



Working Papers

www.cesifo.org/wp

Technology Adoption in Education: Usage, Spillovers and Student Achievement

Peter Bergman

CESIFO WORKING PAPER NO. 6101

CATEGORY 5: ECONOMICS OF EDUCATION

SEPTEMBER 2016

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

CESifo

Center for Economic Studies & Ifo Institute

Technology Adoption in Education: Usage, Spillovers and Student Achievement

Abstract

Given significant expenditures on education technologies, an important question is whether these products are adopted by their end users and are effective in practice. This paper studies the adoption, diffusion, and effects of one type of technology that is increasingly ubiquitous: school-to-parent communication technologies. Using data from a Learning Management System in several hundred schools and a two-stage experiment to study the adoption of this technology by parents, I find: A quarter of parents ever use it; adoption follows an S-shape; significant spillovers occur along intensive but not extensive margins; and there is evidence student grades improve as a result.

JEL-Codes: I200.

Keywords: education, technology adoption.

Peter Bergman
Teachers College
Columbia University
525 W. 120th Street New York
USA – New York 10027
bergman@tc.columbia.edu
Website: www.columbia.edu/~psb2101

I thank Josefa Aguirre and Eric Chan for excellent research assistance. I am particularly grateful to the Seminar for the Study of Development strategies and Abhit Bhandari and Kolby Hanson for thoughtfully replicating this paper and providing comments. I also thank George Bulman, Sue Dynarski, Jay Greene, Macartan Humphreys, Scott Imberman, Brian Jacob, Isaac McFarlin, Richard Murphy, Kevin Stange, and seminar participants at the University of Arkansas, the University of Michigan and the University of Connecticut for their comments and suggestions. All errors are my own.

I Introduction

New technologies in the public sector often aim to improve the quality of government-provided services. This is true in the education sector, where the purchase of technologies may improve curriculum delivery, data management and school-to-parent communication. A number of papers have studied the educational impacts of information technologies such as computers (Machin et al., 2007; Barrera-Osorio and Linden, 2009; Malamud and Pop-Eleches, 2011; Fairlie and Robinson, 2013; Vigdor et al., 2014; Beuermann et al., 2015), access to the Internet (Goolsbee and Guryan, 2006; Belo et al., 2013; Bulman and Fairlie, 2015; Dettling et al., 2015), computer-aided instruction (Angrist and Lavy, 2002; Rouse and Krueger, 2004; Barrow et al., 2009; Banerjee et al., 2007; Linden, 2008; Taylor, 2015), teacher dashboards (Tyler, 2013) and mobile devices (Fryer, 2013; Bergman, 2014; Castleman and Page, 2014; Beland and Murphy, 2015).

Similar to many other contexts however, the end users of education technologies may be distinct from the administrators in control of procurement. For instance, while the end users for local education agencies are often teachers, parents and students, many purchasing decisions are made at the district or school level. Given the growing private-sector investments in new education technologies, from \$600 million in 2009 to \$2.5 billion in the first half of 2015 alone, plus an additional \$11 billion spent by K-12 and higher-education institutions (Adkins, 2016; McCarthy, 2016), an important question is whether the products purchased by local education agencies are adopted by their end users and are effective in practice.

This paper studies the adoption, diffusion, and effects of one type of technology that is increasingly ubiquitous in schools: school-to-parent communication technologies. Unlike computer-aided instructional technologies, which can substitute for teacher instruction (Taylor, 2015), communication technologies can complement instruction in the classroom by informing parents about their academic progress. These technologies also have the potential to remedy the gap in communication quality that exists between low and high-achieving

schools (Bridgeland et al., 2008).

Previous research suggests school-to-parent communication can address significant information asymmetries that exist between parents and their children. These asymmetries can impede human capital investments (Akabayashi, 2006; Bergman, 2014; Bursztyrn and Coffman, 2012; Cosconati, 2009; Hao et al., 2008; Weinberg, 2001). Recent experimental evidence shows that reducing these information problems can improve student achievement, and often at low cost. For instance, Kraft and Dougherty (2013) conducted an experiment in a Boston charter school that shows daily phone calls home to parents from their child’s teachers improve student behaviors. Bergman (2014) randomized the provision of bimonthly text messages to parents detailing their child’s missing assignments and grades, which increased student effort and achievement. Kraft and Rogers (2014) show that messages from teachers to parents significantly reduced dropout from a high school credit recovery program.

Many school districts are leveraging Learning Management Systems (LMS) to improve parental access to student information at scale by placing students’ academic data onto a “parent portal” for parents to view online. This technology allows parents to view performance indicators such as their child’s grades, attendance and missing assignments in real time as teachers update it. Figure 1 shows an example of the parent portal studied in this paper. Parents are provided a website address, a user name and a password either by teachers or the school. Once a parent logs in they see their child’s classes, teachers and the associated grades. Parents may also receive and respond to messages directly from teachers as well. Figure 2 displays the screen a parent sees once they click on a specific class their child is taking. Parents can then view their child’s assignments, assignment scores, the grading scale and scoring codes. Students may also view this same information through their own separate account.

However, these systems are typically purchased at the school or district level, and the adoption, usage and effects of this technology are unknown. As opposed to the experimental evidence on school-to-parent communication described above, which *pushes* information out

to families via text messages and phone calls, this parent-portal technology requires much more of a *pull*. There are several potential barriers to adoption and usage: parents must have internet access, be aware the system exists, keep track of their user name and password, and remember to log in. Like many school-to-parent communication systems, parent user names and passwords must be downloaded from the LMS and distributed to parents. This distribution can occur by mail, email, or at school events.

To examine adoption, diffusion, and efficacy, I use data from a learning management company operating in 15 school districts as well as a two-stage experiment providing families their account information in 59 schools across three districts. This experimental design is similar to that used by Duflo and Saez (2003) to study the role of social interactions in retirement plan decisions: First, schools are randomized to either have a sample of families treated or to have no families treated. Second, families within treated schools are randomly selected to actually receive the intervention. This design permits analysis of the intervention along extensive (whether a parent ever used the portal) and intensive margins (how often the parent used the portal) as well as potential peer effects.

In general, the influence of peers on individuals' behaviors is difficult to estimate due to the reflection problem (Manski, 1993). A number of papers show how peer influence can either encourage or discourage the adoption of health and agricultural-related technologies, particularly in lower-income countries (Foster and Rosenzweig, 1995; Kremer and Miguel, 2007; Conley and Udry, 2010; Foster and Rosenzweig, 2010; Duflo et al., 2011; Oster and Thornton, 2012; Dupas, 2014). Several other papers find that low-cost information interventions can “nudge” the adoption of new behaviors in a variety of contexts (Cialdini et al., 2006; Goldstein et al., 2007; Allcott, 2011; Allcott and Rogers, 2014). This paper contributes to this literature by studying how peer effects and could affect the adoption of an education-related technology in the United States.

I find that parent adoption of this technology follows an S-shape curve over the course of the school year that rises quickly then levels off. Usage is far from complete. Across several

hundred schools, 24% of parents have ever logged into the system by the end of the year and roughly 4% of parents log in at least once per week. School-level adoption rates positively correlate with measures of family income, school-level test scores and teacher usage. At the individual level, families with higher-achieving students adopt. These patterns suggest that this technology, without intervention, may not address the disparities in student achievement or school-to-parent communication that exist across income and performance groups.

The experimental intervention increased adoption and usage among treated families relative to families in schools where no one received the intervention. As parents in the spillover group did not receive their account information directly via the intervention, qualitative interviews with district administrators and parents suggested they used their child’s student account to log in. In the data, there is a significant effect on usage and adoption of the student accounts in the spillover group that is nearly as large as the effects on parent usage for the treated group. Total usage—usage by both parents and students—is nearly equivalent in the treated and spillover groups.

Lastly, access to the portal system modestly improved student grades. For both the treatment and spillover groups, GPA improved by roughly 0.10 points. Overall, the results suggest this technology is capable of a modest improvement in student outcomes and adoption is not widespread without significant intervention.

The rest of this paper is organized as follows. Section II describes the data and patterns of usage. Sections III describes the data and empirical strategy for the experimental intervention. Section IV presents the results and Section V concludes and provides a basic cost analysis.

II Data and Descriptive Results

This study draws data from several sources. The first is deidentified data from a Learning Management System (LMS) company for the 2012-2013 and 2013-2014 school years. This

LMS provider hosts a parent portal, a teacher gradebook, and a student portal. The student portal shows the same academic information to students as the parent portal shows to parents, but the user name and password are distinct from the parent user name and password.

The LMS records parent, student and teacher logins into each of these services by date. During the 2012-2013 school year, there are more than 25,000,000 login-by-day observations across all students and more than 3,000,000 logins-by-week observations across 149,107 students. The LMS also records student grades by marking period and course. Students in elementary school do not receive letter grades, so these marks are excluded from the analysis sample (9.75% of marks).

While the data have the unique aspect of recording portal usage and student grades the data have several limitations as well. First, the LMS data only have a single demographic variable that is recorded across all schools, which is student gender. Second, grade levels for students are missing. Third, there are no standardized test scores in the data. However, GPA is a stronger predictor of college performance than SAT or ACT scores, even unadjusted for high school quality (Rothstein, 2004; Bowen et al., 2009; Hiss and Franks, 2014; Scott-Clayton et al., 2014).

I supplement the LMS data with information from the NCES Common Core Data, which records school-level characteristics for the universe of public schools in the United States. These data describe, at the school level, demographic shares by race, receipt of free/reduced-price lunch, as well as Title I status and location in an urban, suburban, town or rural location.

Lastly, to obtain a unified measure of school performance across school districts, I draw on the decile performance ratings constructed by GreatSchools, a nonprofit organization. In most settings, GreatSchools formulates these ratings by calculating the average share of students who are proficient in math and English per grade and averaging these shares across the grades a school offers. GreatSchools then uses this measure to assign schools their state-

wide decile. Thus if a school receives a rating of 10, that school is in the top-ten percent of the state according to this measure of proficiency. This variable is only used as a covariate.

Table 1 presents summary statistics of the data used to describe parent portal usage. There are 264 schools across 15 school districts. These schools enroll 149,107 students. On average, schools are 78% white, 15% Black, and 4% Hispanic. The majority (55%) receive free or reduced-price lunch. The plurality of the sample is rural (41%) with the remaining sample primarily urban and suburban. While this geographic balance is not representative of the nation, it nonetheless has significant enough variation to find informative correlates of portal adoption and usage across a variety of contexts.

The vast majority of parents have never logged into their parent portal accounts. Table 2 uses data from the LMS to describe basic usage patterns. During the 2012-2013 school year, the share of families who had ever logged into the system was 24%. Overall, 2% of families log in once per day and parents log in .13 times per week for a total of 13 times over the year on average. Figures 3 and 4 show the distribution of total usage for all parents and parents who have logged in at least once, respectively. The latter subgroup is important because it defines those parents who likely knew their account information at one point.

Figure 5 traces out the adoption curve—the share of parents who have ever used the parent portal—by date over the course of 2012-2013 school year. Adoption takes on an “S” shape, similar to that found in the adoption of other types of products and technologies (Rogers, 2010). There is a sharp rise at the start of the school year, but by late November the curve levels off. The share of parents who have ever logged into the system reaches just under 25% by the end of the school year. This level of adoption is not necessarily unique to this system. New York City Department of Education officials stated that many parents had never logged on to their now-defunct ARIS parent portal system. Internal analyses provided to me by a large California network of charter schools also showed similar rates of usage across.

Adoption also correlates with measures of income and test scores. Figure 6 shows a

negative correlation between the share of students receiving free or reduced price lunch and the share of parents who have ever logged in. Figure 7 uses the decile-proficiency measure to chart the relationship between test scores and the share of parents who have ever logged in. For the highest-performing ten percent of schools, roughly half of parents have ever logged into the system. For the lowest-performing ten percent of schools, less than ten percent of parents have ever logged into the system.

To study how usage correlates with achievement at the individual level, I estimate the following regression model:

$$Percent_i = \alpha + \sum_{k=1}^K \beta_k * 1[logins \in [a_k, b_k)] + \varepsilon_i$$

In which $Percent_i$ is the average percent grade of student i . β_j are coefficients on indicator variables for whether a parent has logged in between a_k and b_k times, where the latter take on values such as one to four times or five to ten times. Zero logins is the omitted category.

Figure 8 plots the β coefficients of this regression. This graph shows the percentage-point gain in student grades associated with different levels of portal usage relative to the average percent grade of students whose parents have never logged into the system. For example, the first point on the graph shows that parents who logged in between one and four times had students who score four percentage points higher than students whose parents never logged in. The gradient is remarkably flat, with all associated gains in performance occurring between those whose parents never log in and those whose parents have logged in at least once; further usage is not associated with better or worse student performance.

To study the correlates of adoption rates at the school level, I estimate the following:

$$ShareAdopted_s = \gamma + X_s' \theta + \varepsilon_s$$

The dependent variable is the share of parents who have ever logged in at school s . The independent variables, X_s , also measured at the school level, are indicators for whether a

school is a middle or high school, Title I status, urban, rural or suburban location, as well as variables for share Hispanic, Black, free and reduced-price lunch recipients. Average student-to-teacher ratio and total teacher logins at school s are included as well.

Table 3 presents the results of this regression for the year 2012-2013. The share Hispanic at a given school negatively correlates with adoption, possibly reflecting language barriers, though the portal can present information in Spanish. Interestingly, adoption at the high school level is lowest relative to middle and elementary school students' parents. Though cross sectional, this disparity is in line with other cross-sectional measures of parental monitoring, such as parent teacher conference attendance, which drops sharply from middle to high school (Noel et al., 2013).

The final row of Table 3 measures how the supply of information correlates with demand. The logins-per-teacher variable equals the total teacher logins to the LMS at a given school divided by the number of teachers at the school. This measure of how often teachers use the gradebook positively correlates with parent adoption of the system.¹ Higher student-to-teacher ratios, which may make it more difficult to keep grade information up to date, negatively correlates with adoption. These results highlight how the supply and demand for information are likely determined simultaneously, and the difficulty of recovering the causal effects of the technology on student outcomes. The experiment discussed below identifies the effects of adoption, spillovers and achievement impacts of this technology through an encouragement design leveraging the fact that 75% of families have never logged in to the system.

¹Similar measures of supply, such as the average number of teacher logins per student, also positively correlate with parent adoption.

III Experimental Design and Implementation

Experimental Design

The experimental intervention consisted of a mailer and a phone call. The mailer informed families about the parent portal, that they will be called regarding the parent portal service, and provided the school phone number so parents can obtain their account information directly from the school. The subsequent phone call to parents told families their user name, password and the website URL for the parent portal.

The sample frame for the intervention was comprised of three districts operating 59 elementary, middle and high schools across two states. Within these districts, the sample was restricted to parents who had ever logged into the system five times or less. The latter restriction aims to target the intervention to low-usage parents while retaining 82% of all students' parents.

Figure 9 describes the treatment allocation. The assignment of the intervention was randomized in two stages. First, 29 schools were randomly selected to have a sample of families receive the intervention. The remaining 30 schools had access to the parent portal, but no parent received any form of the intervention by the researchers. Within the 29 selected schools, just under half of the parents in the sample frame were selected to receive the intervention. This allocation mechanism formed a treated group, who was assigned to receive a phone call and a mailer, a spillover group, who was in the same schools as the treated families but did not receive either a mailer or a phone call, and a control group, who attended schools in which no one was treated. School-level treatment assignment was stratified according to indicators for whether more than 25% of families had logged in at baseline, more than 50% of students had received free or reduced-price lunch, and indicators for each school's district. Importantly, all families and teachers were blinded to the study and the intervention was a form of district outreach to parents.

Data and Implementation

Data for this experiment are similar to the data studied above but with a few additions. As above, these data consist of login information by date for parents, teachers and students; student course grades; demographic information from the NCES Common Core; and GreatSchools school quality ratings. Students' GPA is standardized according to the untreated schools' mean and standard deviation.

The additional data come from the phone intervention. Phone agents captured call response rates, whether or not parents or guardians had internet access, and whether or not parents were willing to provide their cell phone and email addresses to use for a future parent-school information technology system. Common Core data could be merged for 58 of 59 schools in the sample. GreatSchools school quality ratings could be merged for 54 of the 59 schools.

5,027 students' parents were assigned to the treatment group. Mailers notifying parents about the parent portal, how to obtain their account information, and the impending phone call were sent to arrive at the start of November 2013. A phone bank contacted families over the course of the second week of November, 2013. Phone contact was made with 61% of students' parents. Of these parents, 11% said they already had their account information, which may have been caused by the mailer or previous usage, and nearly all remaining families took down their account information.

Empirical Strategy

The random assignment of the phone and the mailer intervention across schools, and subsequently across individuals, means that families in the treatment, spillover and control groups have similar potential outcomes with respect to the treatments. By comparing outcomes between each group, we can estimate the impacts on the treatment and spillover groups. I estimate intent-to-treat impacts as follows:

$$Y_{is} = \beta_0 + \beta_1 \text{Treatschool}_{is} + \beta_2 \text{Treatschool}_{is} \times \text{Untreated}_{is} + X'_{is} \Gamma + \eta_{is} \quad (1)$$

Outcomes y_{is} are login and academic outcomes at the individual level for students in school s . The Treatschool_{is} variable indicates whether a student is in a school in which anyone receives the treatment. The Untreated_{is} variable indicates a student who was not assigned to the intervention, though the individual may have been in a treated school. This specification implies that the β_1 coefficient is the effect of the intervention on those families who were selected to receive the treatment. The coefficient on the interaction term, β_2 , therefore estimates the differential impact on the spillover group—those who were in schools with families selected for treatment. The test of significance for this coefficient provides evidence whether we can reject that the spillover group experienced similar effects to the treated group. The X_{is} term is a vector with school and individual-level controls as well as strata indicators. I impute missing values and include indicators for missing data for any schools or students lacking such data. All standard errors are clustered at the school level.

Random assignment also implies background characteristics should be comparable across groups in expectation. Table 4 shows the covariate balance across the three groups, respectively. The average GPA in the sample is 2.5, students miss 8% of their assignments, on average, and average total parent logins from the start of the school year until the second week of October is 0.6.

Student logins are much higher however. Between the start of the school year and the second week of October, students logged in an average of 23 times. The schools are 63% white, 30% Black, and 3% Hispanic. 60% of students receive free or reduced price lunch. At the individual level and the school level there are no significant differences between the treatment, spillover and control groups. The number of schools is small relative to the number of observations however, and results will be shown with and without controls.

One school is a significant outlier relative to the other schools. Figure A.1 shows endline

usage plotted against Greatschools' test-score proficiency rating. One school is far above the regression line. As will be shown, results *are* sensitive to the inclusion of this particular school. Therefore all results will be shown with and without this outlier. Once controls are added to the regressions, the results are no longer sensitive to the inclusion or exclusion of this school. Qualitatively, district administrators said this school ran an extensive campaign encouraging parents to log in. As a robustness check, I show how the results are affected by excluding each school in the sample one by one, as well as excluding the outlier school plus excluding a second school one by one, both of which demonstrate that the outlier severely skews the results relative to any other school excluded from the sample.

Lastly, differential attrition across treatment, spillover and control groups could bias estimates of treatment effects. The login data does not indicate whether a student has left a participating district, but observing no final grades is an indicator of district attrition. Table A.1 tests for differential attrition across treatment and spillover groups by estimating equation (1), without controls, on an indicator for whether or not a student has a final grade. For both analyses with and without outliers, there is no evidence of differential attrition from the sample.

IV Adoption, Spillovers and Efficacy

Extensive and Intensive Margin Effects

Figure 10 plots the treatment effects for the share of parents who have ever logged in by week. This figure shows extensive-margin effects. The vertical red line in the figure indicates when the phone treatment occurred. The treatment effect on ever logging in rises sharply during this period—roughly two-percentage points above a six percent mean in the control schools. The effect persists through the remainder of the school year.

This increase along the extensive margin contrasts sharply with the effects for the spillover group. Figure 11 shows the treatment effects on the share who ever log in for the spillover

group. The spillover group exhibits slightly lower rate of adoption relative to the control group at the outset, though not nearly significantly so, and effects remain around zero through the duration of the year.

Table 5 presents the regression results. Again, the *Treatschool* variable indicates whether a school was treated and the interaction term indicates the differential impact for the spillover group. The effect on the spillover group is the *Treatschool* coefficient minus the interaction-term coefficient. The significance or not of the interaction term tests whether the differential effect is statistically significant. The first column shows the effects with no control variables (except strata indicators) and the control mean. This impact is positive but not significant and the spillover coefficient is negative but not significantly different from the treatment group. The second column adds controls for race and free-reduced price lunch shares as well as school ratings. The coefficient on the treated students is larger and slightly more precise: a 2.4 percentage point increase. The spillover group exhibits a significantly smaller effect—essentially zero impact on adoption rates. Column three excludes the outlier and control variables and shows a similar coefficient to estimates with entire sample and controls. Finally, column four excludes the outlier school and adds control variables. Once the outlier is excluded the *Treatschool* coefficient is not sensitive to the addition of controls.

As mentioned previously, students may also log in to view their grades and assignments through a separate account, user name and password. Qualitative interviews with parents and district administrators suggested that parents, not knowing their own account information since it was not provided to them via the treatment, asked their child to log in for them. Table 6 shows evidence that this is the case. The dependent variable in column one is an indicator variable for any student use and zero parent usage. Relative to the treatment group, the spillover group is significantly more likely to have student usage with no parent usage. Column two shows that any parent usage with no student usage is significantly less likely for the spillover group, which is of a similar magnitude as the increase in student-only usage. For completeness column three examines the effects on an indicator for parent and

student usage only, which are roughly in line with the previous results on adoption. Finally, the dependent variable in column four is an indicator for any usage by either parents *or* students. Overall, there is a seven percentage point increase for the treatment group and a 6 percentage point increase for the spillover group that is statistically indistinguishable from the treatment group's effect. The control mean is also quite high—68%. Though the test is not shown, the effect on any usage for the spillover group is statistically significant from zero.

While the extensive margin is important, the intensive margin of usage may be equally or more important for fostering student achievement, especially given the high control mean for any usage by parents or students. Figure 12 plots the treatment effect on the number of logins per month relative to the control group. Again, the effects increase sharply and significantly during the month of the intervention, dip slightly during the middle of the second semester, and rise slightly by the end of term.

Figure 13 presents the same graph for the spillover group relative to the control group. The pattern of treatment effects is slightly different. There is a gradual rise in logins in the two months after the start of the intervention rather than the sharp and immediate rise for the treatment group. This more gradual rise could reflect a diffusion of information about the intervention followed by increased parental usage.

Table 7 shows regression analysis for the treatment and spillover impacts on total parent logins. Column one, which includes all schools and no control variables, shows a negative, not significant treatment effect, and a small, positive spillover effect. Adding the control variables in column two, this effect increases to a 0.6 points for the treatment group over the 2.78 logins by the control group, which is statistically significant at the 1% level. Usage by the spillover group is not statistically different from the treatment group and is significantly different from zero (test not shown), which indicates significant spillovers in terms of parental usage. Columns three and four remove the outlier. Again, the results are much more stable with and without controls and are similar to the effects when all schools and baseline covariates

are included. This check provides evidence that the outlier is skewing results when there are no baseline controls.

Table 8 shows student logins also increase along intensive margins for the treatment and spillover groups, though the effects are significant at the 10% level. The increases for the treatment group and the spillover group are statistically indistinguishable from each other. As described above and related by principals, this increase may be the result of students monitoring their progress themselves or as a result of parents asking their children to log in for them. The results for students and parents imply the total impact on logins by either parents or students is just above 6 logins post treatment.

The results on both adoption and usage have implications for interpreting any kind of treatment-on-the-treated effect. Viewing the effects as only operating through parent adoption clearly violates the exclusion restriction, as their effects on both adoption and usage for students and parents. If we consider parent or student adoption, a broad base of users have logged previously, and so a substantial amount of the effect may be due to additional logins by parents or their children—roughly equivalent to one additional login per month.

Robustness to Outliers

To analyze how much this outlier skews the effects, I re-rerun the regressions for parent logins 59 times. Each of these 59 regressions excludes a different school from the analysis. I then de-mean each treatment impact using a leave-one-out-mean so that the treatment effects center around zero. In Figure A.2 I plot the change in treatment effect relative to this mean for each regression. Only two schools out of the 59 exert significant changes on the magnitudes of the treatment effects. The most significant of these is the outlier school plotted previously. As shown in the second robustness check below, the inclusion or exclusion of any other school has no effect on the results.

As a second robustness check, and to affirm that the exclusion of any other school does not significantly alter the results, I exclude the previously identified outlier and then re-rerun

the regressions excluding every *other* school, one by one. These effects are not de-meanded in order to show how the magnitudes and significance change. Figure A.3 shows these results. Nearly all treatment impacts are around 0.60 and no treatment impact is below 0.40. All effects are significant at the 5% level.

Student Achievement Effects

Finally, this section examines the impact of the parent portal intervention described in Section III on student GPA. Table 9 presents the results. Overall the effect size is roughly 0.10 standard deviations. These results are significant with the addition of controls, which reduce the standard errors substantially. The magnitudes of the coefficients is less sensitive to the inclusion or exclusion of the outlier school.

Interestingly, the effect on student grades is not differential by treatment or spillover group. This result is consistent with the combined student and parent usage patterns along extensive and intensive margins, which are similar for both treatment and spillover groups. The effect size is roughly half of the effect size found in Bergman (2014), in which information was actively pushed to parents about their child’s academic performance.

Table 10 explores whether the effects on GPA vary by subgroup. For ease of presentation the analysis is conducted with a school-level treatment indicator, which combines treated and spillover groups. There are no differences in heterogeneity between the spillover and treatment groups (results available on request). The results show there are no differential effects by baseline GPA, gender, or school-level demographic and performance characteristics.

Heterogeneity does occur appear to occur according to measures of baseline usage. Parents who used the system more at baseline saw smaller effects, though this result becomes marginally insignificant when outliers are excluded from the analysis (results available on request). The remaining results are robust: higher levels of student usage is associated with larger effects and students whose teachers use the system more frequently also experience larger gains in GPA. A half-standard deviation increase in student usage leads to .02 stan-

dard deviation gain in GPA and a half-standard deviation increase in the average logins by a student's teachers leads to .10 standard deviation increase in GPA. A half-standard deviation increase in parent usage reduces effects by .01 standard deviations. These results highlight the apparent complementarity between parent usage and teacher usage of the portal as discussed further below.

V Discussion and Conclusion

Previous research has shown that school-to-parent communication can improve parental monitoring and a range of student outcomes. This paper documents some of the first evidence on parents' adoption of a school communication technology that aims to scale school-to-parent communication: parent information portals. Adoption is far from universal; three quarters of parents have never logged into the system. Logins through student accounts are much higher however. Schools with higher login rates tend to be higher income and higher performing, which suggests that this technology may not close achievement gaps without active efforts to promote adoption and usage.

A simple intervention providing account information to parents increased adoption and usage by roughly one login per month. Interestingly, there were significant usage spillovers on families who did not receive the intervention. This increase in usage led to a modest increase in grades in both treated and spillover group students. Though these gains are small, the intervention has low marginal cost as well. The mailers cost \$0.70 to print and send across two states. The phone calls cost \$1.36 per student to manage and implement. These results are evidence that a simple nudge can modestly promote adoption, usage, and student achievement.

The results also emphasize the complementarity between parent usage and teacher usage of the portal. Both the usage and the GPA treatment effects are larger for schools in which teachers used the system more frequently. One might hope that the intervention could

generate a demand shock for information sufficient enough to increase the supply information as proxied by teacher logins, but the study is underpowered to detect such effects.

Overall, these results indicate both the promise and pitfalls of these technologies. Merely providing access to information online may not improve outcomes in low-income area schools and low-performing schools. Given the potential importance of this information and the barriers to online access, future research could examine the take up and efficacy of information technologies aimed at actively pushing information to parents at scale.

References

- Adkins, Sam**, “2015 International Learning Technology Investment Patterns,” Technical Report, Ambient Insight 2016.
- Akabayashi, Hideo**, “An equilibrium model of child maltreatment,” *Journal of Economic Dynamics and Control*, 2006, 30 (6), 993–1025.
- Allcott, Hunt**, “Social norms and energy conservation,” *Journal of Public Economics*, 2011, 95 (9), 1082–1095.
- **and Todd Rogers**, “The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation,” *American Economic Review*, 2014, 104 (10), 3003–37.
- Angrist, Joshua and Victor Lavy**, “New Evidence on Classroom Computers and Pupil Learning,” *The Economic Journal*, 2002, 112 (482), 735–765.
- 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,” 2009.
- 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.
- Beland, Louis-Philippe and Richard Murphy**, “Ill Communication: Technology, Distraction & Student Performance,” Technical Report 2015.
- Belo, Rodrigo, Pedro Ferreira, and Rahul Telang**, “Broadband in school: Impact on student performance,” *Management Science*, 2013, 60 (2), 265–282.

- Bergman, Peter**, “Parent-Child Information Frictions and Human Capital Investment: Evidence from a Field Experiment,” *Columbia University Teachers College Working Paper*, 2014.
- 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.
- Bowen, William G, Matthew M Chingos, and Michael S McPherson**, *Crossing the finish line: Completing college at America’s public universities*, Princeton University Press, 2009.
- Bridgeland, John M, John J DiIulio, Ryan T Streeter, and James R Mason**, “One Dream, Two Realities: Perspectives of Parents on America’s High Schools.,” *Civic Enterprises*, 2008.
- Bulman, George and Robert W Fairlie**, “Technology and Education: Computers, Software, and the Internet,” 2015.
- Bursztyn, Leonardo and Lucas C Coffman**, “The schooling decision: Family preferences, intergenerational conflict, and moral hazard in the Brazilian favelas,” *Journal of Political Economy*, 2012, 120 (3), 359–397.
- Castleman, Benjamin L and Lindsay C Page**, “Summer nudging: Can personalized text messages and peer mentor outreach increase college going among low-income high school graduates?,” *Journal of Economic Behavior & Organization*, 2014.
- Cialdini, Robert B, Linda J Demaine, Brad J Sagarin, Daniel W Barrett, Kelton Rhoads, and Patricia L Winter**, “Managing social norms for persuasive impact,” *Social influence*, 2006, 1 (1), 3–15.

- Conley, Timothy G and Christopher R Udry**, “Learning about a new technology: Pineapple in Ghana,” *The American Economic Review*, 2010, pp. 35–69.
- Cosconati, Marco**, “Parenting style and the development of human capital in children,” *Job Market Paper, University of Pennsylvania*, 2009.
- Dettling, Lisa J, Sarena Goodman, and Jonathan Smith**, “Every Little Bit Counts: The Impact of High-speed Internet on the Transition to College,” 2015.
- Duflo, Esther and Emmanuel Saez**, “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment*,” *The Quarterly journal of economics*, 2003, 118 (3), 815–842.
- , **Michael Kremer, and Jonathan Robinson**, “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya,” *American Economic Review*, 2011, 101, 2350–2390.
- Dupas, Pascaline**, “Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment,” *Econometrica*, 2014, 82 (1), 197.
- Fairlie, Robert W. and Jonathan Robinson**, “Experimental Evidence on the Effects of Home Computers on Academic Achievement among Schoolchildren,” *American Economic Journal: Applied Economics*, July 2013, 5 (3), 211–40.
- Foster, Andrew D and Mark R Rosenzweig**, “Learning by doing and learning from others: Human capital and technical change in agriculture,” *Journal of political Economy*, 1995, pp. 1176–1209.
- and – , “Microeconomics of technology adoption,” *Annual Review of Economics*, 2010, 2.
- Fryer, Roland G**, “Information and student achievement: evidence from a cellular phone experiment,” Technical Report, National Bureau of Economic Research 2013.

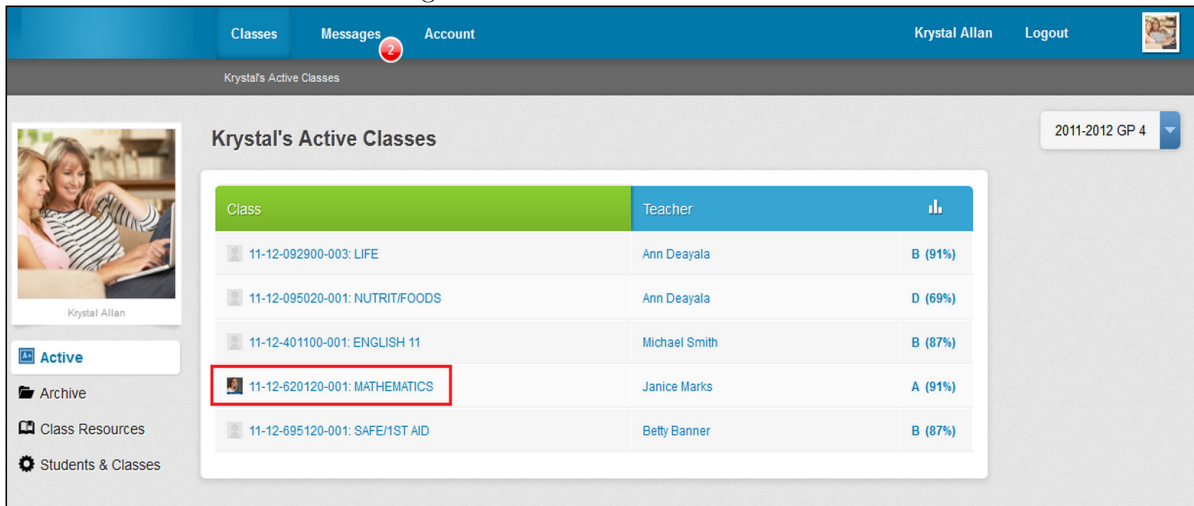
- Goldstein, Noah J, Vldas Griskevicius, and Robert B Cialdini**, “Invoking Social Norms A Social Psychology Perspective on Improving Hotels’ Linen-Reuse Programs,” *Cornell Hotel and Restaurant Administration Quarterly*, 2007, 48 (2), 145–150.
- Goolsbee, Austan and Jonathan Guryan**, “The impact of Internet subsidies in public schools,” *The Review of Economics and Statistics*, 2006, 88 (2), 336–347.
- Hao, Lingxin, V Joseph Hotz, and Ginger Z Jin**, “Games Parents and Adolescents Play: Risky Behaviour, Parental Reputation and Strategic Transfers*,” *The Economic Journal*, 2008, 118 (528), 515–555.
- Hiss, William C and Valerie W Franks**, “Defining promise: Optional standardized testing policies in American college and university admissions,” *Report of the National Association for College Admission Counseling (NACAC)*. <http://www.nacacnet.org/research/research-data/nacac-research/Documents/DefiningPromise.pdf>, 2014.
- Kraft, Matthew A and Shaun M Dougherty**, “The effect of teacher–family communication on student engagement: Evidence from a randomized field experiment,” *Journal of Research on Educational Effectiveness*, 2013, 6 (3), 199–222.
- Kraft, Matthew and Todd Rogers**, “The Underutilized Potential of Teacher-to-Parent Communication: Evidence from a Field Experiment,” 2014.
- Kremer, Michael and Edward Miguel**, “The Illusion of Sustainability*,” *The Quarterly journal of economics*, 2007, 122 (3), 1007–1065.
- Linden, Leigh L**, “Complement or Substitute? The Effect of Technology on Student Achievement in India,” 2008.
- 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 Cristian Pop-Eleches**, “Home Computer Use and the Development of Human Capital,” *The Quarterly journal of economics*, 2011, *126* (2), 987–1027.
- Manski, Charles F**, “Identification of endogenous social effects: The reflection problem,” *The review of economic studies*, 1993, *60* (3), 531–542.
- McCarthy, Shawn**, “Pivot Table: U.S. Education IT Spending Guide, Version 1, 20132018,” Technical Report, International Data Corporation 2016.
- Noel, Amber, Patrick Stark, and Jeremy Redford**, “Parent and Family Involvement in Education, from the National Household Education Surveys Program of 2012. First Look. NCES 2013-028.,” *National Center for Education Statistics*, 2013.
- Oster, Emily and Rebecca Thornton**, “Determinants Of Technology Adoption: Peer Effects In Menstrual Cup Take-Up,” *Journal of the European Economic Association*, 2012, *10* (6), 1263–1293.
- Rogers, Everett M**, *Diffusion of innovations*, Simon and Schuster, 2010.
- Rothstein, Jesse M**, “College performance predictions and the SAT,” *Journal of Econometrics*, 2004, *121* (1), 297–317.
- 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.
- Scott-Clayton, Judith, Peter M Crosta, and Clive R Belfield**, “Improving the Targeting of Treatment Evidence From College Remediation,” *Educational Evaluation and Policy Analysis*, 2014, *36* (3), 371–393.
- Taylor, Eric**, “New Technology and Teacher Productivity,” Technical Report 2015.
- Tyler, John H**, “If You Build It Will They Come? Teachers’ Online Use of Student Performance Data.,” *Education Finance and Policy*, 2013, *8* (2), 168–207.

Vigdor, Jacob L, Helen F Ladd, and Erika Martinez, “Scaling The Digital Divide: Home Computer Technology And Student Achievement,” *Economic Inquiry*, 2014, 52 (3), 1103–1119.

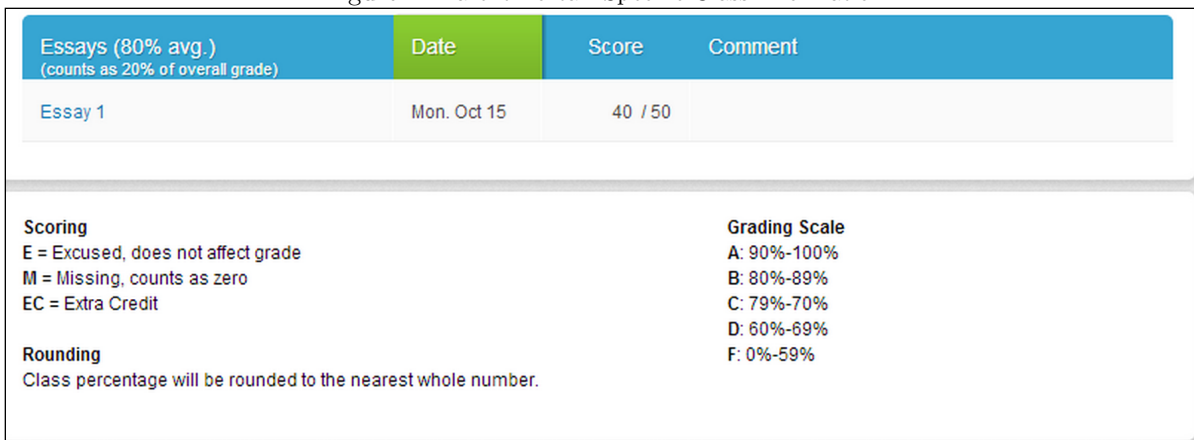
Weinberg, Bruce A, “An Incentive Model of the Effect of Parental Income on Children,” *Journal of Political Economy*, 2001, 109 (2), 266–280.

Figure 1: Parent Portal: Main Screen



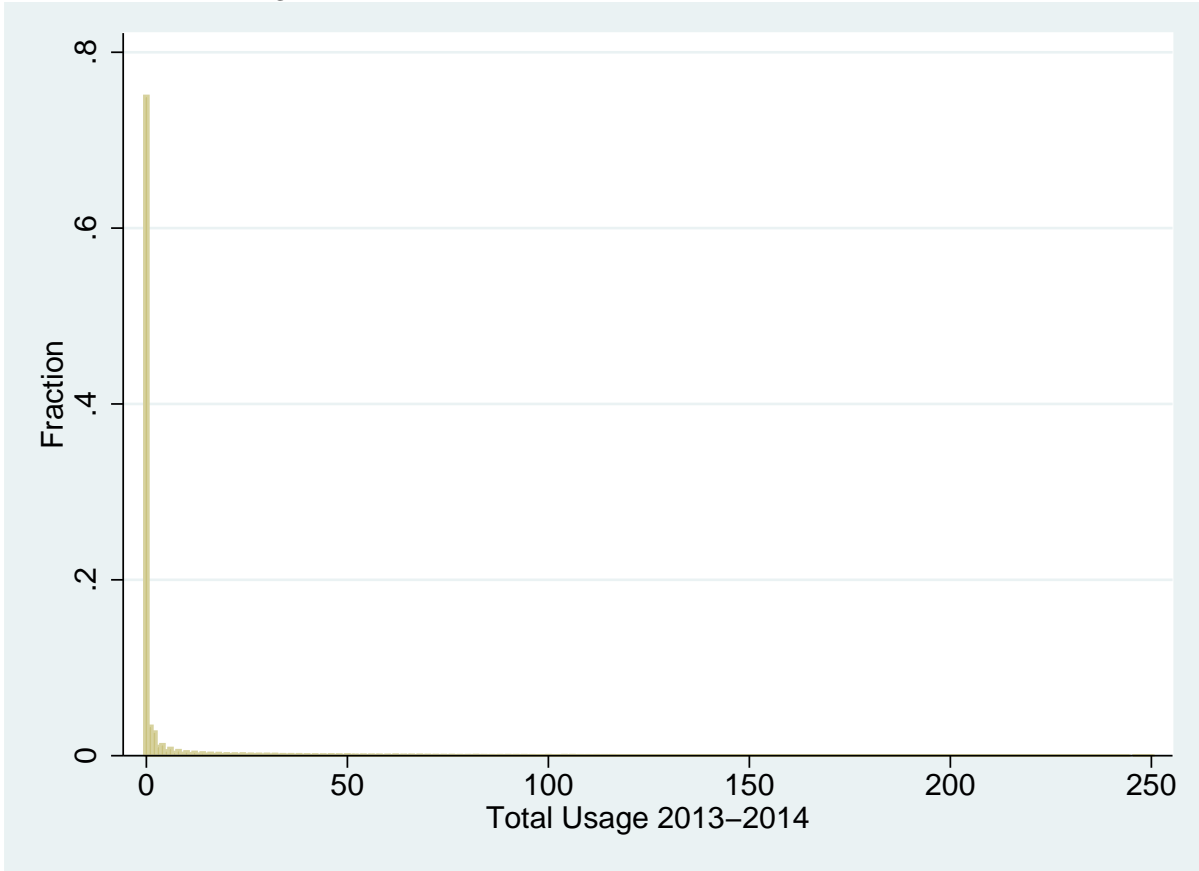
The figure shows an example of the type of academic information that can be found on parent portal. All information on this figure is fictional.

Figure 2: Parent Portal: Specific Class Information



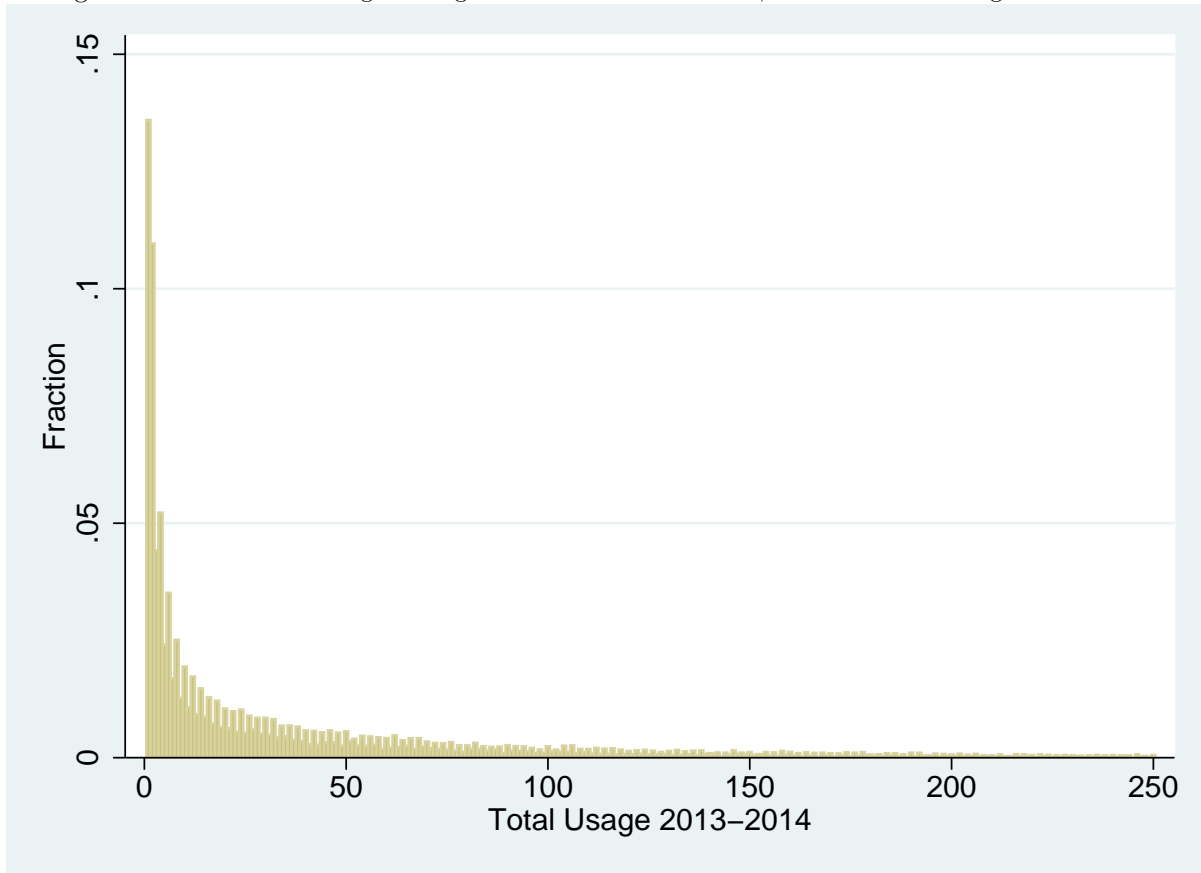
The figure shows an example of the type of academic information that can be found on parent portal once a parent clicks on a specific class. All information on this figure is fictional.

Figure 3: Parent Portal Usage During the 2012-2013 School Year




The figure shows the distribution of portal logins during the 2012-2013 school year. This figure is constructed using data from the Learning Management System and trims the top-most percentile from the data.

Figure 4: Parent Portal Usage During the 2012-2013 School Year, Conditional on Using at Least Once




The figure shows the distribution of portal logins during the 2012-2013 school year conditional on logging in at least once. This figure is constructed using data from the Learning Management System and trims the top-most percentile from the data.

Figure 5: Parent Portal Adoption During the 2012-2013 School Year



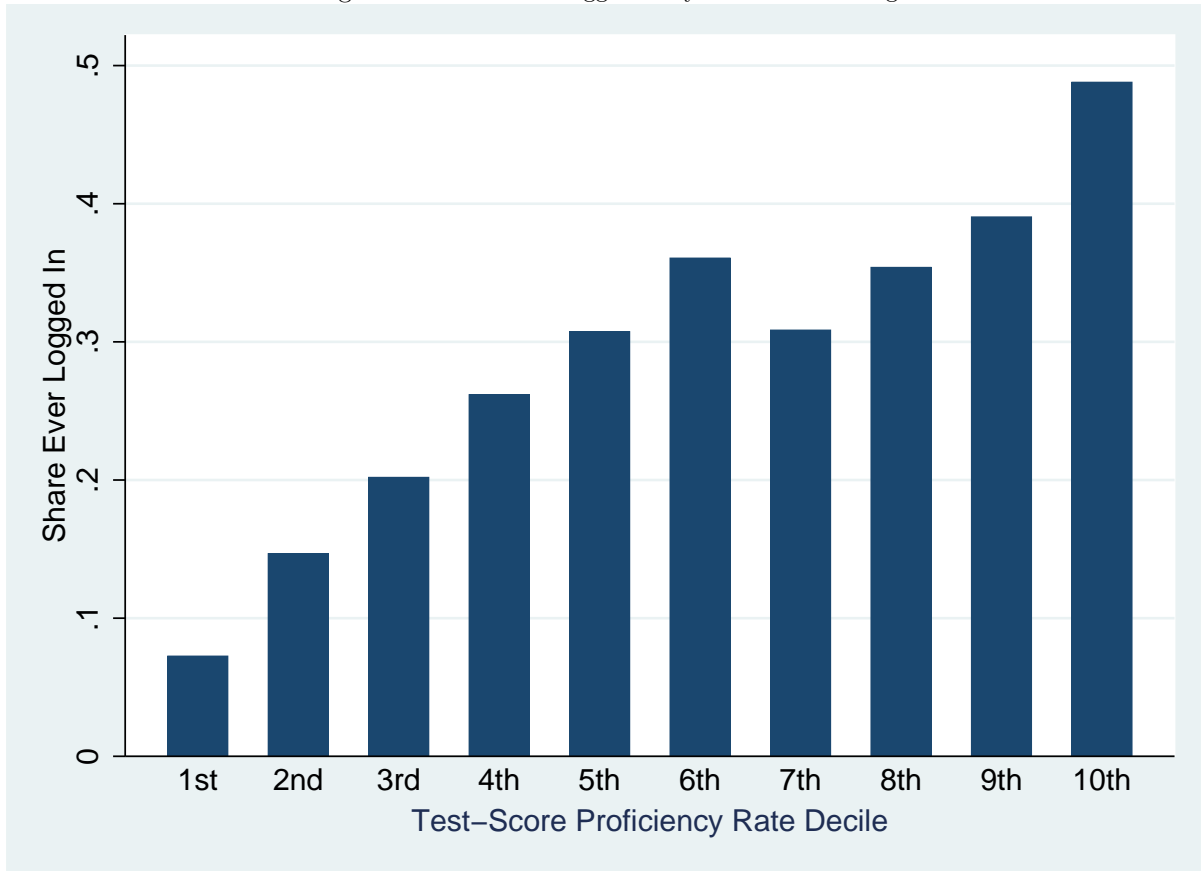
The figure shows the share of parents who have ever logged into the parent portal during the 2012-2013 school year. This figure is constructed using data from the Learning Management System.

Figure 6: Share Ever Logged In by Share Free/Reduced Price Lunch




The figure shows the share of parents who have ever logged into the parent portal plotted against the share of students who receive free/reduced price lunch in each school. This figure is constructed using data from the Learning Management System and excludes one outlier above the 99th percentile of usage.

Figure 7: Share Ever Logged In by Test-Score Rating




The figure shows the share of parents who have ever logged into the parent portal according to the GreatSchools Rating of each school. This figure is constructed using data from the Learning Management System.

Figure 8: Share Ever Logged In by Test-Score Rating




This figure shows the percentage-point gain in student grades associated with different levels of portal usage relative to the average percent grade of students whose parents have never logged into the system. This figure is constructed using data from the Learning Management System.

Figure 9: Experimental Design




This figure shows the experimental design for the account-information intervention. Randomization occurs first at the school level and then at the student level.

Figure 10: Adoption: Treatment v. Control




This figure shows the treatment effect on the share of parents who ever logged into the portal over the course of the school year. Data come from the LMS company.

Figure 11: Adoption: Treatment v. Control




This figure shows the spillover effect on the share of parents who ever logged into the portal over the course of the school year. Data come from the LMS company.

Figure 12: Usage: Treatment v. Control



This figure shows the treatment effect on the number of times parents logged in per month over the course of the school year. Data come from the LMS company.

Figure 13: Usage: Spillover v. Control



This figure shows the spillover effect on the number of times parents logged in per month over the course of the school year. Data come from the LMS company.

Table 1: District Summary Statistics

Variable	Mean	Observations
Districts	N/A	15
Schools	N/A	264
Students	N/A	149,107
Female	49%	149,107
Share Hispanic	5.2%	244
Share Black	16.2%	244
Share White	77.5%	244
Share Free/Reduced Lunch	54.5%	244
Urban	21.5%	244
Suburb	20.7%	244
Town	15.1%	244
Rural	42.6%	244

This table describes school characteristics for the descriptive study. The upper four rows use data from the Learning Management System. The remaining rows use data from the NCES Common Core Data set.

Table 2: Parent-Portal Usage Information: 2012-2013

Variable	Mean	Observations
Share ever logged in	24.3%	149,107
Share who log in per day	1.7%	25,792,800
Average Logins per week	0.13	3,439,040
Average Total logins	13.29	146,060

This table describes school characteristics for the descriptive study. These numbers are constructed using data from the Learning Management System.

Table 3: School-Level Correlates of Adoption

Dependent variable		Ever Logged In	
Black	0.019 (0.052)	Hispanic	-0.237* (0.130)
Middle School	0.036 (0.028)	High School	-0.129*** (0.028)
Share Free/Reduced Lunch	-0.194** (0.078)	Suburban	0.051** (0.031)
Urban	0.016 (0.035)	Test Scores	0.020*** (0.005)
Rural	0.016 (0.028)	Logins/Teacher (thousands)	0.047*** (0.012)
Student/Teacher	-0.009*** (0.002)		
Observations	264		
R-squared	0.54		

This table presents results from a regression of the school-level share of parents who have ever logged into the parent portal on school-level demographic and performance indicators. Student/teacher ratios are coded as missing if larger than 100. Teacher logins are coded as missing if larger than the 99th percentile of all logins. Missing values are imputed and indicators for missing data are included in the regression. Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 4: Balance Table

	Treatment Mean	Control Mean	T-C	P-value	N	Obs.
<u>Treatment v. Control</u>						
GPA	2.43	2.48	-0.05	0.50	59	15,192
Fraction Missing	0.08	0.07	0.01	0.65	59	16,174
Parent Logins	0.60	0.74	-0.14	0.16	59	16,367
Student Logins	23.3	20.6	2.68	0.24	59	16,367
<u>Spillover v. Control</u>						
GPA	2.44	2.48	-0.04	0.53	59	15,680
Fraction Missing	0.08	0.07	0.01	0.64	59	16,639
Parent Logins	0.66	0.74	-0.08	0.43	59	16,827
Student Logins	22.7	20.6	2.71	0.21	59	16,827
<u>School Level</u>						
White	0.63	0.64	-0.01	0.64	58	N/A
Black	0.30	0.31	0.02	0.53	58	N/A
Hispanic	0.02	0.03	-0.01	0.64	58	N/A
Fraction FRL	0.60	0.61	-0.01	0.90	58	N/A
Rating	4.5	5.0	-0.49	0.34	54	N/A

All data are at the student level and are constructed from the learning management company data, with the exception of variables under the "School Level" heading, which are from the NCES Common Core Data and are school-level aggregate variables. Standard errors clustered at the school level are shown in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5: Treatment and Spillover Effects on Adoption

Dependent variable	Ever Logged In			
Treatschool	0.014 (0.011)	0.024** (0.009)	0.023** (0.009)	0.024** (0.009)
Treatschool × Untreated	-0.027 (0.005)	-0.027** (0.005)	-0.028** (0.005)	-0.027*** (0.005)
Control mean	0.07		0.065	
Observations	18,429	18,429	18,121	18,121
Controls	No	Yes	No	Yes
Outliers Excluded	No	No	Yes	Yes

All data are at the student level and are constructed from the learning management company data. Standard errors clustered at the school level are shown in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6: Treatment and Spillover Effects by Parents and Students

Dependent variable	Student Use Only	Parent Use Only	Parent & Student Use	Parent or Student Use
Treatschool	0.047 (0.031)	-0.003 (0.006)	0.026** (0.011)	0.071** (0.032)
Treatschool×Untreated	0.017* (0.009)	-0.011** (0.003)	-0.016*** (0.006)	-0.010 (0.007)
Control mean	0.532	0.049	0.101	0.683
Observations	21,854	21,854	21,854	21,854
Controls	Yes	Yes	Yes	Yes
Outliers Excluded	No	No	No	No

All data are at the student level and are constructed from the learning management company data. Standard errors clustered at the school level are shown in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 7: Effects on Parent Usage

Dependent variable	Total Parent Logins			
Treatschool	-0.275 (0.656)	0.550*** (0.205)	0.571** (0.278)	0.543** (0.201)
Treatschool×Untreated	0.058 (0.213)	-0.161 (0.228)	0.002 (0.212)	-0.156 (0.225)
Control mean	2.78		2.30	
Observations	21,854	21,854	21,453	21,453
Controls	No	Yes	No	Yes
Outliers Excluded	No	No	Yes	Yes

All data are at the student level and are constructed from the learning management company data. Standard errors clustered at the school level are shown in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 8: Effects on Student Usage

Dependent variable	Total Student Logins			
Treatschool	8.13 (6.01)	5.80* (3.27)	10.9* (5.94)	5.82* (3.27)
Treatschool×Untreated	-0.153 (0.763)	-0.171 (0.223)	-0.304 (1.13)	-0.162 (0.764)
Control mean	44.5		44.0	
Observations	21,854	21,854	21,453	21,453
Controls	No	Yes	No	Yes
Outliers Excluded	No	No	Yes	Yes

All data are at the student level and are constructed from the learning management company data. Standard errors clustered at the school level are shown in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 9: Effects on Student GPA

Dependent variable	Grade Point Average Z-Score			
Treatment	0.080 (0.087)	0.098** (0.048)	0.124 (0.084)	0.114** (0.049)
Treatschool×Untreated	-0.007 (0.021)	0.005 (0.017)	-0.010 (0.021)	0.004 (0.017)
Observations	19,218	19,218	18,878	18,878
Controls	No	Yes	No	Yes
Outliers Excluded	No	No	Yes	Yes

All data are at the student level and are constructed from the learning management company data. GPA standardized according to control-group means. Standard errors clustered at the school level are shown in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 10: Subgroup Effects on Student GPA

Dependent variable	Grade Point Average Z-Score							
Treatschool	0.126 (0.116)	0.107** (0.045)	0.135* (0.068)	-0.015 (0.141)	0.171 (0.121)	0.107** (0.046)	0.072 (0.046)	-0.129* (0.065)
Treatschool×Base GPA	-0.010 (0.043)							
Treatschool×Female	-0.013 (0.022)							
Treatschool×Share Black	-0.113 (0.114)							
Treatschool×Share Reduced-Price Lunch	0.214 (0.218)							
Treatschool×GS Rating	-0.013 (0.024)							
Treatschool×Base Usage	-0.008* (0.004)							
Treatschool×Student Base Usage	0.001** (0.000)							
Treatschool×Teacher Base Usage	0.001*** (0.000)							
Observations	19,218	19,218	19,218	19,218	19,218	19,218	19,218	19,218
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Outliers Excluded	No	No	No	No	No	No	No	No

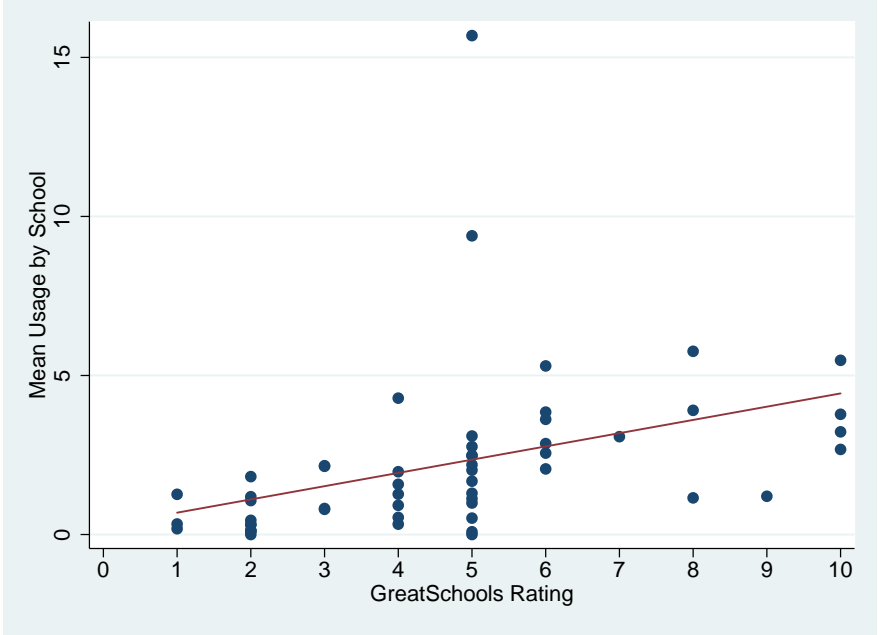
All data are at the student level and are constructed from the learning management company data. GPA standardized according to control-group means. Standard errors clustered at the school level are shown in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Appendix


Figures

Figure A.1: Outlier Usage




The figure shows the distribution of endline usage by GreatSchools' test-score proficiency rating with a fitted line from a regression of usage on rating.

Figure A.2: Outlier Analysis



This figure shows usage treatment effects estimated through 59 regressions with no control variables. Each regression excludes a particular school and each dot on the graph shows a de-meaned treatment effect when that school is excluded. The treatment effects are demeaned using a leave-one-out mean, so that the treatment effects center around zero. Data come from the LMS company.

Figure A.3: Outlier Analysis



This figure shows usage treatment effects estimated through 59 regressions, controlling for baseline usage to increase precision. Each regression excludes the outlier school shown in Figure A.1., and then excludes one additional school one at a time. Each dot represents the treatment effect for excluding a different school. All treatment effects are significant at the 5% level. Data come from the LMS company.

Tables

Table A.1: Attrition

Dependent Variable	Has Final Grade	Has Final Grade
Treatschool	0.020 (0.020)	0.028 (0.020)
Treatschool×Untreated	0.002 (0.007)	0.002 (0.007)
Control mean	0.88	0.88
Observations	21,854	21,453

All data are at the student level and are constructed from the learning management company data. The outcome variable is an indicator for a student having a final grade in the system. Standard errors clustered at the school level are shown in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$