

Contract Teachers: Experimental Evidence from India

Karthik Muralidharan[†]
Venkatesh Sundararaman[‡]

20 October 2013^{*}

Abstract: The large-scale expansion of primary schooling in developing countries has led to the increasing use of non-civil-service contract teachers who are hired locally by the school, are not professionally trained, have fixed-term renewable contracts, and are paid much lower salaries than regular civil-service teachers. This has been a controversial policy, but there is limited evidence on the effectiveness of contract teachers in improving student learning. We present experimental evidence on the impact of contract teachers using data from an 'as is' expansion of contract-teacher hiring across a representative sample of 100 randomly-selected government-run rural primary schools in the Indian state of Andhra Pradesh. At the end of two years, students in schools with an extra contract teacher performed significantly better than those in comparison schools by 0.16σ and 0.15σ , in math and language tests respectively. Contract teachers were also much less likely to be absent from school than civil-service teachers (18% vs. 27%). Combining the experimental reduction in school-level pupil-teacher ratio (PTR) induced by the provision of an extra contract teacher, with high-quality panel data estimates of the impact of reducing PTR with a regular civil-service teacher, we show that contract teachers are not only effective at improving student learning outcomes, but that they are no less effective at doing so than regular civil-service teachers who are more qualified, better trained, and paid five times higher salaries.

JEL Classification: I21, H44, J24, J33, J45, M55, O15

Keywords: contract teachers, decentralization, publicly produced goods, segmented labor markets, teacher pay, teacher quality, class size, primary education in developing countries, India

[†] UC San Diego, NBER, BREAD, and J-PAL; E-mail: kamur@ucsd.edu (Corresponding author)

[‡] South Asia Human Development Unit, World Bank. E-mail: vsundararaman@worldbank.org

^{*} We thank Eli Berman, James Berry, Prashant Bharadwaj, Jeff Clemens, Julie Cullen, Gordon Dahl, Jishnu Das, Xavi Gine, Paul Glewwe, Nora Gordon, Roger Gordon, Namita Gupta, Gordon Hanson, Caroline Hoxby, David Lagakos, Victor Lavy, Michael Kremer, Paul Niehaus, Chris Woodruff, and various seminar participants for useful comments and discussions.

This paper is based on an experiment carried out as part of the Andhra Pradesh Randomized Evaluation Studies (AP REST), which is a partnership between the Government of Andhra Pradesh, the Azim Premji Foundation, and the World Bank. Financial assistance for the project has been provided by the Government of Andhra Pradesh, the UK Department for International Development (DFID), the Azim Premji Foundation, and the World Bank. We thank Dileep Ranjekar, Michelle Riboud, Amit Dar, Samuel C. Carlson, and officials of the Department of School Education in Andhra Pradesh (particularly Dr. I.V. Subba Rao, Dr. P. Krishnaiah, and K. Ramakrishna Rao), for their continuous support and long-term vision for this research. We are especially grateful to DD Karopady, M Srinivasa Rao, and staff of the Azim Premji Foundation for their leadership and meticulous work in implementing this project. We thank Vinayak Alladi, and Ketki Sheth for outstanding research assistance. The findings, interpretations, and conclusions expressed in this paper are those of the authors and do not necessarily represent the views of the Government of Andhra Pradesh, the Azim Premji Foundation, or the World Bank.

1. Introduction

The large scale expansion of primary schooling in developing countries over the past two decades has led to significant improvements in school access and enrollment, but has also created difficulties in maintaining and improving school quality, with learning levels being very low in most developing countries (Pritchett 2004; Pratham 2012, Uwezo 2012). A particularly challenging problem has been that of recruiting enough teachers and staffing them where needed. The challenges include a lack of enough qualified teachers to match the needs of rapidly-expanding school systems, the high cost of hiring them, and the reluctance of qualified teachers to serve in rural areas where the needs of the expanding education system are the greatest.¹

Governments in several developing countries have responded to this challenge by staffing teaching positions with locally-hired teachers on fixed-term renewable contracts, who are not professionally trained, and who are paid much lower salaries than those of regular teachers - often less than one fifth as much (Table 1). The growing use of contract teachers in public schools has been one of the most significant trends in providing primary education in developing countries in the last two decades.² Contract teachers comprise a third of public-school teachers across twelve countries in Africa (Bourdon et al. 2010) and their share among all public-school teachers in India grew from 6 percent in 2003 to 30 percent in 2010 (Muralidharan et al. 2013).

But, the use of contract teachers has been and remains highly controversial. Supporters consider the use of contract teachers to be an efficient way of improving the quality of rural schools while also providing employment to educated young people in rural areas, and argue that contract teachers face superior incentives compared to tenured civil-service teachers. Opponents argue that using under-qualified and untrained teachers may staff classrooms but will not produce learning outcomes, that contract teacher positions may be 'captured' by local elites, that the wide differences in pay and benefits are exploitative, and that contract teachers should be replaced with qualified, trained, and well-paid regular teachers (Kumar 2005; Raina 2006, 2009).

¹ See Duthilleul (2005) for a discussion of how the lack of qualified teachers has been a key challenge in expanding access to schooling in several countries. Teacher salaries comprise ~70-90% of education spending in most developing countries, and the cost of salaries is typically the main fiscal constraint to education expansion (see <http://www.uis.unesco.org/Education> for country-level salary data). Fagernas and Pelkonen (2012) study teacher location preferences and show that more qualified teachers are less likely to desire rural postings.

² Contract teacher schemes have been used in several developing countries including Benin, Burkina Faso, Cambodia, Cameroon, Chad, Congo, Indonesia, Kenya, Madagascar, Mali, Nicaragua, Niger, Senegal, and Togo, (see Duthilleul 2005 and Bourdon et al. 2010 for reviews of contract teacher programs). They have also been widely employed in several states of India, under different titles (see Govinda and Josephine 2004 for a descriptive review).

The perception that using contract teachers contributes to poor school quality in developing countries, and that their use should be eliminated by investing in teacher training and raising education spending to hire qualified civil-service teachers, is deeply embedded in education policy discourse in several developing countries.³ However, while there are several descriptive and observational studies on contract teachers, there is relatively little well-identified evidence on the impact of using contract teachers in primary schools on learning outcomes. The two main questions of interest are (i) whether untrained contract teachers can improve learning outcomes, and (ii) how effective they are relative to qualified civil-service teachers.

In this paper, we answer both these questions using data from an experiment in the Indian state of Andhra Pradesh (AP) that reduced school-level pupil-teacher ratios (PTR) by allowing schools to hire an additional contract teacher. The study was conducted across a representative sample of 200 government-run schools in rural AP with 100 of these schools being selected by lottery to receive an extra contract teacher over and above their usual allocation of teachers.

At the end of two years of the program, we find that students in schools with an extra contract teacher perform significantly better than those in comparison schools by 0.16 and 0.15 standard deviations (σ) in math and language tests respectively. We find no heterogeneity in impact by baseline test scores, suggesting that the gains are broadly distributed among all students. However, students in remotely-located schools benefit more from their school receiving an extra contract teacher. We also find that contract teachers were significantly less likely to be absent from school than regular teachers (18% versus 27%).

Since our experiment induced an exogenous reduction in school-level pupil-teacher ratio (PTR), we can construct an experimental estimate of the impact of reducing PTR with a contract teacher. We estimate that reducing PTR by 10% using a contract teacher would improve mean test scores in the school, across subjects, by $0.03\sigma/\text{year}$.⁴ We also estimate the impact of reducing PTR with a regular civil-service teacher, using five years of panel data from the control schools and estimate that reducing PTR by 10% using a regular civil-service teacher would improve mean test

³ The recently passed 'Right to Education' (RtE) Act in India reflects this thinking and requires the use of untrained teachers to be phased out over a three-year period. In a similar vein, the Indonesian 'Teacher Law' passed in 2005 required all teachers to get certified and offered a *doubling* of salary for certified teachers (Jalal et al. 2009).

⁴ As we discuss in section 3.3, reducing PTR with an extra teacher can help improve learning through multiple channels including a reduction in class size, a reduction in multi-grade teaching, and a reduction in the pedagogical cost of teacher absence (since another teacher is available in the school). However, since decisions on *how* to use teachers within a school are endogenous, we focus our analysis on the impact of reducing PTR, which is defined at the school level, and which is also the main metric used to determine appropriate levels of teacher staffing in education policy in India (as opposed to class size).

scores across subjects by 0.02σ /year. The rich panel data allow us to use value-added estimates as well as school fixed effects (eliminating concerns of cross-sectional omitted variables). We show that these results are extremely robust, and yield perhaps the most credible estimates of the impact of PTR reduction with regular teachers in developing countries. Thus, our two main results are that: (i) untrained contract teachers are able to significantly improve learning outcomes in primary schools and (ii) they are at least as effective at doing so as regular civil-service teachers who are more educated, have formal teacher training credentials, and are paid over five times more.

To place these results in a broader labor market context, we examine the market for rural *private* school teachers in the same districts and find that they are paid even less than contract teachers (though they are more educated) and that their salaries are around *one-eighth* that of regular civil-service teachers. The results on equal effectiveness of contract and regular teachers and the market salary benchmarks for private school teachers strongly suggest that the large wage differential between regular and contract teachers is unlikely to reflect differences in productivity and mostly represents rents accruing to unionized civil-service teachers.

Our results contribute to an emerging literature on understanding the impact of contract teachers in developing countries. In addition to several descriptive and observational papers,⁵ there are two relevant experimental studies to date. Duflo et al (2012) find significant positive effects on grade one test scores from an experimental evaluation of a program in Kenya that reduced class size in grade one using an extra contract teacher. However, Bold et al (2013) experimentally evaluate a similar contract teacher program in Kenya and find that the program had a positive effect when implemented by a non-profit organization (as in the program studied by Duflo et al. 2012) but had no effect when implemented by the government of Kenya.

Our school-level randomization and analysis allow us to make several advances over the existing literature. First, unlike the other experiments, we do not require schools to have fully complied with a within-school randomization protocol to generate unbiased estimates of the impact of adding a contract teacher (this can be difficult to ensure as shown in Bold et al. 2013). Second, our estimates do not depend on the assumption that no other resources in the school were reallocated in response to the provision of an extra contract teacher to one grade. Third, the

⁵ Descriptive papers on contract teachers include Duthilleul (2005), Govinda and Josephine (2004), Pritchett and Pande (2006), Kingdon and Sipahimalani-Rao (2010), and Goyal and Pandey (2011). Papers that use observational data to study the effect of contract teachers include De Laat and Vegas (2005) in Togo, Bourdon et al (2010) in Niger, Mali, and Togo, and Atherton and Kingdon (2010) in India.

literature to date (both observational and experimental) is based on comparisons of students *assigned* to be taught by contract and regular teachers, and cannot account for the bias from unobserved factors that may affect actual teaching *practice*.⁶ This is not a concern for us, since all our analysis is aggregated at the school level. Finally, we evaluate an 'as is' expansion of contract teacher hiring, implemented by the AP Government in exactly the same way as a regular expansion, using random assignment across a representative sample of schools in a state of 80 million people, and thereby provide estimates of program impact that are directly applicable to scaling up on the current policy margin (mitigating concerns of external validity as discussed by Heckman and Smith 1995, Bold et al. 2013, and Pritchett and Sandefur 2013).

In addition to its immediate policy relevance, this paper also contributes to several strands of academic literature. First, there is an old debate on the merits of restricting entry into professions based on credentials (Smith 1776, Friedman and Kuznets 1945), but there is limited evidence on the impact of such restrictions due to a combination of data limitations and identification challenges (Kleiner 2000; Kleiner and Krueger 2013). Our results suggest that formal teaching credentials are poor predictors of primary school teacher effectiveness in India, and highlight the large potential costs of laws that restrict entry into teaching based on these credentials.⁷

Second, while set in the context of teachers and schools, our results also illustrate the polarized two-tier labor market in many developing countries. In particular, the very large differences in pay and benefits between civil-service and contract teachers with no corresponding difference in productivity provides a striking illustration of the 'insider-outsider' theory of two-tier labor markets (Lindbeck and Snower 1988, 2001) and the extent to which politically powerful insiders are able to protect their rents (Kingdon and Muzammil 2001).

Third, the literature on decentralization has highlighted the possibility that locally-hired staff positions may be 'captured' by local elites and become patronage jobs with high absence rates, and limited impact on outcomes (Bardhan 2002, Bardhan and Mookherjee 2000). Our results show that this was not a first-order concern in this setting since an 'as is' decentralization to

⁶ For instance, the effectiveness of contract teachers may be over-stated if regular teachers spend time away from their own students to coach and mentor contract teachers. Conversely, the effectiveness of contract teachers may be understated if they are used to staff the classrooms of regular teachers who are more likely to be absent.

⁷ In related work in the US, Bettinger and Long (2010) show that adjunct faculty (who are less qualified and paid much less than tenure-track faculty) teach slightly better than regular faculty, and Figlio, Schapiro, and Soter (2013) show that students learn more from non-tenure track line faculty than tenure-track faculty in introductory courses. While these results are related to ours, they are less stark and do not imply inefficiency in the status quo because the job description of a tenure track faculty member also includes a substantial focus on research.

school committees of resources to hire contract teachers led to (a) the hiring of teachers who had lower absence rates than regular teachers, (b) significantly improved primary-school learning outcomes, and (c) did so much more cost effectively than the default of civil-service hiring (one rationale for which is to 'professionalize' hiring and to insulate it from pressure by local elites).

Fourth, governments around the world face severe staffing challenges in providing health and education services in rural areas, and try to assign qualified teachers and doctors to rural postings as part of their job rotations. However, since well-educated service providers prefer to live in urban areas, such a policy often results in an internal 'market' for jobs where substantial bribes are paid for more desirable postings (Beteille 2009), and/or a situation where service providers choose to live in further-away towns and commute to the village, which is correlated with higher rates of provider absence (Chaudhury et al. 2006). Our results suggest that hiring local staff from the same village, even with lower levels of education and modest amounts of training, may be a more promising approach to service delivery in rural areas (especially for *primary* services). Our results are also consistent with findings in the health literature, which show that community health workers with limited training can be effective at improving health outcomes in underserved rural areas in developing countries (Bang et al. 1999, Haines et al. 2007).

There are large welfare implications of taking our results seriously. The Government of India is expected to spend an additional \$5 billion/year to fulfill the mandate of the recently passed 'Right to Education' Act (RtE) to reduce pupil-teacher ratio (PTR) from 40:1 to 30:1 by recruiting additional regular teachers. Since it is possible to hire *several* local contract teachers for the cost of a regular civil-service teacher, our results suggest that a substantially larger improvement in education outcomes may be obtained by allocating these funds to hiring more contract teachers at the current margin. Doing so would make it possible to reduce PTR to below 15, to eliminate multi-grade teaching, and to have additional teaching resources to provide supplemental instruction to first-generation learners who are not able to keep pace with the syllabus, for the *same* cost (see Muralidharan 2013a, and 2013b for details).

The rest of this paper is organized as follows: section 2 describes the experimental intervention, and data collection; section 3 presents the results of the extra contract teacher program; section 4 compares the effectiveness of regular and contract teachers; section 5 discusses the results in the larger context of teacher labor markets in rural India; section 6 concludes with a discussion of policy implications, caveats, and extensions to our results.

2. Experimental Design

2.1. Context

While India has made substantial progress in improving access to primary schooling and primary school enrollment rates, the average levels of learning remain very low. Recent surveys show that over 60% of children aged 6 to 14 in rural India could not read at the second grade level, though over 97% of them were enrolled in school (Pratham 2012). Public spending on education has been rising as part of the "Education for All" campaign, but there are substantial inefficiencies in public delivery of education services. Kremer et al (2005) found using a nationally representative survey (conducted in 2003) that 26% of teachers in rural public schools in India were absent on any given day. A more recent survey that revisited the same villages found substantial improvements in school quality as measured by inputs, but found that teacher absence rates in rural India were still around 24% (Muralidharan et al. 2013).

Andhra Pradesh (AP) is the 5th largest state in India, with a population of over 80 million, with around 70% living in rural areas. AP is close to the all-India average on various measures of human development such as gross enrollment in primary school, literacy, and infant mortality (Muralidharan and Sundararaman 2011), and on measures of service delivery such as teacher absence (Kremer et al. 2005). There are a total of over 60,000 government primary schools in AP and around 70% of children in rural AP attend government-run schools (Pratham 2012).

The average rural primary school is quite small, with total enrollment of around 80 to 100 students and an average of 3 teachers across grades one through five. All regular teachers are employed by the state, and their salary is mostly determined by experience and rank, with minor adjustments based on assignment location, but no component based on any measure of performance. In 2006, the average salary of regular teachers was over Rs. 8,000/month and total compensation (including benefits) was over Rs. 10,000/month (per capita income in AP was around Rs. 2,000/month). Regular teachers' salaries and benefits comprise over 90% of non-capital expenditure on primary education in AP. Teacher unions are strong and disciplinary action for non-performance is rare (Kingdon and Muzzamil 2001; Kremer et al. 2005).

2.2 The Extra Contract Teacher Intervention

Contract teachers (also known as para-teachers in India) are hired at the school level by school committees and have usually completed either high school or college but typically have no formal teacher training. Their contracts may be renewed annually and they are not protected

by any civil-service rules. Their typical salary of around Rs. 1000 - 1500/month is less than one fifth of the average salary of regular government teachers.⁸ They are more likely to be younger, to be female, to be from the same village, and live closer to the school (Table 1 – Panel A). Since most primary schools have more grades than teachers, the default teaching arrangement is one of 'multi-grade' teaching (where one teacher simultaneously teaches multiple grades in the same classroom, and also teaches all subjects). As a result, contract teachers almost always teach their own classes and are not 'teacher-aides' who support a regular teacher in the same classroom.

The process by which contract teachers are usually hired in Andhra Pradesh is that schools apply to the district education administration for permission to hire a contract teacher based on their enrollment and teacher strength at the start of the school year. Thus, contract teachers can be appointed both against vacant sanctioned posts (that should have been filled by a regular teacher but have not been filled) and as additional resources to meet the needs of growing enrollment. If the permission (and fiscal allotment) is given, a contract teacher will be hired by the school committee. The authorization of the position is not guaranteed for subsequent years, and since renewal is not guaranteed, the appointment of contract teachers is typically for a 10-month period. New hires are supposed to go through a brief accelerated training program prior to starting to teach, but this is imperfectly implemented in practice. While some states have a system of providing contract teachers with priority for being hired when openings for regular teacher positions are available, there is no such system in place in Andhra Pradesh.

The extra contract teacher intervention studied in this paper was designed to resemble the typical process of contract teacher hiring and use as closely as possible. Schools that were selected for the program by a lottery were informed in a letter from the district administration that they had been authorized to hire an additional contract teacher, and that they were expected to follow the same procedures and guidelines for hiring a contract teacher as they would normally do. Most schools (~80%) reported starting the process of hiring the extra contract teacher within a week of receiving the notification and reported having appointed the selected candidate within a month of receiving the notification. All schools reported having completed the full process within two months.

⁸ The salary of contract teachers was Rs. 1,000/month in the first year of the project (2005 – 06) and was raised to Rs. 1,500/month in the second year (2006 – 07). Unless stated otherwise, all the numbers will refer to those at the time of the study (2005-07 for the most part). The factor of five is a lower bound on the ratio of the total compensation because it does not include the value of medical and retirement benefits paid to regular teachers.

The modal selection committee consisted of three members and typically comprised of the head teacher, a member of the local elected body, and another teacher. The most important stated criteria for hiring the contract teacher was qualification (62%), followed by being from the same village and extent of teaching experience (20% each). In spite of these stated responses, it is possible that local leaders or head teachers may have exercised some favoritism towards their preferred candidates in some cases (as is described by Srivastava 2010 in the context of para-teacher hiring in the state of Madhya Pradesh, and by Duflo et al. 2012 in Kenya). However, as per the experimental protocol, the project staff were not a part of the teacher hiring process in any way, and so the intervention mimicked an 'as is' expansion of the existing contract teacher program in AP to 100 randomly selected schools. We see in Table 1 (Panel B) that the additional contract teachers hired under this program had the same average characteristics as typical contract teachers in the control schools.

The additional contract teachers were allocated to the school and not to a specific grade or pre-specified role, which is also how teachers (regular and contract) are typically allocated to primary schools. We chose to not try to allocate the extra contract teacher to a specific role or grade within the school,⁹ because it was not possible in practice to ensure adherence to a within-school randomization protocol, and it is quite likely that existing teachers would re-optimize their roles and effort in response to the extra contract teacher. It is also likely that the optimal assignment of contract teachers to specific grades and tasks will vary based on unobservable school and teacher-level characteristics, and trying to randomly pre-assign the task for the extra contract teacher may be sub-optimal (and hence less likely to be adhered to). We therefore abstract away from within-school considerations and aggregate our analysis at the school level, which is also the level at which policy decisions on teacher staffing are made and implemented.

2.3. Sampling and Randomization

The extra contract teacher (ECT) program was evaluated as part of a larger education research initiative (across 500 schools) known as the Andhra Pradesh Randomized Evaluation Studies (AP RESt), with 100 schools being randomly assigned to each of four treatment and one control groups.¹⁰ We sampled 5 districts across each of the 3 socio-cultural regions of AP in

⁹ The main value of a within-school randomization protocol would have been to potentially obtain an experimental comparison of the relative effectiveness of regular and contract teachers (as attempted by Duflo et al. 2012).

¹⁰ The AP RESt is a partnership between the government of AP, the Azim Premji Foundation (a leading non-profit organization working to improve primary education in India), and the World Bank, that piloted and evaluated

proportion to population. In each of the 5 districts, we randomly selected one administrative division and then randomly sampled 10 mandals (the lowest administrative tier) in the selected division. In each of the 50 mandals, we randomly sampled 10 schools using probability proportional to enrollment. Thus, the universe of 500 schools in the study was representative of the schooling conditions of the typical child attending a public primary school in rural AP.

The school year in AP starts in mid-June, and baseline tests were conducted in the 500 sampled schools during late June and early July, 2005.¹¹ After the baseline tests were scored, 2 out of the 10 project schools in each mandal were randomly allocated to one of 5 cells (four treatments and one control). Since 50 mandals were chosen across 5 districts, there were a total of 100 schools (spread out across the state) in each cell. The analysis in this paper is based on the 200 schools that comprise the 100 schools randomly chosen for the ECT program and the 100 that were randomly assigned to the control group. Table 2 (Panel A) shows summary statistics of baseline school and student characteristics for treatment and control schools, and the null of equality across treatment groups cannot be rejected for any of the variables.

2.4. Data

The data used in this paper consist of independent learning assessments in math and language (Telugu) conducted at the beginning of the study, and at the end of each of the two years of the experiment. We also use data from regular unannounced visits to the schools made by staff of the Azim Premji Foundation to measure teacher attendance and activity (six visits were made to each school in the first year and four in the second). For the rest of this paper, Y1 and Y2 refer to the first and second year of the two-year experiment. $T(Y_0)$ refers to the baseline tests in June-July 2005; $T(Y_1)$, and $T(Y_2)$ refer to tests conducted at the end of the first and second year of the program in March-April of 2006, and 2007 respectively. All analysis is carried out with normalized test scores, where individual test scores are converted to z-scores by normalizing them with respect to the distribution of scores in the control schools on the same test.

several policy options to improve learning outcomes in AP. While the Azim Premji Foundation (APF) was the main implementing agency for the teacher performance-pay, and school block-grant interventions (described in Muralidharan and Sundararaman 2011 and in Das et al. 2013), the contract teacher intervention was implemented by the AP Government (based on the lottery conducted by the project team). APF was responsible for data collection in all sampled schools, which was identical across all treatment and control schools. This ensures that Hawthorne effects are minimized and that a comparison between treatment and control schools can accurately isolate the treatment effect of interest. The only difference in data collection across treatment schools was a short treatment-specific process survey *after* a full year (that in this case documented the process of hiring the contract teacher).

¹¹ The selected schools were informed by the government that an external assessment of learning would take place in this period, but there was no communication to any school about any of the treatments at this time.

3. Experimental Results

3.1. Teacher and Student Turnover and Attrition

Regular civil-service teachers in AP are transferred once every three years on average. While this could potentially bias our results if more teachers chose to stay in or tried to transfer into the ECT schools, it is unlikely that this was the case since the treatments were announced in August '05, while the transfer process typically starts earlier in the year. There was no statistically significant difference between the treatment and comparison groups in the extent of teacher turnover in Y1, and the turnover rate was close to 33%, which is consistent with rotation of teachers once every 3 years (Table 2 – Panel B, rows 11-12).

As part of the agreement between the Government of AP and the Azim Premji Foundation, the Government agreed to minimize transfers into and out of the sample schools for the duration of the study. The average teacher turnover in Y2 was much lower, and there was no significant difference in teacher transfer rates across the treatment and control groups at the end of 2 years (Table 2 – Panel B, rows 13 - 14).¹² The average student attrition rate in the sample (defined as the fraction of students in the baseline tests who did not take a test at the end of each year) was 7.4% and 20.4% in Y1 and Y2 respectively, but there is no significant difference in attrition across the treatment and control groups (rows 15 and 18). Attrition is higher among students with lower baseline scores, but this is true across all treatments, and we find no significant difference in mean baseline test score across treatment categories among the students who drop out from the test-taking sample (Table 1 – Panel B, rows 16, 17, 19, 20). All our estimates of treatment effects include controls for baseline test scores, which further mitigates any concern of bias from attrition.

3.2. Specification

Our default specification uses the form:

$$T_{ijkm}(Y_n) = \alpha + \gamma_j \cdot T_{ijkm}(Y_0) + \delta \cdot ECT_k + \beta_z \cdot Z_m + \varepsilon_k + \varepsilon_{jk} + \varepsilon_{ijk} \quad (1)$$

The dependent variable of interest is T_{ijkm} , which is the normalized test score on the specific test, where i, j, k, m denote the student, grade, school, and mandal respectively. Y_0 indicates the baseline tests, while Y_n indicates a test at the end of n years of the program (the experiment ran

¹² There was also a court order to restrict teacher transfers in response to litigation complaining that teacher transfers during the school year were disruptive to students. This may have also helped to reduce teacher transfers during the second year of the project.

for two years). Including the normalized baseline test score improves efficiency due to the high autocorrelation of test-scores over time.¹³ All regressions include a set of mandal-level dummies (Z_m) (the unit of stratification) and the standard errors are clustered at the school level. We also run the regressions with and without controls for household and school variables.

ECT_k is a dummy variable at the school level indicating if school k was selected to receive the extra contract teacher (ECT) program, and the parameter of interest is δ , which is the effect on the normalized test scores of being in an ECT school. The random assignment of treatment ensures that the 'ECT' variable in the equation above is not correlated with the error term, and the estimate of the one-year and two-year treatment effects are therefore unbiased.

3.3. Impact of ECT program on Test Scores

Averaging across both math and language, students in program schools scored 0.10σ and 0.16σ higher than those in comparison schools at the end of one and two years of the program respectively (Table 3 – Panel A, columns 1 and 3). The two-year effects of the extra contract teacher are similar across math (0.16σ) and language (0.15σ) (Panels B and C). The addition of school and household controls does not significantly change the estimated value of δ , confirming the validity of the randomization (columns 2 and 4).

Table 4 explores the potential mechanisms of program impact by showing the changes induced by the provision of the extra contract teacher along various dimensions of the education experience of students in treatment schools (by each year of the program). Students in treatment schools had significantly lower pupil-teacher ratios (PTR) than those in control schools (25.4 vs. 32.7 in Y1 and 28.3 vs. 37.8 in Y2), whereas in the absence of the treatment the PTR would have been 35.7 vs. 32.7 in Y1 39.2 vs. 37.8 in Y2.¹⁴ Overall, the program reduced the PTR in treatment schools by 10.3 in Y1, and by 10.8 in Y2. Students in treatment schools also have lower average class sizes (28.0 vs. 33.7 in Y1, and 31.7 vs. 39.1 in Y2), and are less likely to be

¹³ Since grade 1 children did not have a baseline test, we set the normalized baseline score to zero for these children (similarly for children in grades 1 and 2 at the end of two years of the treatment). Note that we allow the coefficient on the lagged test score (γ_j) to vary by grade. This facilitates precision by allowing us to include baseline test scores (for the grades where we have them) without having to drop the grades where we do not have them.

¹⁴ Note that the figures for Y1 here (35.7 vs. 32.7) are different from the PTR of 39.8 and 39.5 shown in Table 2. The figures in Table 2 are based on the start of the school year (and are based on potentially inflated enrollment numbers). The figures in Table 4 are based on student enrollment figures that were verified by the APF field team over the course of the year. The difference between 35.7 and 32.7 is not significant in Table 4, but the difference in magnitude (relative to near identical figures in Table 2) suggests that it is possible that treatment schools may not have received all the teachers that they may have in the absence of the program. To the extent that this is true, the estimates in Table 3 will be a lower bound on the impact of providing an extra contract teacher to a school.

taught in a multi-grade classroom (50% vs. 64% in Y1 and 51% vs. 59% in Y2), with all these differences being significant at the 1% level.¹⁵

Further, when teachers are absent, it is common for grades to be combined (even more than the assigned level of multi-grade teaching). Thus, having an additional teacher in the school can mitigate the pedagogical challenges induced by teacher absences, by increasing the number of teaching days when at least two teachers are present. Based on direct physical verification of teacher attendance during unannounced visits (see Table 7 for summary statistics), we see that treatment schools were significantly more likely to have a minimum of two teachers present on any given day (74% vs. 68% in Y1 and 87% vs. 73% in Y2). Similarly, treatment schools were also significantly more likely to have a minimum of three teachers present on any given day (56% vs. 42% in Y1 and 58% vs. 41% in Y2). Finally, students in treatment schools are significantly more likely to have a contract teacher as their class teacher (27% vs. 7% in Y1 and 31% vs. 8% in Y2), which is consistent with contract teachers typically having full responsibility for a class, and with treatment schools having a significantly higher fraction of contract teachers.

However, while Table 4 provides a useful summary of the potential mechanisms of impact of the extra contract program, these are clearly endogenous at the school level since treatment schools had full flexibility to decide how to deploy the extra contract teacher (as they would in any practical situation where the school received an extra teacher). As a result, we conduct all our test-score analysis aggregated at the school level, and abstract away from any within-school comparisons (or comparisons at the grade-level between treatment and control schools).

3.4. Impact of Reducing PTR with a Contract Teacher

While it is not feasible to credibly identify the specific mechanism of program impact in cases where these mechanisms are endogenous to school decisions on how to allocate teacher quality and quantity across students and grades,¹⁶ our school-level randomization allows us to

¹⁵ The most common way in which treatment schools used the extra contract teacher was to split an existing multi-grade classroom and assign one of the grades to the extra contract teacher. This resulted in a reduction in class size for the students in the grade assigned to the contract teacher, and also for students in the other grades who would have previously been combined with students from this grade. Clearly, this would also lead to a reduction in the fraction of grades experiencing a multi-grade teaching environment in treatment schools. Note that the typical rural primary school in AP did not have a spare classroom for the extra teacher. In practice, if there are more teachers than classrooms, the classrooms are usually divided into 'teaching areas' with students assigned to different teachers facing in different directions and each teacher using his or her own blackboard.

¹⁶ Since we have only one source of exogenous variation (the experimental provision of the contract teacher) and several endogenous mechanisms of impact, an instrumental variable strategy using the random assignment would not credibly satisfy the exclusion restriction for any of the individual mechanisms summarized in Table 4.

experimentally estimate the impact of reducing school-level PTR using a contract teacher. This is precisely the parameter that matters for policy because policy norms on teacher staffing are based on a target school-level PTR and not based on intra-school targets.¹⁷ In particular, since the Right to Education Act calls for reducing school-level PTR from 40:1 to 30:1 by hiring more regular teachers, it is especially relevant to estimate the impact of doing so with contract teachers instead.

We re-estimate the treatment effects shown in Table 3 in terms of PTR by estimating:

$$T_{ijkm}(Y_1) = \alpha + \gamma_j \cdot T_{ijkm}(Y_0) + \beta_1 \cdot \Delta \log_PTR_k + \varepsilon_k + \varepsilon_{jk} + \varepsilon_{ijk} \quad (2)$$

where $\Delta \log_PTR_k$ is the change in \log_PTR (at the school-level) induced by the provision of the extra contract teacher ($\Delta \log_PTR_k$ is zero for the control schools as seen in Table 4). Since the extra teacher is randomly assigned across schools, this change in \log_PTR will be uncorrelated with other determinants of test scores, and β_1 provides an unbiased estimate of the average impact on test scores of reducing PTR with a contract teacher. However, as we see in Table 4, there is a small difference between the initial PTR (and \log_PTR) in treatment and control schools (though not significant). So we estimate (2) both without and with a control for the initial \log_PTR , and our preferred specification includes the control for initial \log_PTR .

We present these results in Table 5, and see that β_1 is equal to -0.34 (averaged across subjects, and controlling for initial \log_PTR) and is significant at the 1% level. Thus, reducing PTR by 10% using a contract teacher would improve mean test scores across subjects and grades by 0.034σ per year. Mechanically, we can see that this is just a decomposition of the mean first year treatment effect of 0.102σ estimated in Table 3 (Panel A – Column 1) into the mean reduction of \log_PTR induced by the treatment in the first year of 0.35 (see Table 4) multiplied by the estimate of β_1 of 0.031σ (from Table 5 – column 1, without controlling for initial \log_PTR as done in Table 3).¹⁸ However, expressing the mean impact of the program (presented

¹⁷ As seen in Table 4, reductions in PTR will be highly correlated with reductions in intra-school indicators like class size and multi-grade teaching. However, PTR can be exogenously manipulated by policy (and by our experiment) whereas variation in class-size and multi-grade teaching will typically also be determined by school and teacher-level unobservable factors that policy-makers cannot observe (and we cannot control for).

¹⁸ Note that the specification in (2) can only be estimated experimentally with the first year of the experiment, because doing so in a value-added specification using data from the second year will require including test scores from the end of the first year on the right hand side of (2), which will not yield an experimental estimate. Also, we do not use the experimental provision of an extra contract teacher as an instrumental variable for PTR because we wish to distinguish between the effects of changes in PTR induced by the provision of contract and regular teachers.

in Table 3) in terms of the marginal impact of reducing school-level PTR using an extra contract teacher allows for a more direct comparison with the impact of doing so with a regular civil-service teacher, and we return to this estimate in section 4.1.

3.5. Heterogeneous treatment effects by other school/student characteristics

We test for heterogeneity of the ECT program effect across student, and school characteristics by testing if δ_3 is significantly different from zero in:

$$T_{ijkm}(Y_n) = \alpha + \gamma_j \cdot T_{ijkm}(Y_0) + \delta_1 \cdot ECT_k + \delta_2 \cdot Characteristic_{ijkm} + \delta_3 \cdot (ECT_k \times Characteristic_{ijkm}) + \beta_z \cdot Z_m + \varepsilon_k + \varepsilon_{jk} + \varepsilon_{ijk} \quad (3)$$

Table 6 shows the results of these regressions on several school and household characteristics, and each column represents one regression testing for heterogeneous treatment effects along the characteristic mentioned (indicated by the coefficients on the interactions).

The first main result is the *lack* of heterogeneous treatment effects by several household and child-level characteristics. In particular, if we consider the baseline test score to be a summary statistic of all prior inputs into the child's education, then the lack of any significance on the interaction of the program with baseline scores (column 8) suggests that all children benefited equally from the program regardless of their initial level of learning and that the gains from the program were quite broad. Similarly, there was no difference in program effectiveness based on household affluence, parental literacy, caste, and gender of the child.

We do see that schools in more remote areas derive greater benefit from the addition of an extra contract teacher. A related (but weaker) result is that schools with poorer infrastructure and with fewer students also benefit more from the extra contract teacher (though the latter are not significant). Since remote schools also typically have fewer students and poorer infrastructure, we include all three variables (remoteness index, infrastructure index, and log of enrollment) and their interactions in an extended version of (2) and find that there is still a significant positive effect of the extra contract teacher in schools that are more remote in both years but not in schools with poorer infrastructure or fewer students (results available on request).

Thus, it appears that the mechanism for finding heterogeneous impacts of the extra contract teacher by school proximity is not just that the marginal reduction in school-level PTR induced by the extra contract teacher is larger in smaller schools (this would be picked up in the interaction with school enrollment) but that locally-hired contract teachers may be especially

effective in remote areas where the social distance between civil-service teachers and the community would be larger, and where civil-service teachers try to avoid getting posted (Fagernas and Pelkonen 2012).

3.6. Teacher Effort by Contract and Treatment Status

Enumerators measured teacher absence, and recorded whether those who were present were actively teaching during unannounced visits to schools. Table 7 – Panel A shows that contract teachers had significantly lower levels of absence compared to regular teachers (17.8% vs. 27.3% on average over the two years). Contract teachers also had significantly higher rates of teaching activity compared to regular teachers (48.4% vs. 42.2% on average over two years).

These differences are even higher with school fixed effects, suggesting that the presence of the contract teachers may have induced regular teachers to shirk more. We test this directly by comparing the absence rates of regular teachers in control schools with those in program schools and see that regular teachers in program schools have higher rates of absence and lower rates of teaching activity than their counterparts in comparison schools (Table 7 – Panel B), though these are not always significant. Thus, our estimate of the impact of an additional contract teacher is a composite estimate that *includes* the reduction in effort of regular teachers induced by the presence of the extra contract teacher, and is therefore likely to be a lower bound on the pure 'production function' effect of providing an additional contract teacher to schools (Das et. 2013).

4. Cost Effectiveness

4.1. Impact of Reducing PTR with a Regular Civil-Service Teacher

The experimental results establish that the marginal product of contract teachers is positive, and refute the conventional view that these teachers will not help improve learning. However, the broader policy question is that of the relative effectiveness of regular and contract teachers, and the ideal parallel experiment would have been to provide a randomly-selected set of schools with an extra civil-service teacher. However, such a design was not feasible because civil-service teacher assignments have to follow administrative rules, and cannot be based on random assignment, and no such experiment has ever been conducted to the best of our knowledge.¹⁹

¹⁹ The Duflo et al (2012) study experimentally provides schools with an extra contract teacher, and then generates the comparison with regular teachers by further randomizing the within-school allocation of students to contract and regular teachers. Similarly, the Tennessee Star experiment (Krueger 1999) also varies class size by randomly varying student assignment within schools as opposed to randomly assigning an extra regular teacher across schools.

Nevertheless, as we show below, the richness of the panel data we collect allows us to construct a robust and credible estimate of the impact of PTR reductions (at the school level) using a regular teacher. While the contract teacher experiment lasted for two years, the parallel experiment on teacher performance pay was continued for five years (Muralidharan 2012) as a result of which we tracked learning outcomes in the control schools for five years (from 2005-06 to 2009-10).²⁰ We use the data from the control schools over these five years (which represent the 'business as usual' school system with no other interventions) to estimate:

$$T_{ijkm}(Y_n) = \alpha + \gamma_j \cdot T_{ijkm}(Y_{n-1}) + \beta_2 \cdot \log_PTR_k(Y_n) + \beta_x \cdot X_{ijk} + \beta_z \cdot Z_{d/m/k} + \varepsilon_k + \varepsilon_{jk} + \varepsilon_{ijk} \quad (4)$$

where the estimation sample is restricted to control schools that *had only regular teachers* in the year concerned (see Appendix A1 for details). Thus, $\log_PTR_k(Y_n)$ is the logarithm of school-level PTR (in school k in year n) in cases where *all* the teachers in the school in that year are regular teachers, and β_2 estimates the correlation between $\log_PTR_k(Y_n)$ (using only regular teachers) and test-score *gains* in school k in year n in a standard value-added specification.

The main challenge in interpreting β_2 as the causal estimate of PTR on test-score gains is the concern that there are omitted variables across schools that could be correlated with both student learning trajectories as well as PTR. We therefore augment (4) with a full set of school and household controls (X_{ijk}), which are the same as in Table 3, and then with district, and mandal (sub-district) fixed effects ($Z_{d/m}$).²¹ Nevertheless, the identifying variation in these specifications is cross-sectional and it is difficult to fully allay concerns regarding omitted variables. Our preferred estimates therefore do not use any between-school variation in PTR and are based on estimating (4) with school fixed-effects (Z_k). The identifying variation in PTR now comes from variation in PTR in the *same* schools over time (which is driven by changes in cohort sizes over time and teacher transfers in and out of schools). We provide more details on

²⁰ The AP RESt experiments featured two input (extra contract teacher, and school block grant) and two incentive (group and individual teacher performance pay) interventions and was initially planned (and budgeted) for two years. At the end of the initially budgeted period of two years, we were not able to continue the input interventions (since these required funding at the *start* of the third school year, which the project did not have), while we were able to continue the incentive interventions for longer (because these only required funds at the *end* of the third school year, which we were assured of being able to secure during the course of the school year).

²¹ Since there is considerable variation in the quality of administration across jurisdictions, specifications that include district (mandal) fixed effects eliminate concerns of omitted variables across administrative jurisdictions that are correlated with both PTR and learning trajectories, and are identified using variation *within* the concerned administrative unit (with the mandal being the lowest unit of government administration in AP).

the estimation sample and robustness checks in Appendix A1 (and Appendix Tables 1-3), but the main point to highlight is that there is no correlation between changes in PTR between years within a school and the end of year test scores of the previous year, suggesting that within-school variation in PTR over time is quasi-random and highly unlikely to be correlated with omitted variables that may also be correlated with test-score gains.

We show each of these estimates in Table 8, and see that that β_2 ranges from -0.25 (with no controls or fixed effects) to -0.18 (in our preferred specification with school fixed effects and household controls). It is worth noting that the estimates in columns 1-8 of Table 8 are quite similar and not significantly different from each other. The most likely reason for this is that our core specification is a value-added one and the inclusion of lagged test scores in (4) probably accounts for most cross-sectional omitted variables that may be correlated with both the PTR and the gains in test scores. Our preferred estimate of β_2 from columns 7-8 (with school fixed effects) is around -0.20, implying that a 10% reduction in PTR induced by hiring more regular teachers would improve mean test scores across schools and grades by 0.02σ per year.²²

It is worth highlighting that the estimates presented in Columns 7 and 8 of Table 8 are probably the most credible estimates of the impact of reducing PTR (or class size) in the literature on education in developing countries. The best estimates to date are from Urquiola (2006) in Bolivia, and Angrist and Lavy (1999) in Israel, but both study the impact of class size on test score *levels* as opposed to test score *gains*. Thus, while both papers find that larger class sizes reduce test scores, the magnitude of the estimates may be confounded by class-size variation in previous grades (these are likely to be correlated within a cohort over time) and are therefore less useful for cost effectiveness comparisons, especially relative to experimental evaluations of interventions that typically report impacts in terms of test-score gains.

Thus, the panel-data estimates of the impact of PTR reductions using regular civil-service teachers on test score gains are an important result in themselves, and provide a useful

²² We also estimate columns 1-6 of Table 8 using control school test score data from only the first two years (corresponding to the duration of the experiment), and obtain very similar estimates of β_2 (from -0.20 to -0.26). However, we prefer to use five years of data since it allows us to include school fixed effects and eliminate concerns of cross-sectional omitted variables. Finally, since the range of variation in PTR used in the panel estimates in columns 7-8 (variation within schools over time) may be considerably narrower than the range of variation in PTR across schools, we plot histograms of the distribution of the identifying variation in de-measured \log_PTR with the full set of fixed effects used in Table 8 (none, district, mandal, and school) and we see that while there are fewer outliers in \log_PTR with school fixed effects, the identifying variation in the specification with school fixed effects spans most of the range of PTR's found in the cross-sectional sample (Figure 1).

benchmark against which to compare other education interventions. However, the main use of these results in this paper is to enable a credible comparison of the impact of reducing PTR using a contract teacher versus a regular teacher. We present a formal test of the equality of the impact of reducing PTR using a contract teacher (β_1 in Table 5 – Column 2) and that of doing so with a regular teacher (each β_2 in Table 8), and we see that we can never reject the null hypothesis that contract and regular teachers are equally effective. If anything, the point estimates suggest that contract teachers may even be *more* effective than regular teachers on the current margin (since $|\beta_1| > |\beta_2|$ for *every* estimate of $|\beta_2|$ in Table 8, and $|\beta_1|$ is nearly twice the value of $|\beta_2|$ in the best-identified specifications in columns 7 and 8).

These results suggest that contract teachers are at least as effective as regular teachers, and it is extremely unlikely that any further omitted variables could reverse this result since $|\beta_1|$ in Table 5 is an experimental estimate, and including a full set of controls and fixed effects barely changes the estimate of $|\beta_2|$ across the columns in Table 8 (with the most credible panel-data estimates with school fixed effects yielding the *lowest* value of $|\beta_2|$).

A final caveat is that $|\beta_1|$ estimates the *marginal* effect of reducing PTR with a contract teacher, whereas $|\beta_2|$ estimates the effect of variation in *average* PTR with a regular teacher, and so the estimates may not be directly comparable. We address this concern by including a quadratic term for \log_PTR in (4) and find that the coefficient on this term is insignificant (table available on request). Hence, the average and marginal effects of PTR reduction are unlikely to be very different for the range of PTR variation that we use in our estimation sample.

Economic theory suggests that optimal production of education will use expensive better-qualified regular teachers and inexpensive less-qualified contract teachers in the proportion where the ratio of marginal costs equals the ratio of marginal productivity. Since regular teachers are paid at least five times higher salaries than contract teachers, we also test if they are five times more productive and strongly reject the hypothesis that $\beta_2 = 5 \cdot \beta_1$ (at the 1% level). In fact, we reject $\beta_2 = 3 \cdot \beta_1$ and even $\beta_2 = 2 \cdot \beta_1$ (both at the 5% level), suggesting that it may be possible to improve learning outcomes *and* reduce costs at the same time by hiring even just two (or three) contract teachers in lieu of an additional regular teacher.

4.2. Comparing Regular and Contract Teachers within schools

Since we can match students in each year to their teacher and know the teacher type in that year, we can also compare the relative impact of being taught by a contract teacher as opposed to a regular teacher on test score gains (using *within* school variation in the type of teacher assigned to a grade). We present these results in Appendix A2 and find that there is no differential effect on learning gains for students taught by contract teachers relative to those taught by regular teachers (Appendix Table 4). As discussed in the introduction and in section 2.2, there are several challenges in interpreting these results causally. Hence, we show these results for completeness, but the most credible comparison of regular and contract teachers is the one shown in Table 8, which abstracts away from how teaching resources are used within schools.

4.3. Absolute Returns to Hiring an Extra Contract Teacher

In addition to the relative cost effectiveness of regular and contract teachers, it is useful to estimate the *absolute* social return to hiring more contract teachers, by comparing the costs of doing so to the present discounted value of the increased earnings that may accrue from improving human capital of students. Doing so requires several assumptions (see Appendix A3), but the estimated returns are large enough (ranging from 12 to 160 times the initial cost) that even if the estimates on the labor market returns to test scores were to be substantially lower, it is likely that hiring additional contract teachers would still have a high social rate of return.

5. Discussion

5.1. Teacher labor markets in rural India

To provide further perspective on our results showing that contract teachers are at least as effective as regular teachers, but at much lower cost, we also briefly discuss the private sector teacher labor market in rural India. India has seen a rapid increase in the enrolment share of private schools in the past decade, with recent estimates showing that around 30% of primary school students in *rural* Andhra Pradesh attend a fee-charging private school (Pratham, 2012). Since fee-charging private schools compete against free public schools as well as against other private schools for both students and teachers, they are likely to face better incentives than public schools to operate close to the efficient frontier of education production, where the desired quality of education is produced at the lowest possible cost.

As part of a different study on school choice (Muralidharan and Sundararaman 2013), we collected detailed data on teachers in private schools in the same five districts where the contract teacher experiment was conducted, and Table 9 compares regular teachers, contract teachers, and private school teachers (sampled from the same villages)²³ on a range of characteristics. The age and gender profile of private school teachers are similar to those of contract teachers (younger and more likely to be female than regular teachers). Private school teachers have higher levels of general education, but even lower levels of teacher training than contract teachers. They live much closer to the school and are more likely to be from the same village relative to regular teachers (though less so than contract teachers).

But, the most relevant comparison is that the salaries of private school teachers are even lower than those of contract teachers and *around an eighth* of regular teacher salaries. Figure 2 plots the salary distribution of teachers in government and private schools, and we see that the distribution of salaries in private schools is around the range of the contract teachers' salaries, and that private school teacher salaries are so much lower than those of regular government-school teachers that there is almost *no common support* between the two distributions. Finally, private school teachers and contract teachers have similarly low rates of absence, which are around half that of the regular teachers in spite of being paid much lower salaries.

The private school data helps clarify the context of teacher labor markets in rural India and provides important guidance for thinking about expanding the use of contract teachers in government schools. First, the employment terms of contract teachers are not 'exploitative' as believed by opponents of their use, but in line with the market clearing wage paid by private schools. While their terms might seem exploitative when working side by side with regular teachers and doing the same work for a fraction of the salary, the distortion is not the 'low' contract teacher salaries but rather the above-market pay of civil-service teachers.

Second, private schools in rural AP pay much lower teacher salaries, but hire more teachers and have significantly lower PTR's, and rates of multi-grade teaching than public schools (Muralidharan and Sundararaman 2013). Thus, it appears that an optimizing producer of

²³ Note that this is a different sample from that used in Table 1. The sample in Table 1 is representative of rural government-run schools, which is the focus of this paper; the sample in Table 9 is from a sample of villages that have private schools (which tend to be larger). The data for Table 9 was also collected 3 years later than the data used for Table 1. AP government policies on contract teacher salaries now provides for some differentiation by education and experience, which accounts for the distribution in Figure 2. The lower absence rates of regular teachers in Table 9 as opposed to in Table 5 are also likely to be because the sample used for Table 9 is drawn from larger villages that are less remote.

primary education services in rural AP would hire teachers whose education, training, demographics, and pay resemble those of contract teachers rather than that of the more qualified regular teachers found in public schools, but would hire more such teachers and have lower PTR's and lower rates of multi-grade teaching than in public schools. To the extent that the input combination used by private schools is likely to be closer to the efficient frontier of education production, expanding the use of contract teachers in government-run schools on the margin may be a way of moving public production of education closer to the efficient frontier.

Third, since private schools are able to fill their teacher positions with salaries that are even lower than those of contract teachers, an expansion of contract teacher hiring is unlikely to hit a supply constraint at current salary levels. None of the 100 treatment schools in our experiment reported any difficulty in filling the contract teacher position and the majority of positions were filled within 2 to 3 weeks from the start of the search. More broadly, the pool of educated but unemployed rural high-school and college graduates from which contract and private school teachers are hired appears to be large enough for the labor supply of contract teachers to be fairly elastic (Kingdon and Sipahimalani-Rao, 2010).²⁴

Finally, it is worth highlighting that the very high salaries paid to civil-service teachers (relative to the market wage) as seen in Figure 2 can have negative consequences beyond the fiscal inefficiency of incurring five times greater costs than necessary. Kingdon (2011) suggests that these high salaries may contribute to the large social distance between communities and teachers and make it more difficult for parents to hold teachers accountable. A further negative consequence of the large rents in civil-service teacher jobs is that a 'market' can emerge for these jobs whereby politicians appoint candidates who are willing to pay for civil-service teacher positions.²⁵ Candidates appointed this way are more likely to be negatively selected on intrinsic motivation (Shleifer and Vishny 1989), and are likely to have poorer norms for effort since they have 'paid' for their jobs (Akerlof 1982).

²⁴ Another contributing factor may be that limited job opportunities within the village for educated rural women (who have cultural and family restrictions against working outside their village) may be providing a subsidy to the teaching sector by lowering the wages of female contract and private school teachers (Andrabi et al, 2012).

²⁵ For instance, a former Chief Minister (elected *head* of government) of the Indian state of Haryana was recently prosecuted and convicted for illegally recruiting over 3000 teachers for financial considerations (Bhardwaj 2013). This is probably the first time in India that a senior politician has faced legal consequences for corruption in teacher hiring, and the threat of such punishment has typically not been a deterrent to markets in the hiring and assignment of government school teachers (Beteille 2009). See Hallak and Poisson (2007) for more examples of corruption in teacher recruitment.

5.2. Decomposing the reasons for contract teacher effectiveness

As the summary statistics in Table 1 indicate, contract teachers are different from regular teachers in *many* ways, some of which may make them less effective (lower levels of education, training, and pay), and some of which may make them more effective (being from the same village with closer connections to the community, living closer to the school, and being more accountable due to the renewable nature of the job contract). Decomposing the relative importance of these factors is clearly important for policy, but we do not attempt to do so here because it is not possible to obtain credible exogenous variation in each of these individual teacher characteristics. Thus, our results do not estimate the impact of the 'contractual' nature of the contract teacher job per se. Rather they compare two different vectors of teacher attributes (as embodied in the contract and regular teachers) and show that the combination of attributes found in a typical contract teacher makes them at least as effective as regular civil-service teachers, but at a small fraction of the cost.

More broadly, our results highlight that teacher quality (measured by their effectiveness in improving learning outcomes) depends on their human capital as well as their effort, and one interpretation of our results is that contract teachers make up for their lower levels of formal education and training with higher levels of effort.²⁶ The results also suggest that a lack of formal teacher training credentials (as currently generated by the teacher education system in India) may not be a binding constraint to teacher effectiveness.²⁷ Of course, it may be possible to achieve even better outcomes by combining features of contract and regular teachers, but estimating the effectiveness of such combinations is beyond the scope of this paper.²⁸

²⁶ Note that unlike some other states in India, where contract teachers are given preference for getting hired as regular teachers, AP had no history of doing so, and there was no provision for such 'regularization' of contract teachers. Thus, it is unlikely that the reason for higher effort of the contract teachers was the prospect of being rewarded for this effort by being hired as a civil-service teacher.

²⁷ Our results are consistent with those from other experimental studies in India showing that supplemental instruction programs delivered at the level of learning of the student (as opposed to the textbook) have been very effective at improving test scores even though these programs have typically been delivered by young locally-hired women on modest stipends (usually even lower than the salary of a contract teacher) who typically have a secondary school or high school qualification and no formal training (Banerjee et al 2007, 2010, 2012; Lakshminarayana et al. 2013). These studies are quite different from ours because they emphasize remedial pedagogy (whereas we do not modify pedagogy in any way), but they also highlight that significant learning gains can be delivered by locally-hired teaching staff who are much less educated and qualified than civil-service teachers.

²⁸ Some possibilities include hiring qualified and trained teachers on renewable contracts at the village level or hiring contract teachers and training them to focus on teaching at the level of the student as opposed to the level of the text book (Banerjee and Duflo 2011; Pritchett and Beatty 2012). These may be especially promising combinations because there is evidence to suggest that teacher knowledge and effort may be complements in this setting (Muralidharan and Sundararaman 2011, Muralidharan 2012).

6. Conclusion

Regular teachers in India are well qualified, but command a substantial wage premium (greater than a factor of five) over contract teachers that can be explained in part by (a) their better education and outside opportunities, (b) a compensating differential to locate to rural and remote areas, and (c) a union and civil-service premium/rent. The hiring of contract teachers can be a much more cost-efficient way of adding teachers to schools because none of these three sources of wage premiums are applicable for them. However, since locally-hired contract teachers are not as qualified or trained as civil-service teachers, opponents of the use of contract teachers have posited that the use of contract teachers will not lead to improved learning.

We present experimental evidence from an 'as is' expansion of the existing contract teacher policy, implemented by the AP Government in a randomly selected subset of 100 schools among a representative sample of schools in rural AP. We find that adding a contract teacher significantly improved average learning outcomes in treatment schools, and that contract teachers are no less effective in improving student learning than regular teachers who are more qualified, better trained, and paid five times higher salaries. Thus, the strong beliefs embedded in the education policies of many developing countries that contract teachers are (a) ineffective, and (b) inferior to civil-service teachers (even if effective), are not supported by our evidence.

Our results suggest that expanding the use of contract teachers could be a highly cost effective way of improving primary education outcomes in developing countries. In particular, expensive policy initiatives to get highly qualified teachers to remote areas may be much less cost effective than hiring *several* local contract teachers to provide much more attention to students at a similar cost. Doing so would enable public schools to have much lower PTR's, much lower levels of multi-grade teaching, and also provide additional teaching resources to support first-generation learners who fall behind relative to the syllabus and appear to not learn much in spite of being enrolled in school for five years (Muralidharan and Zieleniak 2013).²⁹

There are two important caveats to the discussion above. First, our results apply to the current *margin* of the education system given the existing stock of regular and contract teachers. It is possible that the presence of regular teachers (and their support and monitoring) is required

²⁹ Observing the input choices of private schools suggests that this is what a politically unconstrained, optimizing producer of rural education services would do. Note that hiring teachers with a 'contract teacher' profile allows private schools to have lower PTR's, and much lower levels of multi-grade teaching, in spite of having per-student spending that is only *one-third* of the spending in public schools (see Muralidharan and Sundararaman 2013).

for contract teachers to be effective, and thus our results do not imply that all regular teachers should be replaced with contract teachers. Second, our results apply to primary school and in particular to a context where a large fraction of students are first-generation learners without educated parents at home and without having attended pre-school. Thus, while a high-school educated contract teacher may be adequately qualified to teach effectively in this setting, our results may not apply beyond primary school where the teachers' education and training may be more of a binding quality constraint that cannot be compensated for simply with more effort.

Despite our results, opponents of the use of contract teachers may still worry that the expanded use of untrained teachers will erode the professionalism of teaching and reduce incentives to improve teacher training. A second concern is that hiring larger numbers of contract teachers will lead to demands that they all be regularized into civil-service status, which may be politically difficult to resist given the strengths of teacher unions - especially around the time of elections. If such regularization (and the corresponding five-fold increase in pay) were to happen, it would defeat the economic case for hiring a large number of contract teachers in the first place. Finally, courts in India have on occasion ruled that contract teachers should be regularized based on principles of 'equal pay for equal work' (Robinson and Gauri 2010), and thus expansion of contract teacher hiring may not be consistent with existing labor laws.

Many of these concerns can be addressed by hiring new teachers as contract teachers at the village (or school) level, and creating a career ladder whereby bonuses, pay raises, and promotion to regular civil-service rank are based on multiple measures of performance over time. Continuous training and professional development could be a natural component of this career progression, and integrating contract and regular teachers into a career ladder could address most of the concerns above. Pritchett and Murgai (2007) and Muralidharan (2013a and 2013b) discuss implementation details of such a system in policy-focused complements to this paper.

There is also a potentially important extension based on our results. The public school system in India starts in grade one and students in public schools usually have no pre-school or preparation in being 'school ready' when they start school (unlike private school students, who typically have two years of pre-school education before the first grade). The effectiveness of contract teachers in a primary school setting, suggest that there may be high social returns to hiring teachers with similar demographic characteristics to provide pre-school instruction to public school students. This may be a highly promising area for future pilots and evaluation.

References:

- AKERLOF, G. (1982): "Labor Contracts as Partial Gift Exchange," *Quarterly Journal of Economics*, 97, 543-569.
- ANDRABI, T., J. DAS, and A. KHWAJA (2011): "Students Today, Teachers Tomorrow? Identifying Constraints on the Provision of Education," World Bank.
- ANGRIST, J. D., and V. LAVY (1999): "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement," *Quarterly Journal of Economics*, 114, 533-75.
- ASLAM, M., A. DE, G. KINGDON, and R. KUMAR (2011): "Economic Returns to Schooling and Skills – an Analysis of India and Pakistan," in *Education Outcomes and Poverty in the South*, ed. by C. Colclough. London: Routledge.
- ATHERTON, P., and G. KINGDON (2010): "The Relative Effectiveness and Costs of Contract and Regular Teachers in India," Institute of Education, University of London.
- BANERJEE, A., R. BANERJI, E. DUFLO, R. GLENNERSTER, and S. KHEMANI (2010): "Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in Education in India," *American Economic Journal: Economic Policy*, 2, 1-30.
- BANERJEE, A., R. BANERJI, E. DUFLO, and M. WALTON (2012): "Effective Pedagogies and a Resistant Education System: Experimental Evidence on Interventions to Improve Basic Skills in Rural India," MIT.
- BANERJEE, A., S. COLE, E. DUFLO, and L. LINDEN (2007): "Remedying Education: Evidence from Two Randomized Experiments in India," *Quarterly Journal of Economics*, 122, 1235-1264.
- BANERJEE, A., and E. DUFLO (2011): *Poor Economics*. MIT Press.
- BANG, A. T., R. A. BANG, S. B. BAITULE, M. H. REDDY, and M. D. DESHMUKH (1999): "Effect of Home-Based Neonatal Care and Management of Sepsis on Neonatal Mortality: Field Trial in Rural India," *The Lancet*, 354, 1955-1961.
- BARDHAN, P. (2002): "Decentralization of Governance and Development," *Journal of Economic Perspectives*, 16, 185-205.
- BARDHAN, P., and D. MOOKHERJEE (2000): "Capture and Governance at Local and National Levels," *American Economic Review*, 90, 135-139.
- BETEILLE, T. (2009): "Absenteeism, Transfers and Patronage: The Political Economy of Teacher Labor Markets in India," Stanford.
- BETTINGER, E., and B. T. LONG (2010): "Does Cheaper Mean Better? The Impact of Using Adjunct Instructors on Student Outcomes," *Review of Economics and Statistics*, 92, 598-613.
- BHARDWAJ, M. (2013): "Om Prakash Chautala: Rise, Fall, Rise and Downfall," *The Indian Express*
- BOLD, T., M. KIMENYI, G. MWABU, N. A. A. ALICE, and J. SANDEFUR (2013): "Scaling-up What Works: Experimental Evidence on External Validity in Kenyan Education," Washington DC: Center for Global Development Working Paper 321.
- BOURDON, J., M. FRÖLICH, and K. MICHAELOWA (2010): "Teacher Shortages, Teacher Contracts and Their Impact on Education in Africa," *Journal of the Royal Statistical Society, Series A*, 173, 93-116.
- CHAUDHURY, N., J. HAMMER, M. KREMER, K. MURALIDHARAN, and F. H. ROGERS (2006): "Missing in Action: Teacher and Health Worker Absence in Developing Countries," *Journal of Economic Perspectives*, 20, 91-116.
- CHETTY, R., J. N. FRIEDMAN, and J. E. ROCKOFF (2013): "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood," NBER Working Paper 19424.
- DAS, J., S. DERCON, J. HABYARIMANA, P. KRISHNAN, K. MURALIDHARAN, and V. SUNDARARAMAN (2013): "School Inputs, Household Substitution, and Test Scores," *American Economic Journal: Applied Economics*, 5, 29-57.

- DE LAAT, J., and E. VEGAS (2005): "Do Differences in Teacher Contracts Affect Student Performance? Evidence from Togo," World Bank.
- DHALIWAL, I., E. DUFLO, R. GLENNERSTER, and C. TULLOCH (2012): "Comparative Cost-Effectiveness Analysis to Inform Policy in Developing Countries: A General Framework with Applications for Education," MIT.
- DUFLO, E., P. DUPAS, and M. KREMER (2012): "School Governance, Teacher Incentives, and Pupil-Teacher Ratios: Experimental Evidence from Kenyan Primary Schools," NBER Working Paper 17939.
- DUTHILLEUL, Y. (2005): "Lessons Learnt in the Use of 'Contract' Teachers," International Institute for Educational Planning, UNESCO.
- FAGERNAS, S., and P. PELKONEN (2012): "Preferences and Skills of Indian Public Sector Teachers," *IZA Journal of Labor & Development*, 1.
- FIGLIO, D. N., M. O. SCHAPIRO, and K. SOTER (2013): "Are Tenure Track Professors Better Teachers?," NBER Working Paper 19406.
- FRIEDMAN, M., and S. KUZNETS (1945): *Income from Independent Professional Practice*. New York: National Bureau of Economic Research.
- GOVINDA, R., and Y. JOSEPHINE (2004): "Para-Teachers in India: A Review," New Delhi: National Institute of Educational Planning and Administration.
- GOYAL, S., and P. PANDEY (2011): "Contract Teachers in India," *Education Economics*, 1-21.
- HAINES, A., D. SANDERS, U. LEHMANN, A. K. ROWE, J. E. LAWN, S. JAN, D. G. WALKER, and Z. BHUTTA (2007): "Achieving Child Survival Goals: Potential Contribution of Community Health Workers," *The Lancet*, 369, 2121-2131.
- HALLAK, J., and M. POISSON (2007): "Corrupt Schools, Corrupt Universities: What Can Be Done?," International Institute for Educational Planning.
- HECKMAN, J., and J. SMITH (1995): "Assessing the Case of Social Experiments," *Journal of Economic Perspectives*, 9, 85-110.
- JALAL, F., M. SAMANI, M. C. CHANG, R. STEVENSON, A. B. RAGATZ, and S. D. NEGARA (2009): "Teacher Certification in Indonesia: A Strategy for Teacher Quality Improvement," Jakarta: World Bank.
- KINGDON, G. (2011): "Private Versus Public Schooling in India," *Seminar*.
- KINGDON, G. G., and M. MUZAMMIL (2001): "A Political Economy of Education in India: The Case of U.P.," *Economic and Political Weekly*, 36.
- KINGDON, G. G., and V. SIPAHIMALANI-RAO (2010): "Para-Teachers in India: Status and Impact," *Economic and Political Weekly*, XLV, 59-67.
- KLEINER, M. (2000): "Occupational Licensing," *Journal of Economic Perspectives*, 14, 189-202.
- KLEINER, M., and A. KRUEGER (2013): "Analyzing the Extent and Influence of Occupational Licensing on the Labor Market," *Journal of Labor Economics*, 31, S173-S202.
- KREMER, M., K. MURALIDHARAN, N. CHAUDHURY, F. H. ROGERS, and J. HAMMER (2005): "Teacher Absence in India: A Snapshot," *Journal of the European Economic Association*, 3, 658-67.
- KRUEGER, A. (1999): "Experimental Estimates of Education Production Functions," *Quarterly Journal of Economics*, 114, 497-531.
- KUMAR, K., M. PRIYAM, and S. SAXENA (2005): "The Trouble with Para-Teachers," *Frontline*, 18.
- LAKSHMINARAYANA, R., A. EBLE, P. BHAKTA, C. FROST, P. BOONE, D. ELBOURNE, and V. MANN (2013): "The Support to Rural India's Public Education System (Stripes) Trial: A Cluster Randomised Controlled Trial of Supplementary Teaching, Learning Material and Material Support," *Plos One*, 8, 1-13.
- LINDBECK, A., and D. J. SNOWER (1988): *The Insider-Outsider Theory of Employment and Unemployment*. MIT.
- (2001): "Insiders Versus Outsiders," *Journal of Economic Perspectives*, 15, 165-188.

- MURALIDHARAN, K. (2012): "Long Term Effects of Teacher Performance Pay: Experimental Evidence from India," UC San Diego.
- (2013): "Priorities for Primary Education Policy in India's 12th Five-Year Plan," *India Policy Forum*, 9, 1-46.
- (2013): "Proposal for a National Program of Primary School Teaching Assistants to Guarantee Minimum Absolute Learning Standards for All Children by the End of the 12th Plan," UC San Diego.
- MURALIDHARAN, K., J. DAS, A. HOLLA, M. KREMER, and A. MOHPAL (2013): "The Fiscal Costs of Weak Governance: Evidence from Teacher Absence in India," UC San Diego.
- MURALIDHARAN, K., and V. SUNDARARAMAN (2011): "Teacher Performance Pay: Experimental Evidence from India," *Journal of Political Economy*, 119, 39-77.
- (2013): "The Aggregate Effects of School Choice: Evidence from a Two-Stage Experiment," NBER Working Paper 19441.
- MURALIDHARAN, K., and Y. ZIELENIAK (2013): "Measuring Learning Trajectories in Developing Countries with Longitudinal Data and Item Response Theory," UC San Diego.
- PRATHAM (2012): *Annual Status of Education Report*.
- PRITCHETT, L. (2004): "Access to Education," in *Global Crises, Global Solutions*, ed. by B. Lomborg, 175-234.
- PRITCHETT, L., and A. BEATTY (2012): "The Negative Consequences of over-Ambitious Curricular in Developing Countries," Harvard Kennedy School.
- PRITCHETT, L., and R. MURGAI (2007): "Teacher Compensation: Can Decentralization to Local Bodies Take India from Perfect Storm through Troubled Waters to Clear Sailing?," *India Policy Forum*, 3, 123-168.
- PRITCHETT, L., and V. PANDE (2006): "Making Primary Education Work for India's Rural Poor: A Proposal for Effective Decentralization," New Delhi: World Bank.
- PRITCHETT, L., and J. SANDEFUR (2013): "Context Matters for Size: Why External Validity Claims and Development Practice Don't Mix," Washington DC: Center for Global Development Working Paper 336.
- RAINA, V. (2006): "Kill the Bill," *Seminar*.
- (2009): "Right to Education," *Seminar*.
- ROBINSON, N., and V. GAURI (2010): "Education, Labor Rights, and Incentives Contract Teacher Cases in the Indian Courts," World Bank.
- SHLEIFER, A., and R. W. VISHNY (1993): "Corruption," *The Quarterly Journal of Economics*, 108, 599-617.
- SMITH, A. (1776): *The Wealth of Nations*.
- SRIVASTAVA, M. (2010): "Could the Poor Demand and Control Their Own Schools," London School of Economics.
- URQUIOLA, M. (2006): "Identifying Class Size Effects in Developing Countries: Evidence from Rural Bolivia," *Review of Economics and Statistics*, 88, 171-177.
- UWEZO (2012): "Are Our Children Learning? Literacy and Numeracy across East Africa," Nairobi, Kenya.

Table 1: Teacher Characteristics by Type

Panel A: Regular Teachers have Significantly Different Characteristics from Contract Teachers (Control School Sample)			
	Regular Teachers	Contract Teachers	P-value (H0: Diff=0)
Male	63.1%	31.8%	0.005***
Age	38.35	23.81	0.000***
College Degree or Higher	84.3%	45.5%	0.002***
Formal Teacher Training Degree or Certificate	98.3%	9.1%	0.000***
Received any Training in last twelve months	93.5%	54.5%	0.000***
From the same village	7.2%	81.8%	0.000***
Distance from home to school (km)	11.92	1.091	0.000***
Teacher Salary (Rs./month)	8698.1	1000 (1500)	0.000***
Panel B: The Extra Contract Teachers Hired in the Treatment Schools under the Experiment had the Same Characteristics on Average as the Typical Contract Teacher in Control Schools			
	Contract Teachers in Treatment Schools	Contract Teachers in Control Schools	P-value (H0: Diff=0)
Male	32.0%	31.8%	0.99
Age	25.56	23.81	0.18
College Degree or Higher	47.6%	45.5%	0.87
Formal Teacher Training Degree or Certificate	16.5%	9.1%	0.32
Received any Training in last twelve months	36.9%	54.5%	0.10
From the same village	89.3%	81.8%	0.42
Distance from home to school (km)	0.553	1.091	0.15
Teacher Salary (Rs./month)	1000	1000	0.91

Notes:

1. Table reports summary statistics from the first year of the project (2005 - 06). The teacher characteristics were similar in the second year as well (2006 - 07). The main difference was that contract teacher salary was Rs. 1000/month in the first year, but increased to Rs. 1,500 across the entire state in the second year

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 2: Treatment and Control Schools were Balanced on Observable Characteristics

Panel A				
(Mean Pre-program Characteristics)				
	[1]	[2]	[3]	
	Comparison Schools	Extra Contract Teacher Schools	P-value (H0: Diff=0)	
<u>School-level Variables</u>				
1	Total Enrollment (Baseline: Grades 1-5)	113.2	104.6	0.41
2	Total Test-takers (Baseline: Grades 2-5)	64.9	62.0	0.59
3	Number of Teachers	3.07	2.83	0.24
4	Pupil-Teacher Ratio	39.5	39.8	0.94
5	Infrastructure Index (0-6)	3.19	3.13	0.84
6	Proximity to Facilities Index (8-24)	14.65	14.97	0.55
<u>Baseline Test Performance</u>				
7	Math (Raw %)	18.5	17.31	0.34
8	Math (Normalized - in Std. deviations)	0.041	-0.043	0.29
9	Telugu (Raw %)	35.1	34.343	0.66
10	Telugu (Normalized - in Std. deviations)	0.018	-0.019	0.65
Panel B				
(Mean Turnover/Attrition During Program)				
	[1]	[2]	[3]	
	Comparison Schools	Extra Contract Teacher Schools	P-value (H0: Diff=0)	
<u>Teacher Turnover and Attrition</u>				
Year 1 (relative to Year 0)				
11	Teacher Attrition (%)	0.30	0.31	0.76
12	Teacher Turnover (%)	0.34	0.33	0.85
Year 2 (relative to Year 0)				
13	Teacher Attrition (%)	0.36	0.42	0.17
14	Teacher Turnover (%)	0.34	0.34	0.98
<u>Student Turnover and Attrition</u>				
Year 1 (relative to Year 0)				
15	Student Attrition from baseline to end of year tests	0.08	0.07	0.13
16	Baseline Maths test score of attritors	-0.16	-0.17	0.90
17	Baseline Telugu test score of attritors	-0.27	-0.24	0.76
Year 2 (relative to Year 0)				
18	Student Attrition from baseline to end of year tests	0.22	0.19	0.14
19	Baseline Maths test score of attritors	-0.13	-0.03	0.39
20	Baseline Telugu test score of attritors	-0.19	-0.14	0.63

Notes:

1. The school infrastructure index sums 6 binary variables (coded from 0 - 6) indicating the existence of a brick building, a playground, a compound wall, a functioning source of water, a functional toilet, and functioning electricity.
2. The school proximity index ranges from 8-24 and sums 8 variables (each coded from 1-3) indicating proximity to a paved road, a bus stop, a public health clinic, a private health clinic, public telephone, bank, post office, and the mandal educational resource center. A higher value of the index indicates being further away from these facilities.
3. Teacher attrition refers to the fraction of teachers in the school who left the school during the year, while teacher turnover refers to the fraction of new teachers in the school at the end of the year (both are calculated relative to the list of teachers in the school at the start of the year)
4. The p-values for the baseline test scores and attrition are computed by treating each student/teacher as an observation and clustering the standard errors at the school level (Grade 1 did not have a baseline test). The other p-values are computed treating each school as an observation.

Table 3: Provision of an Extra Contract Teacher to Schools Led to a Significant Increase in Student Test Scores in both Math and Language

Dependent Variable (All 3 Panels): Normalized End of Year Test Score				
Panel A: Combined - Math and Language (Telugu)				
	After 1 Year		After 2 Years	
	[1]	[2]	[3]	[4]
Extra Contract Teacher School	0.102*** (0.035)	0.094*** (0.035)	0.156*** (0.045)	0.152*** (0.048)
Observations	27630	24930	26142	22565
R-squared	0.348	0.378	0.203	0.227
Panel B: Math				
	After 1 Year		After 2 Years	
	[1]	[2]	[3]	[4]
Extra Contract Teacher School	0.125*** (0.039)	0.112*** (0.039)	0.158*** (0.051)	0.168*** (0.055)
Observations	13742	12399	13020	11236
R-squared	0.331	0.361	0.197	0.215
Panel C: Language (Telugu)				
	After 1 Year		After 2 Years	
	[1]	[2]	[3]	[4]
Extra Contract Teacher School	0.079** (0.035)	0.076** (0.035)	0.153*** (0.042)	0.135*** (0.045)
School and Household Controls (All 3 Panels)	No	Yes	No	Yes
Observations	13888	12531	13122	11329
R-squared	0.379	0.409	0.218	0.247

Notes:

1. All regressions include mandal (sub-district) fixed effects and standard errors clustered at the school level. They also include lagged normalized test scores (these are set to 0 for students in grade 1, and for students in grade 2 for the regressions showing 2-year effects - since these cohorts did not have a baseline test). All test scores are normalized relative to the distribution of scores in the control schools in the same grade, test, and year.

2. The two year treatment effect regressions include students who entered grade 1 in the second year of the program and who were there in the schools at end of two years of the program, but who have only been exposed to the program for one year at the end of two years of the program.

3. School controls include infrastructure and proximity indices as defined in Table 2. Household controls include a household asset index, parent education index (both defined as in Table 6), student gender, and an indicator for being from a scheduled (disadvantaged) caste/tribe.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 4: Changes In School and Classroom Characteristics Induced by the Provision of an Extra Contract Teacher

	Year 1			Year 2		
	Treatment	Control	Difference	Treatment	Control	Difference
	[1]	[2]	[3]	[4]	[5]	[6]
School Pupil-Teacher Ratio (PTR) WITHOUT the Extra Contract Teacher	35.672 [15.650]	32.663 [11.930]	3.009 (1.950)	39.105 [14.563]	37.774 [13.300]	1.331 (1.978)
School Pupil-Teacher Ratio (PTR) WITH the Extra Contract Teacher	25.409 [11.605]	32.663 [11.930]	-7.254*** (1.664)	28.309 [10.952]	37.774 [13.300]	-9.482*** (1.726)
Change in PTR induced by the Experiment (Using a Contract Teacher)	-10.262 [5.998]	0	-10.262*** (0.600)	-10.814 [6.386]	0	-10.814*** (0.642)
Change in Log_PTR induced by the Experiment (Using a Contract Teacher)	-0.351 [16.485]	0	-0.351*** (0.014)	-0.336 [0.132]	0	-0.336*** (0.013)
Class Size	28.027 [16.485]	33.683 [15.953]	-5.656*** (1.803)	31.724 [16.232]	39.104 [18.263]	-7.380*** (1.886)
Probability of Multi-grade Teaching (at grade level)	0.5 [0.5]	0.64 [0.48]	-0.137*** (0.031)	0.51 [0.5]	0.59 [0.450]	-0.080*** (0.031)
Probability of at least 2 teachers being present in school during unannounced visit	0.74 [0.44]	0.68 [0.47]	0.063** (0.026)	0.87 [0.33]	0.73 [0.44]	0.141*** (0.028)
Probability of at least 3 teachers being present in school during unannounced visit	0.56 [0.50]	0.42 [0.49]	0.142*** (0.029)	0.58 [0.49]	0.41 [0.49]	0.172*** (0.035)
Probability of Being Taught by a Contract Teacher (at grade level)	0.27 [0.44]	0.07 [0.26]	0.199*** (0.023)	0.31 [0.46]	0.08 [0.28]	0.225*** (0.024)

Notes

1. Numbers in brackets are the standard deviations of the underlying variable (in columns 1, 2, 4, and 5), and those in parentheses are the standard errors of the differences (in columns 3 and 6). Differences presented in columns 3 and 6 are based on regressing the concerned indicator on a treatment dummy. Results are at the school/class level and not weighted by enrollment.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 5: Impact of Reducing School-Level Pupil-Teacher Ratio (PTR) using a Contract Teacher

	Dependent Variable: Normalized Student Test Scores - Pooled across Math and Language (First Year Only)					
	Combined		Math		Telugu	
	[1]	[2]	[3]	[4]	[5]	[6]
Change in Log School PTR induced by additional (experimental) Contract Teacher (β_1)	-0.308*** (0.113)	-0.337*** (0.110)	-0.389*** (0.129)	-0.421*** (0.126)	-0.227** (0.108)	-0.253** (0.106)
Original Log School PTR		-0.232*** (0.060)		-0.264*** (0.069)		-0.201*** (0.059)
Observations	27630	27630	13742	13742	13888	13888
R-squared	0.348	0.353	0.332	0.338	0.379	0.383

Notes:

1. All regressions include mandal (sub-district) fixed effects and standard errors clustered at the school level. They also include lagged normalized test scores (these are set to 0 for students in grade 1 who did not have a baseline test). All test scores are normalized relative to the distribution of scores in the control schools in the same grade, test, and year.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 6: Heterogeneous Treatment Effects By School and Household Characteristics

	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
	Log Number of Students	Proximity (8 - 24)	Infrastructure (0 - 6)	Household Affluence (0 - 7)	Parental Education (0 - 4)	Scheduled Caste/Tribe	Male	Baseline Score
After 1 Year								
Covariate	-0.079** (0.039)	-0.009 (0.008)	0.013 (0.021)	0.025** (0.010)	0.063*** (0.011)	-0.007 (0.033)	0.014 (0.018)	0.501*** (0.022)
Interaction	-0.023 (0.067)	0.019* (0.011)	-0.037 (0.035)	0.013 (0.016)	0.003 (0.017)	-0.044 (0.047)	-0.012 (0.030)	0.019 (0.030)
Observations	27630	26788	26788	25651	25638	27630	26207	25979
R-squared	0.350	0.356	0.356	0.368	0.371	0.348	0.361	0.364
After 2 Years								
Covariate	-0.081 (0.054)	-0.005 (0.009)	0.038 (0.032)	0.046*** (0.012)	0.072*** (0.015)	-0.053 (0.041)	0.015 (0.025)	0.435*** (0.023)
Interaction	-0.108 (0.078)	0.043*** (0.016)	-0.124*** (0.041)	0.004 (0.019)	-0.024 (0.022)	-0.011 (0.061)	-0.019 (0.038)	0.032 (0.038)
Observations	26142	26142	26142	22588	22569	26142	22966	24756
R-squared	0.209	0.209	0.209	0.221	0.221	0.204	0.215	0.203

Notes:

1. Each column in each panel reports the results of a regression that includes the covariate in the column title, a binary treatment indicator (not shown), and a linear interaction term testing for heterogeneous effects of the treatment along the covariate concerned.
 2. All regressions include mandal (sub-district) fixed effects and standard errors clustered at the school level. All regressions include baseline test scores.
 3. The school infrastructure and proximity index are as defined in Table 2
 4. The household asset index ranges from 0 to 7 and is the sum of seven binary variables indicating whether the household has an electricity connection, has a water source at home, has a toilet at home, owns any land, owns their home, has a brick home, and owns a television.
 5. Parental education is scored from 0 to 4 in which a point is added for each of the following: father's literacy, mother's literacy, father having completed 10th grade, and mother having completed 10th grade
- * significant at 10%; ** significant at 5%; *** significant at 1%

Table 7: Teacher Absence and Effort

Panel A: Contract Versus Regular Teachers				
Teacher Absence				
	Contract Teachers	Regular Teachers	Difference	Difference with School Fixed Effects
Year 1	17.6%	25.7%	-8.1%***	-9.5%***
Year 2	17.9%	29.2%	-11.3%***	-16.4%***
Combined	17.8%	27.3%	-9.5%***	-12.2%***
Teachers Observed Actively Teaching				
	Contract Teachers	Regular Teachers	Difference	Difference with School Fixed Effects
Year 1	53.5%	47.7%	5.8%*	7.8%***
Year 2	42.9%	35.3%	7.6%***	7.7%***
Combined	48.4%	42.2%	6.2%***	7.3%***
Panel B : Regular Teachers in ECT Schools versus those in Control Schools				
Teacher Absence				
	Regular teachers in ECT schools	Regular teachers in non-ECT schools	Difference	Difference with Mandal fixed effects
Year 1	26.2%	25.3%	0.9%	0.5%
Year 2	31.2%	27.4%	3.8%	4.0%*
Combined	28.4%	26.2%	2.2%	2.0%
Teachers Observed Actively Teaching				
	Regular teachers in ECT schools	Regular teachers in non-ECT school	Difference	Difference with mandal fixed effects
Year 1	44.2%	50.9%	-6.7%*	-6.6%***
Year 2	34.7%	35.8%	-1.1%	-1.0%
Combined	40.1%	44.3%	-4.2%	-4.1%**

Notes

1. All standard errors are clustered at the school level. For school fixed effects, standard errors are clustered at the teacher level.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 8: Impact of Reducing School-Level Pupil-Teacher Ratio (PTR) using a Regular Teacher

Impact of Log School PTR on Test Scores (when taught by Regular Teachers) using 5 Years of Data								
Estimation Sample: Control Schools in Years when they had only Regular Teachers (See Appendix A1 for details)								
Dependent Variable: Normalized Student Test Scores (Pooled across Math and Language)								
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Log School Pupil-Teacher Ratio (PTR) (β_2)	-0.247*** (0.049)	-0.236*** (0.047)	-0.214*** (0.045)	-0.208*** (0.043)	-0.195*** (0.046)	-0.187*** (0.048)	-0.197*** (0.069)	-0.183** (0.070)
Household Controls	No	Yes	No	Yes	No	Yes	No	Yes
School Controls	No	Yes	No	Yes	No	Yes	NA	NA
District Fixed Effects	No	No	Yes	Yes	NA	NA	NA	NA
Mandal Fixed Effects	No	No	No	No	Yes	Yes	NA	NA
School Fixed Effects	No	No	No	No	No	No	Yes	Yes
Observations	45379	37232	45379	37232	45379	37232	45379	37232
R-squared	0.255	0.267	0.263	0.275	0.297	0.309	0.311	0.325
P-value (H0: $\beta_2 = \beta_1$)	0.4639	0.4073	0.3065	0.2828	0.2372	0.2167	0.2824	0.24
P-value (H0: $\beta_2 = 2*\beta_1$)	0.0607	0.0539	0.0417	0.0394	0.0338	0.0317	0.0387	0.0342
P-value (H0: $\beta_2 = 3*\beta_1$)	0.0229	0.0209	0.0172	0.0165	0.0146	0.0139	0.0158	0.0144
P-value (H0: $\beta_2 = 5*\beta_1$)	0.0095	0.0089	0.0078	0.0076	0.007	0.0068	0.0073	0.0069

Notes:

1. All regressions include fixed effects as indicated with standard errors clustered at the school level (clustered at the school-year level in specifications with school fixed effects). They also include lagged normalized test scores (these are set to 0 for students in grade 1 who did not have a baseline test). All test scores are normalized relative to the distribution of scores in the control schools in the same grade, test, and year. See Appendix A1 for details of the estimation sample

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 9: Comparing Regular, Contract, and Private School Teachers

	Regular Teacher	Contract Teachers	Private School Teachers	P-value (Null Hypothesis: Contract Teacher = Private School Teacher)
Female =1	62.5%	80.9%	88.4%	0.047
Age of Teacher	38.39	26.95	26.57	0.626
Teacher Passed College =1	87.0%	31.3%	52.4%	0.000
Received Any Teacher Training =1	99.2%	21.2%	14.1%	0.095
Received Training Within Past Yr =1	78.1%	43.5%	2.8%	0.000
Teacher from the Same Village =1	19.4%	80.2%	54.0%	0.000
Distance to School (km)	11.73	1.01	2.48	0.000
Gross Montly Salary (Rs.)	12,162	1,910	1,527	0.000
Percentage of Absent Teachers	20.7%	11.3%	9.7%	0.487

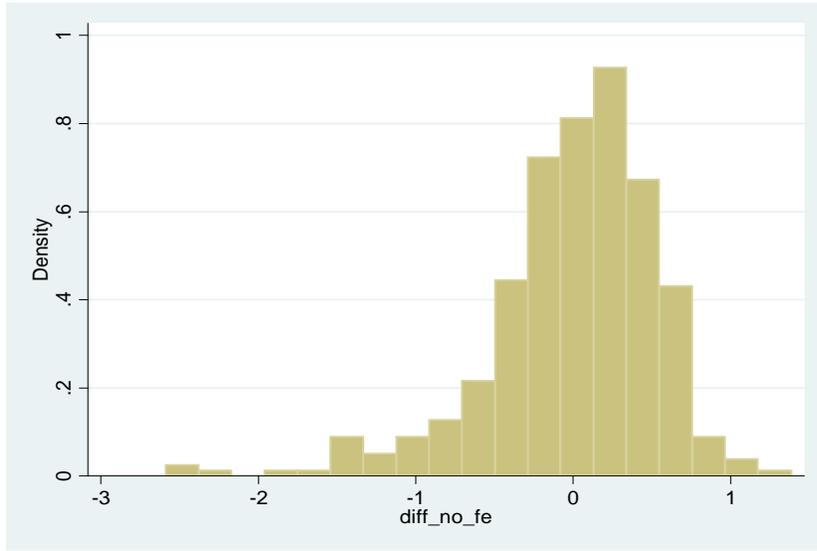
Notes

1. The data used for this table comes from a different study on school vouchers and school choice in different sub-districts of the SAME districts. This data was collected based on teacher interviews in early 2009 (Muralidharan and Sundararaman 2013).

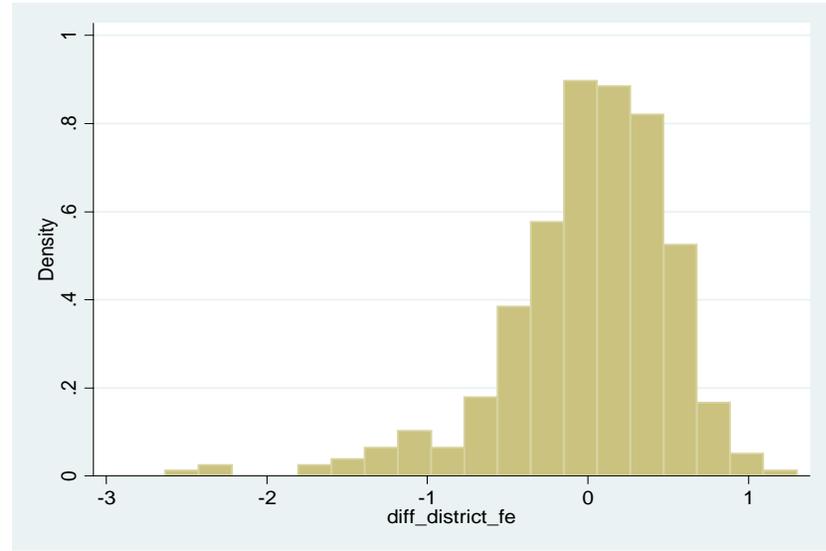
2. Differences in regular and contract teacher characteristics relative to Table 1 reflect (a) the time gap between the 2 sets of data collection of around 3 years, and (b) the fact that the data used for Table 9 comes from villages that had a private school, which tend to be larger than the typical village in AP. The sample in Table 1 is from a representative set of rural government run schools, while the sample in Table 9 is from a sample of villages that have private schools (though the public school data in Table 9 is from the same villages as the private schools in Table 9).

Figure 1: Variation in Log_School_PTR (Using only Regular Teachers) in the Different Specifications in Table 8

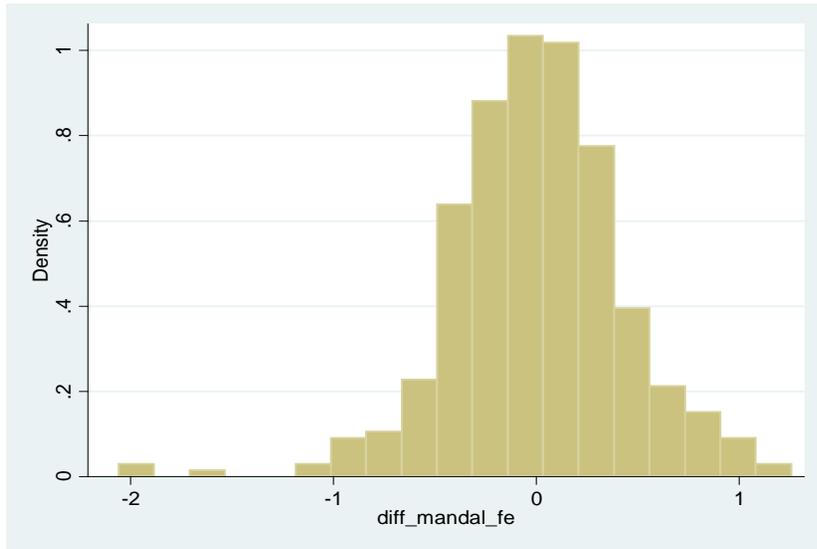
NO FIXED EFFECTS:



DISTRICT FIXED EFFECTS



MANDAL FIXED EFFECTS



SCHOOL-FIXED EFFECTS

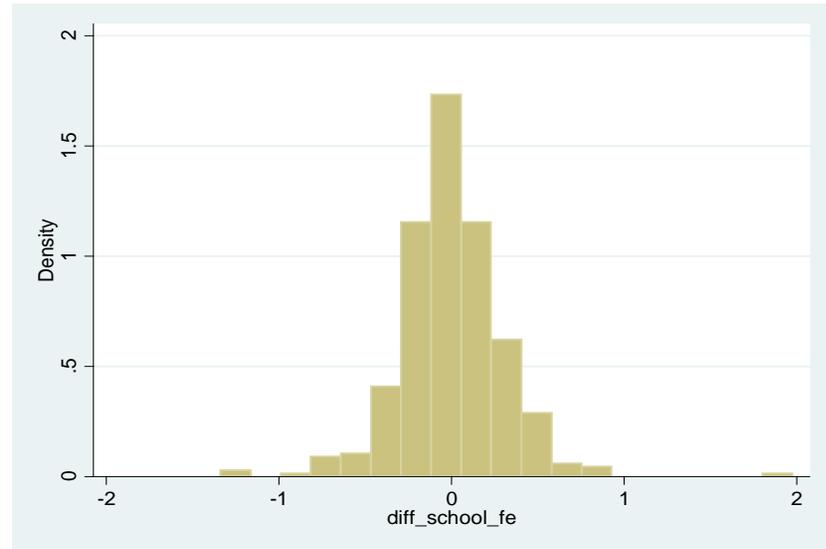
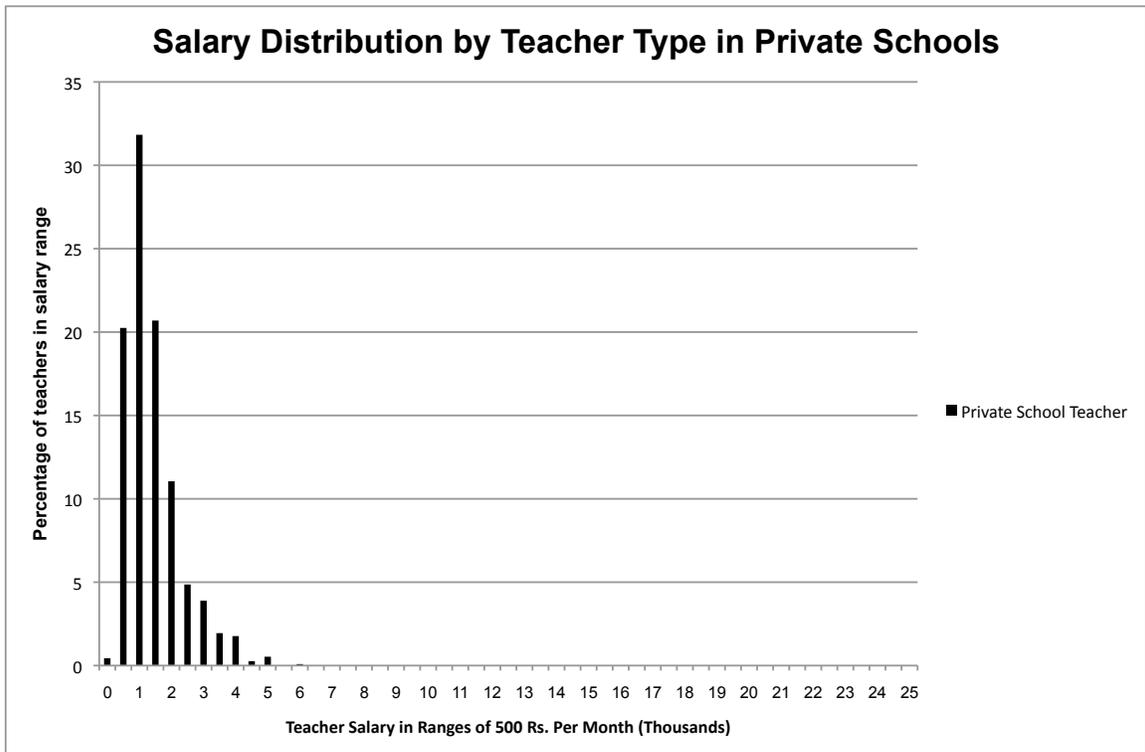
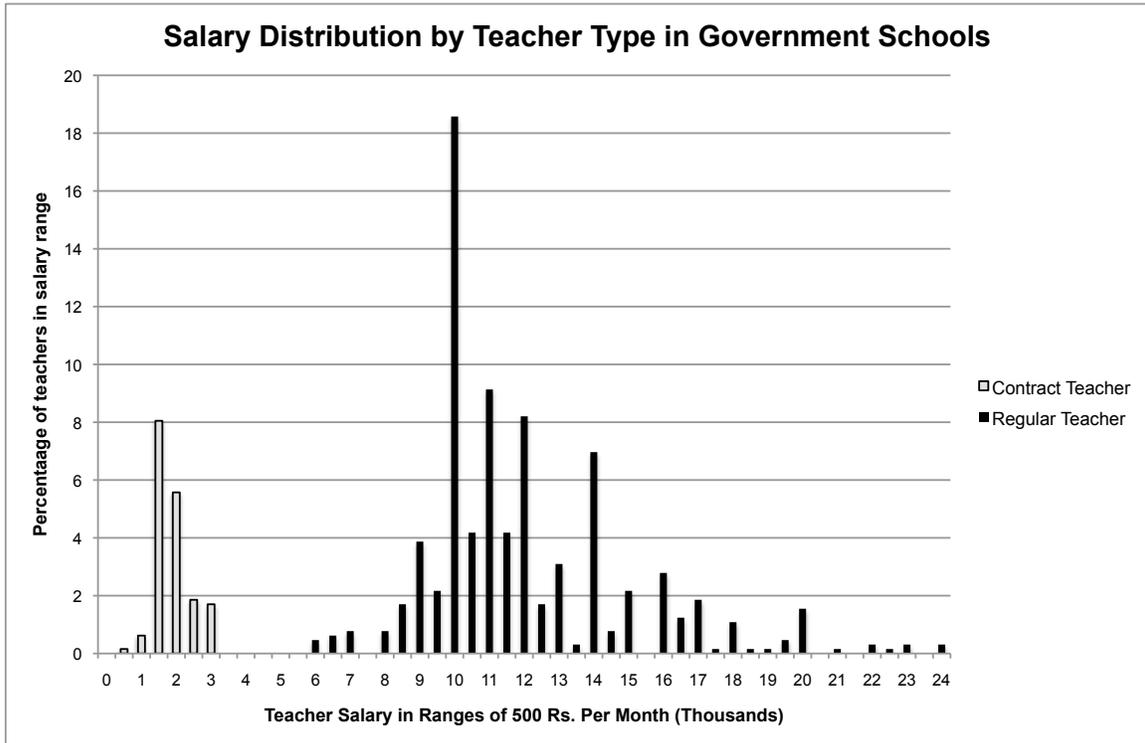


Figure 2: Salary Distribution by School and Teacher Type



Appendices (A1 to A3)

A1. Estimation Sample, and Validity of Identification for Table 8 (Columns 7 and 8)

As noted in the text, our most credible estimate of the impact of reducing school-level PTR with a regular teacher on test score gains are those in Table 8 (Columns 7 and 8). The estimation sample for Table 8 consists of control schools over 5 years, in years in which the schools had only regular teachers (thus the PTR is calculated *only* using regular teachers and is not confounded with the presence of contract teachers). There were a total of 99 control schools over the 5 year period (one of the original 100 got merged with another school during this period) and Appendix Table 1 shows the estimation sample by the number of years in which the concerned control school had only regular teachers (and no contract teachers).

Appendix Table 1: Estimation Sample for Table 8

1	Number of Years with Only Regular Teachers:	5	4	3	2	1	0	Total
2	Number of Schools	43	22	14	12	7	1	99
3	Number of School-Year Level Observations in the Estimation Sample of Table 8	215	88	42	24	7	0	376
4	Number of Adjacent School-Year Level Observations in the Estimation Sample (for Robustness checks in Appendix Tables 2 and 3)	172	51	17	5	0	0	245

We see in Row 2 that 43 control schools were staffed exclusively with regular teachers in all the 5 years, 22 in 4 years, 14 in 3 years, 12 in 2 years, and 7 in only one year. 1 did not have a single year where it had only regular teachers. Row 3 shows the number of school-years in the estimation sample of Table 8 that are contributed by each of these schools (this simply multiplies rows 1 and 2). Thus, out of a total possible number of 495 school-years in the estimation sample (99 x 5), the estimation sample for Table 8 has 376 school-years in it.

The variation in PTR that we use in the specification with school fixed effects comes from (a) changes in student enrollment and cohort sizes over time, and (b) teacher transfers into and out of schools over time. This specification eliminates concerns of unobserved heterogeneity across schools. The main remaining threat to identification in the specification with school fixed effects

is the possibility that school-level PTR changed over time in response to lagged test scores. So for instance, if schools got extra teachers in response to poor test-score performance in the previous year, it is possible that a correlation between PTR reduction and test score gains may simply reflect mean reversion. We show in Appendix Table 2 that there is no correlation between changes in PTR and lagged test scores (with the point estimate being close to zero). This strongly suggests that the identifying variation in PTR is quasi-random, and is uncorrelated with either levels or trajectories of student test scores.

Appendix Table 2: Correlation between changes in Log PTR over time in a school and lagged test scores

Estimation Sample: Control Schools with only Regular Teachers in both the Current and Previous Year		
Dependent Variable: Change in Log PTR		
	[1]	[2]
Lagged Test Score	-0.005 (0.009)	-0.003 (0.008)
Constant	-0.026 (0.019)	-0.026*** (0.000)
Observations	33153	33153
R-squared	0.000	0.241
School Fixed Effects	No	Yes

Notes:

1. All regressions have standard errors clustered at the school level. They also include lagged normalized test scores (these are set to 0 for students in grade 1 who did not have a baseline test). All test scores are normalized relative to the distribution of scores in the control schools in the same grade, test, and year.

2. The sample is limited to those observations for which the current and previous year in the school had only regular teachers.

* significant at 10%; ** significant at 5%; *** significant at 1%

Note that the estimation sample in Appendix Table 2 is smaller than that in Table 8 (Columns 7 and 8) because the table above only uses the school-years where the school had only regular teachers in *adjacent* years (which is necessary to show that there was no correlation between changes in PTR between consecutive years and the test scores at the end of the first of these years). Thus, the test in Appendix Table 2 only uses the sample in Row 4 of Appendix Table 1, while Table 8 (Columns 7, 8) can use the full sample in Row 3 of Appendix Table 1.

We verify that our core results are not affected by this change in sample, by re-estimating Table 8 (Columns 7, 8) using the restricted sample used for Appendix Table 2, and we see below in Appendix Table 3 that the estimates of β_2 in (4) with school fixed effects are unchanged from those in Table 8 (Columns 7, 8). The stability of the estimate of β_2 across all the columns of

Table 8 and in various restricted samples strongly suggest that estimates of β_2 in Table 8 (Columns 7 and 8) can be regarded as the causal effect of changing PTR on annual student test score gains (in a value-added specification).

Appendix Table 3: Re-estimating Table 8 (columns 7 and 8) with the Restricted Sample in Appendix Table 2

Estimation Sample: Control Schools with only Regular Teachers in both the Current and Previous Year		
Dependent Variable: Normalized Student Test Scores (Pooled across Math and Language)		
	[1]	[2]
Log School Pupil-Teacher Ratio (PTR) (β_2)	-0.208** (0.092)	-0.197** (0.089)
Household Controls	No	Yes
School Fixed Effects	Yes	Yes
Observations	27544	22431
R-squared	0.312	0.326
P-value (H0: $\beta_2 = \beta_1$)	0.3587	0.3166
P-value (H0: $\beta_2 = 2*\beta_1$)	0.0477	0.0425
P-value (H0: $\beta_2 = 3*\beta_1$)	0.018	0.0165
P-value (H0: $\beta_2 = 5*\beta_1$)	0.0077	0.0073

Notes:

1. The specification is identical to that in Table 8 - Columns 7, 8; but the estimation sample is restricted to that in Appendix Table 2

* significant at 10%; ** significant at 5%; *** significant at 1%

A2. Comparing Contract and Civil-Service Teachers within Schools

Since we can match students in each year to their teacher and know the teacher type, we can estimate the effect on gains in student learning of being taught by a contract teacher as opposed to a regular teacher using the specification:

$$T_{ijkm}(Y_n) = \alpha + \gamma_j \cdot T_{ijkm}(Y_{n-1}) + \delta_{CT} \cdot CT_{jk} + \beta_X \cdot X_{ijk} + \beta_Z \cdot Z_{k/i} + \varepsilon_k + \varepsilon_{jk} + \varepsilon_{ijk} \quad (5)$$

where the test score variables and error terms are defined as in (1), CT_{jk} is a dummy variable indicating whether the student was in a class taught by a contract teacher, X_{ijk} are a set of

classroom, school and household controls, and δ_{CT} is the parameter of interest, which indicates the extent to which students taught by a contract teacher have different test score gains from those taught by a regular teacher (since the same teacher teaches all subjects within a grade, we pool test scores across math and language in our results).

We estimate (5) using no fixed effects, with school fixed effects, and finally with student fixed effects (using only the sample of students who change teacher type during the course of the two years of the experiment), and find that there is no differential effect on learning gains for students taught by contract teachers relative to those taught by regular teachers, with contract teachers appearing to be more effective in one of the specifications (Appendix Table 4).

These findings further support the main finding of this paper, which is that contract teachers are no less effective than regular civil-service teachers in spite of being less educated, less qualified, and being paid much lower salaries. Nevertheless, as we discuss in the text, there are several challenges in interpreting these results as causal. First, these estimates are based on the *assignment* of teachers to classes, and do not reflect actual day to day teaching *practice*. So it is possible that we may overstate the effectiveness of contract teachers if regular teachers spend time out of their own classrooms in training and coaching them. Conversely, since regular teachers are more absent (as we see in Table 7) it is possible that contract teachers spend time away from their own classrooms covering for absent regular teachers, in which case we may underestimate their effectiveness.

Further, it is also likely that the assignment of teachers to tasks and grades within a school is at least partly determined by factors that are not observable to the econometrician, and which cannot be controlled for (for instance, weaker teachers may be given smaller classes or easier cohorts than stronger teachers). Hence, we show these results in Appendix Table 4 for completeness (since the literature has typically used comparisons based on the type of teacher assigned to students), but our preferred estimates are based on school-level analysis that abstracts away from all within-school comparisons that may be confounded by unobserved factors.

A3. Calculating the Absolute Return to Hiring an Extra Contract Teacher

Recent cross-sectional estimates of the returns to cognitive achievement in India find wage returns of 16% for scoring one σ higher on a standardized math test and 20% for scoring one σ

higher on a standardized language test (Aslam et al. 2011). Assuming that the test score gains in this program correspond to a similar long-term difference in human capital accumulation,¹ the two-year treatment effect would correspond to a 5.5% increase in wages ($0.16\sigma \times 0.16 + 0.15\sigma \times 0.20$). The minimum wage in AP in 2010 was Rs. 112/day. Assuming 250 working days/year yields an annual income of Rs. 28,000 and a 5.5% increase in wage would translate into additional income of Rs. 1,540/year. We treat this as a 40-year stream of *fixed* additional earnings (which is *very conservative* since we don't assume wage growth) and discount at 10%/year to obtain a present value of Rs. 15,060 per student at the time of entering the labor market. Since the average student in our project is 8 years old, we assume that they will enter the labor market at age 20 and further discount the present value by 10% annually for another 12 years to obtain a present value of Rs. 4,800/student. The average school had 65 students who took the tests, which provides an estimate of the total present value of Rs. 312,000. The cost of the program per school for two years was Rs. 25,000, which yields estimated internal rate of return (IRR) of 1250%. If we were to assume that wages would grow at the discount rate, the calculation yields an IRR estimate of 16000%. Thus, depending on assumptions about the rate of wage growth and discount rates, we obtain estimates of an IRR ranging from 1250% to 16000% (or a return ranging from 12.5 to 160 times the cost of hiring an extra contract teacher).

These estimates are clearly a suggestive 'back of the envelope' exercise, and are highly sensitive to the two main assumptions (first, that there will be a similar magnitude of long-term human capital gains as the short-run test score effects; and second, that the cross-sectional wage returns to test scores estimated in Aslam et al. 2011 are causal). This is why our main discussion in the text follows the default approach to cost-effectiveness calculations in the literature on education in developing countries, which has focused on the cost per unit of test-score gains (Dhaliwal et al. 2012).

But we present the estimates on the absolute returns to hiring a contract teacher to show that they are large enough that even if the long-term gains to human capital or the labor market returns to test scores were to be substantially lower, the program would still have a very high rate

¹ Chetty et al. (2013) show that there were significant long-term benefits (in terms of wage gains) to having a 'treatment' of a better teacher (defined as a teacher who produced higher value addition in test scores) even though the test score gains themselves faded out considerably a few years after the intervention. While set in the US context, this work is suggestive of long-term wage gains from interventions that improve test scores in school.

of return. Further, the absolute return is the ideal estimate that a social planner would want because it also helps in thinking about resource allocation across sectors in the economy (not just within education), and whether the investment is worth borrowing for. Thus, even if *existing* education budgets are not reallocated towards many contract teachers in lieu of a single regular teacher, it is likely to be a good public investment to *expand* the education budget to hire more contract teachers over and above the current allocation of teaching resources to schools, because the social returns are likely to be considerably higher than the cost of borrowing.

Appendix Table 4: Comparing the Relative Effectiveness of Regular and Contract Teachers in Increasing Test Scores (within schools)

Dependent Variable: Normalized Student Test Scores (Pooled across Math and Language)						
	No Fixed Effects	School Fixed Effects		Student Fixed Effects		
	[1]	[2]	[3]	[4]	[5]	[6]
Taught by a Contract Teacher	0.026 (0.042)	-0.020 (0.035)	-0.015 (0.034)	0.013 (0.020)	0.023 (0.023)	0.082* (0.047)
Classroom Controls	No	Yes	Yes	No	Yes	Yes
Household Controls	No	No	Yes	N/A	N/A	N/A
Stable Teacher Sample (No Change in teacher grade assignment over the 2 years)	No	No	No	No	No	Yes
Observations	51336	51336	46889	8823	8823	3655
R-squared	0.260	0.349	0.367	0.751	0.753	0.814

Notes:

1. All Regressions include lagged normalized test scores (set to 0 for grade 1 students), with standard errors clustered at the school level
2. Household controls include a household asset index, parent education index, child gender an indicator for being from a scheduled caste/tribe. Classroom controls include class size (in logs), and an indicator for multigrade teaching.
3. Regressions with student fixed effects use the restricted sample of students who switched from being taught by a contract teacher to a regular teacher (or vice versa) over the 2 years (i.e. - the variation is coming from the natural grade progression of students, where they also happen to change teacher type as a result).
4. The Stable Sample refers to the subset of students who switched teacher type over the 2 years, and where the same teacher continues teaching the same class in both years (i.e. - these are student fixed effects estimated off a stable sample of teacher class assignments where the only variation is coming from the grade progression of students).

* significant at 10%; ** significant at 5%; *** significant at 1%