AER: Insights 2025, 7(1): 71–89
https://doi.org/10.1257/aeri.20230594

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

By Erika Deserranno, Philipp Kastrau, and Gianmarco León-Ciliotta*

We study promotion incentives in the public sector. In collaboration with Sierra Leone's Ministry of Health, we introduce exogenous variation in the meritocratic nature of promotions from health worker to supervisor positions and in health workers' perceptions of pay progression upon promotion. Ten months later, our findings reveal that meritocracy leads to a 22 percent increase in health workers' productivity. Greater perceived pay progression in a meritocratic system boosts productivity by 23 percent, whereas in a less meritocratic system, it decreases productivity by 27 percent. We show that this reduction is consistent with a negative morale effect. (*JEL* C93, H51, I11, J24, J31, M51, O15)

Public sector organizations often refrain from directly linking promotions to performance, instead opting for rigid criteria based on seniority or discretionary systems susceptible to favoritism (Meyer-Sahling, Schuster, and Mikkelsen 2018; Shepherd 2003). Does this trend stem from a lack of receptiveness of public sector workers to merit-based promotions? Leveraging a field experiment with a large public sector organization, we show that public sector workers are indeed responsive to merit-based promotions, uncovering considerable potential to enhance public

*Deserranno: Bocconi University and Northwestern University (email: erika.deserrannoeri@unibocconi.it); Kastrau: Open Philanthropy (email: p.kastrau@gmail.com); León-Ciliotta: Universitat Pompeu Fabra, Barcelona School of Economics and IPEG (email: gianmarco.leon@upf.edu). Rohini Pande was the coeditor for this article. We thank the editor for her constructive feedback and insightful comments. Oriana Bandiera, Zoe Cullen, Ernesto Dal BóGeorgy Egorov, Tim Fedderson, Frederico Finan, Rachel Glennerster, Jonas Hjort, Mitch Hoffman, Supreet Kaur, Rocco Macchiavello, Michael Powell, Imran Rasul, Sandro Shelegia, Alessandro Tarozzi, Edoardo Teso, and seminar participants at BREAD, NBER-SI Personnel, NBER-SI Development, CEPR-Labor, SIOE, VDEV/CEPR/BREAD Seminar Series, STEG workshop, USC, UCSD, CAF, University of Boston, Uppsala University, University of Michigan, Northwestern University, for helpful comments. Theo Durandard, Malachy Gavan, Andre Cazor Katz, Mustapha Adedolapo Kokumo, Paulina Koenig, and Eduardo Amaro Furlanetto provided outstanding research assistance. This project would not have been possible without support from the staff at the Sierra Leonean Ministry of Health and Sanitation, the CHW Hub and the International Rescue Committee. We are extremely grateful to Niccolo Meriggi and Rachel Glennerster for their support and feedback and to Luz Azlor for her tireless work in the field. Financial support for this project was provided by UK aid from the UK government (through the Economic Development and Institutions initiative, grant PO_26429), the Rockefeller Foundation (through the International Growth Center, grant PIR_20180220), J-PAL's Governance Initiative (GR-0961), and the Weiss Family Fund (NWU-004). León-Ciliotta thanks the Spanish Ministry of Economy and Competitiveness (through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S) and grants (RYC2017-23172 funded by MICIU/AEI /10.13039/501100011033 and FSE invierte en tu futuro) and the Fundación Ramón Areces. IRB has been approved by the Sierra Leone Ethics and Scientific Review Committee and Universitat Pompeu Fabra (CIREP Approval 107). There are no conflicts of interest. This study was registered in the AEA RCT Registry as AEARCTR-0003993 (Deserranno, Kastrau, and Leon 2019).

[†]Go to https://doi.org/10.1257/aeri.20230594 to visit the article page for additional materials and author disclosure statement(s).

service delivery. We argue that the implementation of a more meritocratic promotion system is particularly important in settings with high salary progression, where higher-level officials earn substantially more than their lower-level counterparts. In such settings, the absence of a merit-based system may be perceived as unfair and demotivate employees at the lower tiers. Overall, this paper underscores the importance for organizations to strike a balance between offering large promotion rewards and adhering to performance-based promotion rules to increase the productivity of lower-tier workers.

We design a large field experiment in collaboration with the Ministry of Health in Sierra Leone in nearly 400 health units across the country. Each health unit comprises an average of eight Community Health Workers (CHWs), who provide basic health services to households in their community, and one Peer Supervisor (PS), who advises and monitors the CHWs. Whenever a PS position becomes vacant, one CHW in that health unit is offered the job. Prior to our experiment, promotion decisions were at the discretion of the local health authority and were subject favoritism. The experimental design creates random variation in the *actual* promotion rule by transitioning half of the health units to a new promotion system, which promotes CHWs based on the quantity of health services provided and their quality (visit length). This change significantly alters workers' perceptions of meritocracy: under the original promotion system, only half of the CHWs viewed the system as meritocratic, compared to 80 percent in the new system. This random variation in the promotion rule is cross-randomized with variation in *perceived* pay progression. Given the low initial awareness of PS pay, where over two-thirds of CHWs were unaware of PS pay, we informed a random subset of CHWs about the actual PS compensation, prompting them to adjust their beliefs toward the truth and influencing their perceived pay progression.

We utilize our 2×2 research design to evaluate how workers' motivation to ascend the organization's ladder—and consequently, their productivity—is influenced by (i) the extent to which promotion rules are performance based (meritocracy), (ii) the size of the expected rewards from promotions (perceived pay progression), and (iii) the interaction between the two. The effects are evaluated by gathering data on worker performance from a random sample of households ten months after the implementation of the new promotion rule. Performance is assessed broadly, encompassing measures incentivized under the meritocratic promotion system, such as the number and length of visits, as well as other aspects like household targeting, which are not directly incentivized.

We present two main sets of results. First, we show that a more meritocratic promotion rule increases the number of visits provided by the average worker by 22 percent. The productivity boost is stronger for high-performing workers with better chances of promotion in a meritocratic regime. The effect is stronger for workers who expect the value of the promotion to be large—that is, those who are likely to see the PS retire soon and those who perceive pay progression to be steep at baseline. Importantly, the increase in visits does not compromise service quality: there is no reduction in visit length or worse household targeting.

Second, we show that steeper perceived pay progression has diverging effects: in the new meritocratic system, it raises the number of visits by 23 percent; in the old (less meritocratic) system, it reduces visits by 27 percent. This indicates that

steeper perceived pay progression motivates the workers to climb the organization's ladder and prompts an increase in effort only when promotions are performance based. When promotions are not performance based, steeper perceived pay progression instead demotivates workers. We provide suggestive evidence that this is because workers perceive the large pay gap as unfair if the system does not reward highly productive workers, leading to a negative morale effect that decreases their motivation.

This paper contributes to different strands of the literature. First, it adds to recent literature on the personnel economics of the state. This literature often attributes low productivity of public sector workers to weak incentives, insufficient monitoring, or inadequate selection (Finan, Olken, and Pande 2017; Deserranno 2019; Xu et al. 2023). Our findings indicate that in the public sector of developing countries, the fact that pay progression is often steeper than in higher-income countries and that promotions are less meritocratic (see Supplemental Appendix Figure A.1) may constrain these governments' capacity to provide high-quality public services. Our results also indicate that, to the extent that meritocratic promotions are infeasible, rules-based systems may dominate discretion (or nonmeritocratic systems) due to morale effects.

Second, this paper extends the literature on promotion incentives, which has largely been theoretical (Lazear and Rosen 1981; Rosen 1986; Gibbons and Murphy 1992; Gibbons and Waldman 1999a, b). Recent empirical studies have observed the positive effects of increased upward mobility on worker performance when a new senior position becomes "attainable," while keeping the promotion rule constant (Karachiwalla and Park 2017; Nieddu and Pandolfi 2022; Bertrand et al. 2020; Li 2020). We complement this literature by assessing the causal impact of a more meritocratic promotion rule on worker productivity and its interaction with perceived pay progression.

Finally, this study contributes to the literature on pay inequality and worker performance. Existing empirical research primarily focuses on *horizontal* pay inequalities among workers in the same layer of an organization while shutting down dynamic incentives (Card et al. 2012; Mas 2017; Breza, Kaur, and Shamdasani 2018). In contrast to these studies, and in line with Cullen and Perez-Truglia (2022), we shift our attention to *vertical* pay inequalities between upper- and lower-tier workers for which the theoretical predictions are less clear. While steeper pay progression can potentially demotivate workers who are averse to vertical pay inequalities, it can also prompt an increase in effort through career incentives. Understanding which of the two effects prevails is of policy relevance given the recent rapid growth of the manager-worker pay ratio (Ashraf and Bandiera 2018; Shepherd 2003).

I. Context and Research Design

A. *The Community Health Worker Program in Sierra Leone*

In 2012, Sierra Leone's Ministry of Health and Sanitation (MoHS) established its first national Community Health Worker Program. The program is organized around Peripheral Health Units (PHUs), small health posts staffed with doctors and nurses.

Each PHU typically has a catchment area of seven to ten villages with one CHW per village and one Peer Supervisor per PHU.[1]

The role of the CHWs is to provide a basic package of healthcare services at the community level through home visits. They support expecting mothers and young children by providing health education, pre- and postnatal check-ups, basic medical care, and referrals to health clinics.

CHWs are hired locally and typically have no experience in the health sector before 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. They do so by training them monthly and accompanying them on household visits. Almost all PSs have previous experience as a CHW and have thus already acquired health knowledge.

CHWs and PSs are part-time employees who typically have a secondary occupation such as farming or small shopkeeping. In our sample, CHWs and PSs report dedicating an average of 18 and 11 hours per week, respectively, to their CHW/PS roles. CHWs are paid Le150,000 per month (US$17.50), while PSs are paid Le250,000 (US$29.20). Despite working fewer hours, PSs earn 67 percent more than CHWs, resulting in a pay gap. Based on self-reported hours, the hourly wage of PSs is 2.7 times higher than that of CHWs.

As with most public sector employees, CHWs and PSs are seldom fired. PSs usually leave their jobs at retirement when they turn 55 years old (Social Security Administration 2019). When a PS position becomes available, one of the CHWs in that PHU is promoted. The competition for a promotion thus happens within the PHU, and PSs are never pushed out by "upstart" high-performing CHWs.

The District Health Management Teams (DHMTs), which oversee the implementation of the CHW program at the district level, are in charge of the promotions, but they typically delegate these decisions to the head of the PHU (the "PHU in charge"), who is responsible for personnel and administrative matters. The system is reportedly subject to patronage and nepotism. Our data indicate that CHWs perceive this system as only partially meritocratic and believe that a connection to the PHU in charge is a key predictor of promotions. Indeed, only 41 percent of the CHWs surveyed reported that the PS was the best-performing CHW when promoted, and 50 percent reported perceiving the system as nonmeritocratic at baseline, a finding that we revisit in Section IIA.

B. *Research Design*

Our experiment took place in 372 PHUs spread across Sierra Leone and covered 372 PSs and 2,009 CHWs. These PHUs were cross-randomized into two treatment arms: (i) the "meritocratic promotions treatment" (T_{merit}), which introduced a more meritocratic promotion regime and (ii) the "pay progression information treatment"

[1] The CHW program was reorganized in February 2017. The updated program effectively reemployed all previously engaged CHWs and PSs from the earlier program and expanded by recruiting additional staff.

(T_{pay}), which provided information to CHWs about the supervisor's pay and which created variation in workers' *perceived* pay progression.[2]

Meritocratic Promotions Treatment.—In November 2018, we collaborated with the MoHS and the DHMTs to transition a random 186 PHUs to a more meritocratic promotion system ($T_{merit} = 1$), while the status quo was unaltered in the remaining 186 PHUs ($T_{merit} = 0$).

In the new promotion regime, the DHMTs committed to promote CHWs based on objective measures of CHW performance collected by the research team. Performance data were collected in $T_{merit} = 1$ and $T_{merit} = 0$ by measuring the number of visits and the average visit length through a household survey and unannounced spot checks with potential patients. 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 was used to decide on whom to promote. No information on performance was shared with DHMTs in the control PHUs ($T_{merit} = 0$).

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

> 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 groups, we reminded CHWs in the 186 control PHUs about the status quo promotion system ($T_{merit} = 0$). The following script was read 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 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.

During the ten months of our study, only 9 of the 372 PS positions in our sample became vacant. Therefore, this paper quantifies the effect of meritocracy on CHW performance in *anticipation* of future promotions. The four CHWs we see promoted to the PS position in $T_{merit} = 1$ ranked ten times higher in terms of performance compared to the five CHWs promoted in $T_{merit} = 0$. Despite the small sample size, this confirms that the DHMTs in $T_{merit} = 1$ used the information we provided to them.[3]

[2] The randomization was performed at the PHU level because promotions are decided at that level, as well as to limit spillovers. We stratified the randomization by district and presence of temporary performance-based incentives in a subsample of the PHUs (which is the focus of Deserranno et al. 2022). See Supplemental Appendix B for details.

[3] See Supplemental Appendix B for details on the implementation.

Pay Progression Information Treatment.—PSs and CHWs are paid Le250,000 and Le150,000 per month, respectively. Importantly, this pay gap was unknown to most CHWs before we revealed the information: only one-third of the CHWs guessed the PS pay correctly, while the remaining two-thirds either over- or underestimated PS pay (see Section IIB). We took advantage of this lack of information to create random variation in *perceived* pay progression. Cross-randomizing by the meritocratic promotions 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 immediately after informing them about the promotion system:

> CHWs are entitled to Le150,000 per month. PSs are entitled to Le250,000 per month, which is Le100,000 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 Le150,000 per month.

C. *Data and Timeline*

The treatments were implemented in November 2018, roughly six years after the CHW program was first established in Sierra Leone in 2012. We leverage three sources of data (Deserranno, Kastrau, and León-Ciliotta 2025):

- *CHW and PS surveys.* 372 PSs and 2,009 CHWs in the 372 PHUs were surveyed on their demographic background and job at two points in time: (i) at baseline in April–May 2018, roughly six months before the implementation of the treatments, and (ii) at endline in July–September 2019, roughly ten months after the implementation.

- *CHW perception surveys.* Two weeks before the implementation of the treatments (November 2018) and two weeks after (December 2018), we surveyed each CHW to assess her perception about meritocracy in the promotion system and pay progression.

- *Household surveys.* A random sample of nearly 10 percent of households' female heads was surveyed at endline in each village (July–September 2019). They were asked about the number of visits received by the CHW and the average length of those visits (which are used to measure CHW performance and are hence an input in the promotion decisions in $T_{merit} = 1$) as well as retrospective questions on their demographic background.[4]

[4] In Supplemental Appendix B, we discuss the sample, the accuracy of the performance measure, and the random spot checks. We argue that households are unlikely to misreport visits, even when connected to the CHW.

Supplemental Appendix Table A.1 reports summary statistics for PS, CHW, and household characteristics and shows that these characteristics are balanced across treatments. Pretreatment CHW beliefs are also balanced.

II. Beliefs Updating

In this section, we establish that our treatments shifted CHWs' beliefs about meritocracy and pay progression.

A. *Beliefs Updating about Meritocracy*

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 ten in terms of performance but does not know the PHU in charge outside of work versus another CHW who ranks X out of ten and knows the PHU in charge outside of work, where $X = \{2, 5, 10\}$. Our measure of perceived meritocracy takes a value of 1 if the CHW perceives the system as meritocratic (i.e., 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), a value of -1 if the CHW perceives the system as nonmeritocratic (i.e., believes that the best-performing worker is never promoted, even when the connected worker is the worst performer), and a value of 0 for intermediate 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).[5]

Figure 1 (panel A) presents the distribution of meritocracy perceptions before and after treatment among CHWs in the meritocratic promotions treatment ($T_{merit} = 1$) and the rest ($T_{merit} = 0$). In line with randomization, perceptions are comparable in $T_{merit} = 1$ and $T_{merit} = 0$ before treatment, with roughly 50 percent of CHWs perceiving the promotion system as meritocratic (prior of 1). Hence, the status quo promotion system is perceived as only partially meritocratic. After the introduction of the new, more meritocratic promotion system, CHWs updated their beliefs upward in $T_{merit} = 1$, with an extra 28.4 percent of CHWs perceiving the system as meritocratic. In $T_{merit} = 0$, CHWs did not significantly update their perceptions.

The corresponding regression results on belief updating are presented in Supplemental Appendix Table A.2 (column 1). They reveal that perceived meritocracy increases by 0.296 (63 percent) in $T_{merit} = 1$ relative to $T_{merit} = 0$. Columns 2–5 show that the meritocratic promotions treatment did not affect perceptions about other aspects of the job, such as the duration until the next promotion and PS pay.

[5] The notes in Figure 1 provide the exact wording of the question.

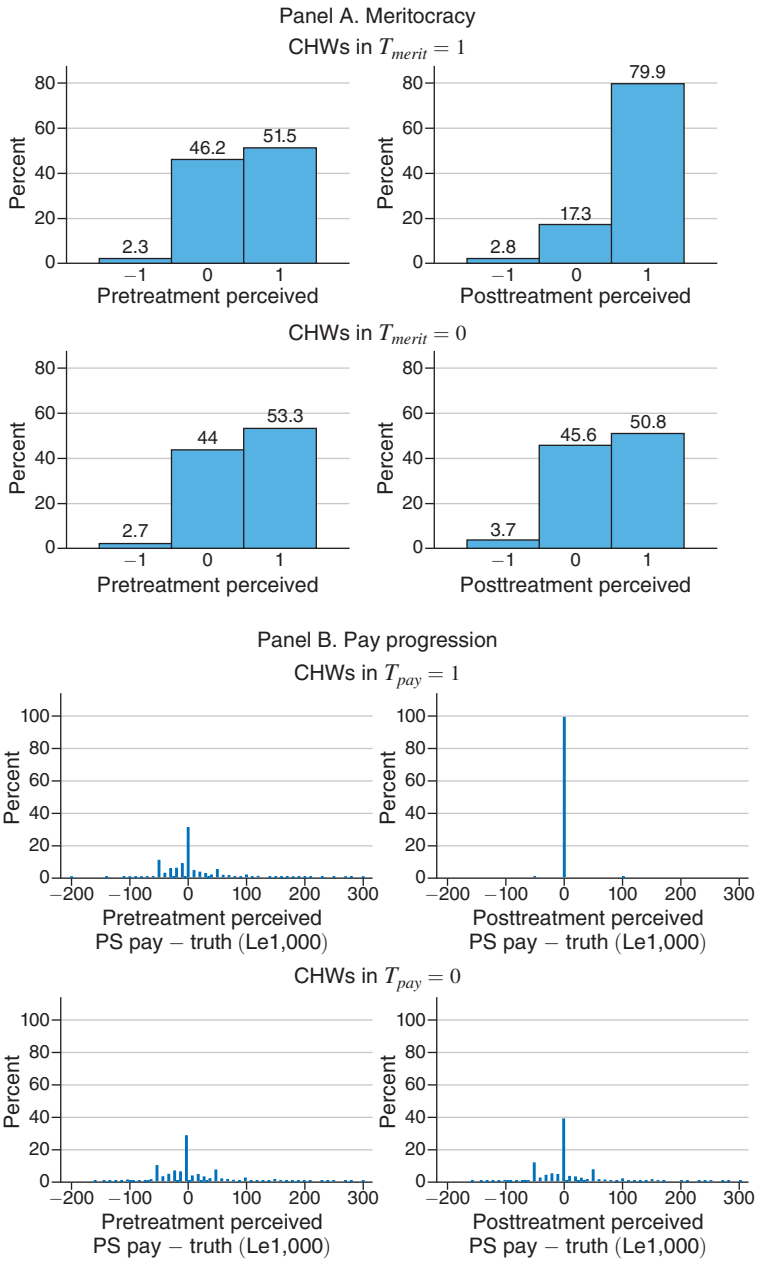

FIGURE 1. BELIEFS UPDATING ABOUT MERITOCRACY AND PAY PROGRESSION

Notes: This figure plots the distribution of perceived meritocracy in the promotion system (panel A) and the distribution of the difference between perceived PS pay and the truth (panel B) before and after treatment. To measure meritocracy, we asked each CHW before and after treatment: "A PHU needs a new PS. Who of the following two CHWs is most likely promoted to PS? (1) Alpha is the best-performing CHW (out of ten). Alpha does not know the PHU in charge outside of work. (2) Foday is the second-best-performing CHW (out of ten). Foday is a very good friend of the PHU in charge." We then repeated the same question but with Foday being the fifth-best and the worst-performing CHW. To measure perceived PS pay, we asked each CHW before and after treatment: "How much does your PS earn from the government each month?" We offered a reward of Le2,000 if the answer was correct. 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. We did not ask CHWs about perceptions of their own pay as this information was revealed to everyone at baseline.

B. *Beliefs Updating about Pay Progression*

Figure 1 (panel B) plots the difference between perceived and true PS pay before and after treatment among CHWs in the pay progression information treatment ($T_{pay} = 1$) and the rest ($T_{pay} = 0$). Perceptions of PS pay are comparable in $T_{pay} = 1$ and $T_{pay} = 0$ before treatment. In both groups, roughly 30 percent of the CHWs estimated correctly that PSs earn Le250,000 monthly, while 37 percent underestimated PS pay and 33 percent overestimated it.[6] After receiving information about PS pay, beliefs converge to the true PS pay in $T_{pay} = 1$. Workers who underestimated (overestimated) PS pay at baseline revised them *upward* (*downward*) in $T_{pay} = 1$, and those who correctly estimated it did not revise them. CHWs in $T_{merit} = 0$ instead barely updated their beliefs. The latter corroborates the lack of information spillover across treatments.

The corresponding regression results on belief updating are presented in Supplemental Appendix Table A.3. Column 1 shows that the mean absolute difference between perceived PS pay and the truth is Le482 in $T_{pay} = 1$ versus Le35,320 in $T_{pay} = 0$. Columns 2–5 show that T_{pay} has no effect on perceptions about the PS workload (hours), PS work expenses (transportation and communication), or meritocracy.

III. The Effect of Meritocratic Promotions on Worker Productivity

This section assesses the causal effect of a more meritocratic promotion regime (induced by T_{merit}) on CHW productivity. The interaction between T_{merit} and T_{pay} is the focus of Section IV.

We start by assessing the effect of T_{merit} on *average* performance using the following specification:

$$Y_{ij} = \alpha + \beta T_{merit,j} + \eta Z_j + \varepsilon_{ij}, \tag{1}$$

where Y_{ij} is the performance of worker i in PHU j, Z_j are the stratification variables, and ε_{ij} are standard errors clustered at the PHU level. The coefficient β captures the effect of the meritocratic promotions treatment ($T_{merit,j}$) for the average worker. 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). To obtain this measure, we take the total number of times a household received a routine visit, prenatal visit, or postnatal visit or was treated/referred for sickness, and then average these data at the CHW level. We will also consider visit length as a measure of visit quality (mean of 15 minutes).

[6] Similarly large misperceptions have been documented in other organizations (Cullen and Perez-Truglia 2022; Card et al. 2012). In our context, misperceptions exist because PS pay is not publicized to CHWs, and discussions between colleagues about pay is not the norm. In the baseline data, the size of the misperception about PS pay is correlated with the experience and the age of the CHW, while it is not correlated with connections with the PS or PHU in charge.

Table 1—Effect of Meritocracy on Worker Performance

Dependent variable	Number of visits				Visit length (in minutes)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T_{merit}	1.497				1.754			
	(0.479)				(0.651)			
$T_{merit} \times$ *High Rank*[i]		2.348				1.676		
		(0.605)				(0.880)		
$T_{merit} \times$ *Low Rank*[ii]		0.965				1.640		
		(0.567)				(0.790)		
$T_{merit} \times$ *Promotion Soon*[i]			3.476				2.559	
			(1.218)				(1.818)	
$T_{merit} \times$ *Promotion not Soon*[ii]			1.260				1.669	
			(0.510)				(0.685)	
$T_{merit} \times \mathbf{1}\{$*Prior PS Pay* $>$ *Truth*$\}$[i]				1.998				2.822
				(1.034)				(1.213)
$T_{merit} \times \mathbf{1}\{$*Prior PS Pay* $\leq$ *Truth*$\}$[ii]				0.403				0.423
				(0.800)				(1.123)
Observations	1,966	1,830	1,966	1,966	1,868	1,735	1,868	1,868
Mean dependent variable if $T_{merit} = 0$	6.749	6.749	6.749	6.749	11.99	11.99	11.99	11.99
p-value H_0: [i] − [ii] $= 0$		0.034	0.095	0.120		0.971	0.644	0.085
p-value MHT correction for [i]		0.004	0.004	0.016		0.024	0.060	0.004
p-value MHT correction for [ii]		0.016	0.004	0.506		0.008	0.004	0.673

Notes: The first column of each outcome variable reports the effect of T_{merit} for the average worker (estimate for β in equation (1)). The second column of each outcome variable reports the effect of T_{merit} for high-rank workers (ranked first, second, or third in terms of performance by the PS at baseline) and for low-rank workers (ranked fourth or more). These correspond to the estimates for β_1 and β_2 in equation (2) when $X_{ij} =$ *High Rank*. The third column of each outcome variable reports the effect of T_{merit} by whether the supervisor of the CHW is within five years of retirement age at baseline. These correspond to the estimates for β_1 and β_2 in equation (2) when $X_{ij} =$ *Promotion Soon*. The last column of each outcome variable presents the effect of T_{merit} by whether the prior about PS pay is above the median (actual salary of Le250,000) or not. These correspond to the estimates for β_{above} and $\beta_{at/below}$ in equation (3). Effectively, this means that we limit the comparisons to workers in $T_{pay} = 0$, who did not receive information on PS pay. All regressions control for the stratification variables and for the uninteracted x-variable (*High Rank*, *Promotion Soon*, *Prior PS Pay* depending on the column). "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. At the bottom of the table, we present p-values adjusted for multiple hypothesis testing across all columns computed using Romano and Wolf's (2016) step-down procedure.

Table 1 (column 1) and the corresponding Figure 2 (panel A) show that the number of visits provided by the average CHW increases by 1.497 (22 percent) in $T_{merit} = 1$ relative to $T_{merit} = 0$.

One possibility is that the CHWs compensate for the higher number of visits by providing shorter visits (by skipping some of the checklist items and thus reducing visit quality), by providing "easier" visits (more routine visits at the expense of fewer natal checks), or by targeting households that are physically or socially close to them (less costly to reach) at the expense of more deserving households. Such a quantity-quality trade-off does not exist in our context: the average visit length increases by 15 percent in $T_{merit} = 1$ relative to $T_{merit} = 0$ (Table 1, column 8), routine visits increase and most of the other types of visits increase too (Supplemental Appendix Table A.4), and household targeting does not change (Supplemental Appendix Table A.5).

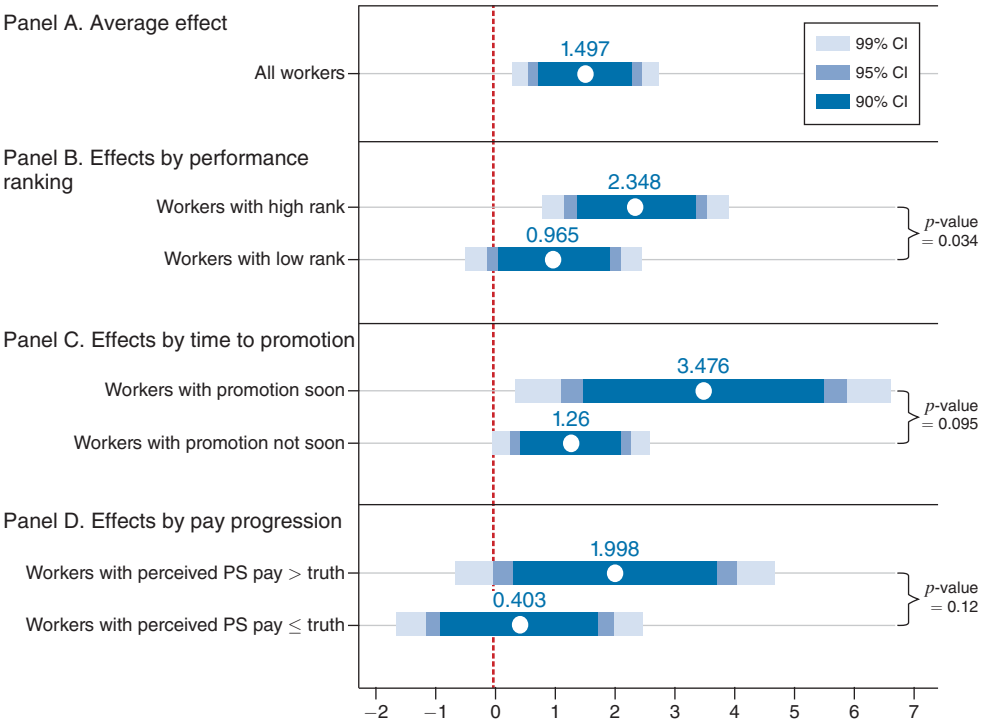

FIGURE 2. EFFECT OF MERITOCRACY ON THE NUMBER OF VISITS

Notes: Panel A plots the effect of T_{merit} on the number of visits provided by the average worker in our sample (estimate for β from equation (1)). Panel B plots the effect of T_{merit} for high-rank workers (ranked first, second, or third in terms of performance by the PS at baseline) and for low-rank workers (ranked fourth or more). These are the estimates for β_1 and β_2 in equation (2) when $X_{ij} =$ *High Rank*. Panel C plots the effect of T_{merit} by whether the supervisor of the CHW is within five years of retirement age at baseline. These are the estimates for β_1 and β_2 in equation (2) when $X_{ij} =$ *Promotion Soon*. Panel D plots the effect of T_{merit} by whether the prior about PS pay is above the median (Le250,000) or not. These correspond to the estimates for β_{above} and $\beta_{at/below}$ in equation (3). "Number of visits" is the average number of household visits provided by the CHW as reported by the households. The *p*-values reported at the right of the figure are for the difference in the treatment effects across worker types.

Having established that meritocratic promotions increase the quantity and quality of the visits provided by the average worker, we now test for heterogeneous productivity responses. In standard tournament theory (Lazear and Rosen 1981; Siegel 2010, 2014), where workers compete for promotions through increased effort, the effect of meritocratic promotions is predicted to be stronger for (i) workers who are highly ranked in terms of performance as they have a higher chance of being promoted in a meritocratic regime, (ii) workers who expect the promotion to materialize soon (higher net present value of the promotion), and (iii) workers with high priors about PS pay (higher value of the promotion).

We test for these heterogeneous effects by estimating

$$(2) \quad Y_{ij} = \alpha + \beta_1 T_{merit,j} \times X_{ij} + \beta_2 T_{merit,j} \times \left(1 - X_{ij}\right) + \delta X_{ij} + \eta Z_j + \varepsilon_{ij},$$

where X_{ij} is an indicator for whether a worker is highly ranked at baseline, expects the promotion soon, or has a high prior about PS pay. The coefficients of interest, β_1 and β_2, capture the effect of T_{merit} on the productivity of workers with $X_{ij} = 1$

and $X_{ij} = 0$, respectively.[7] The estimates of β_1 and β_2 can be visualized in Figure 2 and are presented formally in Table 1. At the bottom of Table 1, we present p-values adjusted for multiple hypothesis testing. In Supplemental Appendix Table A.7, we test for the robustness of the estimates to controlling in equation (2) for the correlates of X_{ij} and their interaction with $T_{merit,j}$.

Effects by Performance Ranking.—Our preferred measure for the ranking of each CHW within the PHU is the one reported by the PS at baseline. The PS has 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).[8]

Figure 2 (panel B) and the corresponding Table 1 (column 2) show that increasing meritocracy boosts the number of visits provided by high-rank workers (top three of their PHU) by 2.348, a 38 percent increase relative to the average for these workers in $T_{merit} = 0$. For "lower-rank" workers, the effect remains positive but is significantly smaller (0.965 visits). The effects are robust to controlling for the correlates of performance ranking (gender, education, wealth) and their interaction with T_{merit} (Supplemental Appendix Table A.7, column 1). The heterogeneity in the treatment effects can thus be attributed to the ranking rather than its correlates.

Supplemental Appendix Figure A.2 (panel A) presents the effect of meritocracy on worker productivity for the full distribution of worker ranking. The effect is positive and significant for workers ranked first, second, and third and converges to zero afterwards.[9]

Overall, the results indicate that a more meritocratic system increases effort for high-rank workers who have a shot at the promotion, while it does not affect the effort of low-rank workers who have no shot and face the same (low) incentives as in the old system.

Effects by Time to Promotion.—We proxy for "CHWs expecting a PS to leave her position soon" with an indicator for whether the supervisor is within five years of the standard retirement age and present robustness to other cutoffs.[10]

Figure 2 (panel C) and the corresponding Table 1 (column 4) show that, for workers who expect a promotion soon, meritocratic promotions increase the number of visits by 3.476 (+45 percent). The effect for workers who do not expect a promotion soon remains positive but is three times smaller (+1.260 visits).

As expected, the results decline when the PS is expected to retire further in the future: Supplemental Appendix Figure A.2 (panel B) shows that the effect of T_{merit} is stronger for workers who expect the PS to retire within two years, while it disappears

[7] Supplemental Appendix Table A.6 shows that workers with $X_{ij} = 1$ revised their perceptions of meritocracy in $T_{merit} = 1$ similarly as those with $X_{ij} = 0$. The estimates of β_1 and β_2 are hence not driven by differential belief updating.

[8] Ranking—as reported by the PS—is positively correlated with health knowledge, education, experience, and the number of household visits self-reported by the CHW. It is also correlated with the number of years the CHW has known the PS, a variable we control for in Supplemental Appendix Table A.7, while it does not correlate with connections to the PHU in charge.

[9] The effect is slightly smaller for workers who are ranked first instead of second or third, perhaps because they do not observe their competitors' effort and underestimate how hard these competitors try to catch up.

[10] Ten percent of the supervisors are within five years of retirement (more than 50 years old).

for workers who expect the PS to retire in ten years. Overall, the results indicate that the worker productivity response to meritocracy intensifies in the years leading up to promotion eligibility.

Effect by Pay Progression.—We now explore the effect of T_{merit} on the productivity of workers whose prior about PS pay is above the median (i.e., above the actual salary of Le250,000) versus below the median. We limit the comparisons to workers in $T_{pay} = 0$ who did not receive information on PS pay.[11]

Figure 2 (panel D) and the corresponding Table 1 (column 6) show that the meritocratic promotions treatment increases the number of visits provided by workers with above-median perceived pay progression by 1.998 (+30 percent). It has no effect on workers with below-median perceived pay progression. The results become even stronger when we control for the correlates of baseline perceived pay progression and their interaction with T_{merit} (column 7).

Supplemental Appendix Figure A.2 (panel C) presents the effect of more meritocracy for different values of priors about PS pay. The figure confirms that the effect of the meritocracy treatment on worker productivity increases with perceived pay progression.

IV. The Effect of Pay Progression on Worker Productivity by Meritocracy

This section studies the impact of pay progression (induced by T_{pay}) on worker productivity in the new meritocratic regime ($T_{merit} = 1$) vis-à-vis the old less meritocratic regime ($T_{merit} = 0$). Unlike other 2×2 experiments, our analysis will *not* rely on a double-interacted specification where productivity is regressed on T_{merit}, T_{pay}, and $T_{merit} \times T_{pay}$. This specification is not informative in our context because we have shown that workers in $T_{pay} = 1$ update their beliefs about pay progression—and hence change their productivity—in opposite directions depending on whether they initially underestimated or overestimated PS pay. Consequently, the average effect of revealing PS pay ($T_{pay} = 1$ versus $T_{pay} = 0$) on CHW productivity is found to be zero.[12] This null effect stems from a similar proportion of workers under- and overestimating PS pay at baseline and the opposing effort responses of these groups that offset each other.

To account for these heterogeneous responses to T_{pay}, we interact T_{merit}, T_{pay}, and $T_{merit} \times T_{pay}$ with indicators for whether workers' priors about PS pay are above, below, or at the truth (q_{ij}):

$$
\begin{aligned}
(3) \qquad Y_{ij} = \alpha &+ \sum_{q=\{below,above,at\}} \gamma_q \left[T_{pay,j} \times T_{merit,j} \times q_{ij} \right] \\
&+ \sum_{q=\{below,above,at\}} \delta_q \left[T_{pay,j} \times \left(1 - T_{merit,j}\right) \times q_{ij} \right] \\
&+ \sum_{q=\{below,above,at\}} \beta_q \left[T_{merit,j} \times q_{ij} \right] + \sum_{q=\{below,above\}} \lambda_q q_{ij} + \eta Z_j + \varepsilon_{ij}.
\end{aligned}
$$

[11] The corresponding comparisons in $T_{pay} = 1$ are uninformative because beliefs converge to the truth in $T_{pay} = 1$ (see Figure 1, panel B), and we would be comparing workers with the same ex post beliefs even though their ex ante beliefs were different. These are the estimates for β_{above}, β_{below}, and β_{at} from equation (3).

[12] See Supplemental Appendix Table A.8, where we use a double-interacted model.

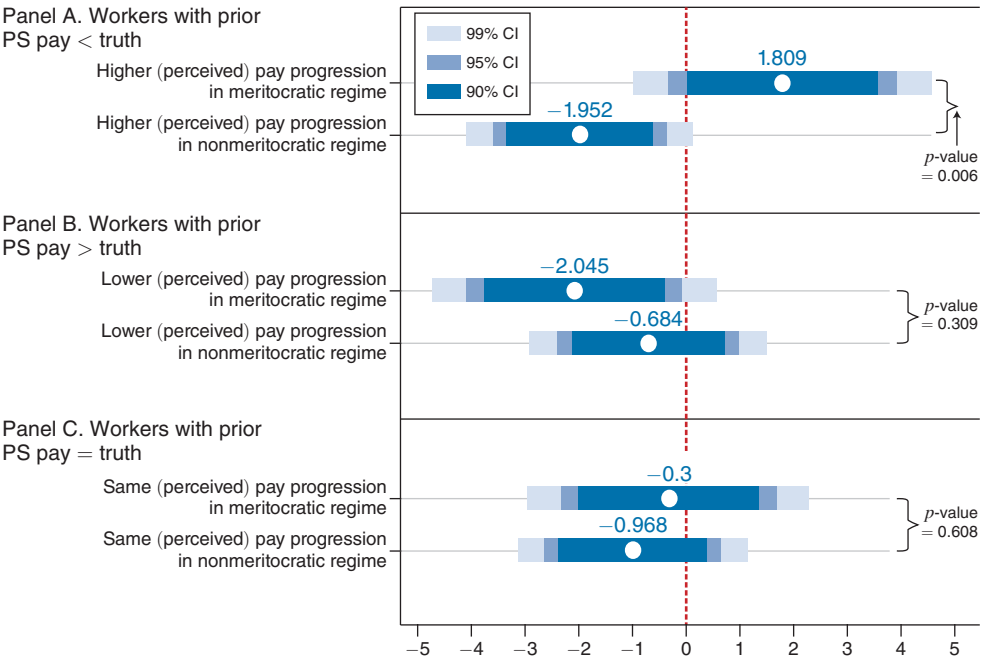

FIGURE 3. EFFECT OF PAY PROGRESSION ON THE NUMBER OF VISITS, BY MERITOCRACY

Notes: This figure plots the effects of T_{pay} on the number of visits in the new meritocratic regime ($T_{merit} = 1$) and in the old nonmeritocratic regime ($T_{merit} = 0$) for three types of workers: those who underestimated PS pay at baseline and for whom perceived pay progression increased (panel A), those who overestimated PS pay and for whom perceived pay progression decreased (panel B), and those who correctly estimated PS pay and for whom perceived pay progression did not change (panel C). Panel A plots γ_{below} and δ_{below} estimated from equation (3). Panel B plots γ_{above} and δ_{above} estimated from equation (3). Panel C plots γ_{at} and δ_{at} estimated from equation (3). "Number of visits" is the average number of household visits provided by the CHW as reported by the households. The *p*-values reported at the right of the figure are for the difference in the treatment effects across regime systems within a worker type.

The coefficients of interest are the γ's and δ's, which capture the causal effect of revealing PS pay (T_{pay}) in the new system ($T_{merit} = 1$) and in the old system ($T_{merit} = 0$), respectively.[13] Throughout the analysis, we refrain from making across-group comparisons—for example, γ_{above} versus γ_{below} or δ_{above} versus δ_{below}—as these could reflect baseline differences across groups. We focus instead on identifying the effect of revealing PS pay *within a worker type*, which is causal.[14]

The results can be visualized in Figure 3. They are presented formally in Table 2. In Supplemental Appendix Table A.9, we present their robustness to extending equation (3) to include the correlates of baseline perceived pay progression (age and experience of the CHW) and their interaction with T_{pay}, T_{merit}, and $T_{pay} \times T_{merit}$.

[13] The heterogeneity analysis by whether the promotion system is more or less meritocratic was prespecified in the AEA RCT Registry. The heterogeneity by whether workers under-, over-, or correctly estimated PS pay was not prespecified because we initially expected that most workers would underestimate PS pay. See Supplemental Appendix A for more details.

[14] CHW characteristics are balanced across treatments within a worker type (Supplemental Appendix Table A.10).

Table 2—Effect of Pay Progression on Worker Performance, by Meritocracy

Dependent variable	Number of visits (1)	Visit length (in minutes) (2)
Panel A. Effects for workers who underestimated PS pay at baseline [*higher pay progression in* $T_{pay} = 1$]		
$T_{pay} \times Meritocratic\ (T_{merit} = 1) \times \mathbf{1}\{Prior\ PS\ Pay < Truth\}$ [i]	1.809	1.330
	(1.075)	(1.291)
$T_{pay} \times Non\text{-}meritocratic\ (T_{merit}=0) \times \mathbf{1}\{Prior\ PS\ Pay < Truth\}$ [ii]	−1.952	−1.846
	(0.822)	(1.243)
Panel B. Effects for workers who overerestimated PS pay at baseline [*lower pay progression in* $T_{pay} = 1$]		
$T_{pay} \times Meritocratic\ (T_{merit} = 1) \times \mathbf{1}\{Prior\ PS\ Pay > Truth\}$ [iii]	−2.045	−2.186
	(1.023)	(1.215)
$T_{pay} \times Non\text{-}meritocratic\ (T_{merit} = 0) \times \mathbf{1}\{Prior\ PS\ Pay > Truth\}$ [iv]	−0.684	−0.639
	(0.860)	(1.316)
Panel C. Effects for workers who correctly estimated PS pay at baseline [*same pay progression in* $T_{pay} = 1$]		
$T_{pay} \times Meritocratic\ (T_{merit} = 1) \times \mathbf{1}\{Prior\ PS\ Pay = Truth\}$ [v]	−0.300	1.308
	(1.018)	(1.460)
$T_{pay} \times Non\text{-}meritocratic\ (T_{merit} = 0) \times \mathbf{1}\{Prior\ PS\ Pay = Truth\}$ [vi]	−0.968	−0.008
	(0.833)	(1.615)
Observations	1,966	1,868
Mean dependent variable if $T_{pay} = 0$	7.965	13.191
p-value H_0: [i] − [ii] = 0	0.006	0.077
p-value H_0: [iii] − [iv] = 0	0.309	0.385
p-value H_0: [v] − [vi] = 0	0.608	0.546

Notes: This table presents the effects of T_{pay} on the number of visits in the meritocratic regime ($T_{merit} = 1$) and in the nonmeritocratic regime ($T_{merit} = 0$), estimated from equation (3). Panel A reports the estimates for γ_{below} and δ_{below} (effects for workers who underestimated PS pay at baseline). Panel B reports the estimates for γ_{above} and δ_{above} (effects for workers who overestimated PS pay at baseline). Panel C reports the estimates for γ_{at} and δ_{at} (effects for workers who correctly estimated PS pay at baseline). All regressions control for the stratification variables, $\mathbf{1}\{Prior\ PS\ Pay < Truth\}$ and $\mathbf{1}\{Prior\ PS\ Pay > Truth\}$, and these last two variables multiplied with T_{merit} (see equation (3)). $\mathbf{1}\{Prior\ PS\ Pay < Truth\}$ ($\mathbf{1}\{Prior\ PS\ Pay > Truth\}$) equals 1 if the pretreatment perception about PS salary is below (above) the actual salary of Le250,000 and 0 otherwise. "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. Differences in the number of observations are due to missing values. Standard errors are clustered at the PHU level.

Workers Who Underestimated PS Pay ($Prior < Truth$).—We start by assessing the effect of revealing the true PS pay (T_{pay}) on the productivity of workers who underestimated PS pay at baseline. These correspond to $\hat{\gamma}_{below}$ and $\hat{\delta}_{below}$ from equation (3) and capture the effects of *increased* pay progression in the more versus less meritocratic regime, respectively.

In the new, more meritocratic regime, higher pay progression increases the number of visits by 1.809 (+23 percent), while it *reduces* the number of visits provided by 1.952 (−27 percent) in the old, less meritocratic regime: see Figure 3 (panel A) and the corresponding Table 2 (panel A, column 1). The results on visit length go in the same direction but are less precise (columns 3 and 4).

The results indicate that steeper pay progression motivates the workers to climb the organization's ladder and prompts more effort when promotions are performance based. When promotions are not performance based, steeper pay progression instead reduces worker performance.

Two potential mechanisms can explain the observed reduction in worker productivity when pay progression increases in a low meritocratic system. The first

is a negative morale effect: the organization may be perceived as more unfair if it increases the pay gap between the PSs and the CHWs *without* promoting workers in a meritocratic way, and this may demotivate CHWs. The second is one of lobbying: when pay progression increases, workers may be more interested in a promotion and may start devoting more time to lobbying (talking with the PHU in charge) to increase their chances of promotion in a nonmeritocratic regime (de Janvry et al. 2023). That said, the context we analyze is one where it is very hard for workers to "lobby" because they are typically located far away from the PHU in charge and do not work in direct contact with them. In fact, half of the CHWs have never even talked to the PHU in charge at baseline (Supplemental Appendix Table A.1, panel B). Moreover, we find no evidence of increased lobbying when pay progression increases: the likelihood that a CHW communicated with the PHU in charge in the past year, and the fraction of time she reports dedicating to non-patient-related activities, which include communications with the PHU in charge, did not increase (Supplemental Appendix Table A.11). Overall, the results provide suggestive evidence that, with low meritocracy, increasing pay progression reduces productivity through negative morale effects.

Workers Who Overestimated PS Pay ($Prior > Truth$).—We now turn to the effect of revealing the true PS pay (T_{pay}) on the productivity of workers who overestimated PS pay at baseline. These correspond to $\hat{\gamma}_{above}$ and $\hat{\delta}_{above}$ from equation (3) and capture the effects of *reducing* pay progression in a more versus less meritocratic regime, respectively.

In the new, more meritocratic regime, lower pay progression reduces the number of visits by 2.045 (-21 percent), while it has no significant effect in the old, less meritocratic regime: see Figure 3 (panel B) and the corresponding Table 2 (panel B, column 1). This suggests that a reduction in perceived pay progression in a low meritocratic system is not perceived as more fair, or at least does not increase fairness by enough to raise worker productivity.

Overall, the results indicate that lower pay progression reduces performance only in the meritocratic regime, where promotions are linked to performance and the marginal returns to effort are higher.

Workers Who Correctly Estimated PS Pay ($Prior = Truth$).—As a placebo check, we look at workers who correctly estimated PS pay at baseline and did not update their perception of pay progression in $T_{pay} = 1$. Revealing the true PS pay has no statistically significant effect on their performance regardless of whether the system is meritocratic: see Figure 3 (panel C) and Table 2 (panel C, column 1). This is reassuring as it indicates that providing information about true PS pay does not affect workers' behavior through channels unrelated to reassessing their beliefs.

V. Conclusion

In a field experiment with the Ministry of Health in Sierra Leone, we show that a more meritocratic promotion system increases the productivity of frontline health workers. This is especially true for highly ranked workers with a shot at promotion or those perceiving the promotion reward as large. Higher pay progression also increases

the productivity of frontline health workers, but only when promotions are meritocratic. When promotions are not highly meritocratic, higher (perceived) pay progression appears to demotivate workers, causing a reduction in their productivity. Overall, these findings underscore the importance for organizations to consider the interaction between two important personnel tools: meritocratic promotions and pay progression.

Our results also indicate additional important implications for further investigation. First, our analysis evaluates the effects ten months after the introduction of the new meritocratic promotion system, recognizing that the impact of meritocratic promotions may evolve, and potentially amplify, over the long term. While we observe few promotions in the span of our study, in the long run, more workers becoming eligible for a meritocratic promotion could intensify their efforts. Additionally, the quality of higher-level staff may improve and increase lower-tier workers' efforts. A meritocratic promotion system may also enhance the quality of workers in the applicant pool, generating positive effects over time through the selection process. Assessing the long-term effects of meritocratic promotions presents a great avenue for future research.

Second, although our research centers on performance-based promotions, organizations have the option of adopting pay-for-performance schemes that do not necessitate worker competition. Our results show that thanks to the tournament structure, promotion incentives convert only a small fraction of gains into higher wages. Nonetheless, it is essential to conduct further studies to evaluate their cost-effectiveness in comparison to non-tournament-based approaches. From a practical perspective, implementing pay-for-performance schemes can be challenging, often facing opposition from public sector unions. Performance-based promotions might be more viable in public sector settings, given that promotions are inevitable and someone will advance.

Finally, the effectiveness of performance-based promotions (or any type of performance-based incentive) depends on an organization's ability to accurately measure worker performance. The less precise the performance measurement, the less likely it is that performance pay will effectively boost worker effort. In our study, performance was assessed through a household survey. Although this method is undoubtedly more costly, it likely provides more accuracy than relying on government administrative data, which often depends on (inflated) self-reports (Singh 2020). Additionally, unlike in our study, government administrative data typically focus on quantitative output measures, which are more observable but overlook quality indicators. This suggests that incentive schemes might reward only certain aspects of performance, potentially leading to diminished effort in other aspects when workers are multitasking. Understanding how to improve the measurement of performance on a large scale is a critical challenge for future research. Addressing this issue could enable the implementation of more performance-based human resource management systems in the public sector.

REFERENCES

Ashraf, Nava, and Oriana Bandiera. 2018. "Social Incentives in Organizations." *Annual Review of Economics* 10: 439–63.

Asiedu, Edward, Dean Karlan, Monica Lambon-Quayefio, and Christopher Udry. 2021. "A Call for Structured Ethics Appendices in Social Science Papers." *Proceedings of the National Academy of Sciences* 118 (29): e2024570118.

Banerjee, Abhijit V., Esther Duflo, and Rachel Glennerster. 2008. "Putting a Band-Aid on a Corpse: Incentives for Nurses in the Indian Public Health Care System." *Journal of the European Economic Association* 6 (2–3): 487–500.

Benson, Alan, Danielle Li, and Kelly Shue. 2019. "Promotions and the Peter Principle." *Quarterly Journal of Economics* 134 (4): 2085–134.

Benson, Alan, Danielle Li, and Kelly Shue. 2021. "Potential and the Gender Promotion Gap." Unpublished.

Bertrand, Marianne, Robin Burgess, Arunish Chawla, and Guo Xu. 2020. "The Glittering Prizes: Career Incentives and Bureaucrat Performance." *Review of Economic Studies* 87 (2): 626–55.

Breza, Emily, Supreet Kaur, and Yogita Shamdasani. 2018. "The Morale Effects of Pay Inequality." *Quarterly Journal of Economics* 133 (2): 611–63.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review* 102 (6): 2981–3003.

Cullen, Zoë B., and Ricardo Perez-Truglia. 2022. "How Much Does Your Boss Make? The Effects of Salary Comparisons." *Journal of Political Economy* 130 (3): 766–822.

de Janvry, Alain, Guojun He, Elisabeth Sadoulet, Shaoda Wang, and Qiong Zhang. 2023. "Subjective Performance Evaluation, Influence Activities, and Bureaucratic Work Behavior: Evidence from China." *American Economic Review* 113 (3): 766–99.

de Ree, Joppe, Karthik Muralidharan, Menno Pradhan, and Halsey Rogers. 2018. "Double for Nothing? Experimental Evidence on an Unconditional Teacher Salary Increase in Indonesia." *Quarterly Journal of Economics* 133 (2): 993–1039.

Deserranno, Erika. 2019. "Financial Incentives as Signals: Experimental Evidence from the Recruitment of Village Promoters in Uganda." *American Economic Journal: Applied Economics* 11 (1): 277–317.

Deserranno, Erika, Stefano Caria, Philipp Kastrau, and Gianmarco León-Ciliotta. 2022. "The Allocation of Incentives in Multi-layered Organizations." Barcelona School of Economics Working Paper 1346.

Deserranno, Erika, Philipp Kastrau, and Gianmarco Leon. 2019. *Vertical Pay Differentials and Worker Productivity: Evidence from the Health Worker Program in Sierra Leone*. https://doi.org/10.1257/rct.3993-1.0.

Deserranno, Erika, Philipp Kastrau, and Gianmarco León-Ciliotta. 2025. *Data and Code for: "Promotions and Productivity: The Role of Meritocracy and Pay Progression in the Public Sector."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E203501V1.

Finan, Frederico, Benjamin A. Olken, and Rohini Pande. 2017. "The Personnel Economics of the Developing State." In *Handbook of Economic Field Experiments*, Vol. 2, edited by Abhijit Vinayak Banerjee and Esther Duflo, 467–514. North-Holland: Elsevier.

Gibbons, Robert, and Kevin J. Murphy. 1992. "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence." *Journal of Political Economy* 100 (3): 468–505.

Gibbons, Robert, and Michael Waldman. 1999a. "A Theory of Wage and Promotion Dynamics inside Firms." *Quarterly Journal of Economics* 114 (4): 1321–58.

Gibbons, Robert, and Michael Waldman. 1999b. "Careers in Organizations: Theory and Evidence." In *Handbook of Labor Economics*, Vol. 3B, edited by Orley C. Ashenfelter and David Card, 2373–437. North-Holland: Elsevier.

Karachiwalla, Naureen, and Albert Park. 2017. "Promotion Incentives in the Public Sector: Evidence from Chinese Schools." *Journal of Public Economics* 146: 109–28.

Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy* 89 (5): 841–64.

Li, Xuan. 2020. "The Costs of Workplace Favoritism: Evidence from Promotions in Chinese High Schools." Unpublished.

Mas, Alexandre. 2017. "Does Transparency Lead to Pay Compression?" *Journal of Political Economy* 125 (5): 1683–721.

Meyer-Sahling, Jan-Hinrik, Christian Schuster, and Kim Sass Mikkelsen. 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." Unpublished.

Nieddu, Marco G., and Lorenzo Pandolfi. 2022. "The Effectiveness of Promotion Incentives for Public Employees: Evidence from Italian Academia." *Economic Policy* 37 (112): 697–748.

Peter, Laurence J., and Raymond Hull. 1969. *The Peter Principle: Why Things Always Go Wrong*. London: Souvenir Press.

Romano, Joseph P., and Michael Wolf. 2016. "Efficient Computation of Adjusted *p*-values for Resampling-Based Stepdown Multiple Testing." *Statistics & Probability Letters* 113: 38–40.

Rosen, Sherwin. 1986. "Prizes and Incentives in Elimination Tournaments." *American Economic Review* 76 (4): 701–15.

Shepherd, G. 2003. "Civil Service Reform in Developing Countries: Why Is It Going Badly?" Paper presented at the 11th International Anti-Corruption Conference, Seoul, South Korea, May 25–28.

Siegel, Ron. 2010. "Asymmetric Contests with Conditional Investments." *American Economic Review* 100 (5): 2230–60.

Siegel, Ron. 2014. "Asymmetric Contests with Head Starts and Nonmonotonic Costs." *American Economic Journal: Microeconomics* 6 (3): 59–105.

Singh, Abhijeet. 2020. "Myths of Official Measurement: Auditing and Improving Administrative Data in Developing Countries." Research on Improving Systems of Education Working Paper 20/042.

Social Security Administration. 2019. *Social Security Programs throughout the World: Africa, 2019 – Sierra Leone*. Washington, DC: Social Security Administration.

Xu, Guo, Erika Deserranno, Diana Moreira, and Edoardo Teso. 2023. "Bureaucracy." *VoxDevLit* 8 (3): 1–18.