Powering Education*

Fadi Hassan and Paolo Lucchino

Abstract:

More than 1.3 billion people worldwide have no access to electricity and this has first-order effects on several development dimensions. In this paper we focus on the link between access to light and education. We randomly distribute solar lamps to 7th grade pupils in rural Kenya and monitor their educational outcomes throughout the year at quarterly frequency. We find that access to lights through solar lamps is a relevant and effective input to education. Our identification strategy accounts for spillovers by exploiting the variation in treatment at the pupil level and in treatment intensity across classes. We find a positive and significant intention-to-treat effect as well as a positive and significant spillover effect on control students. In a class with the average treatment intensity of our sample (43%), treated students experience an increase in math grades of 0.88 standard deviations. Moreover, we find a positive marginal effect of treatment intensity on control students: raising the share of treated students in a class by 10% increases grades of control students by 0.22 standard deviations. We exploit household geolocation to disentangle within-class and geographical spillovers. We show that geographical spillovers do not have a significant impact and within-school interaction is the main source of spillovers. Finally, we provide suggestive evidence that the mechanism through which lamps affect students is by increasing 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's hubs. We want to thank the team of Powering Impact that overviewed the full project for their outstanding support. We are deeply indebted to Givewatts as our supporting NGO on 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, 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: Randomised control trial, solar lamps, education, energy access, spillover effects, randomised saturation design. JEL Classifications: O12, I25, C93.

1 Introduction

More than 1.3 billion people worldwide lack access to electricity and 40% of them live in Sub-Saharan Africa (IEA, 2013). This means that roughly a quarter of humanity lives without lights at home in the evening and without power at the workplace during the day. Energy poverty implies that most people are strongly constrained in their standards of living.

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 to overcome.

The link between energy access and education is an important and underexplored dimension of development. Looking at aggregate data, we can see a strong correlation between electrification rate and the completion of primary schooling (see Figure 1).[1] The lack of access to light limits the possibility of studying after sunset and puts constraints on the time distribution of activities by students. 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. Electrifying 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 distribute solar lamps to 7th grade students in off-grid rural areas, randomising treatment at the pupil

[1]Interestingly, the R-squared of this simple regression of primary education completion on electrification is double than the one of education completion on income per-capita.

[2]The first picture, which made the headlines of major newspapers, refers to Daniel Cabrera, a nine-year-old boy from the Philippines, who is studying under the lights of a McDonald. The second picture is taken in Guinea and has been reported by the New York Times and BBC.

level. This is a novel experiment on a potentially key educational input for development. As lamps can be easily shared, we are interested in identifying both the effects on the individuals who receive them and the spillovers on their peers. To identify these effects we exploit the variation in whether a student was given the lamp, determined through the randomisation of its distribution, as well as the variation in treatment intensity, that is the share of students in a given class who received the lamp.

The randomization at the pupil level and the variation of treatment intensity across classes allow us to use an identification strategy based on the randomised saturation approach as described in Baird et al. (2014) and McIntosh et al. (2014). We find a positive and significant intention-to-treat effect, such that 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, the lamp affects control students too, such that increasing treatment intensity by 10% increases 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 the human capital accumulation of students.

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

There are two critical features in our identification strategy. Firstly, the variation in treatment intensity comes from names' mismatches of students between the baseline survey and the official transcripts of grades provided by the schools. As we explain later in details, this administrative issue led to variation in treatment intensity across classes. We show that such

[3]All results are robust to randomisation inference. We do not find effects on English, Swahili, Science and Social Science.

administrative glitches are uncorrelated with any outcome and covariate of interest and we argue that they are random in nature; so we use this source of variation as part of our identification strategy. Secondly, contrary to the original design of Baird et al. (2014), we do not have a pure control group. Therefore, we address this issue by imposing a functional form for the effect of treatment intensity on control students as in 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 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); and the impact of monetary incentives on teachers' performance (Muralidharan & Sundararaman, 2011). 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 a positive impact on educational outcomes. Our paper is closer 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.[4]

Our paper is also related to the literature on energy access and development, such as Dinkelman (2011), Rud (2012), and Lipscomb et al. (2013). These papers are concerned with the effects of energy access on employ-

[4] Furukawa (2014) investigates the role of solar lamps on educational outcomes in Uganda and he finds no effect. This is likely to be driven by a combination of factors such as a i) significantly lower sample size respect to our paper (more than a third), ii) a randomization structure that involve students in different grades, and iii) importantly he does not account for spillover effects even if randomization is at the student level.

ment, industrialisation, human development and housing values. Our study complements this field by focusing on education. An important distinction, however, is that these studies examine the impact of electrification, which is a large region-wide technology shock, whereas we evaluate the effect of providing solar lamps, which is a smaller and idiosyncratic technology shock that relates to a more easily available and cheaper source of energy access.

Finally, our paper speaks to the broader literature on randomised control trials. Many experiments are likely to fail or have biased results because of the presence of spillovers. Our paper provides methodological guidance on how to use a randomised 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 that. This requires variation in treatment intensity that is as good as random and being able to approximate the functional form underlying the relationship between treatment intensity and the dependent variable.

The paper is structured as follows: Section 2 describes the experiment 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 and identifies the intention-to-treat effect and spillovers disentangling the effects of within-class interaction and geographical proximity; Section 6 provides suggestive evidence on the mechanism underlying the spillover effects; and Section 7 concludes.

2 Project structure, randomisation, and source of variation

The experiment involves about 300 students in 7$^{\text{th}}$ grade across 13 classes in the Loitokitok and Nzaui districts, relatively close to the Tanzanian border and mount Kilimanjaro (see Figure 3). We focus on schools in off-grid rural areas where household electrification is below 2.6%.

The project started with a baseline survey in June 2013 when 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.[5] 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 to 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 and they constitute our core sample, over which we conducted the randomisation. 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.[6] We randomise assignment to treatment at the pupil level so that within each class some students were in the treatment and some in the control group. We chose this level of randomisation to maximise statistical power, given the budget and the size of our sample. In our randomisation strategy, we seek balance between treatment and control groups on grades, which is our dependent variable, gender, classes, and a proxy for wealth.[7]

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 where we selected the allocation of lamps that minimised the statistical difference in means between control and treatment out of 10,000 draws (the so called *MinMax t-stat* method). We prefer this method to stratification, because our sample size would have constrained the number of variables we could stratify on. In this way, we avoid strong imbalance on several variables without forcing close balance on each. Moreover, we chose re-randomisation rather than pairwise matching, because attrition would

[5] The new academic years start 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.

[6] Students in the control group received the lamp in September 2014.

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

posed the risk of leading to the loss of too many observations, potentially invalidating the experiment. We follow recommendations in Bruhn & McKenzie (2009) in the econometric analysis, and account for our randomisation method by including balancing variables in the regression and also running permutation tests to validate our inference following Rosenbaum (2002).

Table 1 reports regressions of the baseline values of the balancing variables on treatment over the sample at the beginning and at the end of our project; the balance between treatment and control was well maintained throughout our study. Moreover, we show that our sample is balanced also on other relevant variables for which we did not explicitly seek balance, like hours of studying at baseline, light source, school attendance, and mother's education. We also run a multiple-hypothesis test as in List et al. (2016) in order to check that balance hold across all these variables simultaneously and this turns to be the case.[8]

In order to identify spillover effects we exploit the variation in class treatment intensity, that is the share of students in a given class who received the lamp. The variation in the treatment intensity between classes arose during the process of matching students' survey at baseline with school transcripts. Starting from the full sample 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 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 name and surname in the transcripts. Figure 4 provides an example of a transcript.[9] Our enumerators did not use the transcripts as a reference for names when interviewing students, but they directly asked the students what their name was. So, for example, a mismatch occurred when we had two "Mwendwa Makenzi" in the survey, but that was a misspelling as

[8]Results available upon request.

[9]School managers agreed to make the transcripts public for the purpose of this study and in most classes the transcripts are publicly available as they are posted on a wall outside the classroom.

one of them should have been “Mwenda Makenzi”; so, 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 misspelling and from inverting name and surname are random. One might wonder whether a student introducing herself with a traditional or a baptismal name is random too. We are not aware of any study that shows that self-reporting of baptismal vs. traditional name is associated to some type of bias. This is especially unlikely in the case of children, who have developed a more limited identity awarness, and in the context of our study, which is characterised 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 analyses whether matched and unmatched students are statistically different across key observable personal and household characteristics like hours of study, wealth, mother’s education, source of light etc. The Table clearly shows that there is no significant difference between the two groups and that the matching does not depend on any specific interviewer. Table 4 reinforces this point by looking at balance at the class level. It shows that treatment intensity across classes is balanced over gender, teacher experience, wealth, and most grades.[10] There is imbalance in grades for science and social studies, but the content of the Kenyan Primary School syllabus generates little complementarity between these subjects and mathematics, which is where we find an effect.[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.

[10] A student in a class whose treatment intensity is 10% higher than another class tends to have 1.9 extra points in mathematics, but this difference is not statistically significant.

[11] 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.

3 Treatment compliance and attrition

We run 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 stayed in their household when they were sleeping; the remaining 16% said that the lamp was kept in school at night. The lamp was resistant and broke in only three cases; in all the other cases the lamp was reported to be in good condition or with only minor problems. Moreover, in all cases students declared that the solar charge was sufficient for either 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 the main activity they used the lamp for.[13]

Despite experimental compliance in terms of lamp appropriation and usage for studying, our experiment exhibits attrition coming from missing grades. Grades 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 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 exam in Term 1, 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 Term 2 and 3, but not all of it. We regress a dummy indicator for those repeating 7th grade on treatment and find an insignificant coefficient.[14] Moreover, in Table 5 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 the students

[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 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.

that had attrition and it turns that students with lower grades at baseline and female are more likely to drop out from our sample; whereas for other key characteristic we do not find a significant pattern consistent through different terms. Finally, notice that Table 1 shows that balance between treatment and control over balancing and additional baseline variables is 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 randomisation, the coefficients of the regressions have a causal interpretation. We show that standard specifications, which do not account for spillover effects, fail to find a significant effect of lamps. This should not be surprising because, given the randomisation of the solar lamp at the pupil level, the stable unit treatment value assumption is likely to fail.

We start our analysis by running an OLS estimation on a pooled cross-section that includes all grades of the end-of-term exams that followed our treatment. Our basic specification is the following:

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

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's age, mother's education, and number of siblings. We also include the balancing variables used in the re-randomisation 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 randomisation, 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 summarises the main findings of these specifications. Given the low number of clusters, we bootstrap the distribution of the test statistics 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 that then share their knowledge with control students; and iii) competition from control students who feel disadvantaged and increase their study effort. We

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

[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.

[17]We present the effect on the average grade across all subjects. We have also run these specifications on each subject and on each end-of-term test separately, but the results are still statistically insignificant.

have evidence of source (i) from students surveys and fieldwork experience. Moreover, we cannot rule out the presence of the other two 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 randomised 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 uniquely treated. This methodology involves a two-step randomisation process: treatment intensity is firstly randomised across clusters; then, individual treatment is randomised within clusters. As argued above, our first step is as good as random and the second step was randomised 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 ITT estimates rely on an out-of-sample prediction that hinges on the functional specification of the model. In Figure 5, we use a local polynomial smoother to analyse the relationship between grade first difference of control students and treatment intensity. We find a positive relation, 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

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

rest on the assumption that the linear specification extends also between 0% and 14% saturation. However, one might expect a concave function of treatment intensity on grades at very low saturations, as providing a few lamps to a class at zero saturation is likely to have a stronger effect than providing additional lamps to a class where there is already a moderate level of treatment saturation. This implies that our ITT estimates are more likely to provide a lower bound of the true effect due to Jensen's inequality.

Therefore, given the linear functional form we assume, our econometric specification is:

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

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

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:

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

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 π_j as the share of treated students in class j, $TUT = E(\Delta y_{ijt}|T_{ij} = 1, \pi_j = 0) - E(\Delta y_{ijt}|T_{ij} = 0, \pi_j = 0)$, where T_{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: $SC(\pi) = E(\Delta y_{ijt}|T_{ij} = 0, \pi_j = \pi) - E(\Delta y_{ijt}|T_{ij} = 0, \pi_j = 0)$. γ is the differential of the saturation slope for the treated and measures the effect of changing saturation on the treated compared to the control, so that $\mu + \gamma$ captures the spillover on treated, defined as $ST(\pi) = E(\Delta y_{ijt}|T_{ij} = 1, \pi_j = \pi) - E(\Delta y_{ijt}|T_{ij} = 1, \pi_j = 0)$. 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 such that $ITT(\pi) = E(\Delta y_{ijt}|T_{ij} = 1, \pi_j = \pi) - E(\Delta y_{ijt}|T_{ij} = 0, \pi_j = 0)$.

The results of this regression are presented in Table 7. 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 randomisation inference as in Rosenbaum (2002). The results are consistent across all procedures.

The results show a positive and significant intention-to-treat effect, such that treated students in a class with average saturation improve grades in mathematics by 0.88 standard deviation.[19] The ITT increases with the level of saturation and it ranges between 0.57 standard deviations at 16% of saturation and 1.1 standard deviations at 62% saturation. Moreover, we can see that there is a positive and significant spillover effect on the control group. The estimates of μ are positive, significant, and large in magnitude such that a 10% increase in saturation raises math grades of the control group by 0.22 standard deviations.

5.2 Geographical spillovers

Externalities may take place not only within the classroom, but also at home. 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, students could interact around the house during daylight or on their way to/from school. Therefore, Equation 5 needs to account for the fact that some of the spillovers that we attribute to class-level interaction, may actually be due to geographical proximity between treated and control students.

We the 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 use this information to construct a measure of the geographical treatment intensity. For each student, we compute the percentage of treated students within a

[19]The average class saturation in our sample is 43%. The intention-to-treat effect is given by the linear combination of $\beta + 0.43 \times (\mu + \gamma)$

radius of 500 meters, one kilometre, and 1.5 kilometres of their home. This is will include both students in the same and in a different class, 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 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 $k = 0.5, 1, 1.5\ km$.

Table 8 reports the results of Equation 6.[20] We can see that a positive and significant ITT is confirmed, such that a treated student in a class and geographic location with average saturation improves her 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 grades of control students by 0.047 standard deviations, but the result is not robust to randomisation inference. Similarly, geographical spillovers on treated students are also not robust to randomisation 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, rather than geographical spillovers, account for the bulk of the spillover effects.

6 Mechanism underlying ITT and spillover effects: suggestive evidence on study habits

The analysis of the survey on study habits and the distribution of student activities over the day provides some insight into the underlying mechanism

[20] We report the results for geographical treatment intensity within 1km. Measures based on a distance of 0.5 or 1.5 kms yield the same results. Details available upon request.

through which lamps can affect treated students and generate spillovers on controls. We find evidence consistent with the lamp 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.

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 which time of the day (before vs. 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 the dependent variable at baseline, 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 probability terms.[21] Given the lack of baseline values of 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 relation between study habits and lamp access.

In Table 9 we report the coefficients of the intention-to-treat and the spillover effects on control 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 for 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, the stronger effect occurs for *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

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

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 the pupils reported to do so at school. Moreover, Figures 6 and 7 show the ITT on treated students at the average saturation in our sample and the spillover effects on control for specific activities over different hours of the day. The Figures report the coefficients estimated using specification 5 and the bands of the standard errors.[22] 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 of work at home before sunset.

There are various reasons that can explain why the lamp allows 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;[23] second, the lamps allow students to walk home safely later in the day, when sunset and darkness are approaching; and, finally, for treated pupils, that 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 better grades we are unable to determine whether this is 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 spillovers. Sharing of knowledge due to interactions between students, as well as a competition effects, where control students increase study effort, are plausible possible candidates. Further investigation using network data on study partners could help to identify the different sources of spillovers.

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

[23] In some case teachers were keeping the lamps in the school and in other cases students occasionally brought the lamp with them at school.

An alternative explanation for the mechanism underlying the impact of the lamps is that it could be related to the income effect that the lamp generates. Das et al. (2013) show 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 households in our sample. Moreover, time use analysis on parents shows that the lamp allows 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 group. 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 find 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 to assess the effect of access to light on education. Through a randomised control 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 size of the technology shock that our experiment provides, all our estimates are likely to be a lower bound to the true effect of lighting, and energy access more in general, on education. Moreover, any experimental issues like lamp appropriation by teachers and lamp sharing with students in different classes, are likely to bias our estimates downwards. However, solar lamps should not be seen as a substitute for electrification, but as a short-term practical solution to limit the drawbacks to human capital accumulation coming from the lack of electricity.

We have also been able to disentangle within-class and geographical spillovers. Most of the spillovers arise from within-class interaction, while geographical proximity between treated and control students does not have a robust effect. The mechanisms through which spillovers arise seem to be related to increased co-studying at school, especially after sunset. Nevertheless, further research into this topic with larger samples and in different settings may help improve our understanding of the effects of light access on education and the mechanisms that can enhance or limit such effects.

References

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.

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*, (3), 286–327.

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

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

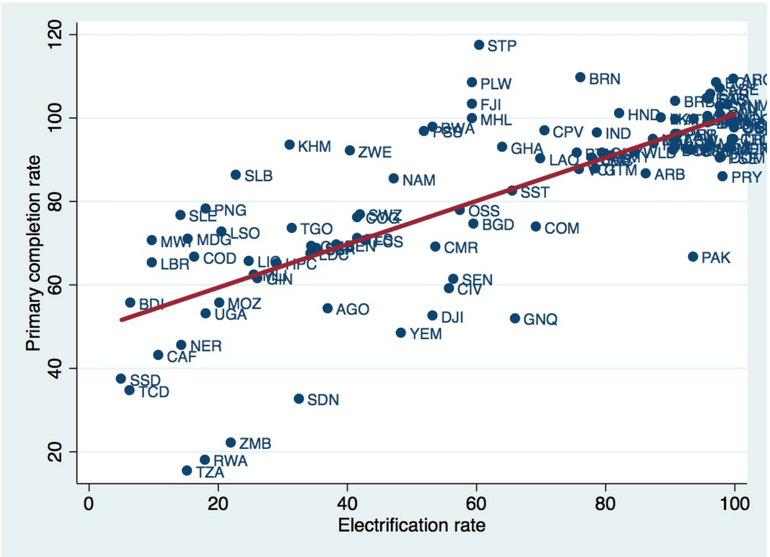

Figure 1: Electricity access and primary schooling (WDI data, electricity<100%)

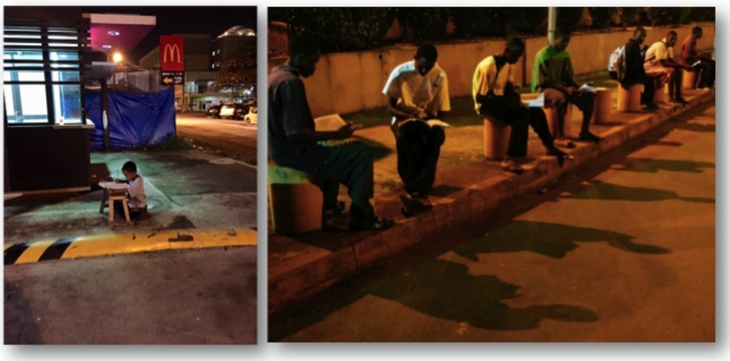

Figure 2: Students and lack of electrification (Associated Press and Facebook).

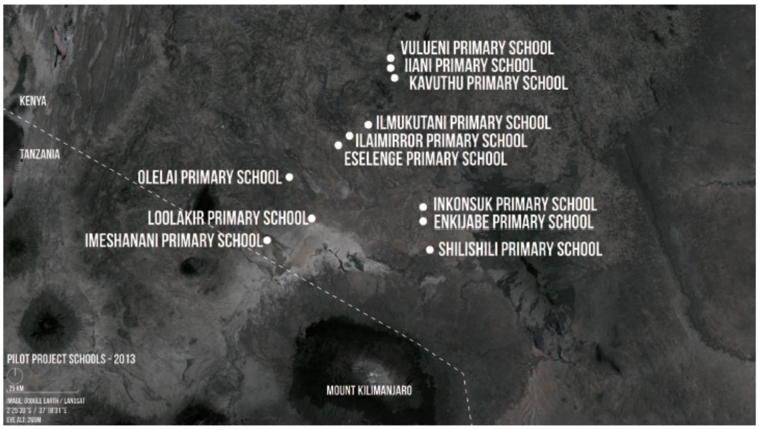

Figure 3: Area of intervention

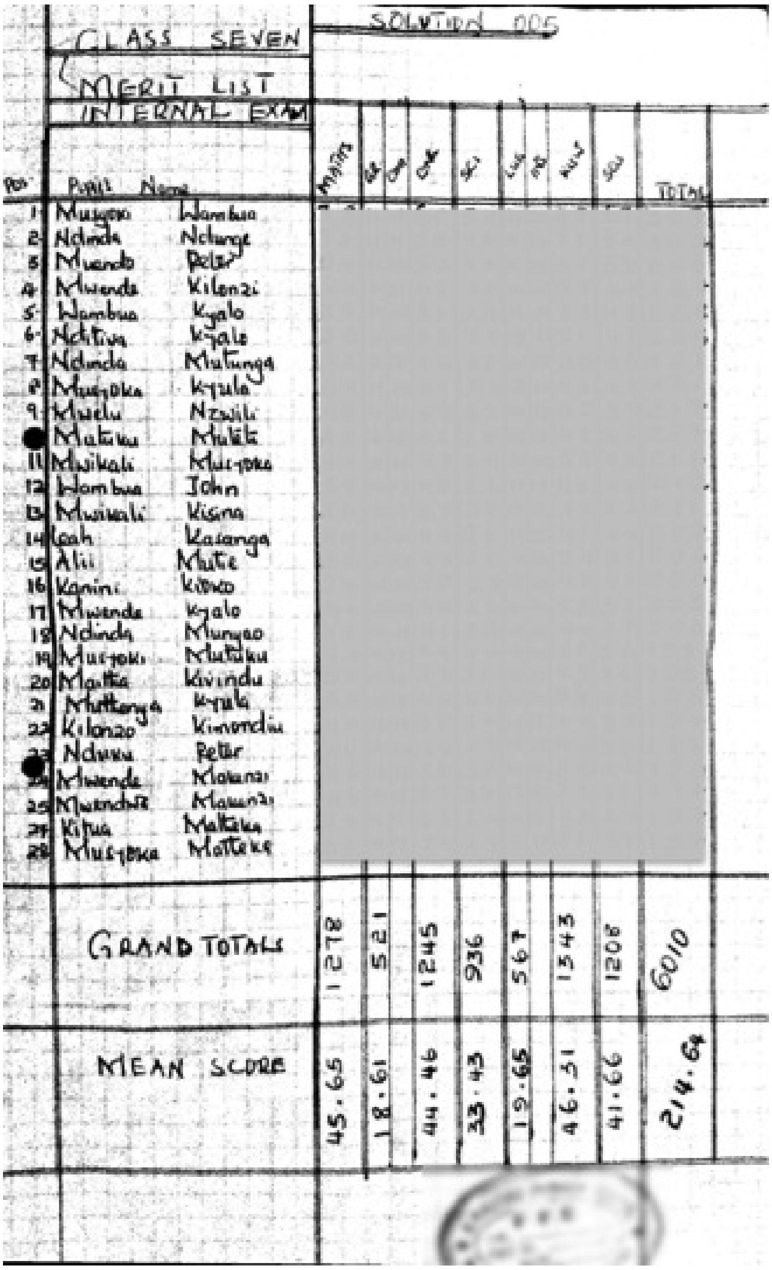

CLASS SEVEN SOLUTION 005
MERIT LIST
INTERNAL EXAM

POS	PUPILS NAME	
1	Musyoka	Wambua
2	Ndinda	Ndunge
3	Muendo	Peter
4	Mwende	Kilonzi
5	Wambua	Kyalo
6	Ndiitu	Kyalo
7	Ndinda	Mutunga
8	Musyoka	Kyule
9	Mwelu	Nzwili
[redacted]	Matiku	Mulele
11	Mwikali	Musyoka
12	Wambua	John
13	Mwikali	Kisina
14	Leah	Kasanga
15	Alii	Mutie
16	Kanini	Kioko
17	Mwende	Kyalo
18	Ndinda	Munyao
19	Musyoki	Mutuku
20	Maltha	Kivindu
21	Muttonga	Kyule
22	Kilonzo	Kimondiu
23	Nduku	Peter
[redacted]	Mwende	Kimanzi
25	Mwendwa	Kimanzi
27	Kitua	Mattaku
28	Musyoka	Mattaku

	MATHS								TOTAL
GRAND TOTALS	1278	521	1245	936	567	1343	1208		6010
MEAN SCORE	45.65	18.61	44.46	33.43	19.65	46.31	41.66		214.64

Figure 4: Example transcript

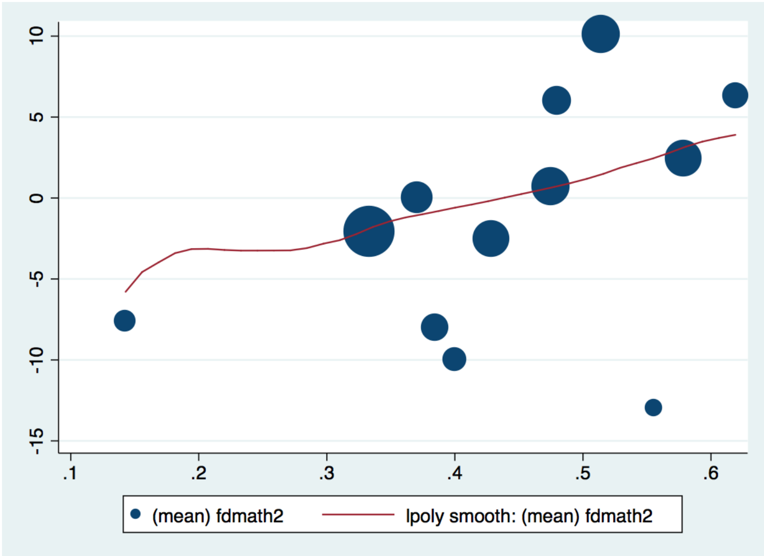

Figure 5: Local polynomial smoother of control groups' grades and class treatment intensity

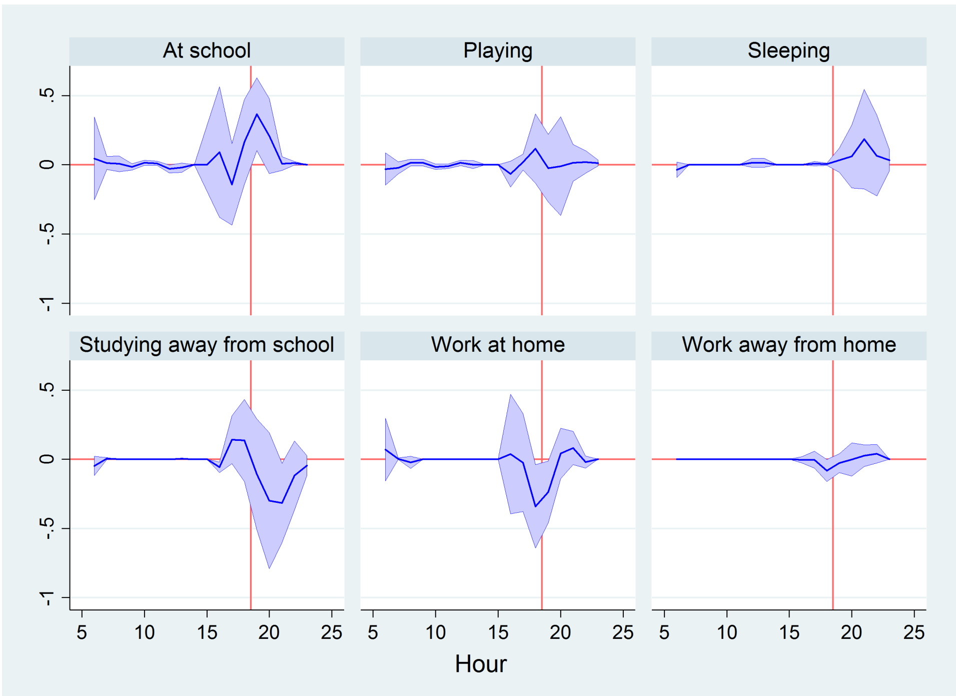

Figure 6: Intention-to-treat by time of day

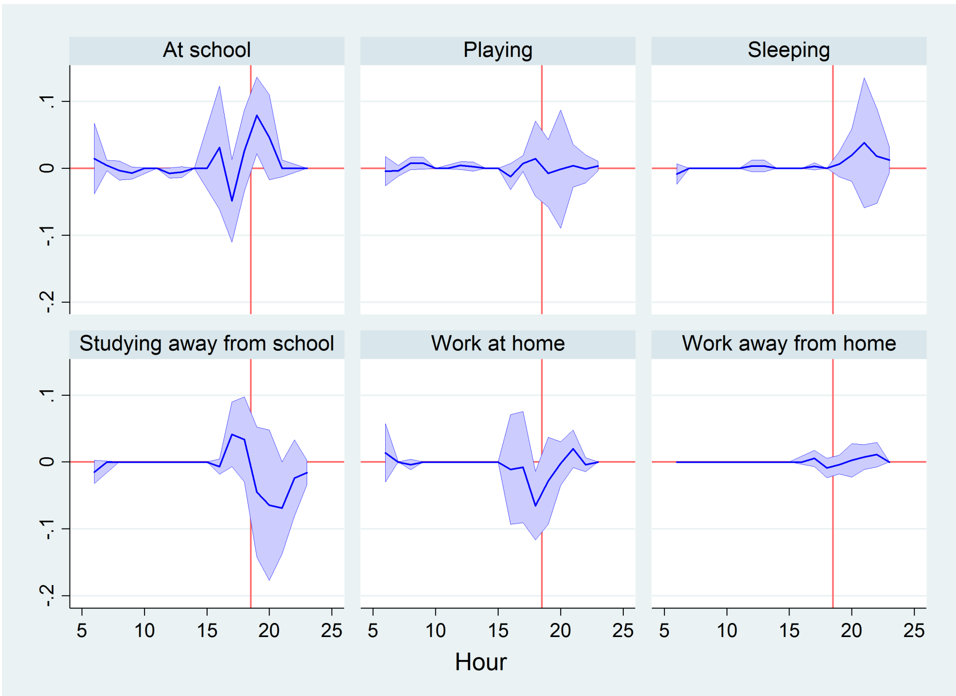

Figure 7: Saturation slope on control students by time of day

Tables

Table 1: Balance between treatment and control on variables at baseline

Explanatory variable: treatment	**Initial randomisation**		**End of Experiment**	
Balanced variables	Coefficient	p-value	Coefficient	p-value
Mathematics	2.26	0.20	3.1	0.19
English	-0.57	0.66	-091	0.60
Kiswahili	-0.26	0.82	-0.91	0.55
Science	0.56	0.75	-0.20	0.93
Social Studies	-1.44	0.29	-1.75	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

***. **, *, significant at the 1%level, 5% level, 10% level. P-values are insignificant also under multiple hypothesis testing as in List et al. (2016).

Table 2: Treatment intensity variation

	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

Explanatory variable: unmatched student	Coefficient	p-value
Hours of study	-0.35	0.48
Missed school days	-0.16	0.55
N. of people in the 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

Explanatory variable: treatment intensity	Coefficient	p-value
Mathematics	19.01	0.41
English	13.04	0.64
Swahili	1.67	0.91
Science	55.61***	0.00
Social studies	-27.7*	0.07
Gender	0.06	0.88
Wealth index	-0.11	0.87
Teacher experience	-5.86	0.38
Teacher education	-0.17	0.81
Hours of study	-1.06	0.71
N. of people in the household	-1.86	0.67
Source of study light: kerosene	-0.12	0.79
Mother education	0.66	0.52

***significant at the 1%level; * significant at the 10% level. P-values are insignificant also under multiple hypothesis testing as in List et al. (2016).

Table 5: Attrition

Y: Missing the exam	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)
N. of children	0.004	0.005	-0.02
in the household	(0.66)	(0.64)	(0.12)
Mother's education	0.14**	0.02	-0.04
	(0.04)	(0.73)	(0.38)
Father's 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	(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}	
Mother's 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 standardised grade in mathematics. All specifications account for class fixed effects and balancing variables.

Table 7: RSD estimates - 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 randomisation inference in braces {}. The dependent variable is the standardised grade in mathematics.

Table 8: Spillover effect with geographical estimates - 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 standardised grade in mathematics.

Table 9: Impact on study habits

	ITT at average saturation (43%)	Saturation slope 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]

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