Journal of Development Economics 99 (2012) 128–138

Contents lists available at SciVerse ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

Grain inflation: Identifying agent discretion in response to a conditional school nutrition program☆

Leigh L. Linden [a], Gauri Kartini Shastry [b,*]

[a] University of Texas at Austin, IZA, J-PAL, NBER, United States
[b] Wellesley College, United States

ARTICLE INFO

Article history:
Received 3 December 2009
Received in revised form 1 November 2011
Accepted 2 November 2011

JEL classification:
H
I
O

Keywords:
School meals
Conditional transfer programs
Decentralization
Nutrition
Education

ABSTRACT

Many incentive programs rely on local agents with significant discretion to allocate benefits. We estimate the degree of discretion exercised by teachers within a conditional transfer program designed to improve nutrition and encourage student attendance in Mumbai, India. The program allocates grain to students every month their attendance exceeds 80%, creating an incentive for teachers to inflate attendance to benefit certain students. We find that teachers manipulate students' records, altering the incentives to attend school. The teachers' response also varies across students. Teachers inflate more for girls, better students, and students from lower castes, but less for Muslim students.

1. Introduction

The efficient provision of social assistance and transfers to the poor hinges on the ability of governments to target those in need of aid. There is a substantial literature looking at whether decentralization can improve such efforts.[1] One common proposal to improve targeting and reduce the cost of gathering the necessary information is to decentralize the task of allocating transfers among potential beneficiaries to local agents.[2] This is particularly common within the context of conditional transfer programs in which beneficiaries' receipt of transfers depends only on the reports of teachers, doctors, or nurses. In education, for example, a number of studies have focused on causal effects of the programs, showing them to be effective, but relatively little research has focused on their design.[3]

In particular, decentralization in conditional cash transfer programs affords local agents significant discretion in the distribution of benefits if they are willing to falsify their reports of potential beneficiaries' behavior. If local agents exercise this discretion, their behavior may have a significant impact on the efficacy of the overall program. Consider transfer programs conditioned on students' attendance in a given month. Teachers' ability to overstate the attendance of a student with poor attendance to make them eligible for a transfer could make such programs more effective if the teachers have information unavailable to policy makers that allows them to relax the attendance incentives

☆ This project could not have been completed without the assistance of many people. We thank Marc Shotland, Mukesh Prajapati, Rajashree Kabare, and Nandit Bhatt for their assistance in collecting the data used in this study. Pratima Bandekar provided particularly valuable assistance both by helping with the collection of the attendance data and also by conducting many of the informational interviews needed to better understand the context of the program. We are also indebted to Abhijit Banerjee, Shawn Cole, David Cutler, Esther Duflo, Caroline Hoxby, Michael Kremer, Sendhil Mullainathan, anonymous referees and seminar participants at the University of Virginia, Virginia Commonwealth University and Wellesley College for their assistance and invaluable suggestions. We are particularly indebted to Esther Duflo and Abhijit Banerjee for providing us with the data used in this paper. All errors are, of course, our own.

* Corresponding author. Tel.: +1 4013383569.
E-mail address: gshastry@wellesley.edu (G.K. Shastry).

[1] See Bird and Rodriguez (1999) and Klugman (1997) for two reviews of this literature.

[2] Alderman (2002) and Coudouel et al. (1998) study specific poverty programs in Albania and Uzbekistan, respectively, to demonstrate that social assistance is better targeted when local agents are responsible for the allocation. On the other hand, Das (2004) demonstrates that subsidies to local governments distributed with clear guidelines and legislated rules reach the targeted schools, while discretionary subsidies do not.

[3] Recent exceptions include Barrera-Osorio et al. (2011), Fernald et al. (2008), and Paxson and Schady (2007) among others.

0304-3878/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jdeveco.2011.11.002

for students for whom they should be relaxed. For example, teachers may know which children failed to attend school due to circumstances outside of their or their families' control or which families will respond least to moral hazard. However, teachers could also relax the incentives too often and weaken the effects of the program, particularly if their preferences are orthogonal to those designing the program.

Unfortunately, simply determining whether agents do, in fact, manipulate incentive schemes is difficult to identify empirically because the only information on student behavior typically available is that provided by the local agents. One option, for example, is to look for unusual spatial and inter-temporal correlations in reported student behavior. Jacob and Levitt (2003) use this method to demonstrate that about 4–5% of US teachers cheat on standardized tests each year.

In this paper, we approach the question of whether local agents use their discretion using a unique data set that allows us to assess teacher behavior within the context of a conditional school nutrition program in Mumbai, India. The nutrition program distributes 3 kg of grain to families of students with attendance rates of at least 80% in a month. Like most conditional transfer programs in education, eligibility is determined solely by the teachers' attendance records, allowing teachers to influence who does and does not receive the benefit by falsifying those records. We match the official daily attendance records with periodic daily attendance data collected once a week by external monitors.[4] This strategy allows for a direct analysis of teacher behavior. In particular, we are able to identify misreporting when the prevailing incentives elicit similar patterns from teacher misreports and true student behavior.[5]

Our results suggest that, over the 2 years for which we have data, teachers do, in fact, use their discretion to inflate attendance records. Students who do not appear to earn the grain based on the rules of the program often receive it anyway. At least 40% of the students (6079 out of 15,519) in our sample received the grain at least once when they did not appear to deserve it.

To ensure that other incentives teachers face to inflate attendance levels do not drive the results, we exploit variation in the benefit of misrepresenting attendance that could only result from the grain program's monthly threshold. First, we use aggregated data to distinguish the patterns in the data from random measurement error and monthly attendance rates to show that children often receive the grain when their attendance patterns suggest they did not earn it. Second, we use monthly attendance data to graphically demonstrate that teachers consistently move children from around 80% to above 80% attendance. Third, we use daily attendance discrepancies between the teachers' and monitors' records to show that teachers respond to variation in the incentives to misrepresent attendance within each month. Specifically, they are more likely to misrepresent attendance for children who are still in the running for the benefit and misrepresent attendance more towards the end of the month when they know which children they can assist. Finally, we also find suggestive evidence that teachers inflate attendance less frequently for boys, children with low test scores, Muslim children, and high caste students, suggesting that teachers are creating different incentives for individual students.[6]

[4] Martinelli and Parker (2009) use a similar strategy to study a program that allows individuals to report the data themselves. They find widespread underreporting in order to qualify despite stigma motives for over-reporting.

[5] Unfortunately, the data does not, however, allow us to assess whether the teachers' response is welfare improving or reducing. In this paper, we demonstrate that such behavior exists and that it must be taken into account in the design of such programs.

[6] It is important to note that differential treatment by student characteristics alone is not evidence of religious or gender discrimination, since these traits may be correlated with other characteristics. For example, one teacher surveyed noted that Muslim students in her school tended to be better off financially than other students. However, in contexts where the potential for discrimination is high, there is a significant concern that teachers may allow their personal prejudices to influence their behavior.

The remainder of the paper is organized as follows. In Section 2, we describe the details of the grain program. In Section 3, we describe the unique data set that allows us to view both actual attendance and teacher-taken attendance. In Section 4, we discuss the various incentives teachers face when recording attendance and the dynamic nature of the incentives created by the monthly threshold. We also describe the empirical specifications used to show that teachers respond to the grain incentives. In Section 5, we describe results using aggregate and monthly data, and in Section 6, we empirically distinguish between responses to the grain program and other incentives to inflate attendance. Finally, in Section 7, we examine how the teachers' behavior relates to student characteristics. We conclude in Section 8.

2. The grain distribution program

The conditional distribution of grain is part of the National Programme of Nutritional Support to Primary Education (NSPE) created in 1995. The NSPE mandated the provision of a daily meal to primary school children all over India. Initially, due to lack of funds, the program was often implemented as 'dry rations', the distribution of uncooked rice to parents whose children surpassed attendance rates thresholds. This program operated in Mumbai in this conditional form during the period of our analysis, but starting in 2005, the government began providing cooked midday meals.

The Mumbai program targeted the poorest households by distributing grain through public municipal schools, where these families sent their children. Children in first through fifth grade were eligible, and there was no explicit poverty target. The program was administered monthly: a child with 80% attendance in a month earned 3 kg of uncooked rice (Mumbai Interviews 2005). Attendance records were individually compiled by teachers. Each teacher was provided with a single sheet of paper that listed each registered child with space to record the child's attendance on each day of the month.

Every month, principals gathered attendance records from the teachers, calculated the number of children who had at least 80% attendance that month, and requested the correct amount of grain from the government. Every 3 months, an administrator would examine the attendance records. If the forms were properly completed, the administrator would give the principal the grain. Teachers then asked parents to collect one, two or three bags of rice depending on the number of months their child had at least 80% attendance. In L Ward, the district of Mumbai where our data was collected, the system seemed to run smoothly. Principals almost always got the number of bags they requested (Mumbai Interviews 2005).

Auditors were responsible for periodically visiting the schools and inspecting attendance records, but audits were infrequent. During an audit, the principal would be responsible for any discrepancies, but this would be limited to reconciling the request for grain with the attendance records. Auditors could not verify actual attendance of the students except by observing suspicious peculiarities in the recorded attendance patterns (such as a significant number of children with exactly 80% attendance).

With such limited monitoring, it is likely teachers could manipulate attendance with significant discretion. Interviews with teachers suggest that this did, in fact, occur (Mumbai Interviews 2005). Teachers were reluctant to admit manipulating records themselves (though one did), but generally suggested that those who did were responding to the extreme poverty in which their students lived. While dishonest, the action was characterized as a compassionate response to a difficult situation, and there were no reports of teachers demanding kickbacks. This is supported by the fact that teachers earn much more than the parents, and the quality of the grain distributed was much lower than what teachers would normally buy.[7]

3. Data and summary statistics

The uniqueness of our data allows us to compare two overlapping records of student attendance. The data were collected during an evaluation of a remedial education program in Mumbai run by the Pratham Mumbai Educational Trust (Banerjee, Cole, Duflo and Linden, 2007).[8] As part of the study, we acquired all attendance records taken by teachers in the third and fourth grades in L Ward during the 2001–02 and 2002–03 academic years; these data were used to allocate the grain. However, due to widespread concern that teachers significantly overestimated the attendance of students, we also hired a team of independent surveyors who visited each class in our sample once a week on a randomly chosen day and time to record attendance.

The attendance data from the teachers' official rosters (hereafter the 'roster' data) include the name of each child enrolled in school and the student's recorded attendance for every day the school was open. The Pratham-collected data (hereafter the 'monitored' data) includes the same student identification information along with the dates the class was observed and the students who were present during the visits. In addition, the data set contains student characteristics such as age, gender and scores on math and language tests administered at the beginning of the school year. We also know the language of instruction for each school.[9]

The sample contains 77 schools with an average of two classes per school in the first year (for grade three) and four classes per school in the second year (two classes for grades three and four each), resulting in data from 436 teachers on approximately 15,000 children. Attendance monitors recorded data from September to February each academic year. This includes almost the entire school year, but excludes August when student enrollment is extremely low as students begin to come back to school and March when students are taking end of the year exams.

Table 1 presents summary statistics. Panel A summarizes student characteristics such as age, religion and caste. Panel B describes the monthly attendance rates. We have data on approximately 75,000 student-months. The average roster attendance rate was 90.5%, while the average monitored attendance rate was only 85.9%. Note that we do not have accurate measures of monthly attendance rates taken by the monitors since monitors only visit a classroom three to five times a month depending on how many weeks school was in session. Around 87% of children attended more than 80% of the month according to the roster data, but we estimate that only 72% of children reached that attendance rate from the monitored data. Of all student-months, 20% of those with recorded attendance greater than 80% had monitored attendance below 80%; in other words, 20% of student-months awarded the grain did not appear to have earned it.[10]

Panel C presents summary information on the variables in the empirical specifications described in Section 4.2. This data is measured at the student-school day level but only for days on which the student was absent, when teachers have the choice of whether to overstate the child's attendance. Overall, teachers misrepresent attendance 43% of the time on average, but the tendency varies by teacher; one teacher in our sample marked 125 student-days incorrectly present (out of an observed 128 absences) while another never exaggerated her students' attendance (out of an observed 57 absences).

The large number of cases in which a teacher erroneously records an absent child as present provides the first piece of evidence pointing towards systematic misrepresentation. Compared to this substantial fraction (43% of monitored absences), the opposite misrepresentation (present children marked absent) occurs only in a small fraction of cases (2% of monitored presences).[11] The bias is clearly towards

Table 1
Summary statistics.

	Obs	Mean	Standard Deviation	Min	Max
	(1)	(2)	(3)	(4)	(5)
Panel A: Student characteristics					
Child in Urdu School	15,519	0.400	0.490	0	1
Muslim child in non-Urdu school	15,519	0.075	0.264	0	1
Brahmin caste	15,519	0.063	0.243	0	1
Kshatriya caste	15,519	0.106	0.308	0	1
Vaishnav caste	15,519	0.013	0.111	0	1
Shudra caste	15,519	0.041	0.199	0	1
Unknown caste	15,519	0.777	0.416	0	1
Male	15,519	0.478	0.500	0	1
Normalized pre-test score	15,519	0.009	0.991	−2.1	3.5
Age	13,785	8.701	1.351	5	15
Panel B: Student-month data					
Roster attendance rate	76,965	0.905	0.196	0	1
Monitored attendance rate	50,310	0.859	0.252	0	1
Child had ≥80% roster attendance	76,965	0.874	0.332	0	1
Child had ≥80% monitored attendance	50,310	0.721	0.448	0	1
Panel C: Student-absence data					
Misrepresented attendance	24,192	0.430	0.495	0	1
Still eligible: still in the running	24,192	0.738	0.440	0	1
Needs help: still in the running but has not already earned the grain	24,192	0.678	0.467	0	1
Already earned: has already earned the grain	24,192	0.060	0.238	0	1
Needs help plus perfect attendance this month	24,192	0.302	0.459	0	1
Already earned plus perfect attendance this month	24,192	0.045	0.207	0	1
Days left in month	24,192	11.476	6.225	1	26

Note: This table displays summary statistics of all the variables used in the following tables. Panel A describes student characteristics. Panel B presents summary information about monthly attendance rates according to both the teacher-taken (roster) and monitored records. Panel C summarizes variables that derive from the variation in the incentives to misrepresent attendance according to the child's past attendance record and the number of days left in the month. See Section 4.2 for more details about the variables in Panel C.

[7] Comparing the two attendance records argues against corruption as teachers' motivation. If teachers intend to extract rents from parents in exchange for awarding their children the grain, we would expect them to concentrate this behavior among families that are amenable to such negotiations. Our data suggest, instead, that almost 70% of the children who receive the grain without appearing to deserve it in at least 1 month do not fall into this category in any other month. Only 7% of those who receive the grain without earning it do so more than twice.

[8] The remedial education program was shown to have no effect on attendance, and even if it had, it would not have caused the monthly attendance patterns observed in this paper.

[9] The dataset does not include the religion of each child, but there is a close correspondence between name and religion in India. Classifying a student's name by religion reveals that almost all the students in schools taught in Urdu are Muslim, while there are Hindu, Jain, Christian, Muslim and Sikh students in non-Urdu schools. We also classified student names by caste between Brahmin, Kshatriya, Vaishnav and Shudra in descending order of status, but are only able to classify 22% of students.

[10] In Section 5.1, we estimate how much of this discrepancy could be due to sampling error from the small number of monitored days.

[11] It is also possible that a student may have arrived after or left before the attendance monitors arrived. Generally, these behaviors should not co-vary with the incentives of the grain program, and thus, should only add noise to the estimates. However, even so, it is unlikely to be a significant problem as the school day is only 4 hours a day – making it unlikely that a child would only attend for part of the day – and even then, the monitors varied the times at which they visited individual schools. Additionally, one would expect such instances to be distributed across teachers, but we observe many teachers for whom the roster and attendance data match almost perfectly. It is also possible that the grain program gives students an incentive only to attend long enough to be marked present and to do this frequently enough to obtain the grain. In this case, recording such students absent seems consistent with the goals of the program since the program was intended to elicit attendance for the full day.

exaggeration and not understatement. This is not evidence of teachers manipulating records in response to the grain program specifically, but it does illustrate the feasibility of doing so. Also for this reason, we will use 'misrepresent attendance' to refer only to marking a child present on a day he is absent unless otherwise noted.

Finally, the method of data collection raises a possible concern. Because the attendance monitors are directly visiting the teachers' classes, there is the possibility that the data collection may have caused teachers to be more cautious than they otherwise would have been (a Hawthorne effect). The very fact, however, that teachers still misrepresent a large fraction of absences (43%) belies this concern. We test for this by comparing teacher-taken attendance records on days the monitor visited to teacher-taken attendance records on days the monitor did not visit. The difference in attendance is miniscule and in the wrong direction: a student is 0.035 percentage points more likely to be recorded as *present* on a monitored day, with a p-value of 0.19. In addition, the data in the first year was collected at the end of the academic year which made it impossible for teachers to know in advance that the attendance data could be used to scrutinize their records (we find similar effects using data from each year separately).[12]

4. Empirical specifications

4.1. Analytic strategy

We present two types of evidence that suggest that teachers are misrepresenting attendance in response to the grain program. First, we provide general evidence that teachers are, in fact, misrepresenting students' attendance and that this misrepresentation is significant enough to cause some students who should not otherwise be eligible for the grain to qualify anyway. Since teachers face general incentives to inflate attendance – for example, to ensure that students are promoted to the next grade, to make the school aggregate administrative data look more positive, or to increase school resources because funds are distributed conditional on attendance – this evidence does not allow us to distinguish between these other incentives and the grain program. However, it does establish the fact that teachers can and do falsely report student attendance.

We next provide evidence that the frequency and timing of misreports is consistent with what one would expect given the intra-monthly variation in the incentives of the grain program. At the monthly level, we show that teachers' misrepresentations occur most frequently for students just below the 80% target rate and that the misrepresentations move children above the target rate. Then, using daily attendance data, we inspect the exact timing of the misreports within the month, testing, for example, whether teachers are more likely to misrepresent attendance at the end of the month when there is less uncertainty about who might benefit from misrepresentation. The analysis of the daily decision to over-report student attendance has two distinct advantages. First, teachers' responses can be attributed to the grain program since the within month incentives created by the grain program are unique from other incentives faced by teachers. Second, it allows us to distinguish the effects of the grain program on teacher behavior even when the program creates the same pattern of incentives for *student* attendance. The incentive to misrepresent attendance for a particular child on a given day if that child is absent is positively correlated with the incentive for that child to attend school. When we compare roster and monitored daily attendance records for an absent child, we look for teacher responses to the grain over and above the student's attendance response.[13]

[12] Even if the teachers did believe that the monitors' data would be used to scrutinize their records, this would simply add a uniform cost to falsifying absences. This cost should not co-vary with the incentives arising from the grain distribution's monthly threshold.

[13] We do find monthly patterns in student attendance: children attend school more towards the end of the month and when they are more likely to earn the grain.

Both sets of analyses assume that teachers have an incentive to give children the grain, even when they have not earned it. This may be because, in line with the goals of the program, they are appropriately modifying the incentives to account for the circumstances of individual students. However, teachers might also do it to prevent certain parents from harassing them for the grain, to respond to outright bribes from parents, to induce the children to work for them in some way,[14] or simply to experience a "warm glow" from giving the students the grain.[15] Both analyses also assume that it is costly for the teachers to misrepresent attendance — either because there is some probability of their getting caught or because they simply prefer not to do so. If falsifying records was costless, the teachers could just always record students as being present.

The daily analysis requires one additional assumption that the monthly analysis does not. We have to assume that it would be costly for teachers to redo or significantly change their attendance records at the end of the month.[16] Violations of this condition would relieve teachers of having to make daily considerations about whether or not to lie for their students. If it was costless to revise the entire attendance roster at the end of the month, then, even if lying itself were costly, they could spread out the misrepresentations across the month to avoid generating the very tell-tale patterns, for which we test.

Since relaxing any of these assumptions would work against finding the patterns for which we test, the empirical results we present in Section 6 suggest that there are costs associated with misrepresenting attendance or recopying the attendance rosters. Furthermore, teachers' apparent practice supports the assumptions as well. Attendance taking is standardized. Teachers take attendance daily and do so (typically in ink) by writing "A" or "P" within a grid listing students on the vertical axis and the days of the given month on the horizontal. They do not, for example, simply add a check mark to an empty cell. Thus, revising a previous day's attendance requires visibly altering the earlier record which would look suspicious if done too often. Further, recopying an entire month's attendance would at the very least take time since it would require recopying the records for every student. These details combined with the fact that such a strategy was never mentioned during the field interviews, suggest that, at the very least, redoing a month's attendance is not costless.[17]

4.2. Empirical tests

We use several empirical specifications that allow us to test three sets of hypotheses. First, we use the aggregated data to distinguish the observed pattern of falsifications from measurement error. Similarly, analyzing the data by month, we show that teachers do, on average, award the grain to students who should not receive it. We also graph monthly attendance rates, showing that teachers seem to be consistently moving children from around 80% to above 80% attendance. Second, within each month, we show that teachers misrepresent attendance more for children who are still eligible for the grain

[14] We are grateful to an anonymous referee for this suggestion.

[15] While we cannot identify the underlying rationale for awarding children with insufficient attendance the grain, the different motivations have important implications for the ultimate effect of allowing teachers such discretion, and since the evidence suggests that teachers do use their discretion, these are important areas for future work.

[16] Note that we do not require the cost to be sufficiently high that this never happens, but rather that it has to be sufficiently high that in some cases, it forces the teachers to make daily decisions about misrepresentation given the child's recorded attendance to date as well as the number of school days left in the month.

[17] Another possible mechanism for ex-post misrepresentation is for teachers to only fill in the attendance roster when a child is present, leaving the roster blank if the child is absent. Then at the end of the month, the teacher could retroactively distribute the misrepresentations by filling in some of the blanks to record the child as having been present. However, even this imposes a cost to the teacher since, if inspected mid-month, such a half-filled attendance roster would seem quite suspicious. We are grateful to an anonymous referee for this suggestion.

and at the end of the month — i.e. students whom the teachers can still help, during a period in time when it is clear who needs help and the cost of helping are possibly lowest. Finally, we show that the propensity of teachers to exaggerate attendance is correlated with certain student characteristics, suggesting that teachers are relaxing incentives for preferred students.

First, we use aggregate monthly data to show that children receive the grain despite having low monitored attendance rates using the following specification:

$$\Pr\left(gotgrain_{ijm}=1\right)=f\left(\alpha+\beta\cdot monitoredattendance_{im}+\nu_{jm}+\varepsilon_{ijm}\right) \tag{1}$$

where $gotgrain_{ijm}$ is an indicator for whether child i achieved 80% recorded attendance from teacher j in month m. The variable $monitoredattendance_{im}$ is a measure of child i's monitored attendance in month m, and ν_{jm} is a teacher-month or student fixed effect. The fixed effects ensure that these results are not driven by teachers who always misrepresent attendance (even if just in certain months) or teachers who always misrepresent attendance for their favorite students. If teachers are following program rules, monitored attendance should explain most of the variation in when the grain is awarded. Since we do not observe the grain distribution itself, we proxy for $gotgrain_{ijm}$ with a dummy variable for whether the child's teacher-taken attendance rate that month was at least 80%. This strategy, however, suffers from a type of sampling error: we only have noisy measures of monitored monthly attendance due to the small number of days per month the monitor visited the classroom. We run simulations to account for this.

Next, we turn to the daily attendance data and estimate the following specification:

$$\Pr\left(misrepresent_{ijt}=1\right)=f\left(\alpha+\beta_0\cdot eligible_{it}+\nu_{jt}+\varepsilon_{ijt}\right) \tag{2}$$

where $misrepresent_{ijt}$ is a dummy variable for whether teacher j exaggerates child i's attendance on day t (conditional on the child being absent on day t), and $eligible_{it}$ is a measure of whether the child is still eligible for the grain given his past attendance this month.

Since there is little information on the cost of overstating a child's attendance, we proxy for this cost with fixed effects at different levels: teacher-day or student. When we include student fixed effects, we also include month fixed effects. These fixed effects account for the possibility that some teachers may face lower costs than others on certain days or for certain students. These costs may be due to a psychic cost of lying, principals who monitor attendance, auditors from the district's administration, or a reputational cost since students may respond by reducing their attendance further. The teacher-day fixed effects also ensure that these results are not driven by teachers who misrepresent attendance for many children on a given day because the teacher herself was absent. Similarly, the student fixed effects ensure that our results are not driven by teachers who misrepresent attendance only for certain students. The results are driven by teachers who misrepresent attendance for a particular child when it might help that child achieve the grain but not at other times.

Note that a response on the part of the teacher to a child's attendance rate in general is not sufficient evidence of a response to the grain: a child's attendance rate may affect the other incentives described above, such as advancement to the next grade, or the cost of misrepresenting attendance may depend on a child's attendance rate. However, teacher behavior in response to child attendance in that particular month, when controlling for the child's average attendance rate with a student fixed effect, strongly suggests a response to the grain.

We first use a broad measure of continued eligibility that includes all children who have not missed more than 20% of the month already. However, if teachers are sophisticated, their propensity to exaggerate attendance should depend on more than this broad measure of eligibility. When making the daily decision to misrepresent attendance, there are several conditions in which teachers experience starkly different incentives. First, if a child's outcome has already been determined – either because the child has missed too many days or because he or she has attended enough days to earn the grain – then the teacher should not misrepresent attendance further. We re-estimate specification (2) with a new measure of eligibility that separates students who have already earned the grain from those who still require help, comparing both groups to those who are already disqualified.

Second, teachers should wait until later in the month to decide whether to misrepresent attendance. Early in the month, students' average attendance for that month is less certain and misrepresentations may be wasted on students who will eventually either receive or become ineligible for the grain despite the teachers' efforts. Waiting until the end of the month allows the teachers to wield their discretion more precisely.[18] To test for this, we estimate the following:

$$\begin{aligned}\Pr\left(misrepresent_{ijt}=1\right)=f\Big(&\alpha+\beta_1\cdot eligible_{it}+\beta_2\cdot daysleft_t\\&+\beta_3\cdot eligible_{it}\cdot daysleft_t+\nu_{jt}+\varepsilon_{ijt}\Big)\end{aligned} \tag{3}$$

where $daysleft_t$ is the number of school days left in the month.[19] We would expect β_3 to be negative since teachers should misrepresent attendance closer to the end of the month and the relationship should be stronger for children who can still be helped.

Finally, in an unpublished appendix to this article, we also build a formal model of the teacher's decision process. We build a Bellman equation that relates the decision on a given day to the future series of decisions that the teacher will have to make if the child is absent given the students' attendance to date and the decisions already made by the teacher. This model, of course, requires much stronger assumptions about a teacher's ability to change information retroactively, but it also allows us to exploit much finer variation in the incentives that teachers face over the course of a month. We find that teachers do respond accordingly.

5. Empirical results: aggregate and monthly attendance data

In this section, we analyze the aggregate and monthly attendance patterns in the data. One advantage to focusing on aggregate data is that we do not need to assume that teachers find it so costly to retroactively revise attendance that they make some of the decisions to misrepresent attendance on that day itself. The primary disadvantage is that we have only 3–5 days of monitored records per month, creating noise in our estimates of monthly attendance rates. First, we formally test the hypothesis that the discrepancies in our two data sets are due to intentional misrepresentation and not simply measurement error. Second, we use the parametric model in Eq. (1), to document that children are receiving the grain when they should not.

[18] As pointed out by an anonymous referee, it may also be more useful for teachers to wait until the end of the month if they want to use the misrepresentations to elicit additional effort from a parent or a child. A deal in which the teacher gets a parent or child to agree to extra attendance the following month or additional effort in class would be more easily struck later in the month when it was obvious that the child had attended too few days and the teacher was in a position to obtain the grain for the family.

[19] We use the number of days left instead of the number of days passed because the total number of school days in each month varies due to public and religious holidays.

5.1. Aggregate trends

It is possible that teachers make random mistakes in recording attendance, as classes are large and potentially full of unruly children. The summary statistics described above suggest this is unlikely to explain the large number of discrepancies. Teachers are much more likely to mark an absent child present (43%) than a present child absent (2%). It could be, of course, that teachers fill in the attendance records at the end of the month, suffer from recall issues and employ a "when in doubt, mark present" strategy. The fact that 93% of the days we can compare match perfectly (171,067 out of 184,662) suggests this is unlikely as well.

Comparing the monthly attendance rates between the teachers' and the monitors' data suggests that teachers award the grain too often. Of the student-months that were awarded the grain, 20% were months in which the monitors' data recorded less than 80% attendance. Some discrepancy between the measures should be expected because there are at most five monitored attendance observations for each month, and we would expect that just through random chance, the monitors could take attendance on days when qualifying students are absent. For example, in a month with 20 days, a child with 80% attendance will be absent on 4 days. If monitors randomly visit on 2 days the child is present and 3 days the child is absent, the child will appear to have not earned the grain when in fact he had. The question is whether the observed discrepancy is too high to be due to such random chance.

To address this, we calculate the distribution of this summary statistic under the null hypothesis that teachers record attendance accurately. Treating each child's monitored attendance record in a month as a binomial variable with 3–5 trials with the probability of success equal to the child's teacher-taken attendance rate that month, we calculate the probability that a child whose teacher-taken attendance rate is greater than 80% attends fewer than 80% of the monitored days. Then, we treat whether this child deserved the grain as a Bernoulli random variable. Using the central limit theorem, we calculate the mean and standard deviation of this summary statistic to be 8.93% and 0.12%. It is therefore extremely unlikely that the 20% discrepancy in whether a child who received the grain actually deserved it is due to sampling error from how the monitor data was collected.[20]

5.2. Empirical regressions using monthly data

We can also analyze the monthly attendance records parametrically using Eq. (1). The results are presented in Table 2. Columns 1–5 estimate linear probability estimates while columns 6–7 estimate conditional logit models (we only show these with student fixed effects to save space — the results from other specifications are very similar). Columns 1, 2 and 6 test the conditional correlation between whether a child earned the grain, as measured by his attendance rate from roster data, and the percent of monitored days the child was present. All three regressions suggest strongly that a child with greater monthly attendance is more likely to receive the grain. The student fixed effects ensure that our results are not driven by students who consistently receive the grain or consistently fail to receive the grain.

We expect the relationship between the attendance rate and receiving the grain to be positive, but the question is whether it is too small. To gauge this, we want to compare the coefficients in columns 1 and 2 with the coefficient we would obtain if teachers are recording attendance accurately. This coefficient depends on the underlying distribution of students' attendance rates, but can be directly estimated by using the attendance rates from the teachers' roster data in Eq. (1) rather than the monitors' data. The resulting coefficient is 1.45 and is presented in column 5. Based on the estimate in column 2, in a month where a particular child has a 10% higher monitored attendance rate, that child will be 6 percentage points more likely to receive the grain relative to other months. If teachers recorded attendance accurately and we had monitored data every day, this statistic would be closer to 14.5 percentage points.

Unfortunately, this benchmark coefficient is too large. The fact that we only have three to five observations of students' true attendance each month means that we have a noisy estimate of the students' true attendance rate. Under the null hypothesis of accurate teacher reporting, this is a classic case of measurement error in the independent variable, and we would expect the resulting coefficient to be attenuated. To separate out the issue of sampling, we re-estimate Eq. (1) but instead of using the actual monitored data we simulate the monitored attendance rate under the null hypothesis that teachers record attendance accurately. Specifically, we calculate the simulated attendance rate using only three to five randomly chosen days of roster attendance. We bootstrap this estimate, along with the R-squared statistic, 1000 times to construct a 99% confidence interval, presenting the results in Panel B.

In columns 1, 2 and 6, the coefficient on the monitors' attendance rates is much lower than what one would expect from simple measurement error alone. Focusing on column 2, the mean coefficient in Panel B is 0.88 and the estimate of 0.63 is lower than the lower bound of the 99% confidence interval. This means that the students' actual attendance rate – formally the only determinant of whether students receive the grain – is not a strong enough determinant of whether students actually receive the grain. Consistent with this, the explained variation is too low — the R-squared statistic of 0.6 is also below the lower bound of the 99% confidence interval of the bootstrapped R-squared statistic.

A similar strategy is to estimate how much of the variation in who receives the grain can be explained by an indicator variable for whether a child's monitored attendance rate is less than 80% (columns 3, 4 and 7 in Table 2). This may, in fact, be a more accurate estimate if the relationship between receiving the grain and the child's true attendance rate were sufficiently non-linear. In this case, the coefficient should be close to negative one since falling below 80% attendance should disqualify the student from receiving the grain, but there is likely to be measurement error. The results demonstrate that not deserving the grain significantly reduces the probability a child receives the grain; however, we see that while a child who does appear to deserve the grain earns it about 95% of the time (the constant term), not deserving the grain only reduces this probability by 23–35 percentage points. From the simulated confidence intervals presented in Panel B, we can conclude that this discrepancy is unlikely to be due to error in our measure of the child's monthly monitored attendance. A child who appears to not deserve the grain should be 39–50 percentage points less likely to actually receive it. These results confirm that teachers respect the rules enough to ensure that children below 80% are less likely to receive the grain, but that the child's attendance alone does not determine whether he or she receives the grain.

6. Empirical results: incentives created by the grain program

The results above indicate that in recording attendance inaccurately, teachers affect the distribution of the grain. We are able to rule out the possibility that the discrepancies are due solely to random, mean zero, mistakes. However, the results above do not rule out the possibility that teachers exaggerate attendance for reasons other than the grain program. To differentiate between these reasons, we assess whether teacher behavior is consistent with the intra-month variation in incentives created by the grain program by analyzing the distribution of attendance rates and by estimating Eqs. (2) and (3).

[20] Another approach is to isolate cases where a child has missed enough of the monitored days that even if he was present all other days in the month, he could not have earned the grain. For example, if the monitor visited a classroom five times in a month with 21 days and the child was absent all 5 days, he would be disqualified for sure. Out of the 261 such instances in our data, 22 (8.4%) student-months received the grain anyway. This is a very weak lower bound on the extent to which misreporting affects the grain distribution.

Table 2
Empirical tests of whether inaccurate attendance records affect grain distribution.

Dependent variable:	Child had ≥80% roster attendance						
	Linear probability					Conditional logit	
					Benchmark		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A							
Monitored attendance	0.795***	0.632***				5.468***	
	(0.017)	(0.024)				(0.194)	
Monitored attendance <80%			−0.345***	−0.229***			−2.383***
			(0.012)	(0.012)			(0.077)
Recorded attendance					1.447***		
					(0.026)		
Constant	0.183***	0.324***	0.964***	0.945***	−0.423***		
	(0.015)	(0.026)	(0.003)	(0.017)	(0.024)		
Additional fixed effects	Teacher × month	Student	Teacher × month	Student	Student	Student	Student
p-value of joint test of f.e.	0.000	0.000	0.000	0.000	0.000		
Number of observations	49,742	49,742	49,742	49,742	49,742	14,013	14,013
R-squared	0.446	0.604	0.330	0.561	0.780		
Panel B							
From simulated trials:							
Coefficient:	0.97	0.88	−0.49	−0.4		7.15	−2.99
99% confidence interval	(0.96, 0.98)	(0.87, 0.9)	(−0.48, −0.5)	(−0.39, −0.41)		(6.75, 7.54)	(−2.86, −3.15)
p-value	0.000	0.000	0.000	0.000		0.000	0.000
R-squared	0.54	0.67	0.43	0.62			
99% Confidence interval	(0.53, 0.55)	(0.67, 0.68)	(0.42, 0.44)	(0.62, 0.63)			
p-value	0.000	0.000	0.000	0.000			

Note: This table displays estimates of Eq. (1). All observations are at the student-month level. The dependent variable is whether or not the student was assigned to receive the grain from his or her teacher while the primary independent variable in columns 1, 2, 3, 4, 6, and 7 is the students' attendance record as measured through the direct monitoring of attendance. Because the exact magnitude of the coefficient on monitored attendance is difficult to predict under the null hypothesis that teachers accurately record attendance, column 5 provides a benchmark regression in which the probability of a student receiving grain is regressed on the students' teacher-recorded attendance record. Columns 2, 4 and 5 also include month fixed effects. Because of the measurement error problems due to the sampling of days for direct monitoring of attendance (see Section 5.2 for a more detailed explanation) Panel B provides estimates of the coefficients on monitored attendance in Panel A from simulated roster data (under the assumption that teachers never misreported attendance and that attendance is taken at the same rate as the monitored attendance data). The estimates are the mean and 99% confidence intervals from a 1000 iteration bootstrap procedure. Columns 1–5 are estimated using a linear probability model while columns 6 and 7 are estimated using a conditional logit model. Robust standard errors clustered by teacher are in parentheses. Observations are used only if there were at least 3 dates matched between monitored and roster attendance data for a particular student-month.
*** 1%, ** 5%, * 10%.

6.1. Density Plots

The analyses presented in the previous section demonstrate that teachers inflate attendance records systematically. To look more closely at the attendance patterns, we first plot distributions of the attendance rates. It is important to note that because the attendance monitors only sample attendance 3 to 5 days a month, there are only 11 possible observable attendance rates in the monitored data (0, 20, 25, 33, 40, 50, 60, 67, 75, 80, and 100%) while the attendance rates in the roster data are much more continuous. The differing levels of discreteness make comparing the distributions a challenge. For example, if we simply plotted continuous distributions of each measure, the monitored data would naturally have a mass at 80% even without the incentives created by the grain program.

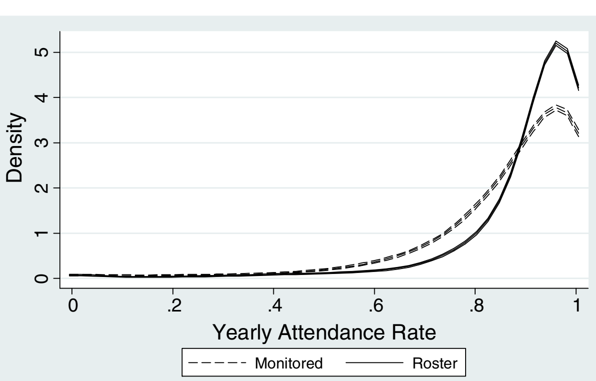

Note: This graph plots the kernel densities of roster and monitored yearly attendance rates separately using a Gaussian kernel and a bandwidth of 0.05. All plots include 95% confidence intervals.

Fig. 1. Density of yearly attendance rates.

To account for the differing levels of discreteness, we adopt two complementary strategies to compare the underlying data. First, we make the monitors' data more continuous. This is not possible on a monthly level, but if teachers have an incentive to exaggerate attendance by month, this will result in exaggerations by year. If we aggregate the monitored data to the student-year level, the underlying distribution of both the monitors' and teachers' data will be similarly continuous. Fig. 1 plots the results of this exercise. The solid line plots the smoothed distribution of attendance rates calculated from the roster attendance, along with a 95% confidence interval created by bootstrapping the plot using a subsample of 10,000 with replacement (the intervals are very small). We use a Gaussian kernel and a bandwidth of 0.05; the shape of each plot is robust to different kernels and bandwidths, as well as the local linear density estimator (Cheng, Fan and Marron, 1997). The dashed line plots the density of the monitored attendance. Compared to the teachers' records, the monitors' attendance rates show that the teachers record too few students as having less than 80% attendance and too many having more than 80% attendance.[21]

[21] Note that while the incentives associated with the grain program could have created a discontinuity at 80% attendance per month, there is nothing that guarantees such a feature in the annual data as such a discontinuity could easily average out over many months and across several thousand students.

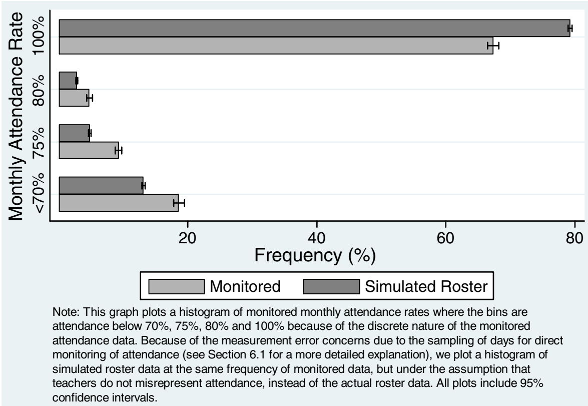

Fig. 2. Histogram of monthly attendance rates.

Another strategy is to use the opposite approach: rather than making the monitor data more continuous, we make the roster data more discrete by reducing the teacher-recorded data to the same number of observations per month as the monitor data. We do this by only using three to five randomly chosen days of roster data so that the structure of the two data sets is identical, making any remaining difference reflect only the teachers' behavior. This process is done for each month and repeated 100 times; the results are plotted as the darker bars in the histogram presented in Fig. 2 along with 95% confidence intervals represented by the bars at the top of each column. The lighter columns represent the distribution of the monitors' attendance data. While the data are more discrete than in Fig. 1, a clear pattern still emerges — we observe many more instances of students with 100% attendance, and proportionately fewer students with attendance rates of 80, 75, and 70% or less in the teachers' simulated data when compared to the monitor's data.

6.2. Daily attendance estimates

We next determine whether or not the daily decisions recorded by teachers correspond to the incentives created by the grain program by estimating Eqs (2) and (3). The results demonstrate that the teachers' behavior does respond to variation in the incentives due to the grain program, although their behavior is not consistent with all of the nuances described in Section 4.2.

Table 3 tests whether teachers misrepresent attendance more for children who are still eligible for the grain, where eligibility is calculated based on roster attendance and the coefficients are estimated using Eq. (2). We estimate these regressions with linear probability models due to the number of fixed effects we want to include.[22] Columns 1–2 use a broad measure of eligibility including all students who have not been recorded absent for 20% or more days of the month before today.[23] The specification with student fixed effects (column 2) demonstrates that teachers are 12 percentage points more likely to misrepresent attendance for children still eligible for the grain.

We next separate students who have already earned the grain (i.e. attended enough days) from those who still need to attend school (i.e. have not missed more than 20% of the month but have not yet been present for 80%) in columns 3 and 4. If teachers are sophisticated in their decisions to misrepresent attendance, they should mark an absence accurately for a child who has already earned it. We find that teachers are 11 percentage points (column 4) more likely to lie for a student when he or she still needs to attend school than when that same student is already disqualified (the omitted group). However, we find that teachers continue to exaggerate attendance for children even after they have earned the grain.

This latter result may be due to many factors. Teachers do, as noted above, have other incentives to inflate attendance. It may also be, however, due to the fact that teachers spend very little time recording attendance and are unlikely to calculate whether each child has achieved 80% attendance. In particular, since student-days can only ever be in this category close to the end of the month, this result could stem from a tendency to misrepresent attendance more at that time. In addition, if there is a cost to misrepresenting attendance, it is likely smaller for children who attend frequently; if anyone asks, other students and other teachers are likely to have seen that student at school. Given the small cost of misrepresenting attendance, it is likely teachers do not believe it is worth their time to prevent some lies from being wasted on children who have already earned the grain. It could also be that these children have more than 80% attendance due to past instances when the teacher has exaggerated attendance for them this month; the teacher may be in the habit of misrepresenting these students' attendance. Note that such a habit could only explain the teacher's behavior in a given month, since the student fixed effects control for a general tendency for teachers to lie for particular students. To ensure that the results in columns 1–2 were not driven by children above the threshold, we re-estimate these regressions excluding those children (columns 5–6), finding similar results.

We next assess whether or not teachers misrepresent more at the end of the month in Table 4. Columns 1–4 estimate variants of Eq. (3).The first column shows that teachers are more likely to misrepresent attendance when there are fewer days left in the month. Columns 2–3 include a measure of eligibility and an interaction with the number of days left in the month; column 3 excludes children already above the 80% threshold. While both main effects are of the expected sign (teachers are more likely to mark absent children present when they are still eligible for the grain and closer to the end of the month), the interaction is not significant. In column 4, we find similar results when we include an indicator for the students who have already earned the grain.

Finally, in column 5 we include an additional measure of past roster attendance (whether the child has perfect roster attendance up to today) and interactions with the number of days left. In column 6, we include fixed effects for the number of days left to ensure that these results are not driven by nonlinearities in the monthly variation in incentives. We find that for a child who still needs to attend in order to earn the grain, perfect past attendance increases the teacher's tendency to exaggerate attendance. This makes sense because it is something teachers can scan easily. However, perfect past attendance does not affect teacher behavior for students who have already earned the grain. In addition, the interaction of eligibility and the number of days left is negative and significant, suggesting that teachers respond more to monthly variation in incentives for children who still need to attend. This interaction is not significant for children who have already earned the grain (even for those with perfect past attendance), although it is negative and large in magnitude. Surprisingly, the interaction with children who still need positive attendance, but have

[22] Conditional logit results are very similar, but we have to drop the month fixed effects for the model to converge.

[23] Note that if teachers attempted to give the grain to all children, then no child should ever have missed more than 20% of the month, rendering this test ineffective. However, this is unlikely given that approximately 40% of student-months do not earn the grain: while it is unclear how large the cost to misrepresenting attendance is, it is certainly large enough to prevent teachers from ensuring that all students receive the grain.

Table 3
Empirical tests of whether teachers respond to grain incentives.

Dependent variable: misrepresent						
	(1)	(2)	(3)	(4)	(5)	(6)
Still eligible	0.387***	0.118***				
	(0.019)	(0.015)				
Needs help			0.378***	0.114***	0.373***	0.104***
			(0.019)	(0.015)	(0.019)	(0.014)
Already earned			0.449***	0.221***		
			(0.030)	(0.026)		
Additional fixed effects	Teacher × day	Student	Teacher × day	Student	Teacher × day	Student
p-value of joint test of f.e.	0.000	0.000	0.000	0.000	0.000	0.000
Number of observations	24,192	24,192	24,192	24,192	22,729	22,729
R-squared	0.65	0.68	0.65	0.68	0.64	0.68

Note: This table displays estimates of Eq. (2). All observations are at the student-day level and are limited to days on which the child is absent. The dependent variable is whether or not the teacher marks the child as present despite the child's absence while the primary independent variables include various measures of the child's eligibility for the grain: "still eligible" is if the child is not yet disqualified, "needs help" is if the child needs further days of attendance to earn the grain and "already earned" is if the child has already qualified for the grain. The omitted group is those who are already disqualified by missing 20% or more of the total number of days in the month. Columns 5–6 exclude children who have already earned the grain, i.e. attended more than 80% of the month. All columns include month fixed effects. All columns estimate a linear probability model; similar results are obtained with conditional logit. See Section 6.2 for a detailed explanation of the results. Robust standard errors clustered by teacher are in parentheses. *** 1%, ** 5%, * 10%.

perfect past attendance, is positive and significant. One possible explanation is that teachers may be more willing to misrepresent attendance for children with perfect past attendance even earlier in the month.

All in all, the results suggest that teachers are more likely to misrepresent attendance when the incentives created by the grain program are higher, but they may not be very sophisticated in their decisions. It is likely the cost to recording attendance inaccurately is so low that teachers do not feel the need to economize their behavior precisely.

7. Do teachers assist certain students in getting the grain more than others?

Once we establish that teachers respond to the grain program, the next natural question is if they use their discretion to help certain students more than others. In Tables 2–4, we tested for the joint significance of all student fixed effects and found them to be significantly different from zero in all cases, suggesting that teachers misrepresent attendance for some children more than others. We explore this question by including demographic characteristics and interaction

Table 4
Empirical tests of whether teachers respond to grain incentives.

Dependent variable: misrepresent						
	(1)	(2)	(3)	(4)	(5)	(6)
Still eligible		0.178***				
		(0.026)				
Needs help			0.155***	0.166***	0.194***	0.219***
			(0.027)	(0.027)	(0.030)	(0.031)
Already earned				0.223***	0.267***	0.307***
				(0.041)	(0.064)	(0.066)
Needs help × Perfect attendance					0.072**	0.061*
					(0.035)	(0.034)
Already earned × Perfect attendance					0.051	0.038
					(0.070)	(0.070)
Days left	−0.002***	−0.007***	−0.007***	−0.007***	−0.007***	
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	
Days left × Still eligible		−0.001				
		(0.002)				
Days left × Needs help			−0.0001	−0.0002	−0.005**	−0.007***
			(0.002)	(0.002)	(0.002)	(0.002)
Days left × Already earned				−0.013	−0.004	−0.030
				(0.018)	(0.049)	(0.048)
Days left × Needs help × Perfect attendance					0.010***	0.011***
					(0.002)	(0.002)
Days left × Already earned × Perfect attendance					−0.010	0.007
					(0.051)	(0.051)
Additional fixed effects	Student	Student	Student	Student	Student	Days left, Student
p-value of joint test of f.e.	0.000	0.000	0.000	0.000	0.000	0.000
Number of observations	24,192	24,192	22,729	24,192	24,192	24,192
R-squared	0.68	0.68	0.69	0.68	0.7	0.7

Note: This table displays estimates of Eq. (3). All observations are at the student-day level and are limited to days on which the child is absent. The dependent variable is whether or not the teacher marks the child as present despite the child's absence while the primary independent variables include the number of days left in the month and various measures of the child's eligibility for the grain: "still eligible" is if the child is not yet disqualified, "needs help" is if the child needs further days of attendance to earn the grain and "already earned" is if the child has already qualified for the grain. The omitted group is those who are already disqualified by missing 20% or more of the total number of days in the month. Column 3 excludes children who have already earned the grain, i.e. attended more than 80% of the month. All columns include month fixed effects. All columns estimate a linear probability model; similar results are obtained with conditional logit. See Section 6.2 for a detailed explanation of the results. Robust standard errors clustered by teacher are in parentheses. *** 1%, ** 5%, * 10%.

Table 5
Empirical tests of whether teachers favor certain types of students.

Dependent variable	Monitored attendance	Child had ≥80% roster attendance			
	(1)	(2)	(3)	(4)	(5)
Monitored attendance		0.785***			
		(0.017)			
Monitored attendance <80%			−0.338***	−0.308***	−0.139***
			(0.012)	(0.042)	(0.040)
Muslim student in non-Urdu school	−0.041***	−0.010	−0.017*	0.002	
	(0.010)	(0.008)	(0.010)	(0.007)	
Brahmin	−0.010	−0.009	−0.014	0.006	
	(0.010)	(0.008)	(0.010)	(0.008)	
Kshatriya	−0.013	−0.003	−0.005	0.011	
	(0.009)	(0.007)	(0.008)	(0.008)	
Vaishnav	0.012	0.007	0.013	0.020*	
	(0.011)	(0.010)	(0.013)	(0.011)	
Unknown caste	−0.008	−0.008	−0.011	−0.012	
	(0.008)	(0.006)	(0.008)	(0.008)	
Male	−0.013***	−0.008***	−0.013***	−0.003	
	(0.003)	(0.003)	(0.003)	(0.003)	
Normalized pre-test score	0.027***	0.015***	0.021***	0.013***	
	(0.002)	(0.002)	(0.002)	(0.002)	
Monitored attendance <80% * Muslim student in non-Urdu school				−0.068**	−0.020
				(0.031)	(0.034)
Monitored attendance<80% * Brahmin				−0.087**	−0.146***
				(0.042)	(0.044)
Monitored attendance <80% * Kshatriya				−0.064	−0.090**
				(0.040)	(0.043)
Monitored Attendance<80% * Vaishnav				−0.035	−0.120*
				(0.064)	(0.069)
Monitored attendance <80% * Unknown caste				0.009	−0.068*
				(0.042)	(0.041)
Monitored attendance <80% * Male				−0.034***	−0.027**
				(0.009)	(0.012)
Monitored attendance <80% * Normalized pre-test score				0.029***	0.024***
				(0.006)	(0.008)
p-values from F-tests of:					
Student characteristics		0.000	0.000	0.000	
Student char. and interactions				0.000	0.001
Additional fixed effects	Teacher × month	Teacher × month	Teacher × month	Teacher × month	Student
p-value of joint test of f.e.	0.000	0.000	0.000	0.000	0.000
Number of observations	50,310	49,742	49,742	49,742	49,742
R-squared	0.152	0.448	0.335	0.338	0.562

Note: Column 1 in this table presents estimates from a regression of monitored attendance rates on student characteristics. Columns 2–5 displays estimates of Eq. (1) with student characteristics. All observations are at the student-month level. The dependent variable in columns 2–5 is whether or not the student was assigned to receive the grain from his or her teacher while the independent variable is the students' attendance record as measured through the direct monitoring of attendance. This table differs from Table 2 because of the inclusion of student characteristic controls and interactions of these characteristics with monitored attendance. All columns include student-age fixed effects and column 5 includes month fixed effects. All columns estimate a linear probability model; similar results are obtained with conditional logit. See Section 7 for a detailed explanation of the results. Robust standard errors clustered by teacher are in parentheses. Observations are only used if there were at least 3 dates matched between monitored and roster attendance data for a particular student-month. *** 1%, ** 5%, * 10%.

terms within Eq. (1) above. We include an indicator for a Muslim child in a non-Muslim school, dummy variables for caste, gender and age, and a normalized score on a test in mathematics and language administered at the beginning of the school year.[24]

It is important to note that any test of how student characteristics affect a teacher's decision to misrepresent attendance will potentially suffer from omitted variables bias since the characteristics are not independent of other factors correlated with a teacher's propensity to lie. This means, for example, that we cannot conclude that teachers are discriminating against particular types of students because they may instead be responding to an unobserved characteristic that happens to be correlated with the observable one. Similarly, one might also make the argument that because the government can also observe the analyzed characteristics, then the question of whether or not the teachers' actions contravene the intent of the program is moot — as the government has clearly chosen not to set incentives as a function of these characteristics. However, the characteristics observable to the teacher that are observable to neither the government nor us may, in fact, be correlated with these characteristics.

Following Table 2, we first include student characteristics in the empirical regressions using monthly data. In column 1 of Table 5, we present the estimates from a regression of monitored attendance rates on student characteristics simply to provide summary information on attendance patterns of students of different religions, castes and genders. Some results are as expected — children who do better on a test administered at the beginning of the school year have higher attendance rates subsequently. Other results are more surprising: boys are less likely to attend school than girls. Muslim children in non-Muslim schools have lower attendance rates than other students in their schools. All regressions in Table 5 also contain age dummies; while some of them are significant, there is no clear pattern.

[24] In results not presented here, we conduct similar tests using the other specifications in Section 4.2. Across these specifications, we find consistent evidence of heterogeneity by test scores and gender, and suggestive (but less consistent) evidence by caste and religion.

Columns 2–3 examine what factors affect whether a child earns the grain, according to roster attendance, conditional on measures of his monitored attendance rate. If the student's actual attendance rate was monitored perfectly for the entire month and if teachers were strictly following the 80% rule, student characteristics would not affect whether a child receives the grain. The p-values at the bottom of the table, however, demonstrate that student characteristics have significant predictive power. Conditional on monitored attendance, Muslim children (in non-Muslim schools) are less likely to receive the grain by about 1.7 percentage points, and boys are less likely to receive the grain by about 1.3. Students with high test scores are also more likely to receive the grain; a child one standard deviation above the mean is 2 percentage points more likely to receive the grain.

Interestingly, column 1 indicates that monitored attendance patterns mirror which children are more likely to receive the grain. If boys are less likely to attend school in general, teachers may not misrepresent attendance for boys as much because they believe it is unlikely these boys will receive the grain and not because of a preference for girls. To examine this possibility, we include interactions of student characteristics and the dummy variable for whether the child does not appear to merit the grain (column 4) and student fixed effects (column 5). Column 4 separates out students who appear to earn the grain and those who do not. Among those who did not deserve the grain (the interacted regressors), both boys and Muslim children are still less likely to earn the grain by 3.4 percentage points for boys and 6.8 for Muslim children in non-Urdu schools.[25] Children of higher castes are also less likely to receive the grain by about 6–9 percentage points. Finally, children with better pre-test scores are more likely to earn the grain even when they appear not to deserve it by about 2.4 percentage points per standard deviation.[26]

Column 5 includes student level fixed-effects. The results regarding boys and children with better test scores remain. The coefficients on higher castes are larger in magnitude and still significant. However, the coefficient on Muslim children is no longer significant. These results confirm that even when we compare 2 months of attendance for the same student, whether the student receives the grain when he or she is less likely to have earned it depends on the student's caste, gender, pre-test score and possibly religion.[27]

8. Conclusion

We assess teachers' use of their reporting discretion within the context of a conditional grain transfer program in Mumbai, India. Using a unique data set that allows us to observe both teachers' official attendance records as well as attendance taken by external monitors over a two-year period, we find strong evidence suggesting that teachers manipulate their attendance records in order to ensure that some ineligible students receive the grain. First, we are able to directly observe that teachers misrepresent absent students as present in their daily attendance records and that this results in ineligible students receiving the grain. Second, we demonstrate that the pattern of observed misrepresentations closely follows the incentives of the grain program, rather than simply a general incentive to overstate attendance. We also find that, conditional on actual attendance, teachers misrepresent attendance more for female students and for students who perform better on a test in mathematics and language at the beginning of the year. There is also some evidence that teachers misrepresent attendance more for children of lower castes and less for Muslim students.

The potential importance of this behavior can be illustrated using the grain program itself. The Central Government of India created the NSPE to improve the nutritional status of school children as well as to increase school attendance rates among poor children — goals shared by almost all conditional transfer programs. The teacher behavior we documented possibly improves the nutritional status of more children, but at the cost of reducing the incentives to attend school. Since teachers have more information about the individual children, however, allowing them some degree of discretion could be efficient. They may, for example, use local information to target children who are in greater need of the grain. They may also be able to award the grain to children in months in which the students' absences were beyond the child's control, thereby allocating more grain while preserving the incentives created by the original requirements. However, they may also be setting the incentives of the program too low or, at worst, discriminating against particular types of students. Either way, however, additional research is needed to determine the implications of local agent discretion and program designers should take such behavior into account.

[25] The fact that we find differential misreporting for girls and boys might suggest that teachers are not lying as a result of parents bribing or pestering teachers. Assuming parents send their daughters and sons to the same school, it is unlikely they bribe or harass the teacher for their daughters but not their sons. Similarly, as argued above, children for whom a teacher is seen to misrepresent attendance tend not to earn the grain in every month; one would expect those parents to continue to be the most likely to harass or bribe the teachers.

[26] The main variables indicate that among children who deserve the grain, those with higher pre-test scores and children of a middle caste are marginally more likely to receive it. This is surprising since teachers rarely mark present children absent, but it could be due to error in the measure of monitored attendance rates. Another explanation is that some of the exaggeration is not for the purpose of helping children receive the grain or that teachers miscalculate. However, both are accounted for in the fixed-effects specification in column 5.

[27] Regressing the fixed effects estimated in all these tables on the student characteristics themselves consistently supports the finding that teachers misrepresent attendance more for children with higher pre-test scores and for girls.

References

Alderman, Harold, 2002. Do local officials know something we don't? Decentralization of targeted transfers in Albania. Journal of Public Economics 83 (3), 375–404.

Banerjee, Abhijit, Cole, Shawn, Duflo, Esther, Linden, Leigh, 2007. Remedying education: evidence from two randomized experiments in India. Quarterly Journal of Economics 122 (3), 1235–1264.

Barrera-Osorio, Felipe, Bertrand, Marianne, Linden, Leigh, Perez, Francisco, 2011. Improving the Design of conditional transfer programs: evidence from a randomized education experiment in Colombia. American Economic Journal: Applied Economics 3 (2), 167–195.

Bird, Richard, Rodriguez, Edgard, 1999. Decentralization and poverty alleviation. International experience and the case of the Philippines. Public Administration and Development 19, 299–319.

Cheng, Ming-Yen, Fan, Jianqing, Marron, James S., 1997. On automatic boundary corrections. The Annals of Statistics 25 (4), 1691–1708.

Coudouel, Aline, Marnie, Sheila, Micklewright, John, 1998. Targeting social assistance in a transition economy: the role of the Mahalla in Uzbekistan. Innocenti Occasional Papers, Economic and Social Policy Series no. 63. UNICEF International Child Development Centre, Florence, Italy.

Das, Jishnu, 2004. Equity in educational expenditures: can government subsidies help? Policy Research Working Paper Series 3249. The World Bank.

Fernald, Lia, Gertler, Paul, Neufeld, Lynnette, 2008. Role of cash in conditional cash transfer programmes for child health, growth and development: an analysis of Mexico's *Oportunidades*. The Lancet 371 (9615), 828–837.

Jacob, Brian, Levitt, Steven, 2003. Rotten apples: an investigation of the prevalence and predictors of teacher cheating. Quarterly Journal of Economics 118 (3), 843–877.

Klugman, Jeni, 1997. Decentralization: a survey from a child welfare perspective. Innocenti Occasional Papers, Economic and Social Policy Series no. 61. UNICEF International Child Development Centre, Florence, Italy.

Martinelli, César, Parker, Susan, 2009. Deception and misreporting in a social program. Journal of the European Economic Association 7 (4), 886–908.

Mumbai Interviews, 2005. Interviews with School Teachers.

Paxson, Christina, Schady, Norbert, 2007. Does money matter? The effects of cash transfers on child health and cognitive development in rural Ecuador. Policy Research Working Paper Series 4226. The World Bank.