Report Cards: The Impact of Providing School and Child Test-scores on Educational Markets

PRELIMINARY DRAFT

Tahir Andrabi (Pomona College) Jishnu Das (World Bank)
Asim Ijaz Khwaja (Harvard University)*

October 2009

Abstract

Recent evidence on the impact of information provision on service delivery has been mixed, with overall outcomes even worsening in certain cases. We examine the market-wide impact of an intervention that provides school and child-level learning report cards in a randomly selected half of 112 educational markets (villages) in Pakistan. We track all 823 public and private primary schools, 12,000 Grade 3 children, 5,000 teachers and a sample of 1,800 households in these villages. Report card provision improves learning by 0.10 standard deviations and decreases private school fees by 21 percent, with very small changes in school switching and moderate increases in overall enrollment. We argue that providing report cards generates credible competitive pressures on schools to increase price-adjusted quality, with the specific tool used - decreasing prices or raising quality - determined by the nature of production costs and market demand. Consistent with this, we find substantial heterogeneity in the impact across schools. Initially "bad" (below median baseline test scores) private schools respond by increasing quality - showing learning gains of 0.34 standard deviations - or shutting down, but show limited fee changes. In contrast, initially "good" (above median) private schools show no learning gains, but drop fees substantially. Government schools see a tenth of a standard deviation increase in learning. Moreover, we find schools increase investments, while there is little evidence of greater (direct) parental investments. The results show the cost of providing information is similar to the school fee drop, and the intervention likely raised child welfare by increasing learning and lowering educational costs.

*tandrabi@pomona.edu, Pomona College; jdas1@worldbank.org; Development Research Group, World Bank; and akhwaja@ksg.harvard.edu, Kennedy School of Government, Harvard University, BREAD, and NBER. The data used in this paper is part of the multi-year LEAPS project. We are grateful to Alexandra Cirone, Ina Ganguli, Sean Lewis-Faupel, Paul Novosad, and Tristan Zajonc for research assistance and to the LEAPS team for all their hard work. We are also grateful to Abhijit Banerjee, Pinar Dogan, Quy-Toan Do, Matthew Gentzkow, Justine Hastings, Nolan Miller, Rohini Pande, Doug Staiger, and seminar participants at Boston University, BREAD (Duke), Brown, Chicago, Cornell, Dartmouth, NBER, Stanford, Syracuse University, UC Berkeley, UCLA, University of Illinois at UC, University of Wisconsin at Madison, and Yale for comments. All errors are our own.

I Introduction

Providing information is often viewed as a panacea for addressing the poor quality of health and educational services. The World Development Report (2004), for instance, argues that information is one pillar through which the users of a service can hold providers accountable and therefore engender quality improvements. In the United States, the "No Child Left Behind" policy requires every state to regularly report and disseminate information on learning and educational achievement.

Nevertheless, the existing evidence that information "empowers users to demand better services from the government" is mixed. Banerjee et al. (2007) report no learning improvements from information dissemination in an Indian state. In contrast, Björkman and Svensson (2008) argue that health provider report cards led to a sharp decline in infant mortality due to an increase in provider effort. Both Rokoff and Turner (2008) and Chiang (2008) examine the impact of accountability and sanctions on public schools in the US and find evidence of learning gains for low performers. These results relate to the performance of public providers of education and health services. When private players are considered, information may provide incentives that are "too strong." Providers may cream skim by selecting better students, a problem that is exacerbated when information is a noisy measure of true quality (Morris and Shin, 2002). In education, examples from Chile (Urquiola and Mizala, 2007) and the United States (Hastings, 2007) suggest increased sorting; in health, Dranove et al. (2003) show that information about hospital outcomes led to greater sorting and worse outcomes in New York.

We document the experimental impact of information on both the private and the public sector in an equilibrium setting. The context is a large-scale experiment conducted in 112 randomly selected rural communities in the universe of 823 public and private schools in the Punjab province of Pakistan. In half of these communities, we tested and disseminated the results of school and child test scores with follow-up surveys to determine the impact of this intervention on learning, school fees, and household behavior. Since the village is our unit of observation and treatment, we are able to look at the impact of information dissemination not only on particular schools but on the entire universe of schooling options within the village. As we will argue below, given the importance of distance for the choice of schooling in Pakistan, rural communities can be viewed as closed primary schooling markets. Therefore, our results show the average impact of information on the schooling market as a whole.

Test-score information had an impact on both learning and school fees. We document a 0.10 standard deviation gain in average child learning – a third of the average yearly gain experienced by children. Moreover, the gain holds two years after the intervention. In addition, we report a 21 percent drop in the fees of private schools (public schools are free). These results appear not to be driven by selective attrition, either of children or schools, in the treatment sample. Finally, while there is some increase in overall enrollment, we find limited additional switching in treatment compared to control villages with initially poorly performing (private) schools seeing slightly more children switch out. Consequently, the gains are almost entirely attributable to test-score improvements among the 95 percent of children who remained in the same school throughout the observation period.

These findings are consistent with optimal pricing and quality choice in a market with imperfect information. The assumptions are that firms compete for market share in an environment where both price and quality are endogenous, and that information provided under the intervention is an exogenous change in the accuracy of the monitoring technology. Under these conditions, greater information can lead to an increase in price-adjusted quality for all (private) schools. The main intuition is as follows: Increased precision in beliefs regarding school quality effectively makes the demand curve facing any school more elastic (since the parental school quality belief distribution gets "compressed," any given change in price will affect a greater number of parents). Thus without change in price or quality, each school faces an enrollment decline, thus decreasing profits. To preserve enrollment, private schools have to increase price-adjusted quality. The specific tool used by a school – decreasing prices or increasing quality – is determined by the nature of production costs and market demand faced by the schools. In particular, one may expect that price drops would be higher among initially high performing schools to the extent that it is costlier for them to generate further quality improvements. In contrast, initially poorly performing schools may find it more effective to raise quality, both because it may be marginally cheaper for them to do so, and also because lower prices may not be a viable option given a free public sector that offers a lower bound on quality.

Further results provide strong evidence for such heterogeneity in the impact across school types. Private schools with baseline scores below the median show the largest learning gains – well over a third of a child-level standard deviation. Private schools with above-median initial scores show no learning gains but see their fees drop by 23 percent - particularly for those with high fees (relative to their quality). Government schools, which are free, show learning gains of a tenth of a standard deviation. These results are suggestive of increased school-level investments arising from competitive pressures induced by the intervention. We find direct evidence of such increased investments by schools. In contrast, we find little evidence for households directly increasing their educational investments in children.

The empirical findings obtained provide several revealing insights. First, at the heart of our conceptualization is the idea that school quality is a function both of an inherent "type" and of school investments. Where quality is entirely determined by school "type," information may lead to worse outcomes in terms of increased sorting and cream-skimming. Second, our results also suggest a concavity in the quality/effort trade-off: The net benefits of increasing quality decrease at higher quality levels. Good private schools respond by dropping price rather than raising quality, while bad private schools raise quality. This concavity could arise from either limited parental demand or due to the nature of the production function. In support of the latter, our previous work suggests that there are significant supply-side constraints that private schools face in providing quality education. Third, our findings suggest that increased information leads to greater pooling in price-adjusted quality across schools. Thus, after the intervention one is left with a lower quality but free public sector and a higher quality and somewhat more expensive private sector. Finally, an interesting implication of the results is that schools and parents commit to relatively forward-looking arrangements. The reason is that there is limited switching in equilibrium, yet the private sector responds as if it faces greater competition.

A couple of characteristics of the Pakistani educational system make it an interesting site for this experiment. Since the mid-1990s, Pakistan has seen an immense boom in non-religious private schooling driven primarily by the establishment of small "mom and pop"enterprises that face little government oversight or regulation (Andrabi et al., 2008). It is the exponential growth of private schools – rather than the oft-cited but incorrect perception of rising religious schooling – that defines a dramatic change in the educational landscape (Andrabi et al., 2007b). Figure 1 provides an example of what this has done for the majority of the population in the country's largest province, Punjab. Here, we overlay GPS readings for private schools and markets with a Google Earth image of a typical village in our sample. The village is roughly two square kilometers in area but is serviced by eight different schools, public and private. The substantial school choice in this village is no exception; we located and followed students in 823 public and private primary schools in 112 villages, with an average of 7.3 schools per village.

While an expansion of school choice is likely beneficial as long as parents are able to accurately judge the quality of schooling imparted, this is not an easy task. Articles in the press, as well as popular current perception, suggest that these private schools are "fly-by-night" operators who try to entice parents on false grounds and collect costly fees while providing low quality education. Our initial experiences in the field resonated with this view. In one case, a school prominently advertised its computer based learning pedagogy but lacked the electricity needed to charge the one laptop that it had. English medium instruction, considered elite and advertised aggressively, is another example often given by critics of the small private schooling movement. These critics argue that parents have no hope of evaluating the progress of their children given their own illiteracy, much less evaluating schools and making informed investment and school choice decisions for their children. If such concerns are true, information should help.

The need to examine how the information intervention impacts the entire education market also means tracking the investments in every school and following children as they (potentially) switch schools across years. In the absence of a regular information system, particularly for private schools, this is possible only if the educational market can be effectively defined within feasibly narrow geographical boundaries. This was feasible in our context since children in our sample do not travel long distances to school. Our strategy was therefore to locate (and track) all schools in the village or within fifteen minutes walking distance of it. A census of the population's schooling choices suggested that this was sufficient: We found that 92 percent of all children from our villages attended the schools in our sample.

The 8 percent who did not were mostly children attending secondary-level classes, at which point travel times increase. We are thus confident that our results are able to capture the impact of this intervention on the schooling market as a whole.

In addition to providing equilibrium effects of an intervention, we feel that the methodology adopted expands the scope of randomized evaluations in low-income countries in multiple ways. In the three districts studied, the villages were chosen randomly from the entire population of villages with at least one private school. Consequently, the results should be valid for the population of these three districts as a whole. We also made an attempt to survey all the main actors in the production of education. Thus, we bring to bear not only data on learning outcomes of over 12,000 children tracked through three years but also information on and responses of the 5,000 teachers, 823 schools, and 1,800 households who were surveyed throughout the duration of the project. As we will argue below, this allows us to better understand the channels through which information could work.

We are cautiously optimistic about the results thus far. A back-of-the-envelope calculation shows that the costs of providing information for the entire population was comparable to the fee reduction in initially well-performing private schools. If we think of the cost-benefit solely in terms of household welfare, the benefits of increased learning are achieved at (essentially) zero cost. Factoring in the utility of the private schools requires us to parse out the fraction of the cost reduction that was a consequence of informational rents relative to pure transfers; this, however, appears not to be a consideration in most cost-benefit calculations (for instance, the welfare losses to teachers of exerting more effort are seldom accounted for). Simultaneously, the children whose scores improved were from the lower part of the achievement distribution. Gains were highest for initially low-performing private schools, but government schools improved as well. At first glance, this appears to be an interesting case of simultaneously improving equity and efficiency with a single intervention. Given that many parts of South Asia and, increasingly, other countries now resemble the kinds of villages we worked in, the broader applicability of these results is appealing.

Nevertheless, caution is warranted. The entire implementation process was in our hands, and this involved the considerable tasks of designing and implementing tests, holding village meetings to explain and distribute report cards, and tracking children over time. If these tasks are in government hands, the results may be very different. As the example from Chicago public schools shows, cheating is very much a part of high-stakes testing, even in high-income countries (Jacob and Levitt, 2003); in low-income countries, the situation may be worse. Furthermore, the relatively small improvements in government schools suggest that information alone will not be sufficient to bring about large quality changes in the public sector. Our causal estimates of the public-private school differential are of the order of 0.2-0.3 standard deviations a year; this intervention is not going to help close the gap (Andrabi et al., 2008). Our results suggest that information, when there is accountability, leads to better outcomes, rather than information itself creating accountability where there is none. This is consistent with the previous findings reported in Banerjee et al. (2007).

The remainder of the paper is structured as follows: Section 2 provides contextual details on the educational environment and describes the data. Section 3 details the report card intervention and outlines the empirical methodology. Section 4 presents the results on how the intervention affected village-level outcomes. Section 5 provides a theoretical discussion in light of these results and generates predictions on the heterogeneous impact of the intervention. Section 6 then presents evidence for such heterogeneity and Section 7 concludes.

II Background and Data

A. The Context

We highlight aspects of Pakistan's educational system that are of relevance to understanding the context of this paper. It is noteworthy that while Pakistan's educational sector faces substantial challenges, both in terms of enrollment and educational quality, there is a robust and increasing private sector presence. The educational marketplace, even in the rural areas we examine, is often characterized by substantial school choice and competition and by active parental decision-making. We summarize some of these features here.[1]

In 2001, Pakistan's overall performance in education was poor with primary enrollment rates are below those expected for its level of income and worse than several of it's South neighbors (Stern 2001). Large gender discrepancies in the provision of education add to the problem. Less than half of those female children who are eligible for primary school are actually enrolled. Further, there are significant differences in primary enrollment rates by household wealth; rich households are almost four times as likely to send their children to school compared to the poor.

However, while enrollment is a clear problem, quality is an even bigger issue. Learning outcomes, similar to many other developing countries, are quite poor. Based on the achievement tests we conducted, we find that by the end of Grade 3 just over 50 percent of children have mastered the mathematics curriculum for Grade 1. They can add double-digit numbers and subtract single-digit numbers, but they cannot subtract double-digit numbers or tell the time. Both multiplication and division skills have not solidified, and advanced topics such as fractions are beyond all but the best students. In Urdu, they cannot form a sentence with the word "school" or the word "beautiful," and less than 20 percent are able to comprehend a simple paragraph. The situation in English is even worse, with most not being able to recognize simple words.

[1] For a more complete description, refer to Andrabi et al. (2007).

Yet despite these low learning levels, there is substantial variation in quality. Unlike enrollment, where richer and more literate villages are more likely to have children in schools, village attributes have almost nothing to do with learning outcomes. Instead, most of the variation in learning is explained by differences across schools in the same village, and a large portion of this is due to differences across public and private schools. Private schools generally score a standard deviation or more higher on tests compared to government schools, and these differences remain robust to an extensive array of socioeconomic parental and child controls (Andrabi et al., 2007b).

Moreover, of particular relevance to this paper is that the educational environment is rich in choice and that this trend is increasing. The median village in our sample has eight schools offering primary education. Figure II provides a stylized illustration of another village in our sample (different from that in Figure I) that illustrates the substantial school choice and variation in school quality and facilities within villages. While the sample is drawn from a universe of villages that had at least one private school in 2000 (see below), such villages constituted over 50 percent of the population in 2003 and this is rapidly increasing each year. Pakistan is, in fact, in the midst of a remarkable private schooling revolution. The last two decades have seen a more than ten-fold increase in the number of private primary schools (3,800 in 1983 to 47,000 by 2005), and currently, over a third of primary-level enrollment is in the private sector (Andrabi et al., 2008). While schools were denationalized in 1979, until 1990 the growth of private institutions was small and scattered, with the rapid growth really picking up after the mid-1990s.

It is also worth correcting a commonly held perception regarding the nature of these schools in Pakistan. They are not religiously motivated or run by non-governmental agencies but instead are small, "mom and pop" enterprises that are purely driven by their own (fee-based) revenues, with little or no government oversight or support. They offer Western education almost exclusively in coeducational settings (at least at the primary level). In fact, contrary to the notion one may have from the popular media, less than 1 percent of children in Pakistan are enrolled in religious schools (madrassas) and there is no evidence of a significant increase in recent years. Moreover, even this enrollment is hard to attribute, as it anecdotally is, to supposed failure of public schools, high costs of private schooling, and religious radicalism. Most of the variation in religious school enrollment occurs within households: three-fourths of the families who send one child to a religious school send another child to a public or private school. And poorer families are more likely to send their children to a madrassa only when there is no other school in their settlement.

What is also noteworthy and is one of the factors that explains the rising prevalence of private schools is that the typical rural private school is very affordable. While public schools charge no tuition fees, private schools are also not that costly. In a nationwide census of private schools in 2000, the fee in the median rural private school was Rs. 60 per month – less than a day's unskilled wage. According to household survey data from the Pakistan Integrated Household Survey (PIHS, 2001), 18 percent of the poorest third sent their children to private schools in villages where they existed. Private schools remain profitable despite charging such low fees by relying heavily on the supply of local (village) female teachers. These women are typically paid Rs. 1,000/month, less than a fourth of pay in the public sector. Other fixed or start-up costs are minimized as the typical private school is often set up as a small venture, usually using rooms in the (village-based) entrepreneur's own house. In a related paper we document this "local-women-as-teachers effect" by showing that (i) private schools are three times more likely to exist in villages which had a preexisting government girls' high school; and (ii) the wages for skilled women in such villages are 20 percent lower (Andrabi et al., 2007a).

Consistent with the substantial school choice and for-profit nature of the private schools, the educational marketplace in villages appears reasonably competitive. Private schools typically locate in the denser settlements in the villages and often in fairly close proximity to and in direct competition with each other. The average private school in our sample reports that half of all other schools in the village are within a 5-minute walk. With 8 schools in every village, the average private school therefore has 4 schools surrounding it. Moreover, the level and variation of school fees and revenues suggest that the market does act relatively competitively. While hard to estimate, given it is based on self-reported revenues and expenditures, we calculate that the median private school earns total profits of Rs. 14,580 annually, roughly consistent with what the school entrepreneur would make were they to teach in another school. Moreover, fees do seem to respond to market conditions. While there is substantial unexplained variation in fees within villages, the results do show that even when isolating such variation (using village fixed effects), schools with better educational quality, more facilities, and convenient locations are also able to charge more (Andrabi et al., 2007b).

These results suggest that parents, while mostly illiterate, are able to make active and reasonably informed educational choices. While they may not always correctly predict school quality, especially at finer levels of distinction and at intermediate quality levels, the picture that emerges is of interested and engaged parents. The typical households in our sample spend up to 3-5% of their overall budgets for each child on schooling expenditures. Even though government schooling is a free option, poor parents are spending substantially on their children's education, both by enrolling their children in private schools and by spending on additional educational investments beyond school fees. In fact, out-of-pocket spending by households on children's education is higher than what the government spends on providing education through public schools for the richest one-third of our sample and is roughly equal for the middle third. Even among the poorest one-third of households, out-of-pocket expenditures, at Rs. 100 per month, amount to 75 percent of government educational spending on this group.

B. The Data

This paper uses data collected by the authors under a multi-year data gathering exercise, the Learning and Education Achievement in Pakistan Schools (LEAPS) Project. The goal of the LEAPS project is to better understand how much learning is actually taking place and to identify what factors determine the quality of the education children receive. The project details and the first year of the data are available at www.leapsproject.org. This section discusses the sample and the relevant survey instruments used in this study.

SAMPLE

The LEAPS data set is collected from 112 villages in the Punjab province, the largest state in Pakistan. Following an accepted geographical stratification of the province into North, Center, and South, these villages were located in the three districts of Attock (North), Faisalabad (Center), and Rahim Yar Khan (South). Villages were randomly chosen from a list of villages with at least one private school according to the 2000 census of private schools. This allows us to look at differences between private and public schools in the same village. While these villages are somewhat bigger and richer than average villages in these districts, in the past few years an increasing number of villages now have private schools. The survey team worked with all schools offering primary level education, as well as a sample of households in each village. The baseline survey in 2004 covered 823 government and private schools, over 12,000 students, 5,000 teachers and 2,000 households.

An important aspect of the project was to gather information and changes in the entire educational marketplace faced by villagers. From the perspective of the average primary-grade eligible child in the village, the goal was to capture her entire school choice set. From the perspective of the schools, the goal was to capture their potential market. Moreover, given that the potential interventions, such as the one examined in this paper, were likely to affect the entire educational marketplace, care was taken that this marketplace be defined in such a way that it was complete and "closed."

We should note that such an exercise – to be able to define a complete and closed marketplace – is generally quite challenging. For example, if one were considering the educational marketplace for college education in the US, it is likely that it would cover all colleges in the US. While primary education in the US is more feasible, especially given that most parents are restricted to their local public school, the presence of private schools does make this more complicated. One of the features in the Pakistani context that made this exercise feasible is that the villages do cover the majority of the educational marketplace. At least at the primary level, our census of all households in the village showed that there was very little enrollment in schools that are outside the village – 92 percent of all enrolled children attended a village school. Where there was such enrollment, it was typically in schools that were located right outside the village but that still primarily catered to the village population. Therefore, our ultimate sampling strategy was to consider boundaries around the villages in our sample that were within a fifteen minute walking distance from any house in the village. All institutions offering formal primary education within this boundary were covered by our study and are considered to be the "village" schools. Figure III illustrates this. The darker/red schools in the diagram are not in our sample (they are more than 15 minutes from any household), while the lighter/blue ones are. This sampling provided us with a total of 823 schools (public and private) in the 112 villages.

SURVEY INSTRUMENTS

The LEAPS project started in 2003 with a complete census of school choice of the over 80,000 households in the sample villages and since 2004 has conducted four annual rounds of surveys consisting of school, teacher, child, and parent surveys, in addition to annual testing of the same children that started in Grade 3 in 2003. These surveys are briefly described here and in more detail at www.leapsproject.org.

The school, teacher, and student surveys were administered at the schools' premises. School surveys, administered to all 823 (primary) schools in the sample, collected information on infrastructure, prices, costs, and other facilities available in the neighborhood of the school. Three types of teacher surveys were administered. A short roster of questions administered to all teachers in the school and to all teachers who had left the school in the previous two years yields information on about five thousand teachers in the LEAPS Project schools. A longer questionnaire administered to the teachers of the tested children includes detailed socioeconomic information about the teacher and yields data on roughly eight hundred teachers. A head teacher questionnaire included questions on management practices and bonus schemes, along with other modules. In addition, for a sample of ten randomly selected children in every class (roughly six thousand), a short questionnaire was administered to collect information on parental literacy, family structure, and household assets (in classes with less than ten children, all children were chosen).

Household surveys were administered by visiting a subsample of village households, usually after the school-based surveys. The survey was a detailed household questionnaire fielded for 1,800 households in the sample villages, with a special focus on households with students eligible by age for Grade 3. A similar stratified approach was used to sample households with school-age children who were not in school to ensure that we could compare the activities of enrolled and out-of-school children.

Finally, an important part of the survey was child achievement tests. These tests were designed by our team, and the first round was conducted in 2004. All children in Grade 3 (totaling approximately twelve thousand) in the sample schools were tested in Urdu (the national language), mathematics, and English. The same tests were administered across all schools and invigilated by the LEAPS team to ensure impartial and comparable test circumstances.

Several points regarding the test are worth highlighting here. First, we chose to administer a norm-referenced test (as opposed to a criterion-referenced test) in order to measure learning with high precision levels at all levels of knowledge. While a criterion-referenced test, more typical of educational data that one usually has, would distinguish sharply between students who meet or do not meet the specified criteria, it may not yield any information of those below or above the critical level. However, because of the huge differences in learning across schools, designing a norm-referenced test poses its own challenges. Therefore, prior to the design and administration of the final LEAPS test, an extensive pilot was used to identify lower and upper limits of learning in the population and provide an analysis of the validity and reliability of the instrument used. The data from this phase was then used to refine the final test used in the LEAPS Project (Andrabi et al., 2002). As a result, all three tests (English, mathematics, and Urdu) start from extremely simple problems and work up gradually in difficulty. For instance, the English test starts with the invigilator saying an alphabet and children writing it down. It ends with questions based on the reading of a paragraph. Similarly, mathematics and Urdu start with number and alphabet recognition and work up to lowest common multiples (mathematics) and reading comprehension (Urdu).

Second, care was taken in the design to ensure that the test was easy to understand, administer, and relatively unbiased in terms of socioeconomic background. Preference was given to formats that were familiar to children, and effort was made to avoid questions requiring a comprehension ability greater than that which is being tested by the question. Bias with respect to socioeconomic status was also reduced by avoiding references to items that children from particular socioeconomic backgrounds would have little or no exposure to.

Finally, in order to facilitate comparisons of the test over time and to better relate the test to underlying student knowledge, we relied on methods derived from Item-Response Theory to examine the validity of each question, as well as the precision of the test taken in its entirety. Instead of using raw scores, which may not correspond to underlying knowledge (for example, an easy test would have the same child scoring higher), Item-Response Theory allows us to map test performance onto underlying knowledge by essentially estimating different weights to correct answers depending on the difficulty of the question. This is the same methodology used for international exams such as TIMMS and most national testing programs such as the United States' SAT. The knowledge score can thus be interpreted as the student's knowledge or ability in a given subject area such as English, mathematics, or Urdu. For ease of interpretation, we present results using a knowledge score which is in units of standard deviations. An additional advantage of using this technique is that we can also determine the relative precision of the test at different points along the children's knowledge distribution. Our analysis shows that scores around the middle of the distribution are more precise than at the ends of the distribution; this is a standard issue with all tests, since items designed for providing information at the extremes of the distribution also add to information for the mean, but not necessarily the other way around (Andrabi et al., 2002).

Table 1 gives summary statistics for baseline values of some variables of interest. As expected given the randomization, all of these variables are balanced between treatment and control villages.

TIMELINE and TRACKING

The first round of the LEAPS surveys was carried out in 2004. Tests and school-based surveys started in January and February and a household survey was implemented between March and April. In September, the treatment group received their report cards and a second round of testing and household surveys followed in early 2005.

The data set is matched across schools, children, and households, allowing us to follow children and teachers even when they switch schools or drop out. The over 12,000 children we tested in 823 public and private schools in Grade 3 in 2004 were retested in Grade 4 in 2005. All children were tracked between surveys since children could (a) drop-out; (b) remain in the same school and be promoted; (c) remain in the same school and not be promoted; (d) switch schools within the village and be promoted/not promoted (in which case they would be tested in another school); or (e) switch to non-sample schools (usually due to household migration). Although close to 1,800 children out of 12,000 were no longer in the same class-school combination that they would have been if they did not switch schools and were promoted, we were able to determine the status of all except 500, giving us a fairly high tracking rate of over 96 percent throughout the LEAPS Project survey period. All subsequent rounds tested the tracked children as long as they were enrolled in some school in the village.

III Report Card Intervention and Empirical Framework

A. The Report Card Intervention

In each of the three districts, we randomly selected half of the villages to receive the report card intervention. Since the report card intervention affects the entire educational marketplace and we were interested in exploring how the overall market would respond, the intervention was carried out at the village level. This section describes the design and delivery of the report cards.

REPORT CARD DESIGN

The purpose of the report card was to provide information to parents and schools regarding the academic performance of children, both on an absolute scale and relative to other children and schools. The actual report card design was based on discussions with parents and schools in sample villages in each of the three districts. These discussions were typically held in a classroom of the school and were well attended by both parents and teachers (ranging from 50 to 125 people). The final decisions following from these focus groups were: (i) Parents wanted both the scores of the child and his/her rank compared to others, (ii) Parents wanted the average score of every school in their village on the card, and (iii) Teachers wanted the scores to also be broken down by sub-categories (for instance, word-recognition and sentence-building in English, etc.) so that they could concentrate on weaker areas.

Based on these criteria, the finalized parent report cards consisted of two distinct cards. The first reports the score of the child in English, mathematics, and Urdu, as well as her quintile rank across all tested children. This card also reports the average village and school score for the child's village and school and a quintile rank for each of these. The second card reports the scores for all schools in the village, with their quintile rank and the number of children tested. A teacher version of the report card includes the breakdown by sub-categories of the subject scores for each child. In addition, every school received a bound booklet that contained the report cards for all children to be used by both the teacher and head teacher and to serve as an extra copy in case parents lose theirs. Figure IV shows both parent report cards. In addition, it was also made clear that a year after the first round of testing there would also be a second test round whose results would also be made publicly available. This is important to keep in mind since part of the report card intervention was not just revealing baseline scores but also included being able to verify how these scores changed over the year. While not the focus of this paper, this also allowed us to provide additional information (such as child and school score gains in the second report card).

REPORT CARD DELIVERY

The report card delivery mechanism was also based on the views expressed in the pilot focus groups. First, given that many illiterate parents needed to have the cards explained to them, it was deemed appropriate to have the cards be delivered through discussion groups rather than have them mailed to parents. Second, it was important that the discussion group focus on the positive aspects of the card rather than using the card to assign blame. We were concerned about the risk that a poor result would lead to blaming the child. This was minimized by spending close to thirty minutes in an open group discussion on what influences test score results (teacher, home environment, school environment, and the child) before distributing the cards to parents. Third, we were very careful in the discussion to not offer any advice to parents or schools. The goal of the meetings was simply to provide the report cards and explain what the information meant but not to advocate or discuss any particular plan of action.

The delivery mechanism therefore involved calling a meeting at the school level of all the parents of children who were in Grade 3 and distributing the cards after a discussion. Along with the distribution, we noted the name and relation of every parent/relative who received a card for a child. The minimum age that a relative had to be to receive the card was 15. The cards that remained after this school-level distribution were then taken to a central point, where the administrator would remain during the afternoon in case parents could not attend the meeting. Any parents who came to the central meeting point were given the card along with a brief explanation of its content. Finally, the cards that were still not claimed were given to a responsible member of the village who would hand them out over time. This person would also receive a roster of the remaining cards and thus maintain a record of what happened to them.

An unanticipated problem that arose during the card delivery was the large amount of churning that happens during a regular school year. Almost 20 percent of the 12,000 children tested were no longer in the same school year cohort that a normal promotion would suggest. This caused some confusion, since parents did not know whether they should turn up at the school where the child was currently enrolled, or at the school where she was enrolled in the previous year. In addition, some children were not promoted, and similar issues arose. However, most of these parents did attend discussions explaining the report cards and therefore received the report card of their child.

In light of the above discussion, we should acknowledge that the report card is not only a provision of child and school level performance information but also includes the discussion during the report card delivery. While these discussions may also have their own impact, it was infeasible to not include them given that parents are unlikely to fully comprehend the information without explanation. Moreover, from the point of view of feasible interventions, it is likely that any such informational intervention would have to be undertaken through such school-level discussions. It is hard to expect people to respond to information unless they also are able to comprehend it and schools to react to information unless they are convinced that parents will react to it.

B. Empirical Framework

We use a standard difference-in-differences approach to estimate the effect of the report card treatment on a range of dependent variables, including test scores at the student and school level, school fees, school inputs, teacher effort, and household inputs. The estimating equation is:

$$\Delta Y_{ijk} = \alpha_d + \ \beta \cdot RC_i + \gamma \cdot X_{ijk} + \epsilon_{ijk} \tag{1}$$

ΔY_{ijk} is, for example, the change in test score from the baseline to the post-intervention year for a child k in school j in village i. X_{ijk} is a vector of village, school, and child level controls, including school and village size, baseline test score, wealth, a school Herfindahl index that measures the level of competition in the village, and district fixed effects (α_d).

Given that the treatment is randomized, β provides an unbiased estimate of the test score gain associated with the report card intervention. The additional controls are included to improve precision of the estimates. Appendix Table 1 shows that, as expected, control variables and baseline values of outcomes are balanced between treatment and control villages. School fees are slightly larger in control villages, but this is due to one large control village. Excluding this villages does not affect our results, and the difference also disappears once we include district fixed effects (the randomization was stratified by district). Since the treatment is at the village level, in all specifications standard errors are clustered at the village level. The above specification is also estimated at the village level by averaging outcome and control variables across all children/schools in a village.

Given that we anticipate there may be treatment heterogeneity across schools, we also estimate a model with treatment interaction effects separately for the school's type (private, government, or NGO) and whether its baseline score was above (good) or below (bad) the median sample school. These specifications are estimated at the school level:

$$\Delta Y_{ij} = \beta_0 RC_i + \beta_1 GOV_{ij} + \beta_2 NGO_{ij} + \beta_3 GOOD_{ij} + \beta_4 GOV_{ij} \cdot GOOD_{ij} + \beta_5 NGO_{ij} \cdot GOOD_{ij} +$$
$$\beta_6 RC_i \cdot GOV_{ij} + \beta_7 RC_i \cdot NGO_{ij} + \beta_8 RC_i \cdot GOOD_{ij} + \beta_9 RC_i \cdot GOV_{ij} \cdot GOOD_{ij} + \beta_{10} RC_i \cdot NGO_{ij} \cdot GOOD_{ij} + \gamma X_{ij} + \epsilon_{ij}$$

So in this model, the treatment effect on, for example, good government schools would be equal to $\beta_0 + \beta_6 + \beta_8 + \beta_9$. We then justify and move on to a simpler specification with only four groups: good private, bad private, government, and NGO schools:

$$\Delta Y_{ij} = \beta_0 RC_i + \beta_1 GOV_{ij} + \beta_2 GOODPRIV_{ij} + \beta_3 RC_i \cdot GOV_{ij} + \beta_4 RC_i \cdot GOODPRIV_{ij}$$
$$+ \beta_5 NGO_{ij} + \beta_6 RC_i \cdot NGO_{ij} + \gamma X_{ij} + \epsilon_{ij}$$

Finally, we also estimate specifications at the household level where we examine whether village households respond to the report card treatment. These results are either estimated by averaging household outcomes at the village level, or in cases where we want to examine differential household response by school type, we estimate:

$$\Delta Y_{ijh} = \beta_0 RC_i + \beta_1 GOV_{ijh} + \beta_2 GOODPRIV_{ijh} + \beta_3 RC_i \cdot GOV_{ijh} + \beta_4 RC_i \cdot GOODPRIV_{ijh}$$
$$+ \beta_5 NGO_{ijh} + \beta_6 RC_i \cdot NGO_{ijh} + \gamma X_{ijh} + \epsilon_{ijh}$$

where the outcome is for a household h in village i whose child attends school j. The variables are defined analogously. Thus, for example, GOV_{ijh} is a dummy variable that indicates whether household i in village j had a child who attends a government school in the village.

IV Results - Overall Impact

A. Learning

Table 2 presents the estimates for the overall effect of the report card intervention on mean test scores in the village. The dependent variable is the change in test score in the year following the provision of report cards. We present the results first by collapsing all of our data at the village level. In these specifications run at the village level, we control for number of village households, their literacy level, a measure of village wealth, and a school-based Herfindahl index to proxy for competitiveness of the village educational market. We include district fixed effects, and standard errors are clustered at the village level, which is the level of the treatment. Given that the treatment was randomized across villages within districts, the only necessary controls for unbiased estimates are district fixed effects. The remaining controls are included to improve the precision of the estimates. We use the same controls and fixed effects in all specifications that follow (and additional ones in the child-level regressions). In all regressions, we also control for the baseline score.

The treatment effect of the report card intervention is positive and ranges from 0.10 to 0.15 standard deviations depending on the specification. Columns 1 through 3 are the effects for individual subjects, while Column 4 is the average across the three tests. While the point estimates are highest for mathematics, in general, we cannot reject equality of effects across the three subjects, and therefore for the remainder of the paper, we will present results for the averaged score. Recall that the scores are constructed using item-response techniques and therefore represent units of standard deviations of underlying child knowledge in the sample population. We also see a significant amount of mean reversion, with a large and significantly negative coefficient on baseline test score. While interesting, given the randomization, this is not an issue since it ensures that our treatment results do not reflect mean reversion effects.

Column 5 examines the persistence of these results. The dependent variable is the change in test scores two years after the report card provision from the baseline test score. We obtain very similar village-level gains, suggesting that the treatment effect persists.[2]

Column 6 estimates a similar specification to Column 4 but is run at the child level, allowing for additional child-level controls for gender and age and for school-level controls of annual fees, total number of students, percentage of children with at least one parent educated beyond elementary, and school wealth. All of these controls are for the baseline year (including any from the year after would be problematic since they could be affected by the treatment) and will be used throughout in the child-level specifications the remainder of the paper. Given that the treatment is at the village level, standard errors in this and all other specifications are always clustered at the village level. The results show that the effect at the child level is comparable to the village-level effect, suggesting that the average child in our sample gains by a tenth of a standard deviation.

[2]Interpreting this persistence is a bit tricky for two reasons. First, following the first report card intervention, we also provided an additional report card between the first and second year after the first report card. This report card had a similar format but also included information on gains in scores. Second, as we show in a paper using the same LEAPS data, there is significant depreciation in learning, i.e., when estimating value-added models, we estimate as much as 50 percent of the gains in a year are dissipated the following year (Andrabi et al., *Value-Added Estimates*, 2008). Thus, estimating the magnitude of the treatment effect between Years 1 and 2 (the two years after the initial report card) are tricky, but the fact that we get a similar gain in Year 2 relative to the baseline (Year 0) suggests that not only did the gains persist but there must have also been an additional treatment effect between Years 1 and 2 to ensure the gains held.

The child-level specifications also address an additional potential issue of the treatment effect possibly capturing the effect of students dropping out or being absent. The village-level regressions shown simply average over all tested kids in both periods, and in general, this set may not perfectly overlap due to drop-outs and absenteeism. The child-level regressions, by definition, only include those kids who were tested in both periods. They therefore do not suffer from such concerns and show that the same child gains on average a tenth of a standard deviation in score. Alternately, we could replicate the same in the village level regressions by restricting the village averages to children tested in both periods. As suggested by the child-level specification results, doing so (regressions not shown) provides very similar estimates. In addition, we can estimate variants of this at the village level by adding a range of different restrictions, such as excluding kids who drop out after the first period, children who switch schools, children who are absent for one of the two test dates, children who are lost after the first year, or children who are new entrants in the second year. The results look very similar under all these specifications since the majority of children in our sample are tested in both years and because there are not that many who switch schools (5 percent).

Column 7 presents an additional check on the fact that our result is not driven by drop-outs, absenteeism, or school switchers. It restricts the sample to children who were tested in both periods (as in Column 6) and who did not switch schools. The results show that the learning gains for these children remain essentially the same at a tenth of a standard deviation.

A potentially useful exercise that illustrates this further is to mechanically decompose the effect in Column 6 into the gains from non-switching children and gains that arise due to changes in school switching. With a bit of algebraic manipulation, one can show that the overall treatment effect in Column 6 can be decomposed into three components: (i) the gain (between treatment and control) if children were not to switch, (ii) the gain from the change in the number of switchers (between treatment and control), and (iii) the additional gain that switchers may experience in treatment villages. In other words, *Average Treatment Effect* $= \delta + \gamma(R_S^T - R_S^C) + \beta R_S^T$ where $\delta =$ treatment gain if non-switcher; $\gamma =$ general gain (common to treatment and control) that switchers experience; $\beta =$ additional gain switchers see in the treatment villages; and R_S^T and R_S^C are the fraction of switchers in the treatment and control villages. Doing this exercise reinforces that the overall treatment effect is almost entirely driven by non-switchers; the above decomposition gives *Average Treatment Effect* $= 0.102 + 0.037 * (0.01) - 0.13 * (0.057) = 0.102 + 0.0003 - 0.0075 = 0.0948$ (the same as the treatment effect in Column 6).[3] Another way to

[3] A couple of notes of caution are warranted in this decomposition. First, this is a mechanical decomposition and should *not* be interpreted as presenting causal effects for each category. The issue is that switching is an outcome of the treatment, and so unless one assumes that the switchers are identical in treatment and control groups, the parameters (δ, γ, and β) are not treatment effects but rather just represent mechanical decompositions. Second, for expositional simplicity, the decomposition assumes equal number of children in treatment and control. In practice the two numbers are slightly different (1 percent difference), but this correction is second order and so is not done. It is interesting to note that the overall treatment effect on switchers in treatment villages seems to be negative ($-0.03 = 0.10 - 0.13$). While as we cautioned above, this cannot be construed as a treatment effect (given we condition on an outcome variable), it does suggest that it does not seem like switchers would gain a lot in treatment villages. If anything, they may lose somewhat

think of this decomposition would be to estimate the maximum gain for switchers needed to generate the observed overall treatment effect if there was no treatment effect on non-switchers. Assuming no differential switching between treatment and control villages and a 0 treatment effect for non-switchers, one would need to have a (implausibly high) gain (that switchers experience in treatment relative to control villages) of $0.0948/0.057 = 1.7$ standard deviations. While one could lower this somewhat by assuming switchers in general also gain, since there are only 1% more switchers in treatment villages (and it is not significantly different from control villages), one would still need to assume implausibly large learning gains from switching to generate the observed overall treatment effect.

Together these results suggest that the mechanism behind the report card intervention is likely to be driven by changes in the inputs provided by the school or household of the child (or the child herself) rather than reflecting gains due to more mechanical aspects, like differential absenteeism and drop-outs, or less mechanical ones, like children simply switching to better schools.

We should add a note about potential concerns arising from differential attrition in the sample. For example, if those children who would have *gained* less systematically drop out or are absent in the post-intervention year, then this could produce biased estimates. *A priori*, while one may expect those with lower-level scores to possibly attrit more in treatment villages, it is not obvious why one would expect *gainers* to attrit differentially. In fact, given the large mean reversion in gains we see in Table 2, if those with lower-level scores were to attrit more in treatment villages, this would lead to an underestimate of the treatment effect. In any case, not having the counterfactual post-intervention test score of those who attrit, we cannot entirely address this concern. However, a couple of facts suggest this is not a significant one. The rate of attrition is not that high (we are able to retest 82 percent of the children tested in the first round) and is consistent with baseline rates of absenteeism and dropouts.[4] Moreover, the attrition rate is not differential across treatment and control villages (there are no differences even if we separately consider absenteeism and dropouts).

B. Market Responses - Fees & Enrollment

Table 3 now examines the impact of the report card on other market-level outcomes such as school fees and child enrollment, switching, and dropout rates.

The first four columns summarize the treatment effect on school fees. The dependent variable is the change in fees in the school year following the report card intervention. We discuss only private school

(at least initially), especially if they switched midway in the class-year and/or because their school closed.

[4]Specifically, of the total children (tested and not tested) in the class roster in Year 0 (13,735), we have test scores in both rounds for 9,890 students (72 percent). In both rounds the rate of absenteeism is 12 percent (this by itself is reassuring since it suggests neither round is an outlier). Since 9 percent of the children drop out/are lost between Year 0 and Year 1 (and therefore with probability 1 cannot be retested) this gives us the expected fraction of (first round class roster) children for whom we have two test scores to be $(0.88 * 0.88 * 0.91 =)$ 71 percent which is very similar to the 72 percent actually obtained.

fees, as government schools are free and while some do charge a "school fund fee" this is so small and infrequent that it is difficult to separate signal from noise.

Column 1 reports the treatment effect on change in the mean fee charged by private schools in each village. The effect is significant and negative, suggesting that private schools have dropped their annualized fees by an average of Rs. 218 in response to the report card intervention. Column 2 is the same regression, weighted by the number of children in private school in each village. The treatment effect is again significant, negative, and of similar magnitude. Column 3 reestimates the specification in Column 1 but uses log fees to ensure that the result is not driven by outliers; we find a statistically significant 23 percent fee drop in private schools.

While the previous results use information on fees as reported by the schools themselves, Column 4 shows that we obtain similar results if we instead use information on school fees paid as reported by the households. We should caution that we only have this information in a subset of the villages where the surveyed household had a child in a private school. The dependent variable in Column 4 is the village mean of reported private schools' annual school fee and shows a statistically significant decline of 142 rupees per year. While the magnitude is somewhat smaller, we cannot reject that it is equal to that obtained in Column 1.

Column 5 now examines overall child enrollment and finds that there is an increase of 26 children in treatment villages - a bit under 5% of baseline enrollment rates. As we will see later, this mostly comes from an increase in enrollment in public schools (and from lower grades). However, Column 6 shows that despite the large variation in baseline test scores, there is little overall increase in children switching schools. This confirms our previous discussion that the learning gains experienced are unlikely a result of children switching schools. Finally, Column 7 shows that the treatment did not result in any change in dropout rates either.

V Discussion

This section outlines a set of general conceptual insights to help understand how the report card intervention affected the schooling market. The hope is to identify features of the market and the nature of the information provision that may explain when and how such provision can work. In doing so we also generate predictions on the heterogeneous impact of the report cards across school type, examined in the following section.

At the outset, we should note that the evidence points towards the role of providing information on school (as opposed to simply child) quality. Although, one can imagine that parents would alter their behavior with better information about their child's ability, the evidence that we present later in Table 8 shows little support of such parental/child input response – either for the average child or for children in the different types and qualities of schools that we examine. There are two potential explanations. One is that there are behavioral responses but that these depend on the direction in which parents revise their estimates of child performance following the information. If parental beliefs were, on average, correct, then compensating behavior by initially overestimating parents would offset that of initially underestimating parents.[5] A second explanation is that the marginal value of pressurizing schools to improve may be higher in this context of low parental education than direct educational investments at home. This is particularly the case if parents perceive that the problem is the schooling rather than the home environment. Moreover, while one can never rule out that children and parents may be changing their effort in unobservable ways or may experience changes in motivation, it is not clear how an explanation which does not rely on school changes will generate the results on school fees that we observe.

We therefore focus on how the report card affects parental views about school quality and how schools compete for market share as a consequence. We use our overall impact results so far - that show learning gains and fee drops with limited child-switching - and further results on how parental beliefs change, to generate predictions on how the treatment effect may vary based on school type.

Table 4 first examines parental beliefs. In both the baseline and post treatment household surveys we asked parents to rank the quality of schools on a 1 to 5 scale. While it is hard to impose cardinality on this scale, the result in Column 1 does suggest that parental beliefs are increasing in school test scores. However, given that most schools are ranked between 3 and 4, this discrimination across schools is somewhat coarse - with a 10 percentage points increase in school only raising average quality perception by 0.1. Moreover, the regression fit is not that high suggesting that other factors may influence parental beliefs about school quality. In particular, Column 2 shows that parents relatively over-estimate the quality of the school their children currently attend. For the same test score an attended school is given a 0.11 higher rank. Column 3 shows furthermore that even after controlling for test scores, parents rank schools with greater fees higher. While this may partly reflect valuing quality dimensions not captured by test scores, the parental beliefs solicited were regarding school learning quality.

Finally, as evidence of how the report card intervention affected parental beliefs, Column 4 regresses parents beliefs on school quality in round 2 (after the intervention) on baseline beliefs and the treatment effect interacted with baseline score. We find that the treatment lowered parental quality perceptions for worse schools and raised them for better schools. A school which obtained a score of 0 saw its quality belief fall by 0.19 whereas a school which obtained a test score of 100 saw its quality belief rise by 0.31. Compared to the baseline effect, the treatment increases belief sensitivity to test scores by over 50%.

[5] In future work, we hope to examine this further, but from the perspective of this paper, such heterogeneity is unlikely to influence the results observed. For example, if bad private schools show large gains due to a differential response in household inputs by their parents, we should see an average effect at least for those subset of parents. As Table 8 will show, we do not.

Taken together these results suggest that while beliefs did respond to school test score in the baseline as well, there was a lot of noise in rankings and the report card improved parents ability to distinguish school quality as measured by test scores.

The nature of this interaction can be modeled in a variety of ways, but any explanation is likely to share several common features. While the appendix presents the basic structure of one such model that we are currently developing, here we focus on the broader features likely to be at play which suggest how the impact of the intervention may vary by school type.

The basis intuition is as follows. Parents have prior beliefs about the quality of schools and choose accordingly. While individual beliefs may be incorrect, we focus on the case where the average (market) belief is correct. Therefore the impact of the report card intervention is introduced as an increase in the precision of this (market) belief distribution.[6] Schools maximize their choices of quality and price given demand induced by parental beliefs. As in any model of simultaneous price and quality determination, equilibrium choices will depend on the elasticity of demand to price and to quality. When the quality of the information increases, the elasticity of demand with respect to quality increases; a likely result is that the price normalized by quality will decline, with this being attained through a drop in prices, an increase in quality or a combination of both. The model in the Appendix highlights simple conditions which highlight the role of the cost and demand functions in determining the comparative static outcomes; of note is that initially low performing private schools also face pressures from the (free) government sector in addition to the competitive pressures from higher quality private schools. This suggests that such schools may be more likely to respond to competitive pressures due to the report card intervention by raising quality rather than dropping prices. To the extent that some of these schools find it too costly to raise quality (are of lower "types") they may be more likely to shut-down.

The model presented does not address public schools and not knowing what incentives they face, it is hard to predict what their response may be. To the extent that public schools in our context do not at least officially have their resources dependent on performance, one may expect them to show little response. However, if the increased atmosphere of competition also has an impact, albeit somewhat diluted, on the public sector, one may expect to see some learning gains.

Regardless of how public schools may respond, this discussion suggests that one would expect impact heterogeneity both by school type - in terms of whether a school is public or private - and school quality, as determined by whether a school initially had high or low test scores. In the next section we will therefore examine the impact separately along these school characteristics.

[6] While the Column 4 results may suggest mean beliefs have shifted as well, as we will see, it is unlikely this would be a salient force behind our results. In particular, the mean shift explanation would suggest that good private schools (for whom the mean parental belief on quality seems to have increased) should be able to charge higher prices - our results however will show the opposite. Our interpretation of the Column (4) results - that the belief responsive to test scores increased after the intervention - is therefore that this is due to a reduction in the noise (attenuation bias) in parental beliefs. However, lacking a ready way to compare beliefs across parents (they may each have a different belief scale) we are unable to directly check for such a an increase in precision.

Before we present these results we should highlight several broader implications of our discussion. First, an important feature already suggested by our aggregate results is that schools can invest in quality and, furthermore, that such commitments are believed by parents. While this may sound like a somewhat benign claim, it has important implications on how information impacts markets. In the extreme case where quality is determined solely by "type," greater information can lead to higher sorting and cream-skimming. This is similar to the findings of Dranove et al. (2003) and Mizala and Urquiola (2007). Even in the case where service providers' outcomes depend on effort, this response has to be large enough and credibly believed in order to convince (most) consumers to stay (without receiving significant price responses). This is consistent with the situation we observe.

The second implication of our discussion is that the specific tools that schools use to increase price-adjusted quality likely depends on their initial quality. More specifically, one would expect a production function where initial quality increases are exceedingly costly (locating and hiring a motivated teacher), but that further improvements are less so. What we have in mind here is the large literature on teachers as a key input into the production function for cognitive achievement, and our own previous work Andrabi et. al. (2007b) on the difficulties of finding teachers in the rural Pakistani context. It is therefore likely that the investment to hire and retain a good teacher is very high, but that further monitoring or provision of schooling materials is less costly.

Finally, an interesting implication of the discussion and results so far is that it appears schools and parents are able to commit to relatively forward looking arrangements. The reason is that one sees limited switching in equilibrium yet the private sector responds as if faced with greater competition by either increasing quality or dropping fees. Therefore, a likely explanation is that schools are able to commit to parents and that parents are willing to accept the commitment. While one may argue that this is because there are high switching costs, this is unlikely to be the sole driver. If switching costs are too high, the schools would effectively be local monopolists and would be unlikely to respond to the threat of increased competition. In this context, it is important to highlight that part of the report card exercise was also a commitment on our part to retest and re-report school outcomes. One can imagine that without such a commitment, schools may not have been able to credibly commit and parents would not have remained in the schools. This highlights that information provision often plays a dual role, both as a means of reducing noise but also as a means of providing third party (future) verifiability.

VI Impact Heterogeneity

As highlighted in the discussion above, while we expect that the impact of the report card is likely to be different for government and private schools and particularly for private schools, it may depend on the initial quality of the school. We now examine treatment heterogeneity along these lines.

A. Learning

Table 5 first examines whether the learning gains vary by school type and quality. In Column 1, we separate the effect by school type (government, private, or NGO), and by school performance. We use two performance categories, one for schools with average test scores above ("good") and another for those below ("bad") the median school score in the entire test sample (all villages). These categorizations are not balanced; three quarters of private schools are above the median, and two thirds of government schools are below the median. We include NGO schools and interaction terms in these regressions but do not report them because there are very few (sixteen) NGO schools in the sample. We therefore have four groups shown: (i) bad private schools (the coefficient on the report card term); (ii) good private schools; (iii) bad government schools; and (iv) good government schools. At the bottom of the tables, for expositional convenience, we report the average treatment effect for each of these groups, along with the p-value from an F-test that tests whether the effect is different from zero.

The results provide fairly strong evidence for treatment heterogeneity. The report card effect on bad private schools is 0.34 standard deviations and significant at the 5 percent level. The coefficient on bad government schools is 0.08 and significant at the 10 percent level. Good government schools have a similar coefficient that is not statistically significant, and good private schools show an effect very close to zero. The standard errors on the reportcard*good and reportcard*government coefficients allow us to reject the null hypothesis that the treatment effect is the same between (i) good and bad private schools and (ii) bad private schools and bad government schools.

Given that the point estimates for good and bad government schools are similar (though less precise for good government schools, which is not surprising given the smaller number of good government schools), it makes sense to group these two school types and to only distinguish between good and bad private schools. Column 2 repeats the specification in Column 1 with this alternative grouping and shows the results are similar. Given this similarity, we now group the two types of government schools in all subsequent specifications in which we examine heterogeneous impact by school type.

Figure 5 illustrates the heterogeneity in treatment effects across private schools. It plots the raw gains (the only controls are district fixed effects) in treatment and control schools against their baseline test score and shows that indeed schools below the median (roughly corresponding to a baseline score of around zero) show much higher treatment effects (the vertical gap between the two lines). The figure illustrates that the largest gains are experienced closer to the median baseline scores with very low schools also not showing significant gains. Moreover, the top private schools, if anything, seem to show slight quality drops although these are not statistically significant differences. The large downward trend in the data is simply a reflection of the mean reversion that both control and treatment villages experience.

Column 3 takes the heterogeneity results a step further by asking whether the learning gains, particularly those experienced by bad private schools, are higher in more competitive markets. We should caution that breaking down the effects further along this dimension is challenging given the additional demands it puts on statistical power. The specification in Column 3 interacts the treatment status with school type *and* a Herfindahl index that measures the level of concentration of schools within a village. We use a dummy variable that is one if the village Herfindahl index is above the median village Herfindahl index (low competition). The results show that competition does seem to have a greater effect for the bad private schools; the treatment effect is 0.46 standard deviations in more competitive villages and 0.16 in less competitive villages. We find similar effects using slightly different measures of the Herfindahl index, including a local index based on distance between schools, and an index that counted all government schools as one school, under the assumption that government schools do not compete for students with one another.

While the previous results show heterogeneity by quality and type of school, one may question whether the quality results are more about child rather than school quality. The fact that bad private schools show large gains while bad government schools do not suggests that this is unlikely to be the case. Column 4 shows this more formally by adding a separate interaction effect in the specification in Column 2 for whether a child is a good student in the overall sample (i.e., above the sample median or not). This interaction effect is small and insignificant, suggesting that the heterogeneity we capture is indeed more about schools rather than children.

We should note that for one to still propose a bad/good child explanation, it would have to be one where the difference between good and bad child varies by school type and quality. Testing between such an explanation and ours is not feasible since there is little variation left if we define good and bad children within each school type and quality grouping.

However, the relative child performance *within* a school is independent of the mean performance of children in that school. In Column 5, we keep the interactions from Column 2, and we add additional interaction terms for whether a child is above (good) or below (bad) her *school's* median scores. We can now see more clearly whether the effect of being at the bottom of the class varies across school type. There is little differential effect between relatively good and bad children in private schools. However, in government schools, bad students have greater treatment effects than good students (this is also true if we run a specification in which we separate out good and bad government schools). This suggests that government schools may have responded to the treatment by either reallocating effort towards the lower-end students or simply that the marginal returns to increased school effort are higher for the under-performing students.

While it is not surprising that the impact on government schools is more muted than that of (bad) private schools (given that we may expect government schools to not face as high of incentives as the private sector), the difference in performance between bad and good private schools does suggest, as outlined in the previous section, that the cost of producing quality is convex in school effort. Thus, it is relatively more costly for a good private school to increase its quality further and it chooses to become more competitive on the price rather than the quality margin. For bad private schools there is no choice but to improve quality since they are bounded below by the free government schools and, consequently, dropping fees without improving performance does not seem to be a viable option. The fact that their gains are large in more competitive markets (Column 3) offers further evidence for these competitive pressures on quality. The next set of results examines this further by looking at how the report card affects school fees across these school types.

B. School Fees

Table 6 summarizes the treatment effect heterogeneity on private school fees (government schools charge no fees).

Column 1 regresses school fees on treatment, controls, and interaction terms to separate the effect for good and bad private schools. The coefficient on bad private schools is -139 but is not statistically significant from zero. The effect on good private schools is -242 and is significant at the 1 percent level. Columns 2 and 3 verify that this treatment effect is robust. Column 4 reestimates the specification in Column 3 but use log fees to ensure that the result is not driven by outliers; we find a statistically significant 23 percent fee drop in good private schools. Column 4 obtains very similar results (a Rs. 263 drop in fees) when using information on school fees paid as reported by the households.

These results shed further light on why we may only see a test score improvement in bad private schools but not in good ones. As discussed before, if the marginal costs of raising quality further is relatively high for good private schools (for example, due to the concavity of the education production function), then good schools will react more to the increased competition induced by report cards by lowering fees rather than raising quality. Moreover, to the extent that some of the good schools may in fact be not that much better than the free government schools, some parents may decide to shift in favor of the government schools, forcing the good private schools to drop fees further to remain competitive. While one may expect this effect to be stronger in more competitive villages (the point estimates are in line with this prediction [regressions not shown]), the differences are not statistically significant partly because the fee data is noisier and it is therefore hard to obtain precise estimates for these additional interaction terms. Similarly, the lack of a consistent result on fees in bad private schools is consistent with quality improvement being the more profitable margin of change for these schools. Given their already low fees, it is unlikely they could decrease them further. While they could do so by dropping quality further, the presence of free government schools creates a lower bound on the quality drop.

To the extent that the fee drops represent "corrections" by private schools, one would expect that the fee drop would be larger for schools that were over-pricing before. Columns 4-5 prevent evidence that this is indeed the case. The challenge is to construct an appropriate measure of over-pricing. In order to do so we regress baseline schools fees on baseline school score and use the residual from this regression as a metric of overpricing (our results are similar if we consider other variants such as adding polynomials in school score of village/school level controls). We should caution that this is really only valid if each school is catering to the overall village (if school markets are segmented then each market may have differing market demand for quality and hence one would need to compute overpricing for each separately). Since our descriptive analysis suggested there is substantial overlap in parental/child background across schools, we believe this is a reasonable assumption. Column 4 shows that indeed, for every rupee of overpricing at baseline, a private school will drops its fees by Rs. 0.44 in the treatment village (we should caution that while the sign of the effect is not, the magnitude estimates are somewhat sensitive to how we compute the residual). Column 5 then shows that this effect is entirely due to overpricing corrections by good private schools - for every rupee overpriced, good private schools drop fees in the treatment group by Rs 0.67; alternatively, a one standard deviation higher overpricing in the baseline (Rs. 740) by good private schools results in a subsequent fee drop of Rs. 494. While bad private schools do not show any such correction, this may be both because they are not as overpriced in the baseline (the standard deviation of the overpricing residual is less than half that of good private schools) and that they partly also respond by raising quality (although we do not find compelling evidence for the latter).

C. School Closure, Enrollment and Switching

We next turn to school closure, enrollment and switching. While we do not find a lot of overall switching across schools, (bad) private schools are more likely to close in treatment villages, suggesting that there may be heterogeneity in switching and competitive pressures exist due to the (credible) threat of switching.

Columns 1 and 2 in Table 5 examine school closures. Since no government schools close in our sample, we exclude them in our estimates. Column 1 shows that bad private schools are 12 percentage points more likely to close in treatment villages. With a baseline closure rate of 3 percent in the control group, this represents a fourfold increase in relative closure rates for bad private schools. Column 2 illustrates the same result, but in a more continuous form, by interacting with school baseline score; this shows that schools with lower scores are generally more likely to close in treatment villages – A school with a baseline score of a (child) standard deviation below the median is 30 percentage points more likely to close.

Columns 3 through 6 use data on school enrollment for each class year from the school surveys. We can measure enrollment change in a variety of ways. Columns 3 and 4 examine the change in enrollment for the tested cohort, and Columns 5 and 6 consider total enrollment in Grades 1 through 5. Even-numbered columns exclude schools that close (and thus have zero enrollment) to see whether our results are driven more by the closure margin.

Column 3 is the change in enrollment for the surveyed cohort. The point estimates show a drop in enrollment in both good and bad private schools and an increase in enrollment in government schools. However, only the increase in government schools is marginally significant.[7] Column 4 shows that the effects remain similar when excluding closed schools. Columns 5 and 6 examine the effect on overall enrollment and obtain a significant (but relatively small) increase of around 5 children in Grades 1 through 5 in government schools. The point estimates suggest that most of these children come from bad private schools and particularly as a result of school closures (the point estimates for the effect on private schools drop substantially in Column 6, which excludes closed schools).

Column 7 examines heterogeneity in impact for school switching and, consistent with the increased school closures, finds that children are 7 percentage points more likely to switch out of bad private schools (over a 60% increase in baseline switching rates for such schools). Since there are fewer such schools, this result is still consistent with little overall switching (Table 3). Column 8 shows there is no heterogeneity in drop-out rates - there is essentially no impact for all three school types. While these changes are not large in magnitude, they do indicate competitive pressures as government schools gain enrollment at the expense of bad private schools. The magnitudes are muted because initially bad but open private schools show quality gains. To the extent they could have, *ex ante*, signaled these gains to parents, this would lower the incidence of switching. Good private schools may have retained enrollment by dropping their fees. However, it does suggest that individuals overestimated the quality of private schools in general, and therefore on the margin, there is a slight increase in enrollment in government schools. In addition, these results show that good private schools are likely to see drops in profitability due to the intervention, assuming their costs do not drop as much as the fee drop (which is likely since they retain quality). This may not be that surprising since, in a sense, the improved precision in beliefs regarding quality is likely to reduce informational rents that these schools may have been charging initially.

D. Household and School Investments

While the above results highlight the the heterogeneity of the report card impact on school scores, fees and enrollment, and support the channel proposed in the discussion section, we now turn to a more direct examination of these channels.

We categorize these channels into two types: (i) a response due to changed effort from the schools, and (ii) a household response whereby parents directly change their input into the child. The distinction is somewhat arbitrary because ultimately a school-level response is also likely driven by parental pressure, but it is useful to examine whether the parent is directly engaging more with the child or not, as this

[7] The results are of similar magnitude and somewhat more precise if we instead use the change in enrollment in Grade 4 (the tested children's class in the post-intervention period).

may partly explain some of the observed changes. This is particularly relevant once we recognize that the report card intervention not only provides parents with the ability to compare quality across schools but also gives them a better sense of their child's performance. This may prompt parents, particularly of under-performing children, to increase their time and attention at home. While it will be generally hard to distinguish between these two broad channels, our results and discussion previously suggested that it is likely that the response is more due to school-level changes rather than direct parental inputs. This is partly alluded to in our results on school fees, which are harder to reconcile with a direct parental response. But it is also not as surprising once we recognize that for most of these parents, being illiterate, it is not obvious what household-level interventions could be effective. Rather, putting more pressure on schools, either directly or through the threat of switching, may be a more effective way for the parents to ensure quality (or fee) changes.

D-1. School Response

Columns 1 to 4 in Table 8 presents treatment effects on select school inputs. We first examine the qualification of the teacher for the tested cohort. Column 1 runs a probit specification on whether the class teacher for the tested school went from below matriculate to above matriculate qualification and reports marginal effects. We find that both good and bad private schools are 17 to 18 percent more likely to have the tested class teacher increase qualifications to above ten years of schooling. This is most likely driven by a change in the teacher rather than increased schooling for the teacher. However, this effect is only statistically significant for good private schools, possibly due to the far fewer observations for bad private schools.

Column 2 reports an analogous exercise for all the teachers in the school since one may think that the school is switching teachers within the school to ensure that the tested class gets the more qualified teacher. The dependent variable is the percentage of teachers in a village with qualification greater than ten years of education. We find that all schools show some increase, with similar (larger) magnitudes for private schools, although the effect in only significant for good private schools.

Columns 3 and 4 examine two other margins of school inputs. Column 3 first examines whether teaching aids have changed. While we find little impact on several other aids (for example, blackboards [regressions not shown]), we do find that bad private schools experience a large and significant effect in teaching material, with a higher probability of seeing an increase in the number of textbooks used (a similar result is obtained if we simply ask whether textbooks are available or not). Column 4 then examines the school schedule and finds that bad private schools also show a decrease in total break time in treatment villages by almost twenty-three minutes during the school day.

These effects suggest that schools are changing some margins of inputs, and some of the results are larger for bad private schools, as one may expect given that they show the largest learning gains.

However, along a variety of other dimensions (for example, school expenditures), we do not find consistent results. This is partly driven by the fact that this data is very noisy in general and that the margin of changes may be more subtle than that which is readily captured in such broad, self-reported measures. Nevertheless, it does suggest treating these results as tentatively indicative of a school input response.

D-2. Household Response

Columns 5 to 8 in Table 8 now examine the report card impact on different dimensions of direct parental inputs into child learning by utilizing data from the household survey. The data is fairly comprehensive and not only asked parents to give a detailed breakdown for their own daily activities but also constructed a detailed time roster for each child in the household. Moreover, we explicitly asked how much time was spent by any member of the household in helping the child learn. The results show that along all these dimensions, we see little to no change in parental inputs (and where we do it is not in the direction that would generate the learning outcomes observed), lending further support to the likely channel for the learning gains being a school-based input change.

Column 5 examine changes in the number of hours per week spent by parents helping their children with schoolwork and reading to them. We find no significant changes across parents who are sending their children to good/bad private or to government schools. Column 6 looks at daily time (minutes) spent by children on schoolwork outside of school. This includes separately asking for time spent doing homework and preparing for school. The only significant effect is that children who attend good private schools spend 27 more minutes per day on schoolwork at home. This makes the lack of change in overall learning in good private schools even more stark given that the children in such schools seem to working harder at home.

Column 7 considers the impact of parental annual spending on their children's education (we separately ask for spending on a list of categories including pocket money, uniform, travel costs, etc.), excluding the amount spent on school fees (we already report those results in Table 6). While there is a change in such spending, it is a decrease (therefore unlikely to generate the learning gains observed) and only significant for government schools (although the magnitude is also large and negative but only significant at 17% for bad private schools).

Our data reveals that children spend a significant amount of "play-time" at home. While we did not see significant changes in children working more at home, one may expect that greater effort at school (as suggested in Table 6) would cut into this margin. Column 8 shows that play time drops by 85 minutes day for children who attend bad private schools. Examining the source of this drop, we find that it mostly comes from an increase in time the child spends at school (37 minutes; p-value=17%) and in sleeping (30 minutes; p-value=17%). The former is very consistent with the decrease in break time of 23 minutes at school we found in Column 4. The latter, while somewhat harder to interpret, may be a result of the child putting in more effort in school or the parent insisting that the child sleep earlier.

VII Conclusion

This paper examined the impact of disseminating school and child test-score information on equilibrium educational outcomes. We were able to provide information on all choices in the educational marketplace and to study the impact on the market as a whole. Perhaps because of this equilibrium approach, our results are fairly encouraging.

Overall learning increased, and these improvements were greatest for initially poorly performing private schools followed by (all types of) government schools. Although learning did not increase for initially better performing private schools, their fees dropped by 20 percent, leading to a substantial cost saving for parents. That the increase in learning came from the bottom end of the distribution suggests that information on learning outcomes improved efficiency and equity simultaneously. The results also suggest that parental demand is somewhat "bunched" in the sense that there is not a large mass of parents that demand differing levels of (price-adjusted) quality. This would limit the potential for vertical segmentation in the market. This may explain why the impact of the report card is essentially an "emptying" out of the middle (price-quality menu) with initially under-performing private schools raising quality. Thus, after the intervention, one is left with a lower quality but free public sector and a higher quality and somewhat more expensive private sector. We should note though that this emptying out of the middle does not lead to a greater segmentation of the market and worsening equity. Recall, that there is limited switching in equilibrium. So it is not the case that better achieving children are increasingly being siphoned off into the private sector. Rather, conditional on the private sector existing, that separation was already in place. In fact, the report cards likely raise equity by reducing the gaps within the private sector.

It is noteworthy that the cost of providing information was similar to the drop in school fees. Specifically, the upper-bound cost of the report card exercise was \$1 per child (this includes the testing, grading, and report card dissemination exercises). The cost savings were approximately \$3 per child in private schools. With one third of all children enrolled in private schools in these villages, the total cost of providing information for all children is comparable to the drop in fees. Since this was a relatively small-scale exercise, the costs of providing information should be lower if the exercise were to be scaled up. As one example, the random sampling of villages often meant that our work sites were scattered and far from each other, leading to substantial logistical and transport expenses. Saturating the district would be cheaper (per child).

The welfare calculation is considerably harder. If we wish to establish a comparison with other cost-benefit calculations in the low-income country educational literature, it appears that we should focus only on the welfare of households and their children. For instance, in cases where improvements have come through greater effort by teachers, the welfare cost for providers is not accounted for in the cost-benefit analysis. Going down the same route would suggest that the entire improvement in learning is free of cost. Any mechanism that recovers the cost of providing information from parents with children in better performing schools would leave such parents no worse off due to the decline in fees. At the same time, learning gains for other children would lead to an absolute gain. However, a complete welfare analysis would recognize that the decline in fees is composed of both a transfer from the schools to parents (which should not be counted) and a decline in informational rents (which should). Parsing out the two requires a more structural approach that is not explored here.

In addition to the welfare gains, the results are important for the ongoing debate on public versus private education, both in the United States and perhaps even more so in low-income countries where the failure of the public system is often dramatic. Many educational interventions in low-income countries (including India, where the growth of the private sector has been truly spectacular during the last decade) have focused on public schools. Yet, private schools are often cheaper; in our context, the cost per student in a public school is twice as high. This raises the natural question of why governments should spend more money on a system that is already more expensive only to improve learning by amounts that still leave them well short of the levels in the cheaper private schools. Following this logic, several commentators have called for voucher systems where money follows the child rather than the school. Perhaps not surprisingly, these demands have been countered by observations that this would lead to a further weakening of the public sector. What we have been able to show here is that the dissemination of information is an intervention that can simultaneously strengthen the public sector and improve performance in the private sector; such interventions must therefore be the backbone of government policy in an environment where the public and private sectors jointly provide a service.

These results argue for a modified role of the government whereby one of its main responsibilities is to act as a provider of information on the quality of services. It also suggests that information, rather than regulation, may facilitate a more efficient and equitable educational system in such contexts.

References

[1] Andrabi, Tahir, Jishnu Das, Asim Khwaja, Tara Vishwanath, and Tristan Zajonc. 2002. *Test Feasibility Survey, PAKISTAN: Education Sector.* Cambridge, MA: Harvard Kennedy School Working Paper.

[2] Andrabi, Tahir, Jishnu Das, Asim Khwaja, 2007a. *Students Today, Teachers Tomorrow? Identifying Constraints on the provision of Education* (with T. Andrabi and J.Das). Cambridge, MA: Harvard Kennedy School Working Paper.

[3] Andrabi, Tahir, Jishnu Das, Asim Khwaja, Tara Vishwanath, and Tristan Zajonc. 2007b. *The Learning and Educational Achievement in Punjab Schools (LEAPS) Report.* Washington, D.C.: The World Bank. Also Oxford University Press, Pakistan (forthcoming).

[4] Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja. 2008. "A Dime a Day: The Possibilities and Limits of Private Schooling in Pakistan" . *Comparative Education Review.* 52 (3): 329-355.

[5] Andrabi, Tahir, Jishnu Das, Asim Ijaz Khwaja, and Tristan Zajonc. 2008. *Do Value-Added Estimates Add Value? Accounting for Learning Dynamic.* Harvard University Center for International Development, Working Paper No. 158.

[6] Banerjee, Abhijit, Rukmini Banerji, Esther Duflo, Rachel Glennerster, and Stuti Khemani. 2008. *Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in Education in India.* Discussion paper series, Centre for Economic Policy Research, 6781 : Development economics. London: Centre for Economic Policy Research.

[7] Björkman, Martina, and Jakob Svensson. 2008. "Power to the People: Evidence from a Randomized Field Experiment of a Community-Based Monitoring Project in Uganda" . *Quarterly Journal of Economics* (forthcoming).

[8] Chiang, Hanley. 2008. *How Accountability Pressure on Failing Schools Affects Student Achievement.* Harvard University Working Paper.

[9] Dranove, D., D. Kessler, M. McClellan, and M. Satterthwaite. 2003. "Is More Information Better? The Effects of 'Report Cards' on Health Care Providers" . *Journal of Political Economy.* 111: 555-588.

[10] Hastings, Justine S., and Jeffrey M. Weinstein. 2007. "Information, School Choice, and Academic Achievement: Evidence from Two Experiments" . *Quarterly Journal of Economics.* 123 (4): 1373-1414.

[11] Jacob, B. and Levitt, S. (2003). "Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating." Quarterly Journal of Economics. 118(3): 843-877.

[12] Mizala, Alejandra, and Miguel Urquiola. 2007. *Parental Choice and School Markets: The Impact of Information Approximating School Effectiveness*. Documentos de Trabajo 239, Centro de Economía Aplicada, Universidad de Chile.

[13] Morris, S., and H. S. Shin. 2002. "Social Value of Public Information" . *American Economic Review*. 92: 1521-1534.

[14] Rockoff, Jonah E. and Turner, Lesley J. 2008. *Short Run Impacts of Accountability on School Quality*. Columbia University Working Paper.

[15] Stern, Nicholas (Senior Vice President and Chief Economist, The World Bank). March 29, 2001. *Investing for Growth and Poverty Reduction: Institutions and People*. Islamabad, Pakistan.

[16] "Making Services Work for Poor People" . 2004. *World Development Report*.

Appendix I: Model (INCOMPLETE)

Illustrative Model

We illustrate the main intuition using a simple illustrative model first and then extend it to a more general formulation. The basic idea is that the report card intervention increases people's precision in their beliefs regarding school quality. This leads to private schools to lose their market share and profitability unless they raise their price adjusted quality.

Setup:

Schools:

Assume there are three types of schools in the economy: an initially high quality private school (H), an initially low/middle quality private school (M) and a government school (G) (the latter has the lowest quality). Normalize the quality of the H school to be 1, and the government school to be 0. The low quality private school has an intermediate quality $q_M \in (0,1)$ The government school is free while the two private schools charge fees P_H and P_M.

School profits are given by $\Pi = p(\#\ students)$. For simplicity, we build in concavity by assuming that $s > 1$ is not feasible/too costly.

Parents:

Suppose that each parent receives signal $s \sim f(.)$ about the quality of the low quality private school. The signal is on average correct i.e. it has a mean of s_M. For simplicity we will use a uniform distribution for this belief, i.e., $f \sim U[s_M - k, s_M + k]$, where k parametrizes the "inaccuracy" around the belief and is reduced by the introduction of report cards. Further, suppose that parents know the quality of the top school (1) and government school (0) with certainty.[8] For each s, there are λ parents who are "quality conscious," who never consider going to government school, and $(1-\lambda)$ parents that are "not quality conscious," who never consider going to the high quality private school.

Each parent/household maximizes price-adjusted quality

$$L = s - p$$

where p is the price charged.[9]

Solution:

Quality Conscious parents choose school M iff $b_M > 1 - (p_H - p_M)$ where b_M is the parents belief about M's quality; otherwise, they choose school H. Not Quality Conscious parents choose school M iff $b_M > p_M$; otherwise, they choose school G. In both cases the intuition is that only parents with a high enough belief about M's quality will choose it. Thus, on average, school M will have parents who overestimate its quality. This also implies that as long as the thresholds at which parents choose school M are greater than q_M, then the mass of parents choosing school M will strictly decrease as k decreases. Appendix Figure 1, Panel A illustrates this choice graphically.

Consider first the simple case where s is not a choice variable for any of the schools. Thus, school $H's$ optimization decision is

$$p_H = \arg\max_p \lambda p F(1 + p_M - p)$$

where $F(1 + p_M - p)$ represents the demand function facing the school, i.e., the fraction of (quality conscious) parents who choose school H given prices of both schools.

The f.o.c is

$$p_H f(1 + p_M - p_H) = F(1 + p_M - p_H) \tag{2}$$

[8] This assumption is not crucial but makes the setup more tractable and allows parents to be categorized simply by their belief about the quality of the low private school.

[9] We assume both quality conscious and not conscious parents have the same utility function. Alternately, we could have the former obtain a higher marginal return from quality. While the solutions are analogous in this case, we prefer using the same utility function since it provides cleaner and more intuitive mathematical expressions.

Schools $M's$ optimization decision is

$$p_M = \arg\max_p \lambda p\left[1 - F\left(1 + p - p_H\right)\right] + \left(1 - \lambda\right) p \left[1 - F\left(p\right)\right]$$

where $1 - F\left(1 + p_M - p\right)$ represents the fraction of quality conscious parents who choose school M and $1 - F\left(p\right)$ the fraction of not quality conscious parents who choose school M.

The f.o.c is

$$1 - \left[\lambda F\left(1 + p_M - p_H\right) + \left(1 - \lambda\right) F\left(p_M\right)\right] = p_M\left[\lambda f\left(1 + p_M - p_H\right) + \left(1 - \lambda\right) f\left(p_M\right)\right] \tag{3}$$

Plugging in for the uniform distribution gives:

$$p_H = \frac{1 + p_M - s_M + k}{2}$$

And (3) is given by:

$$\begin{aligned}
2k - \left[\lambda\left(1 + p_M - p_H - s_M + k\right) + \left(1 - \lambda\right)\left(p_M - s_M + k\right)\right] &= p_M \\
2k - \left[p_M - s_M + k + \lambda\left(1 - p_H\right)\right] &= p_M \\
k + s_M - p_M - \lambda\left(1 - p_H\right) &= p_M \\
\frac{s_M - \lambda\left(1 - p_H\right) + k}{2} &= p_M
\end{aligned}$$

Solving the two best responses by plugging in the expression for p_H gives

$$\begin{aligned}
p_H &= \frac{1 + \frac{s_M - \lambda(1 - p_H) + k}{2} - s_M + k}{2} \\
&= \frac{2 + s_M - \lambda\left(1 - p_H\right) + k - 2s_M + 2k}{4} \\
&= \frac{2 - \lambda\left(1 - p_H\right) - s_M + 3k}{4} \\
p_H^* &= \frac{2 - \lambda - s_M + 3k}{4 - \lambda}
\end{aligned}$$

and

$$\begin{aligned}
p_M &= \frac{s_M - \lambda\left(1 - \frac{2 - \lambda - s_M + 3k}{4 - \lambda}\right) + k}{2} \\
&= \frac{\left(4 - \lambda\right) s_M - \lambda s_M + 3k\lambda - 2\lambda + \left(4 - \lambda\right) k}{2\left(4 - \lambda\right)} \\
p_M^* &= \frac{\left(2 - \lambda\right) s_M + (2 + \lambda)k - \lambda}{4 - \lambda}
\end{aligned}$$

Since we are interested in how school M is affected due to changes in the precision of beliefs, we also compute the profits for school M. Plugging the expressions for p_H^* and p_M^* and after some significant simplification we obtain:

$$\Pi_M^* = \frac{(p_M^*)^2}{2k}$$

Comparative Statics:

Note from the above expressions for optimal prices, conditional on schools M's quality not changing, a increase in precision (drop in k) will lead to a fall in the prices for both schools. An intuition for this results is that a drop in k implies that the distribution of parental beliefs is more “bunched”. Therefore both schools face a more elastic demand curve creating a force towards dropping their price. We are interested in examining what happens to school M's profits as k decreases since we model the provision of report cards as a decrease in k. Intuitively, one would expect that as k decreases, M would lose market share and therefore profitability.

Applying the envelope theorem, what we find is that

$$\frac{\partial \Pi_M^*}{\partial k} = -\frac{(p_M^*)^2}{2k^2} < 0$$

The intuition for this result is partly illustrated in Appendix Figure 1 (Panels A and B) for the case where the threshold is greater than M. In this case, it is easy to see that, conditional on the school not changing s_M as k decreases, there is a smaller mass of both quality conscious and not conscious parents who will choose school M. While school M can respond by dropping its fees, ultimately it cannot raise demand to compensate and so overall profits will fall. From the expression for school M's profits, we can also see that its profits are increasing in s_M.

$$\frac{\partial \Pi_M^*}{\partial s_M} = \frac{p_M^*(2-\lambda)}{k(4-\lambda)} > 0$$

Moreover,

$$\frac{\partial^2 \Pi_M^*}{\partial s_M \partial k} = -\frac{p_M^*}{k^2}\frac{(2-\lambda)}{(4-\lambda)} < 0$$

i.e. the marginal return to having a higher quality is larger if there is less error (smaller k) in people's beliefs about M's quality.

Report Card Impact:

Using the above comparative statics, we can now illustrate how the report card produces the effects we observe, i.e., an increase in quality of initially poor private schools (with some shutting down) and a drop in fees for the high quality private school. We do so by introducing quality as a choice variable in the simplest possible manner. In the subsequent model presented, we endogenize quality more flexibly and formally.

For now we assume that school M can raise its quality to 1 by incurring a fixed cost of b. Moreover, school M can be of two types differentiated by having a low or high cost of raising quality (i.e., low and high b). Thus school M can choose between two quality levels – q_M if it does not invest and 1 if it invests an amount b. In order to simplify the analysis, we also assume that before the report card is introduced, while school M may raise its quality further, this does not affect parental beliefs (since there is too much noise in inferring quality). This implies that, whether the school M is low or high cost, Mschools will choose quality q_M and set the price as given by our analysis above.[10]

[10] The assumption that schools cannot raise the mean belief about their quality in the no report card world beyond q_M is somewhat stark. Alternately, one can assume that schools can raise the belief about their mean quality. In this case, if both schools coexist then one can obtain both separating and pooling equilibria for both types of school M even when no report cards are provided. However, the separating equilibria are likely to be pareto dominated by the pooling (since

Once the report card is introduced, k falls. We assume parents can see shifts in mean quality for school M. While several different cases are possible, it is not difficult to see that there exists one which provides the results we observe, i.e., that low quality private schools either shut down or improve quality while high quality ones drop fees.

From the comparative statics, we know that as k falls school M's profits fall. If there is any cost of retaining quality q_M (we had initially assumed there was none but can readily reintroduce a cost now), one can see that if school M does not increase quality, it will not be profitable. If b is large enough for the low-type of school M, it is easy to see that such a school will not be willing to upgrade quality either and therefore will shut down. However, as long as b is not too large for the high type school M, it will now raise its quality, and in turn, from the expression for P_H^*, we know this will lead to a drop in fees for the high type school. The expression for P_M^* suggests that the effect of M's fees is somewhat ambiguous. While raising its quality means it can raise its price, at the same time, a drop in k implies a lower price. Moreover, while not formally modeled, there may be a further downward pressure on M's prices. These schools still have to convince parents to not switch, and since parents have not seen the school's quality improvement as yet, schools may have to provide a fee "discount" in order to induce parents to stay.

Endogenizing Quality - Single School Case:

The illustrative model presented above is fairly stylized. The following general setup highlights the choice of quality in the case of a monopolist that provides the central intuition relating elasticity to precision of the signal; in ongoing work we are extending this to the case where we have multiple schools. We anticipate that the basic insights of why the report card has the effects we observe remains: By increasing the precision of people's beliefs about quality, the report card creates competitive pressures on all schools. Those schools for whom the return to investing in quality is higher do so, while others (who are already at high quality levels) respond by dropping prices.

We focus on the school's maximization problem and abstract away from the micro-foundations of the market demand function (developed previously). Assume that the market demand facing the firm is given by $Q(x,\sigma)$ where for simplicity we assume $x = \frac{s}{p}$ where s is product quality and p is unit price as before, and σ now more flexibly represents the error in parental beliefs regarding product quality, s. Given a unit-cost cost function of producing quality s, $c(s)$,the firm solves:

$$\max_{p,s} \Pi = [p - c(s)]\, Q(x, \sigma)$$

The f.o.c wrt to p:

$$Q(x, \sigma) + [p - c(s)]\, \frac{\partial x}{\partial p} Q_x(x, \sigma) = 0$$

where the subscripts indicate partial derivatives. Simplifying further and suppressing arguments we obtain:

$$Q = xQ_x[1 - \frac{c(s)}{p}] \tag{4}$$

The f.o.c wrt to s:

$$-c_s(s)Q(x, \sigma) + [p - c(s)]\, \frac{\partial x}{\partial s} Q_x(x, \sigma) = 0$$

Simplifying further and suppressing arguments we obtain:

$\frac{\partial \Pi_H^*}{\partial k \partial q_M} < 0$) and so can be eliminated using that criteria.

$$Q_x = \frac{c_s(s)}{1 - \frac{c(s)}{p}} Q \tag{5}$$

Using (4) and (5) and simplifying we can establish a relationship between x and s: $c_s(s)x = 1$

Comparative Statics & Report Card Impact:

In order to examine the potential impact of the report card intervention we can now consider the impact of a change in variance in parental beliefs. Using the implicit function theorem we obtain:

$$\frac{dx}{d\sigma} = \frac{Q_\sigma - x\left(\frac{p-c}{p}\right) Q_{x\sigma}}{-\partial f/\partial x} \tag{6}$$

where f is the expression in (4). We are interested in signing the expression to examine the conditions under which x increases (either to to a price drop or quality increase) due to the improved precision that the report card intervention induces i.e. whether $\frac{dx}{d\sigma} < 0$. The denominator in the above expression is$-\partial f/\partial x$$=-\partial f/\partial p \partial p/\partial x<0$. (the first term is just the second order condition and is therefore negative). Therefore the conditions under which we would expect the intervention to lead to an improvement in x would be $Q_\sigma - x\left(\frac{p-c}{p}\right) Q_{x\sigma} > 0$. Given that is is plausible that $Q_{x\sigma} < 0$ (an increase in precision raises the demand elasticity), a sufficient condition would be $Q_\sigma \geq 0$ (i.e. as the variance of perception increases the mass of individuals willing to choose the school does not decrease).

This result highlights that the general intuition illustrated in simple model before holds under very general circumstances. However, what is less obvious is whether the increase in x (the inverse of price adjusted quality) in response to improved precision, also raises s. Our setup also gives the specific condition under which this is likely. Given that $c_s(s)x = 1$, a sufficient condition to obtain that s also increase under the increased belief precision is that $c_{ss}(s) > 0$ i.e. the unit-cost is *concave* in s. Note that this does not imply the overall cost function is concave but rather the assumption here is that provided the cost-reduction from dropping quality is not too large to the firm, it will likely increase quality (and not simply drop price) in order to raise x . A plausible scenario under which this is likely to hold is if, conditional on having a teacher, it does not cost the school that much to raise quality by inducing higher teacher effort. To the extent that initially poorly performing schools are more likely to be in such a scenario, this would explain why they display a large quality response.

The more general setting introduced demonstrates conditions on parental demand and production costs under which an improved precision in parental beliefs can lead to an improvement in the quality/price ratio (x) parents obtain. In ingoing work we are extending this general framework to both examine conditions on the underlying parental utility functions which deliver these results, and the extent to which these results differ once we introduce monopolistic competition settings by including other schools.

Table 1 : Summary Statistics

	Mean	Standard Deviation	N
Village Wealth (Median Monthly Expenditure)	4641.495	1575.162	112
Number of Households in Village	631.2857	383.886	112
Herfindahl Index of Schools in Village	0.194306	0.0761075	112
Percent Adults (>24) Literate in Village	37.3252	11.85356	112
School Average Test Score	0.050349	0.7246499	780
School Fees	535.7185	799.7473	803
Number of Students Enrolled at School	166.2201	150.739	804
Percentage of School's Children with 1+ Parent Educated Beyond Elementary	0.555275	0.2818515	804
Mean School Wealth (Child Asset Index)	-0.02485	1.099659	804
Child Average Test Score	-0.01832	0.9130758	12110
Female Child	0.442519	0.496703	13735
Child Age	9.675417	1.474713	13733

Table 2 : Overall Learning Impact

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Round 2 Village-Level Scores				**Round 3 Village-Level Scores**	**Round 2 Child-Level Scores**	
	English	**Urdu**	**Math**	**Average**	**Average**	**Average**	**Average (No Switchers)**
Report Card	0.1	0.101	0.147	0.114	0.123	0.095	0.102
	(0.048)**	(0.041)**	(0.059)**	(0.045)**	(0.043)***	(0.038)**	(0.038)***
Baselines	-0.376	-0.325	-0.322	-0.294	-0.356	-0.335	-0.336
	(0.064)***	(0.059)***	(0.073)***	(0.062)***	(0.075)***	(0.031)***	(0.032)***
Controls	Basic	Basic	Basic	Basic	Basic	Full	Full
Constant	-0.038	-0.113	-0.214	-0.1	0.547	0.093	0.07
	(0.203)	(0.146)	(0.219)	(0.172)	(0.166)***	(0.149)	(0.151)
Observations	112	112	112	112	112	9867	9330
R-squared	0.34	0.55	0.43	0.44	0.4	0.21	0.21

This table presents the estimates for the overall effect of the report card intervention on mean test scores in the village. The dependent variable for the first four columns is the change in test score in the year following the provision of report cards. Columns (1)-(3) are the score effects for individual subjects. Column (4) is the average across the three tests. Column (5) examines the persistence of these results by looking at the change in test scores two years after the intervention. Column (6) estimates a similar specification to Column (4) but is run at the child level. Column (7) is identical to Column (6) but excludes students who switch schools. Robust standard errors are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%. Basic controls include district-level fixed effects and village-level controls (village wealth [median monthly expenditure], number of households in village, Herfindahl index of schools in village, and percent adults [>24] literate in village). Full controls include basic controls as well as school-level controls (school total revenue/student, number of students enrolled at school, percentage of school's children with 1+ parent educated beyond elementary, mean school wealth [child asset index], and school average test score) and child-level controls (child average test score, female child, and child age).

Table 3: Market Reponses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Village Private Fee Change				**Village Enrollment**	**Probability that**	
	From School Survey	**Weighted by School Size**	**Log**	**From Household Survey (short)**		**Child Switches**	**Child Drops Out**
Report Card	-217.96	-190.478	-0.24	-141.714	25.76	0.006	0.006
	(65.090)***	(64.900)***	(0.087)***	(74.351)*	(11.16)**	(0.004)	(0.004)
Baseline	-0.284	-0.233	-0.565	-0.258	0.006		
	(0.094)***	(0.107)**	(0.105)***	(0.083)***	(0.012)		
Controls	Expanded	Expanded	Expanded	Expanded	Expanded	Full	Full
Constant	504.917	408.955	4.146	329.616	-9.364	-0.074	-0.075
	(227.404)**	(249.201)	(0.796)***	(293.395)	(41.69)	(0.021)	(0.021)***
Observations	107	107	106	83	112	12085	12085
R-squared	0.27	0.28	0.38	0.18	0.23	0.02	0.02

This table describes fee and enrollment changes. Column (1) reports the treatment effect on change in the mean fee charged by private schools in each village. Column (2) is the same as Column (1) but weights by the number of kids in private schools in each village. Column (3) repeats Column (1) again but using the log of fees. Column (4) provides a check for Column (1) by using household survey data. Column (5) examines the effect of the treatment on total enrollment in villages. Columns (6) and (7) show the probability of a child switching or dropping out due to the treatment. Robust standard errors are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%. Basic controls include district-level fixed effects and village-level controls (village wealth [median monthly expenditure], number of households in village, Herfindahl index of schools in village, and percent adults [>24] literate in village). Expanded controls include basic controls as well as school-level controls (school total revenue/student, number of students enrolled at school, percentage of school's children with 1+ parent educated beyond elementary, mean school wealth [child asset index], and school average test score). Full controls include expanded controls as well as child-level controls (child average test score, female child, and child age).

Table 4 : Parental Perception of School Quality

	(1)	(2)	(3)	(4)
	Year 1	**Year 1**	**Year 1**	**Year 2**
School Score	0.01	0.01	0.005	0.003
	(0.001)***	(0.001)***	(0.001)***	(0.002)
Attendance (current or previous)		0.111	0.106	
		(0.024)***	(0.023)***	
School Fee			0.000	0.000
			(0.000)***	(0.000)***
Baseline Perception				0.121
				(0.017)***
Report Card				-0.193
				(0.094)**
Reportcard * Score				0.005
				(0.003)*
Constant	2.687	2.663	2.666	2.371
	(0.096)***	(0.096)***	(0.084)***	(0.163)***
Observations	9131	9131	9110	5939
R-squared	0.06	0.07	0.1	0.18

This table summarizes data on school perceptions. Columns (1) - (3) show the correlations between school characteristics, attendance, and parental perception; the dependent variable is perception in year 1, ranked on a five point scale. Column (4) is the regression of perception in year 2 on the treatment indicator, controls, and an interaction between the treatment indicator and school test score. Robust standard errors are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%. All regressions include district-level fixed effects, village-level controls (village wealth [median monthly expenditure], number of households in village, Herfindahl index of schools in village, and percent adults [>24] literate in village), and school-level controls (school total revenue/student, number of students enrolled at school, percentage of school's children with 1+ parent educated beyond elementary, mean school wealth [child asset index], and school average test score).

Table 5: Heterogeneity in Learning Gains

	(1)	(2)	(3)	(4)	(5)
	Child-Level, type interaction	**Child-level, good priv / bad priv / gov / NGO**	**Child-level, type-interaction + herf**	**type, good/bad child in sample, rc**	**type, good/bad child in school, rc**
Report Card (RC)	0.347	0.359	0.462	0.378**	0.415
	(0.145)**	(0.154)**	(0.202)**	(0.160)	(0.229)*
RC * Government School (Gov)	-0.269	-0.271	-0.369	-0.267*	-0.283
	(0.147)*	(0.154)*	(0.200)*	(0.155)	(0.227)
RC * Good School (Good Schl)	-0.317				
	(0.149)**				
RC * Gov * Good Schl	0.348				
	(0.193)*				
RC * Good Private School (Good Priv)		-0.347	-0.436	-0.324**	-0.381
		(0.158)**	(0.211)**	(0.161)	(0.234)
RC * Low Competition			-0.302		
			(0.222)		
RC * Gov * Low Competition			0.295		
			(0.227)		
RC * Good Priv * Low Competition			0.269		
			(0.231)		
RC * Good Student in Sample				-0.045	
				(0.067)	
RC * Good Student in School (Good Stud)					-0.111
					(0.161)
RC * Good Stud * Good Priv					0.077
					(0.168)
RC * Good Stud * Gov					0.03
					(0.162)
Good School	-0.011				
	(0.064)				
Government School	-0.113	-0.16	-0.182	-0.149*	-0.197
	(0.068)	(0.075)**	(0.099)*	(0.076)	(0.100)*
Good Private School		0.024	0.029	0.094	0.059
		(0.065)	(0.081)	(0.083)	(0.094)
Constant	0.209	0.227	0.326	0.042	0.066
	(0.164)	(0.171)	(0.177)*	(0.172)	(0.183)
Observations	9867	9867	9867	9867	9754
R-squared	0.22	0.22	0.22	0.24	0.23

SUBGROUP POINT ESTIMATES, F-TEST p-VALUES IN PARENTHESES

Bad private school	0.347	Bad private school	0.359	Bad private school, high competition	0.462	Bad private school, bad kid in sample	0.378354	Bad private school, bad kid in school	0.415
	(0.019)		(0.021)		(0.025)		(0.020)		(0.072)
Bad government school	0.078			Bad private school, low competition	0.160	Bad private school, good kid in sample	0.333	Bad private school, good kid in school	0.304
	(0.117)				(0.130)		(0.035)		(0.001)
Good private school	0.030	Good private school	0.013	Good private school, high competition	0.026	Good private school, bad kid in sample	0.054604	Good private school, bad kid in school	0.034
	(0.516)		(0.777)		(0.683)		(0.430)		(0.484)
Good government school	0.108	Government school	0.088	Good private school, low competition	-0.007	Good private school, good kid in sample	0.010	Good private school, good kid in school	0.000
	(0.335)		(0.053)		(0.910)		(0.840)		(0.999)
				Government school, high competition	0.092	Government school, bad kid in sample	0.111764	Government school, bad kid in school	0.132
					(0.174)		(0.023)		(0.007)
				Government school, low competition	0.085	Government school, good kid in sample	0.066735	Government school, good kid in school	0.051
					(0.165)		(0.312)		(0.299)

This table examines the differential learning impact across types of schools. Column (1) separates the effect by school type (government, private, or NGO) and by school performance (below/above median). Column (2) repeats the specification in Column (1) with good and bad government schools grouped. Column (3) asks whether the learning gains, particularly those experienced by bad private schools, are higher in more competitive markets. Column (4) interacts the treatment status with school type and whether the child is above or below the sample median score. Column (5) repeats Column (4) with student split relative to the median school score. Robust standard errors are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%. All columns control for district-level fixed effects, village-level controls (village wealth [median monthly expenditure], number of households in village, Herfindahl index of schools in village, and percent adults [>24] literate in village), school-level controls (school total revenue/student, number of students enrolled at school, percentage of school's children with 1+ parent educated beyond elementary, mean school wealth [child asset index], and school average test score), and child-level controls (child average test score, female child, child age) - All columns also include interactions terms with NGO and other interactions and level terms that are necessary given the interaction terms included.
Panel 2 displays the estimated coefficients for relevant subgroups.

Table 6 : School Fees

	(1)	(2)	(3)	(4)	(5)
	School Fees as Reported by Schools		**School Fees as Reported by Households**	**School Fees as Reported by Schools**	
	By School Type	**By School Type with LHS var in logs**	**By School Type**	**By "Residual" Fees**	**By "Residual" Fees & School Type**
Report Card (RC)	-139.901	-0.126	-10.958	-243.92	-142.095
	(116.646)	(0.118)	(137.467)	(51.857)***	(140.993)
RC * Good Private School	-102.934	-0.130	-252.503		-149.713
	(129.963)	(0.142)	(160.36)		(158.998)
Good Private School	49.630	0.114	126.012		119.932
	(134.668)	(0.109)	(99.892)		(134.326)
Resid				-0.024	-0.379
				(0.108)	(0.341)
RC * Resid				-0.438	0.095
				(0.161)***	(0.307)
RC * Resid * Good Private School					-0.668
					(0.326)**
Resid * Good Private School					0.59
					(0.235)**
Baselines	-0.307	-0.818	-0.525	0.044	0.044
	0.416	(0.123)***	(0.134)***	(0.342)	(0.342)
Controls	Expanded	Expanded	Basic		
Constant	580.738	5.402	404.305	546.208	61.357
	(236.027)**	(0.856)***	(141.390)***	(232.143)**	(284.231)
Observations	269	278	878	269	269
R-squared	0.19	0.34	0.25	0.27	0.27
SUBGROUP POINT ESTIMATES, F-TEST p-VALUES IN PARENTHESES					
Bad private school	-139.171	-0.126	-10.958		
	(0.237)	(0.286)	(0.937)		
Good private school	-241.841	-0.257	-263.461		
	(0.000)	(0.001)	(0.002)		

This table summarizes the treatment effect on school fees. Column (1) regresses school fees on treatment, controls, and interaction terms to separate the effect for good and bad private schools. Columns (2) reestimates Column (1) but uses log fees to ensure that the result is not driven by outliers. Column (3) draws on questions in the household survey about school fees paid and separates the effect by school type, with the household mean of reported private monthly school fee as the dependent variable. Column (4) examines whether the fee changes vary by whether (in the baseline year) a school was "over-charging" (a positive value of the fee residual) or not. Column (5) seperately considers this for the different school types. Robust standard errors are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%. Basic controls include district-level fixed effects and village-level controls (village wealth [median monthly expenditure], number of households in village, Herfindahl index of schools in village, and percent adults [>24] literate in village). Expanded controls include basic controls as well as school-level controls (school total revenue/student, number of students enrolled at school, percentage of school's children with 1+ parent educated beyond elementary, mean school wealth [child asset index], and school average test score). Full controls include expanded controls as well as child-level controls (child average test score, female child, and child age).
Panel 2 displays the estimated coefficients for relevant subgroups.

Table 7 : School Closure and Enrollment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School Closure	**School Closure**	**Change in Survey Cohort Enrollment (Year 4 - Year 3)**	**Change in Survey Cohort Enrollment (Year 4 - Year 3, no closed)**	**Change in Total Enrollment, Classes 1-5**	**Change in Total Enrollment, Classes 1-5 (no closed)**	**Probability of child switching**	**Probability of child dropping out**
Report Card (RC)	0.117	0.134	-1.226	-0.769	-5.304	-1.319	0.070	-0.001
	(0.052)**	(0.042)***	(1.008)	(1.02)	(4.16)	(4.00)	(0.037)*	(0.022)
Government School (Gov)			2.100	2.005	7.813	7.496	-0.053	-0.050
			(.96)**	(0.971)**	(3.193)**	(3.232)**	(0.024)**	(0.021)***
Good Private School	0.059		0.421	0.550	0.562	2.195	0.014	-0.019
	(0.032)*		(1.11)	(1.08)	(3.46)	(3.11)	(0.022)	(0.02)
RC * Gov			2.148	1.703	10.380	6.483	-0.073	0.006
			(1.21)*	(1.22)	(5.11)**	(4.79)	(0.038)*	(0.022)
RC * Good Private School	-0.081		0.780	0.540	4.844	1.847	-0.052	0.007
	(0.062)		(1.24)	(1.14)	(5.11)	(4.28)	(0.04)	(0.03)
RC * School Baseline Score		-0.178						
		(0.055)***						
Baseline School Score	0.000	0.003	-0.070	-0.062	-0.058	-0.013		
	(0.001)	(0.001)***	(0.03)**	(0.03)**	(0.08)	(0.08)		
Baseline Enrollment			-0.076	-0.076	-0.138	-0.143		
			(0.19)	(0.19)	(.06)**	(0.06)**		
Constant	0.207*	0.021	-2.901	-2.226	-4.763	-0.317	0.076	-0.023
	(0.115)	(0.103)	(1.66)*	(1.58)	(6.48)	(5.30)	(0.040)*	(0.03)
Observations	303	291	802	779	798	775	12085	12085
R-squared	0.07	0.11	0.17	0.17	0.12	0.11	0.04	0.02
SUBGROUP POINT ESTIMATES, F-TEST p-VALUES IN PARENTHESES								
Bad private school	0.117		-1.225	-0.769	-5.304	-1.320	0.070	-0.001
	(0.030)		(0.227)	(0.451)	(0.205)	(0.742)	(0.06)	(0.95)
Good private school	0.037		-0.446	-0.230	-0.460	-0.528	0.018	0.005
	(0.200)		(0.537)	(0.687)	(0.880)	(0.777)	(0.71)	(0.38)
Government school			0.922	0.933	5.080	5.164	-0.003	0.006
			(0.153)	(0.142)	(0.010)	(0.008)	(0.33)	(0.62)

This table examines the impact on school closure and enrollment. Columns (1) examines whether the report card impacts school closures. Column (2) illustrates the same result but in a more continuous form by interacting with school baseline score. Column (3) examines the change in enrollment for the surveyed cohort. Column (4) shows Column (3) while excluding closed schools. Column (5) examines the change in enrollment for grades 1-5. Column (6) does the same with the exclusion of closed schools. Robust standard errors are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%. All columns control for district-level fixed effects, village-level controls (village wealth [median monthly expenditure], number of households in village, Herfindahl index of schools in village, and percent adults [>24] literate in village), and school-level controls (school total revenue/student, number of students enrolled at school, percentage of school's children with 1+ parent educated beyond elementary, and mean school wealth [child asset index]). In addition, Columns (3)-(6) control for interactions with NGO and all other remaining interaction terms that are implied by the ones shown in the columns.
Panel 2 displays the estimated coefficients for relevant subgroups.

Table 8: School and Household Input Changes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School Inputs (School Level)				**Household Inputs (Household Level)**			
	Class Teachers Improve to Matric	**Percent Change in Matric Teachers**	**Textbook Probit**	**Break Time**	**Parental Time Spent on Education with Kids**	**Kids' Time Spent on School Work Outside of School**	**Parental Spending on Education Not Including School Fees**	**Child Time Spent Playing**
Report Card (RC)	0.165	0.035	0.405	-22.789	-1.492	6.531	-375.684	-85.071
	(0.154)	(0.059)	(0.240)*	(8.151)***	(1.25)	(24.19)	(274.66)	(32.381)***
Government School (Gov)	-0.13	0.042	1.942	-8.814	-1.452	-15.712	-8.621	-27.709
	(0.123)	(0.042)	(0.228)***	-7.957	-1.029	-23.322	-246.007	-25.774
RC * Gov	-0.101	-0.01	-0.359	24.359	1.343	-4.967	202.785	73.548
	(0.164)	(0.060)	(0.286)	(8.314)***	-1.259	-25.195	-285.657	(32.436)**
Good Private School	-0.042	0.019	0.097	-9.758	0.263	-11.15	67.657	-19.216
	(0.112)	(0.046)	(0.204)	-8.264	-1.071	-22.018	-277.341	-26.985
RC * Good Private School	0.018	0.002	-0.414	24.398	0.602	20.482	342.752	73.635
	(0.208)	(0.062)	(0.270)	(8.834)***	-1.413	-27.781	-321.52	(36.557)**
Baseline				-0.96	-0.799	-0.938	-0.737	-0.932
				(0.042)***	(0.033)***	(0.042)***	(0.057)***	(0.028)***
Controls	Expanded	Expanded	Expanded	Expanded	Full	Full	Full	Full
Constant		0.003		51.479	2.495	182.367	1159.677	171.29
		(0.066)		(11.213)***	(1.90)	(40.850)***	(410.462)***	(45.845)***
Observations	340	744	727	778	857	827	810	826
R-squared		0.14		0.56	0.56	0.37	0.36	0.58
SUBGROUP POINT ESTIMATES, F-TEST p-VALUES IN PARENTHESES								
Bad private school	0.165	0.035	0.405	-22.789	-1.490	6.530	-375.684	-85.07***
	(0.29)	(0.56)	(0.09)	(0.01)	(0.234)	(0.788)	(0.174)	(0.010)
Good private school	0.183	0.037	-0.009	1.609	-0.890	27.012*	-32.930	-11.430
	(0.06)	(0.08)	(0.95)	(0.65)	(0.207)	(0.052)	(0.844)	(0.519)
Government school	0.064	0.025	0.046	1.57	-0.149	1.563	-172.899*	-11.520
	(0.39)	(0.87)	(0.77)	(0.47)	(0.625)	(0.873)	(0.055)	(0.108)

This table examines treatment effects on several different school and household inputs. Column (1) runs a probit specification on whether the class teacher for the tested school went from below matric to above matric qualification and reports marginal effects. Column (2) reports an analogous exercise for all the teachers in the school. Column (3) examines whether teaching aids (textbooks) have changed. Column (4) examines the school schedule (break-time). Columns (5) - (7) examine household imputs. Column (5) examines changes in the number of hours per week spent by parents helping their children with school-work and reading to them. Column (6) looks at total daily time in minutes spent by children on schoolwork outside of school. Columns (7) looks at the impact of parental annual spending on their children's education excluding that on school fees. Column (8) looks at total daily time in minutes spent by children playing. Robust standard errors are in brackets with * significant at 10%; ** significant at 5%; and *** significant at 1%. Basic controls include district-level fixed effects and village-level controls (village wealth [median monthly expenditure], number of households in village, Herfindahl index of schools in village, and percent adults [>24] literate in village). Expanded controls include basic controls as well as school-level controls (school total revenue/student, number of students enrolled at school, percentage of school's children with 1+ parent educated beyond elementary, mean school wealth [child asset index], and school average test score). In addition, Columns (2) and (4)-(8) include NGO interactions.
Panel 2 displays the estimated coefficients for relevant subgroups.

Figure 1 : Village Arial View

This figure overlays GPS readings for private schools and markets with a Google Earth image of a typical village in the sample.

Figure 2 : Village Diagram

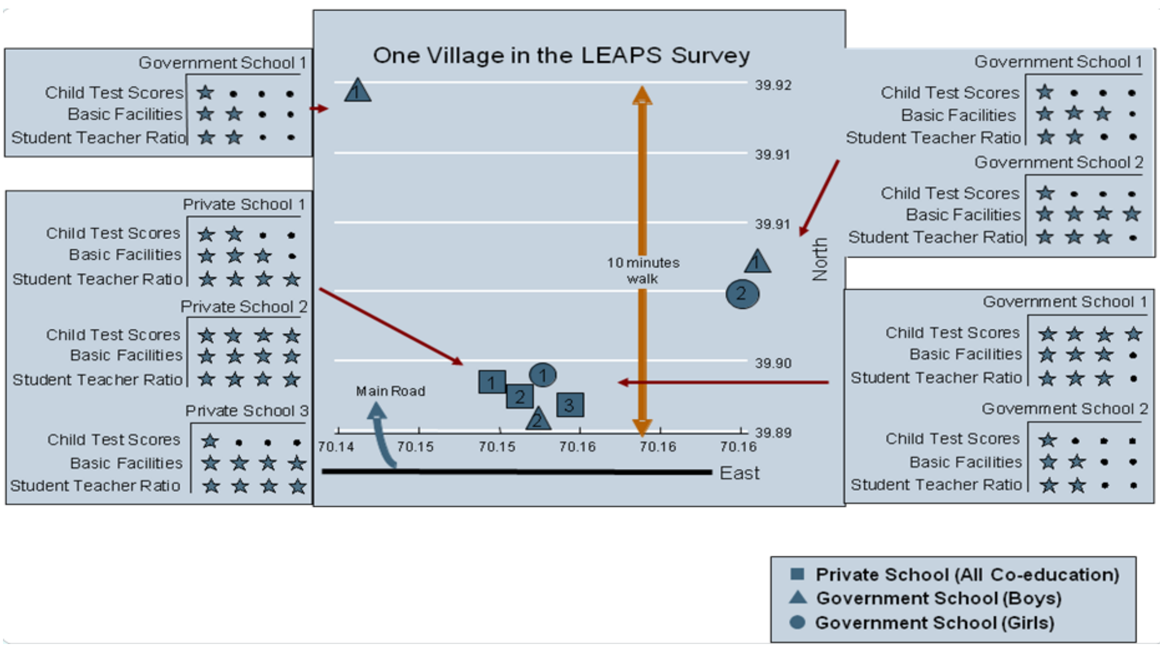

This map shows the actual layout of schools in a sample village (different from Figure 1) and demonstrates the substantial quality and infrastructure variation across schools even within a village.

Figure 3 : School Sample Selection

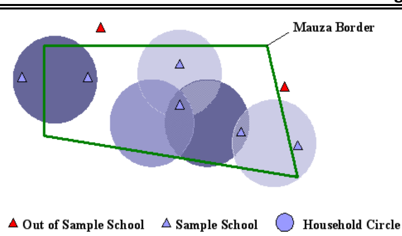

This figure demonstrates how schools were sampled. Boundaries were made around the villages in the sample that were 15 minutes walking distance from any house in the village. All institutions offering formal primary education within this boundary are covered in the study. For instance, the red schools in this diagram are not in the sample (they are more than 15 minutes from any household), while the ones in blue are. The green line represents the village border.

Figure 4 : Report Cards

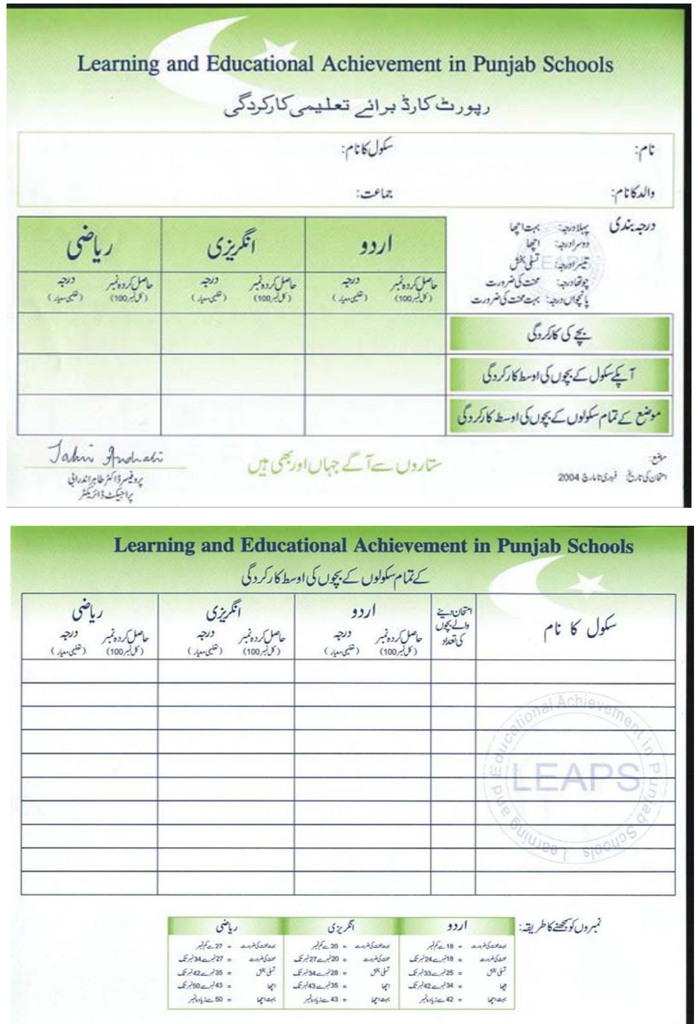

The top image shows the Report Card (part 1 - Child information) for the child with Math, Urdu, and English in each column. For each subject, the absolute score and the quintile (described as "needing a lot of work" to "very good") is given. The three rows display information for the child, her school, and her village. The lower image shows part 2 of the Report Card (school information) and gives information on the village schools, one on each row. The columns display the school name, number of tested children, and school scores and quintiles for each of the three subjects. Each card also had detailed instructions (on the reverse side) on how to read the card and what the rankings meant.

Figure 5 : Score Gains Treatment Heterogeneity - Private Schools

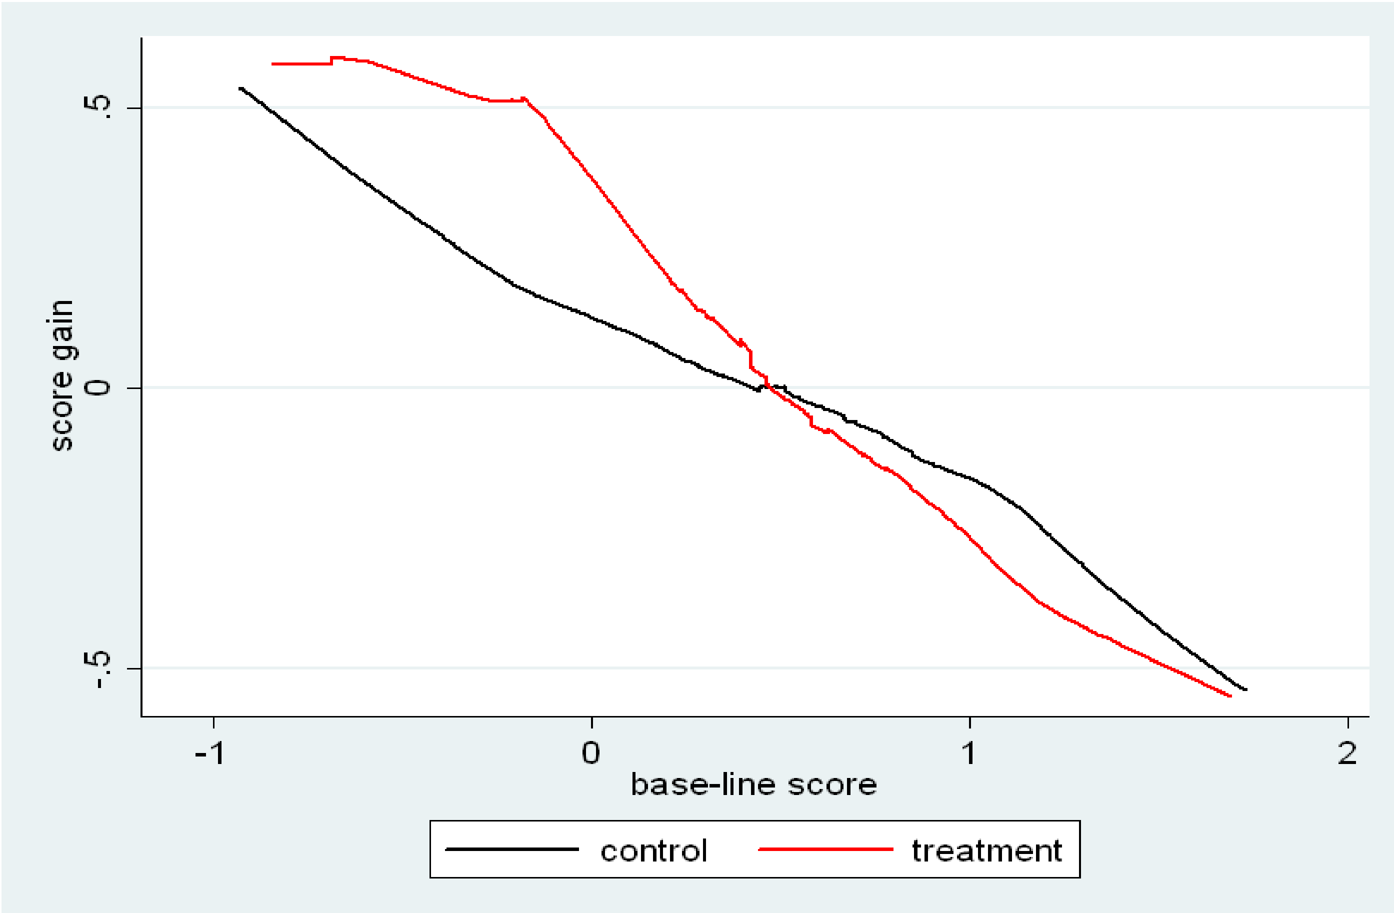

This graph shows a smooted plot of score gains in treatment and control schools restricting the sample only to private schools. The X-axis plots the schools baseline score and the y-axis shows the average score gain for a particular school.

Appendix Table 1 : Randomization Balance

	Treatment	Control	Difference
Village Level			
Village Wealth (Median Monthly	4585.375	4697.615	-112.240
Expenditure)	(219.486)	(202.802)	(298.836)
Number of Households in Village	626.500	636.071	-9.571
	(44.209)	(57.929)	(72.871)
Level of Competition between Schools in	0.197	0.192	0.005
Village (Herfindahl Index)	(0.011)	(0.010)	(0.014)
Village Inequality (Gini Index)	0.808	0.792	0.016
	(0.013)	(0.014)	(0.019)
Number of Government Schools in Village	4.268	4.625	-0.357
	(0.382)	(0.390)	(0.546)
Number of Private Schools in Village	2.875	2.982	-0.107
	(0.316)	(0.302)	(0.437)
Village enrollment % (All)	70.617	71.011	-0.395
	(2.276)	(2.266)	(3.212)
Village enrollment % (Boys)	76.464	75.927	0.538
	(2.099)	(2.097)	(2.968)
Village enrollment % (Girls)	64.106	65.583	-1.477
	(2.710)	(2.565)	(3.731)
No. of Grade 3 Children Tested in Village	103.321	112.929	-9.607
	(8.578)	(9.547)	(12.834)
Village Adult (>24 yrs) Literacy (%)	38.472	36.179	2.293
	(1.746)	(1.403)	(2.240)
School Level			
School Average Test Score	0.048	0.053	0.004
	(0.035)	(0.039)	(0.052)
School Fees	473.650	585.893	112.243
	(35.046)	(44.979)	(56.629)**
Number of Students Enrolled at School	165.179	173.204	8.025
	(7.287)	(8.110)	(10.878)
Percentage of School's Children with 1+	0.551	0.564	0.012
Parent Educated Beyond Elementary	(0.014)	(0.014)	(0.020)
Mean School Wealth (Child Asset Index)	-0.067	0.042	0.109
	(0.052)	(0.058)	(0.077)
Child Level			
Average Test Score	-0.023	-0.013	0.009
	(0.011)	(0.012)	(0.017)
English Test Scores	0.007	-0.006	0.013
	(0.057)	(0.043)	(0.071)
Math Test Scores	0.006	-0.007	0.013
	(0.050)	(0.050)	(0.070)
Urdu Test Scores	0.012	-0.011	0.023
	(0.051)	(0.042)	(0.066)
Female Child	0.439	0.446	-0.007
	(0.014)	(0.012)	(0.019)
Child Age	9.648	9.624	0.024
	(0.085)	(0.055)	(0.101)
Father's Education	2.203	2.139	0.064
	(0.035)	(0.034)	(0.048)
Mother's Education	1.564	1.581	-0.017
	(0.038)	(0.035)	(0.052)
Wealth	0.083	-0.078	0.161
	(0.099)	(0.082)	(0.129)

*This tables shows the raw means and standard errors of given measurements and tests the difference between treatment and control villages. ** denotes significance at 5%. The only significant difference is in School fees (control group is larger). This is due to the presence of one large control village but the difference dissapears once we include district fixed effects (appropriate since the randomization was stratified by district)*

Appendix Figure 1 : Impact of Report Card on Quality Perception

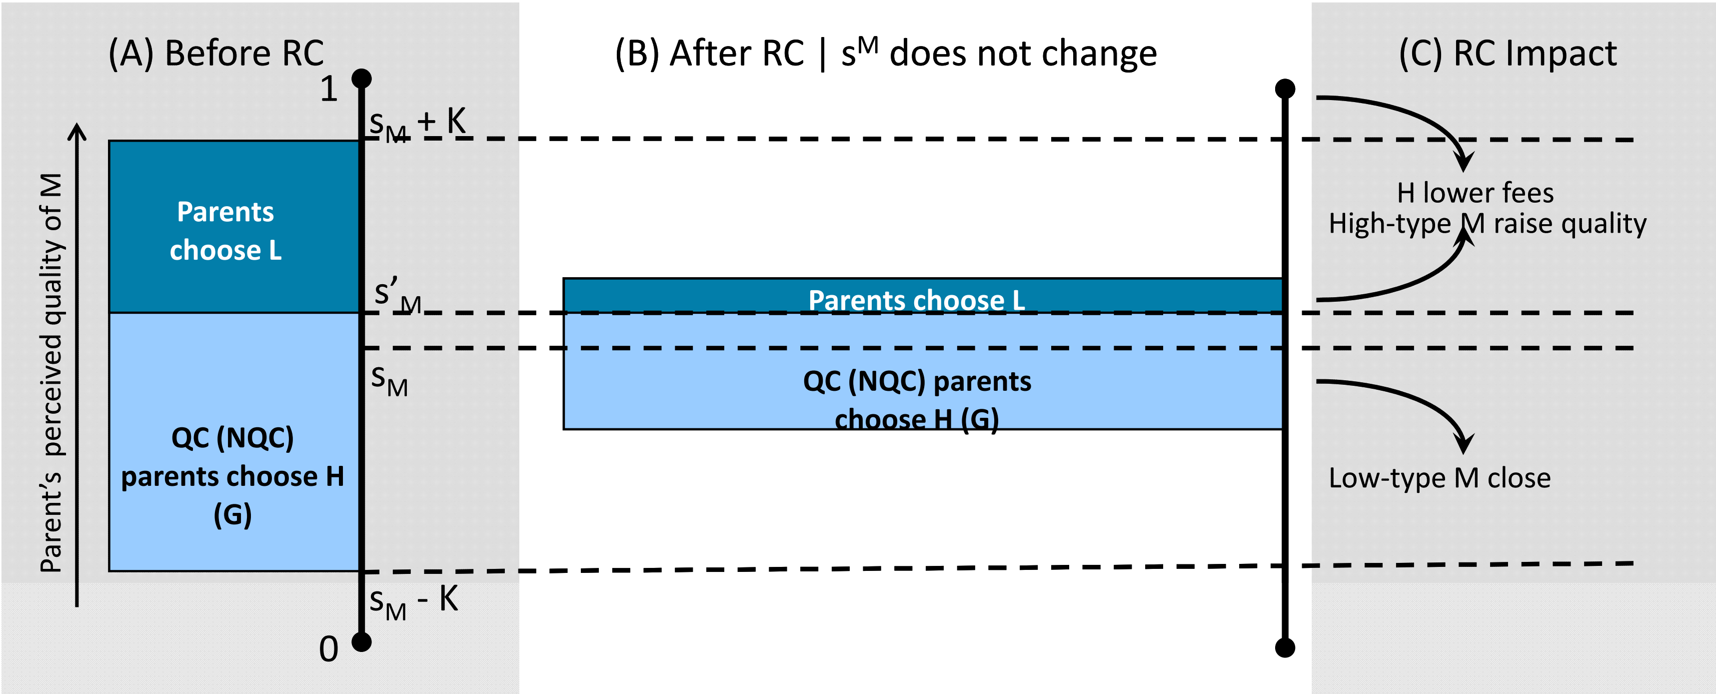

This figure demonstrates a possible impact of the report card intervention for the model presented in the Appendix. Panel A shows the equilbrium prior to the report card. Quality (not) conscious parents above s'_M all choose school M, while those below the threshold choose school H (G). Note that the cutoff s'_M may be different for quality conscious and not conscious parents Panel B demonstrates the change in perception (increases precision) of middle quality schools by parents due to the introduction of the report card. It also illustrates that, conditional on the M schools not changing quality, they will lose market share (one can show the darker shaded area above s'_M is reduced). Panel C then demonstrates the likely response to this market shift i.e. the adjustment school M makes to maintain market share. High type school M's will raise quality (and H types private schools will drop price as a result f the increased competition) while low types will likely close since they are unable to raise quality and cannot compete with the free public school.