CAF - WORKING PAPER #2020/02

April 22, 2020

Self-Selection into Corruption: Evidence from the Lab

Pablo Brassiolo[1] | Ricardo Estrada[2] | Gustavo Fajardo[3] | Juan Vargas[4]

[1]CAF -development bank of Latin America. pbrassiolo@caf.com

[2]CAF -development bank of Latin America. restrada@caf.com

[3]CAF -development bank of Latin America. gfajardo@caf.com

[4]Universidad del Rosario, School of Economics. juan.vargas@urosario.edu.co

We study whether the existence of opportunities to extract rents in a job affects the type of individuals who are attracted to it. We design a laboratory experiment in which individuals choose between two contracts, each offering a payment in return for performing a task, and we experimentally introduce the possibility of graft in one of the contracts. First, we find that the corruptible contract attracts less honest individuals and repels the more honest ones, thus changing the composition of the group that chooses that contract to the detriment of integrity. Second, we observe extensive graft when the opportunity is introduced. Using a double randomization strategy to disentangle pure incentives and selection effects, we find that selection is the fundamental driver of graft in our context.

KEYWORDS

Corruption, selection, rent extraction opportunities, personnel economics

CAF - DOCUMENTO DE TRABAJO #2020/02

Abril 22, 2020

Postulando a la corrupción: Evidencia del laboratorio

Pablo Brassiolo[1] | Ricardo Estrada[2] | Gustavo Fajardo[3] | Juan Vargas[4]

[1]CAF -banco de desarrollo de América Latina.
pbrassiolo@caf.com

[2]CAF -banco de desarrollo de América Latina. restrada@caf.com

[3]CAF -banco de desarrollo de América Latina. gfajardo@caf.com

[4]Universidad del Rosario, Escuela de Economía.
juan.vargas@urosario.edu.co

En este proyecto estudiamos si la existencia de oportunidades para extraer rentas de un trabajo tiene un efecto sobre el tipo de individuos que se ve atraído al mismo. Diseñamos un experimento de laboratorio en el que los sujetos eligen entre dos contratos que ofrecen un pago a cambio de ejecutar una tarea, e introducimos aleatoriamente la posibilidad de apropiarse de rentas en uno de los contratos. Primero, encontramos que los contratos con posibilidades de corrupción atraen a individuos menos honestos y repelen a los honestos, lo que empeora el nivel de integridad del grupo que elige el contrato. Segundo, observamos amplia apropiación de fondos cuando la oportunidad se presenta. Usando una estrategia de doble aleatorización para separar los efectos de incentivos y selección, encontramos que la apropiación de rentas se explica fundamentalmente por la autoselección de personas menos honestas.

KEYWORDS

Corrupción, selección, oportunidades de extracción de rentas, economía del personal

1 | INTRODUCTION

Corruption is one of the most pernicious problems that can affect government (Svensson, 2005; Olken, 2007). Motivated to understand the mechanisms that explain the occurrence of this problem, an important strand of papers finds that corruption flourishes where failures in incentives and monitoring prevail (Olken and Pande, 2012). These findings highlight the key role of institutional features in shaping the decisions of current public officials to engage in illicit behavior. However, incentives might also shape corruption by affecting who joins the public sector. Notably, the potential to extract rents in a job might attract individuals with a higher propensity to profit from these corruption opportunities (Besley, 2005).

In pioneering work, Hanna and Wang (2017) use a measure of dishonesty elicited in a laboratory task and find that in India, more dishonest students are more interested in public sector positions.[1] Barfort et al. (2015) use a similar method to gauge dishonesty among Danish students and find that in the Danish context, it is the more honest who prefer public sector positions. Taken together, these studies suggest that self-selection into the public sector is a function of the existent (or perceived) corruption opportunities. However, to the best of our knowledge, there is no previous work that provides causal estimates on this relationship. This gap is probably due to the challenge of putting together information on integrity and corruption at the individual level coupled with exogenous variation in exposure to jobs with and without rent extraction opportunities. In this paper, we address this gap and design an experiment to study, first, whether the existence of opportunities to extract personal rents in a job affects the number and type of individuals who are attracted to it and, second, to explore how much of actual rent extraction is explained by self-selection vis-à-vis pure incentives (net of selection).

We design a lab experiment in which individuals choose between two contracts that are intended to emulate private and public sector job characteristics, respectively (for simplicity, we will call them "private contract" and "public contract" from now on). Individuals who choose the private contract perform a real-effort task and receive a payment with both fixed and performance-based components. Participants who choose the public contract have to distribute a monetary fund among five NGOs and receive a fixed amount as payment. In the treatment group, participants use real banknotes in their monetary fund, whereas in the control group participants use tokens (toy banknotes). All other characteristics remain the same, including the guarantee of secrecy regarding how much each participant assigns to each NGO. This secrecy allows participants who choose the public contract the opportunity to keep some banknotes for themselves. In this way, we exogenously introduce the possibility of extracting monetary rents for those in the treatment group, which allows us to study how rent-extraction opportunities affect the pool of individuals who are attracted to a job.

We conduct this experiment among 398 students from a private, elite university in Bogota, Colombia. At the beginning of each session, we survey participants to collect information on their socioeconomic and demographic characteristics and individual measures of risk attitude, prosocial motivation, and personality traits. We also run an (incentivized) dice game as in Hanna and Wang (2017) to obtain an individual measure of propensity to cheat.

Our first set of results shows that the possibility of extracting rents changes the number and type of applicants who self-select to a contract. The total number of individuals who select the public contract when it becomes possible to keep part of the money falls by 10 percentage points, with respect to a baseline level of 31 percent in the control group. Crucially, this effect is heterogeneous by integrity (as measured by the dice game). Individuals with a

[1] In a similar vein, Banerjee, Baul and Rosenblat (2015) find that students enrolled in programs preparing for India's national public sector exam are more prone to corruption than MBA students.

higher propensity to cheat are more likely to select the public contract when rent extraction is possible, while the opposite is true for individuals with a lower propensity to cheat. We interpret these results as suggestive evidence that the possibility of rent extraction makes a position more attractive for the dishonest and less attractive for the honest.

We then turn to study actual rent extraction. We find that when the fund to be distributed is made up of real cash, the proportion of individuals who do not allocate the entire fund to the NGOs raises by 64–66 percentage points (from a baseline level of 5 percent in the comparison group), and the proportion of the fund not allocated to any organization increases by 40–42 percentage points. Since the reduction in donations is due to the use of cash instead of tokens, we can interpret it as graft. Moreover, graft is higher among people with lower levels of integrity (as measured by the dice game).

The use of cash constitutes a change in incentives that can affect performance (graft) through two channels: selection and pure incentives (net of selection). In most contexts, it is impossible to disentangle these two channels because the same incentives that affect on-the-job performance also affect selection into the job. To overcome this, one could randomly assign different incentives before and after the self-selection decisions of individuals. Such a double randomization strategy would result in i) individuals who face equal incentives at the moment of choosing a job but different incentives once on the job and ii) individuals who face the same incentives on the job but faced different ones at the moment of selection. This strategy allows for the identification of the pure incentives channel from the first group and the selection channel from the second.

With this in mind, we introduce the second stage of randomization in our experiment. In it, we consider the individuals who self-selected into the token version of the public contract (without rent extraction opportunities) and, after they made that choice, we reassign a random subset of them to play the cash-based version. In this way, we exogenously introduce the opportunity to obtain rents among individuals who self-selected to a contract without this possibility. Using this group, we can isolate the pure incentives channel. Complementarily, we can parse out the selection channel using the individuals who self-selected into different versions of the public contract but ex-post faced the same opportunities for graft. We find that the bulk of the observed graft is driven by self-selection.

Summing up, we show that the availability of rents in a position changes the pool of applicants interested in it. Our results highlight that corruption opportunities might not only attract the dishonest but also repel the honest, perhaps because of the reputation or moral costs of being in a corrupt environment. Using our double randomization design, we find that selection, not pure incentives, explains graft in this context. Of course, this does not imply that the pure incentives channel is irrelevant to curbing corruption. Instead, we interpret our results as evidence of the importance of the selection channel, which implies that to assess the total effect of a policy on corruption, one must pay attention to both channels.

This paper is closely related to the nascent literature on the link between corruption and self-selection to the public sector (Hanna and Wang 2017; Barfort et al. 2015). Our contribution is twofold. We present the first estimates of the causal effect of graft opportunities on the integrity of those who self-select to a task.[2] Second, we show that the selection of low-integrity individuals can be a significant driver of actual corruption.

Our findings also speak to a more general literature on personnel economics of the state that is concerned with estimating the effects of incentives on officials' performance (see a review in Finan, Olken and Pande (2017)). A general challenge faced by this literature is

[2]This is connected to Dal Bo, Finan and Rossi (2013) and Deserranno (2019), who document, respectively, how wages and career incentives affect the attributes of those who join the public sector. We show that the availability of rents is another important factor that might shape entry into the public sector.

that the features of the job that shape the incentives of employees are the same ones that shape employees' selection into that position in the first place. Consequently, it is almost impossible to separate selection and pure incentives effects in a given real-world context. The double randomization in our research design allows us to address this issue in a lab setting. Our results are consistent with those in Ashraf, Bandiera and Lee (2018), who use experimental variation in the recruitment of health care workers in Zambia to show that selection affects job performance.

Together, the results presented here are relevant for recruitment and personnel policies in sectors where positions have—or are perceived to have—some discretion to obtain personal rents (i.e., the public sector in developing countries) and are crucial to sectors where corruption is perceived to be widespread.

The rest of the paper is organized as follows: Section 2 elaborates on the experimental design and the implementation of the experiment, Section 3 describes the empirical analysis, Section 4 presents and discusses the results, and Section 5 concludes.

2 | EXPERIMENTAL DESIGN

The experimental session has three stages. In the first stage, participants fill out a survey questionnaire providing socioeconomic, demographic, and schooling information. The survey also includes questions to elicit values and preferences on several political issues and measures of risk aversion, prosocial motivation, and the Big Five personality traits. Most importantly, at this stage, we obtain an experimental measure of dishonesty for each individual based on a dice game developed by Fischbacher and Föllmi-Heusi (2013) and modified by Hanna and Wang (2017) for use at the individual level. In it, we ask participants to privately roll a die 40 times and record the result of each roll in their workstation. For this task, they receive a monetary payment equivalent to 75 Colombian pesos (COP)—approximately US$ .024—multiplied by the point total they report. This creates an incentive to report higher values than those obtained, which is possible because the die is rolled privately. A reported sum of 140 points (the expected value after 40 rolls without cheating) pays COP 10,500, or approximately US$ 3.40.

In the second stage, individuals are presented with and asked to choose between two contracts designed to capture stylized characteristics of private and public sector employment. A contract consists of a set of instructions for a task and an associated payment. By choosing one, individuals decide to perform the associated task and receive the corresponding payment, foregoing the alternative.

In the private contract, the task is to solve as many sums of five randomly generated two-digit numbers as possible in two minutes. The participants who choose this contract are paid a fixed amount of COP 10,000 (approximately US$ 3.20) plus COP 1,500 for each correct sum. Incorrect answers do not entail any loss.

In the public contract, the task consists in distributing a COP 14,000 monetary fund among five NGOs. The participants who choose this contract receive five empty envelopes labeled with the names of the NGOs and a brief description of the mission and activities of these organizations. The participants distribute the monetary fund among the envelopes according to their allocation decision and deposit all five envelopes (even if some are empty) in a common urn located in the laboratory.[3] The payment for this contract is COP 12,000,

[3]The experimentalist informs the participants that within four weeks from the session, the researchers will transfer to each organization an amount in pesos equivalent to the total amount corresponding to the allocation made to each organization. Participants are offered the possibility of receiving a copy of the vouchers of these transfers via e-mail.

which is received with the rest of the materials to perform the task in an additional envelope labeled "Payment".

Crucially, there are two versions of the public contract: a cash-based one and a token-based one. The only difference between these is that in the cash-based version, the monetary fund comes in the form of seven 2,000 COP legal tender banknotes, while in the token-based version the fund is made of seven tokens (toy banknotes) representing a value of 2,000 COP each. We randomly assign experimental sessions to be offered one of those two versions. Participants read the instructions for the public and private contracts they are offered and make their choice privately.

Key to our identification strategy, the cash-based contract introduces opportunities for graft by creating actual monetary incentives to take part of the fund, while the token-based one does not. Beyond this, there is no difference between the two versions of the contract. The size of the fund, the complexity of the task, and the amount of discretion are all the same in both versions. Moreover, the alternative (private) contract is also equal for both groups.[4] Upon reading the instructions for the public contract, participants can infer that leaving the room with some (or all) of the bills is a distinct possibility, even though it is not suggested at any point. The instructions explicitly say that no one in the room will be able to observe their allocation decision. This statement is credible to the participants because they work in individual stations where nobody can observe how many bills they put in the envelopes that they seal and deposit in the common urn before leaving the room.

We stress the secrecy of the allocation decision to ensure that individuals realize that it is possible to take part of the fund without being caught. However, we do not want to prime subjects to think that graft is a relevant outcome of the experiment. Thus, we justify the secrecy of the decision by including, among the five NGOs to receive funds, some organizations whose mission is controversial in the Colombian context. That is, NGOs with pro-life and pro-choice missions as well as NGOs that support ex-guerrilla combatants and others that support victims of guerrillas.

After participants make their choice between contracts, we perform a second randomized treatment. In half of the sessions offering the token-based public contract, we reassign all participants who chose that contract to play the cash-based public contract. We randomly assign experimental sessions to be offered one of those two public contract versions. We perform this second randomization *à la* Karlan and Zinman (2009), from token-based to cash-based public contracts, to separate selection and incentives channels in graft. Notice that when the participants finally perform the fund distribution task, we as experimenters will be able to observe whether subjects allocated the entire fund, which allows us to infer if they are engaging in graft. However, if we simply compare the fund allocation decisions between individuals who choose the cash-based version of the public contract and individuals who choose the token-based version, we would possibly confound two distinct channels that can affect graft: a selection effect through which more dishonest individuals self-select into the public contract when there are opportunities to extract rents, and an incentive effect through which equally self-selected individuals respond to opportunities for graft that are presented. To disentangle these effects, we create a group of individuals who self-select into one version of the public contract but play the other version.

In the final and third stage, participants perform the tasks corresponding to their contract and receive their payment. First, public contract participants receive their task materials, including the envelope with their payment. Then, everyone in the room starts working on their task. For those with the public contract, a brief explanation of the mission and

[4]The instructions clearly state that the payment for the public contract is COP 12,000, which avoids the concern that individuals may feel legitimately entitled to the monetary fund and allows us to confidently interpret as graft any appropriation of money from the fund by a participant.

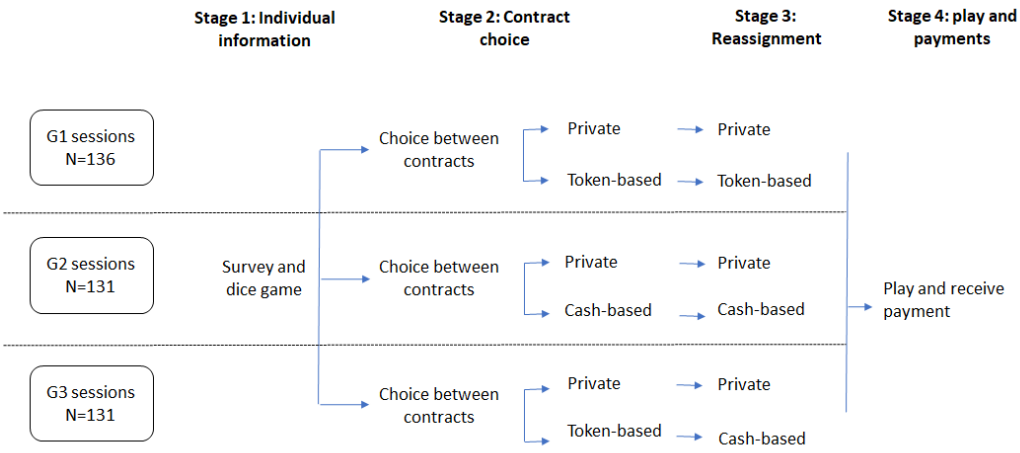

FIGURE 1 Description of experimental sessions

activities of each of the five NGOs is presented on their computer screens. Meanwhile, those with private contracts solve the sums on their computers. To ensure the privacy of public contract participants while they distribute their funds, the experimental team leaves the room for five minutes and asks all participants to remain seated. Once the experimental team returns, public contract participants place their envelopes in the urn. As they leave the session, they receive the payment corresponding to the dice game in the first stage. Only then does the experimental team give private contract participants their earnings (including those corresponding to the dice game).

2.1 | Implementation

We randomize our treatments at the session level. That is, we have three groups of experimental sessions:

- G1: Participants are offered the private contract and the token-based public contract, and they play whichever they choose.
- G2: Participants are offered the private contract and the cash-based public contract, and they play whichever they choose.
- G3: Participants are offered the private contract and the token-based public contract, and the participants who choose the public contract are reassigned to play the cash-based one.

In G3 sessions, after the participants make their contract choice, the experimenter makes the public announcement that instead of tokens, the monetary fund will be made up of real banknotes. By implementing the treatment at the session level, we avoid having participants know that others are receiving a different treatment. Figure 1 offers a description of the timeline of events in each type of session.

Our sample is composed of undergraduate students of Universidad del Rosario. The invitation letter asked students whether they preferred to find a job in the private or public sector. With this information, we stratified preregistered students (all of those who answered our invitation) by gender, field of study, and sector preference and randomly assigned individuals within each stratum to one of the session types. We then conduct a total of 15 sessions (5 of each type) with an average of 25 participants in each.

We mark the five envelopes used to distribute the monetary fund with a unique pattern that allows us to identify how much each participant gave to each organization and, more importantly, how much of the fund the subject distributed in total. We cannot track that allocation decision back to participants' names or personal information, but we can link it to any information provided during the session.

3 | EMPIRICAL ANALYSIS

First, we are interested in studying if the participants' contract choice depends on the version of the public contract they were offered. That is, whether having the monetary fund in real money as opposed to tokens affects the proportion and the type of individuals who choose the public contract. For this, we need information from only the first two stages in the sessions.

Effect of opportunities for graft on contract choice. To test the effects of the offer on the number of individuals who choose the public contract, we estimate the following equation on the full sample of subjects:

$$Y_i = \beta_0 + \beta_1 Cash{-}based_i + \delta_i X_i + \gamma_s S_s + \epsilon_i \quad (1)$$

Our outcome is a dummy that we call $Public\ contract\ choice$ and takes the value 1 if the subject chooses the public contract and 0 otherwise (i.e., chooses the private contract). Our treatment variable is $Cash-based_i$ and takes value 1 if the participant i was offered the cash-based public contract. This means that, for this set of results, individuals from sessions G1 and G3 represent the control group ($Cash-based$=0), while individuals in G2 sessions are the treated ones ($Cash-based$=1). X_i is a vector of individual controls, and S_s represents the stratification variables.

β_1 captures the treatment effect of being offered the public cash contract on the probability of choosing the public contract. We interpret this as the causal effect of introducing graft opportunities in a contract on the probability that individuals self-select into it. We are able to isolate the effect of graft opportunities because our treatment is such that the only difference between groups is the possibility of extracting rents when cash is used.

Composition effects of opportunities for graft. We expect the introduction of opportunities for graft to make the public contract relatively more attractive to more dishonest individuals. This means that the cash-based contract would have consequences on the composition of individuals who self-select into it beyond its effect on the number of people who choose it. To explore this, we start by estimating equation (2), where D_i is the total number of points reported by the individual in the dice task. Since this variable increases with the probability of dishonesty, we expect the coefficient of the interaction term β_3 to be positive.

$$Y_i = \beta_0 + \beta_1 Cash{-}based_i + \beta_2 D_i + \beta_3 Cash{-}based_i * D_i + \delta_i X_i + \gamma_s S_s + \epsilon_i \quad (2)$$

Additionally, we also divide D into its tertiles and use them as a categorical variable that we interact with the $Cash-based_i$ dummy, as in Equation (3), where I_i takes value 1 for observations in the i_{th} tertile of the distributions of total points in the dice game.

$$Y_i = \beta_0 + \sum_{i=1}^{3} \theta_i(I_i = 1) + \sum_{i=1}^{3} \beta_i(I_i = 1)Cash{-}based_i + \delta_i X_i + \gamma_s S_s + \epsilon_i \quad (3)$$

We also assess if there are composition effects on dimensions other than honesty. We do this by estimating Equation 2 but replacing our measure of honesty with variables that capture risk attitudes, prosociality, and Big Five personality traits.

Allocation decisions and graft. We then turn to study actual graft from the information we gather on individuals' allocation of the monetary fund. Two variables capture those allocation decisions. The first one, $Incomplete\ allocation$, is a dummy that takes the value 1 if the amount allocated to the NGOs is lower than the value of the fund and 0 otherwise. The second one, $Not-allocatedshare$ is the percentage of the fund that is not allocated to the NGOs. This variable ranges from 0 to 1 in steps of approximately 0.14. Notice that both of these variables are defined only for those who choose the public contract since only they have the task of distributing the fund.

Note that individuals with the token-based version of the contract can have positive values for the $Not-allocatedshare$ variable (and a value of 1 for the $Incomplete\ allocation$ dummy). This can happen if, for example, the participant forgets to introduce all the tokens in the envelopes by inattention or a mistake. In fact, this gives us a baseline level of losses of the monetary fund that is not motivated by rent extraction. However, we expect both variables to have larger values for those with the cash-based contract, and we interpret as graft the increases in incomplete allocations when opportunities for graft are introduced.

We start by estimating Equation 1 using $Incomplete\ allocation$ and $Not-allocatedshare$ as outcomes for the sample of participants who choose the public contract in groups G1 and G2. This allows us to measure how the introduction of opportunities for graft changes the allocation decisions regarding the monetary fund. That is, it allows us to gauge the extent of graft. However, this specification is silent about the channels that explain graft. Note that individuals from groups G1 and G2 not only *play* different versions of the public contract (token-based and cash-based), but they *self-selected* into these different contracts. This means that there can simultaneously be a selection effect and an incentives effect (net of selection) behind the results. Comparing individuals from groups G2 and G3 does not allow us to disentangle both channels.

Separating selection and incentive effects in graft. G3 sessions are designed precisely to disentangle this effect, as individuals self-select into the token-based version of the public contract but are then reassigned into the cash-based one. Note that the comparison between individuals in G1 and G3 is informative about the pure incentives channel: both groups are equally self-selected (no differences in the composition of honesty), but they randomly end up *playing* different versions of the public contract (one with graft opportunities and one without). On the other hand, the comparison between G2 and G3 is informative about the selection channels: those groups are differently self-selected, which potentially leads to differences in the levels of individual honesty between the groups even though they play under the same conditions.

Using this double randomization strategy, we follow Karlan and Zinman (2009) and separate both channels by estimating Equation 4 for all individuals who chose the public contract in all the sessions. As outcomes, we use both measures of the monetary fund allocation decisions. $Cash-based\ offered$ takes the value 1 if the participant i was offered the cash-based public contract and 0 if token-based. $Cash-based\ played$ equals 1 for those

who play cash-based (either by choice or reassignment) and 0 for those who play token-based. Both variables take missing values for anyone who chooses the private contract.

$$Y_i = \beta_0 + \beta_1 \text{Cash}-\text{based offered}_i + \beta_2 \text{Cash}-\text{based played}_i + \delta_i X_i + \gamma_s S_s + \epsilon_i \quad (4)$$

In this specification, β_1 captures the change in allocation decisions that comes exclusively through changes in the composition of individuals making those decisions (it captures the selection channel). In turn, β_2 captures the effect of the incentives channel, controlling for the changes in composition due to differential selection. We would expect both effects to be weakly positive, but an important question is related to the relative magnitudes of both. The driver behind graft (and other manifestations of corruption) has significant policy implications. If selection is most important, anti-corruption measures should emphasize the mechanisms to recruit, promote, and dismiss officials. But if incentives are the main explanation, efforts should be directed toward the administrative rules and protocols that govern the work of public officials.

4 | RESULTS

4.1 | Balance and Descriptive Statistics

Table 1 shows descriptive statistics and the balance of covariates between the three experimental groups (G1, G2, and G3). The results show that randomization worked well. Differences between groups are small and statistically insignificant.

Regarding the characteristics of individuals in our sample, there are a majority of women and a majority of people from high socioeconomic strata. About a third come from majors with strong quantitative content (mostly economics, but also some from natural sciences and math), and slightly more than half report having a preference for public sector positions.

We also report the results of our experimental measure of dishonesty, collected through the dice game in the first stage of our sessions, which shows significant cheating by participants. Individuals report, on average, a total of 158 points in the 40 dice rolls. This figure is 13 percent larger than the expected value of that variable. We can also compute the proportion of individuals who report point totals with a very low probability of occurrence. In our sample, 30 percent of participants report a total of points equal to or greater than 165, which is the 99th percentile of the theoretical distribution of 40 rolls of a die. These numbers are very close to those obtained by Hanna and Wang (2017) in their sample of Indian students. They find that the average report is 14 percent higher than the expected value, and 34 percent of the reports are above the 99th percentile of the corresponding theoretical distribution.

Figure 2 plots both the empirical distribution of points in the dice task from our sample and a simulated distribution. The former is clearly displaced to the right of the latter. The empirical distribution shows a noticeable spike in the value 240, which is the maximum of the range of possible values of the variable. This is because eight individuals (approximately 2 percent of the sample) reported a six on every roll of the die.

TABLE 1 Summary statistics and balance of covariates between groups

	G1	G1-G2	G1-G3
	(1)	(2)	(3)
Female	0.610	-0.023	-0.068
	(0.489)	(0.060)	(0.061)
Age	20.434	0.062	0.062
	(1.908)	(0.245)	(0.254)
Quantitative Major	0.338	-0.002	0.021
	(0.475)	(0.058)	(0.059)
GPA (Mean=0, SD=1)	-0.005	-0.045	0.028
	(0.917)	(0.120)	(0.114)
High Socioeconomic Stratum	0.618	-0.083	-0.106*
	(0.488)	(0.060)	(0.061)
Preference for Public Sector Job	0.566	-0.017	-0.009
	(0.497)	(0.061)	(0.061)
Points in dice task	158.294	-1.744	-1.836
	(24.660)	(2.892)	(2.779)
High points in dice task ($p \geqslant 99$ of simulated distribution)	0.301	-0.057	-0.019
	(0.461)	(0.055)	(0.056)
Declares bribe is justified	0.103	-0.034	0.019
	(0.305)	(0.034)	(0.039)
Risk Preference (Mean=0, SD=1)	0.026	0.056	-0.114
	(1.031)	(0.122)	(0.127)
Prosociality Index (Mean=0, SD=1)	0.017	0.027	-0.146
	(1.089)	(0.126)	(0.123)
Big Five, Extraversion (Mean=0, SD=1)	0.013	0.018	-0.158
	(1.034)	(0.122)	(0.125)
Big Five, Agreeableness (Mean=0, SD=1)	-0.016	-0.009	-0.066
	(1.046)	(0.128)	(0.118)
Big Five, Conscientiousness (Mean=0, SD=1)	0.061	-0.013	-0.114
	(0.970)	(0.119)	(0.122)
Big Five, Neuroticism (Mean=0, SD=1)	0.009	0.017	-0.032
	(1.075)	(0.126)	(0.128)
Big Five, Openness (Mean=0, SD=1)	-0.035	0.061	0.086
	(0.898)	(0.117)	(0.122)
Observations	136	267	267

Notes: The table shows summary statistics and the balance of covariates between experimental groups. Column (1) presents the mean and standard deviation of each variable for individuals in G1 sessions. Columns (2) and (3) show the mean differences for each variable between individuals in G1 sessions with respect to those in G2 and G3 sessions, respectively, and the corresponding t-statistics in parentheses. * significant at 10%, ** at 5%, *** at 1%. The last row shows the number of observations used for the calculation of each column's statistics: in Column (1) this corresponds to the size of group G1, while in Columns (2) and (3) it indicates the sizes of groups G1 and G2, and G1 and G3, respectively. GPA: grade point average.

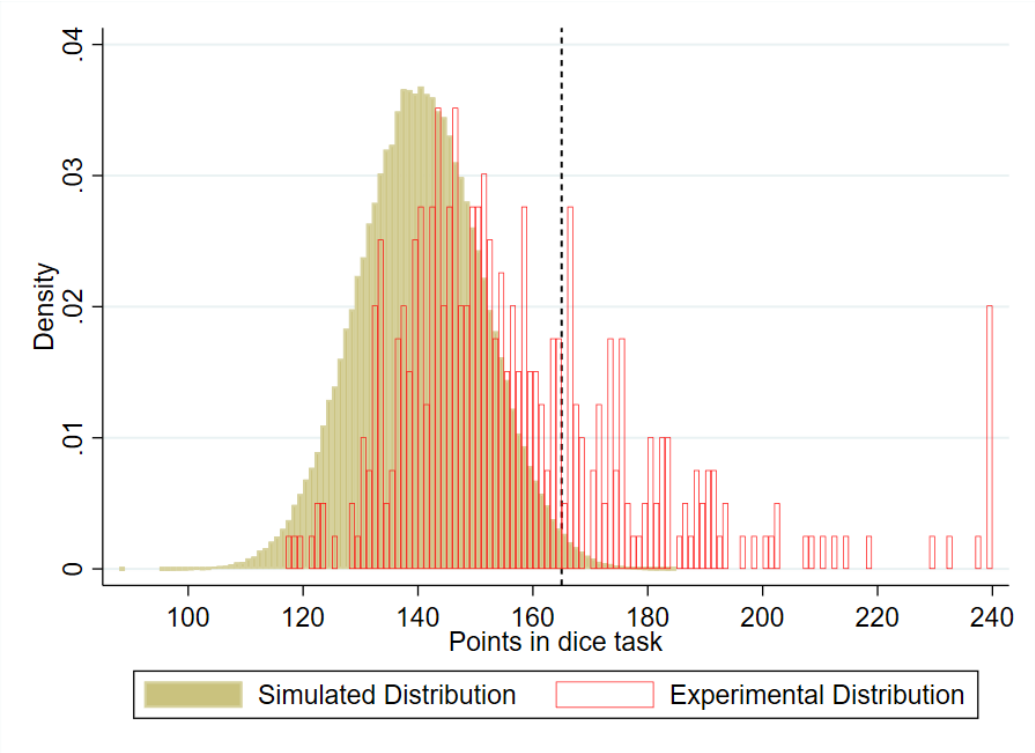

FIGURE 2 Dice points: Simulated and experimental distributions. *Notes:* The figure shows both the experimental distribution of points in the dice task and the simulated distribution. The dashed vertical line shows the value corresponding to the 99th percentile of the simulated distribution.

4.2 | Contract Choice

We want to study whether the existence of opportunities for graft in a contract affects the probability of choosing it. To do this, we turn to the contract choice stage of our experiment and estimate Equation 1. The coefficient for $Cash-based$ captures the causal effect of being offered the cash-based version of the public contract on the probability of choosing the public contract.

The first two columns of Table 2 show the baseline results for this exercise. The introduction of graft opportunities reduces the likelihood of choosing the public contract by 10 percentage points. That means that public contract choice falls by about one-third with respect to the control group, i.e., those who were offered the token-based version of the contract. Column (1) controls for the stratification variables (gender, preference for public sector positions [a dummy], and quantitative major [a dummy]), the main result is unchanged by the inclusion of additional covariates (age, GPA, socioeconomic status, risk preference, prosociality, and the Big five personality traits), as shown in Column (2).

When our measure of dishonesty is included in the baseline regression, it has no effect on contract choice on its own. However, we are interested in examining whether differences in individual honesty trigger different responses to the introduction of graft opportunities. To capture this, we interact the $Cash-based$ dummy with our experimental measure of dishonesty using different specifications. First, we use the total number of points reported by participants in the dice game to estimate Equation 2 (Column 3). We also use a dummy that takes the value 1 for individuals who report total dice points above the 99th percentile of the theoretical distribution (Column 4). In a final alternative, we create a categorical

variable with the tertiles of the distribution of reported dice points, which we interact with the treatment, as in Equation 3 (Column 5).

The results show a noticeably clear pattern: the introduction of graft opportunities in the public contract alienates the most honest individuals while attracting those with a higher propensity to cheat. Column (3) shows that the $Cash-based$ coefficient is unchanged by the introduction of the interaction term, indicating a negative effect of graft opportunities on the probability of choosing the public contract. However, the interaction term is positive and statistically significant, meaning that the public contract with graft opportunities becomes relatively more attractive for those who reported higher numbers in the dice game. In this equation, the variable $Dicepoints$ is normalized to have the mean 0 and a standard deviation equal to 1. The coefficient suggests that the negative effect of graft opportunities on public contract choice disappears for those who report totals in the dice game one standard deviation above the mean.

Results in column (4) have a more straightforward interpretation. They show that the probability of choosing the public contract drops by 16.7 percentage points when opportunities for graft are introduced, but only among those who report reasonable results in the dice game. The effect vanishes for the most dishonest participants. The sum of the $Cash-based$ coefficient and the interaction term is positive but not statistically significant. Column (5) presents consistent results. Graft opportunities reduce the selection of the public contract for those in the first two tertiles of the distribution of reported points in the dice game, but not for those who most likely cheated (third tertile of the distribution).

TABLE 2 Contract choice

	Dependent variable: Public contract choice				
	(1)	(2)	(3)	(4)	(5)
Cash-based	-0.102**	-0.100**	-0.100**	-0.167***	-0.179**
	(0.045)	(0.048)	(0.048)	(0.053)	(0.085)
Dice points		0.003	-0.027		
		(0.024)	(0.027)		
Cash-based x Dice points			0.114**		
			(0.053)		
Unlikely dice points				-0.065	
				(0.064)	
Cash-based x					
Unlikely dice points				0.230**	
				(0.111)	
Tertile 2, dice					-0.145**
					(0.069)
Tertile 3, dice					-0.115
					(0.073)
Cash-based x tertile 2					0.017
					(0.110)
Cash-based x tertile 3					0.213*
					(0.121)
Controls					
Stratification variables	Yes	Yes	Yes	Yes	Yes
All controls	No	Yes	Yes	Yes	Yes
Mean of dependent					
variable (control)	0.315	0.315	0.315	0.315	0.378
Observations	398	398	398	398	398
R-squared	0.040	0.090	0.102	0.084	0.116

Notes: The table presents the ordinary least squares (OLS) estimates of the effects of introducing rent extraction opportunities in the public contract on the proportion and type of individuals who choose it. The dependent variable is a binary indicator that takes the value 1 if the individual chooses the public contract and 0 otherwise. The treatment variable, Cash-based, is a binary indicator that takes the value 1 for participants in group G2 (i.e., those who are offered the cash-based public contract) and 0 for participants in groups G1 and G3 (i.e., those who are offered the token-based public contract). Columns (1) and (2) show the baseline results. Column (3) shows the results when the treatment variable is interacted with the total number of points reported by participants in the dice game. Column (4) shows the results when the treatment variable is interacted with an indicator of having reported total dice points above the 99th percentile of the theoretical distribution. Column (5) shows the results when the treatment variable is interacted with a categorical variable with the tertiles of the distribution of reported dice points. Column (1) controls for the stratification variables: gender, a dummy for preference for public sector positions, and a dummy for quantitative major. Columns (2) to (5) also control for age, GPA, socioeconomic status, risk preference, prosociality, and the Big Five personality traits. Robust standard errors are in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

This implies that opportunities for graft deteriorate the pool of individuals who select the public contract. Figure 3 presents evidence along these lines. First, Panel A shows that our measure of dishonesty is balanced between those who were offered the cash-based public contract (solid line) and the token-based one (dashed line). As expected given the randomization, these distributions overlap almost completely. Then, Panel B shows the effect of introducing graft opportunities on the honesty of those who choose the public contract. The thin solid line is the distribution of reported dice totals over all subjects in our experiment, which is presented as a benchmark. The dashed line is the distribution over those who choose the public contract in the token-based version (no graft opportunities). This distribution is similar to the benchmark but slightly displaced to the left, signaling a very small improvement in selection in terms of honesty. On the other hand, the distribution for those who choose the cash-based public contract (thick solid line) is flatter with much more mass toward higher values, indicating a clearly worse composition of this group in terms of honesty. In fact, 44 percent of those who choose the public contract with graft opportunities report dice results above the 99th percentile of the theoretical distribution. The corresponding number for those who choose the token-based public contract is 22 percent.

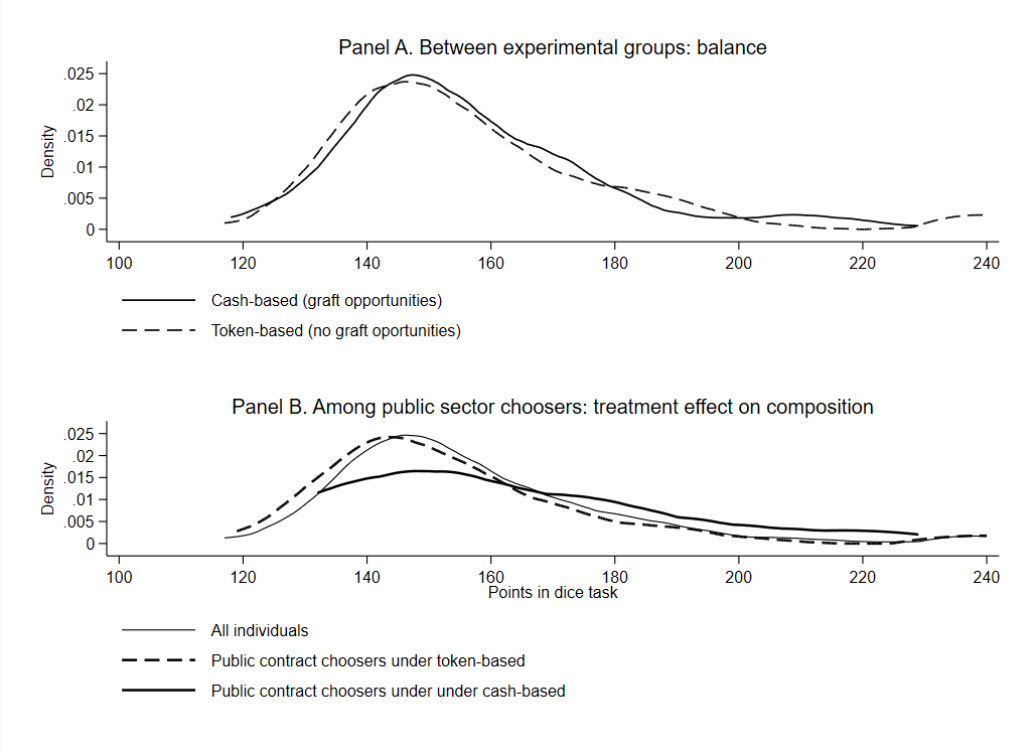

FIGURE 3 Distribution of dishonesty measure. *Notes:* Panel A shows kernel density estimates of total dice points for experimental groups G2 (solid line) and G1-G3 (dashed line). Panel B shows kernel density estimates of total dice points for the full sample (thin solid line), those who choose the public contract in experimental group G2 (thick solid line), and those who choose the public contract in experimental groups G1-G3 (dashed line).

To summarize, our experiment delivers two main results regarding selection. First, the introduction of graft opportunities reduces the number of individuals who self-select into the contract. Second, it simultaneously changes the honesty composition for the worse, meaning that the proportion of dishonest individuals among those who choose the contract increases.

Effects of composition in other dimensions. We also explore whether the introduction of graft opportunities changes the composition of the group who chooses the public sector in other dimensions. To do this, we estimate Equation 2 by substituting our dishonesty measure with variables capturing other relevant attributes. We find no evidence of heterogeneous responses by risk attitudes or prosociality, as captured by our instruments. We also compute scores for the Big Five personality traits and find some action only on the Openness measure: higher values in the Openness index is associated with higher rejection rates of the cash-based version of the public contract.

4.3 | Allocation of the Monetary Fund and Graft

Importantly, our setting also allows us to study actual graft. If the opportunity to extract rents is affecting the decision of individuals regarding contract choice, we would expect to observe individuals cashing in on those opportunities.

We have two variables that capture what individuals do with the resources in the monetary fund they were asked to allocate. The first one is the dummy *Incomplete allocation*, which takes the value 1 for those whose sum of donations across the set of NGOs is less than the amount in the fund. The second one is the difference between the total sum allocated and the fund (as a percentage of the fund), which we call *Not-allocated share*. Note that these variables are defined only for public contract participants.

To study graft, we assess how the cash-based contract (with rent extraction opportunities) affects the fund allocation decisions of individuals who choose the public contract. In Table 3, we regress our fund allocation measures on the type of public contract, which we capture with the dummy *Cash-based* (takes value 1 for those in group G2 and 0 for those in G1). We abstract from individuals in group G3 since the variable cash-based is not well defined for them (they did not play the same version of the contract that they chose). Note that for those who play with tokens (instead of cash), a value of 1 in the *Incomplete allocation* dummy implies that they did not allocate all the tokens among the NGOs. Since the tokens have no value outside the lab, that outcome can capture distraction or a failure to follow the instructions of the experiment. We interpret as actual graft the coefficient of *Cash-based*, as it captures the differences in allocation decisions given by the introduction of graft opportunities through the use of cash.

Practically all the incomplete allocations come from individuals who play the cash-based public contract. Among those in the token-based contract, only 4 percent do not allocate the entire fund, while the figure jumps to 70 percent for those handling cash (Column 1 of Table 3). Something similar occurs with the share of the fund that is not allocated, which goes from 3 percent to 45 percent when cash is introduced (Column 3). Moreover, our experimental measure of dishonesty predicts allocation decisions. An increase of one standard deviation in reported dice points raises the likelihood of an incomplete allocation by 6.7 percentage points (Column 2), and it increases the share of the fund that is not allocated by 7.8 percentage points (Column 4).

4.4 | Separating Selection and Incentive Mechanisms

The previous set of results shows the cash-based version of the public contract leads to graft. This can be a result of both i) worse selection of individuals into the cash-based contract and ii) an incentives channel net of the selection effect. We have already established that offering a cash-based contract leads to worse selection in terms of our measure of honesty (Table 2) and that dishonesty predicts incomplete allocations of the monetary fund (Table 3). Hence, there are very strong clues that the selection channel can be important to explain

TABLE 3 Graft: Differences in monetary fund allocation by type of contract

	Incomplete allocation		Not-allocated share	
	(1)	(2)	(3)	(4)
Cash-based	0.656***	0.635***	0.421***	0.397***
	(0.095)	(0.097)	(0.079)	(0.075)
Dice points		0.067**		0.078**
		(0.032)		(0.033)
Mean of dependent variable (Comparison group)	0.048	0.048	0.034	0.034
Observations	69	69	69	69
R-squared	0.484	0.512	0.366	0.434

Notes: The table presents the OLS estimates of the difference in the fund allocation measures by type of public contract offered. The dependent variable in Columns (1) and (2) is a binary indicator that takes the value 1 if individuals did not fully allocate the fund among the NGOs and 0 otherwise, while in Columns (3) and (4), it is the share of the fund that is not allocated. The treatment variable, Cash-based, is a binary indicator that takes the value 1 for participants in group G2 (i.e., those who are offered the cash-based public contract) and 0 for participants in group G1 (i.e., those who are offered the token-based public contract). Only individuals who choose the public contract in experimental groups G1 and G2 are included in the regressions, since the variable cash-based is not well defined for those in group G3 (they did not play the same version of the contract that they chose). Robust standard errors are in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

observed graft. However, the results in Table 3 do not allow us to disentangle selection and incentives. Remember, in those regressions, we use only individuals from groups G1 and G2, and those two groups self-selected differently (they were *offered* different versions of the contract) and also faced different incentives (they *played* different versions of the contract).

We can separate selection and incentives channels thanks to the second randomization embedded in our experiment: individuals in group G3 are offered the token-based version of the public contract, and if they choose it, they are reassigned to play the cash-based one. Thus, we have observations corresponding to people who are self-selected into one version of the contract but play a different one. Using this variation, we estimate Equation 4 for all individuals who play the public contract.

The results suggest that all the graft we observe comes from the selection channel (Table 4). Those who self-select into the cash-based version are 60 percentage points more likely to make an incomplete allocation of the monetary fund, and the share they allocate is on average 40 percentage points smaller. Controlling for the selection channel, the pure incentives effect captured by the $Cash-based\ Played$ dummy is effectively 0. This result can also be visualized in Figure 4, where we plot the mean of the fund allocation variables by experimental group. The comparison between the first and third groups captures the pure incentive effect, as individuals who play the public contract in those two groups are selected under the same conditions but differ in the version they ultimately play. As can be seen, both the percentage of incomplete allocations and the not-allocated share are higher for the group playing with cash (9.5 against 4.8 and 5.4 against 3.4, respectively), but the differences are small and insignificant. Clearly, the significant increase in graft comes from group G2, where the selection channel is activated.

Our results lead to the conclusion that introducing graft opportunities into a contract

TABLE 4 Separating selection and incentive effects

	Incomplete allocation			Not-allocated share		
	(1)	(2)	(3)	(4)	(5)	(6)
Cash-based Offered	0.608***	0.606***	0.596***	0.401***	0.377***	0.366***
	(0.100)	(0.102)	(0.102)	(0.080)	(0.075)	(0.071)
Cash-based Played	0.048	0.007	0.008	0.020	0.021	0.022
	(0.057)	(0.055)	(0.058)	(0.040)	(0.043)	(0.044)
Dice points			0.038			0.045
			(0.029)			(0.027)
Controls						
Stratification variables	Yes	Yes	Yes	Yes	Yes	Yes
All controls	No	Yes	Yes	No	Yes	Yes
Mean of dep. variable (Comparison group)	0.048	0.048	0.048	0.034	0.034	0.034
Observations	111	111	111	111	111	111
R-squared	0.424	0.550	0.558	0.344	0.474	0.496

Notes: The table presents the OLS estimates of the effect of contract offer and contract assignment on measures of monetary fund allocation. The dependent variable in Columns (1) to (3) is a binary indicator that takes the value 1 if individuals did not fully allocate the monetary fund and 0 otherwise, while in Columns (4) to (6), it is the share of the fund that is not allocated. There are two treatment variables: cash-based offered is a binary indicator takes the value 1 for participants in group G2 (i.e., those who are offered the cash-based public contract) and 0 for participants in groups G1 and G3 (i.e., those who are offered the token-based public contract), while Cash-based played is a binary indicator that takes the value 1 for individuals in groups G2 and G3 (i.e., those who played the cash-based public contract either by choice or re-assignment) and 0 for individuals in group G1 (i.e., those who played the token-based public contract). Columns (1) and (4) control for the stratification variables: gender, a dummy for preference for public sector positions, and a dummy for quantitative major. Columns (2), (3), (5), and (6) also control for age, GPA, socioeconomic status, risk preference, prosociality, and the Big Five personality traits. The regressions include individuals who choose the public contract in all experimental groups. Robust standard errors are in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

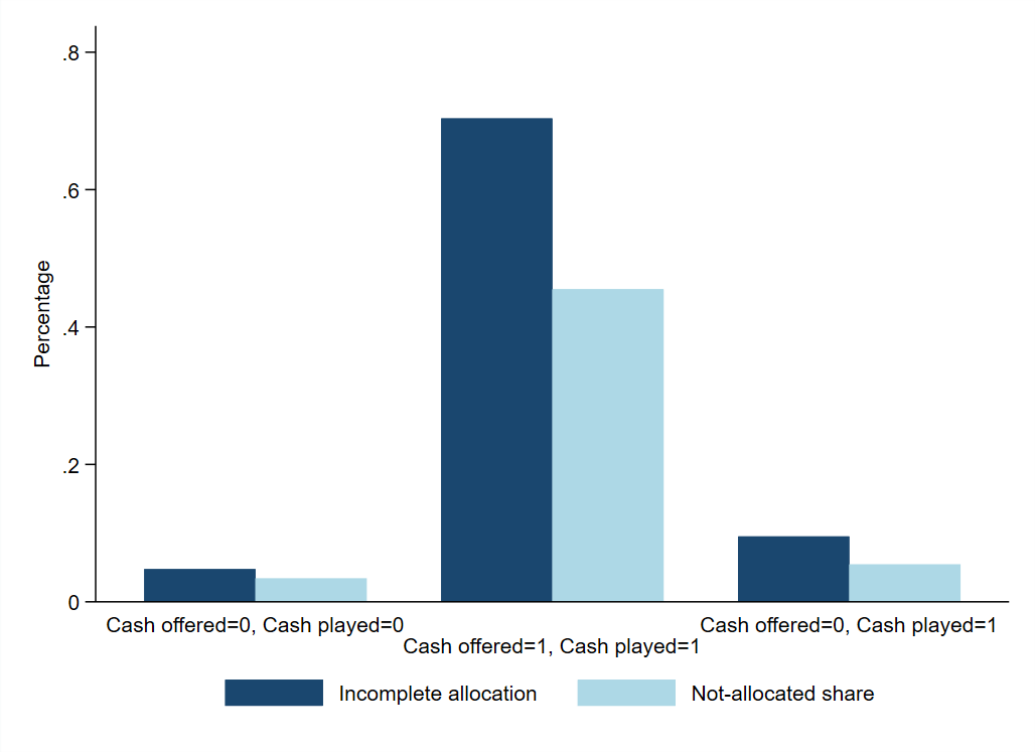

FIGURE 4 Donation by experimental groups. *Notes:* The figure reports the monetary fund allocation measures by experimental group. Dark blue bars show the percentage of participants who do not fully allocate the fund among the NGOs, while light blue bars show the average share of the fund not allocated.

increases observed graft significantly, mostly because individuals with high dishonesty levels self-select disproportionately into that contract. Yet, our finding that there are no important incentive effects must be qualified. We cannot rule out that the importance of incentives is heterogeneous across individuals: for example, more dishonest people could be more responsive to changes in rules that facilitate rent extraction. Our result is that changes in incentives do not lead to important behavioral responses among those who self-select into a contract with no opportunities to extract rents. This suggests that a good selection (in terms of honesty) can mitigate the risks of rule changes that create opportunities for fraud or corruption.

Another important point concerns how the nature of our experiment informs personnel policies in the public sector. A recurring debate in public sector management pertains to trade-offs associated with increasing officials' discretionary powers. On one hand, doing so increases bureaucracies' agility in responding to contingencies; on the other hand, it can increase the risk of corruption. In our experiment, we study a policy change (the use of cash instead of tokens) that creates opportunities for graft without affecting individuals' ability to carry out the task at hand. This allows us to isolate the pure effect of increasing corruption risks. Our results serve as a warning: policies that facilitate rent extraction do lead to more corruption, especially through self-selection among more dishonest individuals. However, most public administration policies that increase discretion (and create opportunities for malfeasance) also increase the flexibility of individuals to complete tasks, and thus it is important to asses how the corruption costs compare to the benefits of that extra flexibility.

5 | CONCLUSIONS

This paper provides the first causal estimates on the effect of graft opportunities on the integrity of those who self-select to a task. Our results highlight that corruption opportunities might not just attract the dishonest but also repel the honest, perhaps because of the reputation or moral costs of being in a corrupt environment. Additionally, our double randomization design allows us to study the contributions of the selection and pure incentives channels to explain graft. We find that selection explains the bulk of graft in this context.

Our findings indicate that the incentives and rules that exist in an organization shape who aspires to join the organization, and the attributes of employees are a big determinant of organizational outcomes.

Two policy implications emerge from this. First, it is important to consider potential effects on selection when designing rules and incentives in an organization. Second, improving public sector recruitment practices should be a focal point of reform, especially where there are high levels of corruption.

Our experiment sheds light on one reason why corruption is so persistent. Organizations where rent extraction is possible disproportionately attract dishonest people, which translates into higher levels of corruption and reinforces these organizations' poor reputations. In short: corruption begets corruption. To break this vicious cycle, policies centered on improving the composition of organizations can be a promising path for reform.

ACKNOWLEDGEMENTS

We thank Mariana Blanco and seminar participants from the Bogota Experimental and Behavioral Economics Seminar (BEBES), Universidad del Rosario and CAF -development bank of Latin America. We are indebted to the Rosario Experimental Economics Lab (REBEL) and especially to Daniel Gomez for support and logistics. Cynthia Marchioni provided

excellent research assistance. This study is registered in the AEA RCT Registry and the unique identifying number is: AEARCTR-0003877.

REFERENCES

Ashraf, N., Bandiera, O. and Lee, S. (2018) Losing Prosociality in the Quest for Talent? Sorting, Selection, and Productivity in the Delivery of Public Services. *LSE Research Online Documents on Economics 88175*, London School of Economics and Political Science, LSE Library.

Banerjee, R., Baul, T. and Rosenblat, T. (2015) On self selection of the corrupt into the public sector. *Economics Letters*, **127**, 43 – 46.

Barfort, S., Harmon, N., Hjorth, F. and Olsen, A. L. (2015) Dishonesty and Selection into Public Service in Denmark: Who Runs the World's Least Corrupt Public Sector? *Discussion Papers 15-12*, University of Copenhagen. Department of Economics.

Besley, T. (2005) Political Selection. *Journal of Economic Perspectives*, **19**, 43–60.

Dal Bo, E., Finan, F. and Rossi, M. A. (2013) Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service. *The Quarterly Journal of Economics*, **128**, 1169–1218.

Deserranno, E. (2019) Financial incentives as signals: Experimental evidence from the recruitment of village promoters in uganda. *American Economic Journal: Applied Economics*, **11**, 277–317.

Finan, F., Olken, B. and Pande, R. (2017) Chapter 6 - the personnel economics of the developing state. In *Handbook of Economic Field Experiments* (eds. A. V. Banerjee and E. Duflo), vol. 2 of *Handbook of Economic Field Experiments*, 467 – 514. North-Holland.

Fischbacher, U. and Föllmi-Heusi, F. (2013) Lies In Disguise—An Experimental Study On Cheating. *Journal of the European Economic Association*, **11**, 525–547.

Hanna, R. and Wang, S.-Y. (2017) Dishonesty and selection into public service: Evidence from india. *American Economic Journal: Economic Policy*, **9**, 262–90.

Karlan, D. and Zinman, J. (2009) Observing unobservables: Identifying information asymmetries with a consumer credit field experiment. *Econometrica*, **77**, 1993–2008.

Olken, B. A. (2007) Monitoring corruption: Evidence from a field experiment in indonesia. *Journal of Political Economy*, **115**, 200–249.

Olken, B. A. and Pande, R. (2012) Corruption in Developing Countries. *Annual Review of Economics*, **4**, 479–509.

Svensson, J. (2005) Eight questions about corruption. *Journal of Economic Perspectives*, **19**, 19–42.