Dynamic Impacts of School-based Internet Access on Student Learning: Evidence from Peruvian Public Primary Schools

Leah K. Lakdawala, Eduardo Nakasone, and Kevin Kho*

Abstract

We investigate the impacts of school-based internet access on second graders' test scores, using over 2 million student observations from a panel of Peruvian public primary schools. We identify effects up to 6+ years after installation on different cohorts of second-grade students, exploiting variation in the timing of internet access induced by the rollout of a national program. We find positive but modest short-run impacts, but importantly, these effects grow for subsequent cohorts. Indeed, short-run estimates alone would have led to different conclusions. These dynamics underscore the value of extended evaluation windows to allow benefits of educational technology to materialize.

In recent decades, developing countries have achieved large increases in school enrollment, particularly at the primary level. However, most remain far behind developed countries in terms of school quality as measured by student achievement (Glewwe and Kremer 2006). New approaches to improving school performance, such as Information and Communication Technologies (ICTs), have garnered increasing interest. The promise of boosting modern-day digital competencies, promoting interactive student-centered teaching models, and providing up-to-date learning materials even in remote areas has encouraged developing countries to invest considerably in ICTs in schools (World Bank, 2018; Escueta et al., 2017; One Laptop per Child, 2016; UNESCO, 2012; Trucano, 2016; International Telecommunication Union, 2014).

Among ICTs, the internet in particular may have important pedagogical uses in developing countries. Internet access can provide underserved students with otherwise unavailable sources of information (Levin and Arafeh, 2002). Similarly, internet can expand teachers' access to references and teaching aids as well as their ability to share information among peers (Jackson and Makarin,

*Leah K. Lakdawala: Wake Forest University, lakdawl@wfu.edu. Eduardo Nakasone: Michigan State University and the International Food Policy Research Institute, eduardo@msu.edu. Kevin Kho: Michigan State University, kevinkho256@gmail.com. We thank Jorge Mesinas at the Peruvian Ministry of Education for access to the data from the *Evaluacion Censal de Estudiantes*. We are very grateful to Chris Ahlin, Diether Beuermann, Prashant Bharadwaj, Julian Cristia, Pascaline Dupas, Steven Haider, Scott Imberman, Emilia Tjernström, and Maisy Wong for their insightful and helpful comments and suggestions. Previous drafts of this paper were circulated under the titles "Impact of Internet Access on Student Learning in Peruvian Schools" and "Dynamic Impacts of School-based Internet Access on Student Learning."

2018; Purcell et al., 2013). However, as with many new technologies, benefits materialize only after a period of learning and adaptation, suggesting the importance of understanding the dynamic effects of ICT interventions over time.

Despite the potential of the internet to improve learning, few studies have rigorously evaluated its impacts on student performance in developing countries. Though previous research in *developed* countries has led to ambivalent conclusions on the effectiveness of internet access as a learning input (Belo, Ferreira and Telang, 2014; Faber, Sanchis-Guarner and Weinhardt, 2015; Gibson and Oberg, 2004; Goolsbee and Guryan, 2006; Machin, McNally and Silva, 2007; OECD, 2015; Vigdor, Ladd and Martinez, 2014), school-based connectivity can be potentially more important in developing countries due to lower levels of teacher skills, larger class sizes, and limited access to other conventional inputs.[1] Additionally, since the broader literature on ICTs (Escueta et al. 2017 and Bulman and Fairlie 2016) typically examines bundles of interventions such as computer access, learning software, and internet expansion[2], it is not yet clearly understood how internet on its own influences learning.[3]

Moreover, most prior studies of internet access — and of ICTs more generally — have been based on short-term evaluation windows, and are thus only designed to detect somewhat immediate treatment effects. Importantly, such studies may overlook potential longer term impacts that may follow from an initial learning period, during which teachers, students, and administrators adapt to new technology. Hence, detecting gains in learning that may arise over such a learning period requires a longer evaluation window.

We examine the impact of internet access on the performance of the universe of students that attended public primary schools in Peru that initially acquired internet between 2007 and 2020 or that remained unconnected by 2020, emphasizing its dynamic effects in schools over time. Over this period, more than 11,300 schools (which jointly enroll about 2 million students per year) gained access to internet. We link administrative data on school-based access to internet with their students' math and reading scores from a large-scale national test that covers nearly the universe of second graders in public schools in Peru between 2007 and 2016. We construct a panel dataset of around 23,300 schools where we observe the scores of about 2.3 million second grade students during

[1] In a recent paper, Malamud et al. (2019) investigate the impact of home-based internet on Peruvian students' school performance, finding no statistically significant effect on standardized test scores 9 months after the implementation of the program. The authors posit that too little time might be spent on computers at home for any educational benefit to materialize. Relatedly, children might use internet as a tool for entertainment rather than learning. Both of these problems might be reduced when internet is provided at school rather than home.

[2] Some notable exceptions analyze the individual impact of computer access (Beuermann et al., 2015; Cristia et al., 2017; Barrera-Osorio and Linden, 2009; Mo et al., 2013; Toyama, 2015; de Melo et al., 2013; Sharma, 2014; Meza-Cordero, 2017; Bai et al., 2016; Malamud and Pop-Eleches, 2011; Meza-Cordero, 2017; Sharma, 2014) or adaptive learning software (Bando et al., 2016; Banerjee et al., 2007; Carrillo, Onofa and Ponce, 2010; He, Linden and MacLeod, 2008; Linden, 2008; Muralidharan, Singh and Ganimian, 2019; Araya et al., 2019) in developing countries. However, there is little evidence on the impact of internet access.

[3] Previous work (e.g., Cristia, Czerwonko and Garofalo, 2014; Bet, Ibarraran and Cristia, 2014; Sprietsma, 2007) has assessed programs providing school-based internet as part of broader ICT expansion schemes. However, these papers do not aim to discern the effect of internet separately from that of other technologies.

our study period. To fully exploit the longitudinal structure of the school data and identify dynamic effects, we employ an event study framework in addition to a trend break analysis — approaches which also allow us to detect and control for pre-existing trends in student performance. Since we observe a large panel of schools over several years, we are also able to assess how other determinants of school performance change over time, tracing out the dynamics of several important student, teacher, and school-level inputs. This allows us to discuss various potential channels through which internet affects student performance, as well as to explore the possibility that other confounding factors drive our results.

We exploit plausibly exogenous variation in the timing of internet access and compare different cohorts of second-grade students who attended these schools before and after they get connected to internet. Using within-school variation in the timing of internet installation, we find that internet access leads to initial modest test score improvements of 0.028 standard deviations in math and 0.017 standard deviations in reading for second-grade students in the first year after installation. Importantly, this advantage grows significantly over time, reaching 0.110 and 0.063 standard deviations five periods after installation for math and reading, respectively. It is important to underscore that our strategies only allow us identify school-level – as opposed to student-level – dynamics. That is, we find that a second grade student in a school that has had internet for 5 years performs 0.110 standard deviations better on math tests than a second grade student at the same school before internet has been installed. The trajectory of estimated effects implies that schools become more efficient in using the internet over time to produce improvements in students' test scores.

We posit that this growth in our estimated impacts over time reflects an adaptation period, during which schools must learn to integrate new technologies. Namely, we observe that schools respond to internet access by hiring teachers with formal training in digital skills, and that this process follows only gradually. In particular, schools are 27% more likely to have a computer-trained teacher by the fifth year after installation relative to before internet access. That the gradual growth over time in test scores shadows growth in the staffing of computer-trained teachers suggests that complementary investment in staff computer proficiency is needed to fully exploit internet-enabled classroom capabilities.[4] This finding aligns with a long line of previous literature examining the impacts of general purpose technologies and the complementary investments and organizational changes that ultimately drive long-run productivity gains (e.g. Brynjolfsson and Hitt 2000).

Furthermore, our data offer evidence for several additional channels through which internet access improves test scores. First, we use a completely separate data source (the Peruvian National Household Survey) to show that internet access at schools leads to meaningful increases in student use of internet. 6 or more years after gaining connections at their schools, public primary school

[4]Similarly, evaluations of laptop provision in the U.S. (Hull and Duch 2019) and computer assisted learning in China (Mo et al. 2015) estimate that the effects of ICT interventions grow over time. More generally, Jackson and Mackevicius (2021) find that the benefits of increased school spending rise with years of exposure.

students report an increase in internet usage of 122% relative to the period before their local schools become connected. This suggests that increased use of internet by students partially explains the improvements in student performance. Second, we present descriptive evidence of teachers' use of internet as a pedagogical tool. For this purpose, we use two nationally representative teacher surveys (the National Survey of Teachers and the National Survey of Educational Institutions). Teachers report that internet is one of the most important materials enhancing teaching. Moreover, teachers at schools with good quality internet connections also report less difficulty performing teacher activities than those without internet connections at school. Taken together, these findings indicate that access to online materials may boost student performance above and beyond the impacts of direct student use of internet-connected computers. We also find suggestive evidence that gains in test scores are larger in schools with high student-to-teacher ratios and that short-run gains in test scores appear largest for schools with relatively high teacher qualifications. However, the estimates in these subsamples are somewhat imprecisely estimated so differences in estimated effects across the groups are not statistically significant. Unfortunately, due to data limitations, we are not able to examine all possible mechanisms. Notably, it might be that internet access prompted differences in student engagement (e.g., attending classes more regularly, exerting more effort, etc.) or teaching practices (e.g., teachers might allocate class time differently when technology is available). We also lack detailed data about students' backgrounds and teachers' characteristics (age, experience, etc.). These limitations prevent us from further exploring some potential sources of heterogeneity in our study.

Our findings are robust to a number of alternative explanations. Concerning potential endogeneity in the timing of internet access, we find that, conditional on year and school fixed effects and a set of time-varying school characteristics (e.g., school enrollments, infrastructure, and resources), schools receiving access to internet do not exhibit systematic pre-trends in performance or different levels of scores prior to access. Second, we also find that our results are not explained by concurrent changes in several other important inputs (infrastructure, textbooks, teachers, or computers) — though we recognize that this is not an exhaustive list of inputs — or by pre-existing trends that differ by geographic areas, administrative units, or initial test performance. Third, while our main specifications are based on an unbalanced sample of schools, our results are very similar when using different sample restrictions (including a sample of non-attriting schools). Fourth, analyzing student composition within schools shows that our findings are not likely to be driven by endogenous sorting of students. Lastly, we show that our results are virtually identical using an alternative estimator that avoids negative weighting and is robust to heterogeneous treatment effects.

We contribute novel insights and perspective to a nascent body of research in developing countries on the educational benefits of school-based internet access, as well as to a wider literature concerning ICTs as schooling inputs.[5] The size and time span of our data present opportunities to

[5]Another study, Hopkins (2014), also examines the relationship between internet access and test performance in Peru using similar data. However, one important difference between our study and Hopkins (2014) is that we imple-

complement and contextualize existing studies, which focus largely on short term effects (usually within one academic year)[6], typically estimated on smaller and more localized samples of schools than ours. In contrast, we use our extended study period to analyze the effects of internet access more than 6 years after it is introduced to schools. Indeed, in the absence of longer-run estimates, our short-run estimates alone would have implied very different conclusions about the efficacy of school-based internet. Thus, our results indicate that this longer evaluation window is highly relevant to understanding the impact of internet access, due to the dynamic effects of internet on learning over time. Additionally, the large size of our sample — containing over 23,000 public schools — and the fact that we study a program that was rolled out nationally (affecting public primary schools serving roughly 2 million children in a given year), allow us to precisely assess the impacts of internet access in conditions relevant for policies implemented on a broad scale.

Finally, we make progress in identifying the gains that internet access produces over hardware resources alone as we study the effect of internet access conditional on other computing resources. Anecdotally, the usefulness of school computers without internet access has been limited by lack of access to information (National Public Radio 2012) and the inability to obtain routine maintenance and software updates — particularly in remote, difficult-to-reach locations (One Laptop per Child 2011). Indeed, previous studies find that computers alone (without internet access) have no discernible impacts on student learning (Bet, Ibarraran and Cristia 2014, Barrera-Osorio and Linden 2009, Cristia et al. 2017, Beuermann et al. 2015 Mo et al. 2013). In this study, we present some correlational evidence that — consistent with previous studies — suggests that the impacts that we find are not driven by an increase in computer access in schools that gained internet connections. Moreover, while we acknowledge that we do not have a credible source of exogenous variation in schools' investments in computers, the patterns of computer access that we observe in our data appear to be inconsistent with the dynamics in test scores.[7]

To the best of our knowledge, the scale, longitudinal length, and setting of this study, along with the comprehensiveness of the available data uniquely address important gaps in the existing literature. More broadly, our work contributes to understanding the role of internet access in economic development. Prior research concludes that the spread of fast internet led to higher employment, incomes, and wealth in African countries (Hjort and Poulsen 2019). Our results imply that increased human capital production may factor importantly in this progress.

---

ment a school-fixed effects strategy whereas Hopkins (2014) compares internet-connected schools to non-connected schools. For this reason, we regard Hopkins (2014) as being important for establishing a statistical relationship between internet access and test scores, but one that is ultimately correlational rather than causal.

[6]However, evaluations beyond one year are becoming less rare; Mo et al. 2015 studies the impact of a 1.5 year computer assisted learning program.

[7]We discuss more fully in Section 3 why we do not find it likely that concurrent changes in computing resources are likely to explain our results.

1 Setting and data

1.1 Education and ICT Access in Peru

Education in Peru is compulsory and free through the public school system beginning at age 3 and continuing until the end of secondary school. In the past few decades, Peru has greatly increased access to primary school (grades 1-6, approximately age 6-11), raising the net enrollment rate from 85.6% in 1980 to 97.9% in 2015 (The World Bank 2016). At the same time, however, the education budget has seen little growth, and thus greater enrollment over time has eroded per-student resources (Saavedra and Suarez 2002). The World Bank (2012) finds that, within Latin America, only the Dominican Republic has a lower education expenditure-to-GDP ratio than Peru.

This dearth of resources has been accompanied by Peru's poor performance in the OECD's Program for International Student Assessment (PISA) — an international standardized test among 15 year olds. In 2012, Peru ranked last out of 65 participating countries in all three evaluated subjects, with results revealing that most Peruvian students have serious deficiencies in math (75% deficient), science (69%), and reading (60%). In 2015, Peru jumped to the 64th place (out of 70 countries in the evaluation), nonetheless demonstrating that substantial progress remains to be made. Widespread under-preparedness is evident as early as primary school. In 2007, the Ministry of Education began administering yearly standardized tests, the National Student Assessment or *Evaluacion Censal de Estudiantes* (henceforth ECE, described below), to all second graders registered in classes with five or more students. The inaugural results of the ECE in 2007 showed that only 7% of students acquired skills mandated by the national curriculum in math and 13% in reading (Appendix Figure A.1). Despite improvement since then in test scores and in the proportion of students meeting expected skill levels, the quality of schooling has continued to prove inadequate for many children; even by 2016, less than 40% of second graders achieved proficiency in math (46% in reading).

In the early 2000s, the Peruvian government launched *Plan Huascarán*, which produced much of the variation in school internet access observed during our sample period. This project aimed to "incorporate information and communication technologies to increase the coverage, quality, decentralization, democratization, and equity of the Peruvian education system." Project planners ambitiously aimed to install hardware and internet in 32,000 schools by 2020.[8] *Plan Huascarán* targeted schools under public management, particularly in rural, peri-urban, and high poverty areas. Officially, selection into the program was rationed, with each Local Educational Management Unit (UGEL) allowed to request installation for a set quota of schools (see Appendix Figure A.2 for an excerpt of the official Ministry of Education flow chart that outlined the specific prioritization protocol under *Plan Huascarán*). As prerequisites for program selection, schools needed to have electricity and a computer lab with anti-theft measures (i.e., perimeter fencing). Within each UGEL

[8] Though teacher training was officially part of *Plan Huascarán*, in practice there was little emphasis on teacher training (Balarin (2013)).

and level, prioritization among qualified schools was officially based solely on the size of the student population, with larger schools receiving higher priority. Lists of eligible schools were aggregated to the regional level and then submitted to *Plan Huascarán* headquarters, accompanied by data sheets on the characteristics of each school listed, a sketch and description of each school's computer facilities, and the discussion minutes from each UGEL.[9] Officially, no school was integrated into the project without all required information.

As a consequence of initiatives such as *Plan Huascarán* and the One Laptop per Child program (OLPC, undertaken by the Peruvian government in 2008)[10], the ratio of students to computers in primary schools fell dramatically from 240 to 6 between 2000 and 2014 . In parallel, the government has steadily increased access to internet in schools (as described in Section 1.2.1). In 2013, the Ministry of Education announced plans to triple the number of schools having internet access.

1.2 Data

Our primary analysis uses school-level data from two sources administered by the Ministry of Education: the *Censo Escolar* (CE), an annual census of schools, and the *Evaluación Censal de Estudiantes* (ECE), an annual standardized test of second graders' skills.

1.2.1 Censo Escolar (CE) and School-based Internet Access

Each year, all school principals are required to submit two forms to the Ministry of Education. Between April and July, principals complete a form on enrollment (by grade and age), teachers (by qualification), available supplies and materials (e.g., books, computers, and laboratories), and infrastructure (e.g., access to utilities, building characteristics, and internet connectivity). Between December and February, another form is completed on year-end pupil outcomes (e.g., promotion and repetition rates, number of pupils transferring to other schools).[11] We refer to the CE for data on school characteristics such as internet access, enrollment, teachers, educational materials and resources, and physical infrastructure.[12] Between 2007 and 2020, around 31,100 public primary

[9] A translated version of the school data sheet is provided in Appendix Figure A.3.

[10] Peru has been the single largest buyer of OLPC laptops and to date has distributed close to one million laptops, mainly targeting school children in poor areas of the country. As our analysis to follow accounts for the total number of computers in a school, including OLPC laptops, we indirectly control for the influence of OLPC. For a discussion of the OLPC program in Peru, see Trucano (2012). In general, impact evaluations of OLPC in Peru suggest that the provision of laptops did not improve student performance (Beuermann et al. 2015; Cristia et al. 2017). We address the issue of concurrent increases in computing resources (including OLPC laptops) explicitly in Section 3.3.

[11] The school year in Peru runs from March to December.

[12] While this information is self-reported by school principals, the Ministry of Education applies different filters and verifies the consistency of the data with secondary data sources. The CE forms are submitted to the Ministry of Education electronically and include consistency rules to avoid reporting errors. Once the electronic forms are submitted, the Ministry of Education validates the information with teacher payrolls, delivery records of materials, and historical information on enrollment. To provide further evidence that schools do not strategically misreport information (for example, inflating enrollment), we compare second grade enrollment as reported in the CE to the number of students scheduled to take the exam in the ECE data. The median discrepancy in these numbers is 0 and the average is 0.33 students (i.e. on average 0.33 more second grade students are reported in the ECE than in the

schools reported administrative information in the CE.

We use information from the CE to determine the timing of initial internet connection among schools in our sample. As detailed in Section 1.2.3 below, we focus on the period 2007-2016 (for which we have student test scores). However we use information from the CE until 2020 (the most recent round currently available) to identify the year in which schools initially install internet and assign them accordingly in our event time specification (described in Section 2.1).[13] This allows us to increase our sample size for periods before internet is installed in schools and better assess any potential pre-existing trends in schools prior to access to internet.

Administrators report in the first semester of every year whether their school currently has access to internet. Though some schools report gaps in internet access, the data do not allow us to distinguish between temporary outages and longer-term disruptions to connectivity.[14] Based on this information, we determine the first year in which a school reports gaining access to internet and interpret this as the time of connection. In our estimation framework, this implies a conservative estimate of the impact of internet access, because we treat schools that might have permanently lost their connections as still being connected. Another benefit of using initial internet connection rather than current access is that we avoid bias due to endogenous changes in access. We estimate that 11,310 schools — and the roughly 2 million primary school students in these schools each year — at some point gained internet connectivity between 2007 and 2020. This implies that the rate of internet connection in schools increased from below 5% to more than 42% and that the share of students with internet connection in their schools jumped from 23% to 81% over this time period (Figure 1).

Most of the observed expansion in internet connectivity during this period was due to *Plan Huascarán*. In Appendix Table A.1 we verify that the official qualification and prioritization rules set by the Ministry of Education do in fact predict actual installation.[15] Schools received priority primarily based on quotas by province (Local Education Management Units, UGELs), high poverty status, location in a rural or marginal urban area, public versus non-public management, the presence of required infrastructure (including electricity, a computer lab, and anti-theft measures such as perimeter fencing), and enrollment. Column 1 includes these characteristics and state-year fixed effects to control for aggregate and state-level trends in internet connectivity. To capture high poverty status, we include district-level fixed effects. We also include UGEL fixed effects to account for the UGEL-specific quotas. As expected, all required/prioritization characteristics positively and significantly predict internet access. In column 2, we add a control for perimeter fencing (available

CE); since the CE is reported earlier in the school year than the ECE, this discrepancy could arise naturally due to students entering after the CE forms have been submitted.

[13]CE data have been collected since 2000. However, the CE has only included variables for internet access since 2006.

[14]Out of the 31,111 public primary schools with at least one year of data in the CE between 2006 and 2020, 29.4% report not having access after having access in a previous year; about 44.1% of those schools regain internet access at a later point.

[15]Note that the official prioritization rank and program status of schools under Plan Huascarán is not available.

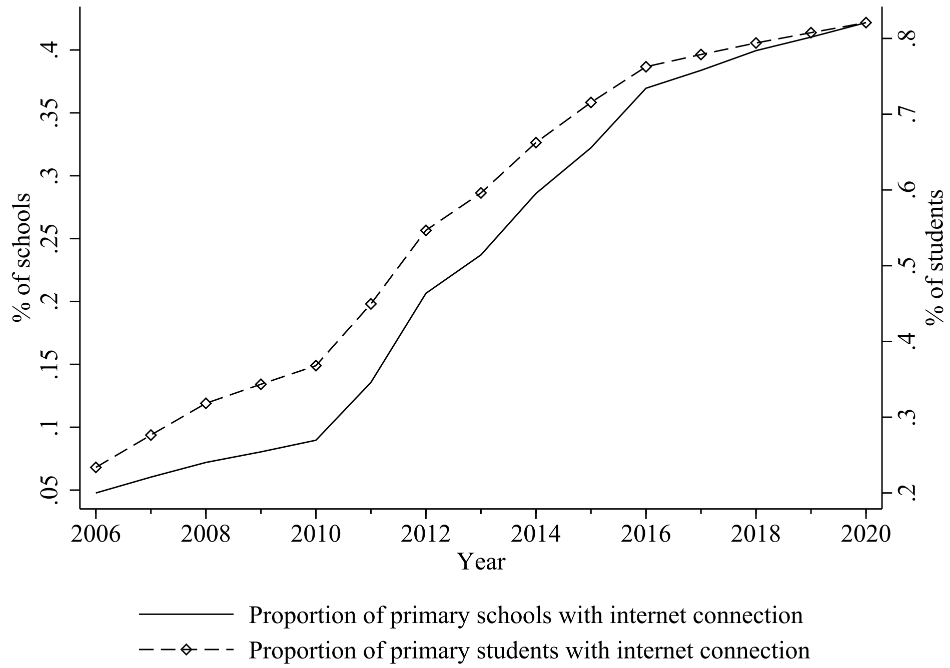

Figure 1: Internet Connectivity in Primary Public Schools, 2006-2020

The stock of schools that gained internet connections is based on the first year in which they report internet access in the Peruvian *Censo Escolar*.

only for 2010 and later). Finally, in column 3, we include information from school data sheets (number of computers used for instruction, number of computers used for administrative purposes, and number of teachers), respectively. Since these factors predict internet access and are also likely to influence student performance directly, we control for all of these measures in our main specifications (except perimeter fencing due to data limitations).[16]

1.2.2 Evaluación Censal de Estudiantes (ECE)

The Ministry of Education also mandates the *Evaluación Censal de Estudiantes* (ECE), a yearly standardized assessment of second graders' skills, which is administered in late November or early December (before the end of the school year). In order to ensure uniform testing environments — and to prevent content leaks or influence from school personnel — the Ministry hires independent staff to administer the test in all schools simultaneously. As the same test is given to all schools, neither the content nor the testing environment varies by school characteristics. Furthermore, the ECE was designed for comparability of results over time: experts defined a set of basic competencies

[16] It is not possible for us to calculate exactly how much of the overall increase in internet access over this period is directly attributable to *Plan Huascarán* versus other efforts. We discuss the robustness of our results to restricted samples of schools that likely gained access to internet only via *Plan Huascarán* in Section 3.4.

prior to the test's first administration. Hence, since its inauguration in 2007, the ECE has assessed the same skill sets with considerable consistency. We use student-level ECE scores from 2007 to 2016.[17] To account for differences in difficulty across cohorts and year-to-year changes in test score dispersion, we standardize ECE scores across the universe of test takers in public schools within each year. It is very important to highlight that the ECE is given to a *different* cohort of second grade students each year, so we have a repeated cross-section of students from a panel of schools. That is, we observe test scores for each student only once (in second grade).

The ECE gauges the academic performance of the vast majority of second graders in Peru, targeting all public and private schools that meet two criteria: 1) having at least five second graders enrolled during the test year, and 2) using Spanish as the primary language of instruction. The rationale for the first criterion is entirely budgetary, as smaller schools are often in remote areas and would take considerable resources to reach. As it stands, the ECE already requires about 40,000 field workers each year. Schools teaching in indigenous languages are covered under a separate testing schedule. In total, 14,000 – 19,000 primary public schools participated per year (48% to 65% of all primary schools; see Appendix Figure A.4a). About 26% – 49% of schools were exempt under the minimum enrollment or language criteria. The remaining schools (between 2% and 11%) were not tested due to logistical problems. The coverage of the test was nonetheless very broad: since the smallest schools were excluded by definition, and since schools in native language tend to have modest enrollments, between 82% and 90% of all second graders in the country were tested in the ECE in a given year (Appendix Figure A.4b).

1.2.3 Estimation Sample

In Appendix Table A.2 we account for the different steps we took to build our estimation sample. Columns 1-3 characterize all public schools that appear in the CE during our study period. We characterize these schools using information from 2007. When information for 2007 is not available, we use the earliest available data for each school. Overall, there were 31,111 schools that reported information in at least one of the rounds of the CE during this period. Our empirical strategies exploit the timing of internet connection within schools. Therefore, in our sample we include all schools that initially installed internet in 2007 or later and all schools that remain unconnected by the end of our CE sample period (2020). We exclude all schools that were connected prior to 2007 (according to the 2006 CE), as we are unable to determine the initial year of access for these schools. This only excludes a small share of public schools, as only 1,366 (4.3%) were internet-equipped by 2007. This leaves us with 29,745 schools, over 95% of all public primary schools in Peru (Appendix Table A.2, columns 2 and 3). Though 18,435 schools (59.3%) remained unconnected by 2020, 11,310

[17] The ECE was not conducted in 2017 because of significant school absences due to El Niño and teacher strikes. Starting in 2018, the test has been given only to a very small set of schools (around 1.7% of the schools tested in 2016) and a new sample of schools is chosen each year. Given that our identification strategy (described in Section 2.1) includes school fixed effects, we can only use ECE scores through 2016.

schools (36.4%) gained access between 2007 and 2020. The sharp expansion in internet connectivity during this period allows us to form insights about the effect of internet access using variation from a large number of schools.

We then merge this information with annual test scores from the ECE, resulting in 26,075 matched schools (about 84% of all public schools; Appendix Table A.2, columns 5 and 6). The sample of schools in columns 2-3 (all schools who gained access to internet between 2007-2020 or did not have access by 2020) and columns 5-6 (for which we observe average scores) appear to be nearly identical on observable characteristics.

All in all, our estimation sample includes 23,318 schools that were tested in the ECE in our sample period (2007-2016), that gained internet between 2007 and 2020 or remain unconnected as of 2020, and that are contain all of the necessary covariate information for our main specification (described in Section 2.1). The final column of Appendix Table A.2 gives the summary statistics for the estimation sample.

While schools that remained without connection to the internet by 2020 aid us mostly in the identification of calendar year effects and covariate coefficients, the main source of variation in our event study comes from schools that adopted internet between 2007 and 2020. Importantly, the data suggest that schools that gained internet from 2007-2020 generally fall "between" the early adopters (who received internet before 2007) and non-adopters in various measures of school quality. Namely, early adopters appear to be schools with higher performance, larger enrollment, and better infrastructure and educational inputs (e.g., piped water, libraries, administrative offices, teachers, classrooms, computers, and textbooks). Conversely, non-adopters systematically appear worse in these areas. Thus, the adopters that provide the variation to identify the effect of internet access focuses neither on the best nor on the worst performing schools.

Within schools that gained access to internet in the study period, we note considerable variation in the timing of access for our analysis. This allows us to implement the event study approach described in Section 2.1. In Appendix Figures A.5 and A.6, we plot each "treatment cohort's" average math performance over time, which we normalize relative to year of internet access. For reference, Appendix Figures A.7a and A.7b represent the performance of schools that gained access prior to 2007 (the start of our sample period) and Appendix Figures A.7c and A.7d similarly display the scores over (calendar) time for schools that had not gained access to internet by 2020 (the end of our CE sample period).[18] Generally, schools that connected later or remained unconnected by 2020 exhibit lower average test scores, indicating that variation in internet access *across* schools is not random.

However, *within* a cohort of schools becoming connected in a given year (2007-2020), there do not appear to be systematic trends in scores prior to internet access. This suggests that within cohorts of treated schools, the timing of access is unrelated to test score trends on average. Furthermore,

[18]We cannot calculate event time for schools that gained access prior to 2007 or were unconnected in 2020 because we do not know the initial year of internet installation for these groups.

Appendix Figures A.7a and A.7b suggest that performance gains among treated schools are very small initially but grow over the medium term. In contrast, the relative performance of schools that had not been connected to the internet by 2020 declines over the period of analysis (Appendix Figures A.7c and A.7d). Thus, if we regard the trends in scores in this group as the counterfactual for performance in the absence of internet connections, the implied effects of internet on test scores are larger than the raw trends in Appendix Figures A.5 and A.6 suggest.

2 Empirical strategies and results

2.1 Event study specification

In order to analyze dynamic impacts of internet access over time, we estimate the following event study specification:

$$Y_{isr} = \sum_{t=-3}^{6+} \beta_t \mathbf{1}\left\{E_{sr} = t\right\} + \gamma X_{isr} + \alpha_s + \theta_{sr} + \varepsilon_{isr}, \quad t \neq -1 \tag{1}$$

Our primary outcomes of interest are standardized math and reading scores for second grade student i in school s observed in year r (Y_{isr}). Scores are normalized across the universe of Peruvian public schools within each year. α_s are school fixed effects, which capture all time-invariant observed and unobserved school-level determinants of performance.[19] X_{ir} is a set of individual and time-varying school characteristics that includes gender, class size, indicator variables for the number of second grade classes within a school, total school enrollment, number of second grade students scheduled for testing, facilities (piped water, library, administrative offices), and resources per student (classrooms, computers and teachers). It is important to control for school size relative to other schools within the UGEL because *Plan Huascarán* explicitly prioritized the schools with the highest enrollment within each UGEL, meaning that enrollment ranking within UGEL is highly predictive of access to internet. Larger schools are not only likely to differ systematically from small schools in level terms, they are also likely to experience different time-variant shocks to test scores. In light of this, our baseline specification includes year fixed effects that are specific to terciles of baseline enrollment within each UGEL, θ_{sr}.[20]

Let I_s denote the year in which school s gains internet connection (the first year in the dataset in which s reports internet access in the CE). E_{isr} represents time relative to internet access for each school; specifically, $E_{sr} = r - I_s$. Each of the event study dummy variables is set to zero for all schools that remain unconnected by 2020. We include these non-adopters in our estimation

[19] Recall that the ECE is given only to second grade students each year, so we observe each student's scores only once (in second grade) though we have a panel of schools. Because of this data structure, we cannot include student fixed effects.

[20] We use baseline enrollment to avoid any potentially endogenous changes in enrollment with respect to school internet access. Since the tercile is based on school enrollment in the first year a school is observed, it is time invariant. Results that include year effects that are specific to quintiles of enrollment are very similar (available upon request).

sample to help identify γ and the calendar year effects.[21] The coefficients on the set of event study dummy variables β_t capture the path of test scores relative to the year before a school receives internet access (i.e., relative to $t = -1$).[22] One important feature in the timing of the two datasets we use is that CE reports internet access in the beginning of the school year, while the ECE is a year-end test. Any school that installs internet after the CE (April-July) does not report internet access until the following calendar year. If internet installation occurs before the ECE exams (end of November - December), students are exposed to internet access during at least part of the year *prior* to reporting initial access in the CE. Therefore in merging internet information from the CE to test scores from the ECE, we match test scores from the ECE to the internet status in the CE of the following calendar year. This means that some schools acquire internet access in $t = 0$ (if installation occurred *before* submitting the CE information) while others acquire it in $t = 1$ (if installation occurred *after* submitting the CE information). Unfortunately, school-level information is not available for either the month of installation or completion of the CE, and so we are unable to tell how many schools receive internet in $t = 0$ versus $t = 1$. Thus, in interpreting estimates of β_t it is important to keep in mind that $t = 0$ is a partially treated year for some schools and a pre-treatment year for others, while $t = 1$ is a partially treated year for some schools and a (fully) treated year for others.

By exploiting variation in the timing of internet access *within* schools (as well as additionally controlling for aggregate year effects and a set of time-varying characteristics), we aim to identify the effects of internet access separately from potential confounders that are fixed at the school level. We consider this a refinement over Hopkins (2014) — who also examines the relationship between internet access and test performance in Peru — but compares internet-connected schools to non-connected schools. We use the event study framework to examine both pre-treatment trends and dynamic effects in a non-parametric fashion from three periods before to more than six periods after gaining internet access. We choose this event window prior to internet access because we have internet installation data from the CE until 2020, allowing us to correctly assign event time only as far back as $t = -3$ for all schools (given that our final year of test score data is 2016). We estimate coefficients for each year relative to internet access up to $t = 5$. Because the number of schools in our sample that had been connected to internet for more than five years declines substantially, we group all periods $t \geq 6$ for our estimation purposes. Standard errors are clustered at the school level to allow for arbitrary serial correlation in ε_{sr}.

Figures 2a and 2b display the results of estimating Equation 1 on our main sample when the outcomes considered are standardized math and reading scores, respectively. The full set of coefficient estimates for the event study dummy variables are reported in Appendix Table A.3. We find

[21] Including non-adopters also avoids the multicollinearity problem identified in Borusyak, Jaravel and Spiess (2021).

[22] Recent work (for example, Goodman-Bacon (forthcoming)) has illustrated issues with estimating treatment effects using a two-way fixed effects model with staggered timing. In light of this, we show our results are robust to alternative estimation strategies in Section 3.

that, prior to school internet access ($t < 0$), math and reading performance was roughly constant from year to year. Importantly, there are no apparent trends in test scores prior to internet access, indicating that the timing of internet access within schools is unrelated to pre-trends in student performance. In particular, we rule out the case in which internet installation is budgeted endogenously as a reward for steadily improving test performance. While relative student performance rises in all years following initial connectivity in both math and reading, immediate gains are small in magnitude: 0.028 and 0.017 standard deviations in $t = 1$ (the first year of partial or full access to internet in our schools) for math and reading, respectively. However, test score growth following internet access is steady in both subjects. By year 5, scores are 0.110 standard deviations higher for math and 0.063 standard deviations higher for reading relative to the year prior to internet installation. Beyond five years of internet access, this effect increases to 0.141 and 0.076 for math and reading, respectively.

Our findings echo those from other studies in developing countries that find limited or no impacts of ICTs on test scores in the short run.[23] Results from Figures 2a and 2b suggest that though classroom internet is beneficial to learning, improvement in the initial years post-intervention is small. The fact that our estimates grow over time, at least through the medium-term, is also consistent with two other longer-term studies of ICTs in education — which have also supported the need for an adaptation period to fully utilize new technologies (Hull and Duch 2019; Mo et al. 2015). Interestingly, the dynamic path of tests scores appears to be linear with respect to school exposure to internet for both subjects. This aligns with recent evidence from Jackson and Mackevicius (2021), who find that the effects of increased school spending on educational attainment grow linearly with years of exposure. It is interesting to note that many of the unguided ICT interventions (including ours) yield zero or small short-term effects, while other evaluations of adaptive ICTs - such as, computer assisted learning and related interventions - typically find much larger short-run effects (0.18 to 0.59 standard deviations); e.g., see Bando et al. 2016, Banerjee et al. 2007, Beg et al. (2019), Carrillo, Onofa and Ponce 2010, He, Linden and MacLeod 2008, Linden 2008, Muralidharan, Singh and Ganimian 2019, and Araya et al. 2019.

In comparing our estimates to the existing literature, it is important to reiterate that we identify school-level - as opposed to student-level - dynamics. Thus the increase in test scores that we observe over time after internet installation are reflection of how schools adapt to a new technology and become more efficient at improving test scores over time. In this way, impacts beyond the short run in our context are not comparable to estimates one might expect from studies that follow individuals over time.

[23] The majority of the studies in this literature focus on impacts within the first 18 months post intervention.

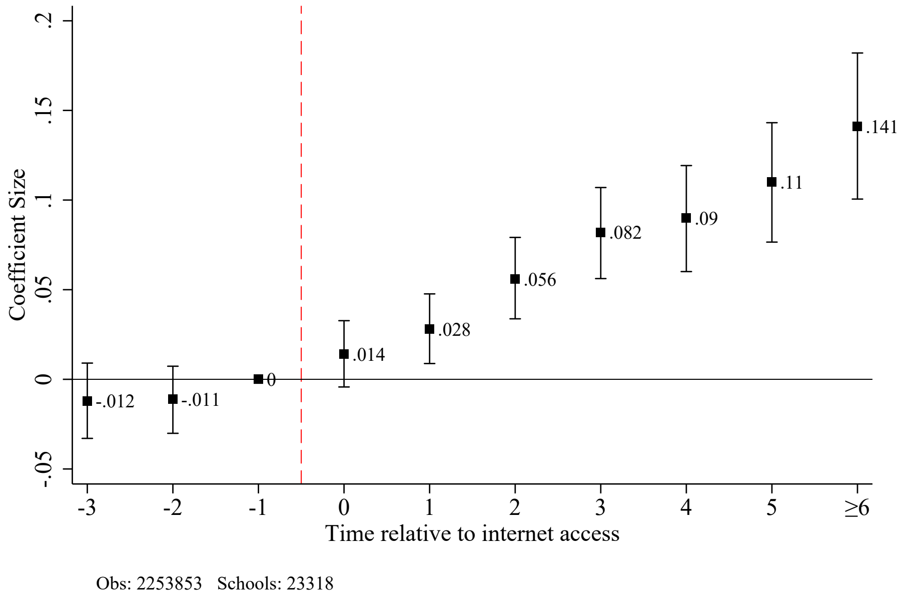

(a) Standardized Math Scores

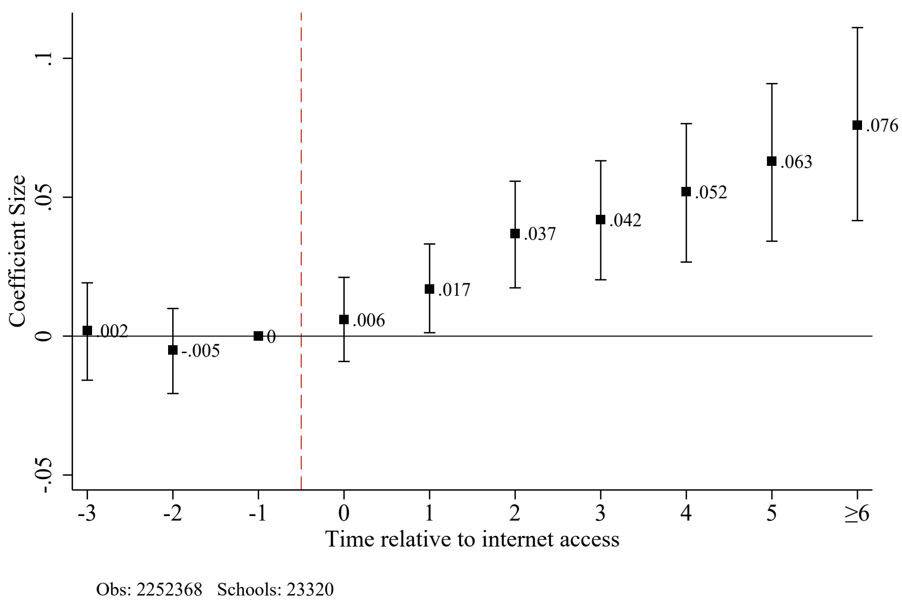

(b) Standardized Reading Scores

Figure 2: Impact of Internet Access on Test Scores

The above figures plot the coefficients and 95% confidence intervals from estimating equation 1. Scores are standardized within each calendar year to have mean zero and standard deviation of one across the universe of test takers in public schools. Coefficients capture the increase in test scores relative to the year before internet installation ($t = -1$). Note that due to the timing of the *Censo Escolar* relative to the ECE exam, some schools receive internet access in $t = 0$ and some receive it in $t = 1$. For more details, see Section 2.1. Control variables include sex, class size, indicator variables for the number of 2nd grade classes at the school, total school enrollment, number of second grade students that took the ECE, facilities (computer room, library, administrative offices), resources per student (classrooms, computers, and teachers), UGEL-specific enrollment tercile by year fixed effects, and school fixed effects. The sample includes all grade 2 students in all public schools that gain internet between 2007 and 2020 or remain unconnected by 2020. Standard errors are clustered by school.

2.2 Trend break specification

Though the shape of Figures 2a and 2b suggest a steadily increasing effect of internet access on test scores over time, it does not explicitly test for a break in the trajectory of scores at the time of internet installation. To do so, we estimate a linear trend break specification as follows:

$$\begin{aligned} Y_{isr} &= \phi_1 \text{Post-internet Access}_{ir} + \phi_2 \text{Event Time}_{ir} \\ &+ \phi_3 \text{Post-internet Access}_{ir} \times \text{Event Time}_{ir} + \gamma X_{isr} + \alpha_s + \theta_{sr} + \varepsilon_{isr} \end{aligned} \tag{2}$$

Here, Post-internet Access is a dummy variable that is equal to one in all periods after internet installation ($t \geq 0$). Event Time is a linear term for time relative to the year prior to access, $t = -1$. The control set (X_{isr}) is otherwise identical to that described in Section 2.1. In this specification, ϕ_1 captures the level shift in test scores in response to internet access; ϕ_3 represents the change in the linear time trend in math scores after schools gain internet access; and ϕ_2 accounts for any pre-existing linear trend. Based on the results in Section 2.1, it is unlikely that there are any existing pre-trends. However, one benefit of this specification is that even in the presence of any linear pre-trends in test scores, ϕ_3 measures the impact of internet access on the growth in test scores *apart* from any such trends.

Table 1: Internet Access & Test Scores: Trend Break Results

	Dependent Variable: Standardized Test Score	
	Math	Reading
	(1)	(2)
Post-internet Access	0.008	0.010
	(0.011)	(0.009)
Post-internet Access X Event Time	0.015***	0.012***
	(0.005)	(0.005)
Event Time	0.006	-0.001
	(0.005)	(0.004)
p-value for PostXEventTime	0.007	0.008
Observations	2,253,853	2,252,368
Number of Schools	23318	23320

Scores are standardized within each calendar year to have mean zero and standard deviation of one across the universe of test takers in public schools. Coefficients capture the increase in test scores relative to the year before internet installation ($t = -1$). Note that due to the timing of the *Censo Escolar* relative to the ECE exam, some schools receive internet access in $t = 0$ and some receive it in $t = 1$. Control variables include sex, class size, indicator variables for the number of 2nd grade classes at the school, total school enrollment, number of second grade students that took the ECE, facilities (computer room, library, administrative offices), resources per student (classrooms, computers, and teachers), UGEL-specific enrollment tercile by year fixed effects, and school fixed effects. The sample includes all grade 2 students in all public schools that gain internet between 2007 and 2020 or remain unconnected by 2020. Standard errors are clustered by school. Significance levels denoted by: *** p< 0.01, ** p< 0.05, * p< 0.1.

Results from estimating equation 2 are displayed in Table 1. Though we observe a level shift in both math and reading scores upon installation (0.008 and 0.010 standard deviations, respectively), the immediate effect is small and we are unable to statistically distinguish these effects from zero. On the other hand, the year on year gain in test scores in both math and reading are statistically significant and the magnitude is meaningful (0.015 and 0.012 standard deviations, respectively). This stands in contrast to the estimated pre-trends, which are close to zero and not statistically significant.

3 Robustness checks

In this section, we address several other potential challenges to identification, namely endogenous changes in sample composition in terms of both students and schools (including non-random attrition), concurrent changes in school resources, pre-existing trends that might differ by school characteristics, and other potentially non-exogenous sources of variation in the timing of internet access.[24] Finally, we show that our results are robust to an alternative estimation method that addresses many of the concerns raised by the recent literature on two-way fixed effects estimation with staggered timing of treatment.

3.1 Unbalanced panel and attrition

As we use students from an unbalanced panel of schools (schools are included when they participate in the ECE) observed over a limited window of time (2007-2016, years for which ECE data is available), it is possible that our estimated treatment effects reflect changes in sample composition. Namely, identification of pre-trends and treatment effects might rely on students from different samples of schools.[25] Appendix Figures A.9a and A.9b and columns 2 and 5 of Appendix Table A.5 suggest that our main findings are not driven by this issue. Specifically, we restrict the sample to students in schools that can appear at least twice prior to and twice following internet installation (i.e., schools that installed internet between 2009 and 2015), for which we can observe *both* pre-trends and treatment effects within the same school. Imposing this restriction drops about 2,000 schools from the sample, but the remaining sample is comparable along many observable dimensions (see Appendix Table A.4). In this restricted sample, we find no statistically significant trends in performance prior to internet access, and the estimated effects are similar in magnitude, statistically significant, and show similar dynamics as those using the full sample.

School-level attrition from the panel may pose another compositional issue. Though school-level

[24] We also perform an exact randomization exercise, in which we randomly reassign schools' initial year of internet access while maintaining the actual distribution of installation. We perform the randomization and estimate equation 1 500 times and plot the median, 5th percentile, and 95th percentile of the resulting coefficients in Appendix Figure A.8. For the pre-internet periods, the coefficients from the baseline specification are close to the median of the coefficients from the placebo exercise (essentially zero), whereas the post-internet coefficients fall well outside the 5th and 95th percentile of the placebo coefficients. We take this as additional evidence that the inference in our baseline specification is appropriate.

[25] Relatedly, we find no evidence that the trend break estimates differ significantly across "early" and "late" adopting schools (split by the median year of adoption).

attrition can happen for several reasons, Appendix Table A.6 shows that over two thirds of overall attrition is likely due to a school dropping below the enrollment threshold.[26] If internet access affects enrollment, it may determine whether a school meets the 5-student criteria to be tested in the ECE.[27] To rule out that our results are driven by selective attrition, we estimate equations 1 and 2 on the restricted sample of schools with ECE scores and the full set of covariates in *every* calendar year 2007-2016 (i.e., the sample of non-attritors). The results are displayed in Appendix Figures A.10a and A.10b and columns 3 and 6 of Appendix Table A.5. Even among schools that are observed in every year, we find that the estimated effects of internet access are sizable and grow over time. The trend break results in columns 3 and 6 of Appendix Table A.5 show that the estimated yearly gain in test scores due to internet access is similar to the baseline results (columns 1 and 3), though the trend break coefficients are not statistically significant. The "non-attritor" sample is somewhat similar to the baseline sample along many observable dimensions, albeit higher achieving and (somewhat naturally) larger in terms of enrollment (Appendix Table A.4).[28]

3.2 Student composition and endogenous sorting

Another related issue is that the composition of students *within* schools may change in response to internet access. A priori, it is hard to tell the direction of the bias that this would entail. For instance, parents who would otherwise not have sent their kids to school might decide to enroll their children in a school connected to internet. If these previously out-of-school students would have otherwise performed poorly, then our estimates of treatment effects are likely conservative. Alternatively, motivated parents seeking learning opportunities for their children may decide to transfer students from schools without internet to schools that gained connectivity. If these new students are better achievers on average, then our findings of positive treatment effects may owe to upward bias from changes in student composition.

It does not appear that an influx of high-achieving transfers or re-entrants affects our main results. We find that grade 2 transfers, re-entry, and total enrollment do not increase in response internet access in columns 1-3 of Appendix Table A.7.[29] If anything, there are fewer transfers over

[26]About 56% of attrition is due to schools having fewer than 5 second grade students, and an additional 18% is explained by having enrollment "near" the threshold (defined as having 5-8 second grade students). The remaining attrition is either due to missing ECE scores for another reason or missing CE (covariate) information. Only a very small portion of attrition is due to school closures.

[27]In Section 3.2, we provide direct evidence that internet access does not affect enrollment.

[28]Our results are also robust to using a very restricted sample that includes only students from schools that appear in all event years in Appendix Figure A.11. With 10 years of data and a 10-period event window, we have no variation in internet timing if we require all schools to appear in all event years. Thus we shorten the event window to allow for this sample restriction. The post-internet point estimates are generally as large or larger in the fully balanced sample as in the baseline sample. The fully balanced sample includes around 2400 schools (around 10% of our baseline sample) because we can only include 3 "treatment cohorts" of schools (our baseline sample contains 14 cohorts). Therefore these estimates are considerably noisier.

[29]Transfers are students enrolled in the current year who were enrolled in a different school in the previous year. Re-entrants are students that are currently enrolled but who were not enrolled in any school during the previous year (i.e., dropouts who come back to school).

time after a school connects to the internet, though the point estimate is small (column 1). We further demonstrate that the makeup of the students taking the test does not appear to change with internet access. Column 4 of Appendix Table A.7 illustrates that there is no effect of internet on the proportion of enrolled students that actually take the ECE. These results are consistent with Cristia et al. (2017) and He, Linden and MacLeod (2008), who find that neither hardware nor CAI/CAL interventions has any significant effects on attendance. In column 5, we illustrate that internet access does not seem to attract more advantaged students to schools, to the extent that the proportion of native Spanish speakers enrolled captures student background.[30][31]

3.3 Concurrent changes in school resources

Timing of internet access may also possibly correlate with changes in other school resources.[32] For example, it might be that internet provision is bundled with other inputs in a multifaceted approach to improve quality of schooling.[33] If this is the case, the improvement of students' performance that we observe might be due to increases in these other resources. Lacking data on school-level spending, our approach is to examine the effects of internet access on specific school inputs.[34]. For the most part, we do not find that the timing of internet access is correlated with increases in other observable inputs (Appendix Table A.9). Classrooms and textbooks, overall teachers (excluding computer teachers, which are separately discussed in Section 4.3.1), and qualified teachers (those with a pedagogical or university degree) per student actually *fall* slightly after internet access (though the point estimate is very small; columns 1-4).[35] However, in column 5, we note an increase in computing resources (which includes OLPC laptops) at the time of internet installation; on the other hand, the estimate of the trend break is negative and significant. More concretely, though the estimates in Appendix Table A.9 imply that there is an increase in computers per student by 0.008 in the year internet is installed (6% relative to the pre-internet mean), computers per student steadily drop to pre-internet levels in year 4 and lower than pre-internet levels thereafter. In other words, the dynamic pattern of computers per student is nearly opposite of the pattern in test score improvements; after an immediate increase, we see computers per student decline over time

[30]The proportion of Spanish-speaking students is positively related to test scores.

[31]In Appendix Figure A.13, we show the results of the event study analysis when considering the outcomes in Appendix Table A.7. Consistent with the trend break results in Appendix Table A.7, we find no statistically significant patterns in the effects on these outcomes.

[32]All of our specifications include school and UGEL-specific enrollment tercile-year effects as well as time-varying school characteristics. We show that our results do not depend on the inclusion of covariates in Appendix Table A.8.

[33]It could also be the case that internet access at schools is correlated with alternative sources of internet. However, we find that only 15% (29%) of students who used internet at school also use it at cyber-cafes (or at home) according to the 2016 Peruvian National Household Survey (ENAHO). Additionally, our results are unchanged if we include a control for whether the town nearest the school has a cyber cafe.

[34]Using state-level data on education budgets from the Peruvian Ministry of Economy and Finance's Integrated System of Financial Administration, we find that state-level spending in computing and telecom equipment is not related to spending in other categories (such as payroll and goods and services). Results available upon request.

[35]We also find no effect on teachers' contract types (permanent or temporary). Results available upon request. Our data do not contain any other information on teachers, aside from gender.

whereas we see very small impacts on test scores at the time of installation that grow larger over time. Moreover, numerous studies (nearly all RCTs) find that non-internet connected computers have no impact on student test performance (Bet, Ibarraran and Cristia 2014, Barrera-Osorio and Linden 2009, Cristia et al. 2017, Beuermann et al. 2015 Mo et al. 2013).[36] Thus, we find it very unlikely that our results are merely a reflection of a concurrent change in computers. However, we acknowledge that we do not have a source of purely exogenous variation in computers.

3.4 Differential pre-trends and other sources of endogenous internet connections

Another possibility is that access to internet is correlated with pre-existing trends in test scores. For example, districts with faster growing economies might be better able to finance internet expansions, increase public spending on education, or otherwise improve student learning. In Appendix Table A.12 we show that our results are robust to allowing for an array of group-specific pre-trends. These include groups that are defined administratively (i.e., by Local Educational Management Unit, UGEL), geographically (district, the finest geographical unit we observe), and by initial academic performance (captured by pre-internet test score deciles). Including UGEL-specific pre-trends (column 2) is important for ruling out endogenous selection for internet installation at higher levels of government (e.g. if the central government allocated more internet funds to UGELs with higher test score growth, even though the official UGEL-level quotas were determined solely by student enrollment). Allowing for pre-trends that are specific to initial test performance (column 4) is also useful in ruling out possible reversion to the mean, as we allow schools with initial poor performance to be on a separate score trajectory than high performing schools. The point estimates of the level shift and trend break in test scores are all stable in terms of sign, magnitude, and statistical significance across specifications; if anything, the magnitudes and statistical significance are slightly higher when we include district- and initial scoring decile-specific pre-trends. Overall, we take the evidence in Appendix Table A.12 to indicate that pre-existing trends in test scores and reversion to the mean do not confound our estimates of the effect of internet access.

Finally, we show that our results are not likely driven by schools that are potentially gaining internet access endogenously. While *Plan Huascarán* is likely to be the most important source of variation in internet access during our study period, schools whose principals become more resourceful or whose parents improve their motivation might have gained access through other sources. To account for this possibility, we perform two additional checks. First, we show that

[36]In Appendix Table A.10, we show that non-internet enabled computers are generally uncorrelated with test scores using the sample of schools that do not gain access to the internet during our sample period. For schools than remain unconnected by 2020, we find no statistically significant correlations between computers and test scores and correlations are small in magnitude. To make this point more directly, we perform some back of the envelope calculations (displayed in Appendix Table A.11). Even after taking into account the increase in computer resources that we find in Appendix Table A.9 and the correlations we document in Appendix Table A.10), increases in computers alone can explain very little of the observed rise in test scores (at most, 0.91% for math and 0.84% for reading).

the results to excluding schools in non-complying UGELs, defined as UGELs in which more public primary schools become connected to the internet than planned under the official *Plan Huascarán*; see columns 2 and 5 of Appendix Table A.13.[37] In this restricted sample, the estimated effects of internet are larger than in the baseline specification and statistically significant. Second, the results are also robust to excluding schools that potentially gained access to internet through means other than *Plan Huascarán*, i.e. those schools that are located in areas with some alternate form of internet access (proxied by the existence of local cyber cafes) *prior* to school-based internet access (columns 3 and 6 of Appendix Table A.13). The idea behind this restriction is that it would be very difficult for schools to install internet other than through *Plan Huascarán* in places where there is no existing internet infrastructure (and thus no cyber cafes). The effect of school-based internet access is similar in this sample of schools (though now it is not statistically significant).

3.5 Two-way fixed effects with staggered treatment timing

Recent work (for example, Goodman-Bacon (forthcoming)) raises some important issues with using two-way fixed effects estimation when there is staggered timing of treatment. We show that these issues do not explain our main estimates by illustrating robustness to the estimator proposed in Sun and Abraham (forthcoming). Similar to Callaway and Sant'Anna (forthcoming), the Sun and Abraham (forthcoming) method identifies cohort- (defined by installation year) and time-specific average treatment effects and then aggregates those individual treatment effects (the weighting method differs across the two papers)[38]. Doing so avoids potential biases that may be present under the standard two-way fixed effects event study methodology. The results using the Sun and Abraham (forthcoming) estimator are displayed in Appendix Figure A.12. The patterns and magnitudes of these estimates are very similar to our main estimates.

4 Explaining dynamics and identifying potential mechanisms

To better understand the mechanisms behind our main results, we now turn to alternate sources of data. We begin by examining student-level mechanisms and show that school-based internet access increases students' internet usage. Next, we use descriptive data to investigate teacher-level mechanisms, i.e., teachers' use of internet over and above students' internet usage in the classroom. We find that public primary school teachers in internet-connected schools find a variety of teaching

[37] We calculate the quotas as follows: first, we obtain annual quotas per UGEL from the Ministry of Education (published in 2004). Second, we multiply the annual quota by 13 (to reflect 13 years between when the quotas were published in 2004 and the end of our sample period, 2016) and then by 0.7 (to reflect that 50% of the quota was for primary (only) schools and 20% was for integrated (primary and secondary) schools). We then define complying UGELs as those with no more connected schools than the imputed quota for 2016. This is likely to overstate the number of non-complying UGELs, as the program expanded in 2007 when it became part of Directorate General of Educational Technologies (*DIGETE*).

[38] Sun and Abraham (forthcoming) weigh these cohort-specific effects by the distribution of cohorts and the relative time indicators.

activities to be less difficult than those without access. We also show that the effects of internet access on school performance appear somewhat larger for schools with high student to teacher ratios and - in the short run - where teachers have more qualifications, though the differences are not statistically significant. Finally, we turn to school-level mechanisms. We analyze the effects of internet installation on the hiring of specially-trained teachers and find evidence that schools make complementary investments in these types of teachers following internet access. We recognize that we are unable to explore all potentially important mechanisms due to data limitations; notably, we lack data on student-level mechanisms such as attendance and engagement and teacher-level mechanisms such as classroom practices.[39]

4.1 Student-level Mechanisms

4.1.1 Does access to internet in schools increase internet use among students?

We show that internet access in school effectively translates into increased use by students. To do this, we turn to the Peruvian National Household Survey (ENAHO) because the CE data do not contain any information on individual students or internet usage. We construct a repeated cross section of all primary school students enrolled in public schools from the 2007-2016 rounds of the ENAHO. The ENAHO collects two crucial pieces of data. First, it collects individual information about internet usage during the 30-day period previous to the survey. Second, it gathers household GPS locations (village location in rural areas or the centroid of the neighborhood block in urban areas). While the ENAHO does not elicit the particular primary school that each student attends, we can infer the school that a student is most likely to attend by assigning each student enrolled in a public primary school to the nearest public primary school.[40] In the ENAHO data, we observe over 55,000 students that are likely to attend 4,693 of the 23,318 schools (20.1%) in our main estimation sample.[41]

We then estimate our event study specification (Equation 1) with an indicator variable for students' internet use (from the ENAHO) as the dependent variable. We regress this variable on event time dummies based on access to internet in the student's closest school (from the CE data). We include the same set of school control variables as in Section 2.1 and add household and individual controls (student's age, grade, native language, household size, and the age, sex, and

[39] We find that dropout and grade completion are not systematically related to internet access (results available upon request). This is consistent with Cristia, Czerwonko and Garofalo (2014), who similarly find no effects of access to technology on dropout in secondary schools.

[40] In matching students to schools, we match only if there is a public primary school within 10km of the household and that is operating in the current survey year. Using this method, the median distance between a student and his/her matched school is 0.33 km. The estimates below are robust to increasing the maximum distance to 20 km or reducing it to 5 km.

[41] The reason we are only able to match 20.1% of schools is that our school data cover the universe of schools, while the household survey covers only a random sample of students each year; by its nature, it will not cover students from all schools in the ECE and CE data. In fact, the ENAHO samples about 0.4% of households in Peru in a given year. Even with this low coverage of households, we are able to match over 20% of schools.

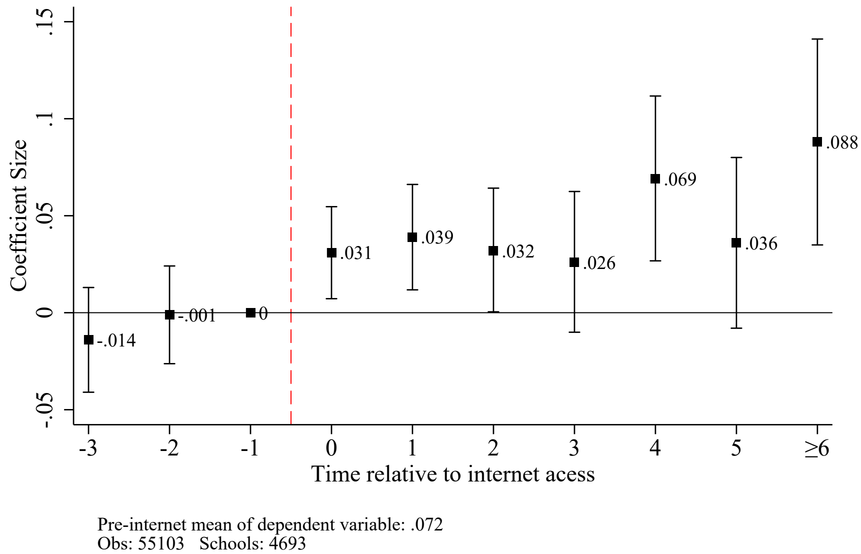

Figure 3: Impact of Internet Access at Schools on Students' Internet Use

The above figures uses data from household surveys (ENAHO), 2007-2016. Control variables include sex, class size, indicator variables for the number of 2nd grade classes at the school, total school enrollment, number of second grade students that took the ECE, facilities (computer room, library, administrative offices), resources per student (classrooms, computers, and teachers), UGEL-specific enrollment tercile by year fixed effects, and school fixed effects. The estimation sample includes all sampled students attending public primary schools that gain internet after 2007 or remain unconnected by 2020. Standard errors are clustered by school.

education of the household head). Importantly, our specification also includes school fixed effects. Thus, our event study coefficients capture whether the changes in internet access *within* schools increase internet usage among students likely to attend them.

We present these estimates in Figure 3. We find that students become much more likely report using the internet after their nearest public school has gained access to the internet. We see no evidence of pre-trends with respect to student internet access. After 6 or more years of in-school access, students' probability of having used the internet in the last thirty days increases by 8.8 percentage points, or 122% over the mean prior to school-based internet connections. This suggests that increased use of internet might be able to explain a portion of the improvements in student performance that we document in Section 2.

4.2 Teacher-level Mechanisms

To complement these results, we present suggestive evidence based on descriptive statistics from two nationally representative surveys of schools in Peru: the 2014 National Survey of Educational Institutions (ENIE) and the 2014 National Survey of Teachers (ENDO). These surveys directly interview teachers and include information about their use of internet use in classrooms and their perceptions of the advantages of ICTs in education. Unfortunately, neither survey provides school identifiers linking these surveys to the CE or ECE. However, they allow us to characterize teachers' approaches to internet use in Peruvian schools by the end of our period of analysis.

The ENIE suggests a considerable degree of student exposure to internet in the classroom by 2014: among those who use it, 76% report using the internet at least once per week (and, on average, for 1.85 hours per week). Information in the ENDO also supports the notion of regular internet use in classrooms. As of 2014, 31% of Peruvian public primary school teachers at internet-connected schools listed internet-connected computers among the top 3 most-used classroom tools. Among schools with internet access, only very basic materials (such as photocopies and flip charts) score higher in this ranking, relative to internet-connected computers. Moreover, teachers view the internet as critical for the success of their students. 83.2% of second grade teachers in public schools with internet access believe internet increases students' access to information that is otherwise unavailable to them and 81.7% state that it improves collaborative learning among students. Among primary school teachers, the most often-cited school factor considered to negatively affect student learning is lack of access to new technologies, including the internet.

4.2.1 How do teachers use the internet?

On the other hand, student use of internet-equipped computers is only one way in which internet access could benefit students; it also can serve as a tool for teachers themselves. For example, teachers can access "off the shelf" lesson plans, repositories of practice questions, instructional aids, etc. According to Sandro Marcone (the former director of *Plan Huascarán*), what teachers demanded most from these programs was not only classroom resources themselves but lesson plans, saying essentially 'tell me what to teach' (Balarin (2013)). In the 2014 ENDO, 63% of teachers at internet-equipped primary schools considered access to internet and technology a top 3 factor in enhancing teaching performance— a larger proportion than those who listed either reference materials (38%) or networking with colleagues (38%). Over 64% of teachers at internet-connected public primary schools use the internet to access virtual pedagogical courses; the only other category with higher reported use is for email correspondence. Furthermore, school-based internet may be very important for teachers who otherwise might not have access to the internet. Among public primary school teachers, nearly half (44%) do not have internet access at home.

We use the data in the ENDO to estimate some suggestive correlations of teachers' ease to perform certain tasks and their access to internet. Public primary school teachers that have access to school-based internet in "good condition" report that teaching activities are less difficult than those without access, even conditional on teacher and school observable characteristics (Appendix Table A.14). These activities include communicating with and motivating students, selecting and making good use of methodology and materials, using time effectively in the classroom, teaching according to different levels of student learning, and addressing the academic problems of students. Thus, teacher access to online materials may boost student performance above and beyond the impacts of direct student use of internet-connected computers.

4.2.2 Heterogeneity by class size

We next examine heterogeneity along the lines of class size, following Barrow, Markman and Rouse (2009). One might expect that the effects of internet (and other ICTs) may be stronger in larger classes for a variety of reasons. For example, previous work suggests that ICTs may provide students with more individualized attention than they would otherwise receive from teachers. If ICTs reduce the time teachers spend in group activities (for which internet might aid through interactive tools), they might be able to increase the time they allocate to individualized instruction. In particular, teachers assigned to larger classes might be more constrained in providing individualized instruction, and thus may be expected to see larger gains from ICTs.

Alternatively, internet access might be especially useful in strengthening the effectiveness of group work. For example, children may be more likely to focus on a group learning activity that involves watching a video or playing an educational game online than more traditional paper- or text-based activities. As mentioned, the overwhelming majority (81.7%) of teachers in internet-connected schools believe that internet access enhances collaborative learning among students (ENDO, 2014). This implies that internet connectivity would be very useful in classrooms with many students, where teachers rely heavily on group work. For both of these reasons, we expect internet access to matter more for schools with high versus low student to teacher ratios (STR).

Splitting schools by the pre-internet median STR, we find that the positive effects of internet access are concentrated among schools with high STRs. We define "high STR" and "low STR" groups as follows. First, we calculate the total number of teachers per second grade student (we do not use the number of teachers exclusively dedicated to second grade, because many smaller schools assign teachers to multiple grades). Then, we calculate each school's pre-internet average STR (time-invariant); this includes all observations for schools that remain unconnected by the end of the CE sample period (2020). Finally, we divide the schools into high and low STR groups based on having a pre-internet average STR above or below the median. In Figures 4a and 4b, the high and low STR trends in test scores prior to internet access are similar, but diverge once internet is introduced. In low STR schools, the effects are much smaller (and close to zero for reading). From Figures 4a and 4b, we can see that the 95% confidence intervals of the event study indicators for low and high STR schools overlap. However, the overall pattern seems to suggest larger average gains for high STR schools. We test this in our trend break specification. Columns 1 and 2 Appendix Table A.15 confirm that the level shift and trend break in test scores are larger in high STR schools (though not statistically significant so). Thus we take these results as suggestive that the effects of internet are overall stronger in larger classes.

4.2.3 Heterogeneity teacher qualifications

Another possibility is that ICTs generate gains in student learning because they compensate for the lack or low quality of other inputs. For example, Jackson and Makarin (2018), determine that the benefits — in terms of math achievement — of providing teachers with online access to "off the shelf" lesson plans were larger among students with weaker teachers. Relatedly, some have found

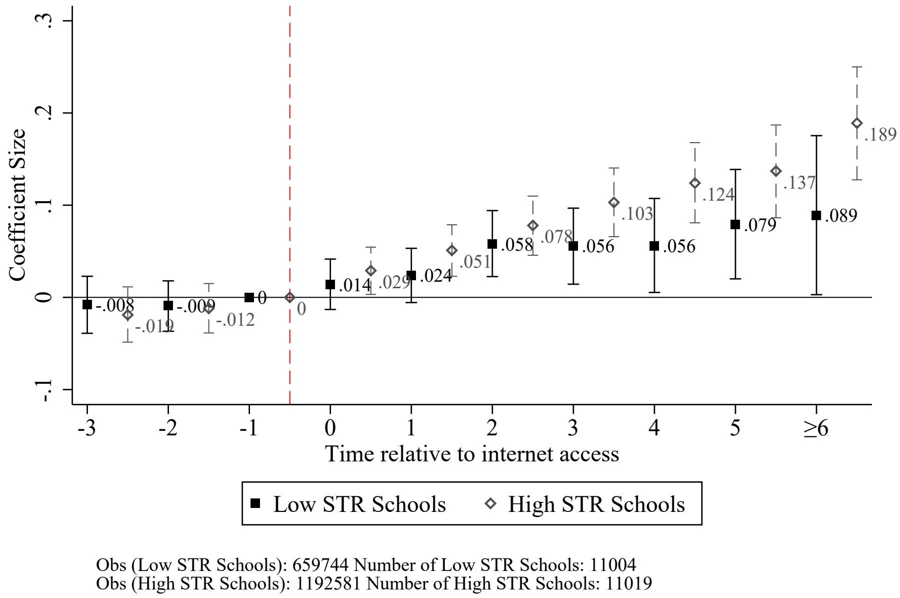

(a) Math

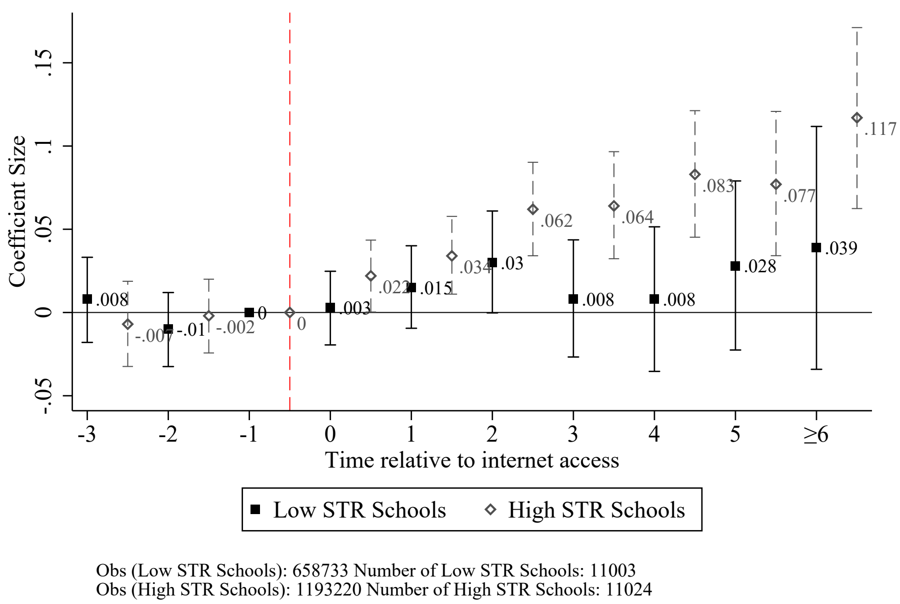

(b) Reading

Figure 4: Heterogeneity by Student to Teacher Ratios: Event Study Results

For each of the above figures, the sample is split based on each school's pre-internet average ratio of second graders to total teachers (STR) — high (low) STR schools fall above (below) the median pre-internet STR. The pre-internet median is calculated using all schools, including ones that do not have internet by 2020. Coefficients capture the increase in test scores relative to the year prior to a school receiving internet access ($t = -1$). Note that due to the timing of the *Censo Escolar* relative to the ECE exam, some schools receive internet access in $t = 0$ while some receive it in $t = 1$. Control variables include sex, class size, indicator variables for the number of 2nd grade classes at the school, total school enrollment, number of second grade students that took the ECE, facilities (computer room, library, administrative offices), resources per student (classrooms, computers, and teachers), UGEL-specific enrollment tercile by year fixed effects, and school fixed effects. The sample includes all grade 2 students in all public schools that gain internet between 2007 and 2020 or remain unconnected by 2020. Standard errors are clustered by school.

that the success of ICT interventions may depend on whether they displace traditional instruction or constitute additional learning activities outside of traditional classroom hours (as part of an after school tutoring program, for example, as in Linden 2008). In cases where ICTs substitute for traditional instruction, impacts may depend on the quality of instruction that the new technology is displacing. Such is hypothesized in Bulman and Fairlie (2016, p. 20), "[...] Interestingly, evidence of positive effects appears to be the strongest in developing countries. This could be due to the fact that the instruction that is being substituted for is not as of high quality in these countries."

Conversely, ICTs could be more effective when teachers are more highly qualified if, for example, teachers with post-secondary degrees may be better able to adapt to and use new technologies. Using the 2014 ENDO, we observe that more qualified teachers (e.g. those holding a college degree or higher) are more likely to report internet-connected PCs as one of the three most often used classroom tools than those without a college degree (34% versus 28%).

To shed some light on which effect is stronger, we examine heterogeneity in results by the level of qualifications that a school's teachers have obtained. In Figures 5a and 5b, we see that there is very little difference in the estimated effects across schools with low and high teacher qualifications. Here, we measure teacher qualification as the per student number of teachers with a pedagogical or university degree. We estimate the average ratio of qualified teachers-to-students by school using the pre-internet period (including all observations from schools that remain unconnected by 2017), and split the sample in two groups based on the sample median across schools. Those with ratios above (below) the sample median are classified as schools with "high" ("low") teacher qualifications. The only noticeable distinction is that the short-run gains in test scores (within the first 3 years of access) appear to be higher when schools have teachers with more qualifications; this can also be seen in columns 3 and 4 of Appendix Table A.15), where the level-shift appears larger for high qualification schools (though not statistically significantly so). This is consistent with the possibility that teachers with more qualifications are better able to immediately use the internet efficiently, while it takes longer for teachers without these qualifications to do so. However, after year 4, gains are similar across schools with high and low teacher qualifications.

4.3 School-level Mechanisms

4.3.1 Complementary investments in trained teachers

One explanation for why we observe delayed impacts of internet access may be that schools require teachers with digital and internet skills in order to incorporate the new technology into the classroom. To investigate this possibility, we study whether schools respond to internet access by hiring teachers with expertise in "computer and information technology." This includes both teachers trained to teach computer skills, as well as teachers who themselves underwent advanced education relating to computers; hereafter, these are referred to as "computer teachers." We estimate equation 1 using an indicator for the presence of a computer teacher as the outcome.

Figure 6 shows that internet access is accompanied by a steady increase in the likelihood a school has a computer teacher that levels off 2 years post-internet installation; by year 4, this

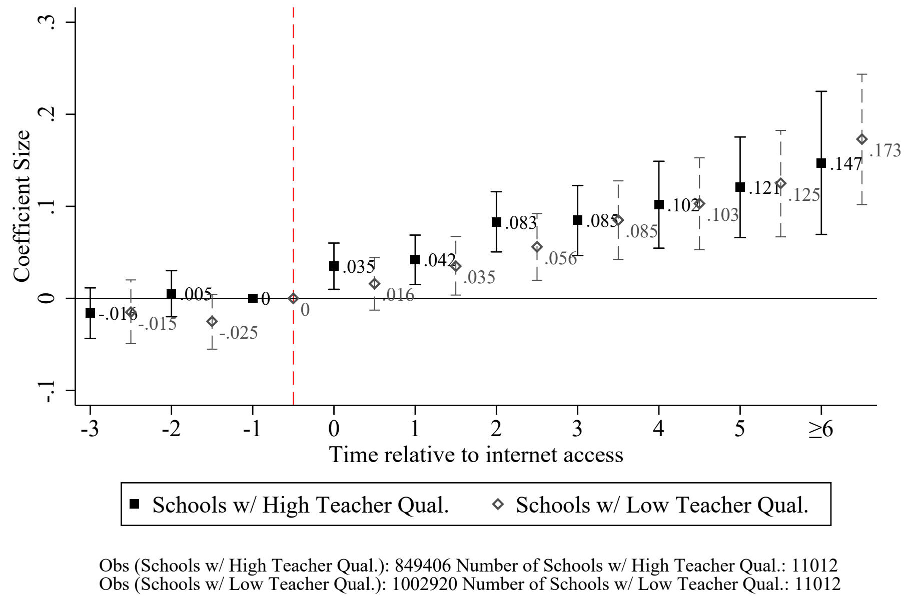

(a) Math

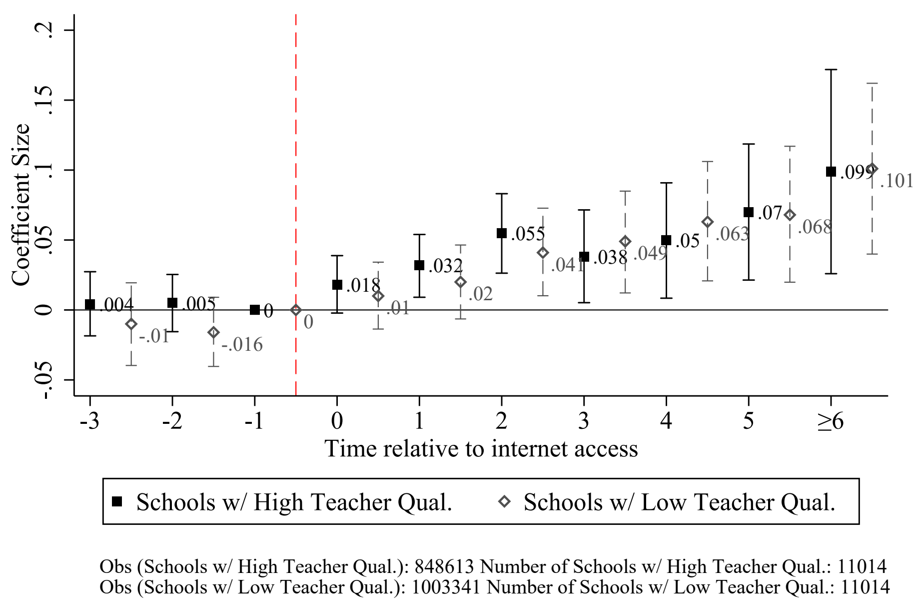

(b) Reading

Figure 5: Heterogeneity by Teacher Qualifications: Event Study Results

For each of the above figures, the sample is split based on each school's pre-internet average number of teachers with a pedagogical or higher education degree per student over the sample period relative to the median of all schools' sample averages. The pre-internet median is calculated using all schools, including ones that do not have internet by 2020. Coefficients capture the increase in test scores relative to the year prior to a school receiving internet access ($t = -1$). Note that due to the timing of the *Censo Escolar* relative to the ECE exam, some schools receive internet access in $t = 0$ while some receive it in $t = 1$. Control variables include sex, class size, indicator variables for the number of 2nd grade classes at the school, total school enrollment, number of second grade students that took the ECE, facilities (computer room, library, administrative offices), resources per student (classrooms, computers, and teachers), UGEL-specific enrollment tercile by year fixed effects, and school fixed effects. The sample includes all grade 2 students in all public schools that gain internet between 2007 and 2020 or remain unconnected by 2020. Standard errors are clustered by school.

results in a 55% increase in the probability that a school has a computer teacher over the pre-internet likelihood. When taken together, the findings for computer teachers and test scores are consistent with the idea that schools may need time to make complementary investments to fully exploit new classroom technologies, such as teachers with computer training. However, because the presence of a computer teacher is a function of internet access (and is not exogenously given), we are not fully able to show the impact of the complementarity of both inputs on school performance.[42]

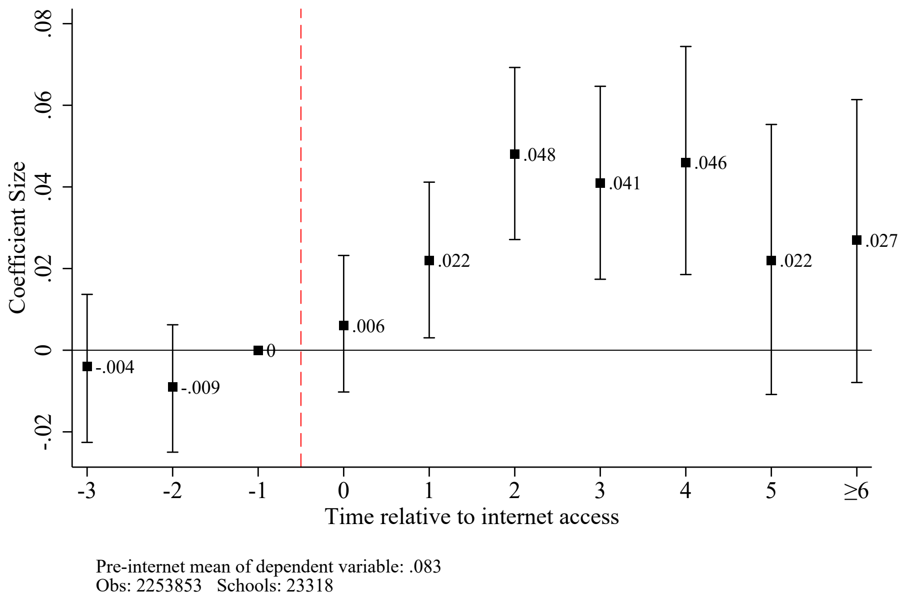

Figure 6: Internet Access and Presence of a Computer Teacher

Coefficients capture the increase in the likelihood of having a computer teacher on staff relative to the year prior to a school receiving internet access ($t = -1$). Note that due to the timing of the *Censo Escolar* relative to the ECE exam, some schools receive internet access in $t = 0$ while some receive it in $t = 1$. Control variables include sex, class size, indicator variables for the number of 2nd grade classes at the school, total school enrollment, number of second grade students that took the ECE, facilities (computer room, library, administrative offices), resources per student (classrooms, computers, and teachers), UGEL-specific enrollment tercile by year fixed effects, and school fixed effects. The sample includes all grade 2 students in all public schools that gain internet between 2007 and 2020 or remain unconnected by 2020. Standard errors are clustered by school.

5 Conclusions

We find evidence that the introduction of internet to Peruvian primary schools produces economically meaningful improvements in student performance as measured by standardized test scores

[42]For example, we do not estimate equation 2 including a triple interaction between post-internet access, event time, and presence of a computer teacher because as shown in Figure 6, schools hire computer teachers *as a lagged response* to gaining internet access.

for grade 2. Gains increase over time, growing from 0.017-0.028 standard deviations in the year of installation to 0.063-0.110 standard deviations 5 years after installation (depending on the subject). Importantly, there are no apparent pre-existing trends in test scores prior to internet access, suggesting little role for reverse causality. Using a trend break specification, we confirm that there is a trend break in test scores that occurs at the time of internet access. In the medium term, the yearly gain in test scores is about 0.012-0.015 standard deviations. These results, based on over 2.2 million students from a large panel dataset of more than 23 thousand schools, are robust to a number of potential confounding factors, including changes in sample composition with respect to either schools or students, changes in school resources, and endogenous timing of installation with respect to prior trends in test performance. In our setting, the nationwide scale of roll-out, large sample of students and schools, and extended time frame uniquely enable the analysis of this technology's application at the farthest-reaching level of policy.

On the one hand, previous research on ICTs has found that providing ICT hardware with few or no complementary learning tools has little immediate impact on student performance (Bet, Ibarraran and Cristia 2014; Barrera-Osorio and Linden 2009; Cristia et al. 2017; etc.). Our short run results (based on up to 1 year after internet installation) confirm that any effects of school-based internet access are small in magnitude — and thus perhaps impossible to detect in smaller samples of schools. On the other hand, medium run gains are sizable, pointing towards the necessity of a longer evaluation window for understanding the effectiveness of ICT interventions. Ultimately, our estimated effects of internet access between the second year and more than five years after internet installation still fall below prior estimates of the impact of computer assisted learning and instruction. While school-based internet does not fully confer the benefits of individualized pedagogical tools, it may provide access to learning resources that are otherwise unavailable to many students in developing countries.

Interestingly, we find that providing internet access in schools is comparable to other interventions in terms of cost effectiveness. Specifically, we find that the cost per student of raising test scores by 0.01 standard deviations ranges between \$0.60 and \$5.20, depending on the subject and the assumptions about the cost of internet installation.[43] This range makes school-based internet

[43] We calculate the lower- and upper-bound cost effectiveness of internet access in Appendix Table A.16. The upper bound of costs are based on costs of two recent large contracts (in 2016 and 2019) that would expand internet access in public schools in Peru, according to the Ministry of Education. There are two reasons why this might overestimate the per-student cost of internet installation for the schools in our sample. First, these contracts reflects the costs of internet connection for those schools who had not received access at least by 2016 (or 2019). It is reasonable to assume that schools that got connected earlier would have had lower per-student connection costs because they were more urban and had higher enrollments (Appendix Table A.2). Second, the recent contracts established that in many cases the service would include a combination of optic fiber connections and other more expensive types of access (i.e., through satellites, radio links, etc.). These more expensive connections likely reflect that schools with recent connections were in more distant and inaccessible areas. The lower bound of costs are based on residential connections with a bandwidth of 500 megabytes per second. These costs are based on Telefonica del Peru's rate, as advertised in their webpage (www.movistar.com.pe) during October 2021. Telefonica del Peru is the largest internet provider in the country and, by November 2019, concentrated 70% of all landline internet access in Peru. While such type of connection can potentially provide access to several computers in classrooms, it is likely to underestimate the

neither the most or least cost-effective within a wide variety of educational interventions in a number of different settings (see Appendix Table A.17). With future potential technological advances leading to reductions in the cost of internet provision, the cost effectiveness of this policy could increase.

Based on two additional nationally representative surveys of households and teachers, we present evidence about students' and teachers' use of internet. On one hand, internet connection in schools does make students more likely to directly access this tool. On the other, it appears that internet helps teachers increase their effectiveness through access to additional online teaching resources (e.g., off-the-shelf materials, teaching materials, etc.). Gains in test scores are somewhat concentrated among schools that have high student-teacher ratios. Hence, school-based internet may generate important gains in learning particularly when the number of teachers per student is constrained below the optimum. We also find that the short-run improvements in test scores are higher in schools where more teachers hold pedagogical or university degrees (though not statistically so), suggesting that teacher qualifications may be an important complementary input to new technology. However, we find that, over time, schools with lower teacher qualifications are able to benefit from access to internet as well.

We provide supporting evidence that achievement gains are slow to emerge because schools need time to adapt to new technologies. Specifically, after installing internet public schools require time to augment their staff with teachers experienced in computers and information technology. We thus concur with several prior studies finding that student achievement begins to increase only as teachers learn to integrate new technology into their curricula (Hull and Duch 2019, Mo et al. 2013, Sprietsma 2007).

However, the interpretation of the results presented is subject to a number of limitations. While we characterize teachers' and students' use of internet through secondary data sources, we are unable to directly incorporate measures of their use of technology as a mediating factor in our analysis. We are largely unable to explore heterogeneity in the effectiveness of school-based internet based on student characteristics. Indeed, previous work suggests that individual heterogeneity - especially with regard to initial achievement - significantly determines how technology affects the learning process (Bai et al. 2016; Barrow, Markman and Rouse 2009; Linden 2008; He, Linden and MacLeod 2008; Muralidharan, Singh and Ganimian 2019). Future research on heterogeneous impacts of internet in education could bear broad implications for inequality within and across learning environments.

Perhaps most notably, our results speak only to school-level dynamics and may mask important individual-level dynamics. Thus it is still an open question whether access to school-based technology might produce longer-term effects – for example, at higher schooling levels or later in the labor market. We consider longer-term individual-level studies (such as follow-ups of previous

true connection costs for many schools.

RCTs which evaluated relatively short-term impacts of technology) as an avenue for relevant future research.

References

Araya, Roberto, Elena Arias Ortiz, Nicolas L Bottan, and Julian Cristia. 2019. "Does gamification in education work? Experimental evidence from Chile." IDB Working Paper Series. 2, 14

Bai, Yu, Di Mo, Linxiu Zhang, Matthew Boswell, and Scott Rozelle. 2016. "The impact of integrating ICT with teaching: Evidence from a randomized controlled trial in rural schools in China." *Computers & Education*, 96: 1–14. 2, 31

Balarin, María. 2013. "Las pol'íticas TIC en los Sistemas Educativos de América Latina: Caso Perú." Programa TIC y Educación Básica, UNICEF, Buenos Aires, Argentina. 6, 24

Bando, Rosangela, Francisco Gallego, Paul Gertler, and Dario Romero. 2016. "Books or Laptops? The Cost-Effectiveness of Shifting from Printed to Digital Delivery of Educational Content." *NBER Working Paper*, , (22928). 2, 14

Banerjee, Abhijit V., Shawn Cole, Esther Duflo, and Leigh Linden. 2007. "Remedying Education: Evidence from Two Randomized Experiments in India." *Quarterly Journal of Economics*, 122(3): 1235–1264. 2, 14

Barrera-Osorio, Felipe, and Leigh L. Linden. 2009. "The Use and Misuse of Computers in Education: Evidence from a Randomized Experiment in Colombia." The World Bank Policy Research Working Paper 4836, Washington, DC. 2, 5, 20, 30

Barrow, Lisa, Lisa Markman, and Cecilia Elena Rouse. 2009. "Technology's Edge: The Educational Benefits of Computer-Aided Instruction." *American Economic Journal: Economic Policy*, 1(1): 52–74. 25, 31

Beg, Sabrin A, Adrienne M Lucas, Waqas Halim, and Umar Saif. 2019. "Beyond the Basics: Improving Post-Primary Content Delivery through Classroom Technology." *NBER Working Paper*, , (w25704). 14

Belo, Rodrigo, Pedro Ferreira, and Rahul Telang. 2014. "Broadband in School: Impact on Student Performance." *Management Science*, 60(2): 265–282. 2

Bet, German, Pablo Ibarraran, and Julian Cristia. 2014. "The Effects of Shared SchoolTechnology Access on Students' Digital Skills in Peru." Inter-American Development Bank IDB Working Paper Series 476, Washington, DC. 2, 5, 20, 30

Beuermann, Diether W., Julian Cristia, Santiago Cueto, Ofer Malamud, and Yyannu Cruz-Aguayo. 2015. "One Laptop per Child at Home: Short-Term Impacts from a Randomized Experiment in Peru." *American Economic Journal: Applied Economics*, 7(2): 53–80. 2, 5, 7, 20

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2021. "Revisiting event study designs." *Unpublished Manuscript*. 13

Brynjolfsson, Erik, and Lorin M Hitt. 2000. "Beyond computation: Information technology, organizational transformation and business performance." *Journal of Economic perspectives*, 14(4): 23–48. 3

Bulman, George, and Robert W. Fairlie. 2016. "Technology and Education: Computers, Software, and the Internet." *NBER Working Paper*, , (22237). 2, 27

Callaway, Brantly, and Pedro HC Sant'Anna. forthcoming. "Difference-in-differences with multiple time periods." *Journal of Econometrics*. 21

Carrillo, Paul, Mercedes Onofa, and Juan Ponce. 2010. "Information Technology and Student Achievement: Evidence from a Randomized Experiment in Ecuador." Inter-American Development Bank IDB Working Paper Series 223, Washington, DC. 2, 14

Cristia, Julian, Alejo Czerwonko, and Pablo Garofalo. 2014. "Does Technology in Schools Affect Repetition, Dropout and Enrollment? Evidence from Peru." *Journal of Applied Economics*, 17(11): 89–111. 2, 22

Cristia, Julian, Santiago Cueto, Pablo Ibarraran, Ana Santiago, and Eugenio Severin. 2017. "Technology and Child Development: Evidence from the One Laptop per Child Program." *American Economic Journal: Applied Economics*, 9(3): 295–320. 2, 5, 7, 19, 20, 30

de Melo, Gioia, Alina Machado, Alfonso Miranda, and Magdalena Viera. 2013. "Impacto del Plan Ceibal en el Aprendizaje. Evidencia de la Mayor Experiencia OLPC." Instituto de Economía, Facultad de Ciencias Económicas y de Administración, Unversidad de la República Serie Documentos de Trabajo DT 13, Montevideo, Uruguay. 2

Escueta, Maya, Vincent Quan, Andre Joshua Nickow, and Philip Oreopoulos. 2017. "Education Technology: An Evidence-based Review." *NBER Working Paper*, , (23744). 1, 2

Faber, Benjamin, Rosa Sanchis-Guarner, and Felix Weinhardt. 2015. "ICT and Education: Evidence from Student Home Addresses." *NBER Working Paper*, , (21306). 2

Gibson, Susan, and Dianne Oberg. 2004. "Visions and Realities of Internet Use in Schools: Canadian Perspectives." *British Journal of Educational Technology*, 35(5): 569–585. 2

Glewwe, Paul, and Michael Kremer. 2006. "Schools, Teachers, and Education Outcomes in Developing Countries." In *Handbook of the Economics of Education*. Vol. 2, , ed. Erik A. Hanushek and Finis Welch, Chapter 16, 945–1017. North Holland. 1

Goodman-Bacon, Andrew. forthcoming. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*. 13, 21

Goolsbee, Austan, and Jonathan Guryan. 2006. "The Impact of Internet Subsidies in Public Schools." *The Review of Economics and Statistics*, 88(2): 336–347. 2

He, Fang, Leigh L Linden, and Margaret MacLeod. 2008. "How to Teach English in India: Testing the Relative Productivity of Instruction Methods within the Pratham English Language Education Program." 2, 14, 19, 31

Hjort, Jonas, and Jonas Poulsen. 2019. "The arrival of fast internet and employment in Africa." *American Economic Review*, 109(3): 1032–79. 5

Hopkins, Alvaro. 2014. “Internet in Schools. The effect on educational performance in Peru: 2007-2011.” 4, 5, 13

Hull, Marie, and Katherine Duch. 2019. “One-to-one technology and student outcomes: Evidence from Mooresville’s digital conversion initiative.” *Educational Evaluation and Policy Analysis*, 41(1): 79–97. 3, 14, 31

Instituto Nacional de Estadistica. 2007-2016. “Encuesta Nacional de Hogares.” The data can be downloaded here: http://iinei.inei.gob.pe/microdatos/. Accessed September 2018.

International Telecommunication Union. 2014. “Final WSIS Target Review: Achievements, Challenges and the Way Forward.” Partnership on Measuring ICT for Development, Geneva, Switzerland. 1

Jackson, C Kirabo, and Claire Mackevicius. 2021. “The distribution of school spending impacts.” *NBER Working Paper*, , (w28517). 3, 14

Jackson, Kirabo, and Alexey Makarin. 2018. “Can online off-the-shelf lessons improve student outcomes? Evidence from a field experiment.” *American Economic Journal: Economic Policy*, 10(3): 226–54. 1, 25

Levin, Douglas, and Sousan Arafeh. 2002. “The Digital Disconnect: The Widening Gap between Internet-Savvy Students and their Schools.” American Institutes for Research, Washington, DC. Report prepared for the Pew Internet and American Life Project. 1

Linden, Leigh L. 2008. “Complement or substitute?: The effect of technology on student achievement in India.” World Bank InfoDev Working Paper 22928, Washington, DC. 2, 14, 27, 31

Machin, Stephen, Sandra McNally, and Olmo Silva. 2007. “New Technology in Schools: Is There a Payoff?” *The Economic Journal*, 117(522): 1145–1167. 2

Malamud, Ofer, and Cristian Pop-Eleches. 2011. “Home computer use and the development of human capital.” *The Quarterly Journal of Economics*, 126(2): 987–1027. 2

Malamud, Ofer, Santiago Cueto, Julian Cristia, and Diether W. Beuermann. 2019. “Do Children Benefit from Internet Access? Experimental Evidence from Peru.” *Journal of Development Economics*, 138: 41–56. Mimeo. Inter-American Development Bank. 2

Meza-Cordero, Jaime A. 2017. “Learn to Play and Play to Learn: Evaluation of the One Laptop per Child Program in Costa Rica.” *Journal of International Development*, 29(1): 3–31. 2

Ministry of Education. 2004. “UGEL-level quotas for Project Huascaran.” The data can be downloaded from Annex 1 here: http://www.minedu.gob.pe/normatividad/directivas/Dir083VMGP2003.php. Accessed September 2018.

Ministry of Education. 2006-2020*a*. “Censo Escolar.” Data for 2006-2016 can be downloaded from http://sistemas02.minedu.gob.pe/anda/archivos_censales/index.html. Data for 2017-2020 can be downloaded from http://escale.minedu.gob.pe/uee/-/document_library_display/GMv7/view/10481. Accessed March 2022.

Ministry of Education. 2007-2016. “Evaluacion Censal de Estudiantes.” Data access can be accessed through a FOIA request here: https://enlinea.minedu.gob.pe/.

Ministry of Education. 2014. "Encuesta Nacional de Docentes." The data can be downloaded here: http://escale.minedu.gob.pe/uee/-/document_library_display/GMv7/view/4385278. Accessed September 2018.

Ministry of Education. 2018,2020*b*. "Mapa de Escuelas." The map can be downloaded here: http://sigmed.minedu.gob.pe/mapaeducativo/. We downloaded a version of this map in September 2018 and an updated version in January 2021.

Mo, Di, Johan Swinnen, Linxiu Zhang, Hongmei Yi, Qinghe Qu, Matthew Boswell, and Scott Rozelle. 2013. "Can One-to-One Computing Narrow the Digital Divide and the Educational Gap in China? The Case of Beijing Migrant Schools." *World Development*, 46: 14–29. 2, 5, 20, 31

Mo, D., L. Zhang, J. Wang, W. Huang, Y. Shi, M. Boswell, and S. Rozelle. 2015. "Persistence of learning gains from computer assisted learning: Experimental evidence from China." *Journal of Computer Assisted Learning*, 31(6): 562–581. 3, 5, 14

Muralidharan, Karthik, Abhijeet Singh, and Alejandro J Ganimian. 2019. "Disrupting education? Experimental evidence on technology-aided instruction in India." *American Economic Review*, 109(4): 1426–60. 2, 14, 31

National Public Radio. 2012. "One Child, One Laptop ... And Mixed Results In Peru." *Article for Weekend Edition Saturday on October 11, 2012.*, Last accessed March 2016. 5

OECD. 2015. *Students, Computers and Learning: Making the Connection.* OECD Publishing. 2

One Laptop per Child. 2011. "Internet Connectivity: The Achilles Heel of the Peru Deployment." Last accessed March 2016. 5

One Laptop per Child. 2016. "2016 Annual Report: A Year of Growth." 1

Purcell, Kristin, Alan Heaps, Judy Buchanan, and Linda Friedrich. 2013. "How Teachers are using Technology at Home and in their Classrooms." PEW Research Center, Washington, DC. 2

Saavedra, Jaime, and Pablo Suarez. 2002. "El Financiamiento de la Educación en el Perú: El Rol de las Familias." Grupo de Análisis para el Desarrollo (GRADE) Documento de Trabajo 38, Lima, Peru. 6

Sharma, Uttam. 2014. "Can Computers Increase Human Capital in Developing Countries? An Evaluation of Nepal's One Laptop per Child Program." Paper presented at the 2014 Annual Meeting of the Agricultural and Applied Economics Association. 2

Sprietsma, Maresa. 2007. "Computers as Pedagogical Tools in Brazil: A Pseudo-Panel Analysis." Centre for European Economic Research ZEW Discussion Paper 07-040, Mannheim, Germany. 2, 31

Sun, Liyang, and Sarah Abraham. forthcoming. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*. 21

The World Bank. 2016. "World Development Indicators." Accessed August 2017. 6

Toyama, Kentaro. 2015. *Geek Heresy: Rescuing Social Change from the Cult of Technology.* . First ed., New York, NY:PublicAffairs. 2

Trucano, Michael. 2012. "Evaluating One Laptop Per Child (OLPC) in Peru." Post in the EduTech blog. 7

Trucano, Michael. 2016. "Technologies in Education Across the Americas: The promise and the peril and some potential ways forward." *World Bank Education, Technology & Innovation: SABER-ICT Technical Paper*, , (12). 1

UNESCO. 2012. "ICT IN EDUCATION IN LATIN AMERICA AND THE CARIBBEAN: A regional analysis of ICT integration and e-readiness." UNESCO Institute for Statistics 22923, Montreal, Canada. 1

Vigdor, Jacob L, Helen F Ladd, and Erika Martinez. 2014. "Scaling the Digital Divide: Home Computer Technology and Student Achievement." *Economic Inquiry*, 52(3): 1103–1119. 2

World Bank. 2012. "Public Expenditure Review for Peru:Spending for Results." Report No. 62586 - PE. 6

World Bank. 2018. "World Development Report 2018: Learning to Realize Education's Promise." The World Bank, Washington, DC:The World Bank. 1