Powering Education*

Fadi Hassan and Paolo Lucchino

Abstract:

More than 1.3 billion people worldwide have no access to electricity. This has first-order effects on several dimensions of development. In this paper, we focus on the link between access to light and education. We randomly distributed solar lamps to 7th grade pupils in rural Kenya and monitored their educational outcomes throughout the year at a quarterly frequency. We find that access to lights through solar lamps is an effective input to education. Once our identification strategy accounts for spillover effects by exploiting variations in treatment at the pupil level, and in treatment intensity across classes, we find a positive and significant effect on treated students as well as a positive and significant spillover effect on control students. In a class with the average treatment intensity of our sample (where 43% of students have solar lamps), treated students experience math grade increases of 0.88 standard deviations. Moreover, we find a positive effect of treatment intensity on control students; increasing a class's proportion of students with lamps by 10% increases the grades of control students by 0.22 standard deviations. We exploit household geolocation to disentangle within-class and geographical spillover effects. We show that geographical spillovers does not have a significant impact, and within-school interaction is the main source of spillover effects. Finally, we provide suggestive evidence that the mechanism through which lamps affect students is by increasing rates of co-studying at school, especially after sunset.

*Fadi Hassan, corresponding author: Trinity College Dublin, TIME, and CEP. Paolo Lucchino: LSE. This project was part of the Global Shapers initiative of the World Economic Forum under the Rome and Nairobi hubs. We want to thank the Powering Impact team, who supervised the project for their outstanding support. We are deeply indebted to Givewatts as our supporting NGO in the field. We thank Enel Foundation and STICERD for funding the project. Paolo Lucchino acknowledges financial support from the ESRC. We have benefited from comments and discussions with Jenny Aker, Oriana Bandiera, John List, Craig McIntosh, Gaia Narciso, Carol Newman, Olmo Silva, Abhijeet Singh, and Chris Udry. We thank seminar participants at the Advances in Field Experiments Conference Chicago University, London School of Economics, Maynooth University, NEUDC-Brown University, Oxford, Trinity College Dublin, and UCD. All remaining errors are ours.

Keywords: Randomized control trial, solar lamps, education, energy access, spillover effects, randomized saturation design. JEL Classifications: O12, I25, C93.

1 Introduction

More than 1.3 billion people worldwide lack access to electricity and 40% of whom live in Sub-Saharan Africa (IEA, 2013). This means that roughly a quarter of humanity lives without electricity to light their homes in the evenings, and are without power in the workplace during the day. This 'energy poverty' strongly constrains the standard of living of such people.

In Africa, the electrical power grid reaches only about 400 million of the continent's 1 billion people. In urban and semi-urban areas, over 30% of people have access to grid electricity. This figure drops to less than 2% in rural areas. The electrical power grid is expanding slowly and unevenly. Governments and the private sector are working to reach deeper into remote areas but financial, political and logistical barriers have proven to be significant obstacles.

The link between energy access and education is an important and under-explored dimension of development. Looking at aggregate data, we can see a strong correlation between rates of electrification and primary school completion (see Figure 1).[1] The lack of access to electric light limits opportunities to study after sunset and it constrains students' allocation of activities throughout the day. In developing countries, it is not uncommon to find students of all ages gathering to read at night under the lights of a gas station or a shop (see Figure 2).[2] However, in rural areas, the lack of such basic infrastructure means that even this option may not be possible. Electrification of rural areas in developing countries is a long and costly process. By the time this occurs, generations of students risk being affected by the lack of lighting, undermining the process of human capital accumulation.

In this project, we evaluate the impact that solar lamps, which are a readily available source of lighting, have on education. We distributed solar lamps

[1] Interestingly, the R-squared of this simple regression of primary education completion on electrification (0.6) is substantially higher than that of educational completion on per-capita income (0.34).

[2] The first picture, which made the headlines of major newspapers, shows Daniel Cabrera, a nine-year-old boy from the Philippines, studying under the lights of a McDonalds restaurant. The second picture was taken in Guinea and was published by the New York Times and the BBC.

to some 7th grade students in off-grid rural areas, randomizing the treatment at the pupil level. This was a novel experiment used to investigate the influence of a potential key educational input on development. Throughout the experiment, we realized that spillovers between treatment and control students were occurring because of co-studying, which was generating knowledge spillovers, and because, in some cases, students were sharing the lamp. To identify these effects, we exploit the variation in whether a student was given a lamp, determined through the random distribution of lamps, as well as the variation in treatment intensity, that is the share of students in a given class who received lamps, which was not part of our randomization, but we argue that it is as good as random.

Randomization at the pupil level and the variation of treatment intensity across classes, allow us to use an identification strategy based on the randomized saturation approach as described in Baird et al. (2014) and McIntosh et al. (2014). We find a positive and significant total effect of treatment. Treated students in a class with the average treatment intensity of our sample, which is about 43%, experience an increase in math grades of 0.88 standard deviations. Moreover, lamp distribution affect control students too - increasing treatment intensity by 10% improves their math grades by 0.22 standard deviations.[3] This implies that access to light through solar lamps has a significant impact on the education and human capital accumulation of students.

We use data on the geolocation of households to analyze geographical spillovers. We exploit the variation in treatment intensity across the geographical areas around pupils induced by the randomization. We do not find robust evidence of geographical spillovers. Spillover effects appear to be driven by within-class interactions between students. Finally, using a survey on students' habits and time use, we find suggestive evidence that the mechanism through which the lamps affected students is by influencing studying habits, especially by increasing co-studying at school after sunset.

[3]All results are robust to randomisation inference. We do not find effects on English, Swahili, science and social science grades. This is not surprising as mathematics was often the focus of study groups at school after regular classes. Our results on mathematics grades are robust to correction for multiple hypothesis testing with the other subjects.

There are two critical features in our identification strategy. Firstly, the variation in treatment intensity was not part of our original design. We exploit the natural variation in treatment intensity across classes that arose from name mismatches between official school transcripts and the baseline school survey that we carried. As we explain later in detail, this administrative issue led to variation in treatment intensity across classes. We show that such administrative glitches are uncorrelated with any outcome or covariate of interest, and we argue that they were random in nature. Therefore, we use this source of variation as part of our identification strategy. Secondly, contrary to the original design of Baird et al. (2014), we did not have a pure control group. We address this issue by imposing a functional form for the effect of treatment intensity on control students, as per McIntosh et al. (2014). We show through non-parametric estimation that a linear functional approximation is appropriate for the interval of treatment intensity in our sample; moreover, a linear approximation is likely to provide a lower bound of our estimates.

Our paper is related to the large body of literature on inputs to education in developing countries. Promoting human capital accumulation is one of the key steps in the development process. The literature shows the importance of building school infrastructure (Duflo, 2001; Burde & Linden, 2013); the relevance of providing free primary education (Lucas & Mbiti, 2012); the effect of subsidies to households and pupils on enrolment (Schultz, 2004; Angrist & Lavy, 2009; Ambler et al., 2015); the impact of monetary incentives on teachers' performance (Muralidharan & Sundararaman, 2011); and the role of learning external technologies such as mobile phones (Aker et al., 2014). Soft inputs like information on schooling returns (Jensen, 2010) and involving parents in school management (Gertler et al., 2012) have also been found to have positive impacts on educational outcomes.

Our paper is similar in spirit to the literature that analyses the role of complementary inputs to education, such as deworming programs (Miguel & Kremer, 2004) and flip-charts (Glewwe, 2002). Given the lack of electrification in developing countries, our contribution is to investigate the importance of access to light as a key input for education and to measure the spillover effects that access to light can have on students. There has been a growing number of studies on the the effects of solar energy on households. In particular, Furukawa (2014) and Kudo et al. (2017) focused on the relation between solar lamps and education in Uganda and Bangladesh, although they do not find an effect on school grades. The main differences between their study and ours, apart from different experimental contexts, is that their design involved students over different grades and their identification strategy does not properly account for spillover effects.

Our paper is also related to the literature on energy access and development, such that of Dinkelman (2011), Rud (2012), Lipscomb et al. (2013), and Lee et al. (2016). These papers are concerned with the effects of energy access on employment, industrialization, human development, housing values, and consumer surplus. These studies examined the impact of electrification (which is a large region-wide technology shock), or an off-grid equivalent solution, whereas we evaluate the effect of providing solar lamps (which is a smaller, idiosyncratic technology shock achieved through a more accessible and cheaper source of energy). Our study complements this field by focusing on education.

Finally, our paper contributes to the broader literature on randomized controlled trials and the role of spillovers. Many experiments are likely to fail or have biased results because of the presence of spillovers, which violates the stable unit treatment value assumption (SUTVA). Our paper provides methodological guidance on how to use a randomized saturation approach, as described in Baird et al. (2014) and McIntosh et al. (2014), in order to account for spillovers even if the experiment was not initially designed for this. This requires variation in treatment intensity that is as good as random, and a reliable approximation of the functional form underlying the relationship between treatment intensity and the dependent variable.

The paper is structured as follows: Section 2 describes the experimental structure; Section 3 discusses lamp usage and attrition; Section 4 shows that a standard identification strategy, that does not account for spillovers, fails to identify a significant intention-to-treat effect; Section 5 focuses on the role of spillovers, by identifying the intention-to-treat effect and spillovers, disentangling the effects of within-class interaction and geographical proximity; Section 6 provides suggestive evidence for the mechanisms underlying the spillover effects; and Section 7 concludes.

2 Project structure, randomization, and source of variation

The experiment involved about 300 students in 7^{th} grade across 13 classes in the Loitokitok and Nzaui districts, relatively close to the Tanzanian border and Mount Kilimanjaro (see Figure 3). We focused on schools in off-grid rural areas where household electrification rates are below 2.6%.

The project started with a baseline survey in June 2013. We interviewed students and collected end of term grades from school transcripts. Lamps were distributed to the treatment group in September 2013, at the beginning of a new school term.[4] We then collected end of term grades for the treatment and control groups in November 2013, March 2014, and June 2014. We also ran extensive interviews of students in November 2013 and March 2014.

The baseline survey covered 341 pupils. We were able to match 286 of these with the transcripts of grades provided by the school. This constituted our core sample, within which we conducted treatment randomization. We distributed solar lamps to 143 pupils and, in order to mitigate resentment and in the interest of fairness, control students were promised that they would receive a lamp at the end of the experiment.[5] We randomized treatment assignment at the pupil level so that each had a mix of treatment and control group. We chose this level of randomization to maximize statistical power, given the budget and the size of our sample. In our randomization strategy, we sought balance between treatment and control groups on marks, which was our dependent variable, and balanced distribution of genders, classes,

[4]The academic year starts in January. This implies that the students in our sample started in 7th and finished in 8th grade. This contributed to attrition, as some students in our sample did not graduate to 8th grade or changed schools. As we discuss in Section 3, attrition is unrelated to treatment status.

[5]Students in the control group received their lamps in September 2014.

and household wealth.[6]

Given our sample size and the number of variables that we wanted to balance, we followed Bruhn & McKenzie (2009) and used a re-randomisation method. We selected an allocation of lamps that minimized the statistical difference in means between control and treatment groups out of 10,000 draws (the so-called *MinMax t-stat* method). We preferred this method to stratification, because our sample size would have constrained the number of variables we could stratify on. In this way, we avoided strong imbalances of several variables without forcing close balance on each. Moreover, we chose re-randomization rather than pairwise matching, because attrition would have posed the risk of losing too many observations, potentially invalidating the experiment. We followed the recommendations of Bruhn & McKenzie (2009) in the econometric analysis, and accounted for our randomization method by including balancing variables in the regression and also running permutation tests to validate our inference, following Rosenbaum (2002).

Table 1 illustrates the regressions of baseline values of the balancing variables on treatment, at the beginning and at the end of our project. The balance between treatment and control groups was well maintained throughout the study. Moreover, we show that our sample is balanced on other relevant variables for which we did not explicitly seek balance, such as hours of studying at baseline, light source, school attendance, and mothers' education. We also run a multiple-hypothesis test as in List et al. (2016) in order to check that balance was maintained across all these variables simultaneously. This is confirmed to be the case.[7]

In order to identify spillover effects, we exploit the variation in class treatment intensity. That is, the proportion of students in a given class who received a lamp. The variation in treatment intensity between classes was not part of our original design, but arose during the process of matching the names of students that we surveyed at baseline, with the names of students listed in the official school transcripts. Starting from the full sample

[6]We constructed a wealth index using a principal component analysis based on household characteristics (e.g. type of walls, water, and toilet facilities) and goods owned (e.g. radio, telephone, bicycle, etc.).

[7]Results available upon request.

of 341 students surveyed at baseline, a match with transcripts was achieved for only 286 students. The match rate differed across classes, leading to a variation in treatment intensity ranging between 14% and 62% (see Table 2).

We argue that the variation in the match rate is as good as random. Mismatches occurred for reasons such as: misspellings of names in the survey; the use of baptismal names in the survey and traditional names in the transcript; and inverting first names and surnames in the transcripts. Our enumerators did not use the transcripts as a reference for names when interviewing students, they directly asked each student what their name was. For example, a mismatch occurred when we had two "Mwendwa Makenzi" names in the survey, due to one being a misspelling of "Mwenda Makenzi". Therefore, we could not distinguish between the two pupils. Similarly, we had a mismatch when "Wambua Kyalo", as reported in the transcript, used his baptismal name, "Jonathan Kyalo", in the survey. Arguably, mismatches from name misspellings and from inverting names and surnames were random. One might wonder whether a student introducing herself with a traditional or a baptismal name was random too. We are not aware of studies that shows that self-reporting of baptismal vs. traditional names is associated to some type of bias. This is especially unlikely in the case of children, who have developed a more limited identity awareness, and in the context of our study, which was characterized by a homogenous ethnic and social setting. From our fieldwork experience, this type of name reporting depends on how many cases of homonymy there are in a class, and has a random nature.

In order to investigate the nature of name mismatches further, Table 3 illustrates whether matched and unmatched students were statistically different in key observable characteristics such as hours of study, wealth, mothers' education, source of light etc.[8] The table clearly shows that there is no significant difference between the two groups and that the matching was unrelated to which interviewer each student had. Table 4 reinforces this point by looking at the balance at the class level. It shows that treatment

[8]We run univariate regressions, over the full sample of students, of the baseline values of a set of dependent variables on a dummy on names' matching.

intensity across classes was balanced over gender, teacher experience, wealth, and most grades.[9] There was a significant imbalance in grades for science, but the content of the Kenyan Primary School syllabus generates little complementarity between science and mathematics, which is the subject for which we find a treatment effect.[10] The coefficient on teacher experience is large in magnitude, although statistically insignificant, but the negative sign work against our finding as classes with higher treatment intensity tend to have less experience.[11] Given the balance across observable characteristics and the random nature of being matched or unmatched, we argue that the source of variation of treatment intensity is as good as random.

3 Treatment compliance and attrition

We ran two student surveys, 3 and 6 months after treatment. During these follow-ups, we asked specific questions about lamp usage and appropriation. In terms of appropriation, 84% of treated students reported that the lamp was kept in their household at night. The remaining 16% said that the lamp was kept in school at night. The lamps were durable and broke in only three cases. In all the other cases the lamps were reported to be in good condition or only had minor problems. Moreover, in all cases, students declared that the solar charge was sufficient for all or most of the activities they wanted to carry out with the lamp. All these elements suggest that compliance was high, implying that intention-to-treat will be very close to the average treatment effect.[12] Additionally, 97% of students declared that studying was

[9] We run univariate regressions of the baseline values of a set of dependent variables on class treatment intensity. Variation in the data is at the class level; hence, given the low number of clusters, both the statistical significance and absolute values of the coefficients matter. The table says that as treatment intensity increases by 10% a class tends to have grades in mathematics that are 0.01 standard deviations lower, with a statistically insignificant value.

[10] Note that science covers topics like vegetation, how to create compost, human diseases and similar issues and not fields like physics or chemistry, which would have complementarities with math. Moreover, if we include grades at baseline in science and social studies as control variables, our results hold.

[11] The coefficient says that moving from a class with 30% to one with 50% of saturation, teachers tend to have about one year less of teaching experience.

[12] We could not systematically check if students sold the lamp. In the first survey, we asked students to bring the lamp to the interview. About 55% of them did, but many declared that the lamp was installed at home in a way that was not easily removable. Indeed, during our field visit, we saw many cases in which the lamp was wired in the

the main activity they used the lamp for.[13]

Despite experimental compliance in terms of lamp appropriation and usage for studying, our experiment had attrition due to missing marks. Marks are our main dependent variable of interest, but these are not always available for all students in our sample. This could be, for example, because they did not sit the end of term exam, or they left the school. Specifically, grades data are missing for 13% of our initial sample in Term 1. This increases to 23% in Term 2, and to 39% in Term 3. After the Term 1 exam, students were promoted from 7th to 8th grade. Unfortunately 16% of students in our sample did not pass the exam and had to repeat 7th grade. This explains a large share of attrition between Terms 2 and 3, but not all of it. We regress a dummy indicator for those repeating 7th grade on treatment and find no statistically significant relationship (Table 5).[14] Moreover, we regress a dummy indicator for students with missing grades on treatment, and find no statistically significant relation. The table also looks into the characteristics of attrition and it finds that students with lower grades at baseline, and females are more likely to drop out of the sample. Whereas for other key characteristics, we did not find a significant pattern that was consistent through different terms. Finally, Table 1 shows that the balance between treatment and control groups over balancing and additional baseline variables was preserved across all terms among students sitting the exam. Therefore, we conclude that attrition is unrelated to treatment and that our results are unlikely to be affected by attrition bias.

4 Intention-to-treat effect: standard estimates

In this section, we run a series of reduced-form regressions to identify the impact of treatment on educational outcomes. Given randomization, the coefficients of the regressions have a causal interpretation. We show that

house and used as a proper lighting fixture. During the fieldwork, we visited households at random and with no notice; in all cases the lamp was in the house. In light of this, we believe lamp resale was minimal, if it happened at all.

[13]40% of the students reported using the lamp to study all subjects equally, 25% to study mainly mathematics, and 20% mainly science.

[14]The coefficient is 0.04 with a p-value of 0.34.

standard specifications, which do not account for spillover effects, fail to find a significant effect of lamp use.

We start our analysis by running an OLS estimation on a pooled cross-section that includes all end-of-term exam scores following treatment. Our basic specification is as follows:

$$y_{ij} = \beta_0 + \beta_1\ Treatment_{ij} + \mathbf{Z_{ij}}\boldsymbol{\gamma} + \lambda_j + \epsilon_{ij} \tag{1}$$

where y_{ij} is the grade of student i in class j; λ_j captures class fixed effects; and $\mathbf{Z_{ij}}$ is a vector of controls that includes student age, mothers' education, and number of siblings. We also include the balancing variables used in the re-randomization as controls.

Then, we extend our analysis to a lagged dependent variable specification. This allows controlling for past grades that, given the cumulative process of education and learning, might influence current grades. We use grades at baseline as the lagged dependent variable of reference. Therefore, we estimate the following regression:[15]

$$y_{ijt} = \beta_0 + \beta_1\ y_{ij0} + \beta_2\ Treatment_{ij} + \mathbf{Z_{ij}}\boldsymbol{\gamma} + \lambda_j + \epsilon_{ijt} \tag{2}$$

Finally, we run a first difference estimation that allows us to control for individual fixed effects. Despite the randomization, this specification offers an important robustness check. The first difference is taken with respect to grades at baseline, so all time-invariant variables between the two periods are controlled for through the transformation.[16] Therefore, we estimate:

$$\Delta y_{ijt} = \beta_0 + \beta_1\ Treatment_{ij} + \epsilon_{ijt} \tag{3}$$

Table 6 summarizes the main findings of these specifications. Given the low number of clusters, we bootstrap the distribution of the test statistics

[15] Also, in this case we include the balancing variables used in the re-randomization as controls.

[16] The controls used in the other specifications are all time-invariant, so they are not included in this case. When controls are included to account for differential trends we still do not find an effect.

using the wild-cluster bootstrap, as in Cameron et al. (2008), and report the associated p-values. Moreover, we also use randomization inference as in Rosenbaum (2002). The results from all of the estimation procedures are consistent. We are unable to detect any treatment effect independently from the specification used and the controls that are added.[17]

5 Accounting for spillover effects

The lack of intention-to-treat effect found in the previous section could be due to the presence of spillovers. Spillovers can arise from: i) lamp sharing, which increases the quantity and/or quality of study time for both treatment and control students; ii) improved learning of treated students, who then share their knowledge with control students; and iii) competition from control students who feel disadvantaged and increase their study effort accordingly. We have evidence of (i), and to some extent (ii), from student surveys and fieldwork experience. Moreover, we cannot rule out the presence of the other spillover sources. This can explain why we do not find evidence of treatment effects by directly comparing the performance of students in treatment and control groups. For this reason, in this section we implement an identification strategy that allows us to account for the presence of spillovers. Moreover, although we cannot distinguish between the sources of spillovers, we are able to disentangle spillovers arising from within-class interaction and from geographical proximity.

5.1 Within-class spillovers

In order to identify spillovers, we use the econometric specification of a randomized saturation design, as proposed by Baird et al. (2014), where saturation is defined as the percentage of students treated within a class (treatment intensity). This methodology allows for the identification of different components of the experimental effect of treatment: spillovers on the control group, spillovers on the treated group, and treatment on the

[17] We also apply Lee bounds to check if attrition can drive the lack of significance in our results, but this is not the case. The lower and the upper bound are below and above zero respectively.

uniquely treated. This methodology involves a two-step randomization process: treatment intensity is firstly randomized across clusters; then, individual treatment is randomized within clusters. As argued above, our first step is as good as random and the second step was randomized explicitly.

However, contrary to the original design of Baird et al. (2014), we do not have a "pure" control group. Therefore, we follow an identification strategy that addresses this limitation as in McIntosh et al. (2014). This involves estimating the pure control outcome by imposing a functional form assumption for the effect of treatment intensity on control students. This means that our estimates rely on an out-of-sample prediction that hinges on the functional specification of the model. As shown in Figure 5, we fit a local polynomial smoother to describe the relationship between grades of control students and treatment intensity. We find a positive relationship, which is what we would expect in the presence of spillovers, and a linear functional form seems to be appropriate for the interval of our data.[18] Our estimates rest on the assumption that a linear specification fits the data between 0% and 14% saturation. However, one might expect that a concave function would capture better the relationship between treatment intensity and grades at very low saturations. This is because providing a few lamps to a class without any lamps is likely to have a stronger effect than providing additional lamps to a class that already has a moderate number of them. This implies that our estimates are more likely to provide a lower bound of the true effect due to Jensen's inequality.[19]

Therefore, given the linear functional form we assume, our econometric specification is similar to that of McIntosh et al. (2014):

$$y_{ijt} = \beta Treatment_{ij} + \mu(TI_j * \delta_t) + \gamma(TI_j * Treatment_{ij} * \delta_t) + \delta_t + s_{ij} + \epsilon_{ijt} \tag{4}$$

where TI_j captures treatment intensity in class j; δ_t is a time dummy variable for the post-treatment period and s_{ij} are individual fixed effects.

[18] As a robustness test we added a squared term on treatment intensity in the main specification presented below, but that delivers insignificant results

[19] In fact Jensen's inequality implies that the concave transformation of a mean is more than or equal to the mean applied after the concave transformation, $E[f(X)] \leq f(E[X])$; hence, with a linear approximation, for any given level of treatment intensity, the probability of a higher grade is lower than or equal to using the true concave function.

Estimating regression (4) as a difference in difference model between a specific term date and grades at baseline is equivalent to estimating this simplified version in first difference:

$$\Delta y_{ijt} = \alpha + \beta \; Treatment_{ij} + \mu \; TI_j + \gamma(TI_j \; * \; Treatment_{ij}) + \epsilon_{ijt} \quad (5)$$

where β is the treatment effect on the uniquely treated (TUT) and captures the theoretical intention-to-treat effect at the point of zero saturation. Defining TI_j as the share of treated students in class j, $TUT = E(\Delta y_{ijt}|Treat_{ij} = 1, TI_j = 0) - E(\Delta y_{ijt}|Treat_{ij} = 0, TI_j = 0)$, where $Treat_{ij}$ indicates if a student i in class j is treated or not. The coefficient μ is the saturation slope for the control group and captures spillovers on the control group: $SC(TI) = E(\Delta y_{ijt}|Treat_{ij} = 0, TI_j = TI) - E(\Delta y_{ijt}|Treat_{ij} = 0, TI_j = 0)$. Variable γ is the differential of the saturation slope for the treated group, and measures the effect of changing saturation on the treated group compared to the control, so that $\mu + \gamma$ captures the spillover on treated students, defined as $ST(TI) = E(\Delta y_{ijt}|Treat_{ij} = 1, TI_j = TI) - E(\Delta y_{ijt}|Treat_{ij} = 1, TI_j = 0)$. This methodology allows us to compute the intention-to-treat measure as the sum of the treatment on uniquely treated and of spillovers on treated students such that $ITT(TI) = E(\Delta y_{ijt}|Treat_{ij} = 1, TI_j = TI) - E(\Delta y_{ijt}|Treat_{ij} = 0, TI_j = 0)$.

The results of this regression are presented in Table 7. As in Burde & Linden (2013), we account for the small number of clusters by i) calculating statistical significance relative to the small sample t-distribution with eleven degrees of freedom while clustering standard errors at the school level; ii) presenting the p-values of the test statistics using the wild-cluster bootstrap procedure as in Cameron et al. (2008); and iii) re-calculating the p-values using randomization inference as in Rosenbaum (2002). The results are consistent across all procedures.[20]

[20] Results are also robust also to multiple hypothesis testing using a Bonferroni method that corrects for correlation across the tested outcomes, as in Aker et al. (2016) and Sankoh et al. (1997). Given the five subjects that students cover and a correlation among other subjects at baseline of 0.43 that was observed at the start of the study, the adjusted p-value is equal to $padj = 1-(1-p(math))^{g(math)}$, where $g(math) = 5^{r(.math)}$ and $r(.math)$ is the average correlation between other subjects excluding math.

The results show a positive and significant effect, such that treated students in a class with average saturation had mathematics score improved by 0.88 standard deviations.[21] The effect increases with the level of saturation and ranges between 0.57 standard deviations at 16% saturation and 1.1 standard deviations at 62% saturation. Moreover, we can see that there are significant, positive spillover effects on the control group. The estimates of μ are positive, significant, and large in magnitude such that a 10% increase in saturation raises the math scorse of the control group by 0.22 standard deviations.

5.2 Geographical spillovers

Potentially, externalities may take place, because interaction occurs not just at school, but also at home or around households. Students live in clusters of houses called *bomas*. There are no roads or illumination to connect bomas, so pupils are unlikely to move between them at night, as they may get lost or encounter wild animals. However, on principle, students could interact around *bomas* during daylight or on their way to and from school. Therefore, in this section we check whether spillovers arose from the geographical proximity of treated and control students' houses.

We then exploit the exogenous variation in the geographical density of treatment across pupils generated by the experiment. We collected the geographical coordinates of the houses where students live, and we used this information to construct a measure of the geographical treatment intensity. For each student, we compute the percentage of treated students within a radii of 500, 1000, and 1500 meters from their homes. This includes both students in the same and in a different classes, with the latter accounting for about 23% of the variation in the data. We rely on the following specification to identify the overall experimental effect accounting for both within-class and

[21]The average class saturation in our sample is 43%. The effect we measure is given by the linear combination of $\beta + 0.43 \times (\mu + \gamma)$. The results are robust to excluding the clusters with the lowest or highest treatment intensity. When dropping the class with the lowest level of saturation the average effect goes down to 0.66 standard deviations and it remains significant at the 10% level. Results do not change when excluding the class with the highest treatment intensity.

geographical externalities, thereby disentangling the two effects:

$$\begin{aligned}\Delta y_{ijt} = \alpha + \beta\ Treatment_{ij} + \mu\ TI_j + \gamma(TI_j\ *\ Treatment_{ij}) + \\ \sigma\ GTI_{ik} + \phi(GTI_{ik}\ *\ Treatment_{ij}) + \epsilon_{ijt}\end{aligned} \tag{6}$$

where GTI_{ik} is the geographical treatment intensity around student i within a radius of $k = 0.5, 1, 1.5\ km$.

Table 8 reports the results of Equation 6.[22] We can see that a positive and significant total effect of treatment is confirmed, such that a treated student in a class and location with average saturation improved their grades by 0.9 standard deviations. The results show a positive but not robustly significant spillover effect on control students arising from geographical proximity to treated pupils (coefficient σ). A 10% increase in geographical treatment intensity within 1km leads to an increase in the grades of control students by 0.047 standard deviations, although the result is not robust to randomization inference. Similarly, geographical spillovers on treated students are also not robust to randomization inference. Finally, the results in Table 8 show that spillovers on control students associated with class treatment intensity remain stable. Overall, we interpret these results as indicating that within-class spillovers, rather than geographical spillovers, account for the bulk of spillover effects.

6 Mechanisms underlying treatment and spillover effects: Suggestive evidence for study habits

The analysis of study habits and the distribution of student activities over the day provides some insight into the underlying mechanism through which lamps can affect treated students and generate spillovers on controls. We find evidence consistent with lamps influencing study habits. Specifically, the availability of lamps appears to trigger increased co-studying at school during the early evenings among both treated and control students.

[22]We have reported the results for geographical treatment intensity within 1km. Measures based on distances of 0.5 or 1.5 kms yielded the same results. Details available upon request.

Our dependent variables of interest refer to study habits. Students were asked: i) if they usually study with other pupils; ii) where they co-study (home vs. school); and iii) at what time of day they study (before or after sunset). We apply the econometric specification in Equation (5) using these responses as the dependent variable. However, in this case we do not have baseline data on the dependent variable, so we run Regression (5) as a single cross-section. Moreover, in this case, the dependent variable is a dummy variable; hence, the regression specification turns into a linear probability model and the coefficients should be interpreted in probabilistic terms.[23] Given the lack of baseline values for the dependent variables and the reliance on a single cross-section with a limited number of observations, we interpret these results only as suggestive evidence of the relationship between study habits and lamp access.

In Table 9 we report the coefficients of the intention-to-treat and the spillover effects on controls for a set of study habits. The results on co-studying are positive and significant in both cases, such that in a class with the average treatment intensity of our sample (43%), the incidence of co-studying in treated students increase by 45 percentage points. Moreover, a 10% increase in treatment intensity raises the incidence of co-studying for control students by 10 percentage points. If we decompose co-studying by location and timing, a stronger effect occurs for the variable *studying with others at school after sunset*. This suggests that an important channel through which the lamp affects student performance is by allowing pupils to study together during a period of the day that was previously less feasible due to the lack of light.

These results are consistent with the responses on lamp sharing and on time use that students gave in our survey. In fact, 48% of treated respondents declared they shared the lamp with other people when studying; 60% of these that they shared the lamps with students of the same class and the remaining shared primarily with siblings. When studying with other students, about 90% of pupils reported doing so at school. Moreover, Figures 6 and 7 show the total treatment effect on treated students in averagely-saturated classes,

[23]As a robustness check, we also run a probit specification and all results are confirmed.

and the spillover effects on controls for specific activities over different hours of the day. The figures report the coefficients estimated using Specification 5 with bands representing standard errors.[24] The results confirm a significant increase in the incidence of studying at school after sunset for both treated and control students, and also a slight reduction in working at home before sunset.

There are various reasons that can explain why the lamps allowed students to spend more time at school in the evening. Anecdotal evidence from experience in the field suggests that the most plausible explanations are that, first, some lamps are used in class, allowing the room to be lit in the darker hours;[25] second, the lamps allow students to walk home safely later in the day, when sunset and darkness are approaching; and, finally, for treated pupils, students are no longer required to go home early to do chores because they, or their parents, can undertake them more efficiently during the evenings with the use of the lamps. The policy implication of the first point is that electrifying schools, so that students can spend more time in school and study together after class, can have a significant impact on human capital accumulation.

As for the impact of co-studying on exam scores, we are unable to determine whether this was due to better lighting itself, or to the benefits of studying together. However, given that only 48% of treated students stated they shared the lamp with other students, lamp sharing is unlikely to account for all the spillover effects. Sharing of knowledge through interaction between students, as well as a competition effects (where control students increase study effort to remain competitive with treated students), are plausible candidates. Further investigation using network data on study partners could help to identify the different sources of spillovers.

An alternative explanation for the mechanism underlying the impact of the lamps is that it could be related to the income effect that the lamps generate.

[24] We report cluster-adjusted standard errors. For ease of illustration we have not reported the standard errors from randomization inference, but the main results are robust to the use of this approach.

[25] In some cases, teachers were keeping the lamps in the school, and in other cases students occasionally brought the lamps with them to school.

Das et al. (2013) showed that increasing school inputs may affect household spending responses and, in turn, learning outcomes. The lamp can help generate savings on other lighting fuels, kerosene in particular. Indeed, evidence from student surveys and household expenditure surveys indicates that families with treated students experience a reduction in fuel expenditure of about 60–90 Ksh ($0.66–$1) per week. This is equivalent to around 10–15% of the median weekly income of the households in our sample. Moreover, time-use analysis of parents shows that the lamps allow mothers to do chores more effectively at night, freeing time for other activities, especially paid work, during the day and in the evenings.

To explore the possibility that improved learning outcomes could be attributed to income effects associated with the lamp, we ran a household expenditure survey after one year at the end of our experiment and we did not find significant differences across expenditure categories between treatment and control groups. So the mechanisms highlighted by Das et al. (2013) do not seem to hold in this context. Our findings are not inconsistent with their results, given that they found an effect on household expenditure in the second year, and only if the input was anticipated. Our survey expenditure was conducted after the first year and the input was not anticipated. Additionally, income effects would only explain spillover effects on grades if the income effect itself spills over onto control households. These considerations strengthen our confidence that the income effect on grades is unlikely to explain the observed effects of the lamps on grades.

7 Conclusions

This study presents a novel experiment that assed the effect of access to light on education. Through a randomized controlled trial, we document an overall positive effect of solar lamps on education in rural Kenya. Once our identification strategy takes into account the potential presence of spillovers, we are able to find a positive and significant intention-to-treat effect and a positive and significant spillover effect on the control group.

Given the small scale of the technology shock that our experiment provided, all our estimates are likely to represent a lower bound of the true effect of lighting, and energy access in general, on education. Moreover, any experimental issues, like lamp appropriation by teachers and lamp sharing by students of different classes, would likely bias our estimates downwards. Despite their advantages, solar lamps should not be seen as a substitute for electrification. However, they are a short-term, practical solution for limiting the drawbacks to human capital accumulation that result from a lack of electricity.

We are also able to disentangle within-class and geographical spillovers. Most of the spillovers arose from within-class interactions, while the geographical proximity of treated and control students did not have a statistically robust effect. The mechanisms through which spillovers arise seem to be related to increased co-studying at school, especially after sunset. This suggests that solar lamps can be particularly effective when there are study groups at school after lectures. Further research into this topic with larger samples and in different settings, could provide more detailed understanding such as the mechanisms that enhance or limit the effects of light access on education.

References

Aker, J. C., Boumnijel, R., McClelland, A., & Tierney, N. (2016). Payment mechanisms and antipoverty programs: Evidence from a mobile money cash transfer experiment in niger. *Economic Development and Cultural Change*, 65(1), 1–37.

Aker, J. C., Ksoll, C., & Lybbert, T. J. (2014). Can Mobile Phones Improve Learning? Evidence from a Field Experiment in Niger. *American Economic Journal: Applied Economics*, 4(4), 94–120.

Ambler, K., Aycinena, D., & Yang, D. (2015). Channeling Remittances to Education: A Field Experiment among Migrants from El Salvador. *American Economic Journal: Applied Economics*, 7(2), 207–32.

Angrist, J. & Lavy, V. (2009). The Effects of High Stakes High School Achievement Awards: Evidence from a Randomized Trial. *American Economic Review*, 99(4), 1384–1414.

Baird, S., Bohren, J. A., McIntosh, C., & Ozler, B. (2014). *Designing Experiments to Measure Spillover Effects*. SSRN Scholarly Paper ID 2505070, Social Science Research Network, Rochester, NY.

Bruhn, M. & McKenzie, D. (2009). In Pursuit of Balance: Randomization in Practice in Development Field Experiments. *American Economic Journal: Applied Economics*, 1(4), 200–232.

Burde, D. & Linden, L. L. (2013). Bringing Education to Afghan Girls: A Randomized Controlled Trial of Village-Based Schools. *American Economic Journal: Applied Economics*, 5(3), 27–40.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics*, 90(3), 414–427.

Das, J., Dercon, S., Habyarimana, J., Krishnan, P., Muralidharan, K., & Sundararaman, V. (2013). School Inputs, Household Substitution, and Test Scores. *American Economic Journal: Applied Economics*, 5(2), 29–57.

Dinkelman, T. (2011). The Effects of Rural Electrification on Employment: New Evidence from South Africa. *American Economic Review*, 101(7), 3078–3108.

Duflo, E. (2001). Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *The American Economic Review*, 91(4), 795–813.

Furukawa, C. (2014). Do Solar Lamps Help Children Study? Contrary Evidence from a Pilot Study in Uganda. *Journal of Development Studies*, 50(2), 319–341.

Gertler, P. J., Patrinos, H. A., & Rubio-Codina, M. (2012). Empowering parents to improve education: Evidence from rural Mexico. *Journal of Development Economics*, 99(1), 68–79.

Glewwe, P. (2002). Schools and Skills in Developing Countries: Education Policies and Socioeconomic Outcomes. *Journal of Economic Literature*, 40(2), 436–482.

IEA (2013). *World Energy Report*. Technical report.

Jensen, R. (2010). The (Perceived) Returns to Education and the Demand for Schooling. *Quarterly Journal of Economics*, 125(2), 515–548.

Kudo, Y., Shonchoy, A. S., & Takahashi, K. (2017). *Can Solar Lanterns Improve Youth Academic Performance? Experimental Evidence from Bangladesh*. Policy Research Working Paper 7954, The World Bank.

Lee, K., Miguel, E., & Wolfram, C. (2016). *Experimental Evidence on the Demand for and Costs of Rural Electrification*. Working Paper 22292, NBER.

Lipscomb, M., Mobarak, A. M., & Barham, T. (2013). Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil. *American Economic Journal: Applied Economics*, 5(2), 200–231.

List, J. A., Shaikh, A. M., & Xu, Y. (2016). *Multiple Hypothesis Testing in Experimental Economics*. Working Paper 21875, National Bureau of Economic Research.

Lucas, A. M. & Mbiti, I. M. (2012). Access, Sorting, and Achievement: The Short-Run Effects of Free Primary Education in Kenya. *American Economic Journal: Applied Economics*, 4(4), 226–53.

McIntosh, C., Alegra, T., Ordez, G., & Zenteno, R. (2014). *Infrastructure upgrading and budgeting spillovers: Mexico's Habitat experiment*. Technical Report 036, UC Berkeley, Center for Effective Global Action.

Miguel, E. & Kremer, M. (2004). Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities. *Econometrica*, 72(1), 159–217.

Muralidharan, K. & Sundararaman, V. (2011). Teacher Performance Pay: Experimental Evidence from India. *Journal of Political Economy*, 119(1), 39–77.

Rosenbaum, P. R. (2002). Covariance Adjustment in Randomized Experiments and Observational Studies. *Statistical Science*, 17(3), 286–327.

Rud, J. P. (2012). Electricity provision and industrial development: Evidence from India. *Journal of Development Economics*, 97(2), 352–367.

Sankoh, A. J., Huque, M. F., & Dubey, S. D. (1997). Some comments on frequently used multiple endpoint adjustment methods in clinical trials. *Statistics in Medicine*, 16(22), 2529–2542.

Schultz, T. P. (2004). School subsidies for the poor: evaluating the Mexican Progresa poverty program. *New Research on Education in Developing Economies*, 74(1), 199–250.

Figures

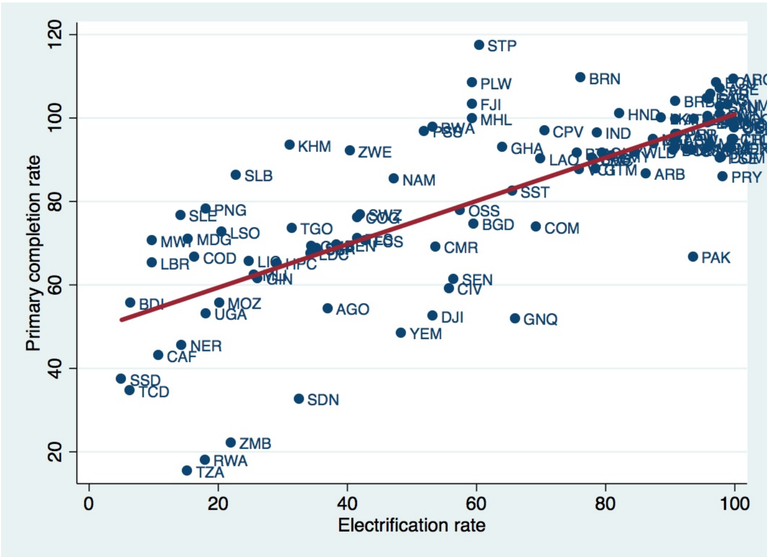

Figure 1: Electricity access and primary schooling completion rate (WDI data, electricity<100%)

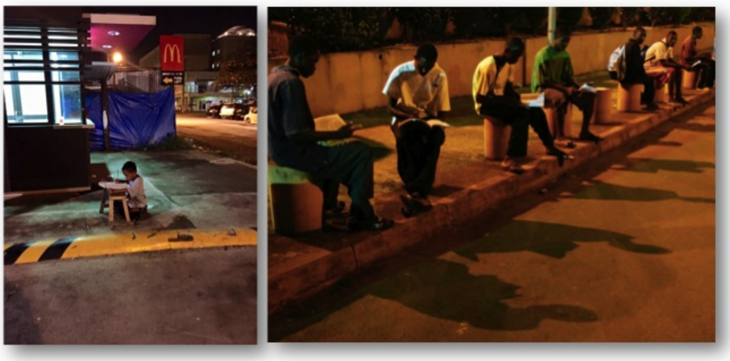

Figure 2: Students studying with a lack of electrification (Associated Press and Facebook).

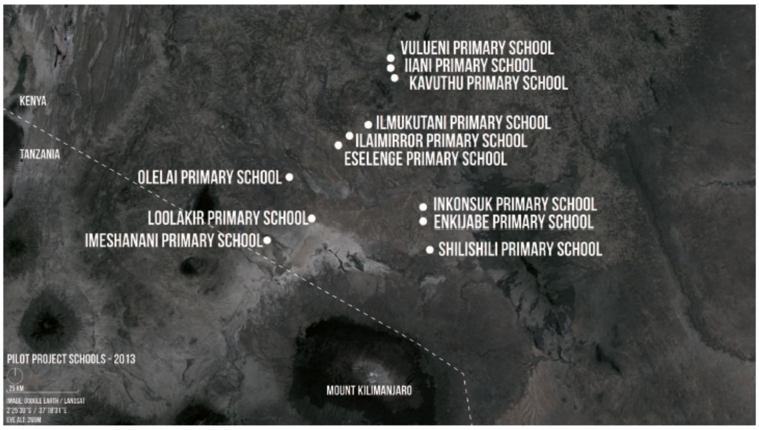

Figure 3: Map of study area

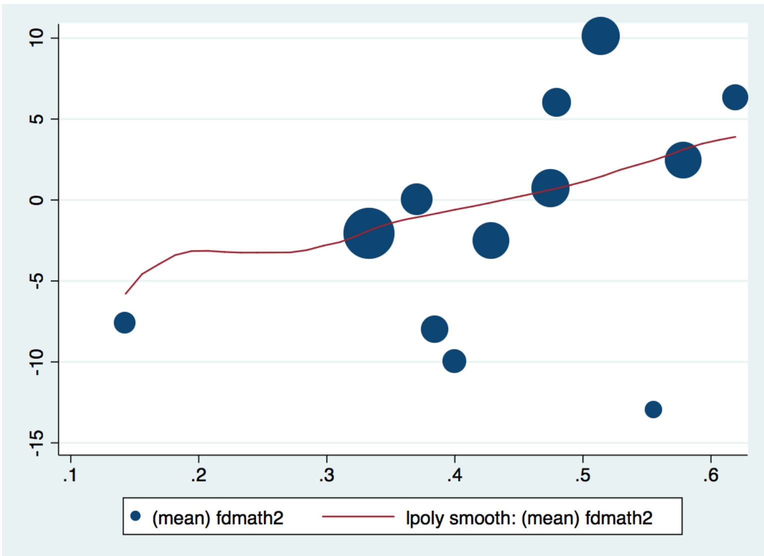

Figure 5: Local polynomial smoother fitted to control groups' grades and class treatment intensity

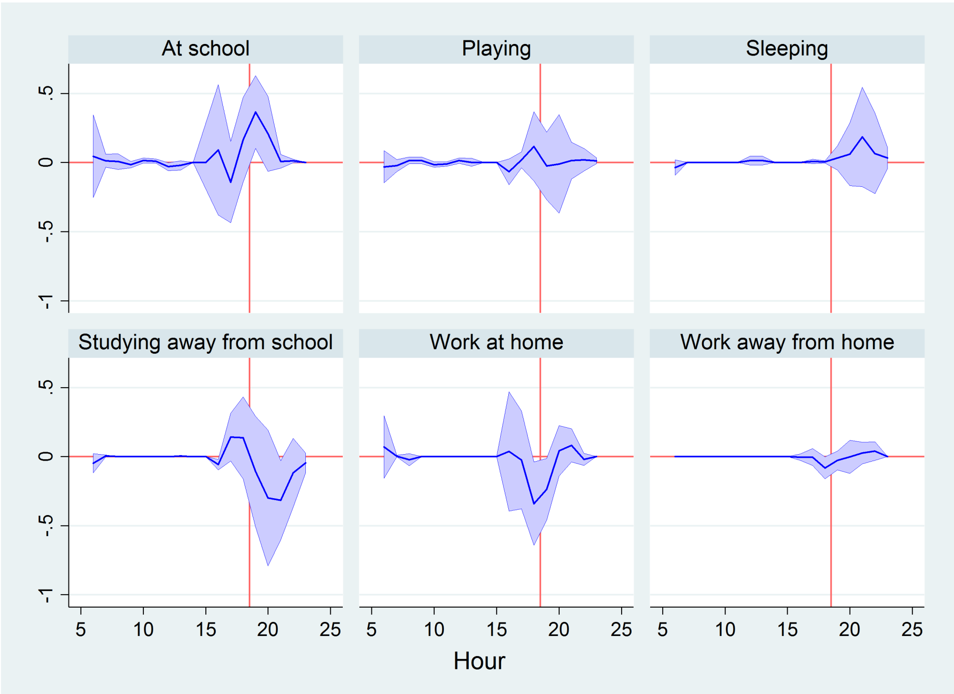

Figure 6: Students' activities by time of the day. Estimates of intention-to-treat effect for treated students at average class saturation (shaded area = clustered standard errors).

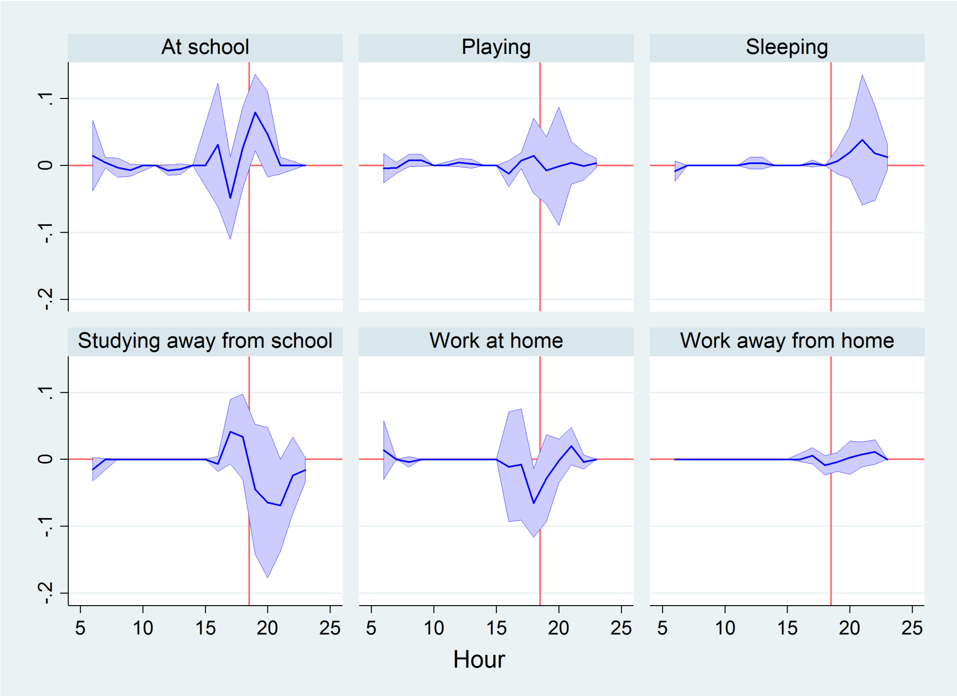

Figure 7: Students' activities by time of the day. Estimates of spillover effects on control students (shaded area = clustered standard errors).

Tables

Table 1: Balance between treatment and control on baseline variables

Explanatory variable: treatment	**Initial randomization**		**End of experiment**	
Balanced variables	Coefficient	*p*-value	Coefficient	*p*-value
Mathematics	0.16	0.20	0.22	0.21
English	-0.05	0.66	-0.07	0.60
Kiswahili	-0.02	0.83	-0.08	0.55
Science	0.04	0.75	-0.08	0.95
Social Studies	-0.12	0.29	-0.15	0.33
Gender	0.00	0.95	0.00	0.99
Wealth index	0.02	0.79	0.02	0.82
School 1	-0.06	0.22	-0.03	0.64
School 2	0.01	0.49	0.03	0.26
School 3	0.04	0.22	0.03	0.58
School 4	0.04	0.17	0.08*	0.08
School 5	0.02	0.70	-0.04	0.37
School 6	-0.03	0.54	0.00	0.97
School 7	-0.03	0.13	-0.04	0.17
School 8	0.02	0.48	0.02	0.61
School 9	0.02	0.60	-0.02	0.49
School 10	0.00	0.99	0.02	0.68
School 11	-0.02	0.40	-0.04	0.29
School 12	-0.04	0.22	-0.02	0.49
Additional variables				
Hours of study	0.11	0.79	0.16	0.76
Missed days of schools (previous month)	0.06	0.78	0.5*	0.07
Source of studying light: wood/candle	0.00	0.97	0.05	0.55
Source of studying light: kerosene	-0.04	0.69	-0.05	0.60
Mother's education	0.04	0.44	0.06	0.38

***. **, *, denote significance at the 1%level, 5% level, and 10% levels, respectively.

Table 2: Treatment intensity variation across classes

	Treatment intensity	Class size
Class 1	14.20%	7
Class 2	33.3%	18
Class 3	33.3%	36
Class 4	37.0%	27
Class 5	38.4%	13
Class 6	40.0%	15
Class 7	42.8%	28
Class 8	47.5%	40
Class 9	48.0%	25
Class 10	51.4%	35
Class 11	55.5%	9
Class 12	57.8%	38
Class 13	61.9%	21

Table 3: Balance between matched and unmatched students' baseline characteristics

Explanatory variable: unmatched student	Coefficient	p-value
Hours of study	-0.35	0.48
Missed school days	-0.16	0.55
Number of people in household	0.31	0.63
Wealth index	0.06	0.12
Source of study light: kerosene	0.02	0.79
Mother education	0.05	0.41
Interviewer 1	0.03	0.56
Interviewer 2	-0.05	0.14
Interviewer 3	0.01	0.69
Interviewer 4	0.06	0.20
Interviewer 5	-0.03	0.57
Interviewer 6	-0.03	0.37
Interviewer 7	0.02	0.70
Interviewer 8	0.03	0.51
Interviewer 9	-0.01	0.72
Interviewer 10	-0.03	0.37

Table 4: Balance of treatment intensity across classes, average values per class at baseline

Explanatory variable: treatment intensity	Coefficient	p-value
Mathematics	-0.13	0.94
English	1.04	0.51
Swahili	-0.30	0.75
Science	3.59***	0.00
Social studies	-1.13	0.25
Gender	0.23	0.54
Wealth index	-0.97	0.17
Teacher experience	-6.5	0.54
Teacher education	-0.17	0.57
Hours of study	-0.79	0.72
Number of children in household	1.04	0.40
Source of study light: kerosene	0.33	0.37
Mothers' education	-1.69	0.32
Observations	13	

***significant at the 1%level; * significant at the 10% level.

Table 5: Attrition by term

Y: Missing the exam (attrition)	Term 1	Term 2	Term 3
Treatment	0.01	-0.03	-0.03
	(0.72)	(0.50)	(0.50)
Grades at baseline	-0.05*	-0.05*	-0.08**
	(0.09)	(0.10)	(0.03)
Gender	-0.08	-0.19***	-0.2**
	(0.14)	(0.00)	(0.03)
Age	-0.02*	-0.02	-0.01
	(0.09)	(0.25)	(0.55)
Number of children in household	0.004	0.005	-0.02
	(0.66)	(0.64)	(0.12)
Mothers' education	0.14**	0.02	-0.04
	(0.04)	(0.73)	(0.38)
Fathers' education	0.14*	0.05	-0.03
	(0.08)	(0.44)	(0.70)
Wealth index	-0.06	-0.05	0.1
	(0.30)	(0.59)	(0.16)
Observations	286	286	286

***significant at the 1%level; ** significant at the 5% level; * significant at the 1% level. Clustered standard errors at the school level. *P*-values in parentheses.

Table 6: Intention-to-treat effect - Pooled regressions

	Cross-section		**Lagged dependent variable**		**First difference**
Y: Grades in mathematics	(1)	(2)	(1)	(2)	(1)
Treatment	0.047	0.048	-0.024	-0.008	-0.01
	(0.60)	(0.54)	(0.77)	(0.91)	(0.92)
	[0.61]	[0.56]	[0.77]	[0.88]	[0.95]
	{0.3}	{0.38}	{0.58}	{0.48}	{0.8}
Age		-0.057		-0.061	
		(0.26)		(0.10)*	
		[0.30]		[0.23]	
		{0.65}		{0.46}	
Mothers' education		0.038		-0.1	
		(0.88)		(0.27)	
		[0.91]		[0.42]	
		{0.91}		{0.44}	
Number of siblings		0.025		0.027	
		(0.15)		(0.18)	
		[0.11]		[0.18]	
		{0.11}		{0.18}	
Grades at baseline			0.61	0.63	
			(0.00)***	(0.00)***	
			[0.00]***	[0.00]***	
			{0.25}	[0.48]	
Observations	646	582	639	575	641

***significant at the 1%level; ** significant at the 5% level; * significant at the 10% level. P-values from clustered standard errors at the school level in parentheses () and p-values from permutation testing in brackets []. The dependent variable is the standardized grade in mathematics. All specifications account for class fixed effects and balancing variables.

Table 7: RSD estimates with class saturation, pooled sample

Y: Grades in Mathematics	
β, treatment on uniquely treated	0.37 (0.25) [0.35] {0.21}
μ, saturation slope on control	2.21 (0.026)** [0.09]* {0.06}*
γ, differential saturation slope on treatment	-1.03 (0.18) [0.26] {0.82}
Intention-to-treat:	
- Min saturation (16.6%)	0.57 (0.06)* [0.06]* {0.08}*
- Average saturation (43%)	0.88 (0.028)** [0.06]* {0.10}*
- Max saturation (62%)	1.1 (0.035)** [0.047]** {0.09}*
Spillover effects:	
Marginal effect of 10% higher treatment intensity on *control* students	0.22 (0.026)** [0.09]* {0.07}*
Marginal effect of 10% higher treatment intensity on *treated* students	0.11 (0.18) [0.21] {0.22}
Observations	641

***significant at the 1%level; ** significant at the 5% level; * significant at the 1% level. *P*-values from clustered adjusted standard errors at the school level in parentheses (), *p*-values from wild-bootstraps in brackets[], *p*-values from randomization inference in braces {}. The dependent variable is the standardized grade in mathematics.

Table 8: RSD with class and geographical saturation, pooled sample

Y: Grades in mathematics	
β, treatment on uniquely treated	0.52
	(0.23)
	[0.30]
	{0.15}
μ, class saturation slope on control	1.86
	(0.067)*
	[0.046]**
	{0.09}*
γ, differential class saturation slope on treatment	-1.33
	(0.27)
	[0.33]
	{0.87}
σ, geo saturation slope on control	0.47
	(0.25)
	[0.28]
	{0.26}
ϕ, geo saturation slope on treatment	-0.05
	(0.9)
	[0.9]
	{0.57}
Intention-to-treat:	
Average class saturation (43%)	0.9
& Average geo saturation (37%)	(0.00)***
	[0.00]***
	{0.087}*
Spillover effects:	
Marginal effect of 10% higher	0.18
class treatment intensity on *control* students	(0.067)*
	[0.046]**
	{0.087}*
Marginal effect of 10% higher	0.04
geo treatment intensity on *control* students	(0.25)
	[0.28]
	{0.087}*
Marginal effect of 10% higher	0.05
class treatment intensity on *treated* students	(0.63)
	[0.69]
	{0.087}*
Marginal effect of 10% higher	0.04
geo treatment intensity on *treated* students	(0.35)
	[0.37]
Observations	521

***significant at the 1%level; ** significant at the 5% level; * significant at the 1% level. P-Values from clustered adjusted standard errors at the school level in parentheses (), p-values from wild-bootstraps in brackets[], p-values from randomisation inference in braces {}. The dependent variable is the standardized grade in mathematics.

Table 9: RSD estimates, effects on study habits

	Total treatment effect at average saturation (43%)	Spillovers on control, marginal effect of a 10% increase in saturation
Dependent variable:		
Study with others	0.45	0.10
	(0.01)**	(0.01)***
	[0.02]**	[0.02]**
Study with others at school after sunset	0.31	0.05
	(0.07)*	(0.19)
	[0.04]**	[0.11]
Study with others at school before sunset	-0.00	0.03
	(1.00)	(0.58)
	[0.52]	[0.28]
Study with others at home after sunset	0.16	0.03
	(0.11)	(0.19)
	[0.13]	[0.19]
Study with others at home before sunset	0.01	0.00
	(0.35)	(.)
	[0.38]	[0.45]
Observations:	254	254

P-values from clustered adjusted standard errors at the school level in parentheses () and p-values from randomization inference in brackets[]. * p<0.1, ** p<0.01, ***p<0.001.