

Information Systems, Service Delivery, and Corruption: Evidence From the Bangladesh Civil Service

Martin Mattsson*

August 22, 2022

Abstract

Government bureaucracies in low- and middle-income countries often suffer from both corruption and slow public-service delivery. Can an information system – providing information about delays to the responsible bureaucrats and their supervisors – improve delivery speed? Paying bribes for faster service delivery is common, but does improving the average delivery speed reduce bribes? To answer these questions, I conduct a large-scale field experiment over 16 months with the Bangladesh Civil Service. I send monthly scorecards measuring delays in service delivery to government officials and their supervisors. The scorecards increase on-time service delivery by 11% but do not reduce bribes. Instead, the scorecards *increase* bribes for high-performing bureaucrats. The results are inconsistent with the predictions of several existing models of service delivery and corruption. Instead, I propose a model where bureaucrats' reputational concerns or shame constrain corruption. When bureaucrats' reputations improve through positive performance feedback, this constraint is relaxed, and bribes increase.

Key words: Public sector management, Service delivery, Corruption, State capacity

*National University of Singapore, martin.mattsson@nus.edu.sg. I would like to thank Nathan Barker, Gaurav Chiplunkar, Anir Chowdhury, Andrew Foster, Eduardo Fraga, Sahana Ghosh, Marina Halac, Ashraf Haque, Enamul Haque, Daniel Keniston, Mushtaq Khan, Ro'ee Levy, Imran Matin, Mushfiq Mobarak, Farria Naeem, Rohini Pande, Dina Pomeranz, Mark Rosenzweig, Nick Ryan, Yogita Shamdasani, Jeff Weaver, Jaya Wen, Fabrizio Zilibotti, and numerous seminar participants for helpful comments and suggestions. I also thank Muhammad Bin Khalid, Mahzabin Khan, and Ashraf Mian for excellent research assistance, IPA Bangladesh for outstanding research support. A randomized-controlled-trial registry and pre-analysis plan are available at: www.socialscisearch.org/trials/3232. This project was approved by Yale University IRB (Protocol ID 2000021565). This work was supported by the JPAL Governance Initiative (GR-0861), the International Growth Centre (31422), the Yale Economic Growth Center, the MacMillan Center for International and Area Studies, the Weiss Family Fund, and the Sylff fellowship.

1 Introduction

A government's capacity to implement its policies, secure property rights, and provide basic public services is paramount for economic development. To have this capacity, states need functioning bureaucracies with government officials motivated to carry out these tasks. While explicit incentive structures such as pay-for-performance contracts can change the behavior of government officials, they are often hard to implement without unintended consequences or political resistance (Banerjee et al., 2008; Dhaliwal and Hanna, 2017). Another approach is to improve systems that measure bureaucrats' performance. This may improve incentives by allowing supervisors to let job performance determine postings and promotions, a strong motivator for civil servants (Khan et al., 2019; Bertrand et al., 2020). Regular performance feedback may also make the performance more salient to the bureaucrats themselves, potentially leveraging government officials' intrinsic motivation (Prendergast, 2007; Banuri and Keefer, 2016).

This paper studies an information system designed to improve processing times of applications for changes to government land records in Bangladesh. Both slow public-service delivery and corruption are substantial problems in Bangladesh. Among households that paid a bribe for a public service, 23% stated that "timely service" was one of the reasons for the bribe (Transparency International Bangladesh, 2018). This suggests that slow service delivery on average may cause corruption, as some citizens and firms pay bribes to avoid having to wait for the service. This is a common hypothesis in the literature (e.g. Myrdal, 1968; Rose-Ackerman, 1978; Kaufmann and Wei, 1999), but the causal relationship has not been established empirically.

In an experiment with the Bangladesh Civil Service, I provide information on junior bureaucrats' performance using monthly scorecards sent to the bureaucrats and their supervisors. The scorecards, designed to reduce delays in processing ap-

plications for land-record changes, are based on data from an e-governance system. Two performance indicators appear on the scorecards: the number of applications processed within a time limit of 45 working days and the number of applications pending beyond that limit. The scorecards also show bureaucrats' performance on these indicators relative to all other bureaucrats in the experiment. The intervention is randomized at the level of the subdistrict land office and involves 311 land offices (60% of all land offices in Bangladesh), serving a population of approximately 97 million people.

The scorecards improve processing times. Using administrative data from more than a million applications, I estimate that they increase the share of applications processed within the time limit by 6 percentage points (11%) and decreased processing times by 13%. The effect is present throughout the 16 months of the experiment and is driven by improvements among offices that were underperforming relative to the median at baseline.

Despite their effect on processing times, the scorecards did not decrease bribe payments. I collect survey data on bribe payments from applicants; the point estimate for the effect on my preferred measure of bribes paid is an increase of BDT 940 (USD 11 or 15%). The lower bound of the 95% confidence interval is a decrease of 4%, thus the result is inconsistent with a substantial reduction in bribes paid. Using a randomized information intervention among surveyed applicants, I rule out that lack of information about the improved processing times is the reason for not seeing a negative effect on bribes paid.

The positive effect of the scorecards on bribe payments is concentrated among the offices that were overperforming relative to the median at baseline, where the scorecards had no effect on processing times. In the underperforming offices, where scorecards improved processing times, there is no effect on bribes. This suggests that there is no causal relationship between average processing times and bribes, because processing times can improve without bribes changing, and bribes

can increase without processing times changing.

I propose a model whose predictions are consistent with the experimental results. In this model, bureaucrats trade off reputation, bribe money, and the utility cost of effort. Their reputation is determined by their visible job performance along two dimensions, processing times and bribes taken. The scorecards are modeled as an increase in the visibility of processing times, making processing times more important for reputation, which in turn incents all bureaucrats to improve their processing times (akin to a substitution effect).¹ However, for overperforming bureaucrats, the increased visibility of their good performance also increases their reputation level, which reduces the marginal importance of reputation (akin to an income effect). For overperforming bureaucrats, the two effects run in opposite directions, and the overall effect predicted by the model on processing times is ambiguous. For underperforming bureaucrats, the two effects run in the same direction, and the model predicts improved processing times.

In the model, bureaucrats refrain from taking more bribes because bribes negatively affect their reputation. When the scorecards improve the reputation level of overperforming bureaucrats they reduce the marginal importance of reputation for their utility and cause them to take more bribes. For underperforming bureaucrats, the decrease in reputation from their low performance being more visible is counteracted by their increased effort, so the overall effect of the scorecards on reputation, and therefore also on bribes, is ambiguous.

This paper contributes to four strands of literature. First, it provides empirical evidence on the causal relationship between the average speed of public-service delivery and corruption. It has been shown that slow public-service delivery is positively associated with corruption (Kaufmann and Wei, 1999; Freund et al., 2016), and that individuals seeking services may have to pay bribes to reduce the

¹The increase in visibility can be interpreted as an improvement in the supervisors' ability to monitor this aspect of the bureaucrats' work, an increase in the salience of the information to the bureaucrats themselves, or both.

time between application and provision (Bertrand et al., 2007). In the theoretical literature, one view is that corruption allows individuals to circumvent excessive bureaucratic hurdles (Leff, 1964; Huntington, 1968). An opposing view is that corruption causes delays and red tape in public services, because making the *de jure* regulation more onerous allows government officials to extract more bribes (Rose-Ackerman, 1978; Kaufmann and Wei, 1999).² According to both views, we could reduce corruption by improving the speed of service delivery for everyone. However, I show that an intervention successfully targeting delays in service delivery did not decrease bribe payments in this context.

Second, this paper contributes to the literature on how incentives shape bureaucratic performance and corruption. There is an extensive literature on both monetary and nonmonetary explicit incentives (e.g., Duflo et al., 2012; Ashraf et al., 2014; Khan et al., 2016; Rasul and Rogger, 2018; Khan et al., 2019), and a growing literature on the effects of information systems within government bureaucracies (Dhaliwal and Hanna, 2017; Muralidharan et al., 2020; Callen et al., 2020; Raffler, 2020; Dal Bó et al., 2021; Banerjee et al., 2021; Dodge et al., 2021; de Janvry et al., 2021). Consistent with this literature, I show that increased transparency about individual civil servants' performance can improve public-service delivery, even without explicit incentives, and that this effect is persistent over time. This could be due to bureaucrats' career concerns (Niehaus and Sukhtankar, 2013a; Bertrand et al., 2020), or to their own sense of shame or pride. Complementing the empirical results, my model suggests that reputational concerns among bureaucrats create incentives for performance and limit corruption.

Third, this paper contributes to the literature on the effects of performance monitoring. The finding that providing performance feedback leads to improvements in performance, especially for underperformers, is consistent with the re-

²In Banerjee (1997), Guriev (2004), and Banerjee et al. (2012), both corruption and red tape emerge from the nature of public-service provision due to the principal-agent problem between the government and its bureaucrats. The experimental results can neither reject nor provide evidence in favor of these models.

sults of several other experiments (e.g. [Allcott, 2011](#); [Ayres et al., 2013](#); [Byrne et al., 2018](#); [Barrera-Osorio et al., 2020](#)), including recent work with Indian bureaucrats by [Dodge et al. \(2021\)](#).³ Most papers in this literature show smaller or even negative effects for high-performers. This paper adds to this literature by showing that these negative effects can spill over into domains not covered by the performance information, for which data is not typically collected. There is a substantial theoretical and empirical literature on the multitasking problem, i.e., how incentives for improving one indicator have negative spillovers by taking attention and resources away from other types of performance ([Holmström and Milgrom, 1991](#); [Finan et al., 2017](#)). My model suggests a different mechanism for negative spillovers, namely the decreasing marginal utility from reputation after receiving positive performance feedback.⁴

Fourth, this paper contributes to the literature on the determinants of bribe amounts. Some models and empirical evidence suggest that bribe payers' outside options and abilities to pay constrain bribe amounts ([Svensson, 2003](#); [Bai et al., 2019](#)), potentially leaving little room for applicant complaints or government monitoring to reduce corruption ([Niehaus and Sukhtankar, 2013b](#)). In other settings, such monitoring has been effective in reducing corruption ([Reinikka and Svensson, 2005](#); [Olken, 2007](#)).⁵ I show that, in my context, individual bureaucrats *can* increase bribes and that bribes are not fully determined by the applicants' willingness to pay. My model highlights how bureaucrats' concern for their reputation constrains bribes, thus explaining why bribes are substantially below applicants' willingness to pay for the service.

³[Ashraf et al. \(2014\)](#) and [Ashraf \(2019\)](#) show that privately provided social comparisons reduced the performance of low-performing healthcare and garment workers, while publicly announcing good performances increases performance for low-performing groups. [Blader et al. \(2020\)](#) show similar effects among truck drivers, but the results are reversed with a management practice establishing a cooperation-based value system.

⁴Also consistent with the literature on moral licensing ([Sachdeva et al., 2009](#); [Clot et al., 2018](#)).

⁵[Do et al. \(2021\)](#) show that doctors' fear of punishment, reputational concerns, or moral obligations cause them to take smaller bribes when treating patients with acute conditions.

The rest of this paper is organized as follows. Section 2 describes the context, experimental interventions, and data. Section 3 describes my empirical strategy. Section 4 presents the effects of the scorecards on processing times and bribes. Section 5 discusses mechanisms and how the results relate to theories of speed of service delivery and corruption. Section 6 describes the model of bureaucratic behavior and how it can explain the empirical results. Section 7 concludes.

2 Context, Experimental Intervention, and Data

2.1 Land Record Changes in Bangladesh

This paper studies land record changes (called *mutations* in Bangladesh) and the time it takes to process applications for them. When a parcel of land changes owners, either through sale or inheritance, the official record of ownership has to be updated and a new record of rights (*khatian*) issued to the new owner. An updated record of land ownership is crucial for maintaining secure property rights. Unfortunately, the burdensome and costly process of applying for land records is causing many new land owners to wait to apply until they need a record of rights. This means that the government land records substantially lag actual ownership, contributing to land disputes, one of the most severe legal problems in Bangladesh—where 29% of adults have faced a land dispute in the past four years ([Hague Institute for Innovation of Law, 2018](#)).

2.1.1 Structure of Bureaucracy

Applications for land record changes are processed by civil servants holding the position of Assistant Commissioner Land (ACL). Throughout the paper, I refer to ACLs as the *bureaucrats*. ACL is a junior position in the Bangladesh Administrative Service, the elite cadre of the Bangladesh Civil Service. Each subdistrict (*Upazila*) land office is headed by a single ACL and processing land record changes is a central duty of the ACL. In qualitative interviews, ACLs estimate spending between

25-50% of their working time on land record changes. Bureaucrats typically hold the position of ACL for one to two years; when an ACL is transferred, it is typically to a non-ACL position.⁶

The ACL is directly supervised by an Upazila Nirbahi Officer (UNO), the most senior civil servant at the subdistrict level. During periods when no ACL is assigned to an office, the UNO is responsible for the ACL's duties. The UNO has substantial power to influence the ACL's future career: through writing an Annual Confidential Report about the ACL's performance with the Ministry of Public Administration. The UNO is in turn supervised by a Deputy Commissioner (DC), the most senior bureaucrat at the district level. Throughout the paper, I refer to the UNOs and DCs as the *supervisors*.

2.1.2 Application Process

The *de jure* process for making a land record change starts when the new owner applies at the subdistrict land office. There is no competition between land offices for applicants, because each parcel of land is under the jurisdiction of a single subdistrict. The application is inspected by the office staff, who verify that the application has the required documents. The application is then sent to the local (*Union Parishad*) land office, which is the lowest tier of land offices. There, a Land Office Assistant verifies the applicant's claim to the land by meeting with the applicant and visually inspecting the land. The Land Office Assistant then sends a recommendation back to the subdistrict land office on whether to accept or reject the application. The application is then verified against the government land record. Finally, the ACL holds a meeting with the applicant where the application is formally approved. The applicant then pays the official fee of BDT 1,150 (USD 14) and receives the new record of rights.⁷ The subdistrict land office also changes

⁶Of the 617 ACLs I observe in my administrative data, only 10% held the position of ACL in more than one land office in the 60% of land offices covered by the experiment.

⁷Throughout the paper, I use a USD/BDT exchange rate of 84.3, the average exchange rate during the experiment. This exchange rate is not adjusted for purchasing power parity.

the official government land record to reflect the new ownership.

The government has mandated that applications should take no more than 45 working days to process, but in practice delays beyond this time limit are common.⁸ In my data, only 56% of applications in the control group were processed within the time limit, and the average processing time was 64 working days.

2.1.3 Bureaucrats' Discretionary Powers and Corruption

In practice, applicants also pay bribes. [Transparency International Bangladesh \(2018\)](#) estimates that the land sector is the second-largest receiver of bribes from Bangladeshi citizens and that the average bribe for a land record change is BDT 4,085. Appendix Figure [A1](#) shows that among the applicants in my survey, the average estimated bribe payment for a "normal person" like themselves is BDT 6,718 (USD 80).

The most common responses to the open-ended question of why a bribe was paid are akin to: "to get the work done" (39%), "to avoid hassle" (38%), and "for faster processing" (9%). This highlights the bureaucrats' power over applications along two dimensions. First, they can decide whether to accept or reject the application. Second, they can speed up or slow down the application, as well as create various hassles for the applicant.

Appendix Figure [A1](#) shows that the average stated valuation of getting the record of rights is BDT 1,610,998 (USD 19,110). This is more than two orders of magnitude larger than the estimate of the average bribe payment. Applicants' stated willingness to pay for having their application processed within the shortest realistic processing time (seven days) is, on average, BDT 2,189 (USD 26). Because this amount is substantially lower than my preferred estimate of the average bribe, it suggests that applicants are paying bribes not just for faster processing, but also for getting the approval.

⁸At the end of the experiment, the limit was shortened to 28 working days, but the design of the scorecards never changed. The scorecards were discontinued in March, 2020.

2.1.4 E-governance System for Land Record Changes

In February 2017, a new e-governance system for land record changes was gradually introduced to simplify the process for both applicants and bureaucrats.⁹ The e-governance system generates administrative data on each application made in the system. However, until the start of the experiment, this data was not used for measuring adherence to the 45-working-day time limit or for evaluating specific subdistrict land offices or ACLs.

2.2 Experimental Intervention: Performance Scorecards

The main experimental intervention consist of monthly scorecards designed to decrease delays in the processing of applications for land record changes. The content of the scorecards is automatically generated using data from the e-governance system. While the intervention is randomized at the land-office level, the scorecards were addressed to the bureaucrats (ACLs) and sent to each bureaucrat's land office, as well as to the offices of the supervisors (the UNO and DC). Two versions of the scorecard, one in English and one in Bengali, were sent out by courier in the first two weeks of each month.¹⁰ Offices in the treatment group were not informed that they would receive a scorecard before the start of the treatment, but the first scorecard was followed by phone calls to the ACLs, where it was confirmed that the scorecard had been received and where the indicators were explained.

Appendix Figure A2 depicts an example scorecard. The scorecard evaluates the bureaucrat's performance using two performance indicators: the number of

⁹Because the system was new at the time of the experiment, even in the subdistrict offices where it had been installed, not all applications were processed through it. The most common reason for not using the e-governance system was that the relevant *local* land office had not yet installed the system. Other reasons cited for using the paper-based system were problems with internet connectivity, new officials not yet trained to use the system, and temporary problems with the e-governance server.

¹⁰An email with a PDF version of the scorecard was also sent to bureaucrats who had listed their email addresses publicly or in the e-governance system. No scorecards were sent in January 2019 due to some bureaucrats' responsibilities the previous month were shifted to the December 2018 elections instead of their regular duties.

applications processed within 45 working days in the past month, where a higher number indicates a better performance, and the number of applications pending beyond 45 working days at the end of the month, where a lower number indicates a better performance.¹¹ The scorecards compare these numbers with the average numbers for all land offices in the experiment.¹² The scorecard also provides the office's percentile ranking for each indicator, with a short sentence and a thumbs-up or thumbs-down symbol reflecting the performance.

At the bottom of the scorecard the number of applications received by the office over the past six months is displayed. Offices vary substantially in terms of their workloads and the two performance indicators were chosen such that neither large nor small offices would have an inherent advantage in receiving a good score on average, as shown in Appendix Figure A3. Appendix A.2 provides more details about the rationale for choosing the performance indicators and analyzes the extent to which bureaucrats matter for the performance score.

2.3 Randomization and Implementation Timeline

Figure 1 depicts the randomized interventions and data collection. The scorecard intervention was carried out in two waves; each wave randomized all offices where the e-governance system had been installed at that time into either the treatment group or the control group. I conducted the first randomization in August 2018: 56 of the 112 land offices where the e-governance system had been installed were randomly selected to receive the scorecards. Appendix A.1 provides details about the stratification of the randomization. The first wave of scorecards started in September 2018. By April 2019, 199 additional offices had installed the e-governance system. 99 of these offices were randomly selected to receive the performance scorecards. The second wave of scorecards started in April 2019. The

¹¹The scorecards contained an explanatory note showing how the numbers in the scorecard are calculated and a phone number to call to ask questions.

¹²The initial comparison group was the 112 offices in the first randomization. After the second randomization, the group was expanded to include all 311 offices.

scorecards were sent out monthly until March 2020, when the COVID-19 pandemic caused the intervention to end.

2.3.1 Additional Intervention: Peer Performance List

To test for peer effects, a Peer Performance List was added to the scorecards for 77 randomly selected offices already receiving the scorecards. The addition took place in September 2019, a year after the first scorecards were sent out. The list contains the percentile rankings of the two performance indicators for all 77 offices. The list informed the bureaucrats that it had been sent to the 76 other offices, making it clear that their peers could observe their performance. Appendix Figure A2 shows an example of such a list.

2.4 Data

I use two main data sources, administrative data from the e-governance system for all land offices in the experiment and data from a survey conducted among applicants in the offices that were part of the first randomization wave. Table 1 shows summary statistics for both data sets.

2.4.1 Administrative Data

The administrative data is based on 1,034,688 applications from 311 land offices. The data contains information about the office in which the application was made, the application start date, the date it was processed, the decision to accept or reject the application, and the size of the land.¹³ The administrative data was downloaded from the e-governance system at the beginning of each month from August 2018 until December 2020.

For the main analysis, I use administrative data for applications made from August 13, 2018, one month before the start of the intervention, until January 20, 2020. I include applications made one month before the intervention if they had not been processed by the arrival of the first scorecard, as these were partially

¹³Information about the applicants is not available for research purposes due to privacy concerns.

treated. I chose to use data until 45 working days before March 26, 2020—this is when the COVID-19 pandemic caused a long general holiday and the end of the intervention.

The processing time is measured as the number of working days between the application start date and the date the processing was finalized. For applications that had not yet been processed by December 2020, I impute the processing time by taking the mean of actual processing times that were longer than the time the application had been pending.¹⁴ Appendix A.3.1 provides more information about the administrative data.

2.4.2 Survey Data

I collected the survey data in two rounds from applicants who applied in the 112 offices that were part of the first randomization wave. To create the sample of applicants, enumerators were stationed outside land offices to interview all applicants entering the office for the purpose of a land record change, regardless of their stage in the application process. The enumerators stayed outside a specific office for at least two days and until they had completed at least 20 interviews. This first-round interview focused on the basic details of the application and applicant, the applicant's expectation for the application processing time, and the applicant's willingness to pay for faster processing.

The follow-up interviews, conducted by phone approximately three months after the initial interview, focused on the outcome of the application and bribe payments.¹⁵

To avoid extreme outliers potentially caused by enumeration errors, all continuous variables from the survey are winsorized at the 99th percentile. Appendix A.3.2 provides more information about the survey data.

¹⁴2% of applications that had not yet been processed by December 2020. The estimate of the scorecards' effect on the share of not-yet-processed applications is a decrease of 0.4 percentage points.

¹⁵Enumerators were not informed about which offices had received the scorecards or whether they were calling a respondent from a treatment or control office.

There are two measures of bribe payments. The first is based on a question of how much the applicant thinks it is "normal for a person like yourself to pay." For the 63% of respondents who were willing to answer this question, the amount is recorded as the variable *typical payment*. The average response is BDT 6,718 (USD 80) or 1.5 months of the sample's average per capita household expenditure. 73% of the responses were nonzero amounts. The second measure is based on a series of questions about each applicant's actual payments to any government officials or agents assisting with the application. The outcome variable *reported payment* is the sum of the reported amounts. The average reported payment is BDT 1,477 (USD 18), and 27% of respondents provided a nonzero value. Among those reporting a nonzero amount, the average amount was BDT 5,283 (USD 63).

The typical payments measure is my preferred measure of bribes. The main problem with the reported payment measure is the large number of zero responses, suggesting that it is an underestimate of bribes paid. However, I have no reason to believe that either of the two measures is biased differently between the treatment and control offices. Throughout the paper I show that the main results are robust to using either of the two measures.

Of the 3,213 applicants from the in-person interviews, 2,869 were successfully interviewed in the follow-up phone call, for a total attrition rate of 11%. The estimated effect of the scorecards on the attrition rate was 3 percentage points (p-value = 0.08). In Appendix A.3.3, I discuss attrition and non-responses in detail and use Lee bounds (Lee, 2009) to show that the differential attrition is not sufficiently large to substantially affect the main findings from the survey data.

2.5 Additional Intervention: Providing Information to Applicants

An additional experimental intervention giving applicants information about processing times was also carried out together with the survey. The motivation for this intervention was to ensure that applicants knew about the improvements in pro-

cessing times. On randomly selected days the enumerators gave applicants leaflets that described how the median processing time for all land offices had been substantially reduced over the past six months and that a new e-governance system had been implemented. The information was the same in treatment and control offices.

Appendix Figure A2 shows an English translation of the leaflet. Due to some non-compliance with the treatment assignment by the enumerators delivering the treatment, I use the median treatment delivered in a land office survey day as the main treatment variable when analyzing the effect of the information intervention. Appendix B.3.1 discusses this choice and shows the robustness of the main results to using alternative treatment variables.

2.6 Balance of Randomization

Panel A of Appendix Table A1 shows balance-of-randomization tests for variables from the administrative data. To exclude all data that the scorecards could have affected, I restrict the data to applications made at least 45 working days before the start of the experiment. Applications not processed by the start of the experiment were assigned an imputed processing time based on the time they had been pending at the start of the experiment, using the imputation procedure described in Section 2.4.1. There are no statistically significant differences between scorecard and control offices before the start of the experiment. Making the same comparisons separately for offices over- and underperforming at the start of the experiment similarly yields no statistically significant differences.

Panel B of Appendix Table A1 shows that the scorecards did not affect the composition of applicants or applications in the survey data. This is not a traditional balance-of-randomization test, since the treatment may have affected which applicants decided to apply and what type of applications to make. However, I find no evidence for any such changes in composition. Comparing the age, gender, and in-

come of the applicants; the size and value of the land the applications are for; and the stages that the applications are in at the time of the first interview there are no statistically significant differences. Making the same comparisons for offices over- and underperforming at the start of the experiment yields one difference statistically significant at the 5% level: among overperforming offices, a smaller share of applications were approved at the time of the in-person interview in the treatment offices. One more difference is statistically significant at the 10% level: among overperforming offices, a smaller share of applicants are female in the treatment offices. Overall, I perform 27 balance tests in the survey data of which only one is significant at the 5% level, and an additional one at the 10% level.

3 Empirical Strategy

3.1 Overall Effects

To estimate the effects of the scorecards, I use the following regression specification:

$$Outcome_{ait} = \alpha + \beta Treatment_i + Stratum_i + Month_t + \varepsilon_{ait} \quad (1)$$

where $Outcome_{ait}$ is an outcome for application a , in land office i , made in calendar month t . $Stratum_i$ are randomization stratum fixed effects. Since no randomization stratum overlap the two randomization waves, these fixed effects also control for randomization-wave fixed effects. $Month_t$ denotes fixed effects for the month the application was made.¹⁶ For the main results, I provide p-values testing the null hypothesis of no effect using conventional standard errors clustered at the land-office level, as well as p-values based on randomization inference.¹⁷ Each observation is weighted by the inverse of the number of observations in each land office.

¹⁶In the survey data, consistent with the cleaning of other continuous variables, the application month variable is winsorized at November 2018, so that all application dates in or before November 2018 take the same value. A separate indicator variable controls for missing start-date values.

¹⁷The randomization inference is implemented using the Stata command *randcmd*, and the reported p-value is from the randomization t-test calculated using 9,999 iterations (Young, 2019).

This has three benefits. First, it makes the regression estimate the average effect of the scorecards on a land office, the unit relevant for studying changes in bureaucrat behavior. Second, using these weights, the analysis of the administrative data and the survey data estimates the same effect. Third, the weighting improves the estimates' precision by weighting each cluster equally in the analysis.¹⁸

The two additional randomized interventions—the peer performance lists and the information intervention to applicants—are excluded from the main specification, because they are not the main treatments being evaluated. For the two main outcomes—processing times and bribe payments—the full specifications, including the scorecard treatment, the relevant additional randomization, and the interaction, can be found in Appendix Tables A2 and A3. These tables show that neither of the two additional experiments have substantial interactions with the scorecard treatments, validating my approach to analyze the scorecard treatment separately.

3.2 Heterogeneous Effects by Baseline Office Performance

To better understand the mechanisms behind the overall effects, I separate offices by their baseline performance and estimate the effect of the scorecards separately for offices performing above and below the median at baseline.¹⁹ I define baseline performance as the average of the two performance indicators' percentile rankings in the first month of treatment. I then separate all offices into *overperformers* (above the median baseline performance) and *underperformers* (below the median baseline

¹⁸As illustrated in Appendix Figure A3a, some land offices receive more applications than others. This causes 57% of the administrative data sample to come from the 25% largest offices, while only 6% come from the 25% smallest offices. For a discussion of why weighting observations by the inverse of the number of observations in a cluster improves precision, see <https://blogs.worldbank.org/impactevaluations/different-sized-baskets-fruit-how-unequally-sized-clusters-can-lead-your-power>.

¹⁹Heterogeneity in the effects of performance-information provision between high and low performers has been recorded in several settings (e.g., Allcott, 2011; Byrne et al., 2018; Ashraf, 2019; Barrera-Osorio et al., 2020; Dodge et al., 2021). This was the only heterogeneity test based on office characteristics specified in the preanalysis plan.

performance).²⁰ Since the classification of offices uses data only from before the first scorecard was delivered, it is not affected by the treatment.

I use the following regression specification to estimate the effect of the scorecards on the two types of offices separately:

$$y_{ait} = \alpha + \beta_1 \text{Treatment}_i \times \text{Overperform}_i + \beta_2 \text{Treatment}_i \times \text{Underperform}_i + \gamma \text{Overperform}_i + \text{Stratum}_i + \text{Month}_t + \varepsilon_{ait} \quad (2)$$

where β_1 is the estimated effect of the scorecards for offices overperforming at baseline, β_2 is the effect for offices underperforming at baseline, and γ is the difference between overperforming and underperforming offices in the control group.²¹ As in the estimation of the overall effects, standard errors are clustered at the land-office level and the regressions are weighted by the inverse of the number of observations in each land office.

4 Results: Effects of the Scorecards

4.1 Effect on Processing Times

Table 2 shows that the scorecards increased the number of applications processed within the government time limit and improved processing times overall. Each column presents the result of a regression using the specification in Equation 1. Column (1) shows the estimated effect of the scorecards on a binary variable in-

²⁰I classify 112 offices in the first randomization wave into over- and underperformers by comparing them to the median performance among these 112 offices. For the offices in the second randomization wave, I compare them to the median performance of all 311 offices in the experiment at the time of their first scorecard. This makes the *overperformer* and *underperformer* classifications correspond to the relative performance presented in the first scorecards.

²¹To test the hypothesis that the treatment had the same effect on offices overperforming and underperforming at baseline, I use a similar regression but where the first treatment variable is not interacted with the indicator variable for whether the office overperformed at baseline. I then test the hypothesis that the coefficient on the treatment variable interacted with with the indicator variable for whether the office was underperforming at baseline is zero. This test's p-value is reported as "P-value subgroup diff." in the regression tables reporting the heterogeneous effects.

dicating whether the application was processed within the 45-working-day time limit. The scorecards increased applications processed within the limit by 6 percentage points (11%). Column (2) shows that the scorecards led to a 13% reduction in overall processing times by estimating the effect on the natural logarithm of the processing time.²²

For column (3), I create a *time index* of the two outcomes used in columns (1) and (2).²³ The estimated effect of the scorecards on the time index is 0.13 standard deviations (conventional p-value = 0.028; randomization inference p-value = 0.037).

4.1.1 Effect on the Distribution of Processing Times

Figure 2 shows the effect using two overlaid histograms with minimally processed data on of processing times. In the treatment offices, more applications were processed within the 45-working-day time limit. The effect is relatively evenly spread over the whole span from 0 to 45 working days, with no substantial bunching just before the 45-day limit. This is to be expected given that the process for approving an application is relatively long and depends on several individuals, as described in Section 2.1. Thus, even if the bureaucrat cares only about maximizing the share of applications processed within 45 working days, the processing time target has to be lower than 45 working days.

The figure also shows that the processing times that are reduced in frequency by the scorecards are in the whole span from 55 working days and up. This is also reasonable, given that the scorecards emphasize both processing applications within the 45-working-day limit and reducing the number of applications pending beyond the limit. Overall the spread of the effect in the distribution of process-

²²The exact effect is -12.6 log points, which is equivalent to an 11.9% decrease. For simplicity, I will describe log point changes as percentage changes throughout the paper.

²³I created the index by first taking the negative of the log processing time so that a higher value indicates a better performance, then recasting the two outcome variables as standard deviations away from the control-group mean, and finally taking the sum of the two standard deviations and rescaling them so that the index has a standard deviation of one in the control group.

ing times alleviates the concern that bureaucrats are "gaming" the scorecards by speeding up the processing only of applications that would otherwise have been processed just outside the time limit.

4.1.2 Effect Over Time

It is possible that the initial effect of the scorecards is different from the long-run effect. Figure 3 shows the 10-working-day moving average of the share of applications processed within the time limit by application date relative to the start of the experiment. The first and second vertical lines indicates the date 45 working days before the first scorecards were sent and the date of sending them, respectively. Applications made between these dates were partially affected by the scorecards, as the bureaucrats received the scorecards while processing these applications.

Starting just before the first scorecards were sent, we see that the treatment group process more applications on time relative to the control group. With a few short exceptions, the treatment offices continued to process a higher share of applications within the time limit until the end of the experiment.²⁴

4.2 Effect on Bribe Payments

Table 3 shows that the scorecards did not lead to a decrease in bribe payments. Instead, the estimated effect on bribes is positive, though this increase is not statistically significant. As described in Section 2.4.2, data on bribe payments was collected using two separate survey questions. The first question asked about the typical bribe payment "for a person like yourself." When this measure is used, the column is marked as *typical*. The second question asked about each payment made by the applicant. When this measure is used, the column is marked *reported*.

Columns (1) and (2) of Table 3 show the effect on the amount of bribes paid.

²⁴Estimating the effect on the time index for applications made after the end of the intervention (March 26, 2020) until 45 working days before the the end of my data (September 28, 2020) yields an estimate of 0.07 standard deviations. This suggests that there is only a small amount of persistence in the effect of the scorecards. However, this estimate should be interpreted carefully—the COVID-19 pandemic drastically changed the operating conditions for the bureaucrats during this period.

Column (1) shows that the effect on the perceived typical payment was BDT 940 (USD 11), a 15% increase (conventional p-value = 0.130; randomization inference p-value = 0.159). The lower bound of the confidence interval is a decrease of 4%, ruling out a meaningful decrease. Column (2) estimates that the scorecards increased reported bribe payments by BDT 297, a 23% increase. Columns (3) and (4) show that there is no effect on the propensity to report a nonzero bribe. This can be interpreted as the scorecards having no effect on the extensive margin of bribe payments. Another interpretation is that the intervention did not affect applicants' willingness to talk about bribe payments in the survey. In columns (5) and (6), the sample is restricted to those who reported nonzero bribe payments. Bribe payments increased by 18% for typical payments and 25% for reported payments.²⁵

4.3 Heterogeneity of Results by Baseline Office Performance

Table 4 uses the empirical strategy from Section 3.2 to show that the scorecards' effect on processing times is driven by offices that were underperforming at baseline, while the effect on bribe payments is driven by offices that were overperforming at baseline. Column (1) shows that for offices that were underperforming at baseline, the estimated effect on the time index is an increase of 0.24 standard deviations (conventional p-value = 0.007; randomization inference p-value = 0.012).²⁶ For offices overperforming at baseline, the effect is just 0.03 standard deviations.

Columns (2) and (3) in Table 4 show that the positive effect on bribe payments is entirely driven by the offices that were overperforming at the start of the experiment. Column (2) shows that the effect of the scorecards on estimated typical bribe payments among offices overperforming at baseline is an increase of BDT 2,069, equivalent to 38% (conventional p-value=0.008; randomization inference p-value

²⁵Again these effects have two interpretations. Either the scorecards affected only the intensive margin of bribe payments, or the scorecards increased bribe payments for at least those applicants who were willing to describe what bribes they paid, but potentially also for other applicants.

²⁶The randomization p-value is 0.023 when using the Westfall-Young multiple-hypothesis testing method to adjust for that I am testing two hypotheses (Young, 2019).

= 0.016).²⁷ The effect on offices underperforming at baseline is close to zero.

Appendix Tables A5 and A6 show that the heterogeneous effects on processing times and bribes by baseline performance remains similar after controlling for other office baseline characteristics interacted with the treatment. However, the association between the effect size and baseline performance should still not be interpreted as a causal relationship. Baseline performance is not randomly assigned, and it is plausible that unobserved office characteristics are associated with both performance and the treatment effect size. Instead, I interpret the results as describing which type of offices reacted to the scorecards.

Even without a causal interpretation, the heterogeneity in the results is surprising: overperforming offices did not change their behavior in terms of processing times, but among these offices bribe payments increased. Sections 5 and 6 investigate this further.

4.4 Effects in Survey Data: Visits, Time, and Satisfaction

In columns (1)-(3) of Appendix Table A4, I use survey data from the 112 offices in the first randomization wave to show that the scorecards reduced the number of visits to land offices, increased the share of applications processed within 45 working days, and decreased overall processing times. The results on processing times are not statistically significant but the heterogeneity between offices over- and underperforming at baseline are consistent with the results from the administrative data. Column (4) shows the estimated effects of the scorecards on applicant satisfaction, which is negative but small and not statistically significant.²⁸ Splitting up this effect between offices that were underperforming and offices overperforming at baseline, the negative effect is driven by offices that were overperforming

²⁷The randomization inference p-value is 0.032 when using the Westfall-Young multiple-hypothesis testing method to adjust for that I am testing two hypotheses.

²⁸Satisfaction was measured in the follow-up phone survey by asking applicants "Overall, how satisfied are you with the processing of your application?" The respondent could answer the question on a five-point scale ranging from very satisfied to not satisfied at all. The response was then transformed into standard deviations from the control group mean.

at baseline, which is consistent with the observation that the scorecards increased bribe payments without improving processing times in these offices.

4.5 Robustness Tests

The main results for the effects of the scorecards on processing times and bribes, as well as the heterogeneity in these effects, is robust to a range of alternative specifications. Appendix Tables A7 and A8 show the results using various combinations of controls, weights, and winsorizations of the bribe amounts, as well as including only bribes given directly to government officials while ignoring fees paid to agents. All alternatives to the main estimates are of the same sign and of similar magnitude, but some of them are not statistically significant. Appendix Table A9 shows the effects, estimated at the office-month level, on the number of applications processed within 45 working days and the number of applications pending beyond 45 working days, as well as those figures' corresponding percentile rankings. The point estimates suggest that the scorecards improved all four of these outcome variables, driven by improvements in offices underperforming at baseline, but only the results for underperforming offices are statistically significant.

Appendix Table A10 shows that the heterogeneity found in the effects is robust to three alternative measures of baseline performance. First, I use three months of baseline data to establish the baseline performance—the results are similar to when using one month of baseline data. Second, I separate offices by the performance quartiles and create four, instead of two, performance groups. The positive effect on the processing time index is monotonically decreasing in being in a better baseline performance group. The effect on bribes is monotonically increasing in being in a better performance group, except for in the offices between the 50th and 75th baseline percentile ranking where the estimated effect is slightly larger than in the group of offices above the 75th baseline percentile ranking. Finally, I show that the continuous baseline performance ranking is negatively associated with the

size of the effect on the time index but positively associated with the effect on bribe payments, though the latter association is not statistically significant.

4.6 Unintended Consequences

A common problem of quantitative performance measures is that they often lead to gaming of the measures or other unintended consequences. In Appendix B.1, I test for four such potential unintended consequences. One, if bureaucrats allow fewer applicants to start applications, then their scorecards may improve, provided that the lower number of applications helps them process a larger share of the applications within the time limit. Two, if bureaucrats allowed applications selectively, such that the average application was easier to process within the time limit, then their scorecards may improve. Three, the scorecards may lead bureaucrats to make worse decisions regarding accepting or rejecting applications. Four, bureaucrats may divert attention from applications not made in the e-governance system, because those applications do not count toward the scorecards. I find no evidence for large unintended consequences, except for suggestive evidence of a higher incorrect rejection rate in offices overperforming at baseline, potentially as a response to applicants not being willing to pay the new higher bribes.

5 Mechanisms and Implications for Existing Theories

5.1 Two Potential Mechanisms for the Effect on Processing Times

The scorecards could improve performance of bureaucrats through two main channels. First, the information that supervisors get may improve their ability to incentivize bureaucrats by facilitating better promotions and more attractive postings for those bureaucrats with good scorecards. This is an example of the widely studied mechanism of increased information enabling better contracts that improve

output (Holmström, 1979).²⁹

Second, bureaucrats may change their behavior due to receiving the scorecards themselves. For bureaucrats, receiving information about their processing times each month may increase this information's salience, causing it to be more important for their personal sense of pride in their work. Since the scorecards were sent to both bureaucrats and their supervisors, I cannot separately estimate the importance of these two mechanisms, and I refer to them collectively as *reputational concerns*.

5.2 Effect of Peer Performance List

It is possible that information flows about performance *between* bureaucrats at the same level in the organizational hierarchy create an additional incentive for improved performance (Mas and Moretti, 2009; Blader et al., 2020). I measure such a peer effect, as an addition to the effect of the scorecard, by estimating the effect of the peer performance list sent to a randomly selected subgroup of the offices receiving scorecards, as described in Section 2.3.1.

Appendix Table A3 shows no substantial effect of the peer performance list on processing times. Using only data from land offices receiving performance scorecards, columns (1) and (3) shows that the effect on processing times is close to zero. Columns (2) and (4) use the full data set and estimate the effect of being in the the scorecard and the peer performance list treatment groups simultaneously, both before and after the performance list intervention started. This is done using an indicator variable for the peer performance list treatment that takes the value of one for applications made in offices receiving the peer performance lists and a *Post* indicator variable for applications made later than one month before the peer performance list intervention started. The point estimate is a minor improvement

²⁹It is also possible that bureaucrats care about their supervisors receiving information about them for other reasons, such as the shaming effect of having a negative performance being shown to a superior.

in performance and not statistically significant.³⁰

5.3 Implications for Theories of Service Delivery and Corruption

There are several reasons why bribes may be causally related to processing times. In the literature on this relationship, two opposing views lead to drastically different policy conclusions. One view is that bribes "grease the wheels" by providing incentives to bureaucrats and allowing for excessively onerous red tape to be circumvented (Leff, 1964; Huntington, 1968). In this view, rooting out corruption would decrease the speed of service delivery and increase inefficiencies of excessive bureaucratic control. On the other hand, improving the speed of service delivery through better technology or more personnel would decrease bribery, because the need for bribes would decrease.

Another view is that opportunities for corruption is the reason for excessive red tape and delays in public services, as making more onerous hurdles allows officials to extract more bribes in exchange for avoiding these hurdles (Rose-Ackerman, 1978; Kaufmann and Wei, 1999). Under these circumstances, providing the bureaucracy with more personnel or resources would not improve processing times. However, service delivery could be improved by eliminating corruption.³¹

A common prediction for both types of theories is that faster average service delivery would reduce bribes. However, the results in Section 4 show that the scorecards did reduce processing times, but not bribes. This is true even for the offices that were underperforming at baseline and improved their processing time the most. Furthermore, as bribes increased in overperforming offices, there was no corresponding change in processing times among these offices. Together these

³⁰The effects from the standard scorecard treatment as measured by the coefficients on *Scorecard* and *Post* × *Scorecard* can account for virtually all of the effect estimated by my main specification in Table 2, showing that the effect of the scorecards is not driven by the inclusion of the Peer Performance Lists.

³¹The scorecard experiment was originally designed to test a model of this type but as that specific model is rejected together with the larger class of models described here I do not include a detailed description of that model in this paper. See pre-analysis for more details.

results are suggestive evidence against a causal relationship between average processing times and bribe payments in this context.³²

5.3.1 Does Knowing About Faster Processing Times Decrease Bribes?

One potential reason for the lack of a negative effect from the faster processing times on bribe payments could be that the information about the improvement in processing times had not yet disseminated among applicants. The information treatment was designed to test this hypothesis, by informing applicants about improvements in processing times, is described in Section 2.5. Column (2) of Appendix Table A11 shows that the effect of this intervention on expected processing times at the time of the in-person survey was a reduction of 5%.³³ Appendix B.3 shows that the information treatment did not affect bribes, neither by itself nor in combination with the scorecards. Together these two results show that even when applicants are informed about the improved processing times, faster processing times do not reduce bribe payments.

5.3.2 Implications for Other Theories

The increase in bribes among the offices that were overperforming at baseline is also inconsistent with models where it is an applicants' outside option or ability to pay that determines the bribe levels (Svensson, 2003; Niehaus and Sukhtankar, 2013b). If the bribe level was fully determined by the applicants' outside option or ability to pay, it could not have been increased by a positive scorecard, without any observable change in service quality. However, this result is most likely dependent on the market structure of the interaction in which the bribe is paid. In this context, the land office is the only institution that can make the required land-record

³²While the scorecards were designed to reduce processing times, it is possible that they also affected other aspects that matters for corruption, such as the bureaucrat's reputation, and that such changes canceled out the causal effect of average processing time. I can therefore not definitively rule out that average processing times have a causal effect on bribe payments.

³³Appendix B.3.1 discusses the noncompliance with the treatment assignment in the delivery of this intervention and shows the robustness of the results to alternative definitions of the treatment variable.

change, and there are no close substitutes for this service. Therefore, the bribe level is expected to be determined mainly by other factors. Had there been competition for applicants between land offices, or a close alternative to a land-record change, it is plausible that these outside options would have been more important in determining the bribe level (Svensson, 2003; Bai et al., 2019).

A different class of models exists, in which government officials could extract more bribes if they wanted to, but they choose not to do so because they face a trade-off between taking bribes and achieving some other objective. This trade-off could be between bribe money and the risk of getting caught (Becker and Stigler, 1974; Olken, 2007; Niehaus and Sukhtankar, 2013a), but it could also be between bribes and altruistic motivations or reputation concerns (Do et al., 2021). In the next section, I will develop such a model and show that it is consistent with the results of the experiment.

6 Model of Bureaucrat Behavior

In this section, I propose a model that explains the results of the experiment.³⁴

6.1 Model Setup

6.1.1 The Bureaucrats' Objective Function

In the model, bureaucrats get utility from a reputation term, which is a function of bribe money taken and visible job performance in terms of processing times.³⁵ Bureaucrats get disutility from effort, but effort is needed to process applications within the time limit, which improves the bureaucrat's reputation. Bribe money has decreasing marginal utility, while effort has increasing marginal disutility. There is decreasing marginal utility from reputation in performance and

³⁴A formal model is available from the author upon request.

³⁵The reputation term captures both possible mechanisms for the scorecards' effect on processing times described in Section 5.1. I.e., the reputation term represents reasons for why the bureaucrats care about what their supervisors think of them, as well as psychological reasons that are internal to the bureaucrats.

increasing marginal disutility through the reputation mechanism in bribe money taken. The intuition for why bribes reduce the bureaucrat's reputation is that for bureaucrats consistently asking for higher-than-normal bribes, a reputation will build over time that this bureaucrat is corrupt. Similarly, if the bureaucrat consistently outperforms peers in terms of processing times, this will build good reputation over time.

The scorecards are modeled as an increase in the visibility of processing times. The scorecards thereby increase the importance of processing times for the bureaucrats' reputation.

I assume that bribe money and visible performance are complements in generating reputation, or equivalently, that honesty (the absence of bribe-taking) and performance are substitutes. In terms of bureaucrats' career prospects, this assumption is based on the idea that some corruption is acceptable as long as bureaucrats are performing their duties well and that poorly performing bureaucrats still have a good career in the bureaucracy, as long as they follow the official rules. If bureaucrats are both corrupt and poorly performing this could endanger their careers, while honest, high-performing bureaucrats still can't be promoted much faster than their colleagues, because most promotions are based on seniority. An alternative rationale for this assumption is moral licensing, which suggests that the internal shame a bureaucrat may feel from immoral behavior in one aspect of their work can be alleviated by acting better in another aspect.

Finally, I assume that bureaucrats differ only in how much they value their reputation. This could be because of differences in the valuation of future career prospects or differences in intrinsic motivation to be honest and perform well.³⁶

³⁶Prendergast (2007) and Hanna and Wang (2017) provide evidence that differences in the intrinsic motivations of government officials can be important for public-service delivery and corruption. Bertrand et al. (2020) provide evidence for the importance of differences of future career prospects.

These differences is what generate over- and underperforming bureaucrats.³⁷

6.1.2 Abstractions From Reality

In the model, the applicants simply pay the bribe amount that the bureaucrats are demanding. While this is clearly a simplification, Appendix Figure A1 shows that even the largest estimate for the average bribe is just 0.1% of the average stated value of the record of rights. This indicates that applicants' willingness to pay for the service is not an important determinant of the bribe value. Instead, what determines the amount of bribes that the bureaucrats extract in the model is the trade-off between bribe money and reputational concerns.

The model also abstracts away from bureaucrats buying reputation or career advancement using money. This avoids bureaucrats taking as-high-as-possible bribes, then use the money to regain their reputation. The results would be the same if bureaucrats would pay for their position in the first place but then could not bribe their way to future career advancement.

6.2 Model Predictions

The theoretical model has two main testable predictions. In what follows, I describe these predictions, the intuition behind them, and how I test them empirically.

6.2.1 Effects of the Scorecards on Processing Times

The scorecards have two effects on processing times, an *incentive effect* and a *reputation-level effect*. These effects are akin to the substitution and income effects from a wage increase in a labor-supply model. The incentive effect leads to an improvement in processing times for all bureaucrats. This is because the scorecards increase the visibility of the bureaucrats' performance in terms of processing times and there-

³⁷The predictions of the model still hold even if the differences are driven by other bureaucrat characteristics such as ability. I choose to model the differences as in the valuation of reputation because it is consistent with the observation that overperforming bureaucrats collect fewer bribes than underperforming bureaucrats in the control group, as shown in Table 4.

fore the marginal effect it has on utility. This causes bureaucrats to provide more effort to improve processing times.

For overperforming bureaucrats, the reputation-level effect from the scorecards on processing times is negative because the scorecards increase their reputation by making their good performance more visible. Since reputation has decreasing marginal utility, this decreases the marginal utility of reputation and reduces the optimal amount of effort overperforming bureaucrats exert to process applications on time. As the incentive effect and reputation-level effect move in opposite directions for overperforming bureaucrats, the overall effect of the scorecards on processing times is ambiguous and depends on which of the two effects is stronger.

For underperforming bureaucrats, the reputation-level effect from the scorecards on processing times is positive because the scorecards highlight their poor performance and lower their reputation level. Hence, the incentive effect and reputation-level effect move in the same direction, and the model predicts that the scorecards will improve processing times among underperforming bureaucrats. The model's prediction is consistent with the scorecards improving processing times for offices underperforming at baseline, while the effect is close to zero for offices overperforming at baseline.

6.2.2 Effects of the Scorecards on Bribes

The second prediction relates to the effect of the scorecards on bribes taken from applicants. In the model, the bureaucrats can increase bribes by simply asking applicants for more money. What constrains bureaucrats from extracting more bribes is the negative marginal effect it has on their reputation. When the scorecards improve overperforming bureaucrats' reputation, the marginal effect bribes have on utility through the reputation channel becomes less negative. This leads to an increase in bribes taken by overperforming bureaucrats when they receive the scorecards.

Underperforming bureaucrats, suffer an initial decrease in reputation from the

information provision, but since effort increases in response to the scorecards, the overall effect on reputation could be positive or negative. Therefore, the effect of the scorecards on bribes is ambiguous for underperforming bureaucrats. The model's prediction is consistent with the scorecards increasing bribes paid in offices overperforming at baseline, while the effect is close to zero for offices underperforming at baseline.

6.3 Alternative Explanations for the Effects of the Scorecards

Appendix [B.2](#) discusses four potential alternative explanations for the experimental results. First, that the scorecards increased applicants' willingness to pay bribes or increased the opportunity cost of bureaucrats' time. This explanation cannot explain the results, because it is not the offices improving their processing times that increases the amount of bribes paid. Second, that supervisors shifted monitoring attention away from overperforming offices leading to increased bribes. This explanation is dependent on an asymmetry where an increase in the monitoring attention on underperforming offices either did not occur or did not decrease bribes there. Third, that the supervisors' reputation level increased, causing them to demand higher bribes. This is unlikely, because land-record changes constitute a much smaller share of supervisors' responsibility than they do for the bureaucrats. Fourth, that bureaucrats use the scorecards in negotiations with applicants as "proof" that they are able to process the application quickly. This is not consistent with applicants in overperforming offices' expectations on processing times not improving, shown in Appendix Table [A11](#). It is also inconsistent with the lack of an effect from the information treatment on bribe payments, shown in Appendix Table [A2](#).

7 Conclusion

I have shown that a management information system – providing information about individual bureaucrats’ performance – can improve this performance. This effect is present despite the absence of explicit performance incentives, and it is persistent, at least over a 16-month period. The system, made possible by an underlying e-governance system, is an example of how new technologies are creating opportunities for improved management in the public sector.

To create a rough estimate for the value of the improved processing times, I multiply the applicants’ average stated valuation of having their application processed one day faster by the reduction in the total number of processing days due to the scorecards.³⁸ For the 155 offices receiving the treatment, the approximate value is USD 9.7 million per year.³⁹ This value should be interpreted carefully—it relies heavily on the value stated by the applicants. However, the number is more than two orders of magnitude larger than the cost to implement the scorecards, which was approximately USD 40,000 per year.

The overall welfare effect of the intervention becomes less clear when taking into account the effect on bribe payments. Multiplying the effect of the scorecards on typical payments with the number of applications in the treatment area results in an estimate of the effect on total bribes paid of USD 6.6 million per year.⁴⁰ Furthermore, the scorecards have a negative but not statistically significant effect on stated satisfaction. Overall, there is no strong evidence that the scorecards had either a positive or a negative effect on average applicant welfare.

More than half of Bangladesh’s land offices took part in the experiment, mak-

³⁸I calculate the value of having the application being processed one day faster using the following formula: $\frac{\text{Value of processing in 7 days}}{\text{Expected processing time from survey date}-7}$. All the information comes from the in-person survey, which took place before the application was actually processed. Applicants with an expected processing time of 7 days or less are excluded from the calculation.

³⁹See Appendix B.4 for details on the number of processing days saved per year.

⁴⁰If instead I use the effect on the reported payment, the total increase is USD 2.1 million per year.

ing it plausible that the results are externally valid within Bangladesh. However, while I designed the scorecards in collaboration with the government, they were produced and distributed by a nonprofit research organization. Hence, one should be cautious when extrapolating the results from the experiment to a potential scale-up by the government itself.

The empirical results and the model have several policy implications. First, they show that improving the speed of public service delivery is not an effective tool for reducing corruption in this context. Two important features of my setting are that there are no close substitutes to the public service and that the bureaucrats can control both the service delivery speed and whether the service is delivered at all. Further research is needed to investigate if service delivery speed is more important for bribes when bureaucrats can only control the speed of service delivery or when closer substitutes to the service are available.

Second, the differential effects of the scorecards on underperforming and overperforming offices suggest that it is especially important to improve information about underperforming bureaucrats. They imply that the type of recognition systems that are common for bureaucrats—where outstanding performances are recognized without addressing inadequate performances—are ineffective. Positive feedback might still have an overall positive effect, because it may motivate underperforming bureaucrats who want to receive better feedback. But due to the reputation-level effect, the positive feedback is less effective than the negative feedback and can, in some cases, even be counter-productive.

Finally, the model points out a general problem when using performance feedback for socially desirable behavior. If the reputations or self-perceptions of some agents are improved, this may have negative spillovers on all other behavior where reputation is a motivating factor. When evaluating such interventions, it is therefore important to measure effects on all domains of performance, not just performances where there could be direct spillovers due to multitasking. This is an especially important insight for government bureaucracies, where compressed wage structures, secure employment, and potential counterproductive incentives due to

corruption often make reputational concerns and a sense of pride in one's work more important motivators.

References

- Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics* 95(9-10), 1082–1095.
- Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association* 103(484), 1481–1495.
- Ashraf, A. (2019). Do performance ranks increase productivity? Evidence from a field experiment. Discussion Paper 196, Ludwig-Maximilians-Universität München und Humboldt-Universität zu Berlin.
- Ashraf, N., O. Bandiera, and B. K. Jack (2014). No margin, no mission? A field experiment on incentives for public service delivery. *Journal of Public Economics* 120, 1–17.
- Ashraf, N., O. Bandiera, and S. S. Lee (2014). Awards unbundled: Evidence from a natural field experiment. *Journal of Economic Behavior & Organization* 100, 44–63.
- Ayres, I., S. Raseman, and A. Shih (2013). Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *The Journal of Law, Economics, and Organization* 29(5), 992–1022.
- Bai, J., S. Jayachandran, E. J. Malesky, and B. A. Olken (2019). Firm growth and corruption: empirical evidence from vietnam. *The Economic Journal* 129(618), 651–677.
- Bandiera, O., M. C. Best, A. Q. Khan, and A. Prat (2021). The allocation of authority in organizations: A field experiment with bureaucrats. *The Quarterly Journal of Economics* 136(4), 2195–2242.
- Banerjee, A., R. Chattopadhyay, E. Duflo, D. Keniston, and N. Singh (2021). Improving police performance in Rajasthan, India: Experimental evidence on incentives, managerial autonomy, and training. *American Economic Journal: Economic Policy* 13(1), 36–66.
- Banerjee, A., R. Hanna, and S. Mullainathan (2012). Corruption. In R. S. Gibbons

- and J. Roberts (Eds.), *The Handbook of Organizational Economics*, pp. 1109–1147. Princeton University Press.
- Banerjee, A. V. (1997). A theory of misgovernance. *The Quarterly Journal of Economics* 112(4), 1289–1332.
- Banerjee, A. V., E. Duflo, and R. Glennerster (2008). Putting a band-aid on a corpse: incentives for nurses in the Indian public health care system. *Journal of the European Economic Association* 6(2-3), 487–500.
- Banuri, S. and P. Keefer (2016). Pro-social motivation, effort and the call to public service. *European Economic Review* 83, 139–164.
- Barrera-Osorio, F., K. Gonzalez, F. Lagos, and D. J. Deming (2020). Providing performance information in education: An experimental evaluation in Colombia. *Journal of Public Economics* 186, 104185.
- Becker, G. S. and G. J. Stigler (1974). Law enforcement, malfeasance, and compensation of enforcers. *The Journal of Legal Studies* 3(1), 1–18.
- Bertrand, M., R. Burgess, A. Chawla, and G. Xu (2020). The glittering prizes: Career incentives and bureaucrat performance. *The Review of Economic Studies* 87(2), 626–655.
- Bertrand, M., S. Djankov, R. Hanna, and S. Mullainathan (2007). Obtaining a driver's license in India: an experimental approach to studying corruption. *The Quarterly Journal of Economics* 122(4), 1639–1676.
- Blader, S., C. Gartenberg, and A. Prat (2020). The contingent effect of management practices. *The Review of Economic Studies* 87(2), 721–749.
- Byrne, D. P., A. L. Nauze, and L. A. Martin (2018). Tell me something i don't already know: Informedness and the impact of information programs. *Review of Economics and Statistics* 100(3), 510–527.
- Callen, M., S. Gulzar, A. Hasanain, M. Y. Khan, and A. Rezaee (2020). Data and policy decisions: Experimental evidence from Pakistan. *Journal of Development Economics* 146, 102523.
- Clot, S., G. Grolleau, and L. Ibanez (2018). Moral self-licencing and social dilemmas: an experimental analysis from a taking game in madagascar. *Applied Eco-*

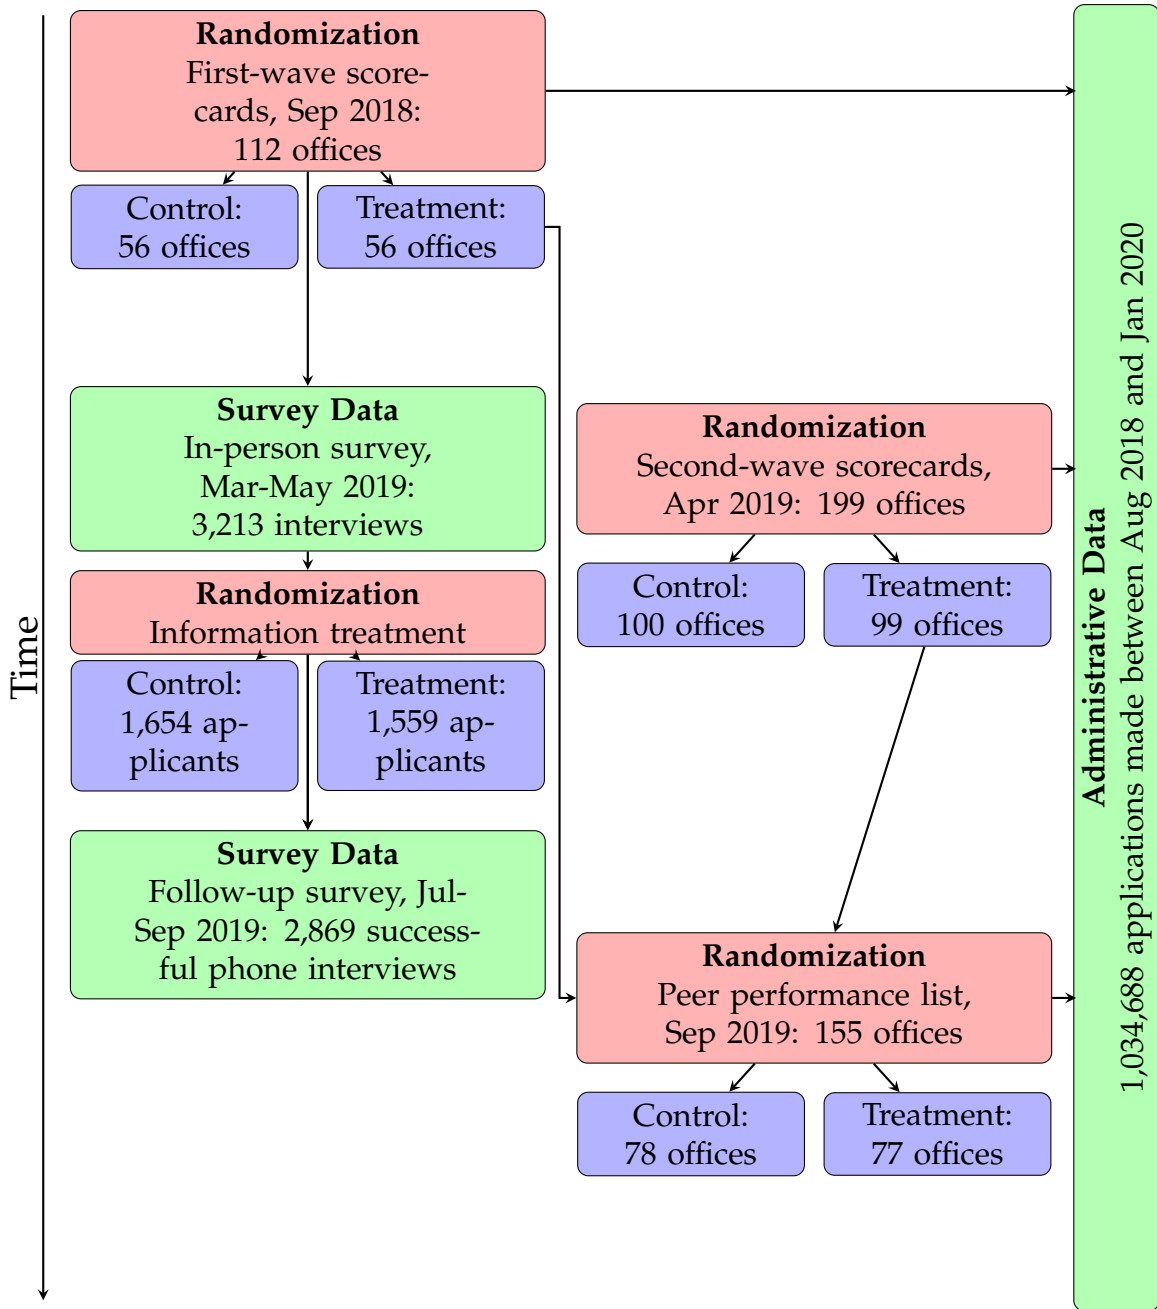
- nomics* 50(27), 2980–2991.
- Dal Bó, E., F. Finan, N. Y. Li, and L. Schechter (2021). Information technology and government decentralization: Experimental evidence from Paraguay. *Econometrica* 89(2), 677–701.
- de Janvry, A., G. He, E. Sadoulet, S. Wang, and Q. Zhang (2021). Performance evaluation, influence activities, and bureaucratic work behavior: Evidence from China. Working paper.
- Dhaliwal, I. and R. Hanna (2017). The devil is in the details: The successes and limitations of bureaucratic reform in India. *Journal of Development Economics* 124, 1–21.
- Do, Q.-A., T. Van Nguyen, and A. Tran (2021). Do people pay higher bribes for urgent services? Evidence from informal payments to doctors in Vietnam. Working paper.
- Dodge, E., Y. Neggers, R. Pande, and C. Moore (2021). Updating the state: Information acquisition costs and public benefit delivery. Working paper.
- Duflo, E., R. Hanna, and S. P. Ryan (2012). Incentives work: Getting teachers to come to school. *American Economic Review* 102(4), 1241–78.
- Finan, F., B. A. Olken, and R. Pande (2017). The personnel economics of the developing state. In *Handbook of Economic Field Experiments*, Volume 2, pp. 467–514. Elsevier.
- Freund, C., M. Hallward-Driemeier, and B. Rijkers (2016). Deals and delays: Firm-level evidence on corruption and policy implementation times. *World Bank Economic Review* 30(2), 354–382.
- Guriev, S. (2004). Red tape and corruption. *Journal of development economics* 73(2), 489–504.
- Hague Institute for Innovation of Law (2018). Justice needs and satisfaction in Bangladesh. Research report, Hague Institute for Innovation of Law.
- Hanna, R. and S.-Y. Wang (2017). Dishonesty and selection into public service: Evidence from India. *American Economic Journal: Economic Policy* 9(3), 262–90.
- Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*,

74–91.

- Holmström, B. and P. Milgrom (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.* 7, 24.
- Huntington, S. P. (1968). *Political Order in Changing Societies*. New Haven: Yale University Press.
- Kaufmann, D. and S.-J. Wei (1999). Does "grease money" speed up the wheels of commerce? Working Paper 7093, National Bureau of Economic Research.
- Khan, A. Q., A. I. Khwaja, and B. A. Olken (2016). Tax farming redux: Experimental evidence on performance pay for tax collectors. *The Quarterly Journal of Economics* 131(1), 219–271.
- Khan, A. Q., A. I. Khwaja, and B. A. Olken (2019). Making moves matter: Experimental evidence on incentivizing bureaucrats through performance-based postings. *American Economic Review* 109(1), 237–70.
- Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies* 76(3), 1071–1102.
- Leff, N. H. (1964). Economic development through bureaucratic corruption. *American Behavioral Scientist* 8(3), 8–14.
- Mas, A. and E. Moretti (2009). Peers at work. *American Economic Review* 99(1), 112–45.
- Muralidharan, K., P. Niehaus, S. Sukhtankar, and J. Weaver (2020). Improving last-mile service delivery using phone-based monitoring. *American Economic Journal: Applied Economics*. Forthcoming.
- Myrdal, G. (1968). *Asian drama, an inquiry into the poverty of nations*. London: The Penguin Press.
- Niehaus, P. and S. Sukhtankar (2013a). Corruption dynamics: The golden goose effect. *American Economic Journal: Economic Policy* 5(4), 230–69.
- Niehaus, P. and S. Sukhtankar (2013b). The marginal rate of corruption in public programs: Evidence from India. *Journal of Public Economics* 104, 52–64.
- Olken, B. A. (2007). Monitoring corruption: Evidence from a field experiment in Indonesia. *Journal of Political Economy* 115(2), 200–249.

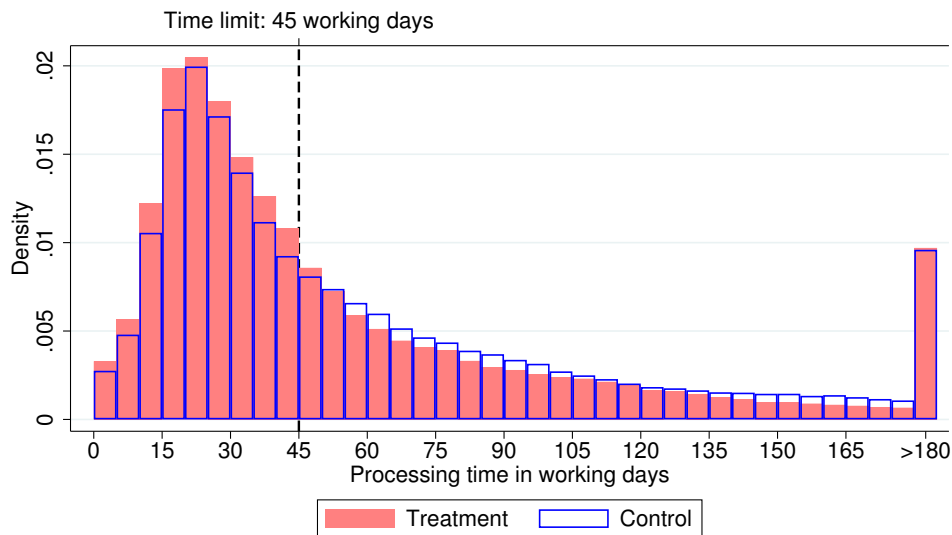
- Prendergast, C. (2007). The motivation and bias of bureaucrats. *American Economic Review* 97(1), 180–196.
- Raffler, P. (2020). Does political oversight of the bureaucracy increase accountability? Field experimental evidence from a dominant party regime. Working paper.
- Rasul, I. and D. Rogger (2018). Management of bureaucrats and public service delivery: Evidence from the nigerian civil service. *The Economic Journal* 128(608), 413–446.
- Reinikka, R. and J. Svensson (2005). Fighting corruption to improve schooling: Evidence from a newspaper campaign in uganda. *Journal of the European Economic Association* 3(2-3), 259–267.
- Rose-Ackerman, S. (1978). *Corruption: A study in political economy*. New York: Academic Press.
- Sachdeva, S., R. Iliev, and D. L. Medin (2009). Sinning saints and saintly sinners: The paradox of moral self-regulation. *Psychological science* 20(4), 523–528.
- Sanchez De la Sierra, R., K. Titeca, A. Amani Lameke, and A. Jolino Malukisa (2020). Corruption in hierarchies. Working paper.
- Svensson, J. (2003). Who must pay bribes and how much? Evidence from a cross section of firms. *The Quarterly Journal of Economics* 118(1), 207–230.
- Transparency International Bangladesh (2018). Corruption in service sectors: National household survey 2017. Research report.
- Young, A. (2019). Channeling fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results. *The Quarterly Journal of Economics* 134(2), 557–598.

Figure 1: Overview of Randomizations and Data Collection




The figure displays the experiment design and data collection. The timeline is chronological from top to bottom.

Figure 2: Histogram of Processing Times by Treatment



The figure displays histograms of processing times for the treatment and control groups separately. The data contains processed applications made between one month before the start of the experiment and 45 working days before the experiment ended.

Figure 3: Percentage of Applications Processed Within 45 Working Days



The figure displays the 10-working-day moving average of the share of applications processed within the time limit. The data is arranged by application start dates. The first and second vertical lines represent the date 45 working days before the first scorecard and the date of the first scorecard, respectively. The third vertical line represents the end of the data from the second randomization wave.

Table 1: Summary Statistics

	(1)	(2)	(3)	(4)
	Mean	Median	St. Dev.	Observations
Panel A: Application-level administrative data				
Process time \leq 45 working days	0.59	1	0.49	1,034,688
Actual process times (working days)	54	35	52	1,006,272
Process time, incl. imputed values (w. days)	61	36	69	1,034,688
Approval rate	0.67	1	0.47	1,006,265
Panel B: Monthly office-level admin. data				
Total applications	282	211	268	4,516
Applications processed	241	136	332	4,516
Apps. disposed within 45 working days	151	79	195	4,516
Apps. pending beyond 45 work. days	358	91	688	4,516
No ACL assigned	0.13	0	0.33	4,516
Female ACL	0.34	0	0.47	3,947
Panel C: Applicant survey data				
Applicant age	47	47	13	2,760
Female	0.06	0	0.23	2,869
Applicant monthly income (BDT)	23,552	20,000	19,811	2,653
Applicant HH per capita expenditure (BDT)	4,396	3,400	3,579	2,869
Land value (BDT 100,000)	19	8	31	2,671
Land size (acre)	0.24	0.09	0.40	2,748
Typical payment amount (BDT)	6,718	5,000	8,416	1,802
Typical payment $>$ 0	0.75	1	0.43	1,802
Reported payment amount (BDT)	1,477	0	3,480	2,869
Reported payment $>$ 0	0.28	0	0.45	2,869

Notes: The table reports summary statistics for applications in the administrative data, offices, and applicants in the survey data. Observations in Panels A and C are inversely weighted by the number of observations in that land office. Observations in Panel B are uniformly weighted. Continuous variables in the survey data are winsorized at the 99th percentile. USD/BDT \approx 84.3. *Reported payment* amount is any payment reported by the applicant above the official fee. *Typical payment* amount is the answer to the question of how much it is "normal for a person like yourself to pay."

Table 2: Scorecards' Effect on Application Processing Times

	(1) ≤45 working days	(2) ln(working days)	(3) Time index
Scorecard	0.060** (0.027)	-0.126** (0.059)	0.131** (0.059)
Start-month FE	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes
Observations	1,034,688	1,034,688	1,034,688
Clusters	311	311	311
Control mean	0.56	63.87	-0.00
Fraction imputed		0.02	

Notes: The table reports the effect of the scorecards on the speed of application processing. Column (1) shows the effect on applications processed within the time limit. Column (2) shows the effect on the log of processing time. Column (3) shows the effect on an index combining the two outcome variables. The data contains all applications made between one month before the start of the experiment and 45 working days before the experiment ended. Standard errors are clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table 3: Scorecards' Effect on Bribe Payments

	Amount		Any bribe		Amount if > 0	
	(1)	(2)	(3)	(4)	(5)	(6)
Scorecard	940 (616)	297 (182)	-0.018 (0.023)	-0.002 (0.022)	1,491* (763)	1,170** (453)
Start-month FE	Yes	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,802	2,869	1,802	2,869	1,324	779
Clusters	112	112	112	112	112	111
Control mean	6,127	1,284	0.75	0.28	8,119	4,660
Bribe measure	Typical	Reported	Typical	Reported	Typical	Reported

Notes: The table reports the effect of the scorecards on bribe payments made for application processing. Columns (1), (3), and (5) show the effect of the scorecards on the response to the question about the value of a typical payment by "a person like yourself". Columns (2), (4), and (6) show the effect on the response to the question about actual payments to government officials or agents. Columns (3) and (4) show the effect on the percentage of nonzero responses to the questions (extensive margin). Columns (5) and (6) show the effect among applicants who reported a nonzero bribe (intensive margin). All monetary amounts are in BDT. USD/BDT≈84.3. All continuous variables are winsorized at the 99th percentile. Standard errors are clustered at the office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table 4: Scorecards' Effects by Office Baseline Performance

	Time index	Bribe payment	
	(1)	(2)	(3)
Scorecard × Overperform	0.034 (0.080)	2,069*** (765)	630*** (233)
Scorecard × Underperform	0.238*** (0.088)	-100 (958)	42 (259)
Overperform baseline	0.375*** (0.108)	-1833* (977)	-815*** (294)
P-value subgroup diff.	0.09	0.09	0.10
Start-month FE	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes
Observations	1,034,688	1,802	2,869
Clusters	311	112	112
Overperformers: q-value	0.51	0.02	0.02
Underperformers: q-value	0.01	0.85	0.77
Overperformers: control mean	0.27	5,455	944
Underperformers: control mean	-0.30	6,726	1,578
Bribe measure		Typical	Reported

Notes: The table reports the effect of the scorecards separately for offices with above- and below-median baseline performance. Column (1) is based on administrative data. Columns (2) and (3) are based on survey data. Column (1) shows the effects on the time index constructed from the two variables: whether the application was processed on time and the log of the overall processing time. Column (2) shows the effects on the response to the question of how much it is "normal for a person like yourself to pay" beyond the official fee. Column (3) shows the effects on reported payments to government officials or agents. USD/BDT≈84.3. Standard errors are clustered at the land-office level. Q-values are sharpened false discovery rate q-values for the hypotheses that the effect is zero for both overperformers and underperformers (Anderson, 2008). Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Online Appendix

A Additional Details on the Experiment and Data

A.1 Randomization of Treatment Assignments

Half of the land offices were randomly assigned to receive the scorecards. The randomization was implemented by the author. The first-wave randomization was done separately for the group of land offices classified by the government as having full implementation of the e-governance system at the start of the experiment and for the group with partial implementation at the start of the experiment. The randomization strata were created using: the number of applications processed within 45 working days in June and July 2018; and applications pending for more than 45 working days at the end of July 2018. The second-wave randomization was done separately for the group of land offices having received above/below the median number of applications in February and March 2019. The randomization strata were created using: the number of applications processed within 45 working days and applications pending for more than 45 working days at the end of March 2019. The peer performance list was randomized in the same way as the second wave of scorecard randomizations using September 2019 data.

The information intervention was randomly assigned to half of the survey days in each land office. The randomization was stratified by the weekday pairs. During the treated days all applicants taking the in-person survey received the intervention. The enumerators were informed about whether they should deliver the information treatment or not using printed schedules, noncompliance with the treatment assignment is discussed in Appendix [B.3.1](#).

A.2 Scorecards Design and Performance Indicators

The two performance indicators were chosen after discussions with both the Ministry of Land and the Access to Information (a2i) agency. The performance score-

cards are based on applications made in the land office, regardless of which bureaucrat was assigned to the office when the application was made. Percentile rankings are calculated using the performance indicators for offices in both treatment and control group. If several offices had the same value of the indicator, they both received the highest percentile in the percentile range covered by the offices. A thumbs-up symbol is shown for rankings above the 60th percentile while a thumbs-down symbol is shown for rankings below the 40th percentile.

Both indicators are absolute and not relative to the total number of applications. I chose to base the performance indicators on absolute numbers for two reasons. First, a relative number could have created incentives to not let applicants apply, and then focus on processing the applications of those allowed to apply. Second, a measure relative to the number of applications received was perceived by the government as being unfair toward larger land offices processing more applications. Instead I choose one absolute measure benefiting larger land offices (the number of applications processed on time) and another benefitting smaller land offices (the number of applications pending beyond the time limit). Appendix Figure A3 shows that this worked in the intended way and that the average percentile ranking is not correlated with the number of applications received.

A.2.1 How Much of the Performance is Determined by the Bureaucrat?

An important question is to which extent bureaucrats can control the performance of their office and how much of it that is predetermined by factors outside the control of the bureaucrat. I address this question using the adjusted R-squared from fixed-effects regressions where the monthly average performance ranking is the outcome variable, for land offices that had more than one bureaucrat during the span of my data. Including fixed effects for the land office and month yields an adjusted R-squared of 0.44 suggesting that less than half of the variation in the performance score is determined by land office characteristics and seasonality. Including bureaucrat fixed effects increases the adjusted R-squared to 0.63.

This suggests that bureaucrats have a significant level of control over their performance score and that there is variation in bureaucrat performance that can explain a substantial part of the variation in the scores.

A.3 Data

A.3.1 Administrative Data From the E-governance System

The updated administrative data was transferred to the author each month from August 2018 until December 2020. The data includes all applications made in the e-governance system since its inception.

To calculate the number of working days between the start of the application and the date it was processed, I use data on holidays in Bangladesh.⁴¹ I then convert the number of calendar days to working days. For applications that were processed within one day, I use the the exact time of the application and processing to create a measure of what fraction of a day it took to process the application.

During the training of bureaucrats in using the e-governance system, example applications were made in two land offices that had not yet installed the e-governance system, making it appear as if the e-governance system was active in these two offices. In September 2018, I found out that these two offices had not yet installed the e-governance system, and I removed all applications from these offices. I also found out that some other applications in the e-governance system are the result of examples created in training. Using information provided about the dates of the training, I removed applications made before the first wave of randomization suspected to be the result of training.

The administrative data contains the size of the land that the land-record change was made for. Some of the values in this variable are unrealistically high, most likely due to mistakes in what unit was used when inputting the land value. I clean this variable by setting any observation with a value of above 10 acres to

⁴¹The data was retrieved from <https://www.timeanddate.com/holidays/bangladesh/>.

missing.⁴²

I also use administrative data to determine which bureaucrat was assigned to which land office. I take user data from the e-governance system to separate ACLs from other users and then assign a particular ACL to an office if that ACL is the ACL making the largest number of updates to land-record change applications that month. If an office has no updates made by any ACL in a month, I do not assign any ACL to that office in that month, unless an ACL was assigned to that office both prior to and after that month.

A.3.2 Survey Data

The survey was carried out in two stages: in the first stage, 3,213 individuals approaching the land office for the purpose of a land-record change were surveyed in person. The follow-up interview was conducted by phone. Enumerators attempted to reach applicants by phone three times in a day, and then made another three attempts the next day. If all attempts failed, the phone number was kept for another round of attempts one to two months later. Out of 3,213 in-person surveys 89% were successfully interviewed in the follow-up survey, yielding a final sample of 2,869 respondents.⁴³

I exclude observations from the same person being interviewed more than once, probably as a result of them visiting the land office multiple times, and multiple applicants providing the same phone number, probably because these individuals were in the same households. In some cases there were multiple interviews where the applicant started providing answers to interview questions but where the interview was redone because it was interrupted. In these cases I keep only the data from the final interview.

I only asked applicants about bribes in the phone interview. I asked about bribe

⁴²The largest land-record change in the survey data is 7.3 acres.

⁴³The average time between the two interviews was 3.3 months. Interviewees were given BDT 50 (USD 0.6) in the form of a mobile phone recharge for a completed in-person interview and BDT 100 (USD 1.2) for a completed phone interview.

payments in two different ways. The first question was phrased "How much do you think it is normal for a person like yourself to pay in order to get the mutation processed and receive the khatian? Include all extra payments or gifts to agents, government officials, and other individuals, but do not include the 1150 taka official fee." This is the question I refer to as the "typical payment". The next question was phrased as, "Did you pay any fees or give any gifts to anyone working for the Upazila Land Office or Union Parishad Land Office?" if the applicant responded yes, the enumerator asked about to whom the payment was made to and "How much taka did you pay or what was the monetary value of the gift that you paid to [recipient]?" The respondent was also asked the same question about agents helping with the application. The "reported payment" outcome variable is the sum of all such reported payments.

Processing times in the survey data for applicants for whom the application was not processed are imputed using the procedure described in Section 2.4.1. For applicants who did not answer the question about monthly income, I impute their income using the income predicted from a regression of income on per capita household expenditure.

A.3.3 Attrition and Non-Responses in the Survey

The attrition from the survey was 11%. Appendix Table A12 provide estimates for the effect of the scorecards and information treatments on the attrition rate. Column (1) shows that the scorecards are estimated to have had a positive effect on the attrition rate by 3%. Column (2) shows that the effect of the scorecards on attrition is somewhat higher in overperforming offices. Column (3) shows that the information intervention did not affect attrition.

An alternative definition of attrition is if a respondent did not answer a specific question in the survey. In columns (4)-(6) of Table A12 I show that when defining attrition in this way for the question about the typical bribe payment, the scorecard treatment did not affect attrition, neither overall nor for over- or underperformers.

Finally, columns (7)-(9) show the effect of the scorecards and information treatments on the non-response rate of respondents to the question about typical bribe payments, conditional on the respondent taking the follow-up survey.

A.3.4 Using Lee Bounds to Address Differential Attrition

If the scorecards caused some applicants to not take the follow-up survey and these applicants, on average, had different values for an outcome variable, this would bias the estimates of the effect of the scorecards on those outcomes. To assess the potential bias stemming from the differential attrition on the estimated effect of the scorecards on bribe payments, I construct lower Lee bounds for the estimated effect (Lee, 2009). Lower Lee bounds are the relevant robustness test, since the effects on bribe payments are positive (overall and for overperforming offices) or non-negative (for underperforming offices). If the estimated effect of the scorecards on attrition for a particular subgroup is positive, I create lower Lee bounds by adding back a random selection of the applicants from treated offices for whom there is no follow-up survey data. The number of observations added back is determined so that there is no longer a difference in the attrition rate between the treatment and control group. I then set the bribe payments for this sub-sample to be zero. If the effect of the scorecards on attrition is instead negative, I add back randomly selected respondents to the control group and assign these with the maximum bribe reported in the whole sample. The results are shown in Appendix Table A13. The lower Lee bounds are not qualitatively different from the main results.

B Additional Empirical Analyses

B.1 Unintended Consequences of the Scorecards

Below I discuss potential unintended consequences of the performance scorecards.

B.1.1 Gaming the Performance Indicators

Land offices may have reduced the number of applications, by refusing to serve some applicants or by processing some applications using the paper-based system. With a smaller number of applications, it may be easier to reach a higher performance. Column (1) in Panel A of Appendix Table A14 shows that the scorecards did not substantially affect applications received in the e-governance system. Column (1) in Panel B shows that the scorecards did not decrease applications more in the offices that were underperforming at baseline.

Another potential problem could be that bureaucrats allow applications only from individuals for whom they know it is easier to process the application within the time limit. The size of the land for which the land-record change is being made is positively associated with the processing time.⁴⁴ Therefore, if bureaucrats intended to avoid accepting complex applications, we would also expect to see a decrease in applications average land size. Column (2) in Panel A of Appendix Table A14 shows that the scorecards did not substantially affect the average land size among applications received. Column (2) in Panel B shows a negative point estimate for the effect of the scorecards on application land size in underperforming offices, but that effect is imprecisely measured and does not provide conclusive evidence for an effect in these offices.

B.1.2 Quality of Decision Making

Another concern is that the quality of the decisions by the bureaucrats was reduced by the scorecards. The main decision the bureaucrat makes with regards to the application is whether to accept or reject it. Column (1) in Panel A of Appendix Table A15 shows that the scorecards did not substantially change the percentage of applications rejected. Column (1) in Panel B shows that for overperforming offices, the point estimate for the effect is positive, while for underperforming offices the

⁴⁴In a regression of log processing time on log land size, controlling for application month and land-office fixed effects, the coefficient on land size is 0.012, statistically significant at the 1% level.

point estimate is negative. However, both coefficients are small.

It is still possible that the quality of the decisions was worse. If an application is wrongfully rejected, applicants typically reapply in the same office. Therefore, the percentage of applicants reapplying after having been previously rejected can be used as an indicator for the percentage of incorrect rejections. Column (2) in Panel A of Appendix Table A15 shows that the percentage of applicants stating that they were reapplying, after previously having been rejected, increased with the scorecards, but the estimate is imprecise and not statistically significant. Column (2) in Panel B shows that the increase is driven by an increase in overperforming offices, while there was no increase in underperforming offices.

Together, these results suggest that the scorecards did not lead to a decrease in the quality of decision making because bureaucrats were pressured to make faster decisions, since both rejections and incorrect rejections were unaffected by the scorecards in underperforming offices, the offices that improved their processing speed. Instead the results are consistent with applicants in overperforming offices not satisfying the new higher bribe demands in these offices and therefore got rejected. However, these results should be interpreted with caution as the low-rate of reapplying applicants causes the estimated effects to be imprecise relative to the control group mean.

B.1.3 Spillover Effects on Applications Not Measured by the Scorecards

Even after the e-governance system had been installed, not all applications were made in the e-governance system, as described in Section 2.1.4. In the survey data, 24% of applications were not made using the e-governance system. If the bureaucrats had reacted to the scorecards by diverting attention away from any task that was not measured by the scorecards, we would expect a negative spillover from the scorecards on the processing times for these applications. In Appendix Table A16, I estimate the effects of the scorecards on applications made outside of the e-governance system. Due to the decreased sample size, the estimates are impre-

cise, but overall there is no evidence for negative spillovers. While these results do not rule out all potential negative spillovers from the scorecards on other tasks performed by the bureaucrats, it suggests that the bureaucrats did not “game the system” by focusing only on applications that would improve their scorecards.

B.2 Alternative Explanations for the Effects of Scorecards

B.2.1 Increased Workload on Bureaucrats

One potential explanation for the scorecards causing higher bribe payments is that scorecards increase bureaucrats’ workload and the opportunity cost of their time. If bribe payments are made for bureaucrats time, this could increase bribe payments. Another alternative explanation is that faster processing times cause applicants to be willing to pay more. However, these explanations are inconsistent with the results that in the underperforming offices, where the changes in processing times are the largest, bribe payments do not change. Instead, bribe payments increase in the overperforming offices where changes in processing times are small.

B.2.2 Monitoring Efforts Shifted Away from Overperformers

If the scorecards led supervisors to decrease the monitoring of bureaucrats in the overperforming offices, that could have increased the bribes taken in these offices. However, if this was the main mechanism, we would have expected underperforming offices to have experienced an increase in monitoring and thus a decrease in bribes taken in those offices. Therefore, the results are only consistent with this mechanism if there was an asymmetric response by supervisors shifting monitoring away from overperforming offices but not towards underperforming offices, or if a decrease in monitoring leads to a larger increase in bribe payments than an increase in monitoring leads to a decrease in bribe payments. While I cannot completely rule this out, I find it less plausible than my preferred interpretation.

B.2.3 Supervisors Demanding Higher Bribes

I cannot observe sharing of bribe money within the bureaucracy, but it is plausible that this happens ([Sanchez De la Sierra et al., 2020](#); [Bandiera et al., 2021](#)). If the scorecards led to the supervisors of overperforming bureaucrats demanding larger amounts, this could drive up bribe payments. The most plausible reason for such a demand would be a mechanism in the model. In other words, the model would be correctly describing the mechanism, but with the wrong person described as the “bureaucrat.” This seems unlikely, since for the supervisors (UNOs), processing land-record applications is a very small share of their responsibilities, while for the bureaucrats (ACLs), it represents 25%-50% of their work.

B.2.4 Bureaucrats Using Scorecards in Negotiations With Applicants

Another alternative explanation is that positive scorecards help bureaucrats prove to applicants that they can process applications quickly. This could then allow the bureaucrats receiving positive scorecards to charge higher bribes. This explanation is implausible for four reasons. First, the coefficients on "Overperform baseline" in columns (4) and (5) of Appendix Table [A11](#), show that the expected processing times are 13% lower in the overperforming offices not receiving the scorecards, suggesting that the applicants are already aware of the faster processing times in these offices. Second, column (4) shows that the point estimate for the effect of the scorecards on applicants' expectations of processing times in overperforming offices is a 5% increase, if the scorecards helped bureaucrats improve applicants' expectations of processing times, the effect on the expectations should be negative. Third, in none of the qualitative interviews done with bureaucrats and applicants was it ever mentioned that the scorecards were shown to applicants. Fourth, the information intervention tried to accomplish what a bureaucrat could have achieved by showing the scorecards to applicants, but Appendix Table [A2](#) shows that the information intervention did not substantially increase bribes.

B.3 Effects of Information Intervention on Bribes

Appendix Table A2 shows that the information treatment did not affect bribes, neither by itself nor in combination with the scorecards. Columns (1) and (4) show that the information treatment by itself did not have a substantial impact on bribes. Columns (2) and (5) show that even together with the information treatment, the scorecards did not reduce bribes. Columns (3) and (6) show the results separately for overperforming and underperforming land offices. In particular the coefficients on $Info \times Scorecard \times Underperform$ show that even for applicants receiving the information intervention in offices that were underperforming at the start of the experiment, where both actual processing times and the expectations of these processing times declined the most, bribes did not decrease.

B.3.1 Noncompliance With the Information Treatment Assignment

Enumerators implemented the information intervention during the in-person interviews. They were given a different schedule for each land office where each weekday was indicated as a day when the treatment would be delivered or not. The enumerator would indicate in the survey if the information had been given or not. According to the enumerators' indications of whether the information was given, in 17% of interviews incorrect information was given. This was due to enumerators misunderstanding the schedule.⁴⁵

The main result related to the information intervention is that it did not decrease bribe payments. In particular, that it did not decrease bribe payments in offices receiving the scorecards. Since this is a null result, I defined the treatment variable such that the information intervention has the highest ex-ante probability to cause a decrease in bribe payments. In the paper's main analysis, I use the median treatment delivered in a particular land office and survey day as the treatment variable. Alternatively, I could use the treatment indicated by the enumerator for

⁴⁵Among the interviews where the treatment assignment was not followed, 91% were on days when the median interview was not given the information as per the treatment assignment.

each individual applicant, but due to likely information spillovers within an office survey day, I prefer the median treatment of the office-survey day.


Appendix Table A17 shows the robustness of the estimated effects of the information treatment to using alternative definitions of the treatment variable. Column (1) uses my preferred treatment variable based on the median treatment delivered in the land-office survey day. Column (2) uses the individual treatment delivered for the applicant. Column (3) uses the treatment assigned, and column (4) uses the treatment assigned but drops observations where the treatment was not correctly implemented. The null result of the effect on bribe payments is robust across all definitions.

B.4 Estimating processing days saved per year

To estimate processing days saved per year, I multiply the estimates for the total number of processing days per year and the improvement in the processing time for the average application. I take the average processing time in the control offices and multiply it by an estimate for the annual number of applications in the treatment offices.⁴⁶ I then multiply this number with the percentage decrease for the average application, estimated by a uniformly weighted regression according to the specification in Equation 1.

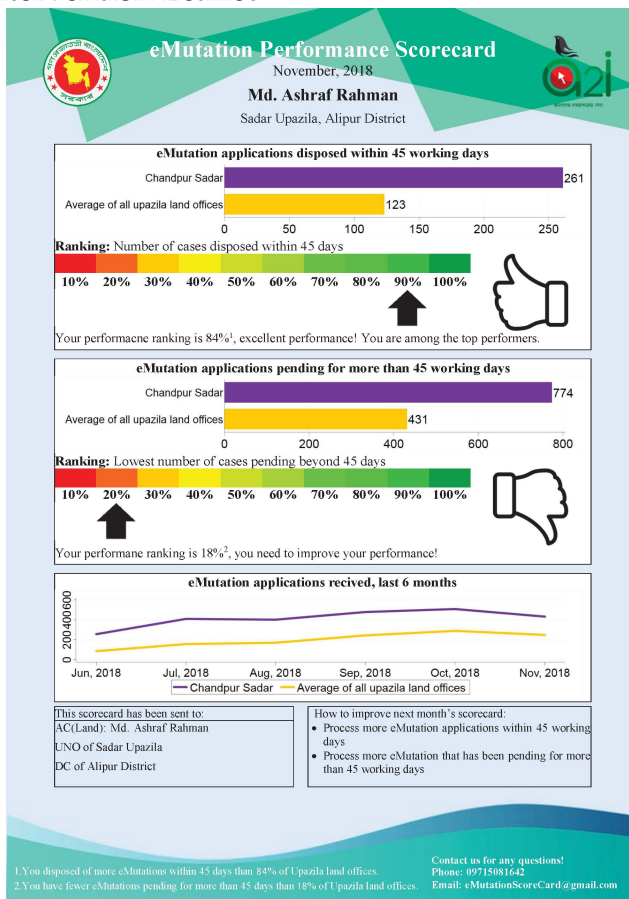
⁴⁶The estimate is two times the number of applications in the treatment offices in the last six months of 2019. During this time period the usage of the e-governance system had stabilized, but it was before the reduction in applications due to the COVID-19 pandemic.

Figure A1: Value of Record of Rights, Faster Processing, and Bribes



The figure displays the applicants' stated valuations of the land-record change and compare these to the bribe payments. The first bar shows the average of applicants' stated value of getting the land record change approved and receiving a record of rights. The second bar shows the average stated value of getting the application processed within seven days from the time of the first survey. The third bar shows the average value of bribe payments reported by the applicant; 73% of the applicants reported having paid no bribes. The fourth bar shows the average value of reported bribe payments among applicants reporting a nonzero bribe. The fifth bar shows the average response to the question about a typical bribe payment by "a person like yourself", 27% of the applicants responding to this question reported that a typical applicant paid no bribes. The first bar is measured on the left axis, the next four bars are measured on the right axis. All variables are winsorized at the 99th percentile. Observations are weighted by the inverse of the number of observations in that land office. USD/BDT \approx 84.3.

Figure A2: Example Performance Scorecard, Peer Performance List and Information Intervention Leaflet



List of eMutation Performance
November, 2018

The performance scorecard is sent monthly and evaluates a land office in terms of the number of eMutation applications disposed within 45 working days and the number of applications pending for more than 45 working days. A ranking from 0% to 100% is then calculated for the number of applications disposed within 45 working days (disposed ranking) and the number of applications pending for more than 45 working days (pending ranking). A higher ranking percentage reflects a better performance.

We are now piloting sending out a list of all Upazila's performance every month. Your Upazila has been selected to be part of this pilot project. The list below has been sent to 774 Upazilas.

Land offices can improve their performance scorecards by:

- Processing more eMutation applications within 45 working days
- Processing more eMutation applications that has been pending for more than 45 working days

Upazila Land Office	District	Disposed ranking	Pending Ranking
(Name of upazila)	(Name of district)	48%	20%
(Name of upazila)	(Name of district)	49%	30%
(Name of upazila)	(Name of district)	51%	21%
(Name of upazila)	(Name of district)	17%	16%
(Name of upazila)	(Name of district)	18%	69%
(Name of upazila)	(Name of district)	52%	29%
(Name of upazila)	(Name of district)	45%	74%
(Name of upazila)	(Name of district)	98%	27%
(Name of upazila)	(Name of district)	29%	70%
(Name of upazila)	(Name of district)	07%	44%
(Name of upazila)	(Name of district)	2%	29%
(Name of upazila)	(Name of district)	26%	57%
(Name of upazila)	(Name of district)	49%	74%
(Name of upazila)	(Name of district)	56%	87%
(Name of upazila)	(Name of district)	20%	74%
(Name of upazila)	(Name of district)	95%	71%
(Name of upazila)	(Name of district)	29%	33%
(Name of upazila)	(Name of district)	22%	49%
(Name of upazila)	(Name of district)	95%	48%
(Name of upazila)	(Name of district)	44%	35%
(Name of upazila)	(Name of district)	52%	41%
(Name of upazila)	(Name of district)	85%	2%
(Name of upazila)	(Name of district)	85%	69%
(Name of upazila)	(Name of district)	81%	4%
(Name of upazila)	(Name of district)	27%	4%
(Name of upazila)	(Name of district)	99%	4%
(Name of upazila)	(Name of district)	9%	3%
(Name of upazila)	(Name of district)	95%	2%
(Name of upazila)	(Name of district)	81%	2%
(Name of upazila)	(Name of district)	93%	2%
(Name of upazila)	(Name of district)	64%	59%

Information for applicants
Land Record eMutation

Over the past 6 months the Government of Bangladesh has taken several steps to reduce the time it takes to process a Land Record eMutation. Before a typical Land Record eMutation took 69 working days (more than 3 months) now a typical Land Record eMutation takes 49 working days (a little more than 2 months).

You can apply for a Land Record eMutation by visiting the Upazila Land Office or from any computer connected to the internet (<http://training.land.gov.bd/mutation/application>).

Steps of Land Record eMutation and timeline:


1. Make application online or in Upazila Land Office
2. Upazila Land Office will check the application and send it to Union Land Office
3. Union Land Office Assistant will make visit to land and write report to Upazila Land Office
4. Upazila Land Office will read report and call you for hearing via text message
5. You will attend hearing (according to text message)
6. Pay fee of 1150 taka and receive your Khatian

This information sheet was prepared by Innovations for Poverty Action in collaboration with a2i and the Land Reforms Board of Bangladesh.
Contact phone number: [Redacted]


On top is an Example of a performance scorecard. The names are changed to preserve anonymity. To the bottom left is an example of the first page of a peer performance list. To the bottom right is an English translation of the information leaflet used in the information intervention.

Figure A3: Land Office Size and Performance

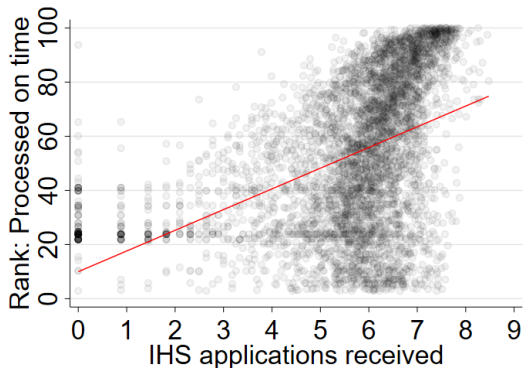
A3a: Offices' Monthly Average of Applications Received



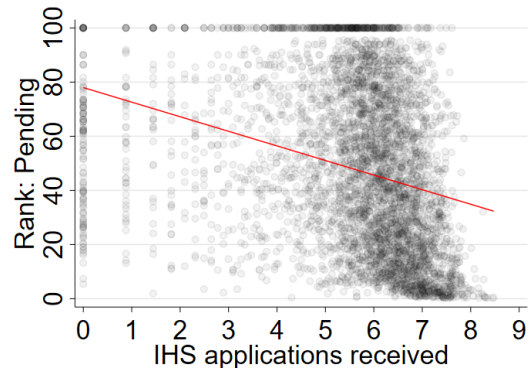
A3b: Average Performance Ranking and Applications Received



A3c: Ranking of Applications Processed Within 45 Working Days and Applications Received



A3d: Ranking of Applications Pending Beyond 45 Working Days and Applications Received



Notes: The figure displays information applications received by a land office in a month and that variables relationship with the performance ranking in that month. Data is from when the office first started using the e-governance system until the end of the experiment. Sub-Figure a) shows a histogram of the offices monthly average of applications received. Sub-Figure b) shows the relationship between the average performance ranking in a month and the inverse hyperbolic sine (IHS) of applications received in that month. Sub-Figure c) shows the percentile ranking of the number of application processed on time and the IHS of applications received. Sub-Figure d) shows the percentile ranking of applications pending beyond the 45-working-day time limit and applications received.

Table A1: Balance of Randomization

	Scorecard		Control		Difference
	Obs. (Cluster)	Mean (SD)	Obs. (Cluster)	Mean (SD)	Diff. (SE)
Panel A. Administrative Data					
≤45 working days	56,703 (146)	0.50 (0.50)	56,564 (146)	0.53 (0.50)	-0.028 (0.04)
ln(Process time)	56,703 (146)	3.77 (1.40)	56,564 (146)	3.69 (1.44)	0.083 (0.13)
Time index	56,703 (146)	-0.06 (0.98)	56,564 (146)	0.00 (1.00)	-0.061 (0.09)
Process time (w. days)	56,703 (146)	77.60 (73.53)	56,564 (146)	76.90 (79.15)	0.700 (6.98)
Approved	36,409 (136)	0.70 (0.46)	38,776 (141)	0.73 (0.44)	-0.030 (0.03)
Panel B. Survey Data					
Applicant age	1,383 (56)	47.34 (13.35)	1,377 (56)	47.39 (13.82)	-0.05 (0.64)
Female	1,460 (56)	0.06 (0.23)	1,409 (56)	0.06 (0.24)	-0.00 (0.01)
Monthly income (BDT 1K)	1,460 (56)	24.49 (17.62)	1,409 (56)	22.85 (17.88)	1.64 (1.22)
App. status: Applying	1,382 (56)	0.23 (0.42)	1,377 (56)	0.22 (0.42)	0.01 (0.03)
Ongoing	1,382 (56)	0.64 (0.49)	1,377 (56)	0.63 (0.49)	0.01 (0.04)
Rejected	1,382 (56)	0.00 (0.07)	1,377 (56)	0.00 (0.06)	0.00 (0.00)
Approved	1,382 (56)	0.13 (0.34)	1,377 (56)	0.15 (0.36)	-0.02 (0.03)
Land value (BDT 100K)	1,335 (56)	20.70 (31.87)	1,336 (56)	17.32 (29.82)	3.38 (2.17)
Land size (acre)	1,379 (56)	0.24 (0.41)	1,369 (56)	0.24 (0.39)	-0.00 (0.02)

Notes: The table reports the balance of randomization. Panel A. uses administrative data from 45 working days before the first scorecard. Due to this restriction, only 292 of the 311 offices are included. For applications not processed by the date of the first scorecard, the processing time is imputed using the procedure in Section 2.4.1. Data on approvals are as per the date of the first scorecard. Panel B uses survey data from the in-person survey of applicants, which was conducted before the conclusion of the processing of the applicants' applications but after the start of the scorecards. Continuous variables are winsorized at the 99th percentile. In both panels observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A2: Effects of Information Treatment and Scorecards on Bribes

	Typical payment			Reported payment		
	(1)	(2)	(3)	(4)	(5)	(6)
Information treatment	212 (490)	-326 (767)	-311 (760)	-31 (150)	-56 (184)	-63 (188)
Scorecard		462 (719)			285 (218)	
Scorecard x Information		987 (1,105)			29 (278)	
Info. x Scorecard x Overperform			1,644* (926)			740*** (280)
No info. x Scorecard x Overperform			2,427** (956)			533* (280)
Info. x Scorecard x Underperform			1,343 (1,411)			-26 (334)
No info. x Scorecard x Underperform			-1,519* (896)			122 (310)
Overperform baseline			-1,712* (983)			-824*** (294)
Start-month FE	Yes	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,802	1,802	1,802	2,869	2,869	2,869
Clusters	539	112	112	570	112	112
Control mean	6,553	6,258		1,508	1,328	

Notes: This table reports the effect of the scorecards and information treatments on bribe payments. Columns (1)-(3) show the effect on the response to the question of how much it is "normal for a person like yourself to pay" beyond the official fee. Columns (4)-(6) show the effect on reported payments to government officials or agents. All outcome variables are winsorized at the 99th percentile. The outcome variables are in BDT. USD/BDT≈84.3. In columns (1) and (4), standard errors are clustered at the land-office-by-day level. In columns (2)-(3) and (5)-(6), standard errors are clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A3: Effect of Peer Performance List

	≤ 45 working days		ln(working days)	
	(1)	(2)	(3)	(4)
Peer Performance List	0.002 (0.045)	-0.007 (0.048)	-0.030 (0.094)	0.037 (0.105)
Post \times Peer Performance List		0.008 (0.045)		-0.060 (0.095)
Scorecard		0.048 (0.038)		-0.124 (0.090)
Post \times Scorecard		0.020 (0.038)		-0.002 (0.083)
Start-month FE	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes
Observations	280,002	1,034,688	280,002	1,034,688
Clusters	155	311	155	311
Control mean	0.68	0.59	49.47	63.87

Notes: This table reports the effect of adding the peer performance list to the scorecard and sharing the performance information about one office with the ACLs, UNOs, and DCs of 76 other offices. Columns (1)-(2) show the effect on the percentage of applications processed within the 45-working-day time limit. Columns (3)-(4) show the effect on the log of processing time. Columns (1) and (3) only use data from offices receiving the scorecards and applications made one month before the first performance list was sent out until the end of the experiment (August 15, 2019 to January 20, 2020). Columns (2) and (4) use data on all applications made between one month before the start of the scorecard experiment started and 45 working days before the experiment ended (August 13, 2018 to January 20, 2020). The indicator variable *Peer Performance List* takes the value of one for applications made in offices receiving the peer performance lists. The indicator variable *Post* takes the value of one for applications made later than one calendar month before the first performance lists were sent out. Standard errors are clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table A4: Scorecards' Effects on Visits, Processing Times and Satisfaction in Survey

	(1)	(2)	(3)	(4)
Panel A: Overall effect	Visits	≤ 45 w. days	ln(w. days)	Satisfaction (SD)
Scorecard	-1.266*** (0.422)	0.040 (0.030)	-0.059 (0.049)	-0.038 (0.060)
Panel B: Heterogeneous effects				
Scorecard \times Overperform baseline	-1.054* (0.606)	0.012 (0.049)	-0.005 (0.078)	-0.080 (0.081)
Scorecard \times Underperform baseline	-1.272** (0.636)	0.055 (0.041)	-0.088 (0.066)	-0.016 (0.086)
Overperform baseline	-1.676** (0.784)	0.113** (0.051)	-0.216** (0.084)	0.181* (0.097)
Start-month FE	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes
Observations	2,869	2,869	2,869	2,869
Clusters	112	112	112	112
Control mean	10.73	0.56	59.81	

Notes: This table reports the effect of the scorecards on visits to land offices, processing times, and stated satisfaction measured in the survey data. Standard errors are clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A5: Heterogeneous Effects on Processing Times, Controlling for Baseline Values Interacted With Treatment

	Time index						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Scorecard × Underperform baseline	0.204* (0.121)	0.200* (0.120)	0.180 (0.121)	0.197 (0.121)	0.219* (0.119)	0.205* (0.121)	0.184 (0.119)
Scorecard × Rank: Applications within 45 working days before baseline		0.001 (0.002)					-0.004 (0.003)
Scorecard × Rank: Applications prior to 45 working days before baseline			-0.002 (0.002)				-0.003 (0.003)
Scorecard × Female ACL baseline				-0.039 (0.141)			-0.125 (0.140)
Scorecard × No ACL baseline					-0.297* (0.154)		-0.388** (0.158)
Scorecard × Land size (st. dev.)						-0.008 (0.013)	0.000 (0.013)
Baseline control	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline control		Yes	Yes	Yes	Yes	Yes	Yes
Start-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,034,688	1,034,688	1,034,688	1,034,688	1,034,688	1,007,014	1,007,014
Clusters	311	311	311	311	311	311	311

Notes: This table reports the heterogeneous effects on processing time by baseline performance, controlling for other baseline variables interacted with treatment. The first row of estimates shows the difference in the effect between under- and overperforming offices. Column (1) uses the main specification from Table 4. Columns (2) and (3) controls for the treatment interacted with two measures of office size: the offices' percentile ranking in terms of applications received in the 45-working-day time period before the treatment and offices' percentile ranking of the total number of applications received before that 45-working-day period. Columns (4)-(6) control for the treatment interacted with the ACL being female at baseline, no ACL being appointed at baseline, and the size of the land the application was for. Column (7) includes all the controls. Standard errors clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A6: Heterogeneous Effects on Bribe Payments, Controlling for Baseline Values Interacted With Treatment

	Typical bribe payment						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Scorecard × Overperform baseline	2,169* (1,252)	2,180* (1,251)	2,594** (1,306)	2,044 (1,289)	2,088 (1,280)	2,181* (1,214)	2,691* (1,390)
Scorecard × Rank: Applications within 45 working days before baseline		-2 (22)					-35 (29)
Scorecard × Rank: Applications prior to 45 working days before baseline			-20 (23)				-48 (31)
Scorecard × Female baseline ACL				775 (1,508)			-38 (1,589)
Scorecard × No baseline ACL					442 (1,708)		-125 (1,758)
Scorecard × Land size (st. dev.)						113 (379)	107 (382)
Treatment control	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline control		Yes	Yes	Yes	Yes	Yes	Yes
Start-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,802	1,802	1,802	1,802	1,802	1,726	1,726
Clusters	112	112	112	112	112	112	112

Notes: This table reports the heterogeneous effects on bribe payments by land-office baseline performance, controlling for other variables interacted with the treatment. The first row of estimates shows the difference in the effect between under- and over-performing offices. Column (1) uses the main specification from Table 4. Columns (2) and (3) control for the treatment interacted with two measures of office size: the offices' percentile ranking in terms of the number of applications received in the 45-working-day time period before the treatment and offices' percentile ranking in terms of the total number of applications received before that 45-working-day period. Columns (4)-(6) control for the treatment interacted with the ACL being a female at baseline, no ACL being appointed at baseline, and the size of the land the application was for. Column (7) includes all the controls from columns (2)-(6). Standard errors are clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A7: Effect on Processing Times: Alternative Specifications

	Time index					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Overall effect						
Scorecard	0.089 (0.089)	0.125* (0.070)	0.140** (0.061)	0.097 (0.076)	0.137** (0.057)	
Scorecard × Post						0.156* (0.080)
Panel B: Heterogeneous effects						
Scorecard × Overperform	-0.009 (0.116)	0.001 (0.089)	0.005 (0.081)	0.031 (0.101)	0.032 (0.077)	
Scorecard × Underperform	0.197* (0.119)	0.264*** (0.094)	0.280*** (0.089)	0.169 (0.110)	0.250*** (0.084)	
Overperform baseline	0.537*** (0.110)	0.572*** (0.088)	0.341*** (0.103)	0.487*** (0.130)	0.261** (0.116)	
Scorecard × Post × Overperform						0.040 (0.074)
Scorecard × Post × Underperform						0.283** (0.114)
P-value subgroup diff.	0.21	0.04	0.02	0.37	0.06	0.13
Start-month FE	No	No	No	Yes	Yes	Yes
Stratum FE	No	No	No	Yes	Yes	No
Weighted by office	No	Yes	Yes	No	Yes	Yes
Baseline controls	No	No	Yes	No	Yes	No
Office FE	No	No	No	No	No	Yes
Observations	1,034,688	1,034,688	1,034,688	1,034,688	1,034,688	1,188,351
Clusters	311	311	311	311	311	306

Notes: This table shows the robustness of the estimated effect of the scorecards on the time index to different regression specifications. The time index is constructed from the two variables: whether the application was processed on time and the log of the overall processing time. Panel A shows the estimates of the overall effect and Panel B shows the estimates of the heterogeneous effects, as in Tables 2 and 4, respectively. The specifications differ from the ones used in those tables in the following ways: column (1) shows the estimate from uniformly weighted regressions with no controls. Column (2) uses no controls. Column (3) uses no strata or month fixed effects, but it controls for baseline month values for the number of applications processed within 45 working days, the number of applications pending beyond 45 working days, the number of applications received, and the percentage of applications received in the month two months before the baseline that were processed within the time limit. Column (4) shows the estimate from a uniformly weighted regression. Column (5) controls for the baseline month controls. Column (6) shows the estimate from a regression, including applications made 45 working days before the start of the experiment (for the 306 offices that have such data) and using land-office fixed effects. Standard errors are clustered at the land-office level. ***p<0.01; **p<0.05; *p<0.1.

Table A8: Scorecards' Effect on Bribes: Alternative Specifications

	Typical payments						Reported payments						To gov. off.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Panel A: Overall effect	1,194*	1,164*	875	688	1,890**	482	388*	386**	307*	238	610**	194	245*
Scorecard	(717)	(698)	(623)	(522)	(853)	(449)	(196)	(195)	(172)	(179)	(264)	(127)	(141)
Panel B: Heterogeneous effects													
Scorecard x Overperform	1,808*	1,602*	2,267***	1,325**	2,811***	1,576***	647***	648***	659***	519**	1,011**	468***	462***
	(920)	(864)	(827)	(612)	(1,067)	(585)	(234)	(219)	(233)	(207)	(418)	(168)	(173)
Scorecard x Underperform	666	811	-393	-215	1,085	-534	175	175	28	75	255	-21	82
	(1,068)	(1,069)	(923)	(813)	(1,320)	(646)	(303)	(300)	(241)	(235)	(348)	(172)	(208)
Overperform baseline	-1,219	-1,272	-2,138**	-1,916**	-2,055*	-1,667**	-455*	-634**	-759***	-1,005***	-634	-636***	-565**
	(837)	(919)	(982)	(843)	(1,149)	(754)	(244)	(255)	(280)	(250)	(404)	(202)	(237)
P-value subgroup diff.	0.42	0.57	0.04	0.06	0.31	0.02	0.22	0.21	0.06	0.10	0.19	0.05	0.16
Start-month FE	No	No	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Yes	Yes	Yes
Stratum FE	No	No	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Yes	Yes	Yes
Weighted by office	No	Yes	No	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes
PDS controls	No	No	No	Yes	No	No	No	No	No	Yes	No	No	No
Percentile Winsorized	99	99	99	99	99	95	99	99	99	99	99	95	99
Observations	1,802	1,802	1,802	1,802	1,802	1,802	2,869	2,869	2,869	2,869	2,869	2,869	2,869
Clusters	112	112	112	112	112	112	112	112	112	112	112	112	112

Notes: This table shows the robustness of the estimated effect on bribes. Panel A shows the estimates of the overall effect and Panel B of the heterogeneous effects. Columns (1)-(6) show the effect on the response to the question of how much it is "normal for a person like yourself to pay" in bribes. Columns (7)-(12) show the effect on reported bribe payments to government officials or agents. The specifications differ from the specifications in Tables 3 and 4 in the following ways: columns (1) and (7) use uniformly weighted regressions with no fixed effects. Columns (2) and (8) use regressions with no fixed effects. Columns (3) and (9) use uniformly weighted regressions. Columns (4) and (10) use post-double-selection (PDS) to select controls among 154 variables generated from the baseline administrative and survey variables unlikely to have been affected by the experiment. Columns (5) and (11) use unwinorized outcome variables. Columns (6) and (12) use outcome variables winsorized at the 95th percentile. Column (13) uses payments made directly to government officials as the outcome variable, excluding payments made to agents. Standard errors are clustered at the land-office level. ***p<0.01; **p<0.05; *p<0.1.

Table A9: Scorecards' Effects on Office-by-Month Level Outcomes

Panel A: Overall effect	(1)	(2)	(3)	(4)	(5)
	IHS dis. ≤45	IHS pen. >45	Rank dis.	Rank pen.	Index
Scorecard	0.211*	-0.104	2.096	2.020	0.090
	(0.121)	(0.143)	(1.684)	(1.771)	(0.061)
Panel B: Heterogeneous effects					
Scorecard × Overperform	-0.022	0.171	-0.706	-0.845	-0.065
	(0.147)	(0.219)	(2.187)	(2.619)	(0.082)
Scorecard × Underperform	0.466**	-0.412**	5.171**	5.226**	0.263***
	(0.195)	(0.186)	(2.614)	(2.414)	(0.091)
Overperform baseline	0.393*	-0.422	6.368*	3.853	0.210*
	(0.226)	(0.280)	(3.321)	(3.539)	(0.125)
P-value subgroup diff.	0.05	0.05	0.09	0.10	0.01
Month FE	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes
Baseline controls	Yes	Yes	Yes	Yes	Yes
Observations	4,516	4,516	4,516	4,516	4,516
Clusters	311	311	311	311	311

Notes: This table reports the effect of the scorecards on office-by-month level outcomes. Panel A shows the estimates of the overall effect. Panel B shows the estimates of the heterogeneous effects. Column (1) shows the effect on the inverse hyperbolic sine (IHS) transformation of the number of applications processed within 45 working days. Column (2) shows the effect on the IHS of the number of applications pending beyond 45 working days. Column (3) shows the effect on the percentile ranking in terms of the number of applications processed within 45 working days. Column (4) shows the effect on the percentile ranking in terms of the number of applications processed within 45 working days; a higher number of pending applications leads to a lower ranking. Column (5) shows the result on an Inverse Covariance Weighted index created with the outcome variables of columns (1)-(4); a higher index value indicates a better performance. Standard errors are clustered at the land-office level. ***p<0.01; **p<0.05; *p<0.1.

Table A10: Alternative Measures of Baseline Performance

	Time index			Typical payments			Reported payments		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Scorecard × Overperform 3m baseline	0.017 (0.080)			1,882*** (710)			594** (229)		
Scorecard × Underperform 3m baseline	0.248*** (0.087)			-9 (946)			38 (257)		
Treat x 75-100th percentile		-0.027 (0.100)			1,695* (962)			604** (264)	
Treat x 50-75th percentile		0.044 (0.117)			2,434** (1,163)			668* (382)	
Treat x 25-50th percentile		0.135 (0.119)			116 (1,047)			241 (311)	
Treat x 0-25th percentile		0.337*** (0.121)			-142 (1,496)			-116 (390)	
Treat x Baseline ranking			-0.007*** (0.003)			38.187 (29)			15* (8)
Scorecard			0.115** (0.057)			958.064 (608)			329* (173)
P-value subgroup diff.	0.05			0.11			0.11		
Start-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline performance control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1034688	1034688	1034688	1802	1802	1802	2869	2869	2869
Clusters	311	311	311	112	112	112	112	112	112

Notes: This table shows the robustness of the results for the heterogeneity of the effect of the scorecards for different baseline performance measures. Columns (1)-(3) show the effects on the index of the two main processing-time outcome variables used in column (3) of Table 2; a higher index value indicates a faster processing time. Columns (4)-(6) show the effect on the estimate for how much is it "normal for a person like yourself to pay" in bribes. Columns (7)-(9) show the effects on the bribe payments reported by the applicant. Columns (1), (4), and (7) show the heterogeneity in the effect of the scorecards based on the land office having an above- or below-median average ranking across the last three months of the baseline period. Columns (2), (5), and (8) show the heterogeneity based on the quartile of baseline ranking. Columns (3), (6), and (9) show the heterogeneity based on the continuous baseline ranking. USD/BDT ≈ 84.3. Standard errors are clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A11: Treatment Effects on Expected Processing Time

	ln(Expected processing time)				
	(1)	(2)	(3)	(4)	(5)
Scorecard	-0.006 (0.035)		-0.000 (0.041)		
Information treatment		-0.049** (0.022)	-0.044* (0.024)		
Scorecard x Information			-0.009 (0.042)		
Scorecard x Overperform				0.046 (0.050)	
Scorecard x Underperform				-0.041 (0.049)	
Info. x Scorecard x Overperform					0.018 (0.056)
No info. x Scorecard x Overperform					0.073 (0.057)
Info. x Scorecard x Underperform					-0.067 (0.051)
No info. x Scorecard x Underperform					-0.014 (0.057)
Overperform baseline				-0.129** (0.058)	-0.129** (0.058)
Start-month FE	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes
Observations	2,467	2,467	2,467	2,467	2,467
Clusters	112	561	112	112	112
Control mean	61	68	62		

Notes: This table reports the effect of the scorecards and information treatments on expected processing times at the time of the in-person interview. The outcomes variable is the log transformation of the sum of the expected future processing time and the processing time already incurred at the time of the in-person interview, winsorized at the 99th percentile. Standard errors are clustered at the land-office level, except for in column (2), where standard errors are clustered at the land-office-by-survey-day level, as that is the level of randomization for the information treatment. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A12: Treatment Effects on Survey Attrition

	Attrition: Survey			Attrition: Typical bribe			No Response: Typical bribe		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Scorecard	0.027*			-0.006			-0.025		
	(0.015)			(0.020)			(0.023)		
Scorecard × Overperform baseline		0.034*			0.005			-0.017	
		(0.017)			(0.029)			(0.033)	
Scorecard × Underperform baseline		0.022			-0.017			-0.034	
		(0.026)			(0.030)			(0.033)	
Overperform baseline		-0.017			-0.014			-0.004	
		(0.022)			(0.031)			(0.036)	
Information treatment			-0.002			0.020			0.023
			(0.012)			(0.024)			(0.025)
Start-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control mean	0.09		0.11	0.44		0.43	0.38		0.36
Overperformers: Control mean		0.08			0.42			0.37	
Underperformers: Control mean		0.10			0.46			0.39	
Observations	3,213	3,213	3,213	3,213	3,213	3,213	2,869	2,869	2,869
Clusters	112	112	112	112	112	112	112	112	112

Notes: This table reports the effect of the scorecards and information treatment on attrition from the survey. In columns (1)-(3), attrition is measured from the in-person survey to the phone follow-up survey. In columns (4)-(6), attrition is measured from the in-person survey until the respondent answered the question about the typical bribe payment. In columns (7)-(9), the outcome variable is an indicator for if the respondent did not reply to the question about a typical bribe payment, conditional on participating in the follow-up survey. Standard errors are clustered at the land-office level. Observations are inversely weighted by the number of observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A13: Lower Lee Bounds for Scorecards' Effects on Bribes

	Typical payment		Reported payment	
	(1)	(2)	(3)	(4)
Scorecard	871 (617)		233 (178)	
Scorecard × Overperform baseline		1,627** (729)		546** (219)
Scorecard × Underperform baseline		-105 (967)		-23 (255)
Overperform baseline		-1,830* (969)		-808*** (292)
Start-month FE	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes
Observations	1,808	1,829	2,922	2,922
Clusters	112	112	112	112

Notes: This table reports the lower Lee bounds for the effects on bribe payments. If the effect on attrition was positive (negative), randomly selected observations are added back to the treatment (control) data and assigned a value of zero (the largest reported bribe). Standard errors clustered at the land-office level. Observations inversely weighted by observations in land office. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table A14: Scorecards' Effects on Applications Received and Land Size

	ln(Applications received)	ln(Land size)
	(1)	(2)
Panel A: Overall effects		
Scorecard	-0.053 (0.073)	-0.028 (0.063)
Panel B: Heterogeneous effects		
Scorecard × Overperform	-0.062 (0.107)	0.025 (0.092)
Scorecard × Underperform	-0.044 (0.106)	-0.087 (0.091)
Overperform baseline	0.147 (0.123)	-0.088 (0.110)
Start-month FE	No	Yes
Stratum FE	Yes	Yes
Observations	311	1,007,014
Clusters		311

Notes: This table reports the effect of the scorecards on the number of applications received and the land size of those applications. In column (1), observations are at the land-office level. In column (2), observations are at the application level. Data contains all applications made between one month before the experiment started and 45 working days before the experiment ended. Standard errors clustered at the land-office level. Column (2) is inversely weighted by the number of observations in land office. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table A15: Scorecards' Effects on Rejection Rates

Panel A: Overall effect	(1)	(2)
	Rejection	Previously rejected
Scorecard	-0.000 (0.021)	0.024 (0.020)
Panel B: Heterogeneous effects		
Scorecard × Overperform baseline	0.014 (0.029)	0.064 (0.040)
Scorecard × Underperform baseline	-0.016 (0.031)	-0.020 (0.013)
Overperform baseline	-0.068* (0.038)	-0.004 (0.029)
P-value subgroup diff.	0.49	0.06
Start-month FE	Yes	Yes
Stratum FE	Yes	Yes
Observations	1,034,688	3,064
Clusters	311	112
Control mean	0.32	0.06

Notes: This table reports the effect of the scorecards on rejections. Column (1) shows the effect on rejections in the administrative data. Column (2) shows the effect on applicants who returned after a rejection. Standard errors clustered at the office level. Observations inversely weighted by observations in land office. ***p<0.01; **p<0.05; *p<0.1.

Table A16: Spillovers on Applications Outside E-Governance System

Panel A: Overall effect	(1)	(2)	(3)	(4)
	Visits	≤45 w. days	ln(w. days)	Satisfaction (SD)
Scorecard	-1.11 (0.79)	0.00 (0.06)	-0.13 (0.10)	0.08 (0.12)
Panel B: Heterogeneous effects				
Scorecard × Overperform baseline	1.10 (1.59)	-0.07 (0.12)	-0.00 (0.22)	0.06 (0.18)
Scorecard × Underperform baseline	-1.47 (1.38)	0.07 (0.06)	-0.19** (0.09)	0.08 (0.15)
Overperform baseline	-2.13 (2.25)	0.02 (0.12)	-0.23 (0.22)	0.08 (0.26)
Start-month FE	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes
Observations	644	644	644	644
Clusters	58	58	58	58

Notes: This table reports spillover effects of the scorecards on applications made outside the e-governance system. Standard errors are clustered at the land-office level. Observations are inversely weighted by observations in that land office. ***p<0.01; **p<0.05; *p<0.1.

Table A17: Robustness of Information Treatment Effect to Different Treatment Variables

Panel A:	Typical bribe payment			
	(1)	(2)	(3)	(4)
Information treatment	212 (490)	120 (493)	211 (501)	221 (548)
Observations	1,802	1,733	1,802	1,448
Clusters	539	536	539	437
Panel B:	Typical payment in scorecard offices			
	(1)	(2)	(3)	(4)
Information treatment	300 (775)	501 (802)	300 (775)	498 (865)
Observations	930	881	930	785
Clusters	271	269	271	234
Panel C:	ln(Expected processing time)			
	(1)	(2)	(3)	(4)
Information treatment	-0.049** (0.022)	-0.044* (0.022)	-0.007 (0.022)	-0.029 (0.024)
Observations	2,467	2,467	2,467	2,064
Clusters	561	561	561	462
Start-month FE	Yes	Yes	Yes	Yes
Stratum FE	Yes	Yes	Yes	Yes
Sample	Full	Full	Full	Correct

Notes: This table shows the robustness of the estimates of the effects of the information intervention to defining the treatment variable differently with respect to noncompliance with the treatment assignment in the treatment delivery. The results investigated is the overall effect on bribe payments, the effects on bribe payments among offices receiving the scorecard treatment, and the effect on the expected processing time at the time of the in-person survey. Column (1) shows the results when using my preferred treatment variable based on the median treatment delivered in a land office survey day. Column (2) shows the results when using the actual treatment delivered for each applicant. Column (3) shows the results when using the assigned treatment for each land office survey day. Column (4) shows the results when using the assigned treatment but restrict the sample to applicants who received the assigned treatment. Outcome variables are winsorized at the 99th percentile. Standard errors are clustered at the land-office level. Observations are inversely weighted by observations in that land office. ***p<0.01; **p<0.05; *p<0.1.