I. Introduction

Ia. Overview

The sphere of international development policy has become increasingly concerned with the “digital divide”—the widening gap in information and communication technology (ICT) availability between wealthy and developing countries. A growing front in the battle against the digital divide has been the education sector, where many policymakers believe computers can improve the quality of education as well as the efficacy of its delivery. But this faith in the ability of ICTs to improve education outcomes is supported by very sparse evidence. There is little consensus on how ICTs can most effectively be used in developing country schools. The purpose of this thesis is to contribute to the understanding of how ICTs can affect educational outcomes in developing countries, and to generate deeper knowledge about the circumstances in which such technology can be most valuable.

I address these questions through an analysis of three large-scale computer and internet deployment projects in southern Africa. In total, I compiled quantitative data for over 2,000 treatment and comparison schools from two provincial education departments in South Africa as well as Namibia’s Ministry of Education. Using separate sources of data on standardized exam scores, school characteristics and computer deployment timetables, I constructed a similar dataset for each of these projects which I use to estimate the effects of computer availability on student achievement. I also use qualitative data, which I collected through interviews with teachers and principals at a subset of schools in each project, to present a more in-depth analysis of project implementation and the mechanisms through which computer availability may impact the teaching and learning process.

This chapter has three purposes. I will first discuss the growing role of ICT use for educational purposes in development policy, then review the literature relevant to the effects of computer and internet availability on student achievement. Finally, I will provide an overview of the three projects on which I focus my analysis. In Chapter II I estimate the causal effect of computer and internet availability on education in all three projects, generally finding effects that are statistically indistinguishable from zero. I begin that chapter by discussing the several data sources I draw upon to construct the datasets and then present relevant descriptive statistics, outline the identification strategy used and present and discuss my results.

Chapter III presents and analyzes qualitative data collected through school visits, classroom observations and teacher interviews within each project, with the goal of better understanding how computer use can be most successfully and effectively adopted in schools in developing countries. I first discuss literature that is generally relevant to the adoption of change in organizations and then turn to specific literature relevant to the adoption of computer use in schools. The proceeding discussion of the technical, human and organizational factors which affected implementation in the projects I study is then grounded in this literature. The next section of Chapter III describes the mechanisms through which computer use could potentially benefit students. I conclude with a discussion of how the similarities and differences between the three projects affect the success of computer adoption in the schools they serve.

Chapter IV begins with a discussion of how this thesis contributes to the understanding of how computers can improve learning outcomes in developing country schools. I argue that my results demonstrate that computers are generally ineffective as a pedagogical tool but promise to improve student achievement when used as a supplement to traditional classroom learning. I further note that my results are consistent with a previously unnoted broader trend in the body of literature which evaluates the effectiveness of instructional inputs to developing country schools: generally those inputs that are intended to more effectively deliver conventional classroom lessons have failed to improve student achievement, while those used independently of teacher-directed instruction have been very effective. I conclude by discussing how developing country policymakers can apply my findings toward the design of more effective computer deployment projects.

Ib. Education Technology in Development Policy

"Bridging the digital divide" has become an important catchphrase of international development efforts in recent years, and the role of ICTs in development has received an increasing amount of attention from governments and major multilateral organizations. The United Nations, for example, formed its ICT Task Force in 2001 with the goal of "putting [ICTs] at the service of development" (UN ICT Task Force, n.d., para. 2). A significant aspect of efforts to bridge the digital divide has been a push for the introduction of computers into developing country schools. The prominence of public, private and non-profit organizations promoting the use of computers for primary and secondary education in developing countries demonstrates the status of such efforts in the development agenda.

The One Laptop Per Child (OLPC) organization has received substantial attention for its goal to build an affordable laptop for use by students in developing countries. The organization has already begun selling its first batches of laptops in significant numbers to parts of the developing world. The government of Uruguay purchased the organization's first major deployment of 100,000 laptops, and Peru is soon to follow with a reported order of 270,000 (Beer, 2008). These nations alone have easily invested millions of dollars in equipment, and OLPC reports that over a dozen other countries have expressed serious interest in purchasing the machines (OLPC Talks, 2007). The OLPC project strongly emphasizes the ability of its product to improve the quality of education in developing countries, maintaining that "…the emerging world must… [tap into] children's innate capacities to learn, share, and create on their own. Our answer to that challenge is the XO laptop…." (OLPC, n.d., para. 4).

Low-cost laptops designed for educational use in developing countries have also been pursued by major private companies. Intel's Classmate PC has been the most successful venture, with its two largest orders to the governments of Pakistan and Libya totaling 850,000 machines (Finkle, 2007). Much like OLPC, Intel bills its Classmate as potentially improving the quality of education in developing countries. The company states that "Intel-powered classmate PCs are designed to improve education and provide economic opportunities." (Intel, n.d.)

Major international organizations have directed efforts toward supporting the use of computers in education in different ways. The United Nations has decided to support ICT in the education sector specifically through its Global e-Schools and Communities Initiative (GeSCI), which works with ministries of education in developing countries. GeSCI currently operates in five countries where its core objective is to "harness the potential of ICTs to improve the quality of teaching and learning" (GeSCI, n.d.). The World Bank has provided loans to support governments in their pursuit of ICT deployment in basic education projects, currently backing such projects in 10 different countries (World Bank, n.d.). The Bank also founded and leads the Information for Development Program, infoDev, which maintains a significant ICT-for-education component. InfoDev recently released a "knowledge map" on ICT use in education which determined that limited data exist to support the notion that ICTs can have a positive impact on education and that further study is needed. The report also stated that, "It is generally believed that ICTs can empower teachers and learners, promote change and foster the development of 21st century skills…" adding, "It is believed that specific uses of ICT can have positive effects on student achievement when ICTs are used appropriately to complement a teacher's existing pedagogical philosophies." (Trucano, 2005, pp 5-6)

This review of ICT use for education in the policy arena is by no means comprehensive. It serves, however, to provide an overview of the importance this topic is beginning to take on in the realm of development policy. International agencies and developing country governments are increasingly seeking ways to bring computers and related technologies into primary and secondary schools across the developing world, and there is no shortage of organizations willing to sell them the requisite equipment. There also exists a widespread belief that these technologies will improve education quality, but the extent to which such improvements can be realized through computer availability is not at all well understood.

Ic. Evidence on The Effects of ICTs on Student Achievement

The body of literature on the effectiveness of computers in improving student achievement is relatively small. I limit my review in this section to those studies which examine the overall impact of computer use on exam scores and which employ methods that lend themselves well to causal interpretation. There is an important although smaller body of literature on other aspects of computer use in education, and this work is reviewed in Chapter III where I discuss its relevance to the qualitative evidence I present.

Most of the studies evaluating the causal impact of computer use for education have been conducted in relatively developed countries, and evaluate large-scale government programs which aim to increase computer access in schools. These studies suggest that, in most cases, computer use for education is at best ineffective and may actually have a slight *negative* effect on exam scores. Angrist and Lavy (2002) evaluated a computer deployment program of the Israeli government and found that greater computer access increased the amount of computer-aided instruction provided by teachers. But the authors found no evidence of a positive effect of this computer-aided instruction on students' exam scores, and some of their results suggest a slightly negative and statistically significant effect on mathematics scores. Their results are robust to several estimation methods, but questions have been raised regarding the assumptions underlying part of the authors' identification strategy (Glewwe & Kremer, 2006). In general, however, the conclusion that computers had no substantial effect on student achievement is well-supported by the evidence the authors present.

Leuven et al (2007) found similar results in their evaluation of a program in the Netherlands which increased funding for computers to schools with a high proportion of disadvantaged students. The study exploits the policy discontinuity of a government program which provided computer funding subsidies to schools with at least 70 percent minority students. The authors use a regression discontinuity approach to compare treated schools just above the cutoff to untreated schools just below the cutoff. The study estimates the effect of computers on exam scores as negative and in some cases significantly different from zero.

Goolsbee and Guryan (2006) examined the effect of computer and internet subsidies in California's public schools. They found that the resulting increase in computer availability did not correspond to an increase in student achievement—

rather, a small negative effect was measured in many cases, although these estimates are relatively imprecise. The authors use a difference-in-difference identification strategy to compare changes in exam scores between schools which did and did not receive the subsidy, controlling for several other variables. Program subsidies were determined based in part on the demographic characteristics of each school's student body. If these characteristics which determined subsidy rates were correlated with time-trends in school exam results, then the authors' estimates would be biased. Since the authors present no evidence that systematic differences in preexisting trends do not exist in the data, these results should be interpreted with some caution.

A fourth study evaluating a large-scale government program in the United Kingdom provides an important exception to the literature. Machin and others (2007) exploit a change in ICT funding allocations to Local Education Authorities (LEAs) to design an instrumental variable strategy for estimating the effect of ICT availability on student achievement. Prior to the change in allocations, schools would bid for ICT funds. Anecdotal evidence suggests that those schools which proposed innovative uses of ICTs received preference. After the change, funding was allocated according to an objective formula based on characteristics such as LEA population and population density. The authors compare "winners" and "losers" under the new funding system and find moderately positive and statistically significant effects on exam scores among primary schools. Evidence is presented that there are no systematic differences between "winners" and "losers" related to several observable characteristics, including preexisting exam score time-trends.

There are a few important points to consider when evaluating the validity of Machin and others' results. The first is potential problems related to the instrumental variable. The comparison group under their identification strategy essentially consists of schools which lost funding, with the instrument naturally weighting observations according to the magnitude of lost funds. An ideal comparison group would consist of schools which were altogether untreated. Part of the effects picked up by the instrument could therefore be that of reducing funding to the "losing" schools below what they previously expected. For example, schools winning large bids with strong ICT programs may feel compelled to redirect funding from other efforts after the policy change. The converse could also be true, with "winners" under the funding scheme diverting any general funds already being used for ICT programs toward other expenditures. Since convincing arguments can be made that the funding scheme change could have affected student achievement through several channels, the authors' estimates may not reflect the effect of increased ICT availability alone.

An important trend to note in these four studies is that the effect of computers on performance tends to be more positive for younger students. Both the Angrist and Lavy and Goolsbee and Guryan studies look at scores for younger and older students separately and find more positive effects on younger students—Angrist and Lavy observe scores for 4^{th} and 8^{th} grade students in Israel and Goolsbee and Guryan make use of results for both primary and secondary schools in California. The Leuven study looks only at 8^{th} grade students and finds negative effects while Machin and others use outcomes of students in the equivalent of 5^{th} grade and find positive effects. This is important to keep in mind when considering the results presented in this thesis; I study the SchoolNet and Gauteng Online projects with regard to secondary-level exam scores, while in Western Cape I focus on primary-level results.

A second set of studies on the effectiveness of computer instruction takes a different approach—using randomized programs to evaluate how a specific use of educational software affects the achievement of individual students. One such study was conducted in a developed country; Krueger and Rouse (2004) used a randomized design to evaluate the effectiveness of an educational computer program designed to improve language and reading skills. The program was relatively popular in American public schools and its authors claimed that the program's methodologies were "scientifically-based." The authors estimate the effects of using the software on a variety of learning outcomes and find small and statistically insignificant or marginally significant effects across the board.

The single widely circulated and methodologically-sound study of computer use for education in a developing country is similar to this Krueger and Rouse evaluation in that it also examines a targeted use of computers for educational purposes. Banerjee and others (2007) examined a randomized NGO program in India which allowed students access to a computerized mathematics program for two hours weekly. The result was a very large and highly significant increase in student math scores in the first and second years of the program. This obviously contrasts starkly with the evidence on computer use for education from developed countries. The authors attribute this to the fact that the return on investment of computer-assisted learning in developing countries is much larger than in developed countries due to the large disparities in educational resources between schools in developed versus developing nations.

Based on this review of the literature, there is some evidence that computer-aided instruction could be a potential boon for education in developing countries. However, the Banerjee study only provides us with insight into a very specific use of computers to improve learning. It remains an open question whether large-scale computer deployment projects such as those I evaluate in this thesis can have an equally large impact on school performance or if they will be as ineffective as they appear to have been in many developed countries.

Id. Overview of the Projects Studied

I study three different projects in southern Africa which focus on deploying computers to primary and secondary schools to be used for educational purposes. In total, these projects serve roughly 2500 schools and over one million students and have spent the equivalent of tens of millions of US dollars over the past several years, mostly in public funds. Two of the projects are administered by provincial governments in South Africa. The project called Gauteng Online is administered within Gauteng Province by that province's department of education. A second South African project, called Khanya, is run under the auspices of the Western Cape Province's education department. The third project I study is located in Namibia and is run by the non-governmental organization (NGO) SchoolNet Namibia.

Gauteng Province encompasses a largely metropolitan area in the northern part of South Africa. It is constituted primarily by the national capital city, Pretoria, and South Africa's most populous city, Johannesburg, making it the country's geographically smallest but most densely populated province. The Gauteng Department of Education administers the province's nearly 2,000 ordinary public schools which serve over 1.7 million primary and secondary students (South Africa Dept. of Education, 2007). It is South Africa's second-wealthiest province according to census estimates, but the population is characterized by severe income inequality. A disproportionate amount of wealth is concentrated among the province's White citizens who tend to utilize private and top-tier public schools, while Gauteng's majority Black, Coloured[1], Indian and other populations are on average much less wealthy and rely almost exclusively on public education. While significant progress has been made in Gauteng and across South Africa to improve the equity of education, the school system remains very segregated along socioeconomic lines.

The Gauteng Online project began in 2002 as an initiative of the Gauteng Department of Education. Its stated goals include improving the quality of education and developing the province's information technology capacity. Each school participating in the project received an internet-connected lab of 25 computers, and teachers received training on using these computers as an educational tool. In general, the project focused on providing computers to schools which serve poorer communities. Gauteng Online is managed by a small staff which works within the department of education. The project management relies heavily on outsourcing tasks such as training and technical support to private contractors in its administration of the project. To date, over 1200 schools have received computers from Gauteng Online. I restrict my research on the project to schools with a Grade 12, which in practice are almost exclusively secondary schools. (A small number are combined schools, which have both primary and secondary grades.)

Western Cape Province has a smaller school system than Gauteng, with approximately 1450 public ordinary schools serving just under one million students (South Africa Dept. of Education, 2007). While it is in part also constituted by a major metropolitan area—Cape Town and its outlying suburbs—there is also a substantial rural population and a major agricultural sector. Western Cape is South

[1] Coloured is a widely accepted term in South Africa used to describe individuals of mixed racial heritage. Throughout this thesis I use the same racial and ethnic categories which South Africa's Department of Education uses for statistical and demographic purposes.

Africa's wealthiest province and is characterized by the same severe income and racial inequalities as Gauteng. The province's education system similarly faces significant disparities related to the socioeconomic groups served by its schools.

Western Cape's Khanya project states that its goal in providing technology to schools is to improve curriculum delivery through the use of computers as a teaching aid. Much like Gauteng Online, Khanya generally focuses its efforts on providing computers to disadvantaged schools. Unlike in Gauteng, however, the ultimate project inclusion decisions are made by Western Cape Education Department officials in each of the province's seven education districts. Khanya also has staff in each district which include a coordinator, a project manager and several facilitators. District coordinators are responsible for all of Khanya's activities in their area while project managers work to deploy computers and provide training to new project schools. Facilitators are each responsible for training and oversight of project activities in a small group of schools, typically about 10-15 schools each. Schools in the Khanya project typically receive one internet-connected computer lab of approximately 25 to 30 machines. The Khanya project also provides ample educational software to schools and creates its own curriculum content in each of the three major languages used in Western Cape's schools. Khanya has deployed computers to over 800 schools and I focus my research on those Khanya schools with a Grade 3.

Namibia is one of the most sparsely populated countries on earth, home to only two million people but encompassing a landmass almost twice the size of California. Much of the population is spread across rural areas, especially in the northern parts of the country, with the remainder residing in a few large towns or the capital city of Windhoek. In general, schools have much fewer resources than those in Gauteng and Western Cape, with even the most basic educational materials often lacking. There are significant disparities between schools in rural areas, those in township areas and those in more affluent parts of major population centers, both in terms of resources and in the racial composition of the student body. Many of the severe income disparities present in Gauteng and Western Cape are also evident in Namibia and tend to similarly break down along racial lines. Namibia's education system is relatively centralized, with most major policy dictated by the Ministry of Education's head office. According to government statistics, the education system includes almost 600,000 primary and secondary school students across over 1600 schools.

SchoolNet Namibia was incorporated in February of 2000 with the goal of deploying computers to the country's schools. Unlike the Gauteng Online and Khanya projects, SchoolNet sets no specific goals for how it hopes its computers to be used by schools, stating generally that its purpose is to provide youth empowerment through access to technology. The NGO focuses on schools in historically disadvantaged areas and to this end it has developed innovative solutions for serving schools in remote locations. For example, schools without electricity receive solar-powered computer labs and those without access to a telephone line may receive a wireless internet connection. SchoolNet is heavily donor-funded and has received substantial support from the Swedish International Development Cooperation Agency and the United States Agency for International Development. The NGO relies solely on open-source software and makes use of donated computers which are refurbished for use in Namibian schools. SchoolNet maintains several support centers across Namibia and has provided training to its project schools through trainers who reside at new project schools for 2-3 months.

Each of these three computer deployment projects operates in its own context and has taken a unique implementation approach. Having been able to collect similar quantitative data on all three projects and education systems, I use these data to devise a strategy for identifying the causal effect of computer availability on student performance as measured by test scores. Having also collected similar qualitative data at a set of schools within each project, I present an additional analysis of how these different contexts and implementation strategies influenced project success. The results of my analyses, taken together, provide a thorough picture of how the Gauteng Online, Khanya and SchoolNet projects have impacted their participant schools and students. They also help answer the important question of how policymakers in developing countries can implement increasingly prevalent ICT-for-education projects in ways which maximize their effectiveness.