

E-governance, Accountability, and Leakage in Public Programs: Experimental Evidence from a Financial Management Reform in India *

Abhijit Banerjee[†], Esther Duflo[‡], Clement Imbert[§],
Santhosh Mathew,[¶] Rohini Pande^{||}

October 19, 2016

Abstract

In collaboration with the Government of Bihar, India, we conducted a large-scale experiment to evaluate whether transparency in fiscal transfer systems can increase accountability and reduce corruption in the implementation of a workfare program. The reforms introduced electronic fund-flow, cut out administrative tiers, and switched the basis of transfer amounts from forecasts to documented expenditures. Treatment reduced leakages along three measures: expenditures and hours claimed dropped while an independent household survey found no impact on actual employment and wages received; a matching exercise reveals a reduction in fake households on payrolls; and local program officials' self-reported median personal assets fell.

*We thank Abhishek Anand, Madeline Duhon, Pooja Khosla, Shweta Rajwade, M.R. Sharan, Niki Shrestha and Pankaj Verma for excellent research assistance. We thank the International Initiative for Impact Evaluation (3ie) and the International Growth Centre (IGC) for financial support. We thank Julien Labonne, Eliana la Ferrara, Rema Hanna, Ben Olken, Debraj Ray, and Sandip Sukhtankar for helpful comments. As the principal secretary of Bihar's Department of Rural Development, Santhosh Mathew was involved in the design of this reform and early implementation efforts. However, all views and errors are solely ours and this paper does not represent the view of the Ministry of Rural Development or any part of the Government of India. The project was approved by the Institutional Review Boards at Harvard, IFMR, and MIT (COUHES Protocol 1207005145). This study is registered in the AEA RCT Registry and the unique identifying number is AEARCTR-0000009.

[†]MIT, banerjee@mit.edu

[‡]MIT, eduflo@mit.edu

[§]Warwick, c.imbert@warwick.ac.uk

[¶]Ministry of Rural Development s.mathew@ids.ac.uk

^{||}Harvard, rohini_pande@harvard.edu

1 Introduction

Implementation bottlenecks constrain the effectiveness of social programs the world over, but their costs – in terms of reducing program inefficiency and creating opportunities for officials to seek rents – are particularly severe in the developing world (?). The theoretical literature on corruption has long emphasized the importance of a lean administrative structure and streamlined organization of tasks in safeguarding against malfeasance (???). Yet, somewhat perversely, implementation bottlenecks are often themselves a consequence of government-instituted accountability mechanisms: more monitoring means a longer pipeline with more joints that can spring leaks.

Empirical studies of corruption have typically focussed on the effects of information disclosure, increased monitoring, and monetary incentives, while holding the administrative structure constant (???). A few papers examine the effect of changes in the number of independent and potentially competing functionaries or jurisdictions (??) and the impact of reducing bureaucratic discretion (??). However, other aspects of the bureaucratic architecture are rarely the subject of study by empirical economists, despite receiving significant attention in the public administration literature (????).

In this paper, we focus on how corruption levels respond to changes in administrative structure of fiscal transfers made possible by innovations in e-governance.

The fund-flow mechanism for decentralized programs traditionally involves a transfer of funds from a higher level of government (the state or the federal government in the case of India) to the local implementing body. The standard practice in low-income countries is for funds to be disbursed as advances. This is because the communications of the needs from local authorities to the center and the physical transfer of funds back out to them (as well as the ancillary mechanisms for checking that it does not get “lost” on the way) can be time-consuming if the communication infrastructure is low quality and distances are large. Without advance financing, the local authorities would need the parties expecting payment to extend credit for what could be long and unpredictable lengths of time. This would considerably constrain the local government’s ability to implement the program.

The downside to the cash-advance system is that the local authorities acquire temporary control rights over the advances and can delay accounting for expenses till a time convenient for them.¹ To reduce the possibility of malfeasance, it is standard public-sector management practice to require that fund-requests by local governments be ratified at immediately higher levels of government (in effect certifying trust that the money will

¹Of course, the higher level of government could (and sometimes does), refuse to send the next tranche until the current one is fully accounted for. However, given the delays in getting all the receipts together, preparing the documents, sending them and then getting the funds released and sent, this creates problematic long gaps between tranches. For this reason, the deadline for the full accounting for a particular tranche of money tends to be loose, which makes it harder to verify these expenses (people move, they forget, receipts get lost, etc). This, in turn makes it easier for local officials to get away with malfeasance.

be spent appropriately). But creating a chain of intermediaries with veto-power over advances both slows down the process (and therefore makes larger advances necessary) *and* increases the number of players who can rent-seek.

If additional intermediaries do indeed translate into greater leakage, then electronic platforms that enable an immediate link between fund transfer and program expenditures should lower the scope for corruption at the ground level. We use an unusually large-scale randomized experiment to examine this possibility in the context of India's Mahatma Gandhi National Rural Employment Guarantee Scheme (MGNREGS), the world's largest workfare program. MGNREGS is supposed to provide employment at a given wage to all those who request work and use that labor to improve local infrastructure. The evaluation was conducted between September 2012 and March 2013 in conjunction with Bihar's Department of Rural Development and spanned 12 districts with a population of 33 million.

In the status quo system, funds flowed through four tiers of administrative hierarchy on their way from the Department of Rural Development to the village authority: state, district, block and Gram Panchayat (GP). The GP could request advance funds without specifying intended purpose, but authorities at intermediate levels of the hierarchy (the block and the district) had to approve the request before it went to the state treasury. In the reformed system, fund disbursement to a GP for labor expenses was based on incurred expenditures. Specifically, GP officials entered the names of those employed and wages owed in a central database, which automatically triggered fund release into the GP account. The GP official no longer required approval from block or district officials for the submission of the fund request (although many block officials remained involved, as the data-entry infrastructure was typically only available at the block level). All other aspects of the fund-flow process remained unchanged.²

Not only did the reform reduce involvement of higher administrative tiers, it also enhanced transparency. The fact that each fund request required a list of beneficiaries enabled a more effective audit process – it eliminated the several-month lag between fund transfer and wage payment and when the names of those purportedly paid were available to the auditor.

Theoretically, the impact of this kind of fiscal reform on leakage and corruption is ambiguous. On one hand, increased transparency should improve monitoring and, thereby, reduce rent-seeking. On the other hand, if rent-seeking by GP officials is constrained by the need to reserve some of the spoils for those higher up the administrative hierarchy, then removing tiers could increase, rather than decrease, rent-seeking by GP officials. The net effect, therefore, depends on the trade-off for the GP official between the threat

²This included fund disbursement from GP account to villagers and the subsequent uploading of information on who was hired, for how many days, and for what payments, as well as expenditures on materials purchases. This information was uploaded by GP officials on a separate public access database (nrega.nic.in).

of retribution through increased monitoring and the promise of a bigger slice of the ill-gotten pie. In situations with weak state capacity, one may be concerned that the latter effect would dominate.

Our empirical analysis focuses on identifying this net impact. A key contribution is to provide robust evidence of declines in leakages by triangulating across multiple data sources, including administrative databases and independent surveys. First, using administrative data on daily GP finances, we show a 17 percent expenditure reduction in treatment GPs relative to control GPs. We corroborate the decline with spending data reported in the MGNREGS public access database where we also observe a corresponding decline in the reported number of hired workers in treatment GPs. Meanwhile, in an independent household survey, we find that the number of beneficiaries, the wage payments received and assets built are statistically indistinguishable across treatment and control GPs.

While this seems to be evidence that the reform reduced leakage, a remaining concern is that the relatively small survey sample size may have limited our ability to identify employment declines in treatment GPs. We, therefore, bolster our analysis with two pieces of direct evidence on the reduction in corruption.

First, we construct a measure of leakage that fully exploits the scale of our experiment and the large amount of data available. To systematically identify “ghost workers” (households who are reported to have worked but, in fact, do not exist), we take the 6,292,307 names that the public database reports as having worked on the program and match them with names from the Socio-Economic Census, which the Government of India conducted in 2012 (and itself yields a database containing 34 million names for the 12 districts of our sample). First, using a Hindi-specific Levenshtein linguistic algorithm we match the names of our approximately 18,000 sample villages across the two databases. Then we use the same method to match household names within a given village.³ The name of a ghost worker should fail to be matched, except in the rare case where there are two persons with the same full name and gender in the same village. This matching-based strategy systematizes and implements at scale the audit approach pioneered by ? and also used by ? where investigators physically track down workers reported in the public database. The matching process is imperfect (there are errors in both direction); however, the scale of our experiment allows it to serve as a statistical test of the impact of the reform on corruption. The fraction of unmatched households is significantly lower in GPs where our reform was implemented: for example, 35.5 percent of single-worker households (which make up the majority in our sample) are unmatched in the control group, compared to 33.6 percent in treatment villages – a reduction of 5 percent. This difference is absent outside the reform period.

³Our analysis is based on the adaptation of a code developed and graciously made public by Paul Novosad.

Our second measure of corruption traces the “missing money” by examining affidavit data on public employee assets reported just after the reform period. We find that median wealth of block and GP officials is 14 percent lower in treatment relative to control areas, and a Kolmogorov-Smirnov test rejects equality of the two distributions. If we use mean wealth as the measure, the decline is of similar magnitude, though more noisily estimated. Taken at face value, the point estimate would imply that this decline in officials’ wealth accounts for half of the savings the reformed program achieved.

Turning to other dimensions of program performance, we observe a decline in idle funds sitting in GP accounts, which represents an implementation efficiency gain from a public accounting perspective, since disbursed funds are considered a government expense.⁴ Specifically, the reform reduced fiscal transfers in treatment GPs by 24 percent, of which the decline in expenditure accounted for two-thirds. The other one-third reflects a decline in idle funds in GP accounts. It is conceivable that treatment, everything equal, increased the budget available for the program in *all* GPs (treatment and control), by freeing up funds that were previously idle balances. On the other hand, the reform did not directly improve program delivery for villages or beneficiaries: we do not see an increase in the number of work-days or constructed assets in treatment GPs, and we do see an initial increase in delays in payment for beneficiaries in the treatment group, though they declined over time.

Our paper contributes to a growing set of studies which evaluate administrative reforms in settings with limited state capacity (????). Several recent papers focus on the use of information technology, or e-governance (???), and of those the most closely related to ours is the ‘smart card’ for MGNREGS project in Andhra Pradesh studied by ?. Under this reform, beneficiaries received biometric smart cards and wage disbursement was shifted from the post office to locally hired bank employees armed with a Point-of-Sale machine for verifying identity. From the beneficiary perspective, the reform gave them more control over the process (since they had to be present to use the smart card) and made it more difficult for local authorities to skim off worker wages by colluding with the post office officials. ? find that the intervention increased worker payments and, consequently, household incomes, with no change in government outlay, indicating lower leakage. This test of disbursement reform complements our intervention which changes the fund flow and leaves disbursement processes constant.

Our paper is also inspired by the literature on ways to objectively estimate corruption (see ? for a review). Using a randomized control trial to examine the impact of administrative reforms on the incidence of corruption, we follow ? and ? in combining the

⁴However since public sector banks handle the money, from the point of view of the government as a whole, only the expenses involved in handling this extra money (which includes interest on extra funds that the government borrowed from the money market) is actually a cost; the rest of it is a transfer from one set of state-owned accounts to another. The fact that more funds are now available to use on other programs is also a potential gain.

“forensic” method of tracking expenditure by comparing official records of funds release with actual receipt by beneficiaries. However to the best of our knowledge, ours is the first study to carry out the forensic exercise with administrative rather than survey data. Such cross-validation across administrative data sources provides a promising avenue to detect corruption and a possible basis for effective auditing. Finally, like ? we use officials’ affidavit data to examine wealth effects attributable to corruption – our innovation is to use these data in the context of a large scale experimental evaluation of an administrative reform.

The rest of the paper is structured as follows. Section ?? presents the context for the reform and its expected impact. Section ?? details the data we use and our empirical strategy. Section ?? presents the results, and Section ?? concludes.

2 Background and intervention

India’s Mahatma Gandhi National Rural Employment Guarantee Scheme (MGNREGS) was created in 2005 by the National Rural Employment Guarantee Act. The Act guarantees every rural household 100 days of unskilled manual labor at the stipulated minimum wage per year. Local GP officials are responsible for registering beneficiaries and providing them work on local infrastructure projects. With close to 50 million beneficiary households in 2013, the MGNREGS is the largest social protection program in the world today, costing 0.5 percent of India’s GDP.

Below, we first describe the relevant program aspects for Bihar and the reform that we evaluate. We then use a simple theoretical framework to identify conditions under which the reform will reduce leakages.

2.1 MGNREGS in Bihar: Performance and program monitoring

From the start, the quality of MGNREGS program implementation has differed across Indian states. ? note “the incidence of unmet demand tends to be higher in poorer states even though demand for the scheme is higher.” This is particularly stark in the case of our study state Bihar, which has one of the highest poverty rates in India, and possibly the highest unmet demand for MGNREGS work. Using National Sample Survey data for 2009-2010, we estimate that 77 percent of households in Bihar who wanted MGNREGS work could not obtain it, and at most 10 percent of households have worked on MGNREGS worksites during the year. By comparison, in the better performing state of Andhra Pradesh, only 27 percent of those who wanted work could not find it, and 39 percent of households participated in MGNREGS.

The quality of MGNREGS implementation has improved over time, likely due, in part, to more regular auditing and the channeling of payments directly to beneficiaries

through banks or post offices.⁵

An important impetus for audit reform was a MGNREGS corruption enquiry conducted by India's federal vigilance authority (CBI) in the neighboring state of Orissa. This led many Indian state governments - including Bihar - to tighten up their internal audit systems. In the case of Bihar, this included official audits conducted by teams of administrators and engineers, as well as social audits where officials convened villagers at a meeting and heard their grievances.

In June 2011, the Bihar principal secretary for Rural Development sent district authorities a letter noting the MGNREGS program requirement that block officers undertake random weekly audits of ongoing and recently concluded works. In November 2011, revised department guidelines clarified that the MGNREGS public database should be used for audit and also that additional MGNREGS documentation should be made available to the official audit teams and during the social audit. Finally, coinciding with the start of our reform, the state government issued an audit reform letter on September 1, 2012. This letter explicitly stated that projects to be audited should be chosen from the set completed in 2011-12 and those ongoing in Fiscal year 2012-13 (according to nrega.nic.in). According to official data, between June 2012 and May 2013, 64% of the GPs in our sample districts were audited at least once (?).

However, leakage of funds remains an important program concern. For our control GPs, a comparison of outcomes in our independent household survey to the public access database shows that the (appropriately weighted) number of households who say they have worked in MGNREGS sites account for only 59 percent of households listed as having worked in that period in the official database.⁶ We also surveyed 346 GP heads (Mukhiyas) and 47 percent of them in control GPs mentioned corruption in the administration as a major implementation issue. On average, they estimated the system of "taxes" extracted by MGNREGS functionaries as making up 21-30 percent of program expenditures. 72 percent of Mukhiyas in control GPs also identified a lack of funds as a reason for poor program implementation.

⁵For an early program year (June 2007 and July 2008), employment estimates from national survey data only account for 42-56% of official figures on MGNREGS employment (?). Four years later (July 2011 to June 2012), the same method shows that about 80% of the reported workdays could be accounted for (?). These national household surveys, however, cannot provide reliable state-level leakage estimates.

⁶Using data from a household survey representative of the whole of Bihar in 2009-10, ? estimate significant, but somewhat smaller leakages of MGNREGS funds (20-30%). A possible explanation is that our survey specifically checked with the respondents which MGNREGS project they had worked on (using the list of MGNREGS projects from nrega.nic.in). Hence, we may be less likely to assign other state-run public works project as MGNREGS work.

2.2 Fund-flow management in MGNREGS

2.2.1 Fiscal architecture

MGNREGS is largely financed by India's federal government but implemented by local GP officials. The first level of fund transfers for MGNREGS are tranche-wise transfers from the central government to the state: the first tranche is provided at the start of the fiscal year on the basis of anticipated demand and expenditure from previous years; additional tranches are supposed to be available upon request by the state. To enable expenditure accountability, the central government releases these subsequent fund tranches only after the state accounts for a minimum fraction of labor expenditures by documenting worker details (and amounts paid) on a publicly accessible electronic data collection system (nrega.nic.in) – this reporting also serves as the basis for audits. Just after our reform began (in September 2012) we saw this policy in action: the central government refused to release the requested tranche of funds to Bihar until 60% of labor expenditures was accounted for in nrega.nic.in.

Turning to within-state fund flows, fund requests originate from GP authorities and are then aggregated up the chain to the state-level at the start of each financial year (we discuss this further below). Historically, once disbursed from the state treasury, funds move down the administrative hierarchy: via districts and blocks to GP accounts. Since the money disbursed to a district (and then block) is typically less than the total requested by the lower level of hierarchy, each administrative tier enjoys significant discretion in resource allocation to the tier below. The lumpy and sporadic nature of transfers also implies that at each level (below the state) some units lack funds while others have large unspent amounts.

In 2010-11, the Bihar government reformed its fiscal architecture to prevent unspent funds from accumulating in districts. It created a single state account to receive central transfers and opened district Zero Balance Accounts such that funds withdrawn from the district account would be automatically replenished. Alongside, an electronic platform called Central Planning Scheme Monitoring System (CPSMS) was created. This both allowed the state government to monitor GP account balances and to directly transfer funds from the state pool to the GP account upon district authorization. Thus, funds no longer transited through district and block accounts before reaching the GPs.

Finally, to reduce the discretion enjoyed by block and district administration in passing on fund requests by the GPs, guidelines were issued requiring districts to transfer funds to a GP whenever its account balance fell below Rs. 100,000. However, in practice these guidelines were not followed and fund requests continued to involve bargaining between the district, the block and the GP. Our analysis of GP accounts fund-flow data for 12 Bihar districts shows that the average time taken to replenish a GP account was about three months between July 2011 and July 2012.

2.2.2 Experimental design of Fund flow reform

Cognizant of continuing frictions in fund flows, in 2012 Bihar’s Rural Development Department decided to further reform MGNREGS fund flow within the state and to evaluate the reform experimentally.

The reform occurred between September 2012 and March 2013. It spanned 12 districts in South, West and North of Bihar, covering a rural population of 33 million, and 905,000 reported MGNREGS workers (see Figure ??). In collaboration with Bihar’s rural development department, we identified 69 treatment blocks. Specifically, in each study district, one-third of the blocks were randomly selected to implement the reformed fund flow system. Overall, the study districts were divided into 69 treatment (1033 GPs) and 126 control blocks (2034 GPs).

Figure ?? summarizes the status quo fund flow system. At the start of the financial year, each GP account receives a first tranche of funds. When these funds are exhausted and if automatic replenishment of GP account fails to occur then the GP makes a fund request to the higher administrative tier (block). This request is typically based on anticipated need and is supported by an utilization certificate for the previous tranche of funds. The block officials, who are supposed to play a monitoring role, ratify and pass the request on to the district administration who then requests a fund transfer from the state treasury to the GP savings account via the CPSMS platform.

Figure ?? describes the reformed fund flow for labor payments that was introduced in treatment blocks: the GP official logs into CPSMS and enters beneficiary details; this, in turn, initiates an automatic transfer of incurred wage expenses from the state account to GP saving account. In practice, since most GPs lack necessary infrastructure and/or knowhow, uploading of beneficiary data typically occurred at the block office with assistance from a block-level data entry operator.

The reform left three important elements of the fund-flow system unaffected. First, the final step of payments from GP to beneficiaries, was unchanged: the GP continued to send a check and a list of intended beneficiaries and amounts due to them to the local bank/post office, which then were supposed to credit the beneficiaries account.⁷ Second, the state continued to disburse payment for materials utilized for MGNREGS through CPSMS, with districts and block authorities acting as intermediaries. To enable this, GPs continued to have a savings account available for transaction use. This created a channel through which GPs could get paid for materials and divert this to pay labor, although this was prohibited through a government instruction for treatment GPs. Partly for this reason, like in ?, the implementation of the reform was both gradual and never quite complete, as we document below. Finally, independent of the within-state fund

⁷As emphasized by ?, direct payment from state treasury into the beneficiary’s account does not necessarily prevent GP authorities from claiming a part of it. For example the bank/post office staff may permit the GP official to act as a stand-in for the actual beneficiary.

flow process, GP officials were still required to document every job spell – including the identity of the beneficiary and the payment – on a public access database (nrega.nic.in).⁸

2.3 How may financial reform affect rent-seeking?

There are two distinct mechanisms through which the reformed financial flow system could impact funds leakage.

First, by directly linking each disbursement to a specific (reported) expenditure the reform facilitated monitoring. Recall that in the status quo, the GP gets an advance, and the district is supposed to replenish the account as soon as funds fell below some threshold, based on an “utilization certificate.” The utilization certificate is, in principle, backed up by the electronic entry of the “muster roll” (the information on each beneficiary, and how much they worked), which can serve as a basis for audit.

Data entry, however, significantly lagged spending. During our intervention period, we observed a delay of six months in getting 60 percent of expenditures entered, and one year to record sufficient expenditures in the public database to match the CPSMS data (see appendix table ??). Lags in data entry on the public database limits its use as a monitoring tool; long lags between purported occurrence of work and audit limits potential cross-checking of information in the field. Migration, for instance, could explain an inability to find individuals listed in the database who cannot be found in the village. And those who can be found may not remember how much they worked.⁹

In the reformed system, fund release to GPs occurred after beneficiary details were documented on the electronic platform. By directly linking fund transfer to expenditure documentation, and by enabling (almost) real time documentation, quicker verification and more effective audits became possible. It made it harder for GP authorities to create fake workers, and auditors had more recent data on who worked.

Indeed, an analysis of the reports of all audits of the MGNREGS program conducted by the Rural Development Department between May 2012 and June 2013 suggests that audits were more likely to pick up irregularities in treatment blocks (see appendix table ??, panel D). During the intervention period, the share of audits finding irregularities was similar in both groups, but in the period immediately after, it was twice as large (from 5% of audits finding irregularities in the control group to 10% in the treatment group). Since audits happen with a lag, this captures irregularities found on projects conducted during the intervention period. As we will show below, the weight of the evidences suggests that corruption in the treatment group actually declined over this period. Thus, this increase

⁸In practice, treatment GP officials entered the same information twice: once to get paid, and once after the fact. An interface between CPSMS and public portal was planned but never implemented.

⁹Santhosh Mathew witnessed “flexible memories” during a field investigation of a few cases of workers who had reported looking for, but not receiving, work and had, therefore, requested unemployment compensation. Within a few hours of his arrival, every worker had produced an affidavit stating that they had been offered, but had refused, work and were withdrawing their compensation request.

in the number of audits finding irregularities is strongly suggestive that there was indeed a greater probability of being caught in the treatment group, conditional on cheating. Table ?? also shows that audits are not infrequent. There was on average 33 projects audited in each block during the intervention period, and 9.5 during the three subsequent months.¹⁰

A second reform feature was a reduction in the number of people involved in fund disbursement. As pointed out by ? and empirically demonstrated by ?, the involvement of multiple uncoordinated agencies, each with its own rent-seeking goals, typically increases rent-seeking and hence inefficiency. However, in our setting many GP officials who were the target of rent-seeking by higher levels of administration may have been themselves engaged in rent-seeking. As a result, more rent-seeking by higher-ups in the hierarchy might, perversely, increase efficiency by discouraging stealing by GP officials (since what they steal gets taxed).

Hence the reform's impact on total leakages is a priori ambiguous: increased monitoring should reduce leakages but reduced rent-seeking by officials higher up in the hierarchy might go the other way. We formalize this argument below.

• **The status quo regime**

We label an official at tier i of the administrative hierarchy in the status quo regime as: P (GP), B (block) , D (district) and S(state). Tier P is responsible for program operation and can skim off amount s if she exerts a non-contractible non-pecuniary effort cost $\frac{1}{2}cs^2$. In expectation, the penalty for skimming is $\pi^T s$.

For P to receive s , B and D have to sign off on the fund claim. Assume, following the literature, that $i \in B, D$ can commit ex ante to a price p_i for approving *every rupee* of funds skimmed by P. Further, B and D choose p_B and p_D non-cooperatively to maximize earnings. Therefore s maximizes

$$(1 - \pi^T)s - p_i s - p_{-i} s - \frac{1}{2}cs^2,$$

which implies that

$$s = \frac{1 - \pi^T - p_B - p_D}{c},$$

i 's earnings in status quo regime is therefore

$$p_i \frac{(1 - \pi^T - p_i - p_{-i})}{c},$$

¹⁰Data on audits was compiled in July 2013 by the Rural Development Department for ?. See Appendix ?? for more details.

and the p_i that maximizes this expression is given by:

$$p_i = \frac{(1 - \pi^T - p_{-i})}{2}$$

which from the (evident) symmetry of the solution, yields

$$p_i = \frac{(1 - \pi^T)}{3}$$

and therefore the amount skimmed under the status quo is

$$s = \frac{(1 - \pi^T)}{3c}.$$

Under the status quo B (and D) therefore earn an amount

$$Y^{BT}(\pi^T) = \frac{(1 - \pi^T)^2}{9c},$$

while P earnings from skimming (which we observe) is $Y^{PT}(\pi^T) = \frac{(1 - \pi^T)(1 + 2\pi^T)}{9c}$.

Note that to compute P's utility we would need to deduct the expected penalties and the cost of her effort from this expression.

• **The new regime**

Two things change: First π^T goes up to π^N . And second, P can, in principle, unilaterally claim the money. However, she lacks the technological capacity to do so. So she needs B to collude with her. We consider two cases:

• Case 1: Assume P and B can collude and entirely cut out D: $p_D = 0$. From above it should be evident that

$$p_B = \frac{(1 - \pi^N)}{2}, \text{ and } s = \frac{(1 - \pi^N)}{2c}$$

which together imply that

$$Y^{BN}(\pi^N) = \frac{(1 - \pi^N)^2}{4c} \text{ while } Y^{PN}(\pi^N) = \frac{(1 - \pi^N)(1 + \pi^N)}{4c}.$$

A comparison of skimmed funds under the two schemes, $Y^{PT}(\pi^T)$ versus $Y^{PN}(\pi^N)$ or $Y^{BT}(\pi^T)$ versus $Y^{BN}(\pi^N)$, shows two countervailing effects: the negative effect of an increase from π^T to π^N and the positive effect of not having to pay D, reflected in the fall in the denominator. The net effect is ambiguous; for the negative effect to dominate, $1 - \pi^T$ needs to be reasonably close to zero or the increase in π^T to π^N must be very large in proportional terms. Otherwise, by reducing the number of officials involved, the

reform increases corruption.

- Case 2: Consider an arguably more realistic scenario, where D retains some leverage, so that she can continue to extract rents, but with probability $\alpha < 1$, D has to be paid a price p^D (per rupee stolen). First, if p^D can be as high as possible, the solution is identical to the status quo, but with a larger penalty:

$$\alpha p_D = p_B = \frac{(1 - \pi^N)}{3} \text{ and } s = \frac{(1 - \pi^N)}{3c}$$

The reason is straightforward: D increases p^D exactly enough to cancel out the effect of $\alpha < 1$, and the problem is solved as before. The only effect of the reform is to change the penalty rate, and skimming will unambiguously decline.

While the point that reducing D's influence encourages her to demand even more when she gets a chance, is general, the *exact neutrality result* relies on the arguably unrealistic ability of D to extract very large bribes. In particular, as $\alpha \rightarrow 0$, $p^D \rightarrow \infty$, which means that P will be, *ex post*, paying large amounts out of pocket to D whenever she can extract rents. It seems more reasonable to define a cap, \bar{p}^D , on how high p^D can go. For α small enough that \bar{p}^D binds, B maximizes

$$p_B \frac{1 - \pi^N - p_B - \alpha \bar{p}_D}{c}.$$

The p^B chosen will be

$$p_B = \frac{(1 - \pi^N - \alpha \bar{p}_D)}{2},$$

and therefore

$$s = \frac{(1 - \pi^N - \alpha \bar{p}_D)}{2c}$$

which implies that

$$Y^{BN}(\pi^N, \alpha) = \frac{(1 - \pi^N - \alpha \bar{p}_D)^2}{4c} \text{ and } Y^{PN}(\pi^N, \alpha) = \frac{(1 - \alpha \bar{p}_D)^2 - (\pi^N)^2}{4c}.$$

Clearly for $\pi^N < 1 - \alpha \bar{p}^D$ (which is the only case that makes sense), an increase in π^N reduces s , Y^{BN} and Y^{PN} , while a fall in α increases all three. Once again, the net effects are ambiguous. Finally:

$$\frac{Y^{BN}(\pi^N, \alpha)}{Y^{PN}(\pi^N, \alpha)} = \frac{1 - \alpha \bar{p}_D + \pi^N}{1 - \alpha \bar{p}_D - \pi^N}$$

This ratio goes up when π^N goes up and down when α goes down. The net effect of changing both, as occurs with the reform we study, is ambiguous: the loss may be greater for P or B.

This model, thus, demonstrates that, despite the added transparency, both the reform's impact on overall corruption and also whether it will relatively favor (or disfavor)

block or GP officials (in terms of their earnings) is ambiguous.

3 Data and experimental design

3.1 Data

Our analysis exploits multiple data sources. We first describe the administrative data sets that we use (these typically cover the universe of GPs in the experiment) and then the survey data we collected.

First, we use the daily financial database associated with the CPSMS system for the period September 2011 to January 2014. This includes all credits and debits in each treatment and control GP savings account, and allows us to monitor daily fund flow. In our analysis, we aggregate these daily transactions to compute total credit and debit for each treatment period. The data does not, however, identify transfer recipients: we can not distinguish between material and labor expenditures, nor can we identify the names of the workers being paid.

Second, we use the public access database, nrega.nic.in, which includes category-wise expenditures aggregated at the fiscal year level (i.e. April 1st to March 31st of every year). The financial year 2012-13 data includes three pre-reform and nine reform (set-up and intervention) months. Four expenditure categories are reported: unskilled labor, material, skilled labor and administrative expenses. In addition, the database includes beneficiary details: who has worked in the household, duration and dates of work and wages paid. This database includes information for all beneficiaries for whom funds have been released: for actual beneficiaries it lists days worked which include both genuine work spells and also days falsely claimed as work days (ghost days) and for ghost workers (those who did not work but against who's name payment was released) it lists names and days purportedly worked.

Third, we obtained data from India's Socio-Economic Caste Census (SECC), which was conducted in 2012, for the 12 study districts. These data cover 16,480 villages across 195 blocks and for each household in the village include name and age of each household member (and relationship to household head). We have data for 34 million individuals, living in more than five million households.

Our matching exercise across SECC and the public access database is a population-level version of the forensic method pioneered by ?, that cross-check administrative data with household survey data.

First, we use an algorithm to match village names across the databases.¹¹ Among

¹¹Since MGNREGS basic administrative unit is the GP, not the census village, the database lacks a village census code. 84% of villages in the MGNREGS database have a match in the SECC census. For the 16% remaining one, we look for matches in all the villages in the GP

matched villages, we use the same algorithm to find a match for each household with a job card in the public access database in 2014 (for more details, see Appendix ??).¹²

Our outcome of interest is the match rate, defined separately for people reported to have worked during and after the intervention period: a household (name) with a job-card in the public access database but missing in the SECC database is more likely to be a “ghost” than a household (name) found in both.

The matching process is probabilistic (based on a threshold), with errors in both directions: individuals may be omitted from the SECC census for example, or the matching could fail because names are spelled too differently to match, or on the other side, two different persons with the same name could be incorrectly matched. That said, there is no reason to expect differential errors across treatment and control groups. Note that this exercise only identifies non-existent workers, not households who report working but in reality never did (and, of course, it does not capture over-reporting of days by working households). On average, in the control villages, we match 50% of the job cards where work was reported during our intervention period to a household in the SECC (67% of the single-worker job cards, and 28% of the job cards with more than one worker).¹³ This is comparable to 59 percent match rate we obtain by comparing the public access database to (population) estimates of workers from our household survey.

Fourth, we use affidavit data on GP and block official assets. In 2012 and 2013, Bihar’s Rural Development Department instructed GP, block and district employees to declare their and their spouse’s personal assets, both movable (cash, jewelery, vehicles) and immovable (land, real estate). Since the data is self-reported it should, of course, be treated with some caution. Recent studies, however, show that the affidavit data contains useful signal.¹⁴

Finally, we use data from surveys we conducted: in May-July 2013 we conducted an independent survey of 10,036 households in 390 GPs to measure MGNREGS participation, employment and payments. We randomly sampled two GPs per block, and 25 households per GP, oversampling poor households, who were more likely to participate in the MGNREGS (see Appendix Table ?? for details). Starting July 2012, each household member was asked about weekly MGNREGS participation and the amount, date

¹²To determine whether two names (village or individual) match, we start from an algorithm developed by Paul Novosad (starting from a standard string matching algorithm, adjusted for language and tested in a large sample), and graciously made publicly available. We adjusted the algorithm for our application. For household job card with one individual, we match the individual based on names (first, last and middle) and gender. When a job card has two or more individuals, we look for a household in the SECC data base with two individuals whose names and gender match that on the job card. The match rate is lower for households with two or more working members.

¹³The difference between single worker and multiple worker households is natural – it is harder for two names to match than one.

¹⁴? use politician affidavit data and show a 3% to 4% higher estimated annual growth rate of wealth for winners than for runner-ups in close election. ? further show that the requirement to disclose discourages several politicians from even running for office.

and payments for each work-spell. As MGNREGS participation was extremely low during our study period, the survey (despite reasonable sample size) only identifies a small number of participants. Hence, estimated treatment effects using the survey are quite imprecise. We also interviewed the elected GP head (the Mukhiya) in 346 of the 390 survey GPs about the main issues they faced in implementing MGNREGS.

Alongside, we surveyed 4,165 MGNREGS infrastructure projects (10 per GP) randomly sampled from the official list of ongoing and recently completed projects (nrega.nic.in). Surveyors recorded whether the asset was found and whether it was completed.

3.2 Reform implementation

A key prerequisite for the reform was IT infrastructure to enable GPs to connect with CPSMS (computers, data entry operators, generator to ensure constant power supply, Internet access, scanner and printer). Appendix Table ?? shows that a minority of blocks had the required facilities in July 2012 but that by January 2013 a majority of treatment blocks had the needed equipment. In large part, this reflected a big push to procure and install IT infrastructure in treatment blocks during the “set up” months of July and August 2012.

The intervention was officially launched on September 8, 2012, but faced multiple implementation hurdles. In October, the central government froze program fund release as less than 60 percent of expenditures incurred since April 2012 in Bihar had been documented on nrega.nic.in. Funds were only released mid-December once data documentation was completed. As soon as the money arrived in December, GP functionaries launched a two-week strike. Figure ?? shows that MGNREGS spending fell sharply in September and rose only slowly in January 2013. This, in part, reflects seasonality: MGNREGS work-sites often close during the peak agricultural season (between July to December (?)). However the dip was longer and stronger that year. Finally, the bank which processed payments entered on CPSMS initially lacked resources to deal with the large number of small invoices sent by treatment GPs, and gave priority to the fewer large invoices coming from Control GPs. By December 2012, the bank increased its capacity and treatment GPs started sending larger invoices.

Thus, the fund-flow reforms really became operational in January 2013. Figure ?? shows that the fraction of treatment GP that used CPSMS at least once increased from less than 20 percent in December 2012 across all districts to 60 percent in April 2013. We observe significant heterogeneity across districts: the best performing district, Begusarai had more than two-thirds of GPs using the system in December 2012, and that proportion reached more than 90 percent in April 2013. By contrast, the fraction of GPs using the system in Madhubani, the worst performing district, only increased to 40 percent by April 2013. Treatment GPs that did not use CPSMS to draw funds were prohibited from

receiving funds for wage payments through another route but could still spend from their savings account. Only 1.5 percent of treatment GPs did not spend any money during the intervention period. This imperfect implementation of an at-scale reform is reminiscent of the difficulties encountered by other evaluation of at-scale government programs (??).

3.3 Randomization check

The random selection of treatment blocks ensures, in principle, that GPs in the 69 treatment blocks are ex ante identical to GPs in the 126 control blocks. To check this, we estimate regressions of the form:

$$X_{pd} = \alpha + \beta T_p + \eta_d + \varepsilon_p$$

where X_{pd} is a vector of baseline characteristics of GP p in district d , T_p is a dummy which is equal to one if GP p is in a treatment block, η_d are district fixed effects, and errors ε_p are assumed to be correlated within each block. The estimated coefficient β represent pre-treatment differences between treatment and control GP.

Table ?? presents the results. We observe very few significant differences: Villages in treatment and control GPs had similar socio-demographic characteristics and had the same level of infrastructures according to 2011 census. Our survey of 390 GPs also shows that households in treatment and control GPs have similar characteristics. Finally, according to the public access database, treatment GPs had 13% higher MGNREGS labor expenditures in the financial year preceding the intervention (April 2011-March 2012), and the difference is significant at the 5% level. However, since total MGNREGS spending between treatment and control GPs was similar at baseline according to CPSMS, and we observe no statistically significant difference in work days, workers, or material expenditure in the public access database for the financial year 2011-12, we conclude this difference in labor expenditures in the database is a reporting error or a fluke, rather than reflecting systematic differences between treatment and control GPs.¹⁵

4 Results

4.1 Financial data

In Table ??, we use GP-level financial data from CPSMS (balances, expenditures, and total debit data) to evaluate the impact of the reform on program finances. Let Y_{pdt} denote the outcome for GP p in district d for period t .¹⁶ As before, T_p is a dummy

¹⁵It is also worth keeping in mind that we will find that labor expenditures go down in treatment GPs relative to control and therefore this baseline imbalance would bias our results towards zero, if anything.

¹⁶CPSMS reports daily transactions, which we aggregate by period for the purpose of the analysis.

variable which equals to one if GP p is in a treatment block and η_d is a district fixed effect. We estimate the following equation:

$$Y_{pdt} = \alpha + \beta T_p + \eta_d + \varepsilon_{pt} \quad (1)$$

where errors ε_{pt} are clustered at the block level. The coefficient β estimates the treatment effect when t is the treatment period (September 2012 to March 2013). We divide the pre-intervention period to consider separately the July-August (set up) period. We also split the interventions period between the September-December 2012 period, when the state pool of funds was dry and the PRS were on strike, and the January-March 2013 period, when MGNREGS was working relatively smoothly. We do not include any control variables in our estimation.

Figure ?? plots average daily spending in treatment and control GPs between July 2011 to January 2014. Since MGNREGS work largely occurs in agricultural lean season, we observe significant seasonality in spending in the fiscal year prior to treatment (?). The pre-reform spending trends are similar across treatment and control GPs (and this is also true for the set-up months of July and August). Between September 2012 and March 2013, spending in treatment GPs is significantly lower than control GPs. Once the intervention is rolled back on April 1, 2013, treatment and control GPs rapidly converge to similar spending levels.

In Panel A of Table ?? we summarize these findings: Spending levels are similar across treatment and control GPs before the reform, and during the set-up period (July-August). Between September to December 2012, spending is 19% lower in treatment GPs, and from January to March 2013 it is 31% lower. After April 2013 treatment and control GPs report similar spending.

In Panel B the outcome variable of interest is the closing balance in GP accounts. This closing balance was similar across treatment and control GPs at the start of treatment in September 2012 and then, reflecting the freeze on funds transfer from the center to the state, similarly declined in both groups as GPs depleted funds until December 2012.

In December 2012, the state account was replenished and control GPs received large inflows corresponding to outstanding tranches, while treatment GPs only received funds corresponding to expenditures they had documented in the electronic system, and which they immediately used to pay wages. As a result, by the end of the reform period in April 2013, the account balance in treatment GPs was 33% lower than that in control GPs.

By April 1, 2013, MGNREGS expenditure in treatment GPs relative to control had declined by 17% and GP account balances were reduced by 30%. Panel C in Table ?? shows that the combination of lower GP spending and a decline in idle funds in the treatment GP accounts, reduced program expenditure by 24% in treatment GPs. This, in turn,

translates into a cost saving of roughly 6 million dollars.¹⁷ An immediate question – which we address below – is whether this reduction in program costs reflected a decline in real outcomes (days of employment offered, and assets built), or a reduction in leakage, or both. The expenditures were not just postponed: in the six months following the intervention, the difference between treatment and control group goes back to zero.

In Table ?? we examine program finance impacts using a different data source: expenditure data from the program’s public data portal (nrega.nic.in). In both treatment and control GPs, officials faced identical requirements on electronically reporting beneficiary details (name, payment received, work spell) that then feature on the public data portal. While data entry occurs with significant lag, eventually it does accounts for close to 100% of the expenditures observed in the CPSMS financial database.¹⁸ As these data are aggregated to the fiscal year (from April to March), we present the results for 2011-2012 (before the intervention), 2012-2013 (which includes the intervention), and 2013-2014 (after the intervention). We continue to report regressions of the form in equation (??).

The public portal expenditure data shows a decline in MGNREGS spending in treatment GPs in line with the CPSMS data. For the fiscal year 2012-13, labor and material expenditures were respectively 16% and 14% lower in treatment GPs. Note that fiscal year includes three pre-intervention months. Accounting for the different time spans of 9 and 12 months respectively, nrega.nic.in data provide slightly more negative treatment estimates on spending than CPSMS data.

It may seem surprising that both labor and material expenditures declined in the same proportion, when the financial reform only affected labor expenditures. However, by law, MGNREGS material expenditure may not exceed 40% of total spending on a project. As Table 3 shows, for the average GP, the rule is close to binding: expenditures on material amounted to 36% and 38% of total expenditure in the financial year 2012-13 and 2013-14, respectively.

4.2 Reported beneficiary outcomes

In Table ??, we use data from the public portal. As we noted, the Government of India insists on the reporting of beneficiary information, and the beneficiary data as reported in the portal matches the total that are reportedly spent on beneficiaries. Therefore, it is not surprising that the treatment effects matches what we found in the financial data. In Panel A, consistent with lower labor expenditures, we observe a negative treatment effect on the number of work days reported during the reform period (Columns 3 - 5).

¹⁷We obtain this figure by multiplying the reduction in expenditure per GP by the number of treatment GPs, and converting the total of $3.19 \times 1003 = 3,204$ lakhs Rupees into million dollars (using the April 1, 2013 INR/USD exchange rate of 0.0183).

¹⁸Appendix Table ?? compares annual expenditures per GP in CPSMS and nrega.nic.in. The discrepancies are only about 8-11% in 2012-13.

In Panel B, we observe no effect on the days per working household and Panel C shows that this decline comes entirely from a reduction in the number of individuals who have supposedly worked. The estimated treatment effect for the intervention period is a 13% decline in the number of days reported, and a 10% decline in the number of working households.

4.3 Real outcomes

Did the reported drop in MGNREGS expenditures and employment in the portal reflect actual changes in program implementation or reduced leakage (was there less work done or just less ghost work)? To find this out, we conducted household and asset surveys.

in Table ??, we consider the household survey data. Let Y_{hdt} denote outcome for household h in district d at period t and T_h is a dummy variable for whether the household lives in a treatment block:

$$Y_{hdt} = \alpha + \beta T_h + \delta Z_h + \eta_d + \varepsilon_{ht} \quad (2)$$

Z_h denotes a vector of household characteristics, which includes religion, caste, gender and literacy of the head of the household, household size, the number of adults in the household, the type of house which the household occupies and a dummy variable for whether the household owns land. Standard errors are clustered at the block level.

We lack baseline data, but estimate separate regressions for the set-up period, the two phases of the reform, and a short post-period. To account for over-sampling of poorer households, our estimation of the treatment effect on household outcomes uses sampling weights, and thus reflects village-level population averages.

Using data from the detailed survey module on MGNREGS employment, which asked about every MGNREGS participation spell between July 2012 and March 2013, we construct three MGNREGS employment measures: first, a binary indicator of MGNREGS participation; second, the number of weeks in which households declares having worked in MGNREGS; and third the number of days worked.

Panel A of Table ?? reports treatment impacts on the probability of participating in MGNREGS during the set-up period (July-August), the two halves of the intervention period, the whole intervention period (September 2012-March 2013) and the post period. The observed MGNREGS participation rates between September and March 2013, while low (below 4%), are consistent with National Sample Survey data: For the year 2011-2012, the NSS reports a participation rate of 9%. Aggregating over the entire year 2012-2013, we find a participation rate of 8%.¹⁹ The lower number during our reform period is likely

¹⁹We also asked the household head whether anybody had participated in the scheme “since the last rainy season,” and 9% of households report that they did. There is no treatment effect on this variable either, see Appendix Table ??

due to the fact that it fell outside the peak season of MGNREGS work.

The treatment effect is significantly negative and large in proportion in the set-up period (July-August), most likely reflecting a sharp drop in work provision while officials were setting up the infrastructure. During the intervention period (columns 2-4) the effect is positive and insignificant. The 95% confidence interval, expressed in fraction of the control mean is $[-5\%; +42\%]$, i.e. we can reject at a 95% confidence level a decline of 5% in NREGA participation. Thus, the significant negative impact we observe on the number of households hired in the NREGA database (10%) appears to be a pure reporting effect, and does not reflect an actual decline in the provision of work. Post intervention, the participation returns to the same level in treatment and control group.

Panel B looks at numbers of days worked (set as zero for households who did not participate during a given period). We also find a negative point estimate during the set up period, positive point estimate during the two interventions period, and overall a positive point estimate for the whole period. We can reject a reduction of 8% at the 95% confidence level, a smaller decline than the 13% we find in the NREGA database. Once again, this suggest that the reduction in workdays reported in the NREGA database is mainly due to a reporting effect.

Panel C considers reported wage payments. For each spell worked in the MGNREGS, the respondents declared whether, when, and how much they had been paid, and we are attributing each payment to the time period where the work happened, regardless of when it was made (Panel D directly looks at delay).²⁰ Unfortunately, the payment data is based on relatively few observations and is quite noisy. Consistent with a lower probability of working, wage payments were significantly lower in the treatment GPs during the set-up period. During the intervention periods, the estimates are imprecise and not significant, but the point estimate suggests a slight decline in payment during the first period, and a slight increase during the second period. Overall, the point estimate is positive (11.96), but the 95% confidence interval, expressed in percentages, is $[-27\%; +52.2\%]$. Thus, we cannot reject at the 5% level the hypothesis that the wages declined by as much (in proportion) as the total debit from the Panchayat accounts, although we cannot reject large increases either: the data seems to be too noisy to be informative.

In Panel D, we examine worker-reported delays in MGNREGS payments. This is based on very few workers, so needs to be taken with some caution. However, as compared to an average delay of 72 days in the control, workers employed during the first phase of the intervention (Sep-Dec 2012) in treatment blocks waited an extra 50 days for their payment. The effect is large, and statistically significant. We observed significant but smaller payment delays during the second phase of the intervention (27 days). These

²⁰If the payment has not happened yet, this is set as zero. Replacing it by missing does not change the estimate very much, though it makes the treatment looks more positive, since delays increased in the treatment group.

results suggest that the intervention slowed down the disbursement of funds to GPs, and delayed payments to workers, especially during the first phase of the intervention.²¹

The increase in payment delays is a significant downside of the intervention, at least initially. The program objective was to speed up payments by reducing steps in the fund flow, but it seemed to have had the opposite effect. We can identify two implementation-related reasons: first, in the early days of the intervention the bank handling CPSMS payment found itself deluged with small payment requests from the treatment GPs. The bank's response was to wait and collect a large batch of invoices before processing them together, which caused delays. The second was delays by GP level functionaries in entering data (since it required traveling to the block office).

The delays could have had an additional negative consequence if GP functionaries exploited the delays in payment to lend workers money (on work completion) and get reimbursed when the funds arrive. The interest is collected in advance by paying the workers less than what they are due. Repayment is enforced by collecting the worker's bank/postal passbook, and taking the money out of their bank/postal account in their name, using pre-signed withdrawal slips in connivance with bank/postal employees. Panel E of Table ?? suggests that this apparently did not happen: instances of advance payment were frequent (a quarter to a third of payments in the control group), but were not increased by the reform. Using our survey, we also compare household consumption levels in the treatment as compared to control GPs and find no evidence of a long-term cost on treatment households (Appendix Table ??).

To the extent that the increase in payment delays were due to a delay in sending money from the Central Bank of India to the GP account, the decline in CPSMS we observe could have been in part due to those delays. However, if this were the case, we would see an increase in expenditure in treatment GPs after the system was discontinued, which is not the case.

Finally, in Table ?? we examine whether the fund flow reform affected the number of physical assets created. In May 2013, after the end of the intervention, we downloaded the list of MGNREGS projects registered in nrega.nic.in. There were on average 14 projects per GP, most of them ongoing, and the numbers are very similar for treatment and control GPs (Columns 1 and 2). We also sent teams to the villages with a list of 10 projects per GP, sampled from nrega.nic.in. The number of projects found is high (12 per GP, or 86% of registered projects), and similar in the treatment and control GPs for all projects as well as for just the ongoing ones (Columns 3 and 4).

²¹Qualitatively, this is corroborated by the Mukhiyas (GP elected leaders) whom we interviewed. Table ??, Panel E shows that twice as many Mukhiyas either spontaneously offered or agreed with the view that the CPSMS created delays in fund flow, in treatment (34%) than in control blocks (17%)—note that this data needs to be taken with a lot of caution, since it is not clear why Mukhiyas in control GPs would report any delay due to CPSMS!

4.4 Did the reform influence fund leakage and corruption?

The financial data – corroborated by data from the public portal - tell us that there was a 17% decline in MGNREGS spending in the treatment GPs, relative to control, and a 10-13% reduction in the number of workdays and workers hired. In contrast, while the public portal data also shows that the entire decline in spending comes from a decline in number of workers, this is not reflected in the household survey. While the wage data is too noisy to come to definite conclusion, the employment data allows us to reject at a 95% confidence interval a decline in number of workers and workdays similar to what we see in the reported database. We also observe no changes in MGNREGA assets - either in the public portal data or in our asset survey. This is suggestive that the reduction in reported expenses and workdays are accounted for by a reduction in corruption. This hypothesis received some support from the GP report: in our survey of GP elected leaders (Mukhiyas), 47% of Mukhiyas in control GPs thought corruption in the administration was a main issue in MGNREGS implementation. This number was significantly lower, by 12 percentage points, among Mukhiyas in treatment GPs (see Table ??, Panel D). The evidence is, however, indirect and based on a sample survey, not administrative data on the universe of our experiment. In this subsection, we present two direct pieces of evidence on a reduction in corruption.

4.5 Leakages: Direct evidence on ghost workers

Fund leakage could occur in two ways: by reporting “ghost” workers on the database and siphoning off the associated payment (people who are reported to be paid but are non-existent, or exist but have never worked) or by reporting “ghost” days (additional days of reported work by people who actually worked under the scheme but for fewer days than what is reported).

The nature of the fund flow reform suggests that the primary accountability impact should be fewer ghost workers: it is now easier to audit and verify that a particular person exists and has been employed. However, conditional on having worked, accurate, verifiable information on how many days someone worked remains as hard to obtain (since audits rely on recall which tends to be imperfect about things like exact numbers, and villagers can easily be intimidated). Consistent with this hypothesis, Table ?? shows that fewer workers, not fewer days per workers, account for the reduction in reported job days during the reform. And the lingering negative effect on the number of workers even after the intervention ends and spending goes back to the same level in treatment and control GPs could come from the fact that once a ghost is added to the roll, he or she stays on them.

We now turn to directly examining the incidence of ghost workers by comparing the match rate of households listed on job cards with villager names in SECC. For each GP,

we compute the fraction of families where one member has a job card in the MGNREGA database, and families where two or more members have a job card, for which we find a match in the SECC census database. For these two variables we then run a GP-level regression of the form:

$$Y_{vd} = \alpha + \beta T_v + \eta_d + \varepsilon_{vt} \quad (3)$$

and cluster standard errors at the block level. We run this specification separately for three different ways of computing the fraction matched variable: first, the match rates for all job cards in the MGNREGA database (as of 2014), then for all job cards who were recorded as working during the intervention period, and finally for all job cards who were in the database and were recorded as working in the post-reform period.

Table ?? reports the results. In the control group, among single-worker households, we match 64% of the job cards listed in the same village (or somewhere in the GP when individual villages could not be matched).²² We observe a significantly higher – by 1.87 percentage points – match rate in the treatment group (Column 1). Restricting to individuals who are reported as having worked during the reform period, we find a match rate of 67% in the control group which increases significantly by 1.81 percentage points because of the treatment (Column 2). Reassuringly, for individuals reported to have worked *after* the reform period, the treatment-induced increase in match rate is smaller and insignificant (Column 3). Among households with two people or more to match, we find lower match rates (since it is more difficult to match two people), but a similar percentage point increase (1.35 percentage points for the entire database, 1.276 for the working job cards).

The increase in match rate is direct evidence of a reform-induced decline of corruption, although it only accounts for a fraction of the 17% reduction in expenditure, perhaps because this exercise only captures pure ghosts (people who do not exist in the village) not people who exist but are in fact not working for MGNREGS.

A remaining question is: why did local officials not react to the reform by over-reporting more on other margins? Possibly, there is some limit on how much over-reporting of workdays can be done in the name of existing workers, e.g. because of the limited number of infrastructure projects carried out. Note also that ? do not find an increase in ghost workers when ghost days decline, which suggests that these are not perfect substitutes.

²²Our survey data is consistent with this number. While the survey did not track people in the database, we can estimate leakage by applying sampling probability to our household sample to estimate the number of people who worked during the reform period, and dividing that, in each GP, by the number of estimated workers according to the MGNREGS database (as in ?). We find that our household survey only accounts for 59% of the workdays in the database. This is comparable to our match-rate of 64%, especially given that some of the ghost workers exist in the village, but are simply not working for MGNREGS.

4.6 Effect on assets of MGNREGS functionaries

Since corruption declined, do we see any evidence of this missing “missing money” in the pockets of the MGNREGS functionaries? To address this question, we now turn to self-reported affidavit data on personal assets of MGNREGS functionaries. While we recognize the limitations of self-reported asset data, we are reassured by previous research that shows a causal link between politicians getting elected and their self-declared assets (?). Moreover, we expect any treatment-induced bias to be towards zero, especially since we are only using the first two years of the affidavit data, which were used for benchmarking: a heightened fear of scrutiny in the treatment group (due to the extra transparency) should reduce under-reporting by officials in order to avoid being caught under-reporting in the future (most prosecutions for “disproportionate assets,” which are becoming more common over time, are based on rapid accumulation since the benchmark year, which gives an incentive to overstate assets in the baseline year).

GP and block functionaries declared personal assets in 2012-13 (a period spanning our intervention) and 2013-14 (at least six months after the intervention had ended). Figures 7 and 8 show the CDF of reported asset for the Block and GP functionaries, taken together. In the year 2012-2013, the asset declaration data are similar in treatment and control groups. However, in 2013-2014, we observe a leftward shift of the treatment distribution, relative to the control distribution. A Kolmogorov-Smirnov test of stochastic dominance presented in Table ?? allows us to reject equality of the distribution at the 5% level in 2013-2014.

Columns (1) and (2) of Table ?? show, on average, a 10% reduction in $\log(\text{wealth})$ of Block and GP functionaries (Panel A) in 2013-2014, but this impact is statistically insignificant. The graph suggests that this reflects the fact that the wealth distribution is highly skewed with large outliers, and the reform had no impact at the high and low ends of the distribution. If we focus on the median instead, we find a significant decline of 13.7% in median wealth (18.9% with control variables). The wealth reduction for GP and block officials is commensurate to the treatment-induced decline in MGNREGS expenditure. Using the estimate from Column 8, Panel A of Table ??, a 19% decline in the median wealth for MGNREGS employees (630,000 INR) scaled up to the whole treatment sample (651 employees) yields a 78 million INR loss, which is equivalent to about a third of total missing expenditures (224 million INR). Using the mean estimates (Column 4 of Table ??) yields a higher loss, 110 million or 48% of the missing MGNREGS expenditure. Since district level officials also presumably lost money (although we do not have an experiment for them), this order of magnitude seems reasonable.

Panel B and C split the results by administrative tier. There is some indication of larger (proportional) effects for GP functionaries, but the results are too noisy to be conclusive.

5 Conclusion

This paper reports on a large-scale field experiment that evaluated a nine-month reform to the within-state fund-flow system for MGNREGS – India’s federal workfare program. Our evaluation covered a population of 33 million in Bihar, one of India’s poorest states. To identify reform impacts, we combine data from a number of sources: rich administrative program data; a survey that covers 10,000 households and assets built in over 300 villages; a set of names matched across the program database and the Indian Socio-economic and Caste Census; and, finally, affidavit data on wealth of GP officials.

The reform linked fund flow to incurred expenditures and reduced the number of intermediaries involved in fund disbursement. It lowered fund leakages in treatment blocks: MGNREGS expenditures declined by 17 percent with no corresponding change in real outcomes, as measured by surveys. A match of official records of MGNREGS workers with a census collected during the same period further demonstrates a reduction in the number of fake beneficiaries (“ghost workers”).

To the extent that the expenditure reductions reflect lower program leakage, we would expect changes in earnings of officials involved in fund flow for MGNREGS. Theory suggests that the direction of change for GP and block level functionaries is ambiguous and depends on their ability to benefit from the exclusion of district functionaries by increasing their own rents. In practice, we estimate a negative effect on the wealth of block and GP officials: the impact on the mean is noisily estimated, but the impact at the median is a significant drop of 19 percent.

This set of results consistent across a number of different sources suggests that corruption in social programs can be reduced through a program of increased transparency in invoicing, that facilitates future audits and clarifies the lines of responsibilities.

On the flip side, contrary to the hypothesis that the red tape induced by corruption can reduce effectiveness, the reform did not improve the program’s ability to respond to villager needs– neither employment nor wages received by households rose and payment delays increased, at least initially. This may well have been due to short-term issues, which would have been solved over time as implementation became smoother. The technical challenges with managing a computer-based system in areas with frequent electricity shortages and limited IT help should not be underestimated. In addition, lack of coordination between the central monitoring system and the program’s own database meant worker details had to be entered twice. Because of these issues, the reform increased the administrative burden on GP officials in the short-run. As expected, resistance was widespread and frequently given voice: dismay at the personnel costs, frustration with lags in infrastructure roll-out, and quiet hostility to reforms that would reduce rents. State officials in the capital city who heard these complaints, lacked information on whether the observed decline in expenditure reflected lower rent seeking or a genuine

decline in employment provision, with the result that by the end of the fiscal year, they were concerned that the reform may have constrained employment under the program.²³

As our experiment was randomized across blocks in a district, we lack direct evidence on any decline in the wealth of the district officials that the reform may have induced. However, there is considerable anecdotal evidence of district officer displeasure. For example, one of our district monitors reported: *“Initially, the POs [Program Officers, block officials] were apprehensive about the system. The DRDA [District Rural Development Agency] Accountant had scared the POs at the beginning and had convinced them that the system was useless. Whenever POs or Mukhiyas would come, the operators and the accountants would scoff them and tell them that they were stuck with a useless system. They would tell them: “Look, you were better under us. Now, you won’t get any money from the state”.* Another district monitor reported that most officials were hostile to the system and that the DRDA fudged the data on IT equipment to show compliance.

Given the uncertainty on benefits at the time, the district officials were able to effectively lobby the state government to end the intervention, which was rolled back in April 2013. A question we often get is whether MGNREGS employees would have been able to figure out a way to circumvent the system over time, and if corruption would have gone back up. Although managing to force a roll back of the system is an extreme approach to circumventing it, it suggests that district officials could not find another way. In general, the MGNREGS experience seems to be one where a series of steps were undertaken to limit corruption, and where corruption effectively went down over time (nationwide, estimates suggest that leakage was halved between 2007 and 2012).

This rollback, however, was not quite the end of the story. Motivated in part by the results of this experiment, in August 2015, MGNREGS officials put in place a nationwide system that combined direct payment to beneficiary bank accounts (though not always based on a smart card) and expenditure-based transfers. The need for better expenditure management models is not exclusive to MGNREGS: the Government of India spends approximately Rs. 4.6 trillion (\$50 billion) every year on Centrally Sponsored Schemes for which money is released to implementing agencies in lumpy installments. Many of these programs have even worse accountability records than MGNREGS, in large part because of inadequate electronic record keeping. In June 2016, the Ministry of Finance issued orders to extend the use of the Public Finance Management System (the successor of the CPSMS, the platform we gained access to) for all Central Sector Schemes and for central assistance for State Plan Schemes. Their announcement emphasized the system as a means to facilitate “just-in-time” (i.e. expenditure based) release of funds and ensure complete monitoring of funds down to the end user. Thus, overall, the advent of e-governance is heralding very significant reforms for the entire government payment architecture in India.

²³The household survey that demonstrated otherwise was conducted starting May.

Figure 1: MGNREGS Fund-flow in Control Blocks

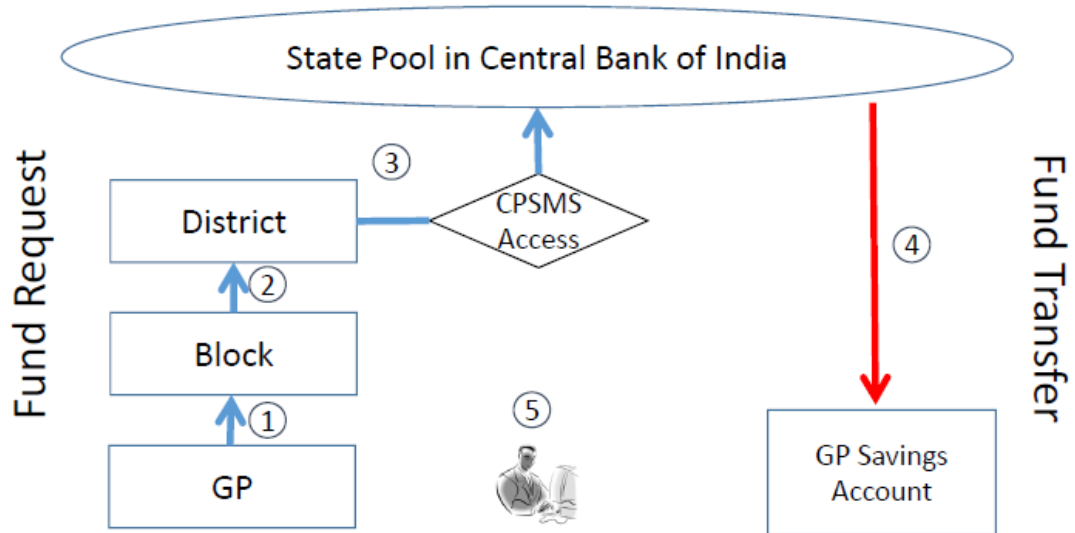


Figure 2: MGNREGS Fund-flow in Treatment Blocks

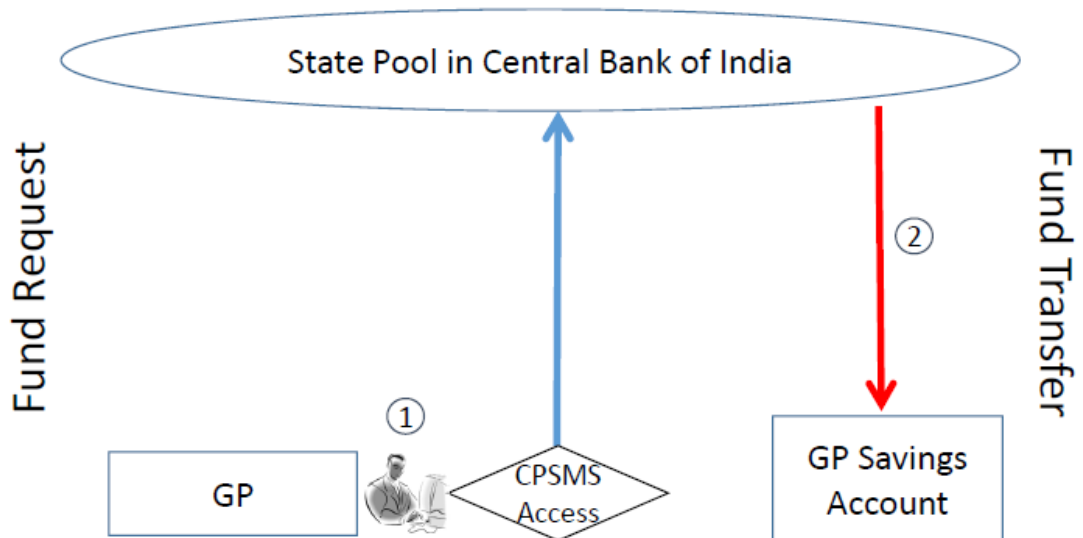
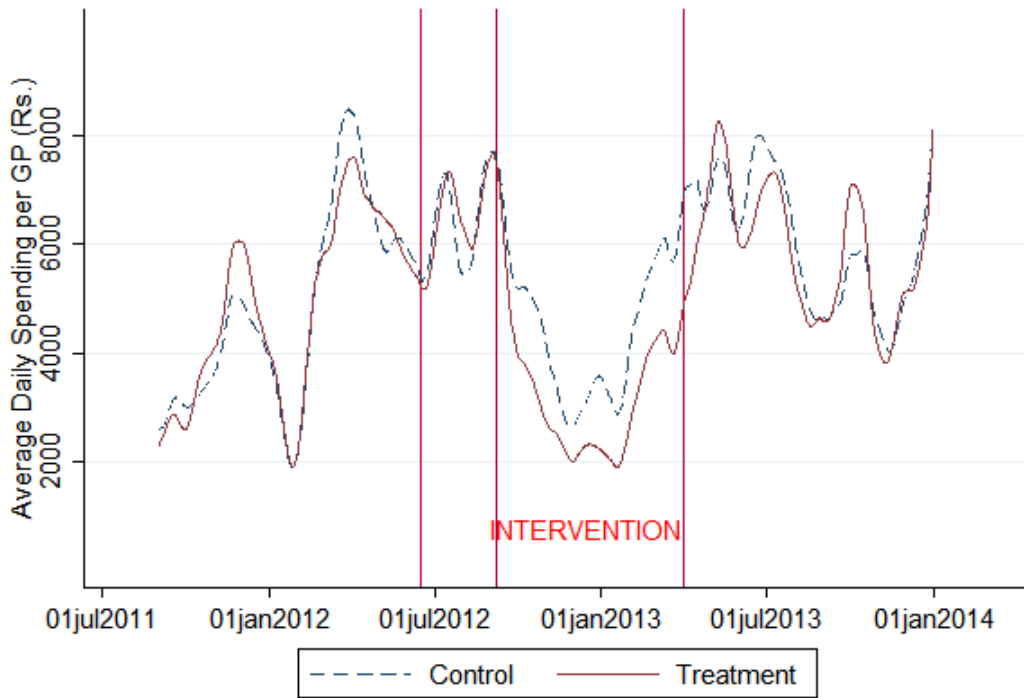


Figure 3: GP daily Expenditures on MGNREGS during the Study Period



Source: CPSMS data on GP savings accounts.

Figure 4: Fraction of Treatment GPs which used CPSMS at least once

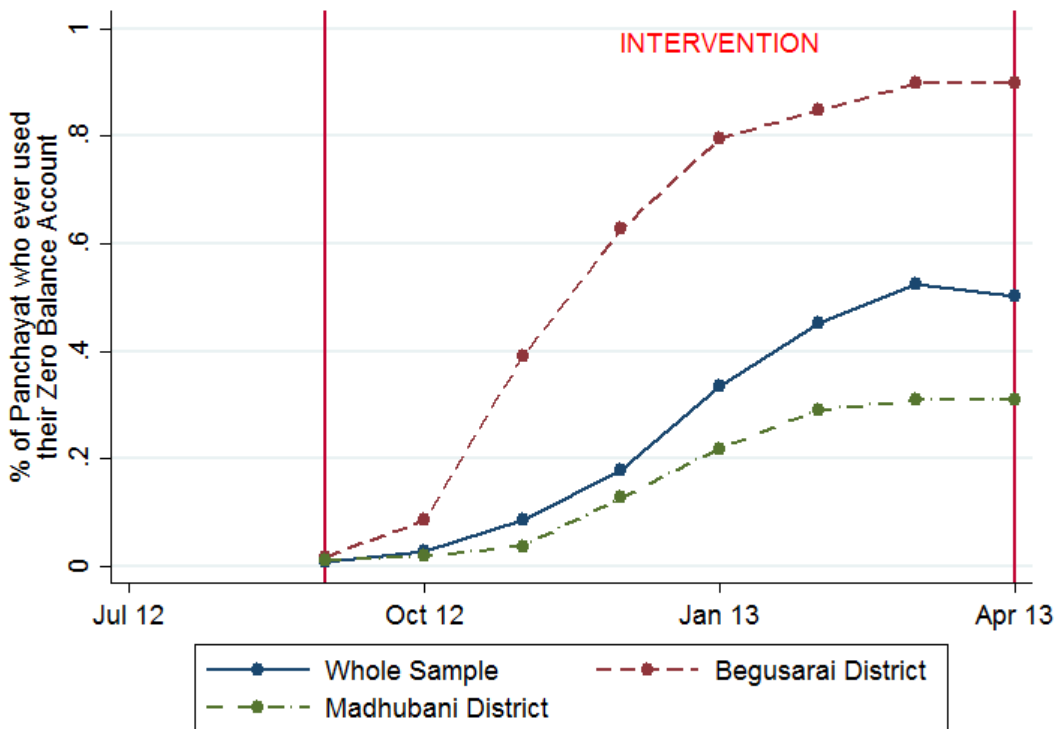
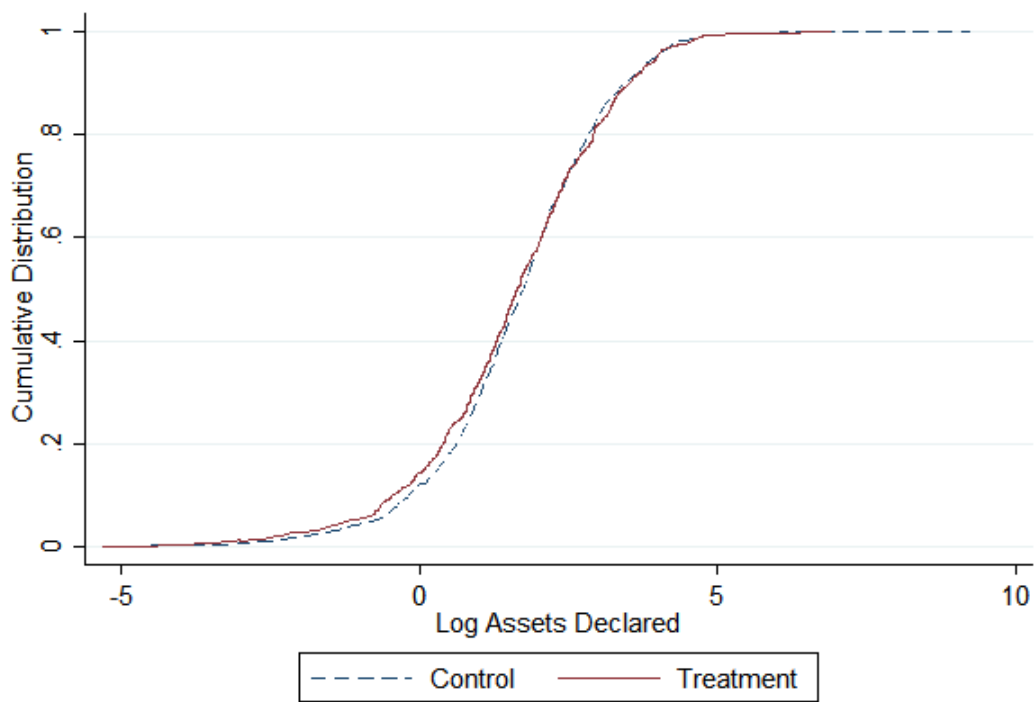
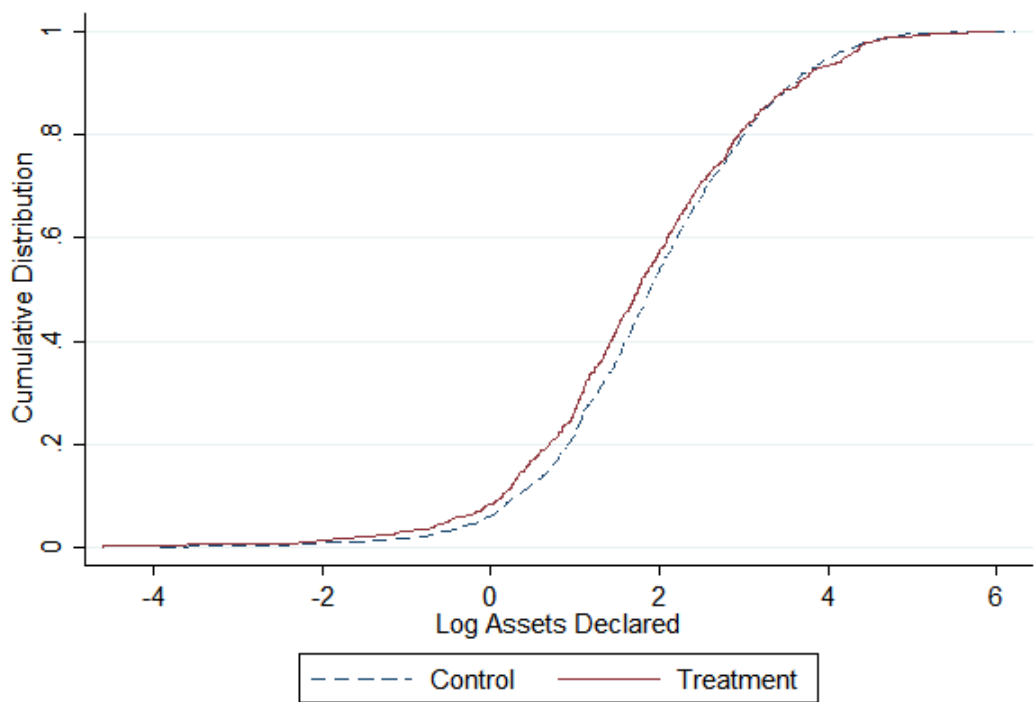


Figure 5: Asset of MGNREGS functionaries: during the intervention



Source: Affidavits of MGNREGS employees 2012-13, Government of Bihar.

Figure 6: Asset of MGNREGS functionaries: after the intervention



Source: Affidavits of MGNREGS employees 2013-14, Government of Bihar.

Table 1: Randomization check

	Control Blocks	Treatment Blocks	Difference	Observations
Panel A: Census 2011				
Area (hectares)	1101	1129	28.38	2,936
Number of households	1860	1845	-15.22	2,936
% SC Population	0.196	0.194	-0.00164	2,936
% ST Population	0.0112	0.0144	0.00320	2,936
Literacy Rate	0.64	0.639	-0.000859	2,936
% With education facility	0.992	0.997	0.00529*	2,936
% With medical facility	0.668	0.679	0.0114	2,936
% With post office	0.0394	0.0357	-0.00367	2,936
% With bank branch	0.352	0.402	0.0496**	2,936
% With electricity supply	0.426	0.46	0.0344	2,936
% Land Irrigated	0.53	0.523	-0.00639	2,936
Panel B: Household Survey				
% Hindu	0.92	0.89	-0.0268**	390
% Scheduled Castes	0.26	0.24	-0.0188	390
% Other Backward Castes	0.59	0.60	0.0162	390
% House without a solid roof	0.38	0.41	0.0246	390
% Owns Land	0.58	0.57	-0.0139	390
% Male Head	0.78	0.76	-0.0129	390
% Literate Head	0.56	0.55	-0.00884	390
Household Size	6.52	6.44	-0.0836	390
Number of adults in the household	3.42	3.36	-0.0664	390
Panel C: nrega.nic.in reports (April 2011- March 2012)				
MGNREGS beneficiary households	187	196	9.283	2,950
MGNREGS work days provided	6290	6673	383.7	2,950
MGNREGS labor expenditures (lakhs)	7.69	8.68	0.996**	2,950
MGNREGS material expenditures (lakhs)	6.57	7.07	0.508	2,950
Panel D: CPSMS reports (Sept 2011- March 2012)				
MGNREGS funds spent (CPSMS)	9.00	8.73	-0.272	3,025
MGNREGS funds received (CPSMS)	9.52	9.59	0.0645	3,025

Note: The unit of observation is a Gram Panchayat (GP). Out of 3067 GP from our sample list, we match 2936 GP with census 2011 data (Panel A), we surveyed 390 GP (Panel B), we match 2950 GP with nrega.nic.in data (Panel C) and 3025 GP with CPSMS data (Panel D). The difference between control and treatment blocks is estimated using a regression of each GP characteristic on a dummy equal to one for treatment blocks and district fixed effects. Standard errors are clustered to take into account correlation at the block level. Stars denote significance levels. *, ** and *** denote significant differences at the 10%, 5% and 1% levels respectively.

Table 2: Impact of the reform on MGNREGS Expenditure: Evidence from CPSMS data

	Before Sept 2011 - June 2012	Set up July- August 2012	Intervention Period			After Apr 2013 - Jan 2014
			Sept-Dec 2012	Jan - Mar 2013	Whole Period	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Total Debit from GP Accounts						
Treatment	-0.502 (0.729)	0.0472 (0.291)	-1.039*** (0.315)	-1.267*** (0.280)	-2.259*** (0.759)	-0.345 (0.895)
Observations	3,025	3,025	3,025	3,025	3,025	3,025
Mean in Control	14.37	4.122	5.394	4.146	13.66	16.03
Panel B: Closing Balance in GP Accounts						
Treatment	-0.0843 (0.245)	0.191 (0.220)	-1.007*** (0.240)	-1.277*** (0.244)	-1.277*** (0.244)	-0.117 (0.235)
Observations	3,025	3,025	3,025	3,025	3,025	3,025
Mean in Control	4.147	4.407	4.099	4.274	4.274	4.236
Panel C: Total Credit to GP Accounts						
Treatment	-0.179 (0.830)	0.251 (0.338)	-2.192*** (0.367)	-1.249*** (0.335)	-3.190*** (0.781)	0.896 (0.883)
Observations	3,025	3,025	3,025	3,025	3,025	3,025
Mean in Control	15.27	4.282	5.146	4.006	13.43	15.97

Note: The unit of observation is a Gram Panchayat (GP). In Panel A the dependent variable is the sum of debits from the savings account of each GP for each period (in lakhs Rupees). In Panel B the dependent variable is the closing balance on the savings account of each GP at the end of each period (in lakhs Rupees). In Panel C the dependent variable is the sum of credits made to the savings account of each Panchayat for each period (in lakhs Rupees). Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects. Standard errors are clustered at the block level.

Table 3: Impact of the reform on MGNREGS Expenditure: Evidence from nrega.nic.in

	Pre- intervention Apr 2011-Mar 2012	Set up and intervention Apr 2012-Mar 2013	Post- intervention Apr 2013-Mar 2014
	(1)	(2)	(3)
Panel A: GP Expenditures on labor from nrega.nic.in			
Treatment	0.996** (0.495)	-2.270*** (0.760)	-0.271 (0.729)
Observations	2,950	2,947	2,954
Mean in Control	7.551	13.83	13.66
Panel B: GP Expenditures on material from nrega.nic.in			
Treatment	0.508 (0.432)	-1.077** (0.526)	0.315 (0.534)
Observations	2,950	2,947	2,954
Mean in Control	6.504	7.717	8.377

Note: The unit of observation is a Gram Panchayat (GP) The dependent variables are expenditures from MIS reports for financial years 2011-12, 2012-13, 2013-14 (in lakhs Rupees). Data was downloaded from the MGNREGS website (nrega.nic.in) in November 2014. The intervention started in September 2012 and ended on March 31st, 2013. Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects. Standard errors are clustered at the block level.

Table 4: Impact of the reform on MGNREGS Employment: Evidence from official reports (nrega.nic.in)

	Pre intervention	Set up	Intervention Period			Post intervention
	April 2011 - June	July-August	Sept-Dec	Jan - Mar	Whole	Apr 2013 - March
	2012	2012	2012	2013	Period	2014
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Days worked (nrega.nic.in)						
Treatment	91.88 (530.3)	-130.3 (111.5)	-404.6* (227.6)	-267.8 (163.3)	-672.4* (363.6)	-859.5 (542.7)
Observations	2,959	2,959	2,959	2,959	2,959	2,959
Mean in Control	10313	1058	2759	2269	5028	10603
Panel B: Days per working household (nrega.nic.in)						
Treatment	-0.0269 (1.010)	-0.712 (0.605)	-0.286 (0.805)	0.187 (0.701)	-0.00410 (0.930)	-0.308 (0.838)
Observations	2,952	2,514	2,728	2,717	2,868	2,945
Mean in Control	36.85	17.35	29.14	25.14	33.65	39.54
Panel C: Number of working households (nrega.nic.in)						
Treatment	2.988 (12.49)	-3.132 (5.151)	-10.02 (6.233)	-8.342 (5.700)	-13.60* (8.150)	-15.03 (10.33)
Observations	2,959	2,959	2,959	2,959	2,959	2,959
Mean in Control	273.6	59.92	91.68	90.37	140.2	257.2

Note: The unit of observation is a Gram Panchayat (GP). In Panel A the dependent variable is the total number of days provided. In panel B the dependent variable is the total number of days provided to households reported to have worked. In panel C the dependent variable is the number of households reported to have worked. In panel D the dependent variable is the number of days worked by households who could not be matched with survey households. In Panel E the dependent variable is the number of days worked by households matched with survey households. The data was extracted from Job card information on the nrega.nic.in server. It covers the period from July 2011 to Sept 2013. Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects.

Table 5: Impact of the reform on MGNREGS Employment: Evidence from household survey

	Set up	Intervention Period			Post-Intervention
	Jul - Aug 2012 (1)	Sept - Dec 2012 (2)	Jan - Mar 2013 (3)	Whole Period (4)	Apr - Jun 2013 (5)
Panel A: MGNREGS Participation					
Treatment	-0.00863*** (0.00282)	0.00306 (0.00321)	0.00379 (0.00331)	0.00699 (0.00445)	0.00101 (0.00474)
Observations	9,436	9,436	9,436	9,436	9,436
Mean in Control	0.0122	0.0168	0.0226	0.0378	0.0390
Panel B: Number of days worked					
Treatment	-0.163*** (0.0560)	0.0639 (0.114)	0.246 (0.165)	0.310 (0.207)	0.154 (0.477)
Observations	9,436	9,436	9,436	9,436	9,436
Mean in Control	0.231	0.470	0.688	1.158	1.821
Panel C: Wages received for MGNREGS employment					
Treatment	-17.95** (7.073)	-1.955 (12.95)	13.91 (13.33)	11.96 (20.01)	-19.75 (24.64)
Observations	9,436	9,436	9,436	9,436	9,436
Mean in Control	24.19	43.29	55.89	99.18	104.7
Panel D: Average delays in payment (days)					
Treatment	-32.96 (25.39)	53.04*** (18.68)	25.66** (10.28)	36.48*** (11.18)	2.344 (9.395)
Observations	112	154	214	355	361
Mean in Control	73.49	71.03	51.14	60.15	38.22
Panel E: Illegal advance payments					
Treatment	0.0512 (0.149)	-0.0656 (0.0831)	0.0876 (0.0812)	-0.0197 (0.0621)	0.0100 (0.0621)
Observations	96	128	170	289	234
Mean in Control	0.376	0.281	0.299	0.295	0.386

Note: The unit of observation is a household. In Panel A the dependent variable is a dummy variable which is equal to one if any household member participated to MGNREGS. In Panel B the dependent variable is the total number of weeks worked by household members under MGNREGS. In Panel C the dependent variable is total wage payments received by each household for MGNREGS employment. In Panel D the dependent variable is the average number of days between the time of work spells and the time of each payment. When payments have not been made at the time of the survey, the delay is set equal to the time between the work spell and the survey date. It is missing for households who did not work for MGNREGS. In Panel E the dependent variable is a binary variable which is equal to one if any household member has received a payment for MGNREGS work in cash within 15 days of the work spell. It is missing when no MGNREGS payment has been made at the time of the survey. The data was collected by a representative survey of 10,036 households in May-July 2013. Households were asked about work spells from July 2012 to the time of the survey. Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects and household controls. Household controls include sets of dummies for religion, caste, type of housing, land ownership, gender and literacy of the household head, household size and number of adults.

Table 6: Impact of the reform on MGNREGS projects: Evidence from asset survey

	Number Registered		Number found	
	All Projects (1)	Ongoing (2)	All Projects (3)	Ongoing (4)
Treatment	0.0494 (0.263)	-0.210 (0.413)	0.309 (0.239)	0.0271 (0.267)
Observations	390	390	385	385
Mean in Control	13.80	11.69	11.79	9.819

Note: the unit of observation is a Gram Panchayat (GP). The dependent variables are the number of projects registered in the public data portal (nrega.nic.in) on May 15, 2013 (1), the number of projects declared as ongoing in nrega.nic.in (2), the number of registered (3) and ongoing (4) projects found by surveyors in June-July 2013. Out of 5390 projects registered in nrega.nic.in for the 390 GP of the survey sample, a random sample of 3900 projects were surveyed (10 per GP). The number of projects found in the survey is scaled up using the number of registered projects divided by the number of sampled projects rate. 5 GP (28 projects) could not be surveyed. All specifications include district fixed effects.

Table 7: Impact of the reform on fake beneficiaries: Evidence from matching of nrega.nic.in job cards with SECC census

	All job cards (as of April 2014) (1)	Job cards with at least one working member	
		Intervention period July 2012-March 2013 (2)	Post intervention Apr 2013 - March 2014 (3)
		Panel A: Match Rate for job cards with one member only	
Treatment	0.0187** (0.00741)	0.0181** (0.00766)	0.0107 (0.00696)
Observations	3,095	2,868	2,922
Mean in Control	0.644	0.673	0.698
Panel B: Match Rate for job cards with two members or more			
Treatment	0.0135** (0.00613)	0.0126 (0.00764)	0.0104 (0.00732)
Observations	3,093	2,836	2,906
Mean in Control	0.243	0.282	0.286

Note: The unit of observation is a GP. The dependent variable is the fraction of job cards from nrega.nic.in matched by name with households from the SECC census. A job card with two members or more is matched when at least to members have been matched by name with a census household. The nrega.nic.in data was extracted from the nrega.nic.in server, it covers the period from July 2011 to March 2014. Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects.

Table 8: Impact of the reform on assets of MGNREGS functionaries: Evidence from affidavit data

	Average Effect (OLS)				Effect at the Median (Quantile Regression)			
	2012-13		2013-14		2012-13		2013-14	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: GP and Block								
Treatment	-0.0754 (0.130)	-0.0659 (0.128)	-0.102 (0.103)	-0.114 (0.102)	-0.117 (0.073)	0.005 (0.068)	-0.137* (0.074)	-0.189*** (0.069)
Observations	2,455	2,455	1,737	1,737	2,455	2,455	1,737	1,737
Kolmogorov Smirnov p-value (for stochastic dominance)					0.105		0.057	
Panel B: GP								
Treatment	-0.0611 (0.144)	-0.0456 (0.143)	-0.102 (0.124)	-0.114 (0.121)	-0.004 (0.08)	0.038 (0.079)	-0.199** (0.081)	-0.19** (0.083)
Observations	1,698	1,698	1,251	1,251	1,698	1,698	1,251	1,251
Kolmogorov Smirnov p-value (for stochastic dominance)					0.344		0.135	
Panel C: Block								
Treatment	-0.128 (0.143)	-0.0832 (0.136)	-0.159 (0.130)	-0.154 (0.126)	-0.185 (0.119)	-0.07 (0.113)	-0.059 (0.132)	-0.115 (0.126)
Observations	757	757	486	486	757	757	486	486
Kolmogorov Smirnov p-value (for stochastic dominance)					0.212		0.29	

Declarations 2012-13 were made from August 2012 to June 2013. Declarations 2013-14 were made from July 2013 to September 2014. The intervention period was September 2012 to April 2013. GP level functionaries are Panchayat Rozgar Sewak. Block level functionaries Program Officers, Accountants, Computer Operators, Junior Engineers, Program Technical Assistants, and Executive Assistants. Functionary Controls include the age, the square of age, and dummies for gender and designation of the functionaries as well as a dummy for whether the functionary is posted in the district she was born in. In Panel A, B and C II specifications include district fixed effects. Standard errors are clustered at the block level.

APPENDIX: FOR ONLINE PUBLICATION ONLY

Data Appendix

In this appendix, we describe the different sources of information we use in the analysis. We first present the official data on expenditures and employment, then turn to the surveys we implemented to assess actual MGNREGS implementation, and finally describe three additional sources we use to measure corruption.

We use two sources of official reports on MGNREGS expenditures and employment: CPSMS and nrega.nic.in.

CPSMS: In July 2014, we were granted access to detailed information MGNREGS expenditures via the Central Planning Scheme Monitoring (CPSMS) Portal. Both treatment and control GPs were monitored in the system from July 2011 onward, and we could observe all credit and debit transactions from GP savings account. We use this information to compute MGNREGS spending per GP for the different periods of interests: from July 2011 to the start of the intervention in September 2012, from September 2012 to December 2012, from January 2013 to March 2013 and from the end of the intervention in April 2013 until July 2014.

NREGA.NIC.IN: The government website nrega.nic.in provides publicly available information on MGNREGS expenditures per GP for every financial year (a financial year start on April 1st). In July 2014, using a newly available facility called the Public Data Portal (jointly produced by the Ministry of Rural Development and Evidence for Policy Design) we downloaded data on GP spending on labor and material for the financial years 2011-12, 2012-13 and 2013-14.

Labor expenditures figures in nrega.nic.in are aggregates of work and payment details of MGNREGS workers which are also entered on the website and made publicly available in the form of job cards. This online jobcard mimics the physical job card delivered to all households who register for MGNREGS work: the rule of one job card per household is not always followed in practice, so that members of a given households may appear on different job cards. We requested access to job card information from the Ministry of Rural Development and were provided with the details of 4,197,904 job cards and 6,292,307 workers in our sample districts for the financial years 2011-12, 2012-13 and 2013-14.

In order to provide independent measures of MGNREGS implementation, we carried out our own survey in the 12 sample districts between May and July 2013. Within each district, we visited every block – in total, we had 69 treatment blocks and 126 control blocks, 195 blocks in total. We surveyed 2 randomly sampled GPs in each block – this gave us a total of 390 GPs. The survey consisted of three main surveys: a household

survey, a survey of MGNREGS assets and a survey of GP head (or Mukhiya).

Household Survey: We conducted a household survey covering 10,036 households. In each GP, we covered at least 25 households. These households were sampled from the list of households obtained from the District Rural Development Authority (DRDA). These lists were initially compiled in 2002 for the purpose of identifying BPL households, so each household was given a poverty score, based on various criteria. From these lists, we sampled 72 per cent of households below the median poverty score and 28 per cent households from above the score. In the case a sampled household had left the village or all its members were defunct, surveyors were asked to interview a replacement household who had been randomly chosen from the initial list. Because the sampling lists were 10 years old and many areas had high migration rates, the proportion of households interviewed as replacements was also high, about 30%.

Asset Survey: We sampled 10 infrastructure projects from each GP. These were randomly sampled from the MIS (www.nrega.nic.in). In total, we sampled a total of 4165 infrastructure projects.

Mukhiya Survey: We attempted to interview the Mukhiya of every single GP we visited. We managed to locate and interview a total of 358 Mukhiyas. Unlike the other two surveys, the Mukhiya survey was conducted on paper and was both quantitative and qualitative in nature.

We use three additional sources of administrative data to provide evidence on corruption in MGNREGS implementation: the Socio-Economic Caste Census, affidavit data and audits data.

SECC and name matching: In order to measure the extent of possible “ghost workers,” we attempt to determine for each working household reported on an nrega.nic.in job card whether or not there is a matched household within the SECC data. The 2011 Socio-Economic Caste Census (SECC) is a national survey of all persons and households in rural and urban India. It is based on the National Population Register from the 2011 Population Census, but was conducted mostly in 2012 due to various implementation issues. The SECC data includes the name, father’s name (or husband’s name for married women), gender, education, and other information for each member of the household and the household overall. In the 12 districts of our sample (inclusive of rural villages only), the SECC data covers 16,480 villages, five million households, and 34 million individuals. The job cards data covers 18,513 villages, 4,197,904 working households, and 6,292,307 working household members.

We proceed in two steps: In the first step, we pair villages in the job cards with corresponding villages in the SECC data to impose the restriction that we search for matching households only within the same village. In the second step, we match households from the job cards data to the SECC data within village pairs based on similarity of name,

gender, and household composition. To calculate the closeness of village names in the first step and individuals' names in the second step, we use a modified levenshtein algorithm (Paul Novosad's lev.py downloaded from <http://www.dartmouth.edu/~novosad/code.html>) as the building block on top of which we add additional alterations that take into consideration alternative spellings, missing/additional portions of names, and abbreviations to quantify the closeness of reported names.

In the first step, we take the following approach to determine village pairs. While the job cards data contains information on block, GP, and village name, the SECC data contains corresponding information for block and village name only. We attempt to match by name each of the 18,513 unique villages in the job cards data within block with a corresponding SECC village. We are able to match 84% of the job cards villages (containing 88% of households). For 16% the job card villages (12% of households), we match them to all SECC villages which are matched with job card villages belonging to the same GP. For about 0.5% of villages (0.7% of households), we are unable to do either and match them with all the villages in the block.

In the second step, we attempt to find a match for each of the job cards from within the paired village or list of villages. We declare a household with one working member listed on the job card as matched if a single matching individual in the SECC data is found, and we declare a household with two or more members listed on the job card as matched if at least two individuals within the same SECC household are matched. The matching rate is thus mechanically lower for household with two working members (37% of households, of which 25% are matched) than for households with one working member (63% of households, of which 64% are matched). Individuals are matched based on two primary criteria: gender, which must match exactly, and name, which must be sufficiently close based on the algorithm described above. Note that once a suitable household match is found according to this process for one or more members, all other members of the job cards household are declared as coming from a matched household. In contrast, the matched SECC household is not removed from the pool of potential matches as the algorithm moves on.

Our outcome of interest is the match rate, separately for people reported to have worked during the period of the intervention and people reported to have worked after the intervention: the idea is that a name or household who is supposed to have a job-card in the MGNREGS data but is not found the SECC database is more likely to be a "ghost" than those who are found in both. This exercise is therefore a population-level version of the forensic method pioneered by ?, using exclusively administrative data. We recognize that the data bases are both imperfect. There are surely errors in both directions (individuals might be omitted from the SECC census for example, or the matching could have failed because the names are spelled too differently to match, or someone could be matched to someone else with the same name), but there is no reason

why these errors would be different in treatment and control groups.

Affidavit data: We also collected affidavits of MGNREGS employees. In the financial years 2012-13 and 2013-14, the Ministry of Rural Development of Bihar made it mandatory for all its employees to declare their personal assets, including cash, movable and immovable assets owned by them or a member of their household. The affidavits were scanned and the pdf files were made available online on the website of each district. Compliance was not perfect, in total we collected 2,463 affidavits for the financial year 2012-13 and 1,741 for the financial year 2013-14. Our measure of MGNREGS employees' personal wealth is constructed by adding the value of movable (cash, bank deposits, bonds, jewellery, other financial assets, vehicles) and immovable assets (land, buildings, other immovables) of the employee and his or her spouse. When the value of the jewellery is missing but the weight of gold or silver owned is given, we impute the value using international prices from <http://www.bullion-rates.com>.

Audits data: Finally, we use reports on MGNREGS audits carried out by the administration of each district between May 2012 and June 2013. These reports were compiled in July 2013 by the Rural Development Department to inform the process evaluation of MGNREGS audits by ?. The data include the date of each audit, the name of the block and GP, the number of MGNREGS projects audited and the number of irregularities found. We aggregate this information and compute the number of audits, the number of projects audited, the number of irregularities found and the number of irregularities per project audited in each block for three periods: May to August 2012 (pre-intervention), September 2012 to March 2013 (intervention period) and April to June 2013 (post-intervention). The completion date of each project audited is not recorded, but the Rural Development Department letter no.120078 (September 1st, 2012) instructs audit teams to select projects undertaken in the financial years 2011-12 and 2012-13. Since the financial year 2012-13 ended in March 2013, projects audited in April to June 2013 had been undertaken during the intervention period.

Figure A.1: Map of Sample Districts

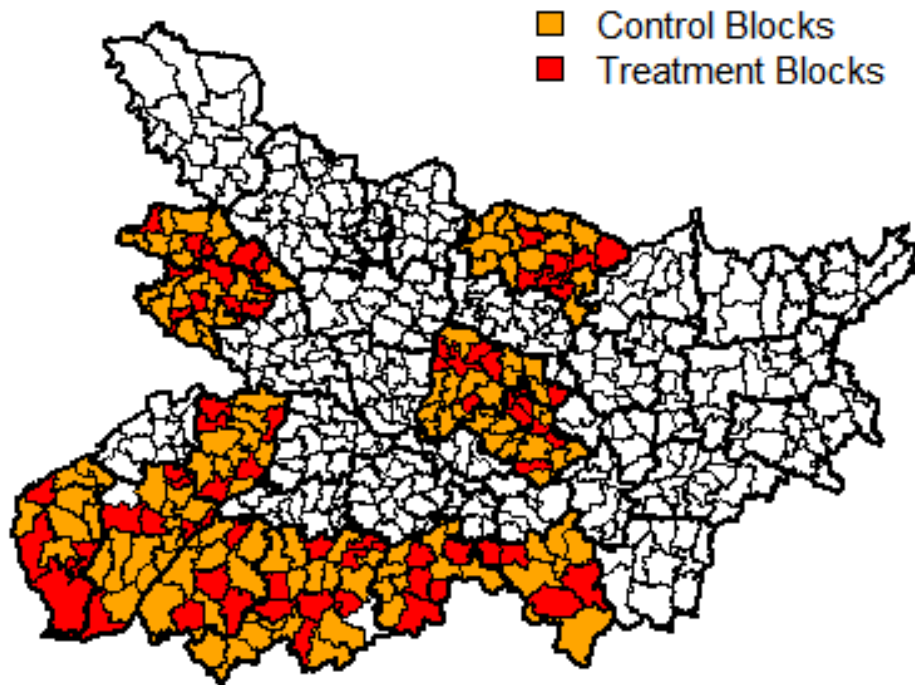


Table A.1: Infrastructure availability

<i>Infrastructure</i>	July '12		Jan '13	Apr '13		Required
	T	C	T	T	C	
Computers (Number)	1.32	1.06	2.48	2.06	1.61	3
Operators (number)	1.22	0.86	2.20	1.75	1.27	3
Generator (1=Yes 0=No)	0.67	0.56	0.97	0.90	0.85	1
Internet (1=Yes 0=No)	0.38	0.33	0.85	0.71	0.60	1
Scanner (1=Yes 0=No)	0.57	0.37	0.73	0.81	0.65	1
Printer (1=Yes 0=No)	0.59	0.43	0.71	0.83	0.76	1
Sampled Blocks	69	126	66	69	123	

Source: Phone surveys of Block Level MGNREGS functionaries (Program officers). The intervention started in September 2012 and ended in April 2013. "T" denotes treatment blocks and "C" denotes control blocks.

Table A.2: MGNREGS Spending levels from different data sources

Panel A		Control	Treatment	Difference	Pvalue
Debit in CPSMS					
	2012-13	19.27	16.84	-2.43	0.11
	2013-14	16.99	16.32	-0.67	0.65
Total Expenditures in MIS					
	2012-13	21.66	18.27	-3.38	0.05
	2013-14	21.48	21.27	-0.21	0.90
Difference CPSMS-MIS					
	2012-13	-2.39	-1.44	0.95	0.15
	2013-14	-4.49	-4.95	-0.46	0.63
Panel B		Control	Treatment	Difference	Pvalue
Payments in Job Cards					
	2011-12	8.30	9.26	0.96	0.24
	2012-13	15.74	14.25	-1.49	0.29
	2013-14	16.27	14.61	-1.66	0.26
Labor Expenditures in MIS					
	2011-12	7.59	9.04	1.45	0.08
	2012-13	13.91	11.66	-2.26	0.06
	2013-14	13.23	12.83	-0.41	0.71
Difference Job Cards-MIS					
	2011-12	0.71	0.22	-0.49	0.21
	2012-13	1.82	2.59	0.77	0.03
	2013-14	3.03	1.78	-1.25	0.02

Source: CPSMS Credit Debit Data, MIS Financial Reports (nrega.nic.in), Job Cards (nrega.nic.in). All amounts are annual panchayat averages in lakhs. CPSMS data is not available for the whole financial year 2011-12. p-values take into account correlation of errors at the block level. Years are financial years (Apr 1st-Mar 31st).

Table A.3: Treatment Effect on MGNREGS audits

	Before Jan 2011 - Aug 2012 (1)	Intervention Period Sep 2012 - Mar 2013 (2)	Post- Intervention Apr - Jun 2013 (5)
Panel A: Number of Audits			
Treatment	0.148 (0.631)	-0.0751 (0.482)	-0.109 (0.240)
Observations	195	195	195
Mean in Control	2.627	6.548	2.167
Panel B: Number of Works Audited			
Treatment	1.268 (4.877)	-0.378 (2.944)	0.580 (1.097)
Observations	195	195	195
Mean in Control	18.13	33.12	9.556
Panel C: Number of Works where irregularities were found			
Treatment	-0.930 (1.777)	-0.0637 (0.816)	0.290 (0.194)
Observations	195	195	195
Mean in Control	4.460	3.222	0.476
Panel D: Share of Works where irregularities were found			
Treatment	-0.0577 (0.0491)	0.00848 (0.0194)	0.0500** (0.0252)
Observations	119	188	148
Mean in Control	0.215	0.0884	0.0514

Source: Rural Development Department, Government of Bihar. The unit of observation is a block. The dependant variables are the number of audits in each period (Panel A), the number of works audite (Panel B) the number of works were irregularities were found (Panel C), and the share of works where irregularities were found (Panel D). Each column present results from a separate regression using data for a different time period. There are missing observations in Panel D for blocks which had no works audited in a given period. Standard errors are clustered at the block level.

Table A.4: Treatment Effect on household MGNREGS participation (household survey)

	Household Participation in MGNREGS	
	Anytime before (1)	Since July 2012 (2)
Treatment	-0.0161 (0.0136)	0.000842 (0.00861)
Observations	10,018	10,007
Mean in Control	0.288	0.0936
Effect as % of Control Mean	-5.608	0.899

Note: The unit of observation is a household. In Column one the outcome is a binary variable equal to one if any member of the household worked for MGNREGS in the past. In Column Two the outcome is a binary variable equal to one if any member of the household did MGNREGS worked since July 2012. The data was collected by a representative survey of 10,036 households in May-July 2013. Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects and household controls. Household controls include sets of dummies for religion, caste, type of housing, land ownership, gender and literacy of the household head, household size and number of adults.

Table A.5: Treatment effect on household consumption (household survey)

	All (1)	Log Monthly Consumption		
		Frequent expenditures (2)	Recurrent expenditures (3)	Rare expenditures (4)
Treatment	-0.00764 (0.0212)	-0.00788 (0.0163)	-0.0400 (0.0261)	0.00104 (0.0393)
Observations	10,033	10,032	10,016	10,009

Note: The dependent variable are the log of household monthly expenditures for different categories of expenditures. Frequent expenditures are expenditures reported every week. Recurrent expenditures are reported every month. Rare expenditures are reported over the past five months. The data was collected by a representative survey of 10,036 households in May-July 2013. Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects and household controls. Household controls include sets of dummies for religion, caste, type of housing, land ownership, gender and literacy of the household head, household size and number of adults.

Table A.6: Impact of the reform on MGNREGS implementation issues: Evidence from GP head (Mukhiya) survey

Panel A: Lack of demand for MGNREGS work	
Treatment	0.0228 (0.0545)
Observations	346
Mean in Control	0.379
Panel B: Mandated price of material lower than market price	
Treatment	0.0279 (0.0386)
Observations	346
Mean in Control	0.833
Panel C: Lack of funds from the government	
Treatment	-0.000833 (0.0498)
Observations	346
Mean in Control	0.718
Panel D: Corruption in the administration	
Treatment	-0.121** (0.0572)
Observations	346
Mean in Control	0.471
Panel E: CPSMS fund-flow creates delays	
Treatment	0.185*** (0.0539)
Observations	346
Mean in Control	0.167

Note: The unit of observation is a Mukhiya (head of GP). The dependent variables are the fractions of Mukhiya who declared that the lack of demand for MGNREGS work (Panel A), the mandated price of material lower than the market price (Panel B), the lack of funds from the government (panel C) corruption in the administration (panel D) and delays in fund-flow created by CPSMS (panel E) were important issues in MGNREGS implementation. The data was collected from a representative sample of 354 Mukhiya from treatment and control blocks in May-July 2013. Treatment is a dummy which is equal to one for the blocks selected for the intervention. All specifications include district fixed effects and Mukhiya controls. Mukhiya controls include sets of dummies for Mukhiya's Religion, caste, gender, education, age, whether any member of the family was elected Mukhiya in 2001 and 2006.

Table A.7: OLS and IV estimates of the main results

	OLS (1)	IV (2)	OLS (3)	IV (4)
Table 2				
Panel A: Total Credit to GP Accounts				
Treatment/use system	-3.194*** (0.790)	-4.871*** (1.261)		
Observations	2,920	2,920		
Mean in Control	13.33	13.33		
Panel B: Average Balance in GP Accounts				
Treatment/use system	-0.242* (0.129)	-0.369* (0.195)		
Observations	2,920	2,920		
Mean in Control	1.999	1.999		
Panel C: Total Debit from GP Accounts				
Treatment/use system	-2.236*** (0.771)	-3.411*** (1.208)		
Observations	2,920	2,920		
Mean in Control	13.57	13.57		
Table 3				
Panel A: GP Expenditures on labor from nrega.nic.in				
Treatment/use system	-2.270*** (0.760)	-3.562*** (1.197)		
Observations	2,947	2,919		
Mean in Control	13.83	13.90		
Panel B: GP Expenditures on material from nrega.nic.in				
Treatment/use system	-1.077** (0.526)	-1.684** (0.806)		
Observations	2,947	2,919		
Mean in Control	7.717	7.737		
Table 4				
Panel A: Days worked (nrega.nic.in)				
Treatment	-827.7* (448.5)	-1,296* (699.9)		
Observations	2,941	2,918		
Mean in Control	6100	6115		
Panel B: Days per working household (nrega.nic.in)				
Treatment	-0.0101 (0.0785)	-0.0202 (0.119)		
Observations	2,887	2,868		
Mean in Control	6.240	6.241		
Panel C: Number of working households (nrega.nic.in)				
Treatment	-138.0* (72.51)	-215.7* (113.3)		
Observations	2,941	2,918		
Mean in Control	992.6	994.8		
Table 5				
Panel A: MGNREGS Participation				
Treatment/use system			-0.00273 (0.00518)	0.00286 (0.00729)
Observations			9,841	9,764
Mean in Control			0.0387	0.03
Panel B: Wages received for MGNREGS employment				
Treatment/use system			-14.75 (21.08)	-23.67 (30.07)
Observations			9,841	9,764
Mean in Control			127.8	128.1
Panel C: Average delays in payment (days)				
Treatment/use system			22.98** (11.56)	32.12* (16.81)
Observations			463	459
Mean in Control			63.15	63.31
Table 6 Fraction of Assets found				
Treatment/use system			0.0176 (0.0176)	0.0254 (0.0244)
Observations			4,165	4,135
Mean in Control			0.855	0.854
Table 7				
Panel A: Match rate of job cards with one person				
Treatment/use system			0.0187** (0.00741)	0.0246** (0.0110)
Observations			3,095	2,897
Mean in Control			0.252	0.262
Panel B: Match rate of job cards with two or more				
Treatment/use system			0.0135** (0.00613)	0.0161* (0.00944)
Observations			3,093	2,897
Mean in Control			0.063	0.064

Note: Column 1 presents the treatment effect for the whole set-up and intervention period estimated with OLS. Column 2 presents the treatment effect for the whole set-up and intervention period estimated using treatment as an instrument for the use of CPSMS system. The panels correspond to the main tables of the paper. The unit of observation is the Gram Panchayat for Table 2, 3, 4, 6 and 7. For Table 5 the unit of observation is a household. The data sources are CPSMS financial data (Table 2), official reports from nrega.nic.in (Table 3 and 4), our own survey data (Table 5 and 6) and the match between nrega.nic.in reports and socio-economic and caste census data (Table 7).