				0.040	0.054		
				(0.067)	(0.066)		
Tmerit × **Promotion not Expected Soon** [ii]				-0.004	-0.006		
				(0.038)	(0.038)		
Tmerit × **High Rank** [i]						0.013	0.019
						(0.044)	(0.045)
Tmerit × **Low Rank** [ii]						-0.013	-0.005
						(0.044)	(0.042)
Observations	1,004	1,004	995	1,004	998	940	929
Mean Dep. Var. if Tmerit=0	0.829	0.829	0.829	0.829	0.829	0.829	0.829
p-value H_0: [i] - [ii] = 0		0.791	0.992	0.572	0.437	0.631	0.676
Extra Controls: Tmerit × Correlates		No	Yes	No	Yes	No	Yes

Notes: The sample is restricted to CHWs in Tpay=0. All regressions control for stratification variables and for the uninteracted x-variable indicated in bold in the table. Columns with odd numbers additionally control for CHW characteristics that are correlated with the uninteracted x-variable (see Table A.5) and their interaction with Tmerit. "Perceived PS Pay > Truth" equals 1 if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. "PS Visited CHW or Accompanied Her to HH Visit" equals 1 if the PS visited or called the CHW at least once or if at least one household reports having received a visit in which the CHW was accompanied by the PS, and 0 otherwise. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.11: Pay Progression and Worker Performance – Fully Interacted Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Dep. Var.:	Number of Visits			Visit Length (in Minutes)			Retention = {0, 1}			PS Visited CHW or Accompanied Her to HH Visit		
Effects for the average CHW:												
Tpay × High Meritocracy (Tmerit=1) [i]	-0.179			-1.049			0.014			0.030		
	(0.719)			(1.074)			(0.018)			(0.031)		
Tpay × Low Meritocracy (Tmerit=0) [ii]	-1.227**			-1.033			-0.001			0.029		
	(0.596)			(1.156)			(0.024)			(0.033)		
Tmerit [iii]	0.978			1.856*			0.031			0.000		
	(0.745)			(1.116)			(0.019)			(0.034)		
Effects for CHWs with Perceived PS Pay < Truth												
Tpay × High Meritocracy (Tmerit=1) [iv]		1.809*	1.729		-0.823	-0.947		0.083***	0.090***		-0.003	-0.016
		(1.075)	(1.150)		(1.700)	(1.629)		(0.030)	(0.030)		(0.038)	(0.037)
Tpay × Low Meritocracy (Tmerit=0) [v]		-1.952**	-1.973**		-0.807	-1.572		-0.061	-0.075**		0.015	0.024
		(0.822)	(0.834)		(1.589)	(1.661)		(0.040)	(0.037)		(0.045)	(0.044)
Tmerit [vi]		0.802	0.979		3.822**	3.396*		-0.004	-0.015		0.020	0.038
		(0.992)	(1.008)		(1.695)	(1.746)		(0.035)	(0.036)		(0.045)	(0.044)
Effects for CHWs with Perceived PS Pay > Truth												
Tpay × High Meritocracy (Tmerit=1) [vii]		-2.045**	-2.298**		-2.379*	-3.316**		-0.044	-0.041		0.018	0.014
		(1.023)	(1.005)		(1.431)	(1.470)		(0.030)	(0.032)		(0.041)	(0.041)
Tpay × Low Meritocracy (Tmerit=0) [viii]		-0.684	-0.756		-1.451	-1.278		0.030	0.033		0.020	-0.000
		(0.860)	(0.842)		(1.673)	(1.679)		(0.040)	(0.038)		(0.051)	(0.048)
Tmerit [ix]		2.006*	1.960*		1.781	2.536		0.075**	0.080**		0.011	-0.011
		(1.035)	(1.041)		(1.524)	(1.562)		(0.032)	(0.033)		(0.045)	(0.041)
Effects for CHWs with Perceived PS Pay = Truth												
Tpay × High Meritocracy (Tmerit=1)		-0.300	-0.322		0.291	-0.026		-0.006	-0.012		0.086*	0.074
		(1.018)	(1.053)		(1.670)	(1.744)		(0.032)	(0.032)		(0.052)	(0.053)
Tpay × Low Meritocracy (Tmerit=0)]		-0.968	-0.281		-0.817	-0.662		0.037	0.045		0.052	0.049
		(0.833)	(0.811)		(1.859)	(1.815)		(0.035)	(0.039)		(0.044)	(0.044)
Tmerit [ix]		-0.060	0.136		-0.467	-0.373		0.020	0.028		-0.039	-0.039
		(0.976)	(0.964)		(1.863)	(1.927)		(0.030)	(0.032)		(0.054)	(0.054)
Observations	1,966	1,966	1,938	1,966	1,966	1,938	2,009	2,009	1,981	2,009	2,009	1,981
Mean Dep. Var.	7.560	7.560	7.560	14.944	14.944	14.944	0.893	0.893	0.893	0.843	0.843	0.843
Mean Dep. Var. if Tpay=0	7.965	7.965	7.965	15.586	15.586	15.586	0.891	0.891	0.891	0.829	0.829	0.829
p-value H_0: [i] - [ii] = 0	0.260	0.006	0.010	0.992	0.994	0.788	0.614	0.004	0.001	0.978	0.758	0.490
p-value H_0: [i] - [iii] = 0	0.373	0.573	0.687	0.127	0.132	0.152	0.593	0.142	0.081	0.604	0.751	0.435
p-value H_0: [ii] - [iii] = 0	0.001	0.002	0.001	0.011	0.012	0.006	0.159	0.167	0.126	0.368	0.914	0.712
p-value H_0: [iv] - [v] = 0		0.309	0.241		0.672	0.359		0.147	0.136		0.968	0.817
p-value H_0: [iv] - [vi] = 0		0.030	0.021		0.092	0.023		0.020	0.023		0.923	0.718
p-value H_0: [v] - [vi] = 0		0.007	0.004		0.027	0.013		0.179	0.143		0.836	0.819
Extra Controls	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes

Notes: All regressions control for the stratification variables. The last two columns of each outcome variables control for 𝟙(Perceived PS Pay < Truth) and 𝟙(Perceived PS Pay > Truth). Additionally, the last column of each outcome variable controls for all CHW characteristics in Table 1 and their interactions with Tpay, Tmerit and Tpay × Tmerit."Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). "Visit Length" is the average visit length as reported by the households. A visit length of zero is inputed to households that are never visited by the CHW. "Retention" equals 1 if CHW self-reported not having dropped out and visited at least one household, and 0 otherwise. "PS Visited CHW or Accompanied Her to HH Visit" equals one if the PS visited or called the CHW at least once or if at least one household reports having received a visit in which the CHW was accompanied by the PS, and 0 otherwise. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.12: Summary Statistics and Balance Checks by PS Pay Priors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mean	S.D.	Tmerit		Tpay		Tmerit × Tpay	
			Coeff	S.E.	Coeff	S.E.	Coeff	S.E.
A. CHW characteristics for CHWs with Perceived PS Pay < Truth (N=738)								
Male = {0, 1}	0.710	0.454	-0.085	(0.052)	-0.082	(0.052)	0.105	(0.075)
Age (in years)	37.10	11.25	-0.855	(1.246)	-0.418	(1.232)	1.489	(1.694)
Completed primary education = {0, 1}	0.706	0.456	-0.077	(0.050)	-0.055	(0.051)	0.077	(0.074)
Completed secondary education or above = {0, 1}	0.081	0.273	0.047*	(0.027)	0.042	(0.028)	-0.049	(0.043)
Wealth score (0 to 8)	2.533	1.224	0.061	(0.123)	0.132	(0.119)	0.069	(0.181)
Health knowledge score (0 to 7)	2.757	1.467	-0.097	(0.173)	-0.082	(0.160)	-0.165	(0.235)
Number of years as CHW	2.001	2.622	0.338	(0.291)	0.319	(0.291)	-0.426	(0.393)
Number of households CHW is responsible for	60.14	69.68	-9.165	(8.201)	3.420	(9.200)	7.861	(11.979)
Number of hours worked as CHW per week	21.83	23.32	3.149	(2.255)	3.927	(3.043)	-3.832	(3.928)
Number of household visits provided per week	19.93	16.20	-1.565	(1.688)	2.292	(1.683)	-0.332	(2.415)
Satisfied with the PS = {0, 1}	0.760	0.427	0.090*	(0.050)	0.064	(0.054)	-0.046	(0.068)
Number of years CHW has known PS for	7.569	8.383	0.621	(1.077)	1.058	(0.974)	0.963	(1.470)
Ever talked to the PHU in-charge = {0, 1}	0.543	0.498	-0.072	(0.061)	-0.038	(0.056)	-0.005	(0.085)
Number of years CHW has known PHU in-charge for	3.126	4.888	-0.916	(0.667)	-1.204*	(0.635)	1.113	(0.851)
PS was the best-performing CHW when promoted = {0, 1}	0.434	0.496	-0.056	(0.083)	-0.092	(0.084)	0.136	(0.122)
B. CHW characteristics for CHWs with Perceived PS Pay > Truth (N=673)								
Male = {0, 1}	0.736	0.441	0.008	(0.048)	-0.023	(0.049)	-0.002	(0.072)
Age (in years)	38.28	11.50	1.052	(1.339)	-0.627	(1.267)	2.042	(1.845)
Completed primary education = {0, 1}	0.689	0.463	0.034	(0.057)	0.054	(0.057)	-0.062	(0.081)
Completed secondary education or above = {0, 1}	0.068	0.253	-0.014	(0.027)	-0.051**	(0.025)	0.048	(0.038)
Wealth score (0 to 8)	2.366	1.064	0.191	(0.121)	-0.010	(0.116)	-0.177	(0.171)
Health knowledge score (0 to 7)	3.007	1.414	0.013	(0.167)	0.050	(0.168)	0.092	(0.231)
Number of years as CHW	2.534	3.041	0.346	(0.374)	0.099	(0.304)	-0.124	(0.512)
Number of households CHW is responsible for	56.39	80.98	6.446	(9.043)	-2.135	(8.216)	0.505	(12.702)
Number of hours worked as CHW per week	23.00	21.58	1.238	(2.496)	2.045	(2.691)	-3.107	(3.611)
Number of household visits provided per week	21.81	21.90	2.667	(2.836)	1.807	(3.120)	-5.510	(3.717)
Satisfied with the PS = {0, 1}	0.761	0.427	0.058	(0.052)	0.022	(0.054)	-0.006	(0.075)
Number of years CHW has known PS for	8.215	8.654	-0.751	(1.048)	-1.454	(0.903)	1.103	(1.411)
Ever talked to the PHU in-charge = {0, 1}	0.508	0.500	-0.024	(0.066)	-0.074	(0.067)	0.031	(0.094)
Number of years CHW has known PHU in-charge for	2.657	4.469	-0.274	(0.615)	-0.330	(0.619)	0.022	(0.802)
PS was the best-performing CHW when promoted = {0, 1}	0.444	0.497	-0.080	(0.090)	-0.006	(0.094)	0.158	(0.128)
C. CHW characteristics for CHWs with Perceived PS Pay = Truth (N=598)								
Male = {0, 1}	0.734	0.442	0.024	(0.053)	0.041	(0.048)	-0.122*	(0.070)
Age (in years)	35.54	10.69	0.018	(1.210)	-1.393	(1.118)	0.699	(1.675)
Completed primary education = {0, 1}	0.747	0.435	-0.032	(0.055)	0.066	(0.057)	0.002	(0.077)
Completed secondary education or above = {0, 1}	0.100	0.301	0.027	(0.044)	-0.053	(0.040)	-0.004	(0.054)
Wealth score (0 to 8)	2.599	1.162	-0.019	(0.141)	-0.104	(0.114)	0.182	(0.186)
Health knowledge score (0 to 7)	2.940	1.373	-0.080	(0.161)	-0.027	(0.154)	0.406*	(0.217)
Number of years as CHW	2.110	2.798	0.271	(0.294)	-0.244	(0.276)	0.218	(0.405)
Number of households CHW is responsible for	53.48	70.71	3.405	(10.761)	-8.216	(6.223)	1.765	(12.681)
Number of hours worked as CHW per week	20.92	19.90	-0.550	(2.466)	-2.585	(2.338)	2.485	(3.447)
Number of household visits provided per week	22.97	21.61	-0.517	(3.418)	-1.949	(2.482)	1.070	(4.138)
Satisfied with the PS = {0, 1}	0.766	0.424	0.063	(0.055)	0.082	(0.056)	-0.064	(0.073)
Number of years CHW has known PS for	7.532	8.225	0.050	(0.943)	-0.581	(0.989)	0.567	(1.328)
Ever talked to the PHU in-charge = {0, 1}	0.538	0.499	0.031	(0.066)	0.001	(0.067)	-0.143	(0.091)
Number of years CHW has known PHU in-charge for	2.981	4.524	-0.994	(0.628)	-1.066*	(0.632)	0.810	(0.775)
PS was the best-performing CHW when promoted = {0, 1}	0.500	0.500	-0.003	(0.100)	0.065	(0.099)	0.024	(0.138)

Notes: This table presents summary statistics of CHW characteristics in the three sub-samples: CHWs who overestimated PS pay at baseline (Panel A), CHWs who underestimated PS pay at baseline (Panel B), and CHWs who estimated PS pay correctly (Panel C). Each row states the sample mean and standard deviation of a variable, as well as the estimates from a regression, where the variable is regressed on an indicator for Tmerit, Tpay and Tmerit × Tpay. All regressions control for stratification variables and cluster standard errors at the PHU level. All variables reported in this table are measured at baseline. *** p<0.01, ** p<0.05, * p<0.1

Table 2.A.13: Pay Progression, Worker Performance (Other Measures) and Supervisor Performance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dep. Var.:	Visit Length (in Minutes)	Retention = {0, 1}	PS Visited CHW or Accompanied Her to HH Visit	Visit Length (in Minutes)	Retention = {0, 1}	PS Visited CHW or Accompanied Her to HH Visit	Visit Length (in Minutes)	Retention = {0, 1}	PS Visited CHW or Accompanied Her to HH Visit
Sample:	*Perceived PS Pay < Truth*			*Perceived PS Pay > Truth*			*Perceived PS Pay = Truth*		
Tpay × High Meritocracy (Tmerit=1) [i]	-0.849	0.087***	-0.003	-2.357	-0.048	0.016	0.274	-0.010	0.081
	(1.698)	(0.030)	(0.039)	(1.429)	(0.030)	(0.040)	(1.661)	(0.032)	(0.051)
Tpay × Low Meritocracy (Tmerit=0) [ii]	-1.136	-0.063	0.013	-1.333	0.029	0.028	-0.878	0.040	0.055
	(1.590)	(0.040)	(0.045)	(1.653)	(0.039)	(0.050)	(1.850)	(0.034)	(0.045)
Observations	701	738	738	668	673	673	597	598	598
Mean Dep. Var.	14.910	0.888	0.852	14.950	0.900	0.840	14.977	0.893	0.836
Mean Dep. Var. if Tpay=0	15.620	0.885	0.848	15.929	0.903	0.831	15.117	0.885	0.801
p-value H_0: [i] - [ii] = 0	0.902	0.002	0.779	0.640	0.125	0.857	0.642	0.288	0.706

Notes: All regressions control for a dummy variable for "High Meritocracy (Tmerit=1)" and for the stratifications variables. Sample described in column headings. The sample of CHWs with Perceived PS Pay < / > / = Truth consist of those with PS salary pre-treatment perception below/above/equal to the actual salary of SLL 250,000. "Visit Length" is the average visit length as reported by the households. A visit length of zero is inputed to households that are never visited by the CHW. "Retention" equals 1 if CHW self-reported not having dropped out and visited at least one household, and 0 otherwise. "PS Visited CHW or Accompanied Her to HH Visit" equals one if the PS visited or called the CHW at least once or if at least one household reports having received a visit in which the CHW was accompanied by the PS, and 0 otherwise. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.14: Pay Progression and Household Targeting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dep. Var.:	% Visits to Households Living Within 30 Minutes Walk of the CHW	Median Distance Between the Visited Households and the CHW	% Visits to Friends/ Family of the CHW	% Visits to Households Living Within 30 Minutes Walk of the CHW	Median Distance Between the Visited Households and the CHW	% Visits to Friends/ Family of the CHW	% Visits to Households Living Within 30 Minutes Walk of the CHW	Median Distance Between the Visited Households and the CHW	% Visits to Friends/ Family of the CHW
Sample:	*Perceived PS Pay < Truth*			*Perceived PS Pay > Truth*			*Perceived PS Pay = Truth*		
Tpay × High Meritocracy (Tmerit=1) [i]	-0.028	-0.189	-0.008	-0.064**	2.427	-0.061	0.051	0.056	-0.056
	(0.038)	(1.426)	(0.042)	(0.032)	(2.496)	(0.044)	(0.036)	(0.593)	(0.041)
Tpay × Low Meritocracy (Tmerit=0) [ii]	0.003	1.233	-0.059	0.011	0.368	-0.015	0.012	0.194	0.012
	(0.043)	(0.965)	(0.041)	(0.034)	(0.506)	(0.043)	(0.043)	(0.702)	(0.045)
Observations	668	525	676	640	483	654	560	433	573
Mean Dep. Var.	0.847	2.609	0.430	0.902	2.256	0.484	0.870	1.754	0.433
Mean Dep. Var. if Tpay=0	0.861	2.360	0.449	0.911	1.715	0.500	0.853	1.627	0.445
p-value H_0: [i] - [ii] = 0	0.597	0.454	0.382	0.113	0.414	0.454	0.481	0.882	0.272

Notes: All regressions control for a dummy variable for "High Meritocracy (Tmerit=1)" and for the stratifications variables. Sample described in column headings. The sample of CHWs with Perceived PS Pay < / > / = Truth consist of those with PS salary pre-treatment perception below/above/equal to the actual salary of SLL 250,000. "Visit Length" is the average visit length as reported by the households. A visit length of zero is inputed to households that are never visited by the CHW. "Retention" equals 1 if CHW self-reported not having dropped out and visited at least one household, and 0 otherwise. "PS Visited CHW or Accompanied Her to HH Visit" equals one if the PS visited or called the CHW at least once or if at least one household reports having received a visit in which the CHW was accompanied by the PS, and 0 otherwise. The sample size varies from one column to another because of missing values in the dependent variable. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.15: Pay Progression and Worker Performance – IV Results

	(1)	(2)	(3)
Dep. Var.:		Number of Visits	
Sample:	*All*	*Perceived PS Pay < Truth*	*Perceived PS Pay > Truth*
Perceived PS Pay Updating × High Meritocracy (Tmerit=1) [i]	0.028***		
	(0.009)		
Perceived PS Pay Updating × Low Meritocracy (Tmerit=0) [ii]	-0.002		
	(0.008)		
Perceived PS Pay × High Meritocracy (Tmerit=1) [i]		0.074*	0.033**
		(0.043)	(0.017)
Perceived PS Pay × Low Meritocracy (Tmerit=0) [ii]		-0.061**	0.010
		(0.025)	(0.015)
Observations	1,966	701	668
Mean Dep. Var.	7.560	7.560	7.560
F-stat 1st Stage (Cragg Donald Test)	181.058	89.894	96.240
p-value H_0: [i] - [ii] = 0	0.007	0.007	0.300

Notes: Sample described in column headings. In column (1), we present an IV regression with four IVs: Tpay × 𝟙(Perceived PS Pay < Truth) × High Meritocracy, Tpay × 𝟙(Perceived PS Pay < Truth) × Low Meritocracy, Tpay × 𝟙(Perceived PS Pay > Truth) × High Meritocracy, Tpay × 𝟙(Perceived PS Pay > Truth) × Low Meritocracy. In columns (2) and (3), we use 2 IVs: Tpay × High Meritocracy, Tpay × Low Meritocracy. "Perceived PS Pay Updating" is the difference between post- and pre-treatment perceived PS pay, and is expressed in thousand of SLL. "Perceived PS Pay" is the post-treatment perceived PS pay, expressed in thousand of SLL. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). All regressions control for a dummy variable for "High Meritocracy (Tmerit=1)" and for the stratification variables. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix 2.B Temporary Incentives Introduced by External Organization

The CHWs and PSs in this study were part of a separate evaluation that is the focus of (Deserranno et al., 2022a) and that involves a temporary performance-based incentive scheme paid by an external organization. The randomization was done at the PHU level. In the Shared Incentives Treatment, CHWs received an incentive of 1,000 SLL for each service performed and the PS received an incentive of 1,000 SLL for each service performed by a CHW under her supervision. In the Worker Incentives Treatment, CHWs received an incentive of 2,000 SLL for each service performed while the PS received no incentives. In the Supervisor Incentives Treatment, the PS received an incentive of 2,000 SLL for each service performed by a CHW under her supervision while the CHWs received no incentives. In the control group, neither the CHWs nor the PS received an incentive. In each treatment, the number of services a CHW provided was measured with an SMS reporting system, which required the CHW to report the date and type of service and the contact information of the patient by sending an SMS to a toll-free number. This reporting system played no role in the main experiment of this paper.

As mentioned in the body of the paper, the randomization of the meritocracy and pay progression treatments was stratified by the above-mentioned incentives treatments. Still, one may be concerned that the main effects shown in the paper are driven by specific interactions between the treatments in the two projects. We address this concern directly in Table 2.B.1, where we first show that the impact of the meritocratic promotion and pay progression treatments on perceptions of meritocracy and pay progression are orthogonal to the presence of these incentives. This is not surprising as these incentives are short-run and are provided by an external organization with no connection with the government, and thus should not affect the perceptions about the promotion criteria or perceptions about the pay PSs receive from the government. Accordingly, Table 2.B.2 shows that the effects of the meritocracy and pay progression treatments on the number of visits do not interact with the incentives treatments (column 2). The effects of the meritocracy treatment by perceived PS pay, promotion expected soon or high rank – which we presented in Section 2.5 – also appear orthogonal to the incentives treatments (columns 3-5).

One may be worried that there may just be too little power to test for these interactions. In that case, one should cautiously interpret the effects of our meritocracy and pay progression treatments as composite treatment effects that include a weighted-average of the interactions with the incentives treatments (Muralidharan et al., 2020). These composite weighted-average treatment effects remain qualitatively informative and policy-relevant.

Table 2.B.1: Incentives and Perceptions

	(1)	(2)	(3)	(4)
Dep.Var.:	Post-Treatment Perceived Meritocracy = {-1, 0, 1}		|Post-Treatment Perceived PS Pay - Truth| (in 1,000 SLL)	
Supv Incentives	0.018	0.043	-1.409	-2.399
	(0.043)	(0.042)	(3.125)	(2.724)
Worker Incentives	0.023	0.042	0.389	3.740
	(0.041)	(0.040)	(3.254)	(2.902)
No Supv/Worker incentives	-0.005	0.027	2.517	4.140
	(0.041)	(0.038)	(3.273)	(2.872)
Tmerit		0.317***		
		(0.044)		
Tmerit × Supv Incentives		-0.007		
		(0.062)		
Tmerit × Worker Incentives		-0.013		
		(0.059)		
Tmerit × No Supv/Worker incentives		-0.035		
		(0.062)		
Tpay				-32.367***
				(2.578)
Tpay × Supv Incentives				2.760
				(3.460)
Tpay × Worker Incentives				-2.899
				(3.500)
Tpay × No Supv/Worker incentives				-2.333
				(3.642)
Observations	1,933	1,933	2,009	2,009
Mean Dep. Var. in Omitted Group	0.615	0.448	18.157	34.405

Notes: All regressions control for district fixed effects and the baseline value of the outcome variable. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1

Table 2.B.2: Main Results, Interactions with Incentives

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:			Number of Visits		
Definition of Z:	-	-	𝟙(Perceived PS Pay > Truth)	Promotion Expected Soon	High Rank
Tmerit	0.998	0.849			
	(0.789)	(1.670)			
Tpay	-1.297**	-1.761			
	(0.637)	(1.474)			
Tpay × Tmerit	1.089	1.312			
	(0.981)	(2.067)			
Tmerit × Supv Incentives		2.772			
		(2.167)			
Tpay × Supv Incentives		0.378			
		(1.786)			
Tpay × Tmerit × Supv Incentives		-3.235			
		(2.675)			
Tmerit × Worker Incentives		-1.920			
		(2.296)			
Tpay × Worker Incentives		1.123			
		(1.967)			
Tpay × Tmerit × Worker Incentives		2.824			
		(2.869)			
Tmerit × No Supv/Worker incentives		-0.755			
		(1.833)			
Tpay × No Supv/Worker incentives		0.546			
		(1.682)			
Tpay × No Supv/Worker incentives		-0.527			
		(2.373)			
Tmerit × Z			1.984	3.822	1.882
			(1.921)	(2.411)	(1.893)
Tmerit × 1-Z			0.190	0.450	0.244
			(1.644)	(1.741)	(1.820)
Tmerit × Z × Supv incentives			3.518	3.498	2.128
			(2.600)	(2.610)	(2.729)
Tmerit × 1-Z × Supv incentives			2.327	2.346	3.203
			(2.257)	(2.289)	(2.290)
Tmerit × Z × Worker incentives			-2.004	3.979	-0.779
			(2.756)	(3.861)	(2.525)
Tmerit × 1-Z × Worker incentives			-2.039	-2.499	-3.470
			(2.251)	(2.299)	(2.385)
Tmerit × Z × No Supv/Worker incentives			-2.494	-3.491	-0.041
			(2.106)	(2.491)	(2.184)
Tmerit × 1-Z × No Supv/Worker incentives			0.138	-0.452	-0.366
			(1.857)	(1.906)	(2.065)
Observations	1,966	1,966	995	995	932
Mean Dep. Var.	7.560	7.560	7.560	7.560	7.560

Notes: Columns (3) to (5) are restricted to Tmerit=0 and control for the uninteracted Z variable, defined in the column heading. "Number of Visits" is the average number of household visits provided by the CHW (as reported by the households). "Perceived PS Pay > Truth" equals one if the PS salary perception of the CHW is above the actual salary of SLL 250,000 and 0 otherwise. "Promotions Expected Soon" equals one if the supervisor of the CHW is within 5 years of retirement age at baseline and 0 otherwise. "High Rank" equals one if the CHW is ranked first, second or third in terms of performance by the PS at baseline and 0 otherwise. Standard errors are clustered at the PHU level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix 2.C Ethics Appendix

Following (Asiedu et al., 2021), we detail key aspects of research ethics.

Pre-Analysis Plan The study was pre-registered on the AEA RCT Registry with the number 0003993. We follow the pre-analysis closely. The outcomes variables we use in the paper, and the heterogeneous treatment effects with respect to perceived pay progression and worker ability were pre-registered.

In the pre-analysis plan, we specified that we would use the number of SMS reports, described in Appendix 2.B, as a secondary measure of worker performance. We ended up not using this variable because the average worker is found to vastly underreport the visits provided: the average reporting rate is 17.7% and is comparable across treatments. This measure is hence uninformative about worker performance. We decided to focus on households' responses in the household survey to measure worker performance.

We also specified that we would study heterogeneous treatment effects by social connections to the PHU in-charge. We did not present these results in the main text due to space constraints and because of the lack of a clear theoretical prediction on this heterogeneity. For transparency, we describe the results here (results available in a table format upon request). We find that higher meritocracy has no significant effect on the productivity of highly ranked workers who are well-connected to the PHU in-charge, and no significant effect on the productivity of low-ranked workers who are well-connected to the PHU in-charge. (A worker is defined as well-connected if she has known the PHU for more years than half of the other CHWs). Making promotions more performance-based significantly increases the number of visits of high-ranked unconnected workers by 4.682 (statistically significant at the 1% level).

IRB and Research Ethics The project received IRB from the University of Pompeu Fabra (CIREP Approval 107) and from the Sierra Leone Ethics and Scientific Review Committee (no IRB number assigned by this local institution).

We obtained informed consent from all participants prior to the study. The consent form described the participants' risks and rights, confidentiality, and contact information. Research staff and enumerator teams were not subject to additional risks in the data collection process. None of the researchers have financial or reputation conflicts of interest with regard to the research results. No contractual restrictions were imposed on the researchers limiting their ability to report the study findings.

On policy equipoise and scarcity, there was uncertainty regarding the net benefits from our treatments for any worker. The interventions under study did not pose any potential harm to participants and non-participants. The intervention rollout took place according to the evaluation protocol.

On potential harms to participants or nonparticipants, our data collection and research procedures adhered to protocols around privacy, confidentiality, risk-management, and informed consent. Participants were not considered particularly vulnerable (beyond some households residing in poverty). Besides individual consent from study participants, consultations were conducted with local representatives at the district levels. All the enumerators involved in data collection were recruited from the study districts to ensure they were aware about implicit social norms in these communities.

The presentation of the findings from the project to district and national level authorities in Sierra Leone was delayed due to COVID-19 but is planned for 2022. No activity for sharing results to participants in each study village is planned due to resource constraints. We do not foresee risks of the misuse of research findings. Policy briefs have been created based on this project and have been distributed to policymakers through IGC and CEGA.

Appendix 2.D Model Appendix

2.D.1 Main Results

This section formally develops the theoretical framework presented in Section 2.4. Throughout we assume that player 2 is willing to participate in the promotion contest but exerts less effort than player 1 such that the costs of effort are equal to the pay progression.

Assumption 1 *The cost functions satisfy* $r_1 > r_2$*, where* $r_1 = bc_1^{-1}(\bar{w} - \underline{w}) = b\frac{\bar{w}-\underline{w}}{c_1}$ *and* $r_2 = \frac{\bar{w}-\underline{w}}{c_2 g_2(b,\bar{w}-\underline{w})}$.[46]

Following (Siegel, 2010), the b-biased promotion tournament with effort costs (c_1, c_2) has a unique equilibrium in mixed strategies. We derive the following lemma, which we prove in Appendix 2.D.2:

Lemma 1 *The average effort, as a function of* $\bar{w} - \underline{w}$*,* c_1*,* c_2 *and* b*, is given by* $\bar{e}_1(\bar{w}-\underline{w}, b, c_{1,}, c_2) = \frac{\bar{w}-\underline{w}}{2bc_2 g_2(b,\bar{w}-\underline{w})}$ *and* $\bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2) = \frac{c_1(\bar{w}-\underline{w})}{2bc_2^2 g_2(b,\bar{w}-\underline{w})^2}$*, for players* 1 *and* 2*, respectively.*

Results without Morale Concerns

This section derives the propositions that underlie the predictions without morale concerns (i.e., $g_i(b, \bar{w} - \underline{w}) = 1$ for $i = 1, 2$) presented in Section 2.4.2. The corresponding proofs are presented in Appendix 2.D.2.

Proposition 1 *Fix* c_1*, and suppose that* $\tilde{c}_2 > \tilde{\tilde{c}}_2$*. Then* $\bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2) > \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)$*, for* $i = 1, 2$.

[46] This assumption does not imply $c_1 < c_2$ or $c_1 > c_2$. In what follows, we do not restrict to either case.

Proposition 2 *Let* $b' > b$*, then* $\bar{e}_i(\bar{w}-\underline{w}, b, c_1, c_2) > \bar{e}_i(\bar{w}-\underline{w}, b', c_1, c_2)$*, for* $i = 1, 2$.

This result implies Prediction 1.

Proposition 3 *Let* $\bar{\bar{w}} - \underline{\underline{w}} > \bar{w} - \underline{w}$*. Then* $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) > \bar{e}_i(\bar{w} - \underline{w}, b, c_1, c_2)$*, for* $i = 1, 2$.

This result implies Prediction 2.

We are also interested in the effect of pay progression on workers' effort at different levels of meritocracy, and the effect of meritocracy at different levels of pay progression. We have that:

Proposition 4 *Let* $\bar{\bar{w}} - \underline{\underline{w}} \geq \bar{w} - \underline{w}$*,* $b' \geq b$*. Then* $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, c_2) \geq \bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, c_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, c_2)$*, for* $i = 1, 2$.

This result implies Prediction 3.

Proposition 5 *Let* $b' > b$*. For* $\tilde{c}_2 > \tilde{\tilde{c}}_2$*, we have that* $\bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2) > \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)$*, for* $i = 1, 2$.

This entails that the result of Proposition 2 is amplified when player 2 is of higher ability.

Proposition 6 *Let* $\bar{\bar{w}} - \underline{\underline{w}} > \bar{w} - \underline{w}$*. For* $\tilde{c}_2 > \tilde{\tilde{c}}_2$ *we have that* $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2) > \bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)$*, for* $i = 1, 2$.

This entails that the result of Proposition 3 is amplified when player 2 is of higher ability.

Proposition 7 *Let* $\bar{\bar{w}} - \underline{\underline{w}} > \bar{w} - \underline{w}$*,* $b' > b$*. For* $\tilde{c}_2 > \tilde{\tilde{c}}_2$ *and* $i = 1, 2$

$$\left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right) - \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)\right) >$$
$$\left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right) - \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\right).$$

This tells us that the result of Proposition 4 is amplified when player 2 is of higher ability. Taken together, Propositions 5, 6, and 7 imply Prediction 4.

Results with Morale Concerns

This section derives the propositions that underlie the predictions of the model with morale concerns presented in Section 2.4.3.

We make three assumptions about the morale cost-shift function g_i. (Section 2.4.3 provides the intuition for each of them):

Assumption 2 *1. $g_1(b, \bar{w} - \underline{w}) = 1$ for all $(b, \bar{w} - \underline{w}) \in \mathbb{R}^2_+$.*

2. $g_2 : \mathbb{R}^2_+ \to \mathbb{R}_{++}$ is strictly increasing in all of its arguments, log supermodular, and $g_2(1, \bar{w} - \underline{w}) = 1 \ \forall \bar{w} - \underline{w}$.

3. Domination of cost-shift for higher pay progression: For $\bar{w} - \underline{w} < \bar{\bar{w}} - \underline{\underline{w}}$, we have that $\lim_{b \to \infty} \frac{g_2(b, \bar{w} - \underline{w})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})} = 0$.

Given these assumptions, we obtain the following propositions, which we prove in Appendix 2.D.2:

Proposition 8 *Let $b' > b$. Then $\bar{e}_i(\bar{w} - \underline{w}, b', c_1, c_2) \leq \bar{e}_i(\bar{w} - \underline{w}, b, c_1, c_2)$, for $i = 1, 2$.*

This result implies Prediction 5.

Proposition 9 *Let $\bar{\bar{w}} - \underline{\underline{w}} \geq \bar{w} - \underline{w}$. Then there exists $\bar{b}, \bar{\bar{b}}$ where $\bar{\bar{b}} \geq \bar{b}$, such that:*
1. If $b \leq \bar{b}$, $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) \geq \bar{e}_i(\bar{w} - \underline{w}, b, c_1, c_2)$, for $i = 1, 2$, and
2. If $b \geq \bar{\bar{b}}$, $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) \leq \bar{e}_i(\bar{w} - \underline{w}, b, c_1, c_2)$, for $i = 1, 2$.

That is, if $b \geq \bar{\bar{b}}$, the equilibrium level of effort decreases as pay progression increases. Instead, if $b \leq \bar{b}$, the equilibrium level of effort increases. From this, we derive Prediction 6.

Proposition 10 *Let* $\bar{\bar{w}} - \underline{\underline{w}} \geq \bar{w} - \underline{w}$, $b' \geq b$ *and* $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, c_2) \geq 0$, *for* $i = 1, 2$. *Then* $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, c_2) \geq \bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, c_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, c_2)$, *for* $i = 1, 2$.

This implies Prediction 7.

Proposition 11 *Let* $b' > b$. *For* $\tilde{c}_2 > \tilde{\tilde{c}}_2$ *we have* $|\bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)| > |\bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)|$, *for* $i = 1, 2$.

This implies that the result of Proposition 8 is amplified when player 2 is of higher ability.

Proposition 12 *Let* $\bar{\bar{w}} - \underline{\underline{w}} > \bar{w} - \underline{w}$. *For* $\tilde{c}_2 > \tilde{\tilde{c}}_2$ *we have* $|\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)| > |\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)|$, *for* $i = 1, 2$.

This implies that the result of Proposition 9 is amplified when player 2 is of higher ability.

Proposition 13 *Let* $\bar{\bar{w}} - \underline{\underline{w}} > \bar{w} - \underline{w}$, $b' > b$, $\tilde{c}_2 > \tilde{\tilde{c}}_2$ *and* $\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2) \geq 0$, *for* $i = 1, 2$. *Then, for* $i = 1, 2$,

$$\begin{aligned}&\left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right) - \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)\right) > \\ &\quad \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right) - \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\right).\end{aligned}$$

We can then say that the result of Proposition 10 is amplified when player 2 is of higher ability. Taken together, Propositions 11, 12, and 13 imply Prediction 8.

2.D.2 Proofs

Lemma 1

Proof 1 *Define the score of player 1 as* $s_1 = be_1$ *and the score of player 2 as* $s_2 = e_2$*. The score indicates how effort maps into the probability of winning. We can rewrite the tournament success function under a biased rule as:*

$$P_i^b(s_1, s_2) = \begin{cases} 0 & \text{if } s_i < s_{-i} \\ p & \text{if } s_i = s_{-i} \\ 1 & \text{if } s_i > s_{-i} \end{cases}$$

where $p \in [0, 1]$.

Mapping to (Siegel, 2010), we have that $v_1(s_1) = \bar{w} - \underline{w} - c_1\left(\frac{s_1}{b}\right)$ *and* $v_2(s_2) = \bar{w} - \underline{w} - g_2(b, \bar{w} - \underline{w})c_2(s_2)$*. Given* $c_i > 0$ *and Assumption 1, (Siegel, 2010)'s assumptions are satisfied. From Theorem 3 in (Siegel, 2010), we conclude that the cdfs of the score are:*

$$E^s{}_1(s) = \begin{cases} \frac{g_2(b, \bar{w} - \underline{w})c_2(s)}{\bar{w} - \underline{w}} & \text{if } y \in [0, r_2) \\ 1 & \text{if } y \geq r_2 \end{cases} \quad \text{and,} \quad E_2^s(s) = \begin{cases} \frac{\bar{w} - \underline{w} - c_1(r_2) + c_1(s)}{\bar{w} - \underline{w}} & \text{if } s \in [0, r_2) \\ 1 & \text{if } s \geq r_2 \end{cases}.$$

We now express the cdfs of the score as cdfs of each player's effort. For any given score where $s_1 = s_2$*, we have that* $\frac{e_1}{b} = e_2$ *and* $be_2 = e_1$*. Therefore,*

$$E_1(e) = \begin{cases} \frac{g_2(b, \bar{w} - \underline{w})c_2(be)}{\bar{w} - \underline{w}} & \text{if } e \in \left[0, \frac{r_2}{b}\right) \\ 1 & \text{if } e \geq \frac{r_2}{b} \end{cases} \quad \text{and,} \quad E_2(e) = \begin{cases} \frac{\bar{w} - \underline{w} - c_1(r_2) + c_1\left(\frac{e}{b}\right)}{\bar{w} - \underline{w}} & \text{if } e \in [0, r_2) \\ 1 & \text{if } e \geq r_2 \end{cases}.$$

We can now compute the average effort as a function of $\bar{w} - \underline{w}$ *and* b*:*

$$\begin{aligned}\bar{e}_1(\bar{w} - \underline{w}, b, c_1, c_2) = \mathbb{E}_{E_1}(e) &= \int_0^{\frac{1}{b}\frac{\bar{w}-\underline{w}}{c_2 g_2(b,\bar{w}-\underline{w})}} \frac{g_2(b, \bar{w} - \underline{w})bc_2}{\bar{w} - \underline{w}} e \quad de \\ &= \frac{g_2(b, \bar{w} - \underline{w})bc_2}{2(\bar{w} - \underline{w})} \left(\frac{\bar{w} - \underline{w}}{bc_2 g_2(b, \bar{w} - \underline{w})}\right)^2 \\ &= \frac{\bar{w} - \underline{w}}{2bc_2 g_2(b, \bar{w} - \underline{w})}\end{aligned}$$

$$\begin{aligned}\bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2) = \mathbb{E}_{E_2}(e) &= \int_0^{\frac{\bar{w}-\underline{w}}{c g_2(b,\bar{w}-\underline{w})}} \frac{c_1}{\bar{w} - \underline{w}} \frac{e}{b} \quad de \\ &= \frac{c_1}{2b(\bar{w} - \underline{w})} \left(\frac{\bar{w} - \underline{w}}{c_2 g_2(b, \bar{w} - \underline{w})}\right)^2 \\ &= \frac{c_1(\bar{w} - \underline{w})}{2bc_2^2 g_2(b, \bar{w} - \underline{w})^2}\end{aligned}$$

Proofs: Model without Morale Concerns

Proposition 1

Proof 2 *We have that* $g_2(b, \bar{w} - \underline{w}) = 1$ *for all* $(b, \bar{w} - \underline{w})$*. Therefore,* $\bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2) = \frac{c_1(\bar{w}-\underline{w})}{2b\tilde{c}_2^2}$ *and* $\bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2) = \frac{(\bar{w}-\underline{w})}{2b\tilde{c}_2}$*, while* $\bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2) = \frac{c_1(\bar{w}-\underline{w})}{2b\tilde{\tilde{c}}_2^2}$ *and* $\bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2) = \frac{(\bar{w}-\underline{w})}{2b\tilde{\tilde{c}}_2}$*. As* $\tilde{c}_2 \geq \tilde{\tilde{c}}_2$*, it immediately follows that* $\bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2) \leq \bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)$ *and* $\bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2) \leq \bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)$*. Without morale concerns, the effort of both players thus decreases as the costs for player* 2 *increases.*

Proposition 2

Proof 3 *We have that* $\bar{e}_1(\bar{w} - \underline{w}, b, c_1, c_2) = \frac{\bar{w}-\underline{w}}{2bc_2}$ *and* $\bar{e}_1(\bar{w} - \underline{w}, b', c_1, c_2) = \frac{\bar{w}-\underline{w}}{2b'c_2}$*, while* $\bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2) = \frac{c_1(\bar{w}-\underline{w})}{2bc_2^2}$ *and* $\bar{e}_2(\bar{w} - \underline{w}, b', c_1, c_2) = \frac{c_1(\bar{w}-\underline{w})}{2b'c_2^2}$*. As* $b' > b$*, it follows that the denominator is strictly larger in both* $\bar{e}_1(\bar{w} - \underline{w}, b', c_1, c_2)$ *and*

$\bar{e}_2(\bar{w}-\underline{w}, b', c_1, c_2)$ than in $\bar{e}_1(\bar{w}-\underline{w}, b, c_1, c_2)$ and $\bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2)$, respectively. Since the numerator is the same in both cases, we conclude that $\bar{e}_i(\bar{w}-\underline{w}, b', c_1, c_2) < \bar{e}_i(\bar{w}-\underline{w}, b, c_1, c_2)$, for $i = 1, 2$.

Proposition 3

Proof 4 *In the model without morale concerns $g_2(b, \bar{w}-\underline{w}) = 1 = g_2(b, \bar{\bar{w}}-\underline{\underline{w}})$. Moreover, as $\bar{w}-\underline{w} \leq \bar{\bar{w}}-\underline{\underline{w}}$, we have that $\bar{e}_1(\bar{w}-\underline{w}, b, c_1, c_2) = \frac{\bar{w}-\underline{w}}{2bc_2} \leq \frac{\bar{\bar{w}}-\underline{\underline{w}}}{2bc_2} = \bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2)$, and $\bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2) = \frac{c_1(\bar{w}-\underline{w})}{2bc_2^2} \leq \frac{c_1(\bar{\bar{w}}-\underline{\underline{w}})}{2bc_2^2} = \bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2)$. If follows that the average effort of both players decreases as pay progression increases.*

Proposition 4

Proof 5 *Note that $\bar{e}_i(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) \lesseqgtr \bar{e}_i(\bar{w}-\underline{w}, b, c_1, c_2)$ if and only if $\bar{e}_i(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) - \bar{e}_i(\bar{w}-\underline{w}, b, c_1, c_2) \lesseqgtr 0$. As morale cost-shifts are normalized to 1, we focus on the following expressions:*

$$\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) - \bar{e}_1(\bar{w}-\underline{w}, b, c_1, c_2) = \frac{1}{2bc_2}\left((\bar{\bar{w}}-\underline{\underline{w}}) - (\bar{w}-\underline{w})\right)$$
$$\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) - \bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2) = \frac{c_1}{2bc_2^2}\left((\bar{\bar{w}}-\underline{\underline{w}}) - (\bar{w}-\underline{w})\right)$$

Because $\bar{\bar{w}}-\underline{\underline{w}} \geq \bar{w}-\underline{w}$, $b \geq 1$, $c_2 > 0$ and $c_1 \geq 0$, it follows that these expressions are strictly greater than zero. Therefore, $\bar{e}_i(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) \geq \bar{e}_i(\bar{w}-\underline{w}, b, c_1, c_2)$, for $i = 1, 2$. As b is only in the denominator of the multiplicative term for both expressions, we conclude that a decrease in b leads to an increase in average effort for $i = 1, 2$.

Note that the relative magnitude of the change in effort for player 1 and player 2 is ambiguous, and ultimately depends on whether $c_1 < c_2$ or $c_1 > c_2$ (both of which are possible).

Proposition 5

Proof 6 *From the expressions of the average effort for each player, we know that:*

$$\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)=\frac{(\bar{w}-\underline{w})}{2\tilde{c}_2}\left(\frac{1}{b}-\frac{1}{b'}\right)$$
$$\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)=\frac{c_1(\bar{w}-\underline{w})}{2\tilde{c}_2^2}\left(\frac{1}{b}-\frac{1}{b'}\right)$$

$$\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)=\frac{(\bar{w}-\underline{w})}{2\tilde{\tilde{c}}_2}\left(\frac{1}{b}-\frac{1}{b'}\right)$$
$$\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)=\frac{c_1(\bar{w}-\underline{w})}{2\tilde{\tilde{c}}_2^2}\left(\frac{1}{b}-\frac{1}{b'}\right)$$

As $\tilde{c}_2$ and $\tilde{\tilde{c}}_2$ only show up in the denominator of each difference in average effort, which is positive by Proposition 2, for $\tilde{c}_2 > \tilde{\tilde{c}}_2$ we have that $\bar{e}_i(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)-\bar{e}_i(\bar{w}-\underline{w},b',c_1,\tilde{c}_2) < \bar{e}_i(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_i(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)$ for $i=1,2$.

Proposition 6

Proof 7 *From the expressions of the average effort for each player, we know that:*

$$\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)=\frac{1}{2b\tilde{c}_2}\left((\bar{\bar{w}}-\underline{\underline{w}})-(\bar{w}-\underline{w})\right)$$
$$\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)=\frac{c_1}{2b\tilde{c}_2^2}\left((\bar{\bar{w}}-\underline{\underline{w}})-(\bar{w}-\underline{w})\right)$$

$$\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)=\frac{1}{2b\tilde{\tilde{c}}_2}\left((\bar{\bar{w}}-\underline{\underline{w}})-(\bar{w}-\underline{w})\right)$$
$$\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)=\frac{c_1}{2b\tilde{\tilde{c}}_2^2}\left((\bar{\bar{w}}-\underline{\underline{w}})-(\bar{w}-\underline{w})\right)$$

As $\tilde{c}_2$ and $\tilde{\tilde{c}}_2$ only show up in the denominator of each difference in average effort, which are positive by Proposition 3, for $\tilde{c}_2 > \tilde{\tilde{c}}_2$ we have that $\bar{e}_i(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-$

$\bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2) < \bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)$ *for* $i = 1, 2$.

Proposition 7

Proof 8 *From the expressions of the average effort for each player, we know that:*

$$\left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right) - \left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{\tilde{c}}_2) - \bar{e}_1(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)\right) = \frac{1}{\tilde{\tilde{c}}_2}\left(\frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b} - \frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b'}\right)$$

$$\left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right) - \left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_1(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\right) = \frac{1}{\tilde{c}_2}\left(\frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b} - \frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b'}\right)$$

$$\left(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right) - \left(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{\tilde{c}}_2) - \bar{e}_2(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)\right) = \frac{c_1}{\tilde{\tilde{c}}_2^2}\left(\frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b} - \frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b'}\right)$$

$$\left(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right) - \left(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_2(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\right) = \frac{c_1}{\tilde{c}_2^2}\left(\frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b} - \frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b'}\right)$$

The term within the brackets $\left(\frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b} - \frac{\left(\bar{\bar{w}} - \underline{\underline{w}}\right) - (\bar{w} - \underline{w})}{2b'}\right)$ *is the same in each expression. Because* $\tilde{c}_2$ *and* $\tilde{\tilde{c}}_2$ *only show up in the denominator of the term outside of the brackets of each of the difference-in-differences of average effort, which are positive from Proposition 4, for* $\tilde{c}_2 > \tilde{\tilde{c}}_2$ *we have that:*

$$\left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right) - \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)\right) > \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right) - \left(\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\right)$$

for $i = 1, 2$.

Proofs: Model with Morale Concerns

Proposition 8

Proof 9 *We have that* $\bar{e}_1(\bar{w}-\underline{w}, b', c_1, c_2) = \frac{\bar{w}-\underline{w}}{2b'c_2g_2(b',\bar{w}-\underline{w})}$ *and* $\bar{e}_1(\bar{w}-\underline{w}, b', c_1, c_2) = \frac{\bar{w}-\underline{w}}{2b'c_2g_2(b',\bar{w}-\underline{w})}$, *while* $\bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2) = \frac{c_1(\bar{w}-\underline{w})}{2bc_2^2g_2(b,\bar{w}-\underline{w})^2}$ *and* $\bar{e}_2(\bar{w}-\underline{w}, b', c_1, c_2) = \frac{c_1(\bar{w}-\underline{w})}{2b'c_2^2g_2(b',\bar{w}-\underline{w})^2}$. *By assumption,* $b' > b$ *implies that* $g_2(b', \bar{w}-\underline{w}) > g_2(b, \bar{w}-\underline{w})$. *It thus follows that the denominator is strictly larger in both* $\bar{e}_1(\bar{w}-\underline{w}, b', c_1, c_2)$ *and* $\bar{e}_2(\bar{w}-\underline{w}, b', c_1, c_2)$ *than in* $\bar{e}_1(\bar{w}-\underline{w}, b, c_1, c_2)$ *and* $\bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2)$, *respectively. As the numerator is the same in both cases, we conclude that* $\bar{e}_i(\bar{w}-\underline{w}, b', c_1, c_2) < \bar{e}_i(\bar{w}-\underline{w}, b, c_1, c_2)$, *for* $i = 1, 2$.

Proposition 9

Proof 10 *Note that* $\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) \lesseqgtr \bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2)$ *if and only if* $\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) - \bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2) \lesseqgtr 0$.
Hence, we focus on the following expressions

$$\begin{aligned}\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) - \bar{e}_1(\bar{w}-\underline{w}, b, c_1, c_2) &= \frac{(\bar{\bar{w}}-\underline{\underline{w}})}{2bc_2g_2(b, \bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{2bc_2g_2(b, \bar{w}-\underline{w})} \\ &= (\bar{w}-\underline{w})(\bar{\bar{w}}-\underline{\underline{w}})\frac{\frac{g_2(b,\bar{w}-\underline{w})}{\bar{w}-\underline{w}} - \frac{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}{\bar{\bar{w}}-\underline{\underline{w}}}}{2bc_2g_2(b, \bar{\bar{w}}-\underline{\underline{w}})g_2(b, \bar{w}-\underline{w})}\end{aligned}$$

$$\begin{aligned}\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) - \bar{e}_2(\bar{w}-\underline{w}, b, c_1, c_2) &= \frac{c_1(\bar{\bar{w}}-\underline{\underline{w}})}{2bc_2^2g_2(b, \bar{\bar{w}}-\underline{\underline{w}})^2} - \frac{c_1(\bar{w}-\underline{w})}{2bc_2^2g_2(b, \bar{w}-\underline{w})^2} \\ &= c_1(\bar{w}-\underline{w})(\bar{\bar{w}}-\underline{\underline{w}})\frac{\frac{g_2(b,\bar{w}-\underline{w})^2}{\bar{w}-\underline{w}} - \frac{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}{\bar{\bar{w}}-\underline{\underline{w}}}}{2bc_2^2g_2(b, \bar{\bar{w}}-\underline{\underline{w}})^2g_2(b, \bar{w}-\underline{w})^2}\end{aligned}$$

We will proceed by showing that there exists a $\bar{b}_2$ *such that* $\frac{g_2(\bar{b}_2,\bar{w}-\underline{w})^2}{\bar{w}-\underline{w}} = \frac{g_2(\bar{b}_2,\bar{\bar{w}}-\underline{\underline{w}})^2}{\bar{\bar{w}}-\underline{\underline{w}}}$ *and a* $\bar{b}_1$ *such that* $\frac{g_2(\bar{b}_1,\bar{w}-\underline{w})}{\bar{\bar{w}}-\underline{w}} = \frac{g_2(\bar{b}_1,\bar{\bar{w}}-\underline{\underline{w}})}{\bar{\bar{w}}-\underline{\underline{w}}}$. *We will equivalently show that* $\frac{g_2(\bar{b}_1,\bar{w}-\underline{w})}{g_2(\bar{b}_1,\bar{\bar{w}}-\underline{\underline{w}})} = \frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}}$ *for some* $\bar{b}_1$ *and* $\frac{g_2(\bar{b}_2,\bar{w}-\underline{w})}{g_2(\bar{b}_2,\bar{\bar{w}}-\underline{\underline{w}})} = \frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}}$ *for some* $\bar{b}_2$.

First, note that $g_2(b, \bar{w} - \underline{w})$ *and* $g_2(b, \bar{\bar{w}} - \underline{\underline{w}})$ *are continuous in* b *and are strictly greater than 1. It follows that* $\frac{g_2(b,\bar{w}-\underline{w})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}$ *is continuous.*

Second, we have that $\frac{g_2(1,\bar{w}-\underline{w})}{g_2(1,\bar{\bar{w}}-\underline{\underline{w}})} = 1 > \frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}}$ *and* $\frac{g_2(1,\bar{w}-\underline{w})}{g_2(1,\bar{\bar{w}}-\underline{\underline{w}})} = 1 > \frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}}$. *Thus, there exists some point such that* $\frac{g_2(b,\bar{w}-\underline{w})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}$ *is above* $\frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}}$ *and* $\frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}}$. *From Assumption 2, we know that in the limit* $\lim_{b\to\infty}\left(\frac{g_2(1,\bar{w}-\underline{w})}{g_2(1,\bar{\bar{w}}-\underline{\underline{w}})}\right) = 0 < \frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}}$ *and* $\lim_{b\to\infty}\left(\frac{g_2(1,\bar{w}-\underline{w})}{g_2(1,\bar{\bar{w}}-\underline{\underline{w}})}\right) = 0 < \frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}}$. *Therefore there exists some point such that* $\frac{g_2(b,\bar{w}-\underline{w})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}$ *is below* $\frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}}$ *and* $\frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}}$. *From the continuity of the function* $\frac{g_2(b,\bar{w}-\underline{w})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}$ *in* b*, there exists some* $\bar{b}_2$ *such that* $\frac{g_2(\bar{b}_2,\bar{w}-\underline{w})}{g_2(\bar{b}_2,\bar{\bar{w}}-\underline{\underline{w}})} = \frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}}$*, and therefore* $\frac{g_2(\bar{b}_2,\bar{w}-\underline{w})^2}{\bar{w}-\underline{w}} = \frac{g_2(\bar{b}_2,\bar{\bar{w}}-\underline{\underline{w}})^2}{\bar{\bar{w}}-\underline{\underline{w}}}$. *There also exists some* $\bar{b}_1$ *such that* $\frac{g_2(\bar{b}_1,\bar{w}-\underline{w})}{g_2(\bar{b}_1,\bar{\bar{w}}-\underline{\underline{w}})} = \frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}}$*, and therefore* $\frac{g_2(\bar{b}_1,\bar{w}-\underline{w})}{\bar{\bar{w}}-\underline{w}} = \frac{g_2(\bar{b}_1,\bar{\bar{w}}-\underline{\underline{w}})}{\bar{\bar{w}}-\underline{\underline{w}}}$.

Finally, take $\bar{b}$ *to be the infimum of all such* $\bar{b}_2$*, ensuring that* $\frac{g_2(b,\bar{w}-\underline{w})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})} > \frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}} > \frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}}$ *for all* $b < \bar{b}$. *Conversely, take* $\bar{\bar{b}}$ *to be the supremum of all such* $\bar{b}_1$*, ensuring that* $\frac{g_2(b,\bar{w}-\underline{w})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})} < \frac{\bar{w}-\underline{w}}{\bar{\bar{w}}-\underline{\underline{w}}} < \frac{(\bar{w}-\underline{w})^{1/2}}{(\bar{\bar{w}}-\underline{\underline{w}})^{1/2}}$ *for all* $b > \bar{\bar{b}}$. *This implies that,* $\frac{g_2(b,\bar{w}-\underline{w})}{\bar{w}-\underline{w}} > \frac{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}{\bar{\bar{w}}-\underline{\underline{w}}}$ *and* $\frac{g_2(b,\bar{w}-\underline{w})^2}{\bar{w}-\underline{w}} > \frac{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}{\bar{\bar{w}}-\underline{\underline{w}}}$ *for all* $b < \bar{b}$. *Therefore,* $\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) > \bar{e}_1(\bar{w} - \underline{w}, b, c_1, c_2)$ *and* $\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) > \bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2)$ *for all* $b < \bar{b}$. *Moreover, we also have that* $\frac{g_2(b,\bar{w}-\underline{w})}{\bar{w}-\underline{w}} < \frac{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}{\bar{\bar{w}}-\underline{\underline{w}}}$ *and* $\frac{g_2(b,\bar{w}-\underline{w})^2}{\bar{w}-\underline{w}} < \frac{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}{\bar{\bar{w}}-\underline{\underline{w}}}$ *for all* $b > \bar{\bar{b}}$*, implying that* $\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) < \bar{e}_1(\bar{w}-\underline{w}, b, c_1, c_2)$ *and* $\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}}, b, c_1, c_2) < \bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2)$ *for all* $b > \bar{\bar{b}}$.

Proposition 10

Proof 11 *Note that* $\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) \lesseqgtr \bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2)$ *if and only if* $\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2) \lesseqgtr 0$. *We, therefore, focus on the following expres-*

sions

$$\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) - \bar{e}_1(\bar{w} - \underline{w}, b, c_1, c_2) = \frac{(\bar{\bar{w}} - \underline{\underline{w}})}{2bc_2 g_2(b, \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{2bc_2 g_2(b, \bar{w} - \underline{w})}$$
$$= \frac{1}{2bc_2}\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})}\right)$$
$$\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2) = \frac{c_1(\bar{\bar{w}} - \underline{\underline{w}})}{2bc_2^2 g_2(b, \bar{\bar{w}} - \underline{\underline{w}})^2} - \frac{c_1(\bar{w} - \underline{w})}{2bc_2^2 g_2(b, \bar{w} - \underline{w})^2}$$
$$= \frac{c_1}{2bc_2^2}\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})^2} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})^2}\right)$$

We proceed by showing that whenever the difference of effort is positive, such difference is decreasing in b.

First, note that $\frac{1}{2bc_2}$ *and* $\frac{c_1}{2bc_2^2}$ *are always decreasing in b.*
Second, we show that $\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})}\right)$ *and* $\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2} - \frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})^2}\right)$ *are decreasing in b. Take any* $b' > b$*. Given the log super-modularity of* g_2*, we have that* $g_2(b, \bar{w} - \underline{w})g_2(b', \bar{\bar{w}} - \underline{\underline{w}}) \geq g_2(b', \bar{w} - \underline{w})g_2(b, \bar{\bar{w}} - \underline{\underline{w}})$ *and therefore* $g_2(b', \bar{\bar{w}} - \underline{\underline{w}}) \geq \frac{g_2(b',\bar{w}-\underline{w})g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{w}-\underline{w})}$*. By substituting this expression into* $\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b',\bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})}\right)$ *we obtain:*

$$\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b',\bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})}\right) \leq \left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{\frac{g_2(b',\bar{w}-\underline{w})g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{w}-\underline{w})}} - \frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})}\right) = \frac{g_2(b,\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})}\right).$$

As $g_2(b, \bar{w}-\underline{w}) \leq g_2(b', \bar{w}-\underline{w})$ *and the difference in effort is positive, i.e.,* $\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})} > 0$*, we have that* $\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b',\bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})}\right) \leq \left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})} - \frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})}\right)$*. The same argument holds for* $\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, c_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, c_2)$*.*

Proposition 11

Proof 12 *From the expressions of average effort we find that*

$$|\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)| = \frac{(\bar{w}-\underline{w})}{2\tilde{c}_2}\left|\left(\frac{1}{bg_2(b,\bar{w}-\underline{w})}-\frac{1}{b'g_2(b',\bar{w}-\underline{w})}\right)\right|$$
$$|\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)| = \frac{(\bar{w}-\underline{w})}{2\tilde{\tilde{c}}_2}\left|\left(\frac{1}{bg_2(b,\bar{w}-\underline{w})}-\frac{1}{b'g_2(b',\bar{w}-\underline{w})}\right)\right|$$
$$|\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)| = \frac{c_1(\bar{w}-\underline{w})}{2\tilde{c}_2^2}\left|\left(\frac{1}{bg_2(b,\bar{w}-\underline{w})^2}-\frac{1}{b'g_2(b',\bar{w}-\underline{w})^2}\right)\right|$$
$$|\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2))| = \frac{c_1(\bar{w}-\underline{w})}{2\tilde{\tilde{c}}_2^2}\left|\left(\frac{1}{bg_2(b,\bar{w}-\underline{w})^2}-\frac{1}{b'g_2(b',\bar{w}-\underline{w})^2}\right)\right|$$

As $\tilde{c}_2$ and $\tilde{\tilde{c}}_2$ only shows up in the denominator of each average effort, and the multiplicative term is the same, for $\tilde{c}_2 > \tilde{\tilde{c}}_2$ we have that $|\bar{e}_i(\bar{w}-\underline{w},b,c_1,\tilde{c}_2) - \bar{e}_i(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)| < |\bar{e}_i(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_i(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)|$ for $i=1,2$.

Proposition 12

Proof 13

$$\left|\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right| = \frac{1}{2b\tilde{c}_2}\left|\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})}\right)\right|$$
$$\left|\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right| = \frac{1}{2b\tilde{\tilde{c}}_2}\left|\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})}\right)\right|$$
$$\left|\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right| = \frac{c_1}{2b\tilde{c}_2^2}\left|\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})^2}\right)\right|$$
$$\left|\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right| = \frac{c_1}{2b\tilde{\tilde{c}}_2^2}\left|\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})^2}\right)\right|$$

Note that $\tilde{c}_2 \geq \tilde{\tilde{c}}_2$ *and thus* $\frac{1}{2b\tilde{c}_2} \leq \frac{1}{2b\tilde{\tilde{c}}_2}$ *and* $\frac{c_1}{2b\tilde{c}_2^2} \leq \frac{c_1}{2b\tilde{\tilde{c}}_2^2}$. *From here,*

$$\left|\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right| = \frac{1}{2b\tilde{c}_2}\left|\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})}\right)\right|$$
$$\leq \frac{1}{2b\tilde{\tilde{c}}_2}\left|\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})}\right)\right| = \left|\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right|$$

and

$$\left|\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right| = \frac{c_1}{2b\tilde{c}_2^2}\left|\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})^2} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})^2}\right)\right|$$
$$\leq \frac{c_1}{2b\tilde{\tilde{c}}_2^2}\left|\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})^2} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})^2}\right)\right| = \left|\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right|$$

We conclude that $|\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)| \geq |\bar{e}_i(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_i(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)|$, *for* $i = 1, 2$. *That is, the impact of pay progression on effort is amplified when player* 2 *is of higher ability, regardless the direction of change.*

Proposition 13

Proof 14 *From Proposition 10, we know that all the difference-in-differences of average effort are positive for all players in this region. For player* 1*, we have that:*

$$\left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\right) - \left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{\tilde{c}}_2) - \bar{e}_1(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)\right) =$$
$$\frac{1}{\tilde{\tilde{c}}_2}\left(\frac{1}{2b}\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})}\right) - \frac{1}{2b'}\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b', \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{g_2(b', \bar{w} - \underline{w})}\right)\right)$$
$$\left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_1(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\right) - \left(\bar{e}_1(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_1(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\right) =$$
$$\frac{1}{\tilde{c}_2}\left(\frac{1}{2b}\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b, \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{g_2(b, \bar{w} - \underline{w})}\right) - \frac{1}{2b'}\left(\frac{(\bar{\bar{w}} - \underline{\underline{w}})}{g_2(b', \bar{\bar{w}} - \underline{\underline{w}})} - \frac{(\bar{w} - \underline{w})}{g_2(b', \bar{w} - \underline{w})}\right)\right)$$

Note that the expression within the brackets,

$\left(\frac{1}{2b}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})}\right)-\frac{1}{2b'}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b',\bar{\bar{w}}-\underline{\underline{w}})}-\frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})}\right)\right)$, *is the same within both* $\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right)-\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)\right)$ *and* $\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right)-\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)\right)$. *Further, it is positive by proposition 10. The multiplicative term outside of the brackets is given by* $\frac{1}{\tilde{\tilde{c}}_2}$ *and* $\frac{1}{\tilde{c}_2}$ *respectively for* $\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right)-\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)\right)$ *and* $\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right)-\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)\right)$. *As* $\tilde{\tilde{c}}_2<\tilde{c}_2$ *we conclude that*

$$\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right)-\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{\tilde{c}}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)\right)>$$
$$\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right)-\left(\bar{e}_1(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)\right)$$

For player 2, we have instead:

$$\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right)-\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)\right)=$$
$$\frac{1}{\tilde{\tilde{c}}_2^2}\left(\frac{c_1}{2b}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})^2}\right)-\frac{c_1}{2b'}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b',\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})^2}\right)\right)$$
$$\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_1(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right)-\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)\right)=$$
$$\frac{1}{\tilde{c}_2^2}\left(\frac{c_1}{2b}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})^2}\right)-\frac{c_1}{2b'}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b',\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})^2}\right)\right)$$

Note that the expression within the brackets, $\left(\frac{c_1}{2b}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b,\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b,\bar{w}-\underline{w})^2}\right)-\frac{c_1}{2b'}\left(\frac{(\bar{\bar{w}}-\underline{\underline{w}})}{g_2(b',\bar{\bar{w}}-\underline{\underline{w}})^2}-\frac{(\bar{w}-\underline{w})}{g_2(b',\bar{w}-\underline{w})^2}\right)\right)$, *is the same within both* $\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right)-\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{c}_2)\right)$ *and* $\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right)-\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)\right)$. *Further, it is positive by proposition 10. The multiplicative term outside of the brackets is given by* $\frac{1}{\tilde{\tilde{c}}_2^2}$ *and* $\frac{1}{\tilde{c}_2^2}$ *respectively for* $\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{\tilde{c}}_2)\right)-\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b',c_1,\tilde{\tilde{c}}_2)-\bar{e}_2(\bar{w}-\underline{w},b',c_1,\tilde{\tilde{c}}_2)\right)$ *and* $\left(\bar{e}_2(\bar{\bar{w}}-\underline{\underline{w}},b,c_1,\tilde{c}_2)-\bar{e}_2(\bar{w}-\underline{w},b,c_1,\tilde{c}_2)\right)-$

$\big(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_2(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\big)$. *As* $\tilde{\tilde{c}}_2 < \tilde{c}_2$, *we can conclude that*

$$\big(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{\tilde{c}}_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{\tilde{c}}_2)\big) - \big(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{\tilde{c}}_2) - \bar{e}_2(\bar{w} - \underline{w}, b', c_1, \tilde{\tilde{c}}_2)\big) >$$
$$\big(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b, c_1, \tilde{c}_2) - \bar{e}_2(\bar{w} - \underline{w}, b, c_1, \tilde{c}_2)\big) - \big(\bar{e}_2(\bar{\bar{w}} - \underline{\underline{w}}, b', c_1, \tilde{c}_2) - \bar{e}_2(\bar{w} - \underline{w}, b', c_1, \tilde{c}_2)\big)$$

Chapter 3

Better Containment but Less Health Access: How Past Exposure to Health Crises Affects the Covid-19 Response

3.1 Introduction

Why did Liberia declare a national health emergency and close down most public places within six days of recording the first Covid-19 case in the country whereas in the United States life continued mostly uninterrupted for weeks after the first Covid-19 case was recorded? Why did South Korea manage to contain the spread of Covid-19 very well early in the pandemic whereas some regions in Italy lost control of the outbreak and were unable to admit patients to already overcrowded hospitals? In this paper I look at the determinants of the demand for containment policies which have played a critical role in slowing down the spread of Covid-19. While there likely is a plethora of reasons for the differential responses mentioned above, I argue in this paper that one part of the story relates to past exposure to similar health crises, e.g. the 2002 SARS outbreak which affected South Korea or the 2014 Ebola outbreak which affected Liberia, which has raised individuals' perceived risk of Covid-19 and has led to a higher demand for containment measures.

In this paper I look at Sierra Leone, which was one of the three countries that were majorly hit by the 2014-16 outbreak of the Ebola Virus Disease (EVD) in West Africa, together with Guinea and Liberia.[1] In particular, I compare households living in rural villages that were directly affected by EVD with households living in villages that were not directly affected by EVD. I find that households in affected villages perceive a higher risk from Covid-19, are more likely to know about the possibility of asymptomatic transmissions of Covid-19, and are more likely to trust health care professionals when it comes to information about Covid-19. These effects translate into affected villages being more likely to have organized the public distribution of face masks within six months of the first confirmed Covid-19 case in the country. However, I also find that the increased caution of households in affected villages comes at the cost of reduced health access stemming from a fear of contracting Covid-19 at health facilities. These results provide novel micro evidence for the way that increased risk perception triggered by past traumatic events can raise the demand for Covid-19 containment measures which in turn have been crucial for slowing down the spread of the virus. Thus, this paper identifies a potentially important channel through which the demand for containment measures is determined, although it cannot speak to its relative importance vis-a-vis other potential channels. This paper also does not touch the determinants of the effective supply of containment measures, which likely differs substantially across countries and regions even if the demand for containment measures is similar.

[1]IRB for this study has been approved by Universitat Pompeu Fabra (CIREP Approval Number 163).

Estimating the causal effect of past exposure to EVD - or any other past event that could shift beliefs about Covid-19 - on Covid-19-related outcomes is not straightforward. Generally speaking, any beliefs which affect individual or collective responses to the Covid-19 pandemic could have affected their response to the 2014-16 EVD epidemic as well. In order to get around this endogeneity issue, one must use plausibly exogenous variation in EVD exposure. In my empirical exercise I use primary data collected for this study from surveys with household respondents and health workers from 535 villages in Sierra Leone during the Covid-19 pandemic in October and November 2020. First, I create a binary EVD indicator for each village which cross-validates responses from households and health workers in the same village. I then compare outcomes related to risk perception, trust in health professionals, Covid-19 prevention, and health access between households in villages directly affected by EVD and those not directly affected by EVD controlling for geographic village characteristics and demographic household characteristics. As EVD-affected and non-affected villages might still differ in other unobserved dimensions, I then use the village road distance to the place of the EVD index case as an instrumental variable for EVD exposure controlling for a set of geographic characteristics. The idea is that geographically comparable villages that are closer to the index case were more likely to be affected by EVD than villages further away. My data confirms that available demographic characteristics of households do not predict EVD exposure.

I find that past exposure to an epidemic significantly affects beliefs about Covid-19: I estimate that household respondents in EVD-affected villages are 50 pp. more likely to believe they can contract Covid-19 and 32 pp. more likely to know that it is possible for asymptomatic individuals to infect others with Covid-19. Furthermore, households in EVD-affected villages are 44 pp. more likely to say that they trust health professionals when it comes to information about Covid-

19. This translates into stronger containment measures in EVD-affected villages which are 44 pp. more likely to have organized the public distribution of face masks. However, households in EVD-affected villages were also 43 pp. more likely to have avoided going to the local health clinic out of fear of contracting Covid-19. These results highlight how past exposure to a health crisis persistently affects individuals' risk perceptions and leads to a higher demand for prevention and containment measures. At the same time, the increase in the perceived risk of Covid-19 reduces health access for households.

This paper contributes to two strands of the literature. First, this study adds to the literature on the demand for policy responses to an emergency and support for these measures. A number of studies suggests that countries affected by the 2002 SARS outbreak, the 2012 MERS outbreak, or the 2014 Ebola outbreak were better prepared for the current Covid-19 pandemic and adopted decisive measures early on (Fotiou and Lagerborg, 2021; Chua et al., 2021), which in turn has been linked with better containment of and a faster recovery from the Covid-19 pandemic (Caselli et al., 2021). The potential reasons mentioned for this link between past exposure to health crises and better containment include the training of primary health care providers, centralized data-driven surveillance systems, widespread testing and contact tracing, and increased citizen demand. While these studies rely on cross-country comparisons or evidence from case studies, I provide micro evidence for a specific mechanism linking past exposure to a higher demand for containment measures: I show that past exposure to the 2014-16 Ebola epidemic on the village-level increases households' perceived risk of Covid-19, increases their trust in health professionals, and increases the likelihood of adopting preventive measures.

Second, this study adds to the literature on the trade-off between fewer Covid-

19 cases and higher socio-economic costs resulting from containment measures.[2] In particular, I add to the literature on the consequences of Covid-19 on health care access. Overall, there is evidence for a great reduction of health care access due to Covid-19 (World Health Organization, 2021). On the supply side, there is evidence for an increase in the morbidity and mortality of health workers which has negatively affected health care access (Bandyopadhyay et al., 2020; Gholami et al., 2021). Furthermore, many health workers and health care facilities were utilized in the Covid-19 response and therefore unavailable for non-Covid-related services (Abelson, 2021). On the demand side, there is evidence that mobility restrictions and lockdowns prevented people from accessing health care services (Cantor et al., 2020). I add to this literature by highlighting another channel limiting health access for individuals, namely decreased demand due to a fear of contracting Covid-19 at the health facility. Furthermore, I show that this effect is stronger for households in villages previously affected by EVD, probably because they perceive Covid-19 to be riskier than households in non-affected villages. There already is some evidence for a deterioration of non-Covid-related health outcomes during the pandemic (Jain and Dupas, 2021) and the predicted negative impact of the current decrease in health access on future health outcomes is substantial (McQuaid et al., 2020; Glaziou, 2020; Sherrard-Smith et al., 2020). Thus, understanding individual beliefs and risk perceptions around health care services is crucial in order to improve health access during the Covid-19 pandemic.

The rest of the paper is structured as follows: Section 3.2 provides some background about the 2014-16 Ebola epidemic and the Covid-19 pandemic in Sierra Leone. Section 3.3 describes the data and variables used in this paper. Section 3.4 lays out the empirical strategy. Section 3.5 discusses the results. Section 3.6

[2] Miguel and Mobarak (2021) provide a great review of the empirical literature on the economic and non-economic consequences of the Covid-19 pandemic, with a focus on low- and middle-income countries.

concludes.

3.2 Background

3.2.1 The 2014-16 Ebola Epidemic

Together with Guinea and Liberia, Sierra Leone was one of the three countries that experienced widespread transmission of EVD following the 2014 EVD outbreak in West Africa. Over the course of the 2014-16 West African EVD epidemic, 28,616 cases and 11,310 deaths were recorded in the three countries, making it the deadliest EVD outbreak in history. Out of these, Sierra Leone accounted for 14,124 cases and 3,956 deaths. Different from other viral diseases, including Covid-19, EVD does not spread through airborne transmission but through people getting in contact with body fluids of an infected person. The index case of the 2014-16 EVD outbreak occurred when a 18-month old toddler came into contact with an infected fruit bat in the town of Meliandou, Guinea, close to Sierra Leone's eastern border.[3] As Fang et al. (2016) show, the subsequent spread of EVD in Sierra Leone closely followed the country's road network starting from the eastern border with a second transmission corridor originating from Freetown, the country's capital city, emerging in the later stages of the outbreak. There are a number of aspects of the 2014-16 EVD outbreak that might have affected the way people react to future disease outbreaks: First, with an average case fatality rate of around 50% EVD is much more deadly than most other common diseases. Second, while EVD transmission happens through body fluids and is not airborne, the set of effective containment measures is very similar to that for other viral diseases. The key to halt the spread of EVD is to limit human-to-human interaction, to

[3] A brief summary about the 2014-16 West African EVD outbreak is available from the Centers for Disease Control and Prevention.

identify, trace, and isolate infected people, and to take accompanying measures that raise awareness and generate buy-in from the population.[4] In Sierra Leone some of these EVD containment measures created tensions as they interfered with deeply rooted burial traditions or because the bodies of the deceased were taken away from villages by health workers in hazmat suits who were unknown to the local population.[5] Third, a disproportionate share of people who became infected with EVD were health workers (Evans et al., 2015). Among other potential negative consequences on the provision of health services, this might have affected people's risk perceptions about interacting with health workers or visiting health care facilities during a disease outbreak.

3.2.2 Covid-19 in Sierra Leone

As of April 2022, Sierra Leone has recorded just under 7,700 Covid-19 cases, a remarkably low number given its population of around 8 million, although part of the reason this number is so low is likely the lack of testing capacities as in many African countries.[6] Compared to most Western governments, the government of Sierra Leone took strong measures to halt the spread of Covid-19 very early relative to the local spread of the pandemic: Even before the first Covid-19 case was recorded in the country on March 31, 2020, the government closed its border for international air traffic and mandated the closure of places of worship as well as schools and universities. Upon recording the first case, the government imposed a 3-day country-wide lockdown at the start of April and after that announced additional containment measures including the ban of inter-district travel and a

[4]A chronology of previous EVD outbreaks, their case fatality rates, and current recommendations regarding disease control for EVD is available from the World Health Organization.

[5]A great ground report on the issues around burial practices from the time of the 2014-16 EVD outbreak in Sierra Leone is available from the National Geographic.

[6]Cumulative case numbers are regularly updated by the World Health Organization. A lack of testing capacities in many African countries has been reported by several media such as the BBC and Scientific American.

nighttime curfew.[7] Not only did the government react quickly and in a determined way, but also people's compliance with and support for government measures has been quite high: Ndubuisi-Obi et al. (2021) find that human movement indeed decreased dramatically during the government-mandated lockdowns and Solís Arce et al. (2020) show that (self-reported) compliance with wearing face masks and hygienic protocols was high among Sierra Leoneans in June 2020.

3.3 Data and Measurement

3.3.1 Data

This study leverages a number of different data sources:

Household Survey

In October and November 2020 enumerators conducted a phone survey with 879 primary care givers from households in 535 villages across 6 districts in Sierra Leone. The survey covered attitudes and beliefs about Covid-19, compliance with containment measures, utilization of health services, and basic demographic information. In addition to that, households were asked whether there had been an Ebola case in their village during the 2014-16 Ebola epidemic.

Village-Level Survey

In the same time period, enumerators conducted a survey with the local health worker responsible for each village. The goal of this survey was to get village-level information both about Covid-19 containment measures that were taken in the village and whether there had been Ebola cases in the village during the 2014-16 Ebola epidemic.

Geographic Data

[7] A graphical illustration of the timeline of government containment measures for Covid-19 in Sierra Leone is available from Reuters.

Village-level geographic data used in this study include GPS coordinates and altitude collected for a previous study by Deserranno et al. (2022a). In addition, the study utilizes a shapefile of Sierra Leone's road network from World Bank Group (2009) in order to calculate the Euclidean distance between each village and the nearest main road.

3.3.2 Measuring Ebola Exposure

Based on the household survey and the village-level survey, I construct a measure of Ebola exposure in the following way: For every village, the binary EVD indicator is equal to 1 if the health worker and at least one household in the village say that there had been an EVD case in the village during the 2014-16 Ebola epidemic. It is equal to 0 if neither the health worker nor any household say that there had been an EVD case in the village during the 2014-16 Ebola epidemic. Villages are removed from the sample if neither of these two conditions hold. I choose this conservative approach to measuring EVD incidence in order to minimize measurement error and ensure having a sufficiently strong first-stage for my IV estimation strategy. Figure 3.1 shows the distribution of villages with and without EVD cases across Sierra Leone. In total, 10 % of the villages in my sample were directly exposed to EVD according to my measure. An alternative measure of EVD incidence would be official cases recorded at the time of the epidemic. However, the number of official cases is not available on the village level. In addition to that, Christensen et al. (2021) show that the official number of cases depends on sick people reporting to health facilities which conditional on being sick is a function of several factors such as trust in the health system and cultural beliefs. To alleviate similar concerns for my measure of EVD incidence I will use an instrumental variable approach that deals with both endogeneity concerns and potential measurement error.

3.3.3 Outcome Variables

Risk Perception and Knowledge Related to Covid-19

My main measure of household risk perception is the household respondents' answer to the question whether or not they think it is possible for them to contract Covid-19. Perceiving a risk of contracting Covid-19 clearly seems a pre-requisite for households to adopt appropriate preventive measures. As Table 3.1 shows, only 37.3% of respondents say that they believe they can contract Covid-19. A substantially larger share of households in villages affected by EVD believe they can contract Covid-19 (54.9%) compared to households in villages not affected by EVD (33.8%).

I also asked households whether or not they know about the possibility of an asymptomatic person to transmit Covid-19 to others. This measure captures both the perceived risk of contracting Covid-19, which should be higher for those who believe they can contract the virus from asymptomatic people, and knowledge about Covid-19. Overall, 59.3% of households say they know about the possibility of asymptomatic transmissions of Covid-19. Again, this number is substantially higher in villages affected by EVD (72.8%) compared to villages not affected by EVD (56.6%).

Trust in Health Professionals

I asked households whether they trust health professionals when it comes to information about Covid-19. During the Ebola epidemic, difficulties in slowing down the spread of EVD have been partially attributed to the lack of locally trusted health workers providing information. At the same time, local health workers are involved in the fight against Covid-19 in Sierra Leone at the community level. Thus, the level of trust households have in health workers is an important factor for their perception and knowledge about Covid-19. In the whole sample, 42.1% of households say that they trust health professionals, with 63.1% in villages affected by EVD and 37.9% in villages not affected by EVD.

Covid-19 Prevention

My first set of measures of Covid-19 prevention is related to wearing face masks. Firstly, I ask household respondents about the occasions on which they wear a face mask, including when using public transport, attending places of worship, going to markets, being at work, and visiting friends or relatives. From that I construct a continuous measure ranging from 0 to 5 counting the number of occasions on which a respondent wears face masks. Households report to wear face masks on 3.1 occasions on average with households in villages affected by EVD wearing a face mask on 0.41 more occasions than households not affected by EVD. Secondly, I use as an outcome the answer of household respondents to the question what percentage of other adults in their village regularly wear a face mask at the local market. The reason for asking about the behavior of others and not only of the respondents themselves is that in a pilot I found that respondents tend to overstate their own compliance with mask mandates. With 85.5% in the whole sample the share of people wearing a face mask at the local market is quite high. There are slightly more people wearing a face mask at the local market in villages affected by EVD (90.5%) compared to villages not affected by EVD (84.5%).

The second set of measures of Covid-19 prevention is related to the willingness to pay for a Covid-19 vaccination. Firstly, I consider a continuous measure of the willingness to pay for a Covid-19 vaccination as reported by households. Given that the survey took place prior to the approval of the first Covid-19 vaccine, the question did not specify any of the currently used vaccines. Since the distribution of this variable is right-skewed with a large share of zero values, I use as my outcome variable the inverse hyperbolic sine transformation of the willingness to pay. As Table 3.1 shows, the willingness to pay for a vaccination is slightly lower in villages affected by EVD compared to villages not affected by EVD. Secondly, however, households in villages affected by EVD are more likely to have a very high willingness to pay as measured by the likelihood that their willingness to pay is above the 77th percentile of the overall distribution.[8] Table 3.1 shows that households in villages affected by EVD are 10 pp. more likely to have a very high willingness to pay for a Covid-19 vaccine compared to households in non-affected villages.

Finally, I use a measure from the village-level survey, namely whether or not there had been a public distribution of face masks in the village between the start of the Covid-19 measures in Sierra Leone in March/April 2020 and the time of the interview. In 44.2% of villages face masks were publicly distributed, but there is a stark difference between villages affected by EVD (67.5%) and villages not affected by EVD (41.6%).

Utilization of Health Services

My measure of health service utilization is whether or not households state that they have actively avoided visiting the local health clinic because they are afraid of contracting Covid-19. Given that health workers were disproportionately affected by EVD during the 2014-16 Ebola epidemic and there is evidence for households avoiding health services during the Ebola epidemic, it is important to understand the determinants of health care utilization during the Covid-19 pandemic. 29.7% of households say that they did avoid their local health clinic out of fear of contracting Covid-19. This share is substantially higher in villages affected by EVD (41.7%) than in villages not affected by EVD (27.2%). Delaying or avoiding necessary health services is a potentially important indirect cost of the Covid-19 pandemic that might be exacerbated by higher risk perceptions related to previous exposure to a health crisis as in Sierra Leone.

[8] The value of the willingness to pay at the 75th percentile of the distribution captures multiple percentiles which is why I choose the 77th percentile as a cutoff point. The results are robust to using different thresholds around the 75th percentile. Households with a willingness to pay above the 77th percentile are willing to pay SLL 10,000 or more for a Covid-19 vaccination.

3.4 Empirical Estimation

In the empirical estimation I always control for a set of geographic village characteristics which are correlated with the EVD indicator and could potentially affect outcome variables. The reason for including these geographic controls is that the transmission dynamics of EVD closely followed human travel routes (Fang et al., 2016). Thus, the probability of an EVD case in a village is expected to be lower in less accessible villages. This is captured by the three variables I use as controls: First, I control for a measure of a village's network centrality.[9] I follow Aggarwal et al. (2021) and measure network centrality of each village by the sum of population-weighted distances to the 25 largest towns in the country.[10] Formally, I measure the remoteness of village v as

$$remoteness_v = \sum_t distance_{vt} * weight_t \tag{3.1}$$

where $distance_{vt}$ is the road distance between village v and town t and $weight_t$ is the relative population of town t (compared to the rest of the 25 largest towns). My data confirms the negative relationship of this control variable with the probability of having an EVD case as the correlation between the EVD indicator and remoteness is -0.36 (significant at the 1% level). Second, I control for a village's distance to the nearest main road. Again the data confirms the relevance of this control variable as the correlation between the EVD indicator and the distance to

[9] In section 3.5 I show that my results are robust to using the distance to Freetown as a control instead. As most of Sierra Leone's population lives between the two central highways connecting Freetown with Gueckedou in Guinea, which is close to the place of the Ebola index case in Meliandou, controlling for the distance to Freetown is not ideal in this context. Conditional on the distance to Freetown, my instrumental variable, the distance to Meliandou, would primarily compare more densely populated places between the two highways with less densely populated places in the periphery. Using a measure of network centrality resolves this problem. I still report the robustness of my results to using the distance to Freetown as a control variable because a similar approach has been used in other settings in the literature.

[10] I use town populations from the 2004 Census as weights.

the nearest main road is -0.20 (significant at the 1% level). Third, I control for a village's altitude. Higher altitude places should be less accessible on average and therefore are expected to be less likely to have EVD cases. This is also confirmed in the data as the correlation between the EVD indicator and altitude is -0.10 (significant at the 1% level).

3.4.1 Ordinary Least Squares (OLS) and Probit

The OLS specification for my continuous outcome variables is

$$Y_{v,i} = \alpha + \beta * T_v + \delta * X_v + \zeta * \Psi_{v,i} + \epsilon_{v,i} \tag{3.2}$$

where i denotes a household and v denotes a village, $Y_{v,i}$ is the continuous outcome variable of interest, T_v is the EVD indicator, X_v is a vector of village-level controls including the geographic village characteristics mentioned above, $\Psi_{v,i}$ is a vector of household-level controls, and $\epsilon_{v,i}$ is the error term. Standard errors are clustered at the village-level. The main coefficient of interest β estimates the effect of village exposure to EVD on the outcome variable $Y_{v,i}$ conditional on controls.

The Probit specification for my binary outcome variables is

$$Pr(Y_{v,i} = 1) = \Theta(\alpha + \beta * T_v + \delta * X_v + \zeta * \Psi_{v,i} + \epsilon_{v,i}) \tag{3.3}$$

where i denotes a household and v denotes a village, $Y_{v,i}$ is the binary outcome variable of interest, T_v is the EVD indicator, X_v is a vector of village-level controls including the geographic village characteristics mentioned above, $\Psi_{v,i}$ is a vector of household-level controls, and $\epsilon_{v,i}$ is the error term. Standard errors are clustered at the village-level. Throughout the paper I report the marginal effect of the Ebola indicator on $Y_{v,i}$ derived from the Probit model evaluated at the means of the covariates.

3.4.2 Instrumental Variables (IV) and Recursive Bivariate Probit

Villages affected by EVD could still differ from villages not affected by EVD in terms of unobserved characteristics such as burial practices, access to information, or trust in health professionals. Therefore, I follow Gonzalez-Torres and Esposito (2020) and Maffioli (2021) and use a village's road distance to the EVD index case as an instrumental variable for the EVD indicator. As Marí Saéz et al. (2015) argue, the index case of the 2014-16 EVD epidemic most likely occurred when a 18 months old toddler came into contact with an infected fruit bat in the town of Meliandou, Guinea, close to Sierra Leone's eastern border. Thus, the location of the index case is plausibly unrelated to the outcome variables considered in this paper. While villages closer to the Guinean border might still differ from villages further away, I argue that my instrumental variable satisfies the exclusion restriction conditional on geographic village characteristics. As Fang et al. (2016) show, the transmission of EVD in Sierra Leone mainly followed human travel routes. Therefore, a village's road distance from the index case should have predictive power for the EVD indicator. In particular, I use a quadratic function of distance to the Ebola index case as an instrument for my Ebola indicator because a quadratic functional form results in a stronger first stage in this context.[11] Overall, my IV strategy captures the idea that between two villages which are similar in terms of remoteness, distance to the nearest main road, and altitude, the one that is closer to the index case was more likely to be exposed to EVD in 2014-16 but was otherwise comparable to the one that is further away.

A complication arises from my endogenous regressor, the EVD indicator, being a binary variable and my instrumental variable, the village road distance to the index case, being a continuous variable and the fact that my endogenous regressor

[11] All results are robust to using a linear function, see section 3.5.

has a high share of zero values. For my continuous outcome variables, I therefore use the 3-step estimation method suggested in Wooldridge (2010) which has been used also in Maffioli (2021): First, I estimate the following Probit model for the EVD indicator

$$Pr(T_v = 1) = \Theta(\alpha_1 + \beta_1 * D_v + \gamma_1 * D_v^2 + \delta_1 * X_v + \epsilon_v) \tag{3.4}$$

where v denotes a village and T_v is the EVD indicator. Θ is the cumulative distribution function of the estimation function in which D_v is the distance of village v to the Ebola index case, X_v is a vector of village-level controls, and ϵ_v is the error term.

Second, I use the predicted values from step 1 as the instrumental variable in the first stage of a conventional 2SLS regression

$$T_v = \alpha_2 + \beta_2 * \hat{T}_v + \delta_2 * X_v + \epsilon_v \tag{3.5}$$

where v denotes a village, T_v is the EVD indicator, $\hat{T}_v$ is the predicted value from the first step, X_v is a vector of village-level controls, and ϵ_v is the error term.

Third, I estimate the second stage of the 2SLS regression using the predicted values from the second step

$$Y_{v,i} = \alpha_3 + \beta_3 * \hat{\hat{T}}_v + \delta_3 * X_v + \zeta_3 * \Psi_{v,i} + \epsilon_{v,i} \tag{3.6}$$

where i denotes a household and v denotes a village, $Y_{v,i}$ is the continuous outcome variable of interest, $\hat{\hat{T}}_v$ is the predicted value from the second step, X_v is a vector of village-level controls, $\Psi_{v,i}$ is a vector of household-level controls, and $\epsilon_{v,i}$ is the error term. The main coefficient of interest β_3 estimates the causal effect of village exposure to EVD on the continuous outcome variable $Y_{v,i}$. In order to account for both the estimation uncertainty from all three steps and the fact that the level of observation is the village in steps one and two whereas it is the household in step three, I bootstrap the standard errors with resampling clustered at the village-level (200 replications).

For my binary outcome variables, I estimate a recursive bivariate probit model that captures the same idea as the IV for my continuous outcome variables. The first equation models the binary outcome variable as a function of the EVD indicator and village-level controls and the second equation models the EVD indicator as a quadratic function of a village's road distance to the EVD index case and the same village-level controls as in the first equation. As the Ebola indicator in the second equation only varies at the village-level I also aggregate the outcome variables in the first equation at the village-level in order to be able to simultaneously estimate both equations. Formally, I estimate

$$Pr(Y_v = 1; T_v = 1) = \Phi(f_v^Y; f_v^T; \rho_v) \tag{3.7}$$

where v denotes a village, Y_v is the binary outcome variable of interest, T_v is the EVD indicator, and Φ is the cumulative distribution function of the bivariate standard normal distribution with

$$f_v^Y = (2 * Y_v - 1) * (\alpha_1 + \beta_1 * T_v + \delta_1 * X_v + \epsilon_v) \tag{3.8}$$

and

$$f_v^T = (2 * T_v - 1) * (\alpha_2 + \beta_2 * D_v + \gamma_2 * D_v^2 + \delta_2 * X_v + \epsilon_v) \tag{3.9}$$

and

$$\rho_v = (2 * Y_v - 1) * (2 * T_v - 1) * \rho \tag{3.10}$$

where X_v is a vector of village-level controls, D_v is the distance of village v to the EVD index case, ϵ_v is the error term, and ρ is the correlation coefficient of the bivariate standard normal distribution. I then proceed to estimate the average treatment effect of EVD exposure on the binary outcome variable as

$$ATE = E[\Theta(\alpha_1 + \beta_1 + \delta_1 * X_v) - \Theta(\alpha_1 + \delta_1 * X_v)] \quad (3.11)$$

where Θ is the cumulative standard normal distribution. I bootstrap the standard errors for the average treatment effect (200 replications).

In order to shed light on the relevance of the instrumental variable, Table 3.2 shows the results from estimating equation 3.4, i.e., the first step of the 3-step generated IV estimation procedure. The χ^2 statistic of the test of joint significance of the coefficients on D_v and D_v^2 is 13.21 (p-value < 0.01). Table 3.3 looks at predictors of the EVD indicator and the results are consistent with the exclusion restriction laid out in this section: The geographic variables have the expected sign and the remoteness measure significantly predicts the EVD indicator while the available household controls do not significantly predict the EVD indicator. As with any instrumental variable approach, I cannot entirely rule out that the Ebola indicator conditional on the instrumental variables and further controls can be predicted by unobserved household or village characteristics.

3.5 Results

Risk Perception and Knowledge Related to Covid-19

Table 3.4 shows the OLS/Probit results. The share of household respondents that believe they can contract Covid-19 is 13.5 pp. higher in villages that were affected by Ebola (significant at the 10% level). Table 3.5 shows the results from the recursive bivariate probit and IV models aimed at getting causal effects of Ebola exposure. The ATE of the Ebola indicator on whether household respondents believe they can contract Covid-19 is 0.50 (significant at the 1% level), quite a bit higher than the estimate of the marginal effect from the Probit model. This suggests that either there is measurement error in the Ebola indicator leading to attenuation bias in the Probit model or the Ebola indicator is indeed endogenous. As a robustness check, Table 3.6 shows the IV estimates also for the binary outcome variables. For the effect of the Ebola indicator on whether household respondents believe they can contract Covid-19, with 0.41 (significant at the 5% level) the IV estimate is similar to the estimate from the recursive bivariate probit model.

For household respondents' knowledge about the possibility of asymptomatic transmissions, the marginal effect from the Probit model is 0.12 (significant at the 10% level). The estimate from the recursive bivariate probit model in column 2 of Table 3.5 is somewhat larger at 0.32 (significant at the 5% level). The IV estimate in column 2 of Table 3.6 is qualitatively similar to that, although less precise.

Overall, these results suggest that people living in villages directly affected by Ebola perceive Covid-19 to be significantly riskier than those living in non-affected villages: They are more likely to believe that they can contract Covid-19 and more likely to know about the possibility of asymptomatic transmissions of Covid-19. While both of these measures likely capture differences in religious and spiritual beliefs, education, and media consumption habits, I show that risk perception about Covid-19 is affected by exposure to previous (health) crises such as the 2014-16 Ebola epidemic.

Trust in Health Professionals

The Probit marginal effect of the Ebola indicator on respondents' trust in health professionals when it comes to information about Covid-19 in column 3 of Table 3.4 is 0.18 (significant at the 5% level). Again the estimate from the recursive bivariate probit model in column 3 of Table 3.5 is larger with 0.44 (significant at the

1% level) and the IV estimate is qualitatively similar to the other two estimates, but less precise. This is related to the finding in Flückiger et al. (2019) that Ebola-affected areas in Sierra Leone exhibit greater trust in government authorities. The authors argue that the mechanism behind this finding is that affected areas value state control more highly. Similarly, it seems plausible that having been affected by a traumatic health-related crisis like the 2014-16 Ebola epidemic increases one's valuation of health care services and providers. More generally, these results point towards (health) crises leading to higher trust in public institutions and a higher demand for public goods. This can be rationalized by the fact that public coordination in general, and a strong public health system in particular, are especially valuable during an epidemic.

Covid-19 Prevention

Having established that exposure to Ebola increases the perceived risk of Covid-19 and the trust in health professionals for household respondents in affected villages, I now turn to the question whether this had an impact on two policy-relevant margins: Attitudes and behaviors related to the prevention of the spread of Covid-19 are crucial in the fight against Covid-19 and understanding their determinants is important for policy makers during the current pandemic and potential future health crises.

My first set of outcomes related to the prevention of Covid-19 covers wearing face masks: Firstly, I look at the number of occasions household respondents report wearing a face mask on. The OLS estimate of 0.34 in column 4 of Table 3.4 is significant at the 10% level and points towards an increase in the number of occasions on which respondents wear face masks of 0.30 standard deviations. The IV estimate in column 4 of Table 3.5 is somewhat large at 0.724 with a slightly smaller p.value of 0.11. Secondly, I consider the percentage of adults wearing a face mask at the local market as reported by household respondents. The OLS estimate in column 5 of Table 3.4 suggests that the share of adults wearing a face mask at the local market is 6 pp. higher in villages affected by Ebola compared to non-affected villages (significant at the 5% level). However, with the IV this becomes an imprecisely estimated zero (column 5, Table 3.5), suggesting a lack of power to detect a relatively small effect. A potential reason for the lack of a stronger effect on face mask usage, especially in light of the results on face mask distribution below, is the timing of the survey: The government initiated containment measures from March 2020, but the survey only took place between October and November 2020, at a time when most of the stricter containment measures had been lifted again and when the number of new daily cases were at an all-time low since the beginning of the pandemic.

My second set of outcomes is related to the willingness to pay for a Covid-19 vaccination. Firstly, I look at (the inverse hyperbolic sine transformation of) the willingness to pay for a vaccination against Covid-19. The OLS (column 6, Table 3.4) and the IV (column 6, Table 3.5) coefficients confirm the unconditional result of no differences between affected and non-affected villages in Table 3.1, but the zero coefficient is very imprecisely estimated both in the OLS and IV specifications. Secondly, the Probit estimate in column 7 of Table 3.4 suggests that households in villages affected by Ebola are 10 pp. more likely to have a willingness to pay for a Covid-19 vaccination above the 77th percentile of the overall distribution (significant at the 5% level). The estimate from the recursive bivariate probit model in column 7 of Table 3.5 is substantially higher at 0.50 (significant at the 1% level) while the IV estimate in column 7 of Table 3.6 lies in between the other two estimates at 0.26 (significant at the 5% level). Taken together, these results suggest that past exposure to EVD does not affect demand for Covid-19 vaccinations on aggregate, although there is evidence for past exposure to EVD to lead to polarization with more households having a very high willingness to pay for a Covid-19 vaccination. While the increase in the share of respondents with a high willingness to pay fits the evidence on increased risk perception and higher trust in health professionals in EVD-affected villages, it is unclear why this effect is offset by a reduction of other respondents' willingness to pay. A plausible explanation in line with the increase in the demand for public goods in EVD-affected villages could be that more households in those villages believe that vaccinations should be provided free of charge and therefore answer that they are unwilling to pay for them. The share of household respondents who would get a Covid-19 vaccination free of charge is not significantly different across households in affected vs. non-affected villages. It is worth mentioning, though, that this share is 0.76 in the whole sample which is in line with the estimates for other low- and middle-income countries found in Solís Arce et al. (2021).

Finally, I look at whether face masks were publicly distributed in a village. The probit marginal effect is 0.21 (significant at the 5% level) as shown in column 8 of Table 3.4. The estimate from the recursive bivariate probit model in column 8 of Table 3.5 is 0.44 (significant at the 1% level). The IV estimate in column 8 of Table 3.6 is 0.55 (significant at the 5% level). This constitutes strong evidence for Ebola-affected villages taking Covid-19 more serious and making sure that face masks are available for households. In addition to a higher risk perception of Covid-19 and higher trust in health professionals, a similar mechanism as in Flückiger et al. (2019) mentioned above might be at play here: People living in villages directly affected by the 2014-16 Ebola epidemic might value public goods more highly and therefore managed to take collective action to prevent the Covid-19 pandemic to have a similarly negative impact. This result also relates to studies by Christensen et al. (2021) and Casey et al. (2021) showing that interventions targeting community mobilization and local public good provision had positive effects on collective action and trust during the Ebola epidemic in Sierra Leone:

While these studies establish a positive link from public good provision to the management of (health) crises, I show that (health) crises themselves generate demand for public good provision which can in turn become beneficial again during future (health) crises.

How is it possible that EVD-affected villages were more likely to distribute face masks but there does not seem to be a substantial increase in the number of households in those villages wearing face masks? Firstly, the delay between the onset of containment measures and the time of the survey corresponds to a period when caution seems to have had decreased as Covid-19 cases had come down and people had become used to the pandemic. Secondly, there is credible evidence from other contexts that face mask distribution alone is insufficient for increasing usage: Jakubowski et al. (2022) show that freely distributing face masks in Uganda had little effect on usage and the findings by Egger et al. (2022) from Kenya suggest that more intensive measures accompanying free distribution of face masks might be required to lead to behavioral changes.

Utilization of Health Services

The last outcome I consider is a measure of health service utilization, namely whether or not households actively avoided visiting their local health clinic out of fear of contracting Covid-19. The probit marginal effect estimate for the EVD indicator on whether respondents say they did avoid visiting their local health clinic is 0.18 (significant at the 5% level) as shown in column 9 of Table 3.4. The estimate from the recursive bivariate probit model in column 9 of Table 3.5 is 0.43 (significant at the 1% level). The IV estimate is qualitatively similar but less precisely estimated (column 9, Table 3.6). These results point towards the avoidance of health services being an important indirect cost of increased caution in the light of Covid-19 induced by previous EVD exposure. While I am not able to assess whether avoiding health clinic visits is consistent with rational expectations of patients, i.e., whether there is a sufficiently high risk of contracting Covid-19 at health clinics that outweighs the risk from not receiving health services, it is clear that increased caution due to the Covid-19 pandemic does come at the cost of reducing health access. This has potentially important implications for the debate on whether to adopt stricter containment measures such as lockdowns (which lead to fewer Covid-19 infections) or more lenient containment measures (which are less disruptive to the economic lives of people but also their health-seeking behavior).

Robustness

I briefly discuss the robustness of the results to two changes to my empirical strategy: First, Table 3.7 reproduces the main results table controlling for a village's distance to the capital city, Freetown, instead of a village's remoteness. Although I outlined above the theoretical advantage of controlling for remoteness due to the particular geography and geometry of this context, the fact that Fang et al. (2016) mention - next to the index case in Meliandou - a second EVD transmission corridor emerging from Freetown during the 2014-16 Ebola epidemic and that other studies in similar contexts - such as the one by Maffioli (2021) in Liberia - control for the distance from the capital city justifies checking whether the results are sensitive to this choice. As Table 3.7 shows, the results are similar in magnitudes to Table 3.5, but slightly less precise for the outcomes in columns 2 and 7.

Second, Table 3.8 reproduces the main results table using as an instrumental variable a village's linear - instead of quadratic - distance to Meliandou. The results are very similar to those in Table 3.5.

3.6 Conclusion

Using novel survey data from Sierra Leonean villages I show that households in villages affected by the 2014-16 EVD epidemic perceive Covid-19 to be riskier and trust health professionals more than households in non-affected villages. This in turn translates into EVD-affected villages to be more likely to adopt a key containment measure, namely the public distribution of face masks. This highlights on the micro level how past exposure to a health crisis affects the response to Covid-19 and could partially explain why countries with recent experience of major disease outbreaks like SARS, MERS, or Ebola appear to have reacted more swiftly and decisively at the onset of the Covid-19 pandemic than other countries. This results also points out an interesting aspect about the dynamic of (health) crises where past exposure generates demand for public goods which in turn help to manage future crises in a better way. However, the increased caution of households in EVD-affected villages comes at the cost of reduced health access due to households avoiding the local health clinic out of fear of contracting Covid-19.

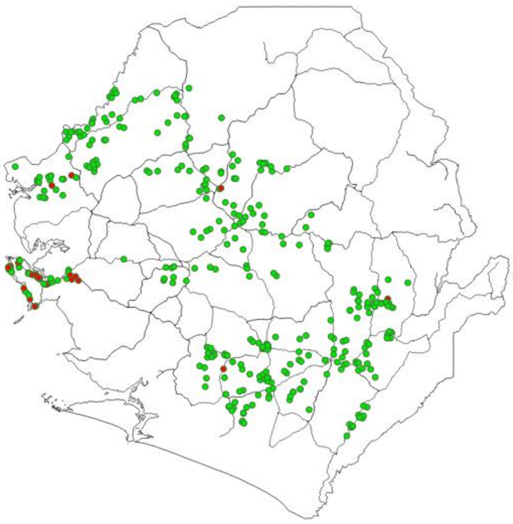

(a) Green dots: Non-EVD Villages Red dots: EVD Villages

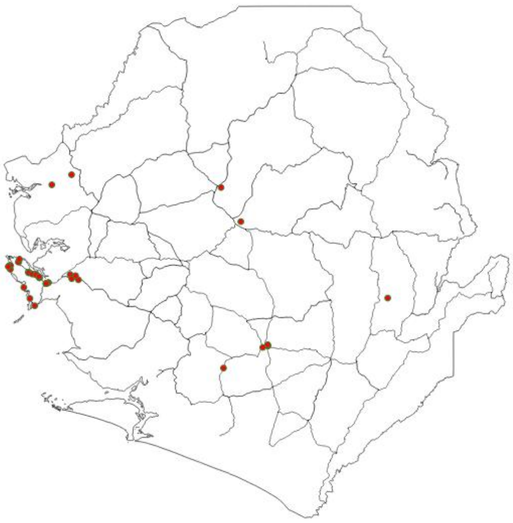

(b) Red dots: EVD Villages

Figure 3.1: All Villages and EVD Villages

Table 3.1: Outcome Variables

	(1) Obs.	(2) Mean	(3) Mean EVD = 0	(4) Mean EVD = 1	(5) (4) - (3)
Outcome Variable					
Can Contract Covid?	614	0.373	0.338	0.549	0.211***
Knows Asymp. Transm.	617	0.593	0.566	0.728	0.162***
Trust Health Prof.	617	0.421	0.379	0.631	0.252***
Occ. Wearing Face Mask	617	3.100	3.027	3.437	0.410***
Perc. Wearing Face Mask	614	85.487	84.483	90.486	6.002***
WTP Vaccination	612	4.222	4.283	3.909	−0.374
WTP Vaccination > 77th pct.	615	0.224	0.206	0.311	0.104**
Face Mask Distribution	391	0.442	0.416	0.675	0.259***
Avoided Health Clinic	617	0.297	0.272	0.417	0.145***

* p<0.1, ** p<0.05, *** p<0.01.

Table 3.2: First Stage

	(1) EVD Indicator = 1
Distance Meliandou	-0.271***
	(0.101)
Distance Meliandou Sq.	0.042***
	(0.014)
Remoteness	-0.001***
	(0.000)
Distance Mainroad	-0.004
	(0.003)
Altitude	-0.005
	(0.005)
Observations	392
Mean EVD Indicator	0.102
Chi2 statistic	13.208
P-value	0.001

* p<0.1, ** p<0.05, *** p<0.01.

Table 3.3: Predictors of Ebola Case in Village

	(1) EVD Indicator = 1
Distance Meliandou	-0.305***
	(0.099)
Distance Meliandou Sq.	0.045***
	(0.014)
Remoteness	-0.001***
	(0.000)
Distance Mainroad	-0.003
	(0.003)
Altitude	-0.005
	(0.005)
Age	-0.001
	(0.001)
Muslim	0.013
	(0.025)
Mende Ethnicity	-0.038
	(0.028)
Temne Ethnicity	-0.018
	(0.025)
Observations	392
Mean EVD Indicator	0.102

* p<0.1, ** p<0.05, *** p<0.01.

Table 3.4: Probit/OLS

	(1) Can Contract Covid?	(2) Knows Asymp. Transm.	(3) Trust Health Prof.	(4) Occ. Wearing Face Mask	(5) Perc. Wearing Face Mask	(6) WTP Vaccination	(7) WTP > 77th pct.	(8) Face Mask Distribution	(9) Avoided Health Clinic
EVD Indicator = 1	0.135*	0.120*	0.176**	0.341*	6.046**	-0.070	0.099*	0.212**	0.181**
	(0.070)	(0.062)	(0.072)	(0.186)	(2.506)	(0.704)	(0.058)	(0.088)	(0.084)
Observations	614	617	617	617	616	612	617	391	617
Mean Dep. Var	0.373	0.593	0.421	3.096	85.487	4.223	0.224	0.442	0.297
Estimation	Probit	Probit	Probit	OLS	OLS	OLS	Probit	Probit	Probit
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes
Geographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table 3.5: Biprobit/IV

	(1) Can Contract Covid?	(2) Knows Asymp. Transm.	(3) Trust Health Prof.	(4) Occ. Wearing Face Mask	(5) Perc. Wearing Face Mask	(6) WTP Vaccination	(7) WTP > 77th pct.	(8) Face Mask Distribution	(9) Avoided Health Clinic
EVD Indicator = 1	0.504***	0.320**	0.440***	0.724	-1.570	0.013	0.504***	0.440***	0.430***
	(0.115)	(0.155)	(0.110)	(0.456)	(7.498)	(1.639)	(0.129)	(0.121)	(0.147)
Observations	390	392	392	617	616	612	392	391	392
Mean Dep. Var	0.330	0.569	0.416	3.096	85.487	4.223	0.221	0.442	0.263
Estimation	Biprobit	Biprobit	Biprobit	IV	IV	IV	Biprobit	Biprobit	Biprobit
Individual Controls	No	No	No	Yes	Yes	Yes	No	No	No
Geographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* p<0.1, ** p<0.05, *** p<0.01.

Table 3.6: Only IV

	(1) Can Contract Covid?	(2) Knows Asymp. Transm.	(3) Trust Health Prof.	(4) Occ. Wearing Face Mask	(5) Perc. Wearing Face Mask	(6) WTP Vaccination	(7) WTP > 77th pct.	(8) Face Mask Distribution	(9) Avoided Health Clinic
EVD Indicator = 1	0.414**	0.205	0.165	0.724	-1.570	0.013	0.263*	0.546**	0.256
	(0.186)	(0.181)	(0.157)	(0.456)	(7.498)	(1.639)	(0.144)	(0.216)	(0.193)
Observations	614	617	617	617	616	612	617	616	617
Mean Dep. Var	0.373	0.593	0.421	3.096	85.487	4.223	0.224	0.472	0.297
Estimation	IV	IV	IV	IV	IV	IV	IV	IV	IV
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* p<0.1, ** p<0.05, *** p<0.01.

Table 3.7: Biprobit/IV Distance to Freetown

	(1) Can Contract Covid?	(2) Knows Asymp. Transm.	(3) Trust Health Prof.	(4) Occ. Wearing Face Mask	(5) Perc. Wearing Face Mask	(6) WTP Vaccination	(7) WTP > 77th pct.	(8) Face Mask Distribution	(9) Avoided Health Clinic
EVD Indicator = 1	0.448**	0.239	0.359*	1.096	7.362	-0.270	0.401	0.455**	0.544***
	(0.180)	(0.297)	(0.202)	(0.816)	(13.955)	(3.032)	(0.259)	(0.185)	(0.129)
Observations	390	392	392	617	616	612	392	391	392
Mean Dep. Var	0.330	0.569	0.416	3.096	85.487	4.223	0.221	0.442	0.263
Estimation	Biprobit	Biprobit	Biprobit	IV	IV	IV	Biprobit	Biprobit	Biprobit
Individual Controls	No	No	No	Yes	Yes	Yes	No	No	No
Geographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* p<0.1, ** p<0.05, *** p<0.01.

Table 3.8: Biprobit/IV Linear Distance

	(1) Can Contract Covid?	(2) Knows Asymp. Transm.	(3) Trust Health Prof.	(4) Occ. Wearing Face Mask	(5) Perc. Wearing Face Mask	(6) WTP Vaccination	(7) WTP > 77th pct.	(8) Face Mask Distribution	(9) Avoided Health Clinic
EVD Indicator = 1	0.512***	0.348**	0.442***	0.939*	-0.406	0.850	0.513***	0.457***	0.464***
	(0.118)	(0.158)	(0.125)	(0.509)	(7.505)	(1.922)	(0.138)	(0.128)	(0.148)
Observations	390	392	392	617	616	612	392	391	392
Mean Dep. Var	0.330	0.569	0.416	3.096	85.487	4.223	0.221	0.442	0.263
Estimation	Biprobit	Biprobit	Biprobit	IV	IV	IV	Biprobit	Biprobit	Biprobit
Individual Controls	No	No	No	Yes	Yes	Yes	No	No	No
Geographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Bibliography

Abelson, R. (2021). 'i just cry all the time': Non-covid patients despair over delayed care. *The New York Times*.

Acharya, A., Blackwell, M., and Sen, M. (2016). Explaining causal findings without bias: Detecting and assessing direct effects. *The American Political Science Review*, 110(3):512–529.

Adhvaryu, A., Gauthier, J.-F., Nyshadham, A., and Tamayo, J. (2020). Absenteeism, productivity, and relational contracts inside the firm. Working Paper.

Adhvaryu, A., Nyshadham, A., and Tamayo, J. (2022). An anatomy of performance monitoring. Working Paper.

Aggarwal, S., Giera, B., Jeong, D., Robinson, J., and Spearot, A. (2021). Market access, trade costs, and technology adoption: Evidence from northern tanzania. Working Paper.

Akerlof, R., Ashraf, A., Macchiavello, R., and Rabbani, A. (2020). Layoffs and productivity at a bangladeshi sweater factory. Working Paper.

Alchian, A. A. and Demsetz, H. (1972). Production, information costs, and economic organization. *American Economic Review*, 62(5):777–795.

Alsan, M. (2015). The effect of the tsetse fly on african development. *American Economic Review*, 105(1):382–410.

Aman-Rana, S. (2020). In self interest? meritocracy in a bureaucracy. Working Paper.

Ashraf, N. and Bandiera, O. (2018). Social incentives in organizations. *Annual Review of Economics*, 10:439–463.

Ashraf, N., Bandiera, O., and Jack, B. K. (2014). No margin, no mission? a field experiment on incentives for public service delivery. *Journal of Public Economics*, 120:1–17.

Asiedu, E., Karlan, D., Lambon-Quayefio, M. P., and Udry, C. R. (2021). A call for structured ethics appendices in social science papers. Working Paper.

Babcock, P., Bedard, K., Charness, G., Hartman, J., and Royer, H. (2015). Letting down the team? social effects of team incentives. *Journal of the European Economic Association*, 13(5):841–870.

Bandiera, O., Barankay, I., and Rasul, I. (2007). Incentives for managers and inequality among workers: Evidence from a firm level experiment. *The Quarterly Journal of Economics*, 122(2):729–773.

Bandiera, O., Best, M. C., Khan, A. Q., and Prat, A. (2021). The allocation of authority in organizations: A field experiment with bureaucrats. *The Quarterly Journal of Economics*, 136(4):2195–2242.

Bandyopadhyay, S., Baticulon, R. E., Kadhum, M., Alser, M., Ojuka, D. K., Badereddin, Y., Kamath, A., Parepalli, S. A., Brown, G., Iharchane, S., et al. (2020). Infection and mortality of healthcare workers worldwide from covid-19: A systematic review. *BMJ Global Health*, 5(12).

Banerjee, A., Duflo, E., Imbert, C., Mathew, S., and Pande, R. (2020). E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in india. *American Economic Journal: Applied Economics*, 12(4):39–72.

Banerjee, A. V., Duflo, E., and Glennerster, R. (2008). Putting a band-aid on a corpse: Incentives for nurses in the indian public health care system. *Journal of the European Economic Association*, 6(2-3):487–500.

Behrman, J. R., Parker, S. W., Todd, P. E., and Wolpin, K. I. (2015). Aligning learning incentives of students and teachers: Results from a social experiment in mexican high schools. *Journal of Political Economy*, 123(2):325–364.

Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678.

Benson, A., Li, D., and Shue, K. (2019). Promotions and the peter principle. *The Quarterly Journal of Economics*, 134(4):2085–2134.

Benson, A., Li, D., and Shue, K. (2021). Potential and the gender promotion gap. Working Paper.

Bertrand, M. (2009). Ceos. *Annual Review of Economics*, 1(1):121–150.

Bertrand, M., Burgess, R., Chawla, A., and Xu, G. (2020). The glittering prizes: Career incentives and bureaucrat performance. *The Review of Economic Studies*, 87(2):626–655.

Besley, T. and Ghatak, M. (2005). Competition and incentives with motivated agents. *American Economic Review*, 95(3):616–636.

Besley, T. J., Burgess, R., Khan, A., and Xu, G. (2021). Bureaucracy and development. Working Paper.

Bose, G. and Lang, K. (2017). Monitoring for Worker Quality. *Journal of Labor Economics*, 35(3):755–785.

Boudreau, K. J., Lakhani, K. R., and Menietti, M. E. (2016). Performance responses to competition across skill-levels in rank order tournaments: Field evidence and implications for tournament design. *RAND Journal of Economics*, 1(47):140—165.

Breza, E., Kaur, S., and Shamdasani, Y. (2018). The morale effects of pay inequality. *The Quarterly Journal of Economics*, 133(2):611–663.

Bubb, R., Kaur, S., and Mullainathan, S. (2018). The limits of neighborly exchange. Working Paper.

Callen, M., Gulzar, S., Hasanain, A., Khan, M. Y., and Rezaee, A. (2020). Data and policy decisions: Experimental evidence from pakistan. *Journal of Development Economics*, 146:102523.

Campbell, J., Dussault, G., Buchan, J., Pozo-Martin, F., Guerra Arias, M., Leone, C., Siyam, A., and Cometto, G. (2013). A universal truth: No health without a workforce. Technical report, Global Health Workforce Alliance and World Health Organization.

Cantor, J. H., Sood, N., Bravata, D., Pera, M., and Whaley, C. M. (2020). The impact of the covid-19 pandemic and policy response on health care utilization: Evidence from county-level medical claims and cellphone data. Working Paper.

Card, D., Mas, A., Moretti, E., and Saez, E. (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review*, 102(6):2981–3003.

Casaburi, L. and Macchiavello, R. (2019). Demand and supply of infrequent payments as a commitment device: Evidence from kenya. *American Economic Review*, 109(2):523–55.

Caselli, F., Grigoli, F., and Sandri, D. (2021). Protecting lives and livelihoods with early and tight lockdowns. *The BE Journal of Macroeconomics*.

Casey, K., Glennerster, R., Miguel, E., and Voors, M. J. (2021). Long run effects of aid: Forecasts and evidence from sierra leone. Technical report, National Bureau of Economic Research.

Castilla, E. (2008). Gender, race, and meritocracy in organizational careers. *American Journal of Sociology*, 113:1479–526.

Christensen, D., Dube, O., Haushofer, J., Siddiqi, B., and Voors, M. (2021). Building resilient health systems: Experimental evidence from sierra leone and the 2014 ebola outbreak. *The Quarterly Journal of Economics*, 136(2):1145–1198.

Chua, A. Q., Al Knawy, B., Grant, B., Legido-Quigley, H., Lee, W.-C., Leung, G. M., Looi, M.-K., and Maurer-Stroh, S. (2021). How the lessons of previous epidemics helped successful countries fight covid-19. *BMJ*, 372.

Cilliers, J., Kasirye, I., Leaver, C., Serneels, P., and Zeitlin, A. (2018). Pay for locally monitored performance? a welfare analysis for teacher attendance in ugandan primary schools. *Journal of Public Economics*, 167:69–90.

Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16):386–405.

Cohn, A., Fehr, E., Herrmann, B., and Schneider, F. (2014). Social comparison and effort provision: Evidence from a field experiment. *Journal of the European Economic Association*, 12(4):877–898.

Coviello, D., Deserranno, E., and Persico, N. (2021). Counterproductive worker behavior after a pay cut. *Journal of the European Economic Association*, (Forthcoming).

Cullen, Z. B. and Perez-Truglia, R. (2019). The old boys' club: Schmoozing and the gender gap. Working Paper.

Cullen, Z. B. and Perez-Truglia, R. (2021). How much does your boss make? the effects of salary comparisons. *Journal of Political Economy*, (Forthcoming).

Dal Bó, E., Finan, F., Li, N., and Schechter, L. (2021). Information technology and government decentralization: Experimental evidence from paraguay. *Econometrica*, 89(2):677–701.

Darmstadt, G. L., Choi, Y., Arifeen, S. E., Bari, S., Rahman, S. M., Mannan, I., Seraji, H. R., Winch, P. J., Saha, S. K., Ahmed, A. N. U., et al. (2010). Evaluation of a cluster-randomized controlled trial of a package of community-based maternal and newborn interventions in mirzapur, bangladesh. *PloS One*, 5(3):e9696.

de Janvry, A., He, G., Sadoulet, E., Wang, S., and Zhang, Q. (2021). Performance evaluation, influence activities, and bureaucratic work behavior: Evidence from china. Working Paper.

De Ree, J., Muralidharan, K., Pradhan, M., and Rogers, H. (2018). Double for nothing? experimental evidence on an unconditional teacher salary increase in indonesia. *The Quarterly Journal of Economics*, 133(2):993–1039.

Deserranno, E., Caria, S., Kastrau, P., and León-Ciliotta, G. (2022a). The allocation of incentives in multi-layered organizations. Working Paper.

Deserranno, E., Kastrau, P., and León-Ciliotta, G. (2022b). Promotions and productivity: The role of meritocracy and pay progression in the public sector. Working Paper.

Deserranno, E., Nansamba, A., and Qian, N. (2020). Aid crowd-out: The effect of ngos on government-provided services. Working Paper.

Dodge, E., Neggers, Y., Pande, R., and Moore, C. T. (2021). Updating the state: Information acquisition costs and public benefit delivery. Working Paper.

Duflo, E., Hanna, R., and Ryan, S. P. (2012). Incentives work: Getting teachers to come to school. *American Economic Review*, 102(4):1241–78.

Egger, D., Jakubowski, A., Nekesa, C., and Walker, M. (2022). Mask up! testing strategies to increase mask usage in kenya. *medRxiv*.

Evans, D. K., Goldstein, M., and Popova, A. (2015). Health-care worker mortality and the legacy of the ebola epidemic. *The Lancet Global Health*, 3(8):e439–e440.

Fang, L.-Q., Yang, Y., Jiang, J.-F., Yao, H.-W., Kargbo, D., Li, X.-L., Jiang, B.-G., Kargbo, B., Tong, Y.-G., Wang, Y.-W., et al. (2016). Transmission dynamics of ebola virus disease and intervention effectiveness in sierra leone. *Proceedings of the National Academy of Sciences*, 113(16):4488–4493.

Fehr, E. and Goette, L. (2007). Do workers work more if wages are high? evidence from a randomized field experiment. *American Economic Review*, 97(1):298–317.

Finan, F., Olken, B. A., and Pande, R. (2017). The personnel economics of the developing state. In *Handbook of Economic Field Experiments*, volume 2, pages 467–514. Elsevier, North-Holland.

Flückiger, M., Ludwig, M., and Sina Önder, A. (2019). Ebola and state legitimacy. *The Economic Journal*, 129(621):2064–2089.

Fotiou, A. and Lagerborg, A. (2021). Smart containment: Lessons from countries with past experience. Working Paper.

Franke, J. (2012). Affirmative action in contest games. *European Journal of Political Economy*, 28(1):105–118.

Franke, J., Kanzow, C., Leininger, W., and Schwartz, A. (2013). Effort maximization in asymmetric contest games with heterogeneous contestants. *Economic Theory*, 52(2):589–630.

Friebel, G., Heinz, M., Krueger, M., and Zubanov, N. (2017). Team incentives and performance: Evidence from a retail chain. *American Economic Review*, 107(8):2168–2203.

Frydman, C. and Jenter, D. (2010). Ceo compensation. *Annual Review of Financial Economics*, 2(1):75–102.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904.

Gholami, M., Fawad, I., Shadan, S., Rowaiee, R., Ghanem, H. A., Omer, A., and Ho, H. S. B. (2021). Covid-19 and healthcare workers: A systematic review and metaanalysis. *International Journal of Infectious Diseases*.

Gibbons, R. (2005). Four formal(izable) theories of the firm? *Journal of Economic Behavior & Organization*, 58(2):200–245.

Gibbons, R. and Murphy, K. J. (1992). Optimal incentive contracts in the presence of career concerns: Theory and evidence. *Journal of Political Economy*, 100(3):468–505.

Gibbons, R. and Waldman, M. (1999a). A Theory of Wage and Promotion Dynamics Inside Firms. *The Quarterly Journal of Economics*, 114(4):1321–1358.

Gibbons, R. and Waldman, M. (1999b). Careers in Organizations: Theory and Evidence. In *Handbook of Labor Economics*, volume 3, chapter 36, pages 2373–2437. Elsevier.

Glaziou, P. (2020). Predicted impact of the covid-19 pandemic on global tuberculosis deaths in 2020. Working Paper.

Glewwe, P., Ilias, N., and Kremer, M. (2010). Teacher incentives. *American Economic Journal: Applied Economics*, 2(3):205–27.

Gonzalez-Torres, A. and Esposito, E. (2020). Epidemics and conflict: Evidence from the ebola outbreak in western africa. Working Paper.

Harris, M. and Holmstrom, B. (1982). A Theory of Wage Dynamics. *The Review of Economic Studies*, 49(3):315–333.

Hines, J. R. and Thaler, R. H. (1995). The flypaper effect. *Journal of Economic Perspectives*, 9(4):217–226.

Itoh, H. (1991). Incentives to help in multi-agent situations. *Econometrica*, 59(3):611–636.

Jain, R. and Dupas, P. (2021). The effects of indias covid-19 lockdown on critical non-covid health care and outcomes: Evidence from a retrospective cohort analysis of dialysis patients. Working Paper.

Jakubowski, A., Egger, D., Mulebeke, R., Akankwasa, P., Muruta, A., Kiwanuka, N., and Wanyenze, R. (2022). Evaluation of a national program to distribute free masks for covid-19 prevention in uganda: Evidence from mbale district.

Karachiwalla, N. and Park, A. (2017). Promotion incentives in the public sector: Evidence from chinese schools. *Journal of Public Economics*, 146:109–128.

Karing, A. (2021). Strengthening state capacity in health with administrative data in sierra leone. Working Paper.

Ke, R., Li, J., and Powell, M. (2018). Managing careers in organizations. *Journal of Labor Economics*, 36(1):197–252.

Kelley, E. M., Lane, G., and Schönholzer, D. (2021). Monitoring in small firms: Experimental evidence from kenyan public transit. Working Paper.

Khan, A. Q., Khwaja, A. I., and Olken, B. A. (2016). Tax farming redux: Experimental evidence on performance pay for tax collectors. *The Quarterly Journal of Economics*, 131(1):219–271.

Khan, A. Q., Khwaja, A. I., and Olken, B. A. (2019). Making moves matter: Experimental evidence on incentivizing bureaucrats through performance-based postings. *American Economic Review*, 109(1):237–70.

Kunze, A. and Miller, A. R. (2017). Women Helping Women? Evidence from Private Sector Data on Workplace Hierarchies. *The Review of Economics and Statistics*, 99(5):769–775.

Lafontaine, F. and Slade, M. (2007). Vertical integration and firm boundaries: The evidence. *Journal of Economic Literature*, 45(3):629–685.

Lameke, A. A., Malukisa, A. J., Sánchez de la Sierra, R., and Titeca, K. (2020). Corruption (with a hierarchy). Working Paper.

Lazear, E. P. (2000). Performance pay and productivity. *American Economic Review*, 90(5):1346–1361.

Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864.

Lee, R. S., Whinston, M. D., and Yurukoglu, A. (2021). Chapter 9 - structural empirical analysis of contracting in vertical markets. *Handbook of Industrial Organization*, 4(1):673–742.

León-Ciliotta, G., Zejcirovic, D., and Fernandez, F. (2022). Policy-making, trust and the demand for public services: Evidence from a nationwide family planning program. Working Paper.

Li, X. (2020). The costs of workplace favoritism: Evidence from promotions in chinese high schools. Working Paper.

Lowes, S. R. and Montero, E. (2021). The legacy of colonial medicine in central africa. *American Economic Review*, 111(4):1284–1314.

Luo, R., Miller, G., Rozelle, S., Sylvia, S., and Vera-Hernández, M. (2019). Can Bureaucrats Really Be Paid Like CEOs? Substitution Between Incentives and Resources Among School Administrators in China. *Journal of the European Economic Association*, 18(1):165–201.

Macchiavello, R., Menzel, A., Rabbani, A., and Woodruff, C. (2020). Challenges of change: An experiment promoting women to managerial roles in the bangladeshi garment sector. Working Paper.

Macchiavello, R. and Morjaria, A. (2015). The value of relationships: Evidence from a supply shock to kenyan rose exports. *American Economic Review*, 105(9):2911–45.

Macchiavello, R. and Morjaria, A. (2021). Competition and relational contracts in the rwanda coffee chain. *The Quarterly Journal of Economics*, 136(2):1089–1143.

Maffioli, E. M. (2021). The political economy of health epidemics: Evidence from the ebola outbreak. *Journal of Development Economics*, 151:102651.

Marí Saéz, A., Weiss, S., Nowak, K., Lapeyre, V., Zimmermann, F., Düx, A., Kühl, H. S., Kaba, M., Regnaut, S., Merkel, K., et al. (2015). Investigating the zoonotic origin of the west african ebola epidemic. *EMBO Molecular Medicine*, 7(1):17–23.

Martinez-Bravo, M. and Stegmann, A. (2022). In vaccines we trust? the effects of the cia's vaccine ruse on immunization in pakistan. *Journal of the European Economic Association*.

Mas, A. (2017). Does transparency lead to pay compression? *Journal of Political Economy*, 125(5):1683–1721.

McKinsey (2015). Women in the workplace. Technical Report.

McMillan, J. and Woodruff, C. (1999). Interfirm relationships and informal credit in vietnam. *The Quarterly Journal of Economics*, 114(4):1285–1320.

McQuaid, C. F., McCreesh, N., Read, J. M., Sumner, T., Houben, R. M., White, R. G., Harris, R. C., et al. (2020). The potential impact of covid-19-related disruption on tuberculosis burden. *European Respiratory Journal*, 56(2).

Miguel, E. and Mobarak, A. M. (2021). The economics of the covid-19 pandemic in poor countries. *Annual Review of Economics*, 14. Forthcoming.

Mishel, L. and Wolfe, J. (2019). Ceo compensation has grown 940% since 1978. *Economic Policy Institute*, 14.

Muralidharan, K., Niehaus, P., Sukhtankar, S., and Weaver, J. (2021). Improving last-mile service delivery using phone-based monitoring. *American Economic Journal: Applied Economics*, 13(2):52–82.

Muralidharan, K., Romero, M., and Wüthrich, K. (2020). Factorial Designs, Model Selection, and (Incorrect) Inference in Randomized Experiments. Working Paper.

Muralidharan, K. and Sundararaman, V. (2011). Teacher performance pay: Experimental evidence from india. *Journal of Political Economy*, 119(1):39–77.

Ndubuisi-Obi, I., Ran, Z., Li, Y., Navalkha, C. A., Williams, S., and Tsai, L. (2021). Using public-private data to understand compliance with mobility restrictions in sierra leone. In *Urban Informatics and Future Cities*, pages 33–49. Springer.

Nieddu, M. G. and Pandolfi, L. (2018). The Effectiveness of Promotion Incentives for Public Employees: Evidence from Italian Academia. Working Paper.

Nti, K. O. (2004). Maximum efforts in contests with asymmetric valuations. *European Journal of Political Economy*, 20(4):1059–1066.

Nyqvist, M. B., Guariso, A., Svensson, J., and Yanagizawa-Drott, D. (2019). Reducing child mortality in the last mile: Experimental evidence on community health promoters in uganda. *American Economic Journal: Applied Economics*, 11(3):155–92.

Peter, L. J., Hull, R., et al. (1969). *The Peter Principle*, volume 4. Souvenir Press London.

Rasul, I. and Rogger, D. (2018). Management of bureaucrats and public service delivery: Evidence from the nigerian civil service. *The Economic Journal*, 128(608):413–446.

Ray, D., Baland, J.-M., and Dagnelie, O. (2007). Inequality and inefficiency in joint projects. *The Economic Journal*, 117(522):922–935.

Rosen, S. (1986). Prizes and Incentives in Elimination Tournaments. *American Economic Review*, 76(4):701–715.

Sahling, J. H. M., Schuster, C., and Mikkelsen, K. S. (2018). Civil service management in developing countries: What works? evidence from a survey with 23,000 civil servants in africa, asia, eastern europe and latin america. Technical report.

Shapira, G., Kalisa, I., Condo, J., Humuza, J., Mugeni, C., Nkunda, D., and Walldorf, J. (2018). Going beyond incentivizing formal health providers: Evidence from the rwanda community performance-based financing program. *Health Economics*, 27(12):2087–2106.

Shepherd, G. (2003). Civil service reform in developing countries: Why is it going badly? Technical Report.

Sherrard-Smith, E., Hogan, A. B., Hamlet, A., Watson, O. J., Whittaker, C., Winskill, P., Ali, F., Mohammad, A. B., Uhomoibhi, P., Maikore, I., et al. (2020). The potential public health consequences of covid-19 on malaria in africa. *Nature Medicine*, 26(9):1411–1416.

Siegel, R. (2010). Asymmetric contests with conditional investments. *American Economic Review*, 100(5):2230–60.

Siegel, R. (2014). Asymmetric contests with head starts and nonmonotonic costs. *American Economic Journal: Microeconomics*, 6(3):59–105.

Soeters, R. and Griffiths, F. (2003). Improving government health services through contract management: A case from cambodia. *Health Policy and Planning*, 18(1):74–83.

Solís Arce, J. S., Humphreys, M., Meriggi, N., Voors, M., and Yam, E. (2020). Sierra leone locked down early to contain covid, but at a high price. *LSE Blog*.

Solís Arce, J. S., Warren, S. S., Meriggi, N. F., Scacco, A., McMurry, N., Voors, M., Syunyaev, G., Malik, A. A., Aboutajdine, S., Adeojo, O., et al. (2021). Covid-19 vaccine acceptance and hesitancy in low-and middle-income countries. *Nature medicine*, 27(8):1385–1394.

Tirole, J. (1986). Hierarchies and bureaucracies: On the role of collusion in organizations. *Journal of Law, Economics, and Organization*, 2(2):181–214.

Tirole, J. (1992). Collusion and the theory of organizations. In Laffont, J.-J., editor, *Advances in Economic Theory*, volume 2, chapter 3. Cambridge University Press.

Voth, J. and Xu, G. (2021). Discretion and destruction: Promotions, performance, and patronage in the royal navy. Working Paper.

Wade, R. (1985). The market for public office: Why the indian state is not better at development. *World Development*, 13(4):467–497.

Waldman, M. (1984). Job assignments, signalling, and efficiency. *The RAND Journal of Economics*, 15(2):255–267.

Weaver, J. (2021). Jobs for sale: Corruption and misallocation in hiring. *American Economic Review*, 111(10):3093–3122.

Wilson, J. Q. (1989). *Bureaucracy: What government agencies do and why they do it.* Basic Books New York.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data.* MIT Press.

World Bank (2014). Pay flexibility and government performance: A multicountry study. Technical report.

World Bank (2016). Public employment and governance in middle east and north africa. World Bank Other Operational Studies 25181.

World Bank (2018). Merit, discrimination, and democratization: An analysis of promotion patterns in indonesia's civil service. Technical report.

World Bank Group (2009). Africa infrastructure country diagnostics (aicd). Technical report, World Bank Group.

World Health Organization (2016). *Health in 2015: From MDGs, Millennium Development Goals to SDGs, Sustainable Development Goals.* Geneva: WHO Press, 2016.

World Health Organization (2017). Global health observatory.

World Health Organization (2021). Second round of the national pulse survey on continuity of essential health services during the covid-19 pandemic: Interim report. Technical report, World Health Organization.

Xu, G. (2018). The costs of patronage: Evidence from the british empire. *American Economic Review*, 108(11):3170–98.

Xu, H. and Adhvaryu, A. (2020). Education and the meritocratic recruitment of bureaucrats. Working Paper.