

Disrupting Education? Experimental Evidence on Technology-Aided Instruction in India*

Karthik Muralidharan[†]
UC San Diego

Abhijeet Singh[‡]
University College London

Alejandro J. Ganimian[§]
J-PAL

July 20, 2017

Abstract

We present experimental evidence on the impact of a personalized technology-aided after-school instruction program on learning outcomes. Our setting is middle-school grades in urban India, where a lottery provided winning students with a voucher to cover program costs. We find that lottery winners scored 0.36σ higher in math and 0.22σ higher in Hindi relative to lottery losers after just 4.5-months of access to the program. IV estimates suggest that attending the program for 90 days would increase math and Hindi test scores by 0.59σ and 0.36σ respectively. We find similar absolute test score gains for all students, but the relative gain was much greater for academically-weaker students because their rate of learning in the control group was close to zero. We show that the program was able to effectively cater to the very wide variation in student learning levels within a single grade by precisely targeting instruction to the level of student preparation. The program was cost effective, both in terms of productivity per dollar and unit of time. Our results suggest that well-designed technology-aided instruction programs can sharply improve productivity in delivering education.

JEL codes: C93, I21, J24, O15

Keywords: computer-aided learning, productivity in education, personalized learning, teaching at the right level, post-primary education, middle school, secondary school

*We thank Abhijit Banerjee, Peter Bergman, Prashant Bharadwaj, Gordon Dahl, Roger Gordon, Heather Hill, Priya Mukherjee, Chris Walters and several seminar participants for comments. We thank the staff at Educational Initiatives (EI)—especially, Pranav Kothari, Smita Bardhan, Anurima Chatterjee, and Prasad Sreeprakash—for their support of the evaluation. We also thank Maya Escueta, Smit Gade, Riddhima Mishra, and Rama Murthy Sripada for excellent research assistance and field support. Finally, we thank J-PAL’s Post-Primary Education initiative for funding this study. The operation of Mindspark centers by EI was funded by the Central Square Foundation, Tech Mahindra Foundation and Porticus. All views expressed are those of the authors and not of any of the institutions they are affiliated with.

[†]Dept. of Economics, University of California San Diego (UCSD); NBER; J-PAL. E-mail: kamurali@ucsd.edu.

[‡]Dept. of Economics, University College London (UCL). E-mail: abhijeet.singh@ucl.ac.uk.

[§]Abdul Latif Jameel Poverty Action Lab (J-PAL); E-mail: aganimian@povertyactionlab.org.

1 Introduction

Developing countries have made impressive progress in improving school enrollment and completion in the last two decades. Yet, their productivity in converting education investments of time and money into human capital remains very low. For instance, in India, over 50% of students in Grade 5 cannot read at the second grade level, despite primary school enrollment rates over 95% (Pratham, 2017). Further, there have been very limited improvements in learning outcomes in the past decade despite substantial increases in education spending in this period (Muralidharan, 2013). More generally, even in developed countries, productivity growth in the education sector lags the rest of the economy, perhaps because the ‘technology’ of schooling (classroom-based instruction) has changed very little over time compared to rapid technological progress in other fields (Bosworth, 2005; Pritchett, 2013).

Thus, it is not surprising that increasing the use of technology in instruction is seen as a leading candidate for ‘disrupting’ the status quo and improving productivity in education (Negroponte et al., 2006; Khan, 2012; Mead, 2016).¹ Yet, the evidence to date appears rather mixed: A recent review of evidence from high-quality studies on the impact of using technology in education globally reports “mixed evidence with a pattern of null results” (Bulman and Fairlie, 2016). Overall, the evidence thus far suggests that realizing the potential of technology-aided instruction to improve education will require paying careful attention to the *details* of the specific intervention, and the extent to which it alleviates binding constraints to learning.

In this paper, we present experimental evidence on the impact of a technology-led instructional program (called Mindspark) that aimed to leverage technology to improve education by paying sustained attention to such design details. Developed by a leading Indian education firm, the Mindspark software reflects over 10 years of product development; it has been used by over 400,000 students, has a database of over 45,000 test questions, and administers over a million questions across its users every day. A key feature of the software is its ability to use these data to finely benchmark the learning level of every student and dynamically customize the material being delivered to match the level and rate of progress made by each individual student. A second noteworthy feature is its ability to analyze these data to identify patterns of student errors, and precisely target content to alleviate conceptual ‘bottlenecks’ that may be difficult for teachers to diagnose or address at the individual student level in a classroom setting. Mindspark can be delivered in a variety of settings (in schools, in after-school centers, or through self-guided study); it is platform-agnostic (can be deployed through computers, tablets, or smartphones); and it can be used both online and offline.

¹A non-exhaustive list of posited channels of impact include using technology to consistently deliver high-quality content that may circumvent limitations in teachers’ own knowledge; delivering engaging (often game-based) interactive content that may improve student attention; delivering individually customized content for students; reducing the lag between students attempting a problem and receiving feedback; and, analyzing patterns of student errors to precisely target content to clarify specific areas of misunderstanding.

We evaluate the after-school Mindspark centers in this paper. The centers scheduled six days of instruction per week, for 90 minutes per day; each session was divided into 45 minutes of individual self-driven learning on the Mindspark software and 45 minutes of instructional support from a teaching assistant in groups of 12-15 students.² The centers aimed to serve students from low-income neighborhoods in Delhi, and charged a modest fee.³ Our evaluation was carried out in a sample of 619 students recruited for the study from public middle schools in Delhi. Around half of these students were randomly selected to receive a voucher offering free attendance at the centers. We measure program impacts using independently-conducted paper-and-pencil tests of student learning in math and Hindi (language) before and after the 4.5-month long intervention. These tests were linked using item response theory (IRT) to be comparable on a common scale across both rounds of testing and across different grades.

We use the detailed data from the computer-based assessments to present two key facts about the context. First, we show that average student achievement in our sample (measured at baseline) is several grade-levels behind grade-appropriate standards and that this gap grows by grade. The average grade 6 student is around 2.5 grade levels below grade 6 standards in Math; by grade 9, this deficit increases to 4.5 grade levels. Second, we show that there is considerable heterogeneity in within-grade student learning levels. Students enrolled in the same grade typically *span five to six grade levels* in their preparation, with the vast majority of them being below grade-level standards. Thus, our setting is likely to feature considerable mismatch between the level of student preparation and the default pattern of classroom instruction using textbooks based on grade-appropriate curricula.

We report three main sets of results based on the experiment. First, we find that students winning a program voucher scored 0.36σ higher in math and 0.22σ higher in Hindi relative to students who applied for but did not win the lottery. Relative to the control group, lottery winners experienced twice the test score value-added in math and 2.5 times that in Hindi during the study period of 4.5 months. These are intent-to-treat (ITT) estimates reflecting an average attendance rate of 58%. Using the lottery as an instrumental variable for attendance (and additional assumptions discussed in Section 4.5), we estimate that attending the Mindspark centers for 90 days (which corresponds to 80% attendance for half a school year), would raise math and Hindi test scores by 0.59σ and 0.36σ respectively.

²The teaching assistant focused on helping students with completing homework and with exam preparation, while the instruction was mostly provided by the Mindspark software (see sections 2.1.1 and 5.1 for details).

³The online and school-based models require fees that are not affordable for low-income families. The Mindspark centers were set up with philanthropic funding to make the platform more widely accessible, and were located in low-income neighborhoods. However, the funders preferred that a (subsidized) fee be charged, reflecting a widely-held view among donors that cost-sharing is necessary to avoid wasting subsidies on those who will not value or use the product (Cohen and Dupas, 2010). The subsidized fee of Rs. 200 per month (USD 3 per month) was benchmarked to that charged by providers of private tutoring in the vicinity.

Second, the ITT effects do not vary by students' baseline test scores, gender, or household socioeconomic status. Thus, consistent with the promise of customized technology-led instruction, the intervention was equally effective at improving test scores for *all* students. Further, consistent with the hypothesis of mismatch between classroom instruction and student preparation (Banerjee and Duflo, 2012), we find that the absolute test-score value-added in the control group is close to zero for students in the bottom-third of the within-grade baseline test-score distribution, and we cannot reject that these students made no academic progress through the school year. Thus, while the absolute ITT effect was similar for all students, the *relative* effect was much greater for academically-weaker students.

Third, we examine ITT effects at different levels of test-question difficulty. Since student learning levels were far below grade level in math, the Mindspark system (which customized content to each student's learning level) mainly provided students with content at below grade-level difficulty. In Hindi, where student learning gaps relative to curricular standards were smaller, students were provided with content both at and below grade-level difficulty. The test-score results reflect this pattern of instruction: In math, the test-score gains are only seen in questions of below grade-level difficulty; whereas, in Hindi test-score gains are found in questions both at and below grade-level. Consistent with this, we find significant improvements in test scores on (grade-level) school exams in Hindi but not in Math.⁴

The test score value-added in the treatment group was over 100% greater than that in the control group, and was achieved at a lower cost per student than in the public schooling system. Thus, the program was cost effective even at the very small scale evaluated in this study, and is likely to be highly cost effective at a larger scale (since marginal costs are much lower than the average cost in our study). Further, given large learning deficits in developing countries and finite years of schooling, it is also worth considering productivity per unit of time. For instance, Muralidharan (2012) finds that providing individual-level performance bonuses to teachers in India led to test score gains of 0.54σ and 0.35σ in math and language after five years of program exposure. This is one of the largest effect sizes seen to date in an experimental study on education in developing countries. Yet, we estimate that regularly attending Mindspark centres could yield similar gains in one tenth the time (half a year).

The effects presented above represent a combination of the Mindspark computer-aided learning (CAL) program, group-based instruction, and extra instructional time (since we study an after-school program), and our study design does not allow us to experimentally distinguish between these channels of impact. However, a contemporaneous experimental study on the impact of an after-school group tutoring program that was also in Delhi, also targeted

⁴These results also highlight the importance of ensuring that tests used for education research are informative over a wide range of student achievement (especially in developing country settings with wide variation in within-grade student learning). Using only grade-appropriate tests (or school tests) would have led to incorrect inference regarding program impact (see discussion in Section 4.4).

middle-school students, and featured an even longer duration of after school instruction found *no impact* on test scores (Berry and Mukherjee, 2016). These results suggest that extra after-school instructional time or group-based tutoring on their own may have had limited impact on student learning without the CAL program. Thus, while our experimental estimates reflect the composite impact of a ‘blended learning’ program, they are most likely attributable to the CAL component and not the group instruction (see discussion in section 5.1).

Our results contribute to the growing literature on the impact of technology in education, where our findings are closest to those of Banerjee et al. (2007), who find that primary-school students in urban India scored 0.47σ higher in math after two years of a math CAL program. We make three main advances over the existing CAL literature. First, we use the rich CAL data to get into the ‘black box’ of the classroom and document key facts on constraints to effective teaching and learning. Second, we use this data to show how the CAL program is able to alleviate these constraints by personalizing instruction to the learning level and rate of progress of each student. Third, we focus on post-primary grades where evidence on effective ways of improving learning outcomes is both scarce and increasing in importance (Banerjee et al., 2013). Improving learning outcomes in post-primary grades is more challenging than in primary school because constraints of low teacher and parent human capital become more binding at higher grade levels. Teachers also face increasing heterogeneity in student preparation, partly due to “automatic grade promotion” policies in many countries. This is a setting where the potential for computer-aided instruction to improve teaching and learning is high, but there is limited evidence of notable successes to date (Banerjee et al., 2013).

We also document three key facts on education in developing countries, which have been conjectured to be true, but have not been shown *directly* so far (to the best of our knowledge). Specifically, we use the granular assessment data from the CAL system mapped into grade-level standards (without test ceiling or floor effects), and vertically-linked IRT-scaled test scores, to show that: (a) there are large gaps between student preparation and grade-level standards that grow by grade, (b) there is remarkably-large heterogeneity in student preparation in the same grade, and (c) students with low initial learning levels *make no progress* in learning under the status quo despite being enrolled in school. While our sample is not representative of developing countries, all three facts are consistent with existing experimental results from several contexts across South Asia and Africa (such as Banerjee et al. (2007); Glewwe et al. (2009); Duflo et al. (2011)).⁵ Our analysis demonstrates the value of CAL-data for descriptive education research and provides a template for future work to similarly document facts on learning levels, heterogeneity, and trajectories in other settings.

⁵Indeed, it is the evidence from these experimental studies that has contributed to several researchers of education in developing countries conjecturing the facts that we present here (Banerjee and Duflo, 2012; Pritchett and Beatty, 2015; Glewwe and Muralidharan, 2016). However, there has been no direct demonstration of these facts to date.

To help place our results in the context of the existing evidence, we conducted an extensive review of existing studies with attention to the *details* of CAL interventions (see Appendix B). Our review suggests that some clear patterns are starting to emerge. Hardware-focused interventions that provide computers at home or at school seem to have no positive impact on learning outcomes.⁶ Pedagogy-focused CAL programs that allow students to review grade-appropriate content at their own pace do better, but the gains are modest and range from 0.1σ to 0.2σ .⁷ Finally, the interventions that deliver the largest gains (like the one we study and the one studied in Banerjee et al. (2007)) appear to be those that use technology to also personalize instruction. Thus, our results suggest that personalization may be an important ingredient for achieving the full potential of technology-aided instruction.

Our study also contributes evidence on policy options to address the challenge of large variation in student preparation, which can make the effective delivery of *any* curriculum challenging even for highly motivated and trained teachers. The most promising approaches to date have involved either partial or complete tracking of classrooms to facilitate teaching closer to the learning level of students (Banerjee et al., 2007; Duflo et al., 2011; Banerjee et al., 2016). However, tracking is controversial, and may also not be feasible in many practical settings. Our results suggest that well-designed CAL programs may be able to deliver the pedagogical advantages of tracking while mitigating several of its challenges (see discussion in section 5.3).

More broadly, our evidence on the ability of technology-aided instruction to help circumvent constraints to human capital accumulation in developing countries, speaks to the potential for new technologies to enable low-income countries to leapfrog constraints to development. Examples from other sectors include the use of mobile telephones to circumvent the lack of formal banking systems (Jack and Suri, 2014), the use of electronic voting machines for better enfranchisement of illiterate citizens (Fujiwara, 2015) and the use of biometric authentication to circumvent literacy constraints to financial inclusion (Muralidharan et al., 2016).

The rest of this paper is organized as follows. Section 2 describes the intervention, and experimental design. Section 3 describes our data. Section 4 presents our main results. Section 5 discusses mechanisms, costs, and policy implications. Section 6 concludes.

⁶See, for example, Angrist and Lavy (2002); Barrera-Osorio and Linden (2009); Malamud and Pop-Eleches (2011); Cristia et al. (2012); Beuermann et al. (2015). These disappointing results are likely explained by the fact that hardware-focused interventions have done little to change instruction, and at times have crowded out student time for independent study.

⁷See, for example, Carrillo et al. (2010); Lai et al. (2015a, 2013, 2012); Linden (2008); Mo et al. (2014b); Barrow et al. (2009); Rouse and Krueger (2004). Anecdotal evidence suggests that pedagogy-focused CAL interventions have typically focused on grade-appropriate content in response to schools' and teachers' preference for CAL software to map into the topics being covered in class and reinforce them.

2 Intervention and Study Design

2.1 The Mindspark CAL software

The Mindspark CAL software, developed by Educational Initiatives (EI), a leading Indian education firm, is the central component of the program we study. The software is interactive and includes continuous student assessment alongside instructional games, videos, and activities from which students learn through explanations and feedback. It reflects over a decade of iterative product development and aims to leverage several posited channels by which education technology may improve pedagogy. We highlight some of the key design features of the software here, and provide a more detailed description with examples for each of the points below in Appendix C.

First, it is based on an extensive corpus of *high-quality instructional materials*, featuring an item bank of over 45,000 test questions, iterated over several years of design and field testing. The design of the content tries to reflect current research in effective pedagogy that is relevant to low-income settings, such as the use of same-language subtitling for teaching literacy (Kothari et al., 2002). Further, the software allows this material to be *delivered with uniform consistency* to individual students, thereby circumventing both limitations in teacher knowledge as well as heterogeneity in knowledge and teaching ability across teachers.

Second, the content is *adaptive*, with activities presented to each student being based on that student’s performance. This adaptation is dynamic, occurring both at the beginning based on a diagnostic assessment, and then with every subsequent activity completed. Thus, while the Mindspark content database is mapped into the grade-level curricular standards of the education system, an essential feature of the software is that the content presented to students is not linked to the curriculum or textbook of the grade that the student is enrolled in. In other words, it enables dynamic “Teaching at the right level” for each individual student and can cater effectively to very wide heterogeneity in student learning levels that may be difficult for even highly trained and motivated teachers to achieve in a classroom setting.

Third, even students at approximately similar levels of understanding of a topic, may have very different specific areas of conceptual misunderstanding. Thus, the pedagogical approach needed to alleviate a student-specific conceptual ‘bottleneck’ may be different across students. Mindspark aims to address this issue by using its large database of millions of student-question level observations to identify patterns of student errors and to classify the type of error and target *differentiated remedial instruction* accordingly (see Appendix C.4.2 for examples). This attention to understanding patterns in student errors builds on an extensive literature in education that emphasizes the diagnostic value of error analysis in revealing the heterogeneous needs of individual students (see Radatz 1979 for a discussion). However, while the value of

error analysis is well-known to education specialists, implementing it in practice in classroom settings is non-trivial and the use of technology sharply reduces the cost of doing so.⁸

Finally, the interactive user interface, combined with the individualization of material for each student, facilitates children’s *continuous engagement* with the material. The software makes limited use of instructional videos (where student attention may waver), choosing instead to instruct with steps that require students to constantly interact with the system. This approach aims to boost student attention and engagement, to provide feedback at the level of each intermediate step in solving a problem, and to shorten the feedback loop between students attempting a problem and learning about their errors and how to correct them.

As the discussion above makes clear, Mindspark aims to use technology to *simultaneously alleviate multiple constraints* to effective teaching and learning in a scalable way. In future work, we hope to run micro-experiments on the Mindspark platform to try to isolate the impact of specific components of the software on learning outcomes (such as personalization, differentiated feedback, or the impact of specific pedagogical strategies). However, from an economists’ perspective, we are more interested in studying the extent to which technology-aided instruction can *improve productivity* in delivering education. Thus, our focus in this paper is on studying the “full potential” impact of technology-aided instruction on education outcomes (which includes all the channels above), and we defer an analysis of the relative importance of specific components of Mindspark to future work.

2.1.1 The Mindspark centers intervention

The Mindspark CAL software has been deployed in various settings: private and government schools, after-school instructional centers and individual subscription-based use at home. Here, we evaluate the supplementary instruction model, delivered in stand-alone Mindspark centers that target students from low-income households. Students sign up for the program by selecting a 90-minute batch, outside of school hours, which they are scheduled to attend six days per week. The centers charged a (subsidized) fee of INR 200 (USD 3) per month.⁹

Scheduled daily instruction in Mindspark centers was divided into 45 minutes of computer-based instruction and 45 minutes of supervised instructor-led group-based study. In the time allotted to the computer-based instruction, each student was assigned to a Mindspark-equipped computer with headphones that provided him/her with activities on

⁸The emphasis on error analysis reflects EI’s long experience in conducting similar analyses and providing diagnostic feedback to teachers based on paper-and-pen tests (Muralidharan and Sundararaman, 2010). Thus, the Mindspark development process reflects the aim of EI to use technology to improve productivity in implementing ideas that are believed by education specialists to improve the effectiveness of pedagogy.

⁹The intensity of the program, as well as the fee charged, was designed to be comparable to after-school private tutoring, typically conducted in groups of students, which is common in India. According to the 2012 India Human Development Survey, 43% of 11-17 year olds attended paid private tutoring outside of school.

math, Hindi and English. Two days of the week were designated for math, two days for Hindi, one day for English, and students could choose the subject on one day each week.

The group-based instruction component included all students in a given batch (typically around 15 students) and was supervised by a single instructor. Instructors were locally hired and were responsible for monitoring students when they are working on the CAL software, providing the group-based instruction, facilitating the daily operation of the centers, and encouraging attendance and retention of enrolled students.¹⁰ Instruction in the group-based component consisted of supervised homework support and review of core concepts of broad relevance for all children without individual customization.

Thus, the intervention provided a ‘blended learning’ experience that included personalized one-on-one computer-aided instruction along with additional group academic support by an instructor. As a result, all our estimates of program impact and cost effectiveness are based on this composite program. Further, to the extent that the presence of an adult may be essential to ensure student adherence to the technology (both attendance and time on task), it may not be very meaningful to try to isolate the impact of the technology alone. In section 5.1, we discuss results from a parallel experimental evaluation in the same context showing *no impact* on student learning from an after-school group tutoring program (with no technology). Hence, one way to interpret our results is as an estimate of the extent to which using technology *increased the productivity of an instructor*, as opposed to technology by itself.

2.2 Sample

The intervention was administered in three Mindspark centers in Delhi focused on serving low-income neighbourhoods. The sample for the study was recruited in September 2015 from five public middle schools close to the centers. All five schools had grades 6-8, three of these schools had grade 9, and only two had grades 4-5. Three were all-girls schools and the other two were all-boys schools. Therefore, our study sample has a larger share of girls in grades 6-8. In each school, staff from EI and from J-PAL South Asia visited classrooms from grades 4-9 to introduce students to the Mindspark centers and to invite them and their parents to a demonstration at the nearby center (information flyers were provided to share with parents).

At the demonstration sessions, students and their parents were introduced to the program and study by EI staff. Parents were told that, if their child wanted to participate in the study, he/she would need to complete a baseline assessment and that about half of the students would be chosen by lottery to receive a voucher which would waive the usual tuition fees of

¹⁰These instructors were recruited based on two main criteria: (a) their potential to interact with children; and (b) their performance on a very basic test of math and language. However, they were not required to have completed a minimum level of education at the secondary or college level, or have any teacher training credentials. They received initial training, regular refresher courses, and had access to a library of guiding documents and videos. They were paid much lower salaries than civil-service public-school teachers.

INR 200 per month until February 2016 (i.e. for nearly half of the school year). Students who were not chosen by lottery were told that they would be provided free access to the centers after February 2016, if they participated in an endline assessment in February 2016. However, lottery losers were not allowed to access the program during the study period. These two design features helped to reduce attrition, and increase statistical power respectively.

Our study sample comprises the 619 students who completed the baseline tests and surveys. About 97.5% of these students were enrolled in grades 6-9.¹¹ To assess the representativeness of our self-selected study sample (and implications for the external validity of our results), we compare administrative data on school final-exam scores in the preceding school year (2014-15) across study participants and the full population of students in the same schools. Study participants have modestly higher pre-program test scores (of around 0.15σ) than non-participants (Table A.1). However, there is near-complete common support in the pre-program test-score distribution of participants and non-participants (Figure A.1), suggesting that our results are likely to extend to other students in this setting (especially since we find no heterogeneity in impact by baseline test scores; see Section 4.3).

2.3 Randomization and Compliance

The 619 participants were individually randomized into treatment and control groups with 305 students in the control and 314 in the treatment group. Randomization was stratified by center-batch preferences.¹² The treatment and control groups did not differ significantly on any observable dimension at baseline (Table 1, Panel A). Of the 314 students offered a voucher for the program, the mean attendance rate was 58% (around 50 days out of a maximum possible of 86 days). The full distribution of attendance among lottery-winners is presented in Figure A.2, and we present both ITT estimates of winning the lottery and IV estimates of the dose-response relationship as a function of days of attendance in Section 4.

Of the 619 students who participated in the baseline test, 533 (86%) also attended the endline test. The follow-up rate was 84% in the treatment group and 88% in the control group. This difference is significant at the 10% level and so we will present inverse probability weighted estimates of treatment effects as well as Lee (2009) bounds of the treatment effect (section 4.6.1). We also find no significant difference between treatment and control groups in mean student characteristics (age, gender, SES, or baseline test scores) of those who attend both the baseline and endline test, and comprise our main study sample (Table 1, Panel B).

¹¹589 students were enrolled in grades 6-9, 15 were enrolled in grades 4-5 and, for 15 students, the enrolled grade was not reported. Our focus on Grades 6-9 reflects our funding from the JPAL Post Primary Education Initiative, which prioritized studying interventions to improve post-primary education (after fifth grade).

¹²Students were asked to provide their preferred slots for attending Mindspark centers given school timings and other commitments. Since demand for some slots was expectedly higher than others, we generated the highest feasible slot for each student with an aim to ensure that as many students were allocated to their first or second preference slots as possible. Randomization was then carried out within center-by-batch strata.

3 Data

3.1 Student achievement

The primary outcome of interest for this study is student test scores. Test scores were measured using paper-and-pen tests in math and Hindi prior to the randomization (September 2015, baseline) and near the end of the school year (February 2016, endline).¹³ Tests were administered centrally in Mindspark centers at a common time for treatment and control students with monitoring by J-PAL staff to ensure the integrity of the assessments.

The tests were designed independently by the research team and intended to capture a wide range of student achievement. Test items ranged in difficulty from “very easy” questions designed to capture primary school level competencies much below grade-level to “grade-appropriate” competencies found in international assessments. Test scores were generated using Item Response Theory models to place all students on a common scale across the different grades and across baseline and endline assessments. The common scale over time allows us to characterize the *absolute* test score gains made by the control group between the two rounds of testing. The assessments performed well in capturing a wide range of achievement with very few students subject to ceiling or floor effects. Details of the test design, scoring, and psychometric properties of individual test questions are provided in Appendix D.

3.2 Mindspark CAL system data

The Mindspark CAL system logs all interactions that each student has with the software platform. This includes attendance, content presented, answers to each question presented, and the estimated grade level of student achievement at each point in time. These data are available (only) for the treatment group. We use these data in three ways: to describe the mean and distribution of learning gaps relative to curricular standards in each grade at baseline; to demonstrate the personalization of instruction by Mindspark; and to characterize the evolution of student achievement in the treatment group over the period of the treatment.

3.3 School records

At the school level, we collected administrative records on test scores on school exams of all study students and their peers in the same schools and classrooms. This was collected for both the 2014-15 school year (to compare the self-selected study sample with the full population of students in the same schools) and the 2015-16 school year (to evaluate whether the treatment affected test scores on school exams).

¹³It was important to test students in a pen-and-paper format, rather than computerized testing, to avoid conflating true test score gains with greater familiarization with computer technology in the treatment group.

3.4 Student data

At the time of the baseline assessment, students answered a self-administered written student survey which collected basic information about their socio-economic status, and household characteristics. A shorter survey of time-varying characteristics was administered at endline. We also conducted a brief telephone survey of parents in July 2016 to collect data on use of private tutoring, and their opinion of the Mindspark program.

4 Results

4.1 Learning levels and variation under the status-quo

Data from the Mindspark CAL system provides an assessment of the actual grade-level of each student’s learning level regardless of grade enrolled in. We use these data to characterize learning levels, gaps, and heterogeneity among the students in our sample. The main results are presented in Figure 1, which shows the full joint distribution of the grades students were enrolled in and their assessed learning level at the start of treatment.¹⁴

We highlight three main patterns in Figure 1. First, most children are already much below grade level competence at the beginning of post-primary education. In grade 6, the average student is about 2.5 grades behind in math and about half a grade behind in Hindi.¹⁵ Second, although average student achievement is higher in later grades, indicating some learning over time, the slope of achievement gains (measured by the line of best fit) is considerably flatter than the line of equality between curricular standards and actual achievement levels. This suggests that average student academic achievement is progressing at a lower rate than envisaged by the curriculum — by grade 9, students are (on average) nearly 4.5 grades behind in math and 2.5 grades behind in Hindi. Third, the figure presents a stark illustration of the very wide dispersion in achievement among students enrolled in the *same* grade: students in our sample span 5-6 grade levels in each grade.

While it is possible in theory to plot the equivalent of Figure 1 in any setting, in practice we are not aware of other studies that have done this to date. The most likely reason is that paper-and-pencil tests conducted over a fixed test duration are subject to ceiling and floor effects. In contrast, the dynamically adaptive CAL system can rapidly adjust the difficulty of questions asked till students get some correct and others wrong to enable a precise calibration of the learning level of *each* student. Characterizing and understanding variation

¹⁴Note that these data are only available for students in the treatment group. However, Figure 1 uses data from the *initial* diagnostic test, and does not reflect any instruction provided by Mindspark.

¹⁵While most patterns across grades are similar in the two subjects, the computer system’s assessment on grade-level competence of children may be more reliable for math than for language (where competencies are less well-delineated across grades). Baseline test scores on our independent tests in both subjects are consistently higher for students assessed by the CAL program as being at a higher grade level of achievement, which helps to validate the grade-level bench-marking by the CAL program (See Figure A.4). Further details of the diagnostic test and bench-marking by the software are presented in Appendix C.

of the sort seen in Figure 1 is essential for designing effective instructional strategies, and our analysis highlights the value of CAL data for descriptive education research over and above the (potential) impact of CAL on learning, which we turn to next.

4.2 Program Effects (Intent-to-treat estimates)

The main treatment effects can be seen visually in Figure 2, which presents mean test scores in the baseline and endline assessments in math and Hindi for lottery-winners and losers. While test scores improve over time for both groups, endline test scores are significantly and substantially higher for the treatment group in both subjects.

We estimate intent-to-treat (ITT) effects of winning the lottery (β) using:

$$Y_{iks2} = \alpha_s + \gamma_s \cdot Y_{iks1} + \beta_s \cdot Treatment_i + \phi_k + \epsilon_{iks2} \quad (1)$$

where Y_{ikst} is student i 's test score, in randomization stratum k , in subject s at period t (normalized to $\mu=0, \sigma=1$ on the baseline test); $Treatment$ is an indicator variable for being a lottery-winner; ϕ is a vector of stratum fixed effects; and ϵ_{iks2} is the error term.¹⁶

We find that students who won the lottery to attend Mindspark centers scored 0.36σ higher in math and 0.22σ higher in Hindi compared to lottery losers after just 4.5 months (Table 2: Cols. 1-2). In Cols. 3 and 4, we omit strata fixed effects from the regression, noting that the constant term α in this case provides an estimate of the absolute value-added (VA) in the control group over the course of the treatment.¹⁷ Expressing the VA in the treatment group ($\alpha+\beta$) as a multiple of the control group VA (α), our results indicate that lottery-winners made twice the progress in math, and 2.5 times the progress in Hindi, compared to lottery-losers. These are ITT results based on an average attendance of about 58% among lottery-winners. We present IV results and estimates of a dose-response relationship in Section 4.5.

In addition to presenting impacts on a normalized summary statistic of student learning, we also present impacts on the fraction of questions answered correctly on different domains of subject-level competencies (Table 3). The ITT effects are positive and significant across all domains of test questions. In math, these range from a 12% increase on the easiest type of questions (arithmetic computation), determined by the proportion correctly answered in the control group, to a 36% increase on harder competencies such as geometry and measurement. Similarly, in Hindi, ITT effects range from a 7% gain on the easiest items (sentence completion)

¹⁶We use robust Huber-White standard errors throughout the paper rather than clustered standard errors because of the individual (as opposed to group) randomization of students to treatment status. Common shocks from test day and venue effects are netted out by the inclusion of strata fixed effects since all students in the same stratum (both treatment and control), were tested on the same day in the same location.

¹⁷This interpretation is possible because the baseline and endline tests are linked to a common metric using Item Response Theory. This would not be possible if scores were normalized within grade-subject-*period* as is common practice. Note that treatment effects are very similar (0.36σ in math and 0.21σ in Hindi) when test scores are normalized relative to the within-grade distribution in the control group at the endline (Table A.2).

to a 19% gain on the hardest competence (answering questions based on interpreting and integrating ideas and information from a passage).

4.3 Heterogeneity

We investigate whether ITT effects vary by gender, socio-economic status, or initial test scores, using a linear interaction specification and find no evidence of heterogeneity on these dimensions (Table 4). Since baseline test scores are a good summary statistic of prior inputs into education, we also present non-parametric estimates of the ITT effect as a function of baseline scores. We do this by plotting kernel-weighted locally-smoothed means of the endline test scores at each percentile of the baseline test-score distribution, separately for the treatment and control groups (Figure 3). In both math and Hindi, we see that the test scores in the treatment group are higher than those in the control group at every percentile of baseline test scores, and that the gains appear similar at all percentiles.

Next, we test for equality of treatment effects at different points of the *within-grade* test-score distribution. We do this by regressing endline test scores on the baseline test scores, indicator variables for treatment and for within-grade terciles at baseline, and interaction terms between the treatment variable and two terciles (the regression is estimated without a constant). We see no evidence of heterogeneity here as well (Table 5). The coefficient on the treatment dummy itself is statistically significant, but the interaction terms of treatment with the tercile at baseline are statistically indistinguishable from zero.

Note, however, that we see considerable heterogeneity in student progress by initial learning level *in the control group*. While students in the top third of the baseline test-score distribution show significant academic progress between baseline and endline, it is striking that we cannot reject the null of *no increase* in test scores for the bottom-third of students in the control group over the same period (with coefficients close to zero in both subjects) suggesting that lower-performing students appear to make no academic progress under the status quo.¹⁸

Thus, winning a voucher appears to have benefited students at all parts of the achievement distribution significantly and relatively equally, suggesting that the Mindspark software could teach all students equally well. However, since students in the lowest tercile of the within-grade baseline test score distribution did not make any academic progress in the control group on either subject, the *relative* gains from the treatment (measured as a multiple of what students would have learnt in the absence of treatment) was much larger for the weaker-performing students even though absolute gains are similar across all students (Figure A.3).

¹⁸This result is consistent with evidence on slow test score growth within cohorts over time in repeated cross-sections in India (Pritchett, 2013), and with the patterns of results observed in experimental evaluations of education interventions in developing countries in the past decade (Glewwe et al., 2009; Duflo et al., 2011). However, to the best of our knowledge, we are the first to present direct evidence using panel data and vertically-linked IRT scores, that weaker students make no gains in absolute learning levels under business-as-usual curriculum and instruction in a developing country setting.

4.4 Grade-level impact decomposition, and impacts on school tests

One consequence of the substantial deficits in student preparation (Figure 1), is that even large absolute increases in learning may not be sufficient for raising test scores on grade-level assessments since the gains in learning could be taking place below grade-level. We therefore use the CAL system data to directly examine the grade-level distribution of content presented by the software to students in the treatment group (see Figure A.5). In math, most of the content presented to students by Mindspark was below grade level, with very little content at the level of the grade the student is enrolled in. However, in Hindi, in addition to lower-grade content, a substantial portion of the Mindspark instruction in each grade was at grade level.

We find that distribution of test-score gains are consistent with the pattern of instruction on the CAL platform described above. Table 6 presents separate estimates of treatment effects on the proportion of test questions answered correctly at and at below grade level.¹⁹ We see that while there were large treatment effects in math on items below grade level, there was *no impact on grade-level questions*. In Hindi, on the other hand, we find that the treatment effect is significant for both questions at and below grade level.

These patterns in our data are also replicated in the independent data we collected on test scores on school exams. Table 7 presents the treatment effect of being offered a voucher on scores on the annual end of year school exams held in March 2016.²⁰ Mirroring the results on grade-level items on our own tests, we find a significant increase in test scores of 0.19σ in Hindi but no significant effect on math. We also do not find significant effects on the other subjects (science, social science, or English), although all the point estimates are positive.

Both sets of results (on our tests and on the school tests) are consistent with the data on grade-level distribution of questions presented by Mindspark shown in Figure A.5. They also provide a stark illustration of the importance of conducting education research with well-calibrated tests that are informative over a wide range of student achievement (especially in developing country settings with wide variation in within-grade student learning). In our case, relying on grade-level assessments would have led to *strikingly incorrect inference* regarding program impacts, and have led to a conclusion that the program had no impact on math despite the very large gains in test scores seen on a properly constructed test. See Appendix D for further details on test design for our study, and Muralidharan (2017) for a detailed discussion on test construction for education research in general.

¹⁹Items on our tests, which were designed to capture a wide range of achievement, were mapped into grade-levels with the help of a curriculum expert.

²⁰In Delhi, test papers for the annual exam are common across schools for each subject in each grade. In our regressions, we normalize test scores to $\mu=0$, $\sigma=1$ in each grade/subject in the control group.

4.5 IV estimates of dose-response relationship

All the results presented so far are ITT estimates, which are based on an average attendance of about 58% among lottery-winners.²¹ In this section, we present LATE estimates of the impact of actually attending the Mindspark centers, and (with further assumptions) estimates of predicted treatment effects at different levels of program exposure. We estimate the dose-response relationship between days of attendance and value-added using:

$$Y_{is2} = \alpha + \gamma \cdot Y_{is1} + \mu_1 \cdot Attendance_i + \eta_{is2} \quad (2)$$

where Y_{ist} is defined as previously, $Attendance$ is the number of days a student logged in to the Mindspark system (which is zero for all lottery-losers) and η_{ist} is the error term. Since program attendance may be endogenous to expected gains from the program, we instrument for $Attendance$ with the randomized offer of a voucher.

The IV estimates suggest that, on average, an extra day of attending the Mindspark centers increased test scores by 0.0065σ in math and 0.004σ in Hindi (Table 8: Cols. 1-2). These estimates identify the average causal response (ACR) of the treatment which “captures a weighted average of causal responses to a unit change in treatment (in this case, an extra day of attendance), for those whose treatment status is affected by the instrument” (Angrist and Imbens, 1995). Using these IV estimates to predict the effect of varying the number of days attended requires further assumptions about (a) the nature of heterogeneity in treatment effects across students (since the ACR is only identified over a subset of compliers, and not the full sample) and (b) the functional form of the relationship between days attended and the treatment effect (since the ACR averages causal effects over different intensities of treatment).

We present three pieces of suggestive evidence that constant treatment effects across students may be a reasonable assumption in this setting. First, the ITT effects were constant across the full distribution of initial achievement, which is a good summary measure for relevant individual-specific heterogeneity (Figure 3, Tables 4, 5). We also found no significant evidence of treatment heterogeneity across observed pre-treatment characteristics (Table 4).

Second, we cannot reject the equality of the IV estimates of Eq.(3) and the OLS estimates using a value-added (VA) specification (Table 8: Cols. 3-4), which suggests that the ATE and the LATE may be similar here. For both math and Hindi, the p-value from the difference-in-Sargan

²¹About 13% of lottery-winners attended for one day or less. The mean attendance among the rest was 57 days (around 66%). Figure A.2 plots the distribution of attendance among lottery winners, and Table A.3 presents correlations of attendance among lottery winners with various baseline characteristics.

test (similar to a Hausman test, but allowing for heteroskedasticity) testing equivalence of OLS and IV results is substantially greater than 0.1 (Cols. 1-2).²²

Finally, the constant term in the OLS VA specifications (corresponding to zero attendance) is identical when estimated using the full sample and when estimated using only the data in the treatment group (Table 8: Cols. 3-6). The constant term is identified using both the control group and “never-takers” when using the full sample, but is identified over only the “never-takers” when the sample is restricted to lottery-winners. Thus, the similarity of outcomes for the “never takers” and the control group, suggests equality of potential outcomes across different compliance groups.²³

We next explore the functional form of the relationship between days attended and the treatment effect both graphically (by plotting value-added against attendance for the lottery winners) and analytically. The graphical analysis suggests a linear relationship in both subjects (Figure 4). Further, while test-score value added is strongly correlated with the number of days attended in a linear specification (Table 8: Cols. 3-6), adding a quadratic term does not improve fit, and the quadratic term is not significant (see Table A.4). A linear dose-response is additionally plausible when considering the adaptive nature of the intervention which allows it to be equally effective regardless of the initial learning level of the student or the rate of academic progress. Thus, diminishing returns to program exposure may not apply over the relatively short duration of treatment in this study (which is consistent with the pattern seen in Figure 4).

Under the assumptions of constant treatment effects and a linear dose-response relationship, both of which appear reasonable in this context, our IV results suggest that attending Mindspark centers for 90 days, which roughly corresponds to half a school year with 80% attendance, would lead to gains of 0.59σ in math and 0.37σ in Hindi (last row of Table 8). We extrapolate results to 90 days, rather than a full school year, to keep the predictions near the range of the program exposure provided by our experiment (the maximum was 86 days). Similar or longer durations of program exposure would be feasible, even at observed attendance rates, if for instance the intervention started at the beginning of the school year rather than midway as in this study.

These estimates are conservative and likely to understate the dose-response relationship because the *Attendance* variable includes time spent in the Mindspark centers on instruction

²²Note that this close correspondence between the OLS VA and IV estimates is consistent with much recent evidence that VA models typically agree closely with experimental and quasi-experimental estimates (see, for instance Chetty et al. (2014); Deming et al. (2014); Singh (2015); Angrist et al. (2016))

²³This test is similar in spirit to tests suggested by Bertanha and Imbens (2014) and Brinch et al. (2016), for extending the validity of RD and IV estimates beyond LATE to average treatment effects.

in other subjects that we do not test (especially English).²⁴ In Table A.5, we present analogous IV and value-added estimates which only account for days spent by students on the subjects that we test (math and Hindi). Using these results, and the same assumptions as above, we estimate that 90 days of Mindspark attendance, split equally between the two subjects, would lead to test score gains of 0.76σ in math and 0.5σ in Hindi (last row of Table A.5).

4.6 Robustness

4.6.1 Attrition

Since the difference in attrition between the treatment and control groups is significant at the 10% level (Table 1), we test the robustness of our results to attrition by modeling selection into the endline based on observed characteristics, and present inverse probability weighted treatment effects: the estimated ITT effects are almost unchanged (Table A.6). We also compute Lee (2009) bounds for the ITT effect: although bounds are wide, the treatment effects are always positive and significant (Table A.7).

4.6.2 Familiarity with test questions

Our independent tests used items from several external assessments, some of which (in the Indian setting) were designed by EI; this raises the possibility that results on our assessments are overstated due to duplication of items between our tests and the Mindspark item bank. Note that this item bank contains over 45,000 items and so mere duplication in the database does not imply that a student would have been presented the same item during the intervention. Nevertheless, we test for this concern by computing the treatment effect expressed as the proportion correct on items from EI assessments and items from other assessments. The ITT effects are positive, statistically significant and of similar magnitude for both sets of items in math and Hindi (Table A.8).

4.6.3 Private Tutoring

Our results may also be confounded if winning a Mindspark voucher led to changes in the use of private tutoring. To test for this possibility, we collected data from parents of study students, using phone surveys, on whether the student attended paid extra tutoring (other than Mindspark) in any subject for each month from July 2015 to March 2016. Dividing this period into “pre-intervention” (July to September 2015) and “post-intervention” (October 2015 to March 2016), we test whether winning a Mindspark-voucher affected the incidence of private tutoring in the “post-intervention” period. We present these results in Table A.9. While there is a modest increase in private tutoring for all students in the post-treatment period (consistent with increased tutoring closer to annual school exams), we find no evidence of any differential use of private tutoring among lottery winners.

²⁴See Muralidharan and Sundararaman (2015) for an illustration of the importance of accounting for patterns of time use across subjects for inference regarding the productivity of education interventions.

5 Discussion

5.1 Mechanisms

The estimates presented above reflect a combination of the CAL software, group teaching, and additional instructional time, and we cannot experimentally identify the relative contribution of these channels. In this section, we present four sets of additional evidence that each point to the CAL system being the critical factor driving the large test-score gains we find.

The first and most important piece of evidence comes from a contemporaneous study conducted in the same location and student age group: Berry and Mukherjee (2016) report results from a randomized evaluation that studied the impact of after-school private tutoring on learning outcomes of middle-school students (in grades 6-8) in Delhi at the same time as our study. The program also provided six days of instruction per week, for two hours per day (versus 1.5 hours per day at Mindspark centers), and also charged INR 200 per month. The tutoring program was run by a well-respected and motivated non-profit organization, Pratham, who have run several education programs in India that have been found to have significant positive impacts on student learning at the primary level (see, for example, Banerjee et al. (2007, 2016)). Despite several similarities, there were two key differences between this program and the Mindspark centers. First, this program focused on reinforcing grade-level curriculum and was not customized to students' academic preparation; second, the instruction was delivered in person by a tutor in groups of up to 20 students (a similar ratio of instructor to student as seen in Mindspark centers), but did not make use of any technology for instruction.

At the end of a year of the program, Berry and Mukherjee (2016) find *no impact* on student test scores in independent assessments of either math or language despite the program having spent more than twice the after-school instructional time provided by the Mindspark centers during our evaluation (33% more instruction/week, and evaluated after a full year as opposed to 4.5 months). These results suggest that additional instructional time with group-tutoring (the other two components of our intervention in addition to the CAL) on their own may not have had much impact on learning. They also suggest that the binding constraint to student learning in this setting was not instructional time, but the (likely) ineffectiveness of additional instructional time spent on the default of teaching at a grade-appropriate level in a setting where most students are several grade levels behind (as seen in Figure 1).

Second, we provide direct evidence that the CAL software effectively addressed this constraint to effective pedagogy by targeting instructional material at the level of each individual student, and thereby accommodating the wide variation in student preparation documented in Figure 1. We see this in Figure 5, where the horizontal axis on each subgraph shows the assessed level of academic preparedness of each student enrolled in a given grade, and the vertical axis shows that the CAL software presented students with material that is either *at their grade*

level or at adjacent grade levels.²⁵ Further, the CAL system not only accommodates variation in initial learning levels, but also in the pace of learning across students. Figure 6 presents non-parametric plots of the average difficulty level of the math items presented to students over the course of the intervention, documenting that the software updates its estimate of student achievement levels in real time and modifies instruction accordingly. The individualization of the dynamic updating of content is highlighted further in Figure A.7 where we use student-level data to plot similar trajectories separately for *each student* in the treatment group.

Teaching effectively in a setting with such large heterogeneity in the levels and trajectories of student learning within the same grade would be very challenging even for well trained and motivated teachers. In contrast, once the CAL software is programmed to present content based on a student’s assessed learning level and to adjust content at the rate of student progress, the software can handle additional heterogeneity at zero marginal cost, which is not true for a teacher.²⁶ Thus, the CAL software was likely to have been the key enabler for *all* students to be able to learn relative to the default of grade-appropriate pedagogy in a standard classroom setting (or in an after-school group tutoring setting).

Third, data on assignment of students into Mindspark batches (who would attend group instruction together) strongly suggests that teaching was mainly taking place on the CAL platform, with the role of the instructor being to promote adherence. We see this clearly in Figure A.6, which shows that the students in our study (who are mainly in grades 6-9), were assigned to Mindspark batches that often included students enrolled in *grades 1-5 in the same batch*. This is because EI’s main consideration in assigning students to batches was timing convenience of students and parents and so EI was not concerned about having students ranging from *grades 1-9 in the same batch*, which is a classroom set up that would make very little sense for group instruction.²⁷

Finally, note that the patterns of test score results we present in Section 4.4 are also consistent with instruction being driven mainly by the software. All the gains in math test scores were seen on below grade-level questions (which is what the CAL software taught) and not on grade-level questions (which were not taught by the CAL software).

²⁵In both math and Hindi, we use data from a single day which is near the beginning of the intervention, after all students would have completed their initial assessment, and when Mindspark computer-aided instruction in the relevant subject was scheduled in all three centers.

²⁶Note that the strength of the software lies not just in its ability to personalize the level of instruction, but to do so with uniformly high-quality content at all levels (with the features described in Section 2.1). Even if a teacher wanted to review lower-grade materials in class, it would be very challenging to effectively prepare material spanning several grades and present differentiated content across students in a classroom setting.

²⁷Note that prior evidence on positive impacts of group-based instruction has highlighted the importance of *homogenizing* the groups by learning level for effective instruction (Banerjee et al., 2007, 2016). Thus, it is highly unlikely that EI would have chosen to have batches that spanned so many grades unless they believed that the group instruction was second order to the instruction on the CAL system.

These four pieces of evidence all suggest that the CAL software was the key driver of the results we find. Yet, according to EI the instructor did have an important role in promoting adherence by encouraging regular student attendance at the centers, ensuring time on task while students were in front of the computer, and supervising school homework completion and exam preparation during the group-instruction period (which parents demanded). This discussion suggests that there may be complementarities between teachers and technology. So, our results should not be interpreted as the impact of CAL software by itself, but rather as an estimate of the effect of CAL in a setting where there was also an instructor to support adherence to the CAL. Alternatively, given the null results of instructor-led after-school group tutoring found by (Berry and Mukherjee, 2016)), our results can also be interpreted as showing the extent to which using technology in education can raise the productivity of an instructor.

5.2 Cost-effectiveness

Since we evaluate an after-school program, a natural comparison of cost effectiveness is with after-school private tutoring, which is widespread in our setting. The direct comparison with the results in Berry and Mukherjee (2016) suggest that after-school group-based tutoring on grade-level materials had no impact on learning in the same context even with over double the duration of exposure relative to the program we study.

A second policy-relevant comparison is with the productivity of government-run schools (from where the study subjects were recruited). Per-pupil monthly spending in these schools in Delhi was around INR 1500 (USD 22) in 2014-15; students spend 240 minutes per week on math and Hindi; and we estimate that the upper-bound of the value-added in these schools was 0.36σ in math and 0.15σ in Hindi over the 4.5-month study period. Specifically, this was the *total* value-added in the control group in Table 2, which also includes the effects of home inputs and private tutoring, and therefore likely over-estimates the value-added in public schools.

Using our ITT estimates, we see that Mindspark added 0.36σ in math and 0.22σ in Hindi over the same period in around 180 minutes per week on each subject. The Mindspark program, as delivered, had an unsubsidized cost of about INR 1000 per student (USD 15) per month. This includes the costs of infrastructure, hardware, staffing, and pro-rated costs for software development. Thus, even when implemented with high fixed costs and without economies of scale, and based on 58% attendance, providing access to the Mindspark centers delivered greater learning at lower financial and time cost than default public spending.

Steady-state costs of Mindspark at policy-relevant scales are likely to be much lower since the (high) fixed costs of product development have already been incurred. If implemented in government schools, at even a modest scale of 50 schools, per-pupil costs reduce to about USD 4 per month (including hardware costs). Above a scale of 1000 schools, the per-pupil marginal costs (software maintenance and technical support) are about USD 2 annually, which is a small

fraction of the USD 150 annual cost (over 10 months) during our pilot.²⁸ The program thus has the potential to be very cost-effective at scale.

Further, while education spending can increase continuously over time, student time is finite. Thus, it is also useful to evaluate the effectiveness of education interventions per unit of time, *independent* of financial cost. A useful point of comparison is provided by Muralidharan (2012), who finds that providing individual-level performance bonuses to teachers in India led to test score gains of 0.54σ and 0.35σ in math and language for students exposed to the program for five years. This is one of the largest effect sizes seen to date in an experimental study on education in developing countries. Yet, we estimate that regularly attending Mindspark centers for half a year would yield similar gains (in one tenth the time).²⁹

Figure 6 suggests that students who received access to the Mindspark centers improved a full grade-level in math over just 4.5 months (even with only 58% attendance). Thus, using Mindspark regularly in schools may be an especially promising option for helping to bridge the large gaps in student readiness *within time frames* that may make it feasible for lagging students to catch up to grade-level standards of instruction. Testing this possibility is an important topic for future research.

5.3 Policy Implications

Despite the large test-score gains we find, parental demand for Mindspark centers was low in the absence of the (fee-waiving) vouchers. In fact, all three centers in our study closed down soon after the conclusion of our experiment in the face of low parental willingness to pay (even at the subsidized price).³⁰ Thus, models of technology-aided instruction that charge fees may limit the ability of low-income students to access the programs. As a result, effectively deploying education technology in public schools is likely to be important for providing access to CAL programs to the most disadvantaged students.

This belief is reflected in the growing policy interest around the world in using technology in public education. However, policy makers (especially in developing countries) have mainly

²⁸These numbers are based on an actual budget for deploying Mindspark in government schools that was prepared and submitted by EI in 2017.

²⁹Of course, it is likely that some of these gains will fade out over time as was seen in Banerjee et al. (2007). However, it is now well-known that the effects of *all* education interventions decay over time (Jacob et al., 2010; Andrabi et al., 2011). This is why we do not claim that extending the Mindspark program for 5 years will lead to ten times greater test score gains, but simply note that the gains observed over 5 years in Muralidharan (2012) were achieved in one-tenth the time here.

³⁰The donors who subsidized the fees at Mindspark centers stipulated that they would only continue funding the subsidies if the centers could operate at or above 80% capacity (and thereby demonstrate parental willingness to pay at least the subsidized price). In practice, enrolment levels were considerably below this target, and the centers had to shut down because philanthropic funding for the subsidies ended. Since the centers shut down in March 2016, control group students who had been offered free access to the centers after the endline test, were instead offered free educational materials as compensation for participating in the study. However, Mindspark as a product is doing well and EI continues to operate and improve the full-fee Mindspark models for higher SES families, where the demand continues to be strong.

concentrated on providing computer hardware without commensurate attention to using technology to improve pedagogy.³¹ Our results (combined with the review of evidence in Appendix B), suggest that these hardware investments are likely to yield much greater returns in terms of improved learning outcomes if attention is also paid to deploying Mindspark (or similar) software to improve pedagogy in public schools.

Our results are also relevant for policy debates on the best way to teach effectively in settings with large variation in student preparation. One widely considered policy option is tracking of classrooms, but this may not be feasible in many developing-country settings.³² Further, even when feasible, tracking is controversial and the global evidence on its impact is mixed (Betts, 2011). Our results suggest that well-designed CAL programs may be able to deliver the pedagogical advantages of tracking while mitigating several limitations, as listed below.

First, CAL instruction allows *complete* personalization, whereas tracked classrooms still have to cater to variation in student learning levels and trajectories with a common instruction protocol. Second, by allowing students to work at their own pace, it avoids potential negative effects of students being labelled as being in a weaker track. Third, the dynamic updating of content mitigates the risk of premature permanent tracking of ‘late bloomers’. Fourth, it allows instruction to be differentiated without changing peers in the classroom. Fifth, relative to policies of grade retention or accelerated grade promotion, using CAL programs in classrooms makes it possible to preserve the age-cohort based social grouping of students (which may allow for better development of socio-emotional skills), while allowing for variation in academic content presented.

6 Conclusion

We present experimental evidence on the impact of a technology-led supplementary instruction program in post-primary grades in urban India, and find that gaining access to the program led to large and rapid test-score gains in both math and language. The program was effective at teaching students at all levels of prior achievement, including academically-weaker students who are left behind by business-as-usual instruction. Using detailed data on the material presented to students in the treatment group, we show that the software was successful at targeting instruction precisely to each students’ level of achievement and in handling wide variation in the academic preparation of students enrolled in the same grade.

³¹For instance, various state governments in India have distributed free laptops to students in recent years. Further, progress on implementing the national-level policy on technology in education is typically measured by the number of schools with computer labs.

³²Unlike in developed countries where students in middle and high schools can choose their subjects and can take easier and more advanced courses, most developing-country education systems are characterized by preparing students for a single high-stakes school leaving examination. Thus, the default organization of schools is to have all students in a given grade in the same classroom with the teacher focusing instruction on completing the curriculum mandated by official text books for the corresponding grade.

These substantial effects reflect, in our opinion, the ability of the intervention to effectively target *multiple* constraints that lead to the low productivity of instructional time in Indian schools. The high quality of content, combined with effective delivery and interface, may help circumvent constraints of teacher human capital and motivation. Personalized instruction makes it possible to accommodate large deficits in initial student preparation and wide variation within a single grade. Algorithms for analyzing patterns of student errors and providing differentiated feedback and follow up content that is administered in real-time, allows for feedback that is more relevant and much more frequent. These features all reflect continuous and iterative program development over a long period of more than a decade.

These effects may plausibly be increased even further with better design. It is possible that in-school settings may have greater adherence to the program in terms of attendance. It may also be possible to improve the effectiveness of teacher-led instruction in a ‘blended learning’ environment by using the extensive information on student-performance to better guide teacher effort in the classroom. This ‘big data’ on student achievement also offers much potential of its own. In particular, such a setting may enable high-frequency randomized experiments on effective pedagogical techniques and approaches (which may vary across students) and build a stronger evidence base on effective teaching practices. This evidence may then be used to further optimize the delivery of instruction in the program and, plausibly, also for the delivery of classroom instruction. Finally, the detailed and continuous measures of effort input by the students can be used directly to reward students, with potentially large gains in student motivation, effort, and achievement.³³

However, there are also several reasons to be cautious in extrapolating the success of the program more broadly. The intervention, as evaluated in this paper, was delivered at a modest scale of a few centers in Delhi and delivered with high fidelity on part of the providers. Such fidelity may not be possible when implementing at scale. Additional issues relate to the mode of delivery. We have only evaluated Mindspark in after-school centers and it is plausible that the effectiveness of the system may vary significantly based on whether it is implemented in-school or out-of-school; whether it is supplementary to current classroom instruction or substitutes away current instructional time; and whether it is delivered without supervision, under the supervision of current teachers or under the supervision of third parties (e.g. the Mindspark center staff). Identifying the most effective modes of delivery for the program at larger scale is an important area for future research.³⁴

³³Direct evidence that this may be possible is provided by Hirshleifer (2015) who uses data from a (different) computer-aided instruction intervention to reward student effort and documents large effects of 0.57σ .

³⁴A useful example of such work has been the literature that followed the documenting of the efficacy of unqualified local volunteers, who were targeting instruction to students’ achievement levels, in raising achievement in primary schools in two Indian cities by Banerjee et al. (2007). Subsequent studies have looked at the effectiveness of this pedagogical approach of “Teaching at the Right Level” in summer camps,

A further point of caution is that our results should not be interpreted as supporting a de-emphasis of the role of teachers in education. Rather, since the delivery of education involves several non-routine tasks that vary as a function of individual students and situations, and requires complex contextually-aware communication, it is more likely that technology will complement rather than substitute teachers (as shown more generally by Autor et al. (2003)). So, it may be possible to improve teacher and school productivity by using technology to perform routine tasks (such as grading) and data-analysis intensive tasks (such as identifying patterns in student answers and providing differentiated feedback and instruction to students), and enabling teachers to spend more time on aspects of education where they may have a comparative advantage - such as supporting group-based learning strategies that may help build social and other non-cognitive skills that may have considerable labor market returns (Cunha et al., 2010; Heckman and Kautz, 2012; Deming, 2016).

Overall, our present study is best regarded as an efficacy trial documenting that well-designed and implemented technology-enabled learning programs can produce large gains in student test scores in a relatively short period of time. Our results suggest that the promise of technology to sharply improve productivity in the delivery of education is real, and that there may be large returns to further innovation and research on effective ways of integrating technology-aided instruction into classrooms, and on effective ways of delivering these benefits at a larger scale.

in government schools and delivered alternately by school teachers and by other volunteers (Banerjee et al., 2016). The approach is now being extended at scale in multiple state education systems.

References

- Andrabi, Tahir, Jishnu Das, Asim I. Khwaja, and Tristan Zajonc**, “Do value-added estimates add value? Accounting for learning dynamics,” *American Economic Journal: Applied Economics*, 2011, 3 (3), 29–54.
- Angrist, Joshua and Guido Imbens**, “Two-stage least squares estimation of average causal effects in models with variable treatment intensity,” *Journal of the American Statistical Association*, 1995, 90 (430), 431–442.
- **and Victor Lavy**, “New evidence on classroom computers and pupil learning,” *The Economic Journal*, 2002, 112 (482), 735–765.
- **, Peter Hull, Parag Pathak, and Christopher Walters**, “Leveraging lotteries for school value-added: Testing and estimation,” *The Quarterly Journal of Economics*, 2016, *Forthcoming*.
- Autor, David, Frank Levy, and Richard J. Murnane**, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 2003, 118 (4), 1279–1333.
- Banerjee, A. V., R. Banerji, J. Berry, E. Duflo, H. Kannan, S. Mukerji, M. Shotland, and M. Walton**, “Mainstreaming an effective intervention: Evidence from randomized evaluations of Teaching at the Right Level in India,” 2016. *Journal of Economic Perspectives*, forthcoming.
- Banerjee, Abhijit and Esther Duflo**, *Poor economics: A radical rethinking of the way to fight global poverty*, New York, NY: Public Affairs, 2012.
- **, Paul Glewwe, Shawn Powers, and Melanie Wasserman**, “Expanding access and increasing student learning in post-primary education in developing countries: A review of the evidence,” Technical Report, Abdul Latif Jameel Poverty Action Lab 2013.
- Banerjee, Abhijit V, Shawn Cole, Esther Duflo, and Leigh Linden**, “Remedying Education: Evidence from Two Randomized Experiments in India,” *The Quarterly Journal of Economics*, 2007, 122 (3), 1235–1264.
- Barrera-Osorio, Felipe and Leigh L Linden**, “The use and misuse of computers in education: evidence from a randomized experiment in Colombia,” 2009. (World Bank Policy Research Working Paper No. 4836.) Washington, DC: The World Bank.
- Barrow, Lisa, Lisa Markman, and Cecilia Elena Rouse**, “Technology’s edge: The educational benefits of computer-aided instruction,” *American Economic Journal: Economic Policy*, 2009, 1 (1), 52–74.
- Berry, J. and P. Mukherjee**, “Pricing of private education in urban India: Demand, use and impact,” 2016. *Unpublished manuscript*. Ithaca, NY: Cornell University.
- Bertanha, Marinho and Guido Imbens**, “External Validity in Fuzzy Regression Discontinuity Designs,” Technical Report 20773, National Bureau of Economic Research, Inc 2014.

- Betts, Julian**, “The Economics of Tracking in Education,” in Eric Hanushek, Stephen Machin, and Ludger Woessmann, eds., *Handbook of the Economics of Education*, Elsevier, 2011, pp. 341–381.
- Beuermann, Diether W, Julian Cristia, Santiago Cueto, Ofer Malamud, and Yyannu Cruz-Aguayo**, “One Laptop per Child at home: Short-term impacts from a randomized experiment in Peru,” *American Economic Journal: Applied Economics*, 2015, 7 (2), 53–80.
- Borman, G. D., J. G. Benson, and L. Overman**, “A randomized field trial of the Fast ForWord Language computer-based training program,” *Educational Evaluation and Policy Analysis*, 2009, 31 (1), 82–106.
- Bosworth, B.**, “The Internet and the university,” 2005. In Devlin, M., Larson, R. & Meyerson, J. (eds.) *Productivity in education and the growing gap with service industries*. Cambridge, MA: Forum for the Future of Higher Education & Boulder, CO: EDUCAUSE.
- Brinch, Christian, Magne Mogstad, and Matthew Wiswall**, “Beyond LATE with a discrete instrument,” *Journal of Political Economy*, 2016, *Forthcoming*.
- Bulman, G. and R.W. Fairlie**, “Technology and Education: Computers, Software and the Internet,” in Eric Hanushek, Stephen Machin, and Ludger Woessmann, eds., *Handbook of the Economics of Education*, Elsevier, 2016, pp. 239–280.
- Buswell, Guy Thomas and Charles Hubbard Judd**, *Summary of educational investigations relating to arithmetic* number 27, University of Chicago, 1925.
- Campuzano, L., M. Dynarski, R. Agodini, K. Rall, and A. Pendleton**, “Effectiveness of reading and mathematics software products: Findings from two student cohorts,” 2009. *Unpublished manuscript*. Washington, DC: Mathematica Policy Research.
- Carrillo, Paul E, Mercedes Onofa, and Juan Ponce**, “Information technology and student achievement: Evidence from a randomized experiment in Ecuador,” 2010. (IDB Working Paper No. IDB-WP-223). Washington, DC: Inter-American Development Bank.
- Chetty, Raj, John N Friedman, and Jonah E Rockoff**, “Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates,” *The American Economic Review*, 2014, 104 (9), 2593–2632.
- Cohen, Jessica and Pascaline Dupas**, “Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment,” *The Quarterly Journal of Economics*, 2010, 125 (1), 1–45.
- Cristia, Julian, Pablo Ibarrarán, Santiago Cueto, Ana Santiago, and Eugenio Severín**, “Technology and child development: Evidence from the One Laptop per Child program,” 2012. (IDB Working Paper No. IDB-WP-304). Washington, DC: Inter-American Development Bank.
- Cunha, Flavio, James J. Heckman, and Susanne M. Schennach**, “Estimating the Technology of Cognitive and Noncognitive Skill Formation,” *Econometrica*, 2010, 78 (3), 883–931.


- Deming, David J.**, “The Growing Importance of Social Skills in the Labor Market,” 2016. Cambridge, MA: Harvard University.
- , **Justine S. Hastings, Thomas J. Kane, and Douglas O. Staiger**, “School choice, school quality, and postsecondary attainment,” *American Economic Review*, 2014, *104* (3), 991–1013.
- Duflo, E., P. Dupas, and M. Kremer**, “Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in Kenya,” *American Economic Review*, 2011, *101*, 1739–1774.
- Dynarski, M., R. Agodini, S. Heaviside, T. Novak, N. Carey, L. Campuzano, B. Means, R. Murphy, W. Penuel, H. Javitz, D. Emery, and W. Sussex**, “Effectiveness of reading and mathematics software products: Findings from the first student cohort,” 2007. *Unpublished manuscript*. Washington, DC: Mathematica Policy Research.
- Fairlie, R. W. and J. Robinson**, “Experimental Evidence on the Effects of Home Computers on Academic Achievement among Schoolchildren,” *American Economic Journal: Applied Economics*, 2013, *5* (3), 211–240.
- Fujiwara, Thomas**, “Voting technology, political responsiveness, and infant health: Evidence from Brazil,” *Econometrica*, 2015, *83* (2), 423–464.
- Glewwe, Paul and Karthik Muralidharan**, “Improving School Education Outcomes in Developing Countries: Evidence, Knowledge Gaps, and Policy Implications,” in Eric Hanushek, Stephen Machin, and Ludger Woessmann, eds., *Handbook of the Economics of Education*, Elsevier, 2016, pp. 653–744.
- , **Michael Kremer, and Sylvie Moulin**, “Many children left behind? Textbooks and test scores in Kenya,” *American Economic Journal: Applied Economics*, 2009, *1* (1), 112–135.
- Goolsbee, Austan and Jonathan Guryan**, “The impact of Internet subsidies in public schools,” *The Review of Economics and Statistics*, 2006, *88* (2), 336–347.
- Heckman, James J. and Tim Kautz**, “The Economics of Human Development and Social Mobility,” *Labour Economics*, 2012, *19* (4), 451–464.
- Hirshleifer, Sarojini**, “Incentives for effort or outputs? A field experiment to improve student performance,” 2015. *Unpublished manuscript*. Cambridge, MA: Abdul Latif Jameel Poverty Action Lab (J-PAL).
- Jack, W. and T. Suri**, “Risk sharing and transactions costs: Evidence from Kenya’s mobile money revolution,” *The American Economic Review*, 2014, *104* (1), 183–223.
- Jacob, Brian A, Lars Lefgren, and David P Sims**, “The persistence of teacher-induced learning,” *Journal of Human resources*, 2010, *45* (4), 915–943.
- Khan, Salman**, *The one world schoolhouse: Education reimagined*, Twelve, 2012.

- Kothari, Brij, Avinash Pandey, and Amita R Chudgar**, “Reading out of the ?idiot box?: Same-language subtitling on television in India,” *Information Technologies & International Development*, 2004, 2 (1), pp–23.
- , **Joe Takeda, Ashok Joshi, and Avinash Pandey**, “Same language subtitling: a butterfly for literacy?,” *International Journal of Lifelong Education*, 2002, 21 (1), 55–66.
- Lai, Fang, Linxiu Zhang, Qinghe Qu, Xiao Hu, Yaojiang Shi, Matthew Boswell, and Scott Rozelle**, “Does computer-assisted learning improve learning outcomes? Evidence from a randomized experiment in public schools in rural minority areas in Qinghai, China,” 2012. (REAP Working Paper No. 237). Rural Education Action Program (REAP). Stanford, CA.
- , – , **Xiao Hu, Qinghe Qu, Yaojiang Shi, Yajie Qiao, Matthew Boswell, and Scott Rozelle**, “Computer assisted learning as extracurricular tutor? Evidence from a randomised experiment in rural boarding schools in Shaanxi,” *Journal of Development Effectiveness*, 2013, 52 (2), 208–231.
- , **Renfu Luo, Linxiu Zhang, and Scott Huang Xinzhe Rozelle**, “Does computer-assisted learning improve learning outcomes? Evidence from a randomized experiment in migrant schools in Beijing,” *Economics of Education*, 2015, 47, 34–48.
- , – , – , **Xinzhe Huang, and Scott Rozelle**, “Does computer-assisted learning improve learning outcomes? Evidence from a randomized experiment in migrant schools in Beijing,” *Economics of Education Review*, 2015, 47, 34–48.
- Lee, David**, “Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects,” *The Review of Economic Studies*, 2009, 76, 1071–1102.
- Leuven, Edwin, Mikael Lindahl, Hessel Oosterbeek, and Dinand Webbink**, “The effect of extra funding for disadvantaged pupils on achievement,” *The Review of Economics and Statistics*, 2007, 89 (4), 721–736.
- Linden, L. L.**, “Complement or substitute? The effect of technology on student achievement in India,” 2008. Unpublished manuscript. Abdul Latif Jameel Poverty Action Lab (J-PAL). Cambridge, MA.
- Machin, Stephen, Sandra McNally, and Olmo Silva**, “New technology in schools: Is there a payoff?,” *The Economic Journal*, 2007, 117 (522), 1145–1167.
- Malamud, Ofer and C. Pop-Eleches**, “Home computer use and the development of human capital,” *The Quarterly Journal of Economics*, 2011, 126, 987–1027.
- Mead, R.**, “Learn different: Silicon Valley disrupts education,” 2016. *The New Yorker*. March 8, 2016.
- Mo, Di, Johan Swinnen, Linxiu Zhang, Hongmei Yi, Qinghe Qu, Matthew Boswell, and Scott Rozelle**, “Can one-to-one computing narrow the digital divide and the educational gap in China? The case of Beijing migrant schools,” *World development*, 2013, 46, 14–29.

- , **L. Zhang, J. Wang, W. Huang, Y. Shi, M. Boswell, and S. Rozelle**, “The persistence of gains in learning from computer assisted learning (CAL): Evidence from a randomized experiment in rural schools in Shaanxi province in China,” 2014. *Unpublished manuscript*. Stanford, CA: Rural Education Action Program (REAP).
 - , **Linxiu Zhang, Renfu Luo, Qinghe Qu, Weiming Huang, Jiafu Wang, Yajie Qiao, Matthew Boswell, and Scott Rozelle**, “Integrating computer-assisted learning into a regular curriculum: Evidence from a randomised experiment in rural schools in Shaanxi,” *Journal of Development Effectiveness*, 2014, 6, 300–323.
 - , **Yu Bai, Matthew Boswell, and Scott Rozelle**, “Evaluating the effectiveness of computers as tutors in China,” 2016.
- Morgan, P. and S. Ritter**, “An experimental study of the effects of Cognitive Tutor Algebra I on student knowledge and attitude,” 2002. Pittsburg, PA: Carnegie Learning.
- Muralidharan, Karthik**, “Long-term effects of teacher performance pay: Experimental evidence from India,” 2012. *Unpublished manuscript*. San Diego, CA: University of California, San Diego.
- , “Priorities for primary education policy in India’s 12th five-year plan,” *India Policy Forum 2012-13*, 2013, 9, 1–46.
 - , “Field Experiments in Education in Developing Countries,” in Abhijit Banerjee and Esther Duflo, eds., *Handbook of Field Experiments*, Elsevier, 2017.
 - **and Venkatesh Sundararaman**, “The impact of diagnostic feedback to teachers on student learning: Experimental evidence from India,” *The Economic Journal*, 2010, 120 (F187-F203).
 - **and –** , “The aggregate effect of school choice: Evidence from a two-stage experiment in India,” *The Quarterly Journal of Economics*, 2015, 130 (3), 1011–1066.
 - , **Paul. Niehaus, and Sandip. Sukhtankar**, “Building state capacity: Evidence from biometric smartcards in India,” *American Economic Review*, 2016, 106 (10), 2895–2929.
- Murphy, R., W. Penuel, B. Means, C. Korbak, and A. Whaley**, “E-DESK: A review of recent evidence on the effectiveness of discrete educational software,” 2001. *Unpublished manuscript*. Menlo Park, CA: SRI International.
- Negroponte, Nicholas, Walter Bender, Antonio Battro, and David Cavallo**, “One laptop per child,” in “Keynote address at National Educational Computing Conference in San Diego, CA. Retrieved April,” Vol. 5 2006, p. 2007.
- Pearson, P.D., R.E. Ferdig, R.L. Blomeyer Jr., and J. Moran**, “The effects of technology on reading performance in the middle-school grades: A meta-analysis with recommendations for policy,” 2005. *Unpublished manuscript*. Naperville, IL: Learning Point Associates.
- Pratham**, *Annual Status of Education Report 2016*, Pratham, New Delhi, 2017.

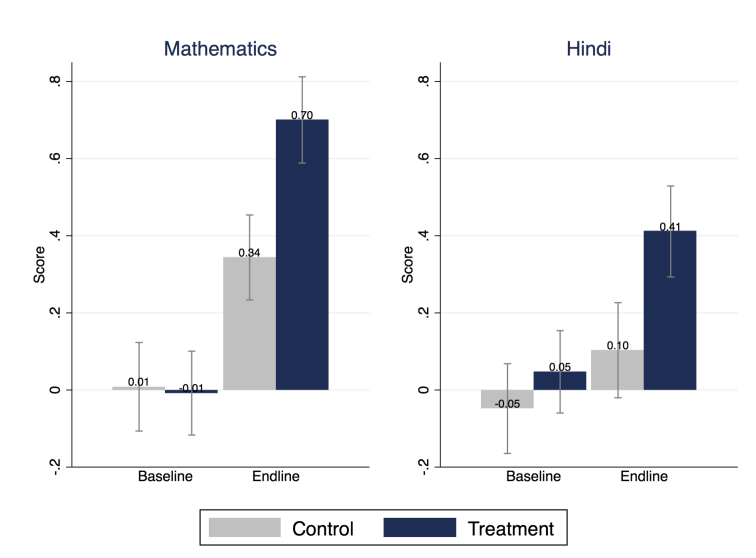
- Pritchett, Lant**, *The rebirth of education: Schooling ain't learning*, Washington, DC: Center for Global Development, 2013.
- **and Amanda Beatty**, “Slow down, you’re going too fast: Matching curricula to student skill levels,” *International Journal of Educational Development*, 2015, *40*, 276–288.
- Radatz, Hendrik**, “Error analysis in mathematics education,” *Journal for Research in mathematics Education*, 1979, pp. 163–172.
- Rockoff, Jonah E**, “Evaluation report on the School of One i3 expansion,” 2015. *Unpublished manuscript*. New York, NY: Columbia University.
- Rouse, Cecilia Elena and Alan B Krueger**, “Putting computerized instruction to the test: A randomized evaluation of a “scientifically based” reading program,” *Economics of Education Review*, 2004, *23* (4), 323–338.
- Singh, Abhijeet**, “Private school effects in urban and rural India: Panel estimates at primary and secondary school ages,” *Journal of Development Economics*, 2015, *113*, 16–32.
- Waxman, H.C., M.-F. Lin, and G.M. Michko**, “A meta-analysis of the effectiveness of teaching and learning with technology on student outcomes,” 2003. *Unpublished manuscript*. CambridgeNaperville, IL: Learning Point Associates.
- Wise, B. W. and R. K. Olson**, “Computer-based phonological awareness and reading instruction,” *Annals of Dyslexia*, 1995, *45*, 99–122.

Figure 1: Assessed levels of student achievement vs. current grade enrolled in school




Note: This figure shows, for treatment group, the estimated level of student achievement (determined by the Mindspark CAL program) plotted against the grade they are enrolled in. These data are from the *initial* diagnostic test, and do not reflect any instruction provided by Mindspark. In both subjects, we find three main patterns: (a) there is a general deficit between average attainment and grade-expected norms; (b) this deficit is larger in later grades and (c) within each grade, there is a wide dispersion of student achievement.

Figure 2: Mean difference in test scores between lottery winners and losers




Note: This figure shows mean of test scores, normalized with reference to baseline, across treatment and control groups in the two rounds of testing with 95% confidence intervals. Test scores were linked within-subject through IRT models, pooling across grades and across baseline and endline, and are normalized to have a mean of zero and a standard deviation of one in the baseline. Whereas baseline test scores were balanced between lottery-winners and lottery-losers, endline scores are significantly higher for the treatment group.

Figure 3: Non-parametric investigation of treatment effects by baseline percentiles




Note: The figures present kernel-weighted local mean smoothed plots which relate endline test scores to percentiles in the baseline achievement, separately for the treatment and control groups, alongside 95% confidence intervals. At all percentiles of baseline achievement, treatment group students see larger gains over the study period than the control group, with no strong evidence of differential absolute magnitudes of gains across the distribution.

Figure 4: Dose response relationship




Note: This figure explores the relationship between value-added and attendance in the Mindspark program among the lottery-winners. It presents the mean value-added in bins of attendance along with a linear fit and a lowess smoothed non-parametric plot.

Figure 5: Precise customization of instruction by the Mindspark CAL program



Note: This figure shows, for treatment group, the grade level of questions administered by the computer adaptive system to students on a single day near the beginning of the intervention. In each grade of enrolment, actual level of student attainment estimated by the CAL software differs widely; this wide range is covered through the customization of instructional content by the CAL software.

Figure 6: Dynamic updating and individualization of content in Mindspark



Note: This figure shows kernel-weighted local mean smoothed lines relating the level of difficulty of the math questions administered to students in the treatment group with the date of administration. The left panel presents separate lines by the actual grade of enrolment. The right panel presents separate lines by the level of achievement assessed at baseline by the CAL software. Please note 95% confidence intervals are plotted as well but, given the large data at our disposal, estimates are very precise and the confidence intervals are narrow enough to not be visually discernible.

Table 1: Sample descriptives and balance on observables

	Mean (treatment)	Mean (control)	Difference	SE	N (treatment)	N (control)
<u>Panel A: All students in the baseline sample</u>						
<i>Demographic characteristics</i>						
Female	0.76	0.76	0.00	0.03	314	305
Age (years)	12.68	12.48	0.20	0.13	306	296
SES index	0.00	0.05	-0.05	0.14	314	305
<i>Grade in school</i>						
Grade 4	0.01	0.01	-0.00	0.01	305	299
Grade 5	0.01	0.02	-0.01	0.01	305	299
Grade 6	0.27	0.30	-0.04	0.04	305	299
Grade 7	0.26	0.26	0.00	0.04	305	299
Grade 8	0.30	0.28	0.02	0.04	305	299
Grade 9	0.15	0.13	0.02	0.03	305	299
<i>Baseline test scores</i>						
Math	-0.01	0.01	-0.02	0.08	313	304
Hindi	0.05	-0.05	0.10	0.08	312	305
Present at endline	0.838	0.885	-0.048*	0.028	314	305
<u>Panel B: Only students present in Endline</u>						
<i>Demographic characteristics</i>						
Female	0.77	0.76	0.01	0.04	263	270
Age (years)	12.60	12.46	0.13	0.14	257	263
SES index	-0.10	0.04	-0.14	0.14	263	270
<i>Grade in school</i>						
Grade 4	0.01	0.01	-0.00	0.01	255	266
Grade 5	0.01	0.02	-0.01	0.01	255	266
Grade 6	0.29	0.31	-0.02	0.04	255	266
Grade 7	0.25	0.25	0.00	0.04	255	266
Grade 8	0.30	0.29	0.02	0.04	255	266
Grade 9	0.14	0.12	0.02	0.03	255	266
<i>Baseline test scores</i>						
Math	-0.03	-0.02	-0.02	0.09	262	269
Hindi	0.06	-0.07	0.13	0.08	263	270

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Treatment and control here refer to groups who were randomly assigned to receive an offer of Mindspark voucher till March 2016. Variables used in this table are from the baseline data collection in September 2015. The data collection consisted of two parts: (a) a self-administered student survey, from which demographic characteristics, details of schooling and private tutoring are taken and (b) assessment of skills in math and Hindi, administered using pen-and-paper tests. Tests were designed to cover wide ranges of achievement and to be linked across grades, as well as between baseline and endline assessments, using common items. Scores are scaled here using Item Response theory models and standardized to have a mean of zero and standard deviation of one in the baseline. The SES index refers to a wealth index generated using the first factor from a Principal Components Analysis consisting of indicators for ownership of various consumer durables and services in the household.

Table 2: Intent-to-treat (ITT) Effects in a regression framework

	(1)	(2)	(3)	(4)
	Dep var: Standardized IRT scores (endline)			
	Math	Hindi	Math	Hindi
Treatment	0.36*** (0.063)	0.22*** (0.076)	0.36*** (0.062)	0.22*** (0.064)
Baseline score	0.54*** (0.047)	0.67*** (0.034)	0.55*** (0.039)	0.69*** (0.039)
Constant	0.36*** (0.031)	0.15*** (0.038)	0.36*** (0.043)	0.15*** (0.045)
Strata fixed effects	Y	Y	N	N
Observations	529	533	529	533
R-squared	0.392	0.451	0.392	0.465

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ Treatment is a dummy variable indicating a randomly-assigned offer of Mindspark voucher till March 2016. Tests in both math and Hindi were designed to cover wide ranges of achievement and to be linked across grades, as well as between baseline and endline assessments, using common items. Scores are scaled here using Item Response theory models and standardized to have a mean of zero and standard deviation of one in the baseline.

Table 3: Treatment effect by specific competence assessed

(a) Mathematics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Dep var: Proportion of questions answered correctly</i>							
	Arithmetic computation	Word problems - computation	Data interpretation	Fractions and decimals	Geometry and Measurement	Numbers	Pattern recognition
Treatment	0.078*** (0.016)	0.071*** (0.016)	0.044** (0.020)	0.072*** (0.020)	0.14*** (0.026)	0.15*** (0.023)	0.11*** (0.029)
Baseline math score	0.13*** (0.0070)	0.11*** (0.0095)	0.080*** (0.013)	0.090*** (0.011)	0.050*** (0.014)	0.067*** (0.012)	0.094*** (0.013)
Constant	0.66*** (0.0080)	0.50*** (0.0077)	0.38*** (0.0098)	0.33*** (0.010)	0.39*** (0.013)	0.45*** (0.011)	0.36*** (0.015)
Observations	531	531	531	531	531	531	531
R-squared	0.365	0.227	0.095	0.153	0.092	0.134	0.109

(b) Hindi

	(1)	(2)	(3)	(4)
<i>Dep var: Proportion of questions answered correctly</i>				
	Sentence completion	Retrieve explicitly stated information	Make straightforward inferences	Interpret and integrate ideas and information
Treatment	0.047* (0.024)	0.046*** (0.016)	0.064*** (0.022)	0.055*** (0.016)
Baseline Hindi score	0.13*** (0.016)	0.14*** (0.0079)	0.14*** (0.011)	0.064*** (0.013)
Constant	0.73*** (0.012)	0.59*** (0.0078)	0.52*** (0.011)	0.31*** (0.0079)
Observations	533	533	533	533
R-squared	0.186	0.382	0.305	0.132

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The tables above show the impact of the treatment on specific competences. The dependent variable in each regression is the proportion of questions related to the competence that a student answered correctly. Baseline scores are IRT scores in the relevant subject from the baseline assessment. Treatment is a dummy variable indicating a randomly-assigned offer of Mindspark voucher till March 2016. All regressions include randomization strata fixed effects.

Table 4: Heterogeneity in treatment effect by sex, socio-economic status and baseline score

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dep var: Standardized IRT scores (endline)</i>						
COVARIATES	<u>Female</u>		<u>SES</u>		<u>Baseline score</u>	
	Math	Hindi	Math	Hindi	Math	Hindi
Treatment	0.43*** (0.14)	0.22** (0.10)	0.36*** (0.063)	0.24*** (0.067)	0.36*** (0.064)	0.22*** (0.076)
Covariate	-0.032 (0.15)	0.17 (0.16)	0.0095 (0.029)	0.088*** (0.020)	0.51*** (0.057)	0.67*** (0.044)
Interaction	-0.082 (0.14)	-0.0037 (0.13)	-0.0011 (0.044)	0.016 (0.042)	0.058 (0.075)	-0.0025 (0.078)
Observations	529	533	529	533	529	533
R-squared	0.393	0.453	0.393	0.472	0.393	0.451

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ Treatment is a dummy variable indicating a randomly-assigned offer of Mindspark voucher till March 2016. The SES index and test scores are defined as in Tables 1 and 2 respectively. All regressions include strata fixed effects and control for baseline subject scores.

Table 5: Heterogeneity in treatment effect by within-grade terciles

	(1)	(2)
<i>Dep var: Standardized IRT scores (endline)</i>		
VARIABLES	Math	Hindi
Bottom Tercile	0.14 (0.091)	-0.11 (0.10)
Middle Tercile	0.35*** (0.073)	0.11 (0.078)
Top Tercile	0.57*** (0.086)	0.46*** (0.079)
Treatment	0.36*** (0.11)	0.34*** (0.13)
Treatment*Middle Tercile	0.081 (0.15)	-0.21 (0.17)
Treatment*Top Tercile	-0.040 (0.16)	-0.16 (0.15)
Baseline test score	0.41*** (0.058)	0.53*** (0.061)
Observations	529	533
R-squared	0.555	0.516

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Treatment is a dummy variable indicating a randomly-assigned offer of Mindspark voucher till March 2016. Test scores are scaled as in Table 2.

Table 6: Treatment effect on items linked to grade levels

	(1)	(2)	(3)	(4)
	<i>Dep var: Proportion of questions answered correctly</i>			
VARIABLES	Math		Hindi	
	At or above grade level	Below grade level	At or above grade level	Below grade level
Treatment	0.0023 (0.039)	0.082*** (0.012)	0.069** (0.024)	0.051*** (0.013)
Baseline math score	0.044 (0.025)	0.095*** (0.0056)		
Baseline Hindi score			0.11*** (0.016)	0.13*** (0.0065)
Constant	0.31*** (0.018)	0.49*** (0.0058)	0.44*** (0.012)	0.58*** (0.0065)
Observations	286	505	287	507
R-squared	0.025	0.341	0.206	0.379

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The table shows the impact of the treatment (winning a randomly-assigned voucher) on questions below or at/above grade levels for individual students. The dependent variable is the proportion of questions that a student answered correctly. Our endline assessments, designed to be informative at students' actual levels of achievement, did not include many items at grade 8 level and above. Therefore students in Grades 8 and 9 are not included in regressions on items at/above grade level. Baseline scores are IRT scores in the relevant subject from the baseline assessment. All regressions include randomization strata fixed effects.

Table 7: Treatment effect on school exams

	(1)	(2)	(3)	(4)	(5)	(6)
	<i>Dep var: Standardized test scores</i>					
VARIABLES	Hindi	Math	Science	Social Sciences	English	Aggregate
Treatment	0.19** (0.089)	0.058 (0.076)	0.077 (0.092)	0.10 (0.11)	0.080 (0.10)	0.097 (0.080)
Baseline Hindi score	0.48*** (0.094)		0.28*** (0.064)	0.41*** (0.098)	0.29*** (0.069)	0.33*** (0.061)
Baseline math score		0.29*** (0.039)	0.10** (0.036)	0.25*** (0.052)	0.11** (0.049)	0.16*** (0.037)
Constant	0.40 (1.01)	0.14 (0.50)	0.88** (0.39)	0.69 (0.69)	1.11 (0.66)	0.68 (0.56)
Observations	595	594	593	592	595	595
R-squared	0.188	0.069	0.117	0.173	0.137	0.202

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. This table shows the effect of receiving the Mindspark voucher on the final school exams, held in March 2016 after the completion of the intervention. Treatment is a dummy variable indicating a randomly-assigned offer of Mindspark voucher till March 2016. Test scores in the school exams are normalized within school*grade to have a mean of zero and a standard deviation of one in the control group. All regressions include grade and school fixed effects.

Table 8: Dose-response of Mindspark attendance

	(1)	(2)	(3)	(4)	(5)	(6)
	<i>Dep var: Standardized IRT scores (endline)</i>					
VARIABLES	IV estimates		OLS VA (full sample)		OLS VA (Treatment group)	
	Math	Hindi	Math	Hindi	Math	Hindi
Attendance (days)	0.0065*** (0.0011)	0.0040*** (0.0011)	0.0068*** (0.00087)	0.0037*** (0.00090)	0.0075*** (0.0018)	0.0033* (0.0020)
Baseline score	0.53*** (0.036)	0.67*** (0.037)	0.54*** (0.039)	0.69*** (0.039)	0.57*** (0.062)	0.68*** (0.056)
Constant			0.35*** (0.040)	0.16*** (0.042)	0.31*** (0.12)	0.18 (0.13)
Observations	529	533	529	533	261	263
R-squared	0.422	0.460	0.413	0.468	0.413	0.429
Angrist-Pischke F-statistic for weak instrument	1238	1256				
Diff-in-Sargan statistic for exogeneity (p-value)	0.26	0.65				
Extrapolated estimates of 90 days' treatment (SD)	0.585	0.36	0.612	0.333	0.675	0.297

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ Treatment group students who were randomly-selected for the Mindspark voucher offer but who did not take up the offer have been marked as having 0% attendance, as have all students in the control group. Columns (1) and (2) instrument attendance in Mindspark with the randomized allocation of a scholarship and include randomization strata fixed effects, Columns (3) and (4) present OLS value-added models in the full sample, Columns (5) and (6) present OLS value-added models using only data on the lottery-winners. Scores are scaled here using Item Response theory models and linked across grades and across baseline and endline assessments using common anchor items. Tests in both math and Hindi are standardized to have a mean of zero and standard deviation of one in the baseline.