RATIONALITY & COMPETITION
CRC TRR 190

Cooperation in the Workplace

Hoa Ho (LMU Munich)
Maren Mickeler (ESSEC)
SIlvia Castro (LMU Munich)

Discussion Paper No. 531

April 25, 2025

Collaborative Research Center Transregio 190 | www.rationality-and-competition.de
Ludwig-Maximilians-Universität München | Humboldt-Universität zu Berlin
Spokesperson: Prof. Dr. Klaus M. Schmidt, Ludwig Maximilian University of Munich, 80539 Munich, Germany
info@rationality-and-competition.de

Cooperation in the Workplace: *Experimental Evidence from Knowledge Workers*

Silvia Castro[*], Hoa Ho[†], Maren Mickeler[§]

April 25, 2025

Abstract

Organizations rely on peer-to-peer knowledge exchange among employees, yet incentivizing cooperative behaviors is a challenge. This study evaluates an intervention designed to encourage peer support in the largest bank in Uganda. Using a cluster randomized controlled trial, we introduced a public recognition incentive—awarding employees identified as the most supportive by their peers and supervisors. The intervention increases employees' willingness to help by 21% in expertise-sharing and 12% in mentoring. The incentive's effectiveness stems from its role in enhancing professional reputation and career prospects. A replication exercise in a second bank confirms the findings and the external validity of the results.

[*]Department of Economics, Ludwig-Maximilians-Universität, Munich, Germany

[†]Department of Accounting, Ludwig-Maximilians-Universität, Munich, Germany

[‡]Department of Management, ESSEC Business School, Cergy, France

[§]We would like to thank Fabian Kasi, Solomy Luyombo, Charles Kabande, Sebastian Mulindwa, Fred Akena, and Nelson Luvule for hosting this project. We thank Martina Sanin and Robert Canwat for their excellent research assistance and project management and Karen Bernhardt-Walther, Florian Englmaier, Guido Friebel, Svenja Friess, and Nathalie Romer for their feedback. Support from the Deutsche Forschungsgemeinschaft through CRC TRR 190 (Project-ID 280092119) and TRR 266 (Project-ID 403041268) is gratefully acknowledged. In addition, this project has received financial support from ESSEC, the Joachim Herz Foundation, and the Hans Riegel Foundation. This experiment was pre-registered as AEARCTR-0012191 and received ethics approval from the Ethics Review Board of ESSEC, of Makerere University (CoBAMS-REC-2024-40), and of the Ugandan National Council for Science and Technology (SS3282ES).

1 Introduction

Organizations depend on employee collaboration. Much of the expertise within firms is tacit —difficult to codify and often acquired only through experience. This tacit knowledge is transferred through peer-to-peer exchange (Battiston et al., 2021). When experienced employees share job-specific knowledge with their peers, they reduce inefficiencies, enhance productivity, and accelerate learning (Hamilton et al., 2003; Sandvik et al., 2020; Battiston et al., 2021).

Yet, fostering cooperative behaviors is challenging. In many organizations, meaningful support to peers is difficult for management to observe, measure, or integrate into standard contracts. Consequently, firms incentivize easily quantifiable individual productivity, leading employees to focus their efforts on these tasks (Holmstrom & Milgrom, 1991; Auriol et al., 2002), even when less quantifiable actions are valuable to the organization. This misalignment creates a tension between individual and organizational goals. Addressing this tension would require firms to measure and incentivize both individual performance and cooperative behaviors.

In this paper, we evaluate a human resources intervention designed to incentivize peer collaboration in a large banking corporation. First, we developed a novel tool to assess the *cooperativeness* of bankers within the organization. Subsequently, in partnership with the bank's top management, we introduced an incentive for cooperation. In particular, top supportive employees received an award acknowledging their *outstanding support to peers*. The award was identified as a suitable incentive, as certificates were considered a valued signal of cooperation—an important skill in the organization and industry. Moreover, monetary bonuses were dismissed due to their non-negligible cost and uncertainty about unintended consequences.[1]

We evaluate the intervention through a pre-registered cluster randomized controlled trial (RCT) conducted across 78 bank branches. To measure employees' *cooperativeness*, we relied on peer voting: each banker identified colleagues in their network from whom they frequently received help and selected the single most supportive person. To mitigate strategic voting, supervisors also assessed their subordinates' cooperation, contributing one-third of the final score. Branches were then randomly assigned to either the treatment or control group. In treated branches, employees ranked in the top 20% of cooperativeness received public recognition during a branch-wide meeting. Control branches continued business as usual without the recognition program. Our RCT design overcomes concerns about omitted variable and selection biases by randomly awarding the most supportive employees in half of the branches, allowing us to isolate the causal impact of the recognition incentive.

[1]Economic research has shown that monetary incentives can have adverse effects, particularly when used to promote altruistic or prosocial tasks (Alfitian et al., 2023; Ashraf et al., 2020; Wagner et al., 2020).

The organization involved in this study employs 3,000 staff across various sections of commercial and corporate banking. Our study involved their credit division, which comprises over 1,000 bankers engaged in knowledge-intensive, non-routine lending tasks. Their roles encompass both independent and collaborative activities, with employees frequently needing help from their peers for analyzing complex credit proposals, recovering disbursements from non-compliant borrowers, formulating loan cases for committee adjudication, and appraising large projects. Compensation includes a fixed base salary and a sizable variable incentive tied to individual performance.

Our study evaluates the impact of the intervention on workplace cooperation. Measuring daily cooperative behaviors, however, presents a challenge for researchers, as these interactions occur in person at the branches and vary depending on the task. To address this, in collaboration with the firm's leadership, we identified two systematically measurable and organizationally valuable cooperation metrics: (1) *expertise contributions*—employees' decisions to provide expert feedback in two structured calls for input, and (2) *peer mentoring*—the decision to volunteer for a junior mentoring program. These metrics were collected at two and six months post-intervention. *Expertise contributions* were measured twice, while *peer mentoring* was measured once, both concurrently with a wave of workplace surveys. After completing the survey, employees received an on-screen prompt inviting them to volunteer for one of these cooperative tasks. They saw a brief task description and could choose to participate, with their decision recorded in the system. For one task, we also tracked the time spent and the content of the contributions. Thus, we are able to assess both the decision to help and the quality the help provided. These metrics capture a broad willingness to contribute to the organization's collective good. To complement these measures, we will analyze workplace surveys on perceived help and support within the organization.

We find that the intervention significantly increases employees' willingness to help their peers. In treated branches, employees are 21% (p-value < 0.05) and 25% (p-value < 0.01) more likely to participate in *expertise contributions* at two and six months post-intervention, respectively, and 12% more likely to volunteer as *peer mentors* (p-value < 0.05). Peer ratings strongly predict cooperation—employees with higher peer ratings have higher values of expertise contribution, reinforcing the validity of our proxy measures for helping behavior. The intervention's effects do not differ between awarded and non-awarded employees, suggesting that the incentive promotes helping behaviors even among those who were not directly recognized but worked in a branch where such behaviors were acknowledged.

The incentive does not compromise contribution quality. For the *expertise contribution* task measured two months post-intervention, employees provided strategies for a repository supporting low performers. Evaluating both effort (proxied by time spent) and quality (assessed through treatment-blind expert evaluations), we find that while effort levels are similar across treated and control branches, the novelty and helpfulness of contributions are over 6%

higher in treated branches. Lastly, self-reported survey metrics indicate a significant improvement in perceived assistance in treated branches post-intervention. Employees in these branches report 0.17 standard deviations higher agreement (p-value < 0.05) with the statement that they received assistance from peers in the past month. Thus, the intervention increases cooperative behaviors without compromising the quality of help.

Having established that the intervention increased cooperation on average, we next examine the mechanisms behind this effect. The changes could be driven by shifts in individual motivation, horizontal relationships (among peers), or vertical perceptions (toward management). To explore this, we embedded a variation within the first *expertise contribution* measurement, conducted two months post-intervention. When employees were invited to share expertise, we experimentally manipulated the visibility of their participation along three dimensions. One group was informed that their participation would remain *private*. A second group was told that their involvement would be known to their *peers in the branch*. The third group was informed that *top management* would be aware of their decision to contribute. This variation was randomized at the individual level within both the treatment and control arms. This feature allowed us to assess how visibility influences helping behaviors independently of the incentive and how it interacts with the intervention.[2]

The intervention influences cooperation through two mechanisms: peer reputation and career concerns, while leaving intrinsic motivation unchanged. Under the private condition, 43% of employees contribute in the control group, compared to 44% in the treatment group, suggesting that the intervention does not affect intrinsic motivation (p-value = 0.82). However, under peer visibility, contributions in treated branches increase by 15% (p-value $<$ 0.05). Employees may now feel a stronger obligation to help or become more aware that their contributions are noticed and valued by colleagues. Survey data further support a shift in social expectations: six months after the intervention, treated employees report a 0.33 standard deviation increase, relative to the control group, in how difficult they find it to decline requests for assistance from peers (p-value $<$ 0.01). This effect is particularly pronounced among award recipients, who report a 0.6 standard deviation increase. Lastly, when contributions are visible to management, cooperation rises by 30% (p-value $<$ 0.01). Plausibly, the award serves as a clear managerial signal that cooperation is valued. Employees may adjust their behavior to align with organizational expectations, either to enhance career prospects or because they recognize that helping others is rewarded.

Our article offers two main contributions. First, we provide a proof-of-concept that it *is* possible to incentivize peer support, the sharing of tacit knowledge, and cooperation without relying on team incentives. This is an important contribution, as team incentives may not be feasible in some organizations, or a firm may need to enhance cooperation not only within

[2]For details on ethical considerations and the handling of potential deception, see Subsection 2.5 on study design, which outlines measures ensuring fairness and transparency.

teams but also across teams. Second, we make a methodological contribution, introducing a simple, effective, and scalable scoring tool for organizations to measure cooperation and implement incentives accordingly. Lastly, this study empirically unpacks the complex interpersonal dynamics of tacit knowledge exchange, showing how both social image toward peers and management-related reputation serve as strong incentives for cooperative behavior.

Our findings have broad applicability. Work settings characterized by complex, interdependent tasks are prevalent not only in corporate environments but also in the public sector and academia (Autor et al., 2003). Our partner organization, a large financial sector firm, shares structural features with commercial banks worldwide, including standardized hiring processes, performance-based incentives, and industry-standard practices for managing productivity and collaboration. These shared characteristics make our results relevant to similar organizations in financial and other knowledge-intensive industries.

Replication is essential for the credibility of the sciences (Camerer et al., 2018) and particularly affects field experiments, where results are often context-dependent and tied to a group of researchers carefully designing and implementing the interventions. To address these concerns, we conducted a pre-registered replication study in a different bank to test the robustness of our findings. This institution, an international corporate bank with headquarters in the United States and 30 branches in Uganda, differed from our original partner in size, culture, and geographic scope. Using the same *Most Supportive Employee* award procedure and outcome metrics, we observed a 13% increase (p-value = 0.08) in willingness to help in treated branches, consistent with the main study. Baseline levels of helping behavior and treatment effects were strikingly similar, reinforcing confidence in the internal and external validity of our main results and suggesting the intervention's effectiveness across organizations. Due to the smaller size of the second collaborating partner (30 branches and a total of 200 bankers), we could only pre-register the replication of the main average effects of treatment versus control; heterogeneity analysis and mechanisms were not possible due to insufficient power. While these findings support the intervention's scalability, further studies in diverse industries and contexts are necessary to fully generalize these results.

Our paper relates closely to two strands of the literature. First, we contribute to the literature on knowledge flows within organizations. Prior theoretical and empirical research demonstrates that efficient information exchange and communication between peers positively affects productivity (Sandvik et al., 2020; Battiston et al., 2021; Guillouet et al., 2024). The literature has highlighted the role of hierarchy as a barrier to employee voice and cooperation in organizations (Auriol et al., 2002; Castro et al., 2022). It has examined barriers from the demand side—why employees may hesitate to seek help—and found that concerns about reputational damage are particularly powerful, creating frictions that discourage help-seeking (Mickeler et al., 2023; Heursen et al., 2024). Our study complements this research by focusing on the supply side—the willingness of employees to *contribute* their

knowledge to their *colleagues*.

Second, our work contributes to the literature on awards as influential social signals (Gallus & Frey, 2017; Gallus et al., 2023). Research has extensively discussed the power of different incentives to change behavior, and their unintended effects (Gneezy et al., 2011). We contribute not only by focusing on recognition as a particular type of incentive but also by identifying the specific channels through which it operate in the workplace. Additionally, we contribute to this literature by examining the effect of the incentives on cooperative behavior over time. Previous experimental research has studied cooperation and incentives in one-shot public goods games (Deversi et al., 2020), our study extends this analysis to a longitudinal workplace setting.

In Section 2, we detail the experimental design and empirical strategy. In Section 3, we present the results, and Section 4 discusses internal and external validity. Finally, Section 5 concludes.

2 Experimental Design

2.1 Setting and Sample

The study was conducted in collaboration with one of the largest commercial banks in East Africa. This firm has its headquarters in Kampala, Uganda. In 2023, the organization employed over 3,000 individuals across various sectors of commercial banking and served over 2.5 million customers, with assets of 6.3 trillion UGX (about 1.6 billion USD). The bank's core business is divided across two major divisions: banking and credit. Our study involved their credit division, which comprises over 1,000 bankers engaged in knowledge-intensive, non-routine lending tasks.

The organization operates across the country through branches. In the credit section, each branch is structured into three hierarchical levels: branch management, supervisors, and bankers. The bank employs approximately 80 branch managers, 200 supervisors, and 1,000 bankers.

At the top of the hierarchy are the *branch managers*. They are responsible for the strategic direction, performance oversight, and reporting of the branch. Their role includes managing all levels of staff, providing leadership to supervisory personnel, and reporting to Head Office. Their incentives are tied to the overall performance of the branch. At the next level are the *supervisors*. They are promoted from the ranks of bankers after at least six years of experience. Supervisors usually oversee a team of 5-6 bankers, and their core responsibilities are approving loan appraisals and coordinating their team's work. Their incentives depend on the performance of the bankers that they supervise. Supervisors rotate subordinates every 6-12 months.

At the lowest level of the hierarchy are the *bankers*, who are our study sample. They are the main field staff responsible for all client-facing work. Their tasks include appraising loan applications, presenting loans to the approval committee, monitoring the contractual use of funds, and ensuring repayment throughout the loan period. Approximately 60% of a banker's salary depends on variable performance and is tied to the number of loans disbursed in the current month within certain parameters of net performing ratio and volume of the portfolio. There are two main categories of bankers: *business bankers*, who handle large business loans and have a target of at least four disbursements per month, and *personal bankers*, whose loan targets vary between 20 to 40 loans per month depending on the subcategory. As part of its employee incentive structure, the bank offers a small financial reward each month to the most productive employee in each branch.

Of the 80 branches that the bank has across the country, 78 participated in our study.[3] We include all bankers in these 78 branches in our sample. Branches have on average 13 bankers. As of November of 2023, the bankers have been on average for 5 years at the organization and at the particular branch for 2.4 years, 39% are female. Table 1 presents summary statistics of our sample in column (1).

2.2 Evaluation Design

For this study, we collaborated with the bank's personnel department to collect primary data. Before the study began, employees were informed about the research collaboration under the broad theme of *personnel and productivity*. The department routinely gathers employee feedback through surveys, conducted either internally or externally, depending on the topic and area. In this case, employees were informed that survey responses would be collected externally by researchers via Qualtrics and that bank management would not have access to individual-level responses; only aggregated results would be generally shared unless explicitly stated otherwise.

We conducted three survey waves, timed one before and two after the intervention. The Chief Human Resources Officer (CHRO) distributed all surveys via email to employees. Employees could complete them on work computers or personal phones. Each survey took an average of 10 minutes to complete, and participants received 10,000 UGX (approximately USD 2.50) in phone credit as compensation. Figure 1 summarizes the timeline and the data collected at each stage.

The baseline survey followed a standard workplace climate design used in human resources, with the addition of the *peer cooperativeness ratings*. Following baseline data collection, we randomly assigned 38 bank branches to the treatment group and 38 to the control

[3]Two branches were excluded: one is located within a university campus, presenting unique operational characteristics, and the other was established after the study commenced.

Table 1: Sample Descriptive Statistics and Balance

	(1)	(2)	(3)	(4)
		Means		Difference
Variable	Full sample	Award Branches	Control Branches	Treat - Control
Share of Women	0.39	0.38	0.41	-0.04
	(0.49)	(0.49)	(0.49)	(0.03)
Tenure in Branch	2.40	2.38	2.44	-0.16
	(2.46)	(2.32)	(2.60)	(0.16)
Tenure in Firm	5.14	5.02	5.33	-0.37
	(3.57)	(3.69)	(3.45)	(0.23)
Survey metrics at baseline (0-100)				
Engagement	75.37	76.49	74.06	2.56*
	(22.66)	(22.81)	(22.66)	(1.43)
Turnover Intentions	25.70	24.99	26.47	-1.55
	(28.01)	(28.95)	(27.23)	(1.77)
Trust in Leadership	68.01	68.90	66.88	2.58
	(25.49)	(26.01)	(25.00)	(1.61)
Feeling Recognized	68.01	68.83	66.93	2.33
	(23.80)	(23.38)	(24.29)	(1.51)
Receive Assistance from Peers	74.12	74.60	73.38	0.78
	(23.31)	(23.43)	(23.36)	(1.48)
Productivity Metrics				
Portfolio Case Load	286.95	283.50	291.36	-10.37
	(306.45)	(296.65)	(315.54)	(19.37)
Number of Disbursements	21.28	22.25	20.57	1.67*
	(14.25)	(14.85)	(13.66)	(0.90)
Non Performing Ratio	0.03	0.03	0.02	0.00
	(0.06)	(0.06)	(0.05)	(0.00)
Observations	1,033	482	551	1,033

Notes. Summary statistics at baseline for the full sample of employees that participated in the workplace climate survey. Columns (1)-(3) report the means and standard deviations (in parentheses). Column (4) reports the difference in means by regressing the variable on treatment assignment and computing the difference along with the associated standard error. The variable *Share of Women* is the ratio of female to male employees. The variables *Tenure in Branch* and *Tenure in Organization* are measured in years. The survey metrics correspond to responses to the baseline survey, ranging from 0 to 100, with 100 representing the highest level of agreement. The *Portfolio Case Load* represents the number of loans handled by bankers as of November 2023. The *Number of Disbursements* refers to the number of loans disbursed in November 2023, and the *Non-Performing Ratio* indicates the share of the portfolio that is in late repayment as of November 2023. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

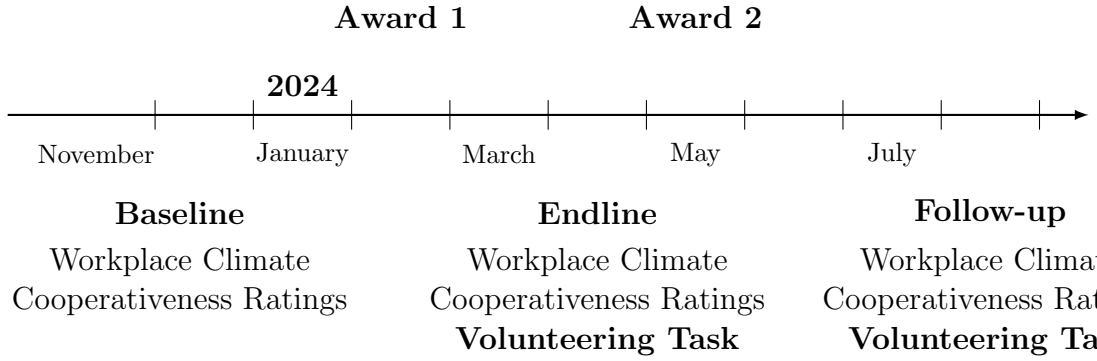

Figure 1: Timeline

group.[4] Table 1 shows that we achieved balance in observables. Column (2) shows average values for the treatment group, column (3) for the control, and column (4) tests the differences between both groups.

Between December 2023 and January 2024, branch management in the treatment group was briefed about the awards and given the names of the selected recipients to verify their appropriateness. No manager raised objections, and they were instructed to keep the names confidential until the award ceremony. The certificates were sent from headquarters to the treatment branches at the end of January, and employees received their awards in February 2024. This first round, which constitutes our main intervention, was presented as a surprise; employees were unaware beforehand that an award would be given. Therefore, the intervention was an unannounced, public recognition of employee cooperative behavior. Following the terminology of Bénabou & Tirole (2003, p. 504), unannounced public recognition corresponds to "discretionary" or "ex post" rewards, as opposed to "promised" or "ex ante" contingent incentives (e.g., innovation prizes).

In April 2024, the endline survey took place, followed by a second round of awards in May and a follow-up survey in August 2024. At the end of both the endline and follow-up surveys, employees were presented with *volunteering options* within the organization. These tasks, typically completed separately from personnel surveys, were timed concurrently with the workplace climate survey to allow matching responses to individual participants. Detailed descriptions of these tasks, which serve as primary outcome variables, along with other measures, are provided in Section 2.5. Each survey wave achieved a participation rate of 95%.

[4]Given the limited number of branches, we employed a matched-pair cluster randomization approach. Branches were paired based on key characteristics using an optimal greedy algorithm (King et al., 2007; Bruhn & McKenzie, 2009), which minimizes the Mahalanobis distance across pre-selected covariates. Matching relied on administrative and baseline survey data, including branch productivity, branch size, the share of business bankers, gender composition, and employees' perceptions of workplace recognition and peer support. Once pairs of similar branches were formed, one branch in each pair was randomly assigned to the treatment group and the other to the control group.

We pre-registered the first award as our main intervention and the volunteering task in the endline survey as our primary outcome. At the time of pre-registration, we were uncertain whether a second round of awards would be feasible. Additionally, while the first round was unexpected by employees, the second was not. As a result, we did not expect the second round to perfectly replicate the first but rather to be influenced by it. Nevertheless, examining the effects of a second award provides valuable insights into mid-term impacts and helps assess whether the surprise element drove any observed effects or if they persist over time.

2.3 Intervention: Most Supportive Employee Award

In this setting, cooperative behaviors are inherent to the nature of the job and are an integral part of daily work. Employees frequently seek help from peers for tasks such as analyzing complex credit proposals, recovering disbursements from non-compliant borrowers, formulating loan cases for committee adjudication, or conducting in-person appraisals of large loans. The organization actively encourages collaboration as part of its workplace culture. However, despite this cultural emphasis on support, cooperative behaviors fall outside the firm's incentive scheme. Incentives are tied to individual performance through variable pay structures, and a *Most Productive Employee of the Month* monetary reward.[5] As a result, while cooperation is both common and valuable, it is neither formally tracked nor explicitly incentivized or recognized.

The intervention introduces a direct and formalized incentive for supportive behaviors. To implement this, we first measured employees' cooperativeness using both peer and supervisor ratings. For peer ratings, we used two survey questions inside the HR workplace climate survey. First, we recorded employees' support networks by asking: *From which credit officers do you receive frequent help?* Employees could select multiple colleagues from the full list of bankers in their respective branches. If they selected more than one peer, they were then asked a follow-up question to identify which one stood out as the most supportive. So we will know for each employee in the organization how many times a peer has marked him/her as supportive and as uniquely supportive. To prevent gamification of the metric, supervisors also rated each employee under their direct supervision. They were asked to rank on a scale from 0 (very rarely) to 100 (very frequently) how often their subordinates engaged in supportive activities towards their peers that went *above and beyond*.

The above information was then aggregated, and a score was calculated for each employee. Peer ratings were normalized by dividing the absolute number of times an employee was mentioned as helpful and the most supportive by the total number of colleagues in the

[5]This reward includes a small monetary bonus of approximately 60,000 UGX (after tax), equivalent to around USD 16.00 at the time of the study, along with public recognition through the display of the employee's name and photo in the branch.

branch minus one, as voting for oneself was not allowed. Supervisor ratings were normalized between 0 and 1 by dividing them by 100. A weighted score was then calculated, assigning 2/3 of the weight to peer ratings and 1/3 to supervisor ratings. This resulted in a final rating for each banker in the organization.

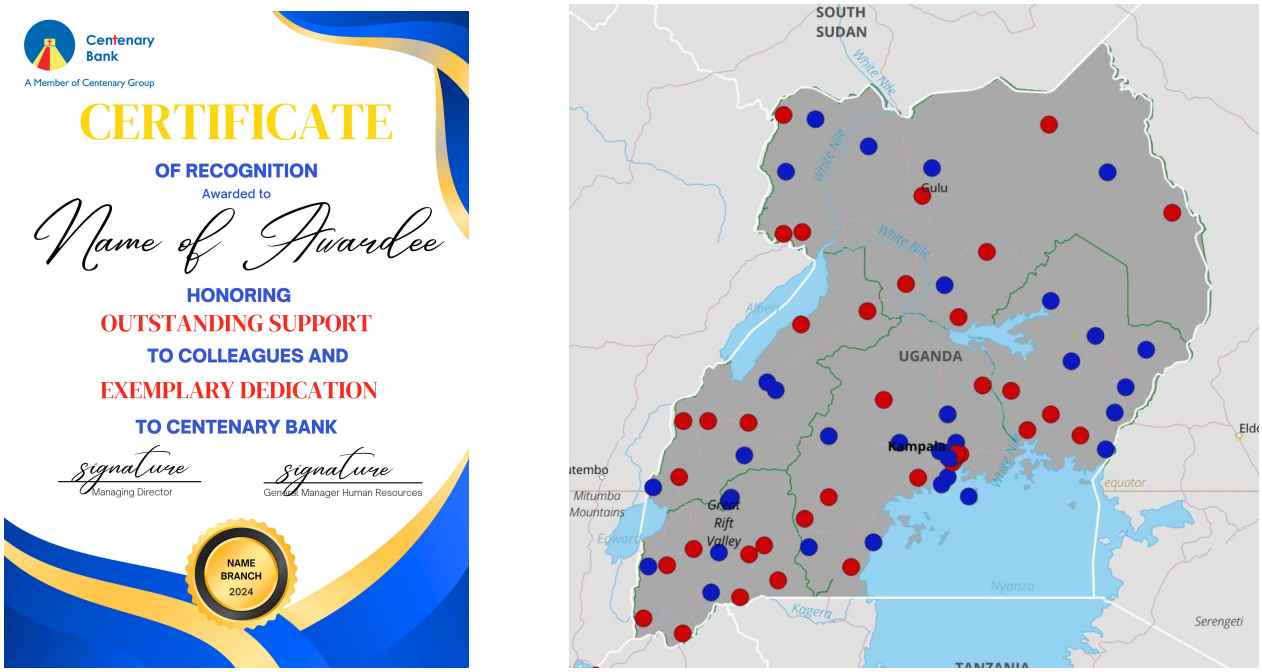

(a) (b)

Figure 2: (a) Template of Certificate. (b) Map of Intervention. Treated Branches in Red.

In the treatment branches, employees ranked in the top 20% in their branch, based on the aforementioned scoring criteria, were selected to receive the award.[6] The 20% threshold was determined in consultation with senior management to balance making the award scarce enough to be meaningful and attainable enough to motivate employees. Additionally, by fixing the proportion of winners at 20%, we ensured consistent intervention intensity across branches, regardless of their size.

Figure 2, part (a), shows a template of the certificate that the employees received. The certificates were hand-signed by the Chief Executive Officer of the bank and the Chief Human Resources Officer and framed. They explicitly mention *outstanding support* and *exemplary dedication* to peers and the firm. Part (b) of Figure 2 shows in red the geographic locations of the branches where the intervention took place. The awards were presented by branch managers during the February and May monthly meetings, which were attended by all branch employees. During the event, the managers emphasized the importance of cooperation and support and acknowledged how the award recipients had personally contributed to fostering

[6]In branches with less than five employees, we selected the employee with the highest score to ensure all treated branches have at least one awarded employee. Only three branches have less than five bankers.

a cooperative culture.

2.4 Conceptual Framework and Hypotheses

Our intervention—the introduction of a peer-nominated award—can influence cooperation through multiple channels. First, the certificate may directly affect behavior due to its intrinsic value or by enhancing intrinsic factors such as employee morale. Second, the award process might shape reputations among peers and influence social expectations around helping behaviors. Moreover, since the awards are signed by top organizational management and presented by branch management, they serve as a visible signal of the importance the organization places on cooperation. This signaling effect may reinforce employees' perceptions of managerial priorities and workplace norms, further shaping their willingness to help.

In this section, we develop a simple conceptual framework that formalizes these mechanisms and derives testable hypotheses to examine how the award influences employees' decisions to help.

Intrinsic Predisposition and the Cost of Helping An employee's decision to help[7] depends on intrinsic factors such as altruism, internalized norms, or personal identity, as well as the cost of providing help (e.g., time or effort). Let α_i represent an individual's intrinsic predisposition to help, e.g., the utility an individual gets from providing help to others. Additionally, let c denote the cost of helping. In the absence of external incentives, the utility derived from helping or cooperating (*h=1*) can be expressed as:

$$u(h = 1) = \alpha_i - c \tag{1}$$

The presence of the award could increase α_i by enhancing morale, reinforcing self-image, or increasing the perceived importance of helping.

Hypothesis 1 (Intrinsic Predisposition Mechanism): Employees are more likely to cooperate because the intervention increases their intrinsic predisposition to help.

To test this hypothesis, we need to compare helping decisions in treatment and control groups when the decision is observable only to the individual, in the absence of external incentives.

$$P(h = 1|\text{Award, Anonymous}) > P(h = 1|\text{No Award, Anonymous}) \tag{2}$$

[7]We use the terms help, cooperation, and support to peers interchangeably. We refer to actions that assist others in improving their productivity, even though these actions do not directly contribute to the productivity of the individual providing the help.

would confirm *Hypothesis 1*, where $P(h = 1)$ denotes the probability that an individual chooses to help under a given condition.

Besides intrinsic motivations, in workplace settings, employees' utility from helping may also be influenced by two important *external factors*: peer reputation and career concerns.

Peer Reputation Individuals care about how they are perceived by others, which can affect cooperation. In workplace interactions, employees may feel social pressure to conform to cooperative norms or seek to signal positive traits, such as being helpful, to their peers. Let R_p represent the social or reputational benefits of being seen as cooperative and adhering to these norms. The utility function can then be expressed as:

$$u(h = 1, \text{Peer Visibility}) = \alpha_i + R_p - c \tag{3}$$

where $R_p > 0$ if peer visibility increases the value of helping due to social expectations or reputational benefits. Introducing a public award may influence R_p by amplifying peer reputation concerns, making social expectations more explicit, or increasing the desirability of being seen as helpful.

Hypothesis 2 (Peer Reputation Mechanism): Employees are more likely to cooperate because the intervention increases the reputational value of helping among peers.

If the award heightens peer reputation concerns, we expect the effect of peer visibility on helping behavior to be stronger in award branches than in non-award branches. We test this by comparing helping behavior under peer visibility:

$$P(h = 1|\text{Award, Peer Visibility}) > P(h = 1|\text{No Award, Peer Visibility}) \tag{4}$$

To isolate the peer reputation mechanism, we compare the effect of the award under peer visibility (Equation 4) to its effect under anonymity (Equation 2). If the award affects intrinsic motivation, the incremental effect of peer visibility will be the difference between the two.

Career Concerns Employees may also consider the career implications of their helping behavior. If management values cooperation, employees might help to align with managerial expectations or to signal attributes that are beneficial for career advancement. Let R_m denote the career concerns effect, which captures the utility derived from being observed as helpful by senior management. The utility function becomes:

$$u(h = 1, \text{Mgmt Visibility}) = \alpha_i + R_m - c \tag{5}$$

where $R_m > 0$ if management visibility increases the career-related benefits of helping. Introducing a public award may heighten career concerns by signaling that management explicitly values and rewards cooperative behavior. Employees may thus view helping as a strategic action to improve their career prospects.

Hypothesis 3 (Career Concerns Mechanism): Employees are more likely to cooperate because the intervention increases the career value of helping.

If the award increases career-related incentives for helping, we expect the effect of management visibility on helping behavior to be stronger in award branches than in non-award branches. If this is the case, then:

$$P(h = 1|\text{Award, Mgmt Visibility}) > P(h = 1|\text{No Award, Mgmt Visibility}) \quad (6)$$

Similarly, to isolate the career concerns mechanism, we compare the effect of the award under management visibility (Equation 6) to its effect under anonymity (Equation 2). If the award affects intrinsic motivation, the incremental effect of management visibility will be the difference between the two.

If the decision to help is observed simultaneously by both management and peers, we will not be able to disentangle these two external mechanisms. To do so, we will have to compare helping choices in treatment and control branches where decisions are only visible to peers or only visible to managers.

2.5 Outcome Variables

Our primary conceptual outcome is collaboration. In particular, the supply side of help decisions. The ideal measure would be to observe real-time interactions where various opportunities to help arise and track employees' choices. However, this is challenging for researchers, as we cannot continuously observe all employees, and we need a standardized measure that allows for individual comparisons. Therefore, in consultation with the firm's leadership, we identified three key measures of cooperation that are both valuable to the organization and systematically measurable: two knowledge-sharing options that we will call *expert contribution I* and *II* and one enrollment option to become a *junior mentor*. In order to be able to systematically measure the decisions to help, these metrics were timed to concur with the endline and follow-up surveys so that they could pop up at the end of the surveys and be recorded. These provide us with revealed measures of cooperative behavior. In addition, we will examine self-reported survey measures that directly capture employees' perceptions of

help and cooperation within the organization.[8]

Revealed Choices: Knowledge Sharing and Mentoring

When employees complete the workplace climate surveys, an option appears on their screen informing them about a call for volunteers for a specific task. Employees read a brief description of the task and they can decide to sign up or decline. There is one call for volunteers after the endline survey and two at follow-up.

Expertise Contribution I. This is our primary pre-registered outcome, collected two months after the intervention. Employees are informed that the organization is compiling a repository of productivity-enhancing strategies to support bankers across the organization, particularly low performers. Human Resources will compile the contributions and share them by the end of 2024. Employees can briefly read that the task would entail providing successful strategies from professional experience, including general advice they find particularly useful, as well as specific guidance when struggling with borrower repayment. They are informed that this task will take approximately 10 minutes of their time.

Our primary outcome will be the *extensive margin*—whether employees select "yes" or "no" to contribute. If the respondent selects "yes" , then it opens a detailed description of specific current issues that underperformers face and the employee has an open box to introduce their approach to solving them. If an employee selects "no" the answer is recorded in our system and the form closes.

We have access to the individual contributions made and we will use them to examine the *intensive margin*, analyzing both the effort and quality of contributions. To assess *effort*, we will examine the time spent on the contribution and its length (measured in number of characters). To assess quality, we rely on expert evaluations conducted blind to treatment. We engaged two independent evaluators with relevant expertise in banking: one is a senior manager from the bank's head office with over 15 years of experience, and the other is a banking consultant with knowledge of commercial banking in Uganda but external to the bank. These experts assessed each contribution across three dimensions: (1) *helpfulness*—how beneficial the advice would be for other officers in the bank, (2) *innovativeness*—whether the contribution presents unique or creative insights, and (3) *actionability*—whether the advice provides clear and implementable steps. Both evaluator scores will be averaged over each contribution to generate a final quality assessment that is the mean average of both evaluations.

To disentangle the mechanisms laid out in subsection 2.4, we introduced a small variation when eliciting this response. It will allows us to decompose the average effect and examine its underlying drivers.

[8]We have also requested access to administrative records. Results from these records will be available 12 months post-intervention.

Mechanisms: We implemented individual-level randomization across treatment and control, varying one sentence in the survey that specified the visibility of the decision to help. Employees were informed about the visibility of their involvement based on the following conditions:

Anonymous: Only the employee knows that they have contributed ("*Your contributions will be anonymous, and no names will be included in the repository*").

Peer Visibility: Employees were informed that their names would be shared with colleagues in their branch ("*Your contributions will be anonymous, and no names will be included in the repository. However, we will circulate a list of contributor names to the bankers in the branch so that the efforts of those who shared can be recognized by their direct peers"*).

Management Visibility: In this condition, following the same structure as above, employees were told that their names would be shared with senior management (*"Your contributions will be anonymous, and no names will be included in the repository. However, we will share a list of contributor names with Head Office –HR and Microfinance teams- so that the efforts of those who shared can be recognized by management"*).

From an ethical perspective, while we aimed to examine mechanisms, we also wanted to avoid creating unfair advantages within the organization as a result of the study. For example, if the visibility of contributions were to provide a career benefit due to top management's awareness, employees in the *Anonymous* and *Peer Visibility* conditions would be at a disadvantage, as they would have contributed, but their names would not be known to management. A related and conflicting goal is to avoid deceiving respondents by not following through on promised behaviors. To address these concerns, we follow the strategy of Bursztyn et al. (2023). We inform respondents in the *Peer Visibility* and *Management Visibility* conditions that the list of names will be shared only if we receive at least one contribution from each branch in the bank. In practice, since at least one branch is unlikely to participate in the survey,[9] we ensure that the likelihood of the list being shared is very low. Due to this design, our outcomes should be interpreted as a behavioral choice that is conditional on respondents believing it sufficiently probable that all branches will be included and that at least one response will be received per branch.[10]

[9]One branch is being set up during the study period, and its systems are not yet fully operational, making it unlikely to participate. However, the survey email is sent across the entire credit division, so they still received the survey.

[10]It is possible that some respondents may believe that the list is unlikely to be shared. However, for this to bias our estimated treatment effects, it would require not only that this belief differs systematically across

Expertise Contribution II. This metric also captures employees' willingness to share their expertise but is measured at a later stage—during the follow-up survey—under a distinct setup and focus. Specifically, employees were invited to volunteer for a consultation with an external expert to provide feedback on challenges in *customer relationship management*. The consultant aimed to gather insights from frontline bankers and propose recommendations to top leadership for improving systems and addressing inefficiencies. As the bank has expanded significantly in recent years, some bankers now manage portfolios of more than 300 customers. To adapt to this growth, the organization is continuously refining its processes. This consultation task serves as a measure of willingness to cooperate by sharing knowledge, as it represents an additional responsibility beyond contractual obligations. Improved processes would not only benefit the organization but also directly support bankers, as their performance depends on the effective management of their portfolios. We observe only the *extrinsic margin decision*—whether employees select "yes" or "no" to sign up for the consultation; the actual meeting with the consultant is outside the scope of our research. Bankers were informed that the expected time commitment was 60 minutes.

Volunteering as a Junior Mentor. When a new employee joins a branch, an experienced banker is assigned as a second point of contact alongside the supervisor. The bank aims to pair new staff with experienced bankers during their initial months to provide guidance and support with integration into the workplace and community. Uganda has more than 32 languages and a diverse culture. Bankers must quickly adapt not only to a new job but also to new customs and, in some cases, a different language. Having a peer mentor during these first months can help facilitate this transition. At the end of the follow-up survey, the bankers saw a call for volunteers for this role. While junior mentors do not need to be volunteers—management can directly assign them—having a roster of volunteers who specify their availability and preferences is beneficial for the organization. Again, we observe the decision on the *extrinsic margin* - whether the respondent decided to sign up for the next round of volunteers. Typical official pairings last three months. Whether voluntary or assigned by management, mentoring is not formally monitored and falls outside the contractual obligations of the banker. However, once assigned, it is expected that the mentor will fulfill this role. It falls within the tasks of the branch management to take care of the final matching of mentors to juniors.

These two measures of cooperative behavior at endline do not allow for an examination of mechanisms, as they inherently involve some degree of visibility—either to management or peers. This reflects the common dynamics in organizations, where volunteering tasks, whether promotable or not, are typically observed by at least some colleagues. However,

treatment conditions but also that respondents who hold this belief are more or less likely to contribute in the visibility arms compared to the non-visibility condition.

the endline measure, *Expertise Contribution I*, provides a unique opportunity to disentangle mechanisms. Since this metric captures voluntary knowledge sharing and can be directly entered into the system under an anonymous option, it allows us to isolate the effects of the incentive independently of the visibility component, which might otherwise influence behavior, especially once we add the award for cooperation.

Workplace Personnel Survey: Additional Measurement

The measures described in the previous section capture employees' decisions to offer help, and while it is logical to expect that an increase in help given would result in more help received, they do not provide a direct measure of whether bankers actually *received* more assistance. To assess whether the intervention led to more received assistance, we leverage self-reported metrics from the workplace climate survey as suggestive evidence. This survey covered various HR-related topics such as leadership quality, role clarity, and organizational transparency, which, while relevant, were not the primary focus of this study. However, it included a key question on peer assistance: *"In the last month, other bankers in my branch assisted me with challenges, even if it was not part of their job description"*. This measure allows us to test whether the intervention increased reported help received from colleagues over time, as the question was included in baseline, endline, and follow-up surveys.

Additionally, the workplace climate survey conducted six months post-intervention included measures to assess the mechanisms discussed in Section 2.4. One question captures peer pressure to provide help, evaluating whether the award influenced perceived expectations to assist others. Respondents rated their agreement with: *"I find it difficult to decline a request for help even when I have other tasks that should take priority"*. Comparing responses between treatment and control groups will provide evidence on whether the introduction of the award increased perceived pressure to help.

To assess whether employees viewed the intervention favorably—an important consideration for organizations introducing new programs—we included a set of perception measures in treated branches. One such measure captures beliefs about career benefits: *"Receiving the award is beneficial for career advancement"*. Since this question was asked only in treatment branches, we cannot compare it to a control group, but it provides insight into whether employees perceive the intervention as linked to career progression. Additional perception questions evaluate the broader organizational implications of the award.

2.5.1 Robustness Checks: Spillovers

Since treated and control branches are located in close proximity in some regions, particularly in Kampala, and employees are occasionally transferred between branches, there is a possibility that those in control branches were exposed to information about the interven-

tion through informal communication, inter-branch interactions, or staff transfers. To assess potential spillover effects, we included a question in the endline and follow-up survey asking employees whether they had observed initiatives in other branches that they would like to adopt in their own workplace. Specifically, respondents were asked: *"Are there initiatives in other branches that you would like to adopt in your own branch? If so, which ones?"*

We will analyze the open-ended responses individually, searching for references to the intervention, particularly mentions of *awards* or *recognition*. If employees in control branches report awareness of the award, we will conduct robustness checks by excluding pairs of branches where spillovers are detected to ensure that our estimates are not biased by cross-branch contamination.

3 Results

3.1 Empirical Specification

To test the null hypothesis that the intervention has no effect on cooperation, we estimate the average treatment effect using the following specification of an Ordinary Least Squares regression:

$$y_{ib} = \beta_0 + \beta_1 T_b + \varepsilon_{ib}, \tag{7}$$

where y_{ib} represents the cooperation decision of employee i in branch b, taking a value of 1 if the employee cooperates and 0 otherwise. The variable T_b is a dummy equal to 1 if branch b is in the treatment group and 0 otherwise. ε_{ib} is the error term. β_1 captures the causal effect of the intervention on cooperation, representing the difference between employees in treated and control branches. Standard errors are clustered at the branch level.

For the intensive margin, the regression above will compare only those who have decided to contribute, so the dependent variable y_{ib} will represent the quality of the contribution, conditional on having chosen to contribute. Since baseline values for most primary variables are not available, we rely on a simple treatment-versus-control comparison.

3.2 Treatment Effect on Cooperation

Sign-ups to Help Table 2 presents the treatment effects on the probability of signing up for cooperative tasks across different volunteer options. Column (1) reports the pre-registered primary outcome: *expertise contribution I*, measured two months after the intervention. In the control group, 42.8% of employees contributed. This rate increased by 8.9 percentage points (a 21.2% increase) in treated branches, indicating that the intervention significantly enhances employees' willingness to engage in cooperative tasks (p-value < 0.05).

Columns (3) and (5) present results from cooperation measures collected six months after the intervention. The treatment effects remain consistent across these metrics. Sign-ups for the second expertise contribution are 24.8% higher (p-value < 0.01) in the treatment group compared to the control, while sign-ups for peer mentoring are 12.3% higher (p-value < 0.05). These findings suggest that the intervention induces a sustained increase in cooperation over time, corroborated by different measures across different periods.

Table 2: Effect of the Intervention on Volunteer Sign-Ups

	(1)	(2)	(3)	(4)	(5)	(6)
	Endline		*Follow-up*			
	Expertise	*Contribution I*	*Expertise*	*Contribution II*	*Volunteer*	*Mentor*
Award Branch	0.089**	0.096**	0.118***	0.118***	0.066**	0.066*
	(0.03)	(0.04)	(0.03)	(0.04)	(0.03)	(0.03)
Most Supportive R1		0.113**				
		(0.06)				
Award Branch × Most Supportive R1		-0.044				
		(0.07)				
Most Supportive R2				0.044		0.094
				(0.05)		(0.06)
Award Branch × Most Supportive R2				0.001		0.002
				(0.08)		(0.08)
Mean Dep. Var	.428		.475		.534	
Observations	908	908	908	908	908	908

Notes: The table presents linear probability model (OLS) estimates of the treatment effect on sign-ups for different cooperation tasks: willingness to provide expertise to the repository (Columns (1) and (2)), willingness to meet with a consultant to improve customer management processes (Columns (3) and (4)), and willingness to sign up as a junior mentor (Columns (5) and (6)). The dependent variable is binary, coded as 1 if the person chose to sign up and 0 otherwise. Columns (2), (4), and (6) include an interaction with being selected as the most supportive employee in the previous survey. If a person was selected and was in the award branch, they received the award. Standard errors are clustered at the branch level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Interestingly, the treatment effect is not driven exclusively by employees who received the award. Column (2) introduces an interaction term between treatment assignment and being selected as the "Most Supportive Employee". Employees identified as the most supportive in the control group serve as a counterfactual for those in the treatment group who received the award. Across Columns (2), (4), and (6), the coefficient on treatment assignment remains large and significant, while the interaction term is consistently insignificant and close to zero. This indicates that receiving the award is not the driver of the treatment effect. Instead, the intervention also appears to motivate the broader employee base—those who participated in the voting process and observed the award—leading to increased willingness to help among

non-recipients in the treatment group.

Additionally, Column (2) reveals that employees selected as the most supportive are more likely to sign up to contribute, with a large and significant coefficient (p-value < 0.05). This suggests that the voting score effectively identifies employees predisposed to helping. In additional analyses, we further test the correlations between supervisor ratings and peer ratings with the likelihood of contributing in the control group (see Appendix Table A4).[11] We observe that peer ratings are a strong predictor of singing up to the cooperation task. Employees with higher peer ratings at baseline showed a significantly stronger likelihood of signing up to contribute (p-value < 0.01). Specifically, a 10% increase in the number of peers selecting an individual as frequently providing support raised that individual's likelihood of volunteering to share knowledge by 2%. In contrast, supervisor rankings showed no correlation with actual helping behavior in this task (p-value $= 0.71$).[12].

The results in Table 2 are based on the most restrictive sample: employees who participated in both the endline and follow-up surveys and were not transferred between treatment and control branches [13] and without using any controls. The effects and coefficients remain unchanged when including pair-matched fixed effects (Table A1), incorporating the full sample (Table A2), and under alternative specifications such as logit estimation (Table A3). Furthermore, the results are robust to replication in a different organization (see Section 4).

Effort and Quality of Contributions We examine whether the observed increase in contributions affects the effort and quality of the knowledge shared. A potential concern is that while more employees opt to contribute, the additional contributions might be of lower quality if the new contributors possess less expertise. Table 3 presents the results. Columns (1) and (2) examine effort, measured by the length of contributions (in characters) and the time spent composing them (in seconds). The estimates indicate no significant difference in contribution length between treatment and control groups. However, treated contributors spent 25% more time composing their responses, corresponding to an increase of 145.7 seconds (2.4 minutes). This effect is not statistically significant (p-value $= 0.13$). Note that due to a coding error, contribution times were recorded only from the second day of data collection, leading to a reduced sample size in Column (2).

Columns (3)–(5) assess contribution quality based on expert evaluations blind to treatment. The quality measures range from 0 to 6, with higher scores indicating better quality. The coefficients across all quality dimensions—$helpfulness$, $actionability$, and $innovativeness$—are

[11] We check only for the control group as their behavior was not affected by the intervention.

[12] We observe similar trends in the subsequent survey wave. Supervisor ratings remain a poor predictor of signing up to volunteer as a mentor (p-value $= 0.62$). Peer ratings continue to perform better, with coefficients of similar magnitude to those at baseline, though they are not significant (p-value $= 0.15$)

[13] Employees are typically transferred after three years in a branch; approximately 6% of the sample was transferred during the study period.

Table 3: Intervention Impact on Effort and Quality of Contributions

	(1)	(2)	(3)	(4)	(5)
	Effort		*Expert Evaluation of Quality*		
	Length	*Time*	*Helpful*	*Actionable*	*Innovative*
Award Branch	22.314	145.726	0.149*	0.100	0.087**
	(45.89)	(99.51)	(0.08)	(0.07)	(0.04)
Mean Dep. Var	413	586	2.62	2.31	1.61
Observations	407	277	407	407	407

Notes: The table presents OLS estimates of the effect of treatment assignment on different intrinsic margin measures. Columns (1) and (2) show the effect of treatment on response length (measured in characters) and response time (measured in seconds). Column (2) has a reduced number of observations as the timer only started counting on the second day of data collection. Columns (3)–(5) present results for quality measures, evaluated by expert evaluators blind to treatment. Quality scores range from 0 to 6, with 6 indicating the highest quality. All expert-evaluated scores are averaged across evaluators. Standard errors are clustered at the branch level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

positive, though not always statistically significant. Notably, the intervention significantly increased the innovativeness of contributions (Column 5), with treated contributors scoring 5.4% higher than their control counterparts.

Overall, these findings suggest that the increased participation in knowledge sharing does not come at the expense of quality. If anything, contributions in the treatment group tend to be more innovative, implying that expanding participation can enhance the variety and originality of shared knowledge without diluting its value. However, we do not have intrinsic margin measures at follow-up to examine whether these effects persist over time.

Self-Reported Help We now examine whether employees report receiving more help as a result of the intervention. Figure 3 presents the estimated impact on self-reported perceptions of receiving assistance from peers. Compared to the control group, employees in treated branches reported a 0.14 standard deviation statistically significant increase (p-value < 0.10) in agreement with having received help from colleagues in the previous month at endline (coefficient from Award Branch # Endline). This effect further increases to 0.17 standard deviations (p-value < 0.05) six months after the intervention (coefficient from Award Branch # Follow-up). These results prove that the intervention not only encouraged employees to offer assistance but also translated into an increase in the self-reported reception of peer support over time.

These results complement the primary findings and provide suggestive evidence that the intervention is associated with an increase in reported peer assistance. While the self-reported nature of the measure introduces potential reporting biases, the consistency of the effects over time suggests a sustained perception of greater support among employees in

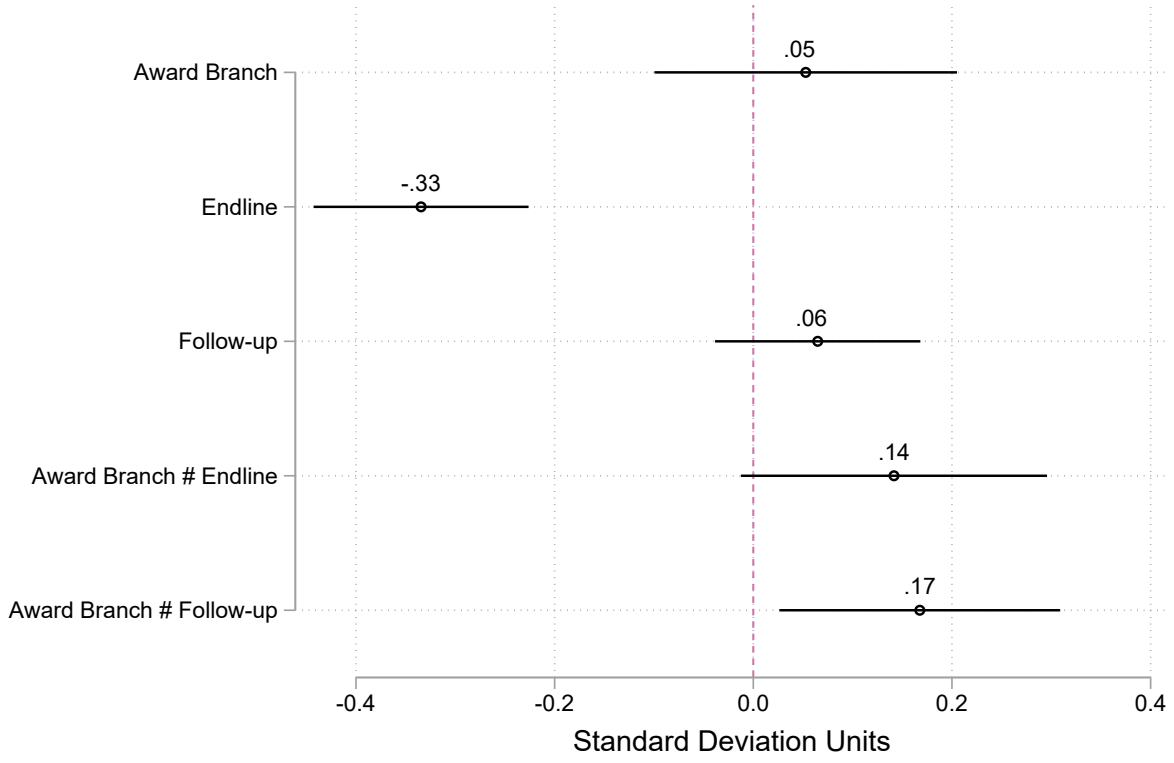

Figure 3: Impact of the Intervention on Perceived Frequency of Peer Help

Notes: This figure presents coefficient estimates from a linear regression analyzing the impact of being in a treated branch compared to a control branch on self-reported perceptions of receiving peer assistance over time. Employees rated their level of agreement on a scale from 0 to 100 with the statement: *"In the last month, other bankers in my branch assisted me with challenges, even if it was not part of their job description."* The dependent variable has been standardized (demeaned and divided by its standard deviation), so coefficients are interpreted in standard deviation units. Confidence intervals are at the 95% level. Standard errors are clustered at the branch level.

treated branches. Although there are several reasons why employees might indicate that they "provide more help" after the introduction of the award, self-reports of having *received* more help are likely subject to less bias. This is because individuals have fewer incentives to overstate the support they receive from colleagues, as doing so does not enhance their self-image. If anything, research suggests that individuals have reputational incentives to avoid being seen as someone who needs help (Heursen et al., 2024). This finding aligns with the observed increase in help offered and reinforces the idea that recognition-based incentives may shape workplace cooperation.

3.3 Mechanisms: What Drives Employees to Cooperate More?

Having established that the intervention increased cooperation on average, we now examine the underlying mechanisms driving this effect. Does the intervention alter individual pre-

disposition to help, does it change the pressure or expectations to cooperate among peers, or does it signal from management that cooperation is valued, leading employees to behave strategically for career gains? Figure 4 disaggregates the decision to help in the *Expertise Contribution I* under different levels of visibility.

The first hypothesis is that the award could alter *individual predisposition* to help. To test this, we compare the percentage of employees who choose to help in the treatment branches to those in the control branches when their decision remains private. As shown in the leftmost bars of Figure 4, there is no statistically significant difference between the treatment and control groups under anonymity (p-value = 0.82). Around 43% of employees contribute in the control group, compared to 44% in the treatment group. Therefore, we reject the hypothesis that $P(h = 1|\text{Award, Anonymous}) > P(h = 1|\text{No Award, Anonymous})$ and conclude that the intervention does not affect intrinsic motivation or purely altruistic behavior. If an organization seeks to increase support for public goods without additional incentives, the award alone may be insufficient.

Second, we examine whether the intervention influences *peer reputation*. If this mechanism is at play, contributions under peer visibility should differ between treatment and control groups. As shown in the two bars in the center of Figure 4, contributions in treated branches increase by 15% (p-value < 0.05) under peer visibility, a statistically significant difference. This confirms our second hypothesis that peer reputation drives cooperation ($P(h = 1|\text{Award, Peer Visibility}) > P(h = 1|\text{No Award, Peer Visibility})$). While we cannot pinpoint which specific dimension of professional reputation is at play, this effect could reflect employees feeling more obligated to help or becoming more aware that their peers notice and value this behavior. Self-reported survey data further suggest that expectations around helping have shifted.

Figure 5 presents the comparison between treatment and control groups regarding the statement: *"I find it difficult to decline a request for help even when I have other tasks that should take priority"*. Six months after the intervention, employees in the treatment group reported a 0.33 standard deviation increase in difficulty declining help requests compared to the control group. This effect is even stronger for award recipients, who report a 0.64 standard deviation increase relative to their counterfactuals—employees in the control group identified as the most supportive.[14]

Overall, these results suggest that the intervention directly influences peer interactions by shifting norms around cooperation. However, this is not the only way the intervention could affect cooperation through this channel. Anecdotal evidence from post-study qualitative interviews with awardees suggests that the award increases the attractiveness and awareness of being seen as a supportive colleague. Helping behaviors in professional settings often con-

[14]The wider confidence intervals for this group reflect the smaller sample size, as it comprises only the top 20% of supportive employees.

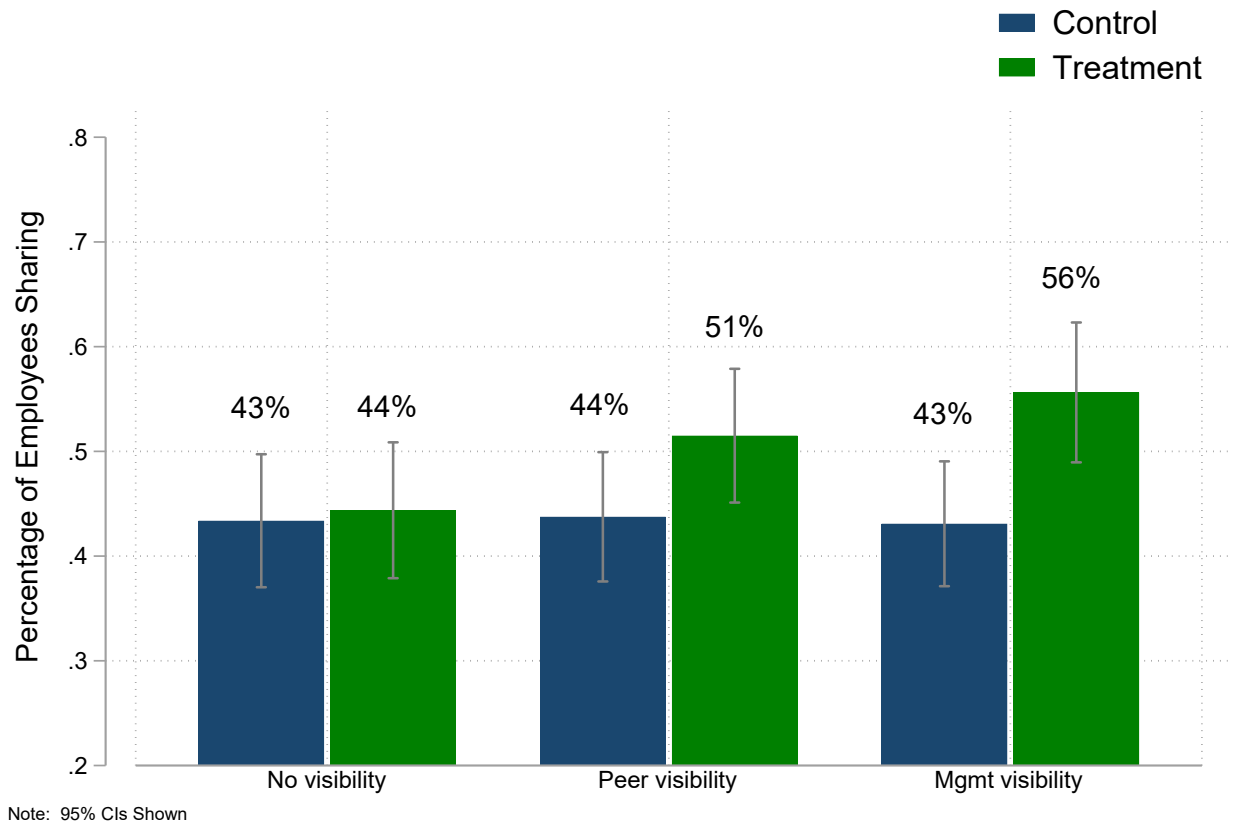

Figure 4: Sign-ups to Expertise Contribution I under Visibility Conditions

Notes: This figure presents the percentage of employees who signed up to contribute to the *Expertise Contribution I* task, disaggregated by visibility condition. The leftmost bars represent the condition where contributions were only known to the individual. The two bars in the middle correspond to conditions where contributions might be known to peers, while the two bars on the right correspond to conditions where contributions were visible to management. The green bars indicate treatment branches where an award was introduced, while the control branches are shown in blue.

sist of minor yet meaningful gestures—answering a colleague's question, sharing advice, or offering quick support—that accumulate over time. Without formal recognition, these acts may lack visibility, leading employees to underestimate their significance or doubt whether peers notice and value them. The award may serve as a signal that such contributions are observed and appreciated, reinforcing the perception of who consistently engages in supportive behavior. For example, one awardee remarked:

> *"I was very surprised when I received the award. I've always helped others because that's how I was raised, and I know I'm a supportive person. What surprised me most was learning that so many people in the branch noticed my efforts and saw me as someone who stands out in supporting others."*

Lastly, an alternative explanation is that overall cooperation did not increase, but rather that employees strategically participated in this particular task to signal their helpfulness

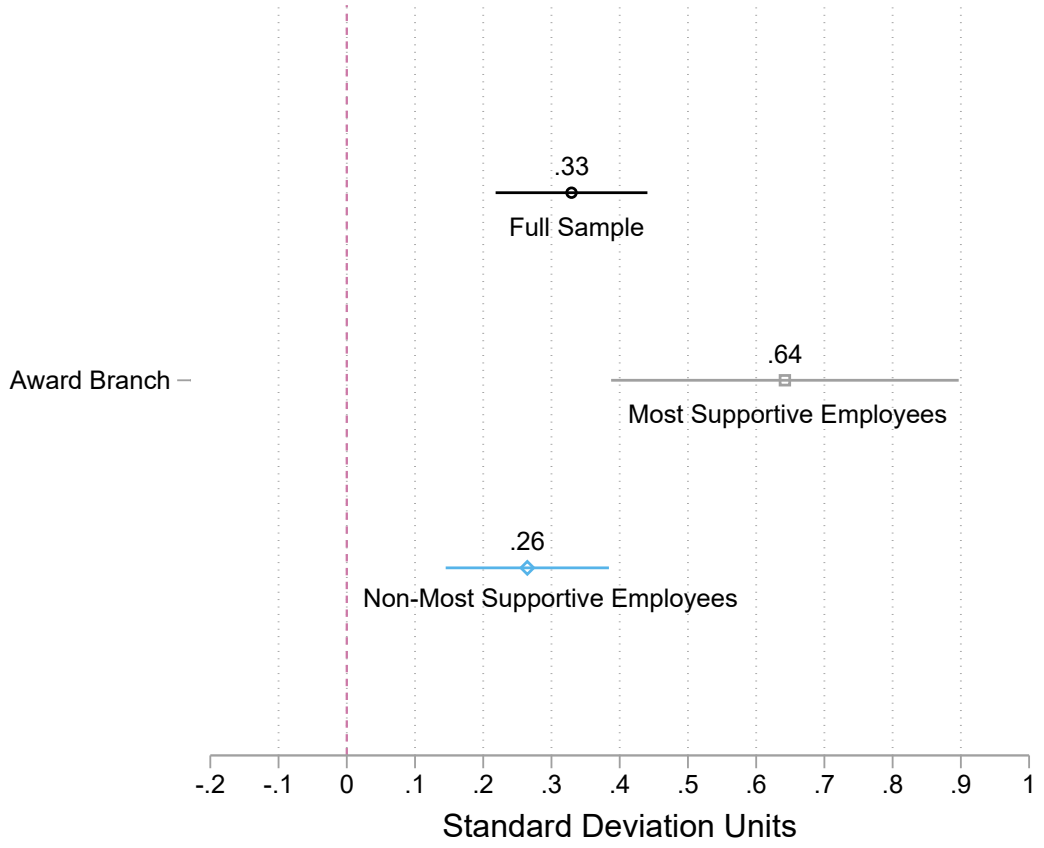

Figure 5: Impact of the Intervention on Perceived Difficulty of Declining Help

Notes: This figure presents coefficient estimates from linear regressions analyzing the impact of the intervention on employees' self-reported difficulty in declining help requests. The dependent variable is based on responses to the statement: *"I find it difficult to decline a request for help even when I have other tasks that should take priority."* The analysis compares treated branches to control branches, with results reported for the full sample as well as separately for *Most Supportive Employees* (top 20% in supportiveness) and *Non-Most Supportive Employees* (remaining 80%). The dependent variable has been standardized (demeaned and divided by its standard deviation), so coefficients are interpreted in standard deviation units. Confidence intervals are at the 95% level, and standard errors are clustered at the branch level.

and improve their chances of winning the award in the next round. Under this scenario, the observed increase in contributions would not reflect a genuine shift in cooperative behavior within the organization but rather a temporary response driven by award incentives. However, strategic voting is unlikely to account for this effect, as the call for contributors occurred after the voting for the next round of awards had already taken place, and the names of those who contributed were scheduled to be shared nine months later. While we cannot completely rule out this channel, we consider it the least plausible explanation.

The third and final hypothesis is that the intervention affects behavior by influencing *career concerns*. By introducing the award, management sends a clear and direct signal that helping others is valued. This may motivate employees to align their behavior with

organizational expectations, either to enhance their career prospects or simply because they have become more aware that such behavior is rewarded. When comparing the two rightmost bars in Figure 4, we observe a 30% increase in cooperation under management visibility (p-value < 0.01). This finding supports the hypothesis that career concerns play a role in the observed behavioral change and therefore confirms the third hypothesis that $P(h = 1|\text{Award, Mgmt Visibility}) > P(h = 1|\text{No Award, Mgmt Visibility})$. Figure 6 further examines whether employees perceive career-related benefits from receiving the certificate. Two and six months after the intervention, employees in treated branches were asked about their perceptions of the award. Figure 6 presents results at two months (Appendix Figure A1 shows similar results at six months). On a scale from 0 to 100, where 100 represents the highest level of agreement, the average agreement level was 82 for the statement that winning the award has positive career implications. If employees believe that a behavior is valued by management, it is plausible that they will seek to comply with it.

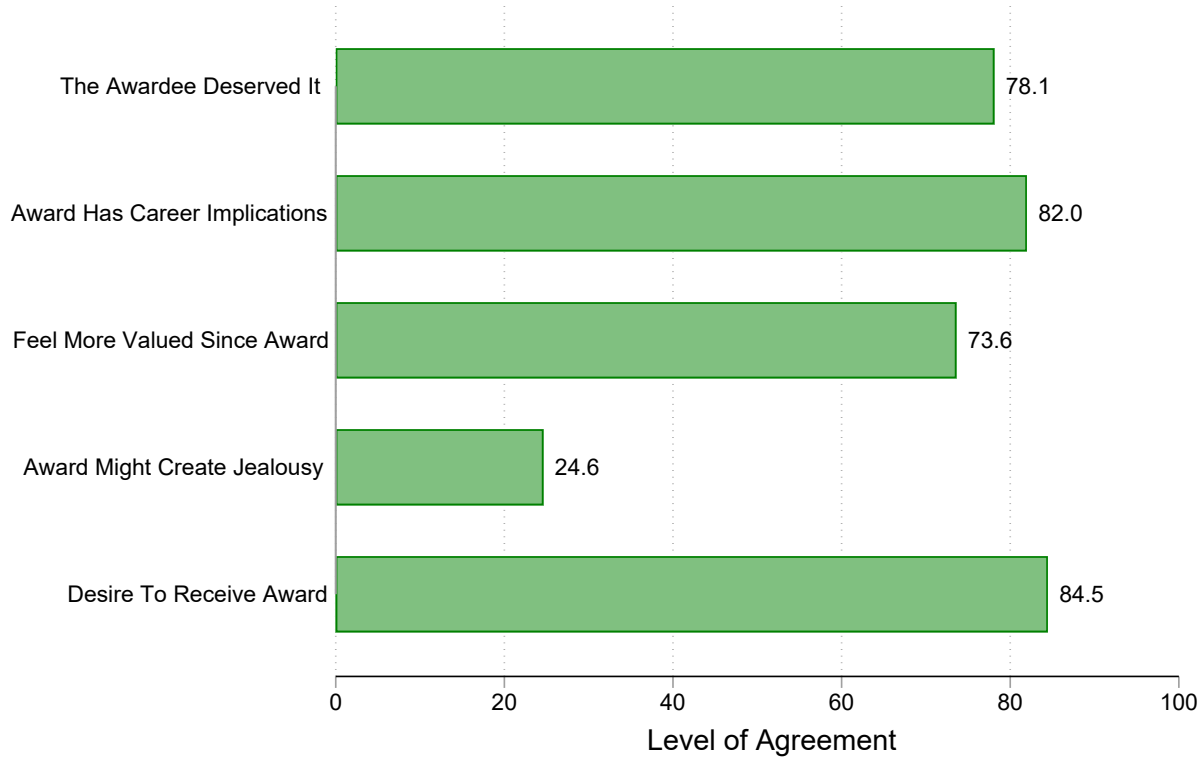

Figure 6: Perceptions of the Intervention

Notes: This figure presents employees' perceptions of the award six months after the intervention. Employees rated their agreement with various statements on a scale from 0 to 100, where 100 represents the highest level of agreement.

3.4 Employee Perceptions of the Intervention

To assess the award's value from the organization's perspective, it is worthwhile to consider employees' direct opinions about it. Understanding how the award is perceived can provide insight into whether it serves as a meaningful incentive, fosters positive workplace dynamics, or creates unintended negative consequences, such as resentment or competition.

To evaluate this, Figure 6 presents the level of agreement with several statements about the award, where employees rated their agreement on a scale from 0 (completely disagree) to 100 (completely agree). Employees in treatment branches overwhelmingly report positive perceptions of the award. Specifically, employees express an average agreement level of 84.5 with the statement that they would like to receive the award at some point, suggesting that the incentive is seen as highly desirable. Additionally, employees report an average agreement level of 73.6 with the statement that the award makes them feel more valued in their role, reinforcing the idea that public recognition enhances employees' sense of appreciation and organizational belonging. The award is also perceived as having professional value: employees report an agreement level of 82 when asked whether the award has career implications. This suggests that employees believe the recognition carries weight beyond symbolic acknowledgment and may influence professional development or advancement opportunities. Furthermore, employees reported an average agreement level of 78.1 with the statement that the awardee truly deserved the recognition, suggesting that the selection process was perceived as fair and legitimate.

A potential concern with workplace awards is that they might generate competition, resentment, or jealousy among employees. However, the average level of agreement with the statement that the award creates jealousy within the branch is only 24.6 out of 100. This suggests the award does not lead to negative interpersonal dynamics or workplace conflict. Instead, the responses imply that employees see the award as a constructive and meritocratic mechanism rather than a divisive or exclusionary one.

Taken together, these findings suggest that the award is broadly well-received and is viewed as both a desirable and meaningful form of recognition. Employees believe that it increases their sense of being valued at work, has professional relevance, and is awarded fairly, while concerns about jealousy or resentment appear to be minor.

4 External Validity and Replication

To assess the robustness of our findings, it is important to consider potential limitations arising from both our study's internal and external validity.

From an internal validity perspective, a concern when designing the study was potential spillovers between branches. Given the proximity of some treated and control branches, par-

ticularly in Kampala, and occasional staff transfers, employees in control branches may have been indirectly exposed to the intervention through informal communication or inter-branch interactions.

To assess potential spillover effects, we included a question in the endline and follow-up surveys asking employees whether they had observed initiatives in other branches that they would like to adopt in their own workplace. Analyzing over 200 open-ended responses from control branches, we found no direct references to the intervention. This is likely due to the broader organizational context in which the study took place. At the time of the intervention, the company was undergoing significant growth and cultural transformation, with multiple concurrent initiatives aimed at reducing hierarchical distance and improving communication. Consequently, the award was perceived as one of many human resource initiatives rather than a distinct program, making it less likely to be singled out and mentioned to other branches. The most frequently cited initiatives employees wished to adopt were unrelated to the intervention. Employees often referenced activities available in a previous branch but no longer present after transfer, such as team-building activities and management strategies like group recoveries. Given that the intervention lasted from February to August, the extent of inter-branch interactions and spillovers appears insufficient to compromise our results. Overall, these findings suggest that the stable unit treatment value assumption was preserved, and spillover effects are unlikely to threaten internal validity.

From the point of view of external validity, an active debate has arisen around the credibility and replicability of empirical findings in various scientific fields (Camerer et al., 2018; Baker, 2016), including psychology, management, and the biological sciences. Concerns that numerous published results may constitute false positives have led some scholars to describe the situation as a "credibility crisis" for modern science. Replication allows researchers to help in mitigating these issues. When multiple studies consistently corroborate an original finding—even under different contexts or with new samples—the claim that the result holds in a broader set of conditions is substantially strengthened.

In organizational economics and management, field experiments are typically conducted in partnership with specific organizations (Levitt & List, 2009; Chatterji et al., 2016). As a result, questions often arise about the degree to which their findings can be generalized—or "externally validated"—to other contexts. Even well-powered field experiments with robust designs can face challenges when tested elsewhere, as organizational culture, size, geographic location, and workforce characteristics could all shape outcomes differently. Replication, therefore, is essential not only for academic rigor but also to help decision-makers in firms determine whether a given research result is replicable in their specific setting.

In light of the concerns about generalizability, we conducted a replication study[15] in a

[15]The pre-registration for this study can be found in AsPredicted #153656. The study was pre-registered separately from the main study to ensure they were treated as two distinct studies. The replication, as

second financial institution to determine whether our main results would hold under different conditions. We deliberately chose an organization in the same industry—a bank employing knowledge workers—but with distinct characteristics from our original partner. Specifically, this second entity is an international corporate group headquartered in the United States rather than a local bank, operating at roughly one-tenth the size (30 branches) of the initial institution. Although employees' daily tasks—assessing borrowers and managing funding—are broadly similar, these tasks are allocated differently among staff, and while incentives remain largely individualized, peer support continues to play a significant role. Consequently, this smaller corporate environment provides a meaningful test of whether the results observed at the larger, more centrally managed bank can be replicated in an alternative institutional setting.

The replication followed the same "Most Supportive Employee" award procedure as in our main study. Specifically, among the 30 branches, half were randomly assigned to receive the award treatment using pair-matching randomization, while the others remained in the control group. Because each branch had only around five employees, one individual per branch was recognized as the top supporter in both rounds of awards, which took place during the first two quarters of 2024—mirroring the timeline in our primary study. As in the main experiment, employees answered a general workplace climate survey and identified their "Most Supportive Peers". After the second round of awards, they were asked to share their knowledge with their colleagues [16] Employees were informed that both their peers and the head office would be aware of their contributions. However, due to logistical constraints in this smaller organization, the key questions related to our mechanisms were only asked at the second round. Of the 215 officers invited, 204 completed the survey. To preserve confidentiality, all communication was handled by the company's marketing and communications team, rather than through the human resources department.

Table 4 presents the primary results of the replication study. Despite the smaller scale of the replication, the findings are consistent with those from the main study, reinforcing the robustness of our original results. Specifically, we observe a 5.7 percentage point increase in employees' willingness to invest additional time in helping colleagues in treated branches compared to the control group. This effect is statistically significant at the 5% level. Standard errors are clustered at the branch level, and pair-matched fixed effects are included to address

detailed in the pre-registration, was limited to only one treatment condition—awarded versus non-awarded branches—due to power constraints.

[16]Note that we administered the knowledge-sharing question after the second round of awards—rather than the first, as in the other organization, due to differences in survey timing across the two organizations, which were dictated by internal departmental schedules. In the smaller organization, the survey was conducted two weeks earlier. Still, the final phrasing of this key outcome metric had not yet been approved by top management in the original organization. At that time, researchers only had a draft version of the phrasing. Since the core goal of the replication was to re-run the exact same metrics with minor contextual variations, we decided to include the knowledge-sharing question after the two award runs.

the small cluster size.

Interestingly, the baseline level of willingness to help in the control group, measured as the proportion of respondents signing up to contribute was 43%. This pre-existing level mirrors the control group result in the original, larger-scale study, suggesting a striking consistency across organizational contexts. The magnitude of the treatment effect in the replication is also comparable to that documented in the main study.

Table 4: Willingness to Share - Replication

	(1) Willingness to Share	(2) Knowledge
Award Branch	.057**	.111***
	(.03)	(.04)
Most Supportive		.215
		(.19)
Award Branch × Most Supportive		-.474*
		(.24)
Mean Dep. Var	.433	.405
Observations	204	204

Note: Linear probability model (OLS) estimates the effect of the treatment on the willingness to provide productivity tips. Column (1) includes the main average effect and Column (2) adds the interaction with being selected in the previous round as most supportive. Regressions control for demographics, including tenure at the branch and firm, as well as gender. Both models include pair-matched fixed effects. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

In Column (2), we observe that individuals previously selected as "Most Supportive" exhibit higher levels of helping behavior, although the coefficient is statistically insignificant. The positive and substantial magnitude of the coefficient suggests that the nomination metric effectively captures a propensity for helping behavior. Furthermore, the treatment effect remains significant and increases in magnitude when an interaction term for receiving an award is included, consistent with the main study. The interaction term is slightly negative, indicating that the observed changes in the awarded branches are not driven by the recipients of the award. Instead, the results suggest that the introduction of the award within the branch fosters a broader behavioral shift among employees, underscoring the treatment's overall positive effect.[17]

These replication findings offer two important lessons. First, for organizations, the results suggest that interventions introducing publicly recognized awards for supportive behaviors can reliably produce an effect size of approximately 13% in employees' willingness to help. This enables organizations to assess whether this effect size justifies the intervention's costs.

[17]Due to the smaller sample size of 200 employees, the study lacks sufficient power to replicate and examine any underlying mechanisms fully.

The consistency of the results, even in an institution with a different culture and size, underscores the robustness of the intervention's impact and its relevance for companies in knowledge-intensive sectors.

Second, for academics, the replication strengthens confidence in the main effect of the intervention. Replication serves as a powerful tool to reduce the risk of false positives, providing greater reliability in results. However, further testing in organizations from different industries or countries is needed to understand how the intervention might perform in varying contexts. Naturally, research findings are always context-specific and depend on the unique characteristics of the setting.

5 Conclusion and Implications for Organizations

This study provides causal evidence on how social recognition incentivizes cooperation and tacit knowledge exchange in organizations. We show that a simple, low-cost intervention—a public award for peer support—significantly increases employees' willingness to help colleagues. The recognition incentive led to a sustained rise in voluntary expertise contributions and peer mentoring, while also improving the quality of help provided. Moreover, our findings suggest that workplace cooperation was driven by peer reputation effects and managerial signaling rather than intrinsic motivation. Notably, this study empirically unpacks the complex interpersonal dynamics of tacit knowledge exchange, demonstrating how both social image—shaped by peers—and management-related reputation and career concerns act as strong incentives for cooperative behavior.

This research contributes empirical evidence to the literature on workplace cooperation and knowledge flows by providing evidence on the supply-side barriers to helping behavior and the mechanisms through which social incentives shape cooperative norms. Additionally, our study adds to the growing body of research on awards and recognition, highlighting their potential as scalable tools for fostering collaboration in professional settings.

Future research could explore the long-term effects of social recognition on workplace dynamics, particularly whether repeated incentives sustain or diminish cooperative behaviors over time. Additionally, while our study focuses on financial institutions, testing similar interventions in other industries—such as healthcare, education, or technology—would further our understanding of how cooperation can be systematically incentivized in different organizational contexts.

Our findings suggest that organizations seeking to enhance cooperation can benefit from simple, well-designed recognition programs that acknowledge and reinforce prosocial behavior. By aligning incentives with the firm's broader goals, such initiatives can help bridge the gap between individual and organizational objectives, ultimately fostering a more collaborative and productive work environment.

6 Appendix

Table A1: Sign-Ups for Cooperative Tasks

	(1)	(2)	(3)	(4)	(5)	(6)
	Endline			*Follow-up*		
	Expertise	*Contribution I*	*Expertise*	*Contribution II*	*Volunteer*	*Mentor*
Award Branch	0.086***	0.089***	0.112***	0.112***	0.068**	0.067**
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Most Supportive R1		0.101*				
		(0.05)				
Award Branch × Most Supportive R1		-0.026				
		(0.07)				
Most Supportive R2				0.041		0.088
				(0.05)		(0.06)
Award Branch × Most Supportive R2				0.002		0.006
				(0.08)		(0.08)
Mean Dep. Var	.418		.475		.534	
Branch FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	908	908	908	908	908	908

Notes: The table presents linear probability model (OLS) estimates of the treatment effect on sign-ups for different cooperation tasks: willingness to provide expertise to the repository (Columns (1) and (2)), willingness to meet with a consultant to improve customer management processes (Columns (3) and (4)), and willingness to sign up as a junior mentor (Columns (5) and (6)). The dependent variable is binary, coded as 1 if the person chose to sign up and 0 otherwise. Columns (2), (4), and (6) include an interaction with being selected as the most supportive employee in the previous survey. If a person was selected and was in the award branch, they received the award. All models include pair-matched fixed effects. Standard errors are clustered at the branch level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2: Sign-Ups for Cooperative Tasks

	(1)	(2)	(3)	(4)	(5)	(6)
	Endline			*Follow-up*		
	Expertise	*Contribution I*	*Expertise*	*Contribution II*	*Volunteer*	*Mentor*
Award Branch	0.072**	0.073**	0.124***	0.123***	0.069**	0.061**
	(0.03)	(0.04)	(0.03)	(0.04)	(0.03)	(0.03)
Most Supportive R1		0.068				
		(0.05)				
Award Branch × Most Supportive R1		-0.012				
		(0.07)				
Most Supportive R2				0.043		0.083
				(0.05)		(0.06)
Award Branch × Most Supportive R2				0.005		0.012
				(0.08)		(0.08)
Mean Dep. Var	.4355		.470		.531	
Observations	1016	1016	1004	1004	1004	1004

Notes: The table presents linear probability model (OLS) estimates of the treatment effect on sign-ups for different cooperation tasks: willingness to provide expertise to the repository (Columns (1) and (2)), willingness to meet with a consultant to improve customer management processes (Columns (3) and (4)), and willingness to sign up as a junior mentor (Columns (5) and (6)). The dependent variable is binary, coded as 1 if the person chose to sign up and 0 otherwise. Columns (2), (4), and (6) include an interaction with being selected as the most supportive employee in the previous survey. If a person was selected and was in the award branch, they received the award. The table includes participants that have been transferred. Standard errors are clustered at the branch level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A3: Sign-Ups for Cooperative Tasks

	(1)	(2)	(3)	(4)	(5)	(6)
	Endline			*Follow-up*		
	Expertise	*Contribution I*	*Expertise*	*Contribution II*	*Volunteer*	*Mentor*
Award Branch	0.359***	0.391***	0.476***	0.474***	0.268**	0.265*
	(0.14)	(0.15)	(0.14)	(0.15)	(0.13)	(0.14)
Most Supportive R1		0.460**				
		(0.22)				
Award Branch × Most Supportive R1		-0.182				
		(0.30)				
Most Supportive R2				0.175		0.384
				(0.21)		(0.25)
Award Branch × Most Supportive R2				0.012		0.032
				(0.33)		(0.35)
Observations	908	908	908	908	908	908

Notes: Logit estimates of the treatment effect on sign-ups for different cooperation tasks: willingness to provide expertise to the repository (Columns (1) and (2)), willingness to meet with a consultant to improve customer management processes (Columns (3) and (4)), and willingness to sign up as a junior mentor (Columns (5) and (6)). The dependent variable is binary, coded as 1 if the person chose to sign up and 0 otherwise. Columns (2), (4), and (6) include an interaction with being selected as the most supportive employee in the previous survey. If a person was selected and was in the award branch, they received the award. Standard errors are clustered at the branch level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A4: Correlation of Survey Helpfulness Rankings with Helping Measurement

	(1)	(2)	(3)
	Signing up to Share Knowlege		
Supervisor rating	0.000		
	(0.00)		
Peer rating		0.002**	
		(0.00)	
In top 20% Most Supportive			0.095*
			(0.06)
Mean Dep Var	0.426		
Observations	471	471	471

The linear probability model (OLS) estimates the predictive power of various rankings on the willingness to provide productivity tips. The dependent variable takes the value of 1 if the employee signed up to share knowledge and 0 if they chose not to. Column (1) shows the predictive power of the supervisor ranking, which ranges from 0 (employee very rarely helps) to 100 (employee helps very frequently). Column (2) presents the linear relationship with peer ratings, which range from 0 (if no one chose the person as supportive) to 100 (if everyone in the branch chose the person as frequently helping). Column (3) uses a binary variable equal to 1 if the person is in the top 20% most supportive employees in the branch (and would therefore have received the award if in the treatment group), and 0 otherwise. Results are shown only for the control group. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

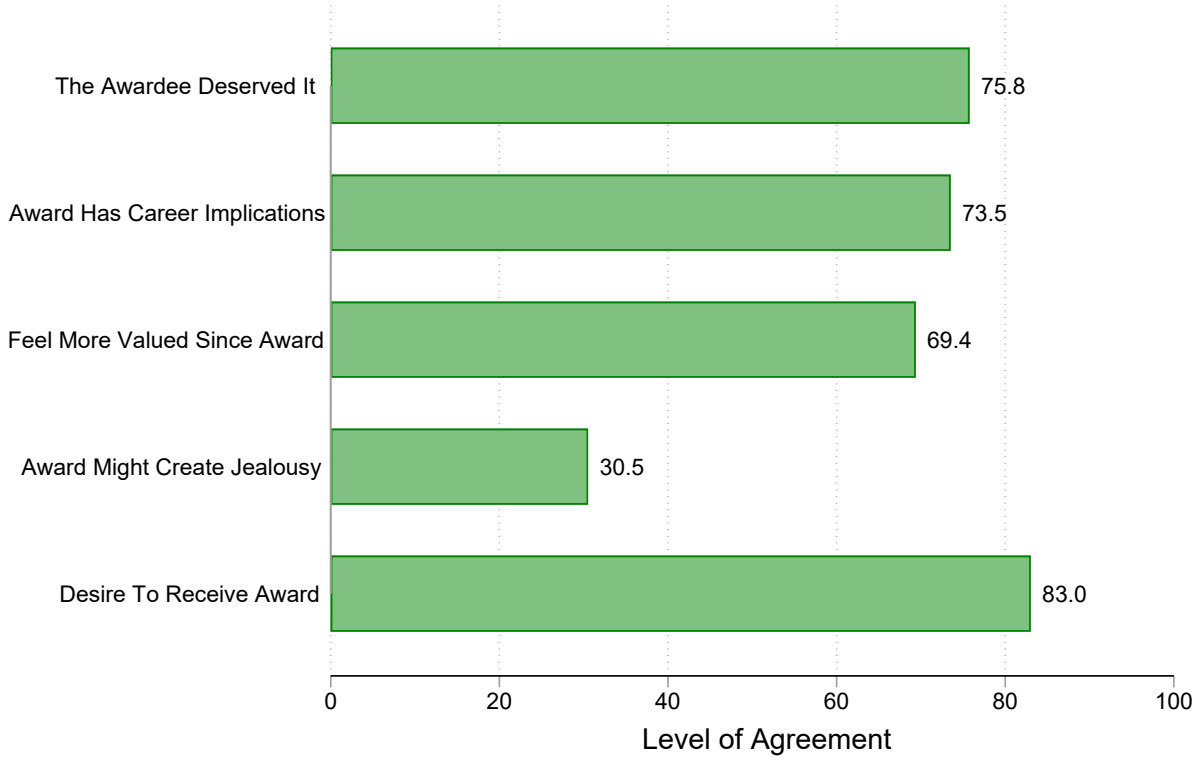

Figure A1: Perceptions of the Intervention at Six Month

References

Alfitian, J., D. Sliwka, & T. Vogelsang (2023): "When bonuses backfire: Evidence from the workplace," *Management Science*.

Ashraf, N., O. Bandiera, E. Davenport, & S. S. Lee (2020): "Losing prosociality in the quest for talent? Sorting, selection, and productivity in the delivery of public services," *American Economic Review*, 110, 1355–1394.

Auriol, E., G. Friebel, & L. Pechlivanos (2002): "Career Concerns in Teams," *Journal of Labor Economics*, 20, 289–307.

Autor, D. H., F. Levy, & R. J. Murnane (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118, 1279–1333.

Baker, M. (2016): "1,500 Scientists Lift the Lid on Reproducibility," *Nature*, 533, 452–454.

Battiston, D., J. B. i Vidal, & T. Kirchmaier (2021): "Face-to-Face Communication in Organizations," *Review of Economic Studies*, 88, 574–609.

Bénabou, R. & J. Tirole (2003): "Intrinsic and extrinsic motivation," *The Review of Economic Studies*, 70, 489–520.

Bruhn, M. & D. McKenzie (2009): "In pursuit of balance: Randomization in practice in development field experiments," *American Economic Journal: Applied Economics*, 1, 200–232.

Bursztyn, L., G. Egorov, I. Haaland, A. Rao, & C. Roth (2023): "Justifying Dissent," *The Quarterly Journal of Economics*, 138, 1403–1451.

Camerer, C. F., A. Dreber, F. Holzmeister, & et al. (2018): "Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015," *Nature Human Behaviour*, 2, 637–644.

Castro, S., F. Englmaier, & M. Guadalupe (2022): "Fostering Psychological Safety in Teams: Evidence from an RCT," *SSRN Electronic Journal*.

Chatterji, A. K., M. Findley, N. M. Jensen, S. Meier, & D. Nielson (2016): "Field Experiments in Strategy Research," *Strategic Management Journal*, 37, 116–132.

Deversi, M., M. G. Kocher, & C. Schwieren (2020): "Cooperation in a company: A large-scale experiment," .

Gallus, J., S. Campbell, & U. Gneezy (2023): "Awards: Tangibility, self-signaling and signaling to others," *Self-Signaling and Signaling to Others*.

Gallus, J. & B. S. Frey (2017): "Awards as Strategic Signals," *Journal of Management Inquiry*, 26, 76–85.

Gneezy, U., S. Meier, & P. Rey-Biel (2011): "When and Why Incentives (Don't) Work to Modify Behavior," *Journal of Economic Perspectives*, 25, 191–210.

Guillouet, L., A. K. Khandelwal, R. Macchiavello, M. Malhotra, & M. Teachout (2024): "Language Barriers in Multinationals and Knowledge Transfers," *Review of Economics and Statistics*, 1–56, forthcoming.

Hamilton, B. H., J. A. Nickerson, & H. Owan (2003): "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation," *Journal of Political Economy*, 111, 465–497.

Heursen, L., S. Friess, & M. Chugunova (2024): "Reputational Concerns and Advice-Seeking at Work," Research Paper 23-17, Max Planck Institute for Innovation & Competition, available at SSRN: `https://ssrn.com/abstract=4538216`.

Holmstrom, B. & P. Milgrom (1991): "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, & Organization*, 7, 24–52, special Issue: Papers from the Conference on the New Science of Organization, January 1991.

King, G., E. Gakidou, N. Ravishankar, R. T. Moore, J. Lakin, M. Vargas, M. M. Téllez-Rojo, J. E. Hernández Ávila, M. H. Ávila, & H. H. Llamas (2007): "A "politically robust" experimental design for public policy evaluation, with application to the Mexican universal health insurance program," *Journal of Policy Analysis and Management*, 26, 479–506.

Levitt, S. D. & J. A. List (2009): "Field Experiments in Economics: The Past, the Present, and the Future," *European Economic Review*, 53, 1–18.

Mickeler, M., P. Khashabi, M. Kleine, & T. Kretschmer (2023): "Knowledge seeking and anonymity in digital work settings," *Strategic Management Journal*.

Sandvik, J. J., R. E. Saouma, N. T. Seegert, & C. T. Stanton (2020): "Workplace Knowledge Flows," *The Quarterly Journal of Economics*, 135, 1635–1680.

Wagner, Z., J. B. Asiimwe, & D. I. Levine (2020): "When financial incentives backfire: evidence from a community health worker experiment in Uganda," *Journal of Development Economics*, 144, 102437.