

PRELIMINARY DRAFT

Do Public Agencies Learn from the Experiences of Others? Evidence from Contracting out Comprehensive School Reform

Silke Forbes, UC San Diego

Nora Gordon, Georgetown University and NBER

This version: January 28, 2011

Abstract

When public agencies outsource complex services to private contractors, they are faced with the problem that the quality of these contractors can be hard to evaluate. We study an empirical setting in which administrators of public schools choose among private consulting firms, and investigate the potential for social learning in a bureaucratic setting. Do administrators have a network of peers upon whose experiences they draw in making their own choices? If so, who are these peers, and how do administrators interpret data from them? In particular, do they consider the correlations between choice of contractor and raw (or statistically adjusted) achievement levels (or changes)? We find that schools do appear to be influenced by the experiences of peers within their own districts, in close geographic proximity (in a district within a 50 mile radius of their district), or in a common regional district responsible for professional development and technical assistance. We find no evidence suggesting that schools make use of the full set of publicly-available data statewide.

I. Introduction

Public agencies often rely on private contractors. Examples of this include cities outsourcing trash collection, the military hiring contractors to support combat operations, and states contracting with private firms to create standardized tests for public schools. The literature on public contracting has suggested that complex services whose quality is difficult to assess *ex ante* may be better suited for public provision rather than private provision through contracts (see Poterba (1996), Blank (2000), and Levin and Tadelis (2010)). Nonetheless, many public agencies—in particular educational agencies—have increasingly been choosing “buy” over “make” in such markets. For example, Burch (2006) documents major increases in revenues for for-profit educational services providers in the years immediately following the passage of No Child Left Behind.¹ Public agencies face a challenge in identifying high-quality producers in these markets, especially when they are relatively new, as may occur if demand for the services is prompted by a new law or funding source. Even when quality is difficult to assess *ex ante*, market participants may learn about the quality of individual contractors over time, either from their own experiences or from the experiences of others. In this paper, we study to what extent learning from others takes place and, by implication, to what extent information asymmetries in such markets are overcome.

We examine an empirical setting in which learning about the quality of private contractors can come from two sources: communications with other consumers (social learning), or from accessing and analyzing large publicly-available comprehensive databases about other users’ past experiences. Specifically, we examine decisions of

¹ Using firms’ annual reports, she finds annual revenue increases of 77 percent for firms specializing in test development and preparation, 46 percent for data management and analysis, 300 percent for remedial services, and 150 percent for content area specific programming over the period from 2001 to 2004.

public schools which choose among private consulting firms for a product called Comprehensive School Reform (CSR). We investigate how these choices are shaped by other schools' previous experiences (using changes in test scores as proxies for experiences) with those contractors. The data allow us to compare different sets of peer groups from which individual users could potentially learn: we estimate one set of models limited to learning from schools that are likely "peers" and another including all schools for which information is available from public statewide databases, but only through a concerted effort. We compare empirical models of learning by examining which sources of information influence schools' choices. The question of how schools obtain and process information is particularly timely given the current policy emphasis on encouraging schools to make evidence-based decisions.

Our setting is the Comprehensive School Reform (CSR) program in Texas public schools from 1997 to 2005. CSR was a federal program which awarded one-time grants to public schools to purchase consulting services from private firms or to develop their own local reform plans. Once the program was established many firms, most of them very small, entered the market and began selling their products to public schools.² These products involved a bundle of relatively complex services, such as teacher training and curriculum development, whose quality would have been difficult to assess prior to purchase. As a result, the first wave of schools receiving grants had to choose among products with close to unknown quality. Each school only received a single grant so there were no repeat purchases. Schools in later waves of the grant disbursement, however, could potentially learn from the experience of other schools in previous waves. Schools

² Many of these firms did exist previously, but sold component products separately rather than bundled together as CSR.

chose specific CSR plans, and each plan is provided by its sole firm (the contractor). While we discuss the choice as a contracting decision, it could similarly be viewed as a choice of technology or policy, and the aggregate effects of these choices could be viewed as describing the diffusion of innovation.

We exploit several interesting features of our setting. First, for the state of Texas, schools' student achievement data as well as their choice of CSR provider were public information throughout the entire period of the program. This means that schools had the potential to respond to the perceived effectiveness, as measured by changes in student achievement, of the plans that were chosen by their peers. This is critical in a market where consumers only make one purchase so we observe some transactions that would not recur and exist in equilibrium: schools that are looking for information from their peers may well find out that their peers were dissatisfied with a plan. It also is critical in addressing the "reflection" problem: consumers may choose the same product as their peers not because they are learning from them, but because they have common preferences which are unobservable to the researcher. Our approach reveals whether a school is more likely to adopt a plan, *conditional on past usage by a neighbor*, if it perceives that plan as being useful for the neighbor. Second, because grants were provided in waves, we allow schools only to learn from the *past* experience of their peers. Finally, our setting allows us to examine multiple plausible characteristics which might define peer groups with the potential to influence a school's behavior. We observe CSR choices and measures of plan effectiveness for all schools in Texas; we test for learning across schools within school districts, as well from schools in geographically neighboring districts, demographically similar districts, and districts with which they share

professional development and technical assistance services. While empirical tests of social learning are best structured when the researcher observes the structure of the peer group, in our case the structure of the relevant peer group is unknown and a key question of interest.³

Our first finding is that consumers (here, schools) are influenced by outcomes serving as indicators for the perceived satisfaction of their peers with their choices—that is, with the *experiences* associated with the goods—but only for a narrow definition of peer group, limited to schools in their same district. This suggests both a capacity for social learning within the peer group definition, and that consumers either are not obtaining information from a wider set of peers or that they strongly discount its relevance. The former possibility is consistent with the limited research available on networks of school administrators (Daly and Finnigan, 2009), revealing little communication across school sites. Second, there is a positive relationship between past peer choice of a product and current consumer choice—even an average or somewhat below average experience is associated with a greater likelihood of a peer choosing the same product. We estimate plan choice with controls for plan “fit” with observable school-level demographics, but at least part of this relationship likely reflects the role of unobservable correlated preferences or market conditions within the peer groups of schools as we define them. We cannot rule out, however, the role of “word of mouth”—which both school and plan staff we interviewed suggested was the most important factor

³ It is possible that particular definitions of peer groups are in fact meaningful for schools but that these peer groups do not affect their choices of private contractors (e.g., schools could share information on some programs but not CSR, or could share information on CSR but not use it in making their own decisions). We cannot distinguish such outcomes from the irrelevance of a particular characteristic--such as geographic proximity--in defining a peer group. Because of the substantive empirical literature documenting social learning over known peer groups, however, we are comfortable interpreting our results as informing which characteristics determine boundaries of peer groups that matter for policy adoption.

in choice of plan. Finally, our results indicate that schools are naïve in how they process data about peer satisfaction, with only the rawest outcome measures having an impact on product choice. This too is consistent with existing views of school administrator decision making; see Lang (2010) for a compelling first-person account.

Given the existing research base finding limited cross-campus communications, we interpret our results as supporting a very narrow definition of peer group. Within that definition, our results are consistent with the literature on social learning showing that agents are influenced by relevant peer experiences or outcomes following choices of products or technologies. In the most convincing social learning literature, researchers typically observe peer groups or have strong priors from which to infer them. For example, Munshi (2004) shows that wheat growers in India respond strongly to the experiences of their geographic neighbors; the effect is weaker for rice, which is a crop that is more sensitive to unobserved farm characteristics so that information from neighbors is less valuable. Conley and Udry (2010) observe communications patterns across pineapple farmers in Ghana, allowing them to identify relevant peers; they then demonstrate that farmers adopt the input choices of their “surprisingly successful” (documented) peers. We believe that our paper is the first in this literature which studies a setting in which learning can potentially happen both through a likely social peer group and through information from publicly available databases, and which explicitly attempts to distinguish learning from these different sources of information.⁴ We show that in such a setting, learning occurs through a narrow peer group rather than through more

⁴ Other settings with publicly available databases are ones which have government-mandated quality-disclosure programs in place, such as health care, restaurant hygiene and airlines. See Dranove and Jin (forthcoming) for a review. These settings pose greater challenges for investigating the role of peer behavior.

costly analyses of large databases; without sufficient statistical expertise, the cost to administrators may seem infinite (while we have not directly observed a market for consultants to help schools choose a consultant, one may well be emerging).

The remainder of the paper is organized as follows. Section II provides institutional background on the Comprehensive School Reform Program. Section III outlines our theoretical framework and section IV presents the data sources and descriptive statistics. Section V presents the estimation and results. A final section concludes.

II. Background on Comprehensive School Reform

The Department of Education (2007) states that “[t]he Comprehensive School Reform program is designed to foster coherent schoolwide improvements that cover virtually all aspects of a school's operations, rather than piecemeal, fragmented approaches to reform.”⁵ CSR was promoted by New American Schools (NAS), a joint public-private venture beginning in 1991 (see Berends *et al.* (2001) for more on NAS). A federal effort to bring CSR to scale began with the Comprehensive School Reform Demonstration grants in 1997, and became the Comprehensive School Reform grants program in 1998. Under these programs, public schools could apply for competitive federal grants to purchase CSR plans from private CSR providers or for school districts to develop programs locally. The program was last funded in FY2005; in most years, about \$200-300 million was appropriated for the entire program nationwide.⁶

⁵ Department of Education (2007) further specifies eleven elements that are required in a CSR program.

⁶ Funds were appropriated for data collection through the CSR Clearinghouse in FYs 2007 and 2008, but not for new program implementation in those years.

In practice, a wide variety of programs has been considered CSR and received federal funding. These programs range from highly specific interventions to those in which it is difficult to detect what exactly the CSR treatment is. At one end of the CSR spectrum are programs like the Coalition for Essential Schools: on its website, it describes “school reform as an inescapably local phenomenon” and acknowledges that “no two Essential schools are alike.” It instead emphasizes the shared commitment to a set of principles that are both relatively uncontroversial and vague, such as personalized instruction and an atmosphere of trust and high expectations. In contrast, Success For All (SFA) is an example of a tightly scripted program, with teachers adhering to a prescribed curriculum and methods, rhyming mnemonics included. In order to create this uniformity, SFA must provide many more curricular, assessment, and training materials than would a plan emphasizing locally specific needs. These examples show that not only do CSR plans differ in their philosophies and methods, but that even the components constituting a CSR plan vary across providers.

Such great variety in CSR program design naturally leads to the question of the *heterogeneity* in efficacy of these programs. There is little existing large-scale quantitative evidence on this question, and none from anything resembling random assignment of providers to schools (see Borman *et al.*, 2003). Mason (2005) investigates the achievement effects of five CSR programs that were implemented in Los Angeles Unified School district. He finds that, of the five programs in his study, none led to uniform improvements in achievement.⁷ We emphasize that we cannot and do not attempt to establish the true efficacy of CSR programs in this paper; rather, we are

⁷ He and others (Vernez *et al.*, 2006) also found that very few schools fully implemented their CSR plans, so we emphasize that the literature on CSR efficacy, including this paper, is studying CSR as implemented in practice rather than the ideal plans as initially developed by the New American Schools research efforts.

interested in how schools *perceive* the efficacy of these programs and how those perceptions shape their consumption choices.⁸

Schools, districts, and states historically have purchased CSR from a variety of funding sources, including but not limited to the federal grants program. There is no publicly available comprehensive accounting of which schools use CSR or of how many schools use particular CSR plans (providers do not wish to share their client lists and public information requests are difficult for contracts older than one or two years).⁹ There is, however, a complete database of all schools purchasing CSR plans with federal funds from 1997 to 2005, so we restrict our attention to these plans. Federal funds were awarded to State Education Agencies, which distributed the funds through competitive grants to school districts that applied on behalf of specific schools within their districts. Applications from schools had to specify the CSR provider that the school intended to use, and grants were awarded for use with this provider only. Award criteria were set by the state. Some of the federal grants were reserved for Title I schools, but other funds could be used for CSR efforts at any public school. Schools could apply for CSR grants only once, and grants were given for one year with the possibility of being extended to up to three years with the same provider (in practice, this was almost always the case). The federal legislation mandated a minimum grant amount of \$50,000 per year. Most states awarded grants in three distinct waves from 1998 and 2005. Each school could only receive a single grant, for up to three years of CSR services.

⁸ We also expect a strong match-specific component to true efficacy of plans, with plans differently suited to different types of schools and student populations. We control for this by interacting plan dummies with school characteristics.

⁹ Datnow (2000) reports over 6,500 schools in 45 states implanting CSR plans (with any funding source) at the time of her study.

The number of firms that were active in the CSR market is very large and the market concentration very low. Nationwide, 755 different plans provided federally financed CSR services to schools over the period from 1998-2005. 443 of these plans (almost 60 percent) only served a single school. Another 80 (10.6 percent) provided CSR services to two schools and 40 (5.4 percent) provided CSR services to three schools. Nationally, the largest CSR plan was Success For All with a market share of 6.9 percent, the second-largest plan was Lightspan with a market share of 4.5 percent, and the third largest was Accelerated Schools with a market share of 3.8 percent. For the state of Texas, there were a total of 124 different plans, and the distribution of market shares was similarly skewed as at the national level. 74 of the Texas plans (again, almost 60 percent) provided CSR to only a single Texas school, 18 plans (14.5 percent) provided CSR to two Texas schools and 6 plans (4.8 percent) provided CSR to three Texas schools. As Table 2A shows, the two largest plans operating in Texas, Accelerated Schools and AVID, had of 17.1 and 13.5 percent respectively of the market for federal CSR grants in the state, and the next largest plan, Co-nect, had only 3.4 percent of the market.

The relevant measure of plan concentration, however, varies by grade level because many CSR plans are geared towards elementary, middle, or high school levels. Table 2B lists the market shares, by level, of the largest plans in Texas. Accelerated Schools, the largest plan, has an even larger market share of the elementary market (21.4 percent) than of the market overall, but the next largest elementary plan, Success For All, has only 5.7 percent of the market. Success For All and Lightspan specialize in elementary grades, while AVID, with over one quarter of the high school market, does

not operate in elementary schools. We therefore define a school's choice set by the grade levels it serves.

Finally, in order to understand factors influencing CSR plan adoption, we must consider the process by which plans were chosen. The most complete academic study of this process is Datnow (2000), with case studies of 22 schools in three states. Her research was conducted from 1995 through 1999, so included several years in which plans were chosen without federal funding. She found that in most districts, the choice of CSR plan was made at the school level, with varying levels of guidance or pressure from the district, while in a few districts, the district essentially dictated the choice of the school-level plan by offering district-level funds for purchasing that plan but not other plans. Because we study schools purchasing CSR with federal grants, the district does not have the ability to shape school choices so strongly in our context. In our own interview with an administrator of a large school district in California, we were told that the district provided guidance on a small number of plans, with the suggestion that schools choose among these plans. Our data show that, in Texas, the largest districts—which are the districts with multiple schools choosing CSR in each wave—typically have a small number of different plans chosen across schools (see Appendix Table 1). An interesting exception is Austin Independent School District, in which all schools choose the same plan in the first wave, but schools in later waves choose a number of different plans.

Because we observe only those grants which are funded, rather than all applications, we have investigated the process by which applications were selected in this policy setting. From the perspective of our methodology, the results are heartening: the

state education agency staff member in charge of the program reports that no particular models of school reform were given preference in the selection process (based on scientific evidence, perceived efficacy within the state, or other factors). While federal requirements post-No Child Left Behind mandated that funds be used for models with positive evidence from “scientifically-based research” the dearth of such research on these plans makes the requirement irrelevant in practice. For example, in the 2004 request for applications for the CSR – Texas High Schools Initiative Program, the application specifies that two out of 100 points are to be awarded based on the extent to which “The design of the proposed project reflects up-to-date knowledge from scientifically-based research and effective practice.”¹⁰ Overall 15 of 100 points that year were awarded based on the larger category of “quality of project services,” further supporting anecdotal evidence that applications were chosen based on the *school’s* level of need and capacity to implement CSR plans from any given provider, rather than based on the particular provider selected.

II. Theoretical Framework

We assume that school administrators have the objective to improve their students’ achievement. This measure is consistent with an objective function formed by altruism as well as with one influenced by incentives imposed by accountability regimes. We assume that a given CSR plan has an average (however noisy) effect on achievement, as well as a school-specific match component to its efficacy. Administrators form an expectation of both effects of each CSR plan of which they are aware. In our model, the

¹⁰ Instructions in other rounds of the grant do not break down points in such a refined way, but the allocation of points across larger categories is constant.

match quality is based on observable characteristics of the school, most importantly the school's initial achievement level. We do not claim to know how administrators form these expectations; to the contrary, our goal is to identify patterns in the data consistent with both various forms of inference and various sources of data. One interesting feature of our setting is that, in principle, school administrators in the later waves have access to a rich set of information that could potentially inform their choice of CSR provider. The timing of the CSR grants in Texas was such that more than 60 percent of the schools ever awarded grants could observe at least one year of test score data from other schools implementing CSR with earlier grants.¹¹ Achievement data and accountability ratings for each school in Texas for each year were public information, as was the allocation of CSR providers to schools. While the association between an endogenously-chosen CSR plan and raw variation in pass rates or accountability ratings over the first year of CSR plan adoption is far from a causal estimate of the plan's impact on achievement, it is the information that is most readily available to school administrators, if they choose to use any quantitative data at all.¹²

We focus on two dimensions of inference: First, what defines a “peer” school—e.g., one from whom administrators at a given school may learn? We compare models in which (i) peer schools are defined as those in the same district, (ii) we define multiple types of peer schools based on a variety of geographic, demographic and administrative

¹¹ The first year of implementation (for first-wave schools) was the 1999-2000 school year, and accountability ratings for that year were released on August 17, 2000. The deadline for applications for the second wave of grants was December 15, 2000, and the corresponding first year of implementation was 2001-2002. Accountability ratings from the first year of implementation for second-wave schools were released on August 1, 2002, and the applications for the final wave of grants were due March 15, 2004 and October 7, 2004 (the later deadline was for high schools).

¹² There are some education policy groups which rank CSR plans. Because there are so many plans, and the studies have typically examined plans based on national market share, there are plans which are major suppliers in Texas but have no formal rating in such schemes. We therefore do not test for the relevance of such outside ratings.

characteristics, and (iii) all schools in Texas are considered peers. We are also interested in *what* administrators learn from their peers. While still not a causal estimate, we do think that the correlation between plan choice and achievement levels is likely to be more informative when it is based on more observations (statewide data, not just from a potentially very small number of peers) and adjusted for school demographics. We create such an estimate and allow schools to observe it if and only if their peer schools have used the plan in the past. We compare this model to one in which schools observe only the actual experience of the peer schools with the plan.

IV. Data Sources and Descriptive Statistics

We focus our analysis on a single state, Texas, for several reasons. Most importantly, because plans align themselves to some extent with state educational standards, there are significant cross-state differences in plan popularity. We assume that any school in Texas might choose a plan which is used with high frequency and geographic variation elsewhere in the state, but we would not want to assume that any school might choose a plan popular in the rest of the country but not in its own state. Potential differences in how CSR grants were awarded further motivate focusing on a single state to reduce the noise in our estimations. Given that Texas is a large state that received many CSR grants, restricting our attention to this state still leaves us with a sufficient number of observations. In addition to its size, Texas has the benefit of an early standardized testing and accountability regime. The state began testing students and collecting achievement data before the start of the federal CSR program. We can thus establish a school's initial achievement level before CSR intervention. Furthermore,

Texas has accountability standards during this period, which simplifies our interpretation by making achievement improvements an explicit objective; the accountability ratings themselves are also an additional useful measure for our analysis.

We use three main data sources. The first of these is the Comprehensive School Reform Awards Database from which we obtained our data on federal CSR grants in 2005. This database was collected and maintained by the Southwest Educational Development Laboratory (SEDL), a largely federally funded not-for-profit research firm. SEDL collected these data from state CSR administrators and maintained an online database.¹³ The federal Department of Education later took over responsibility for this database, and then removed it from the Internet.¹⁴ The CSR Awards Database is organized at the school level of observation; the grants were three years in duration, and no school received more than one grant during the seven year span of the program, so this is essentially a school-grant level of observation. For each school, we use data on the “models used” (equivalent to the CSR firm with whom the school contracted, and in most cases, only one model) and the first year of the grant.

Each school is identified with its unique seven-digit National Center for Education Statistics (NCES) identification number, allowing us to link these data with the NCES Common Core of Data. From the Common Core, we use data on the racial and ethnic composition of each school, as well as data on school enrollment. Table 1A provides descriptive statistics comparing school demographics for schools that receive CSR to schools that do not receive CSR. We see that CSR schools have a lower

¹³ For some of the schools in this database, the name of the CSR provider was missing. We obtained this information through a Public Information Request from the Texas Education Agency.

¹⁴ See <http://csrprogram.ed.gov/> for details on accessing the offline data. All data in the paper are public and available from the authors upon request.

percentage of white students and higher percentages of black, Hispanic and Limited English Proficient (LEP) students. CSR schools are similar to non-CSR schools in their percentage of special education students. Finally, CSR schools have more students and are in larger districts (as measured by the number of schools in the district).

Our final data sources are from state agencies. The Texas Education Agency's Academic Excellence Indicator System contains annual information on each school in Texas, including the percentage of participating students at each school who passed the statewide achievement test in that year ("pass rate")¹⁵ and the percentage of students at the school who are exempt from the test.¹⁶ Pass rates are reported separately for math and reading. We use the average of those two pass rates as our measure of school performance. Table 1B compares the distribution of pass rates in 1994 of schools which received CSR funds to schools that did not receive CSR funds. We observe that CSR schools have lower pass rates than other schools at all points of the distribution. However, there is substantial variation in pass rates even among CSR schools. While CSR schools at the fifth percentile of the distribution have pass rates of 38.2 percent, compared to 43.5 percent for non-CSR schools, CSR schools at the ninety-fifth percentile have pass rates of 81.3 percent, compared to 89.4 percent for non-CSR schools.

To more fully explore the mechanism for dissemination of information, we have gathered data on the composition of Education Service Centers (ESCs) within Texas. These centers are essentially a regional layer between the state and local education agencies, and primarily provide professional development and technical assistance.

¹⁵ Pass rates describe only one part of the distribution, and administrators may well care about other points in the distribution more (for example, if the parents of higher achieving students are more politically vocal). The pass rate is the most consistently reported measure in the data, however.

¹⁶ Cullen and Reback (2006) provide evidence suggesting that schools use exemptions strategically to improve their accountability ratings.

There are 20 centers statewide, with each center serving multiple (contiguous) counties. They therefore constitute networks of administrators that are both plausible (teachers and administrators may attend professional development programs physically provided at the ESC) and observable.

As in most states, CSR grants in Texas were awarded in three waves. While some CSR firms provided services (bundled together as CSR or more likely as separate components) prior to the first wave, no publicly-available information links individual schools to providers of any such services in the state.¹⁷ As a result, schools in the first wave have no observable information about any of the providers. Schools in the subsequent two waves have the potential to learn about the quality of CSR providers from the experience of other schools within Texas in previous waves. Recall that each school only receives a single grant and thus has no opportunity to react to the information that they learn from their own experience with a provider. There are a total of 433 schools in our data set that get CSR grants. 121 of them get a grant in the first wave (1999-2000 academic year), 166 get a grant in the second wave (2001-2002 academic year) and 146 get a grant in the last wave (2004-2005 academic year).

In constructing our regression sample, we start with all schools in Texas which received a federal CSR grant. We drop any schools that are special education, vocational, or alternative. We also drop schools which are missing in one of our three databases. Next, we drop all schools in four districts in El Paso because El Paso had its own CSR initiative which schools in these districts were very likely to choose. Finally,

¹⁷ Other states had already awarded grants in 1998 and 1999 to schools using some of the same firms that were active in Texas. However, the largest plans in Texas were represented sparsely among those early awards, and those schools choosing the major Texas plans were spread across states making data analysis difficult. In particular, no schools chose AVID before Texas' grants program began.

we only allow plans to be in a school's choice set in years in which we see at least one Texas school choosing this plan.

Our model includes explanatory variables that are specific to the school making the choice and explanatory variables that are specific to the plan that is chosen. The former group consists of demographic characteristics of the school, described in Table 1, Panel A. The average school was 29 percent white, 19 percent limited English proficient, and enrolled about 870 students; its pass rate on the statewide test in 1994 (or in the first period that we observe the school after 1994) was 60 percent. A third of the grants went to schools in the largest 15 districts in Texas.¹⁸

The values of our plan-specific explanatory variables vary by school, conditional on plan, and relate to the different concepts of peer or neighboring school. We include these variables to reflect both the transmission of information (if "peer" administrators share access to qualitative or quantitative data) and the issue of "fit" (even if administrators have complete information, they may weight it more heavily if generated in a context similar to their own) but do not have any way of disentangling these two effects. In this vein, we have indicator variables for whether the plan was used in the past by another school in the same district (this was the case for 14.8 percent of plans chosen) or a school in a geographically neighboring district, defined as one with its centroid within a 50-mile radius of the centroid of the school's own district (true for 23.5 percent of plans chosen). Panel A of Table 3 shows the overlap between existing experience in the school's own district and the neighboring districts and illustrates that schools are more likely to choose plans used in their own district or a geographically neighboring district. They are also appear more likely to use plans used by "peers" in their same ESC

¹⁸ The 15th largest district has 41,000 students; the 16th largest district has 32,000 students.

district. In alternate definitions of peers or neighbors, we limit our sample to schools in very large districts (this allows Dallas to care most about Houston, for example) or those with unusually high shares of English-language learners, and allow them to learn only from schools in districts meeting those same criteria.

In our preferred specifications, we also control for the plan’s perceived effect on pass rates for school(s) in the peer group that used the same plan before. Panel B of Table 3 shows that the direction of these effects (schools are more likely to choose plans that have increased achievement among the relative peer group) is consistent with our hypothesis, but they are quite noisy.

V. Empirical Approach and Results

We estimate school plan choice using a conditional logit specification. The available set of options depends on the grade level of the school, as well as the year of the grant. We exclude plans from the choice set in years in which they are not chosen by any school in the state. Because we have a large number of plans that are chosen by very few schools, we do not include all of them as separate choices in the conditional logit model that we estimate. Instead, we define a category that we call “other” for plans that are infrequently chosen at that grade level. The choice set of plans for a given school thus consists of a number of individual large plans and “other.” We construct our definition of a large plan, somewhat (necessarily) arbitrarily, to conform with our intuition about which plans (1) are not supply-constrained and therefore likely constitute a viable option for a school with a new grant,¹⁹ and (2) have had the opportunity to generate results

¹⁹ Nearly all of these “large” providers operate in multiple states, and each of them operates across multiple non-contiguous Texas counties.

which new grantees may observe and use to inform their plan selection. In the estimation, we treat “other” essentially like an outside good. Table 2 shows the market share for each of the 21 sufficiently “large” plans, and reveals that 34 percent of all schools and 36 percent of elementary schools in our sample chose plans in the “other” category.

Plan choice is determined by both school and plan attributes. Because the plan attributes are difficult to observe and quantify for all plans, we use plan-level fixed effects. We allow plan choice to be affected by school demographics, as well as school-level potential exposure to information about the plan’s observed correlation with achievement changes in other schools.²⁰ Depending on our specification, we allow exposure to information to come from various sets of schools, starting with other schools in the same district, then schools in neighboring districts, schools in common administrative regions, and schools in demographically similar districts.

Panel A of Table 4 begins by allowing information to flow only within districts. All specifications throughout the table include the school’s initial pass rate as a control, but to save space we do not report the coefficients in the table. Column A1 begins with a simple specification that includes a dummy for whether another school in the district has used this plan in the past. We find that the coefficient is positive and highly significant, indicating that schools are more likely to choose plans that have been used in the district before. Of course, such an effect could arise for several reasons. Besides the fact that the school may have better information about a plan that has been used in the district before, the district itself may also have a preference for contracting repeatedly with the same

²⁰ We treat the prior experience with a plan as a plan attribute that varies across schools in the conditional logit model.

provider. For example, there may be bureaucratic processes which need to be established by the district only once per provider, or the district may find itself in a better bargaining position if it can offer the CSR provider business with multiple schools. We also might observe this effect if schools within a district share unobservable characteristics influencing their choice of plan. (Recall that the school-level characteristics we do observe and control for include not only basic demographic data but also the school's initial pass rate on the state test.)

In Columns A2 through A5, we include various measures of perceived plan effectiveness for plans used in the same district before, in addition to the dummy for past use of the same plan. This allows us to separate the district's possible incentive to contract repeatedly with the same provider from the type of information that may have been learned from the plan's past presence in the district. In column A2, we control for the raw change in pass rates at schools in the district that used a plan in the past, capturing the potential for administrators to learn from their peers about their specific experience. In this specification, we find the plan's past effect on raw pass rates has a positive and highly significant effect. This result suggests that schools are more likely to choose plans that were used by other schools in the district in the past because of what they perceive to be the efficacy of these plans rather than simply name recognition. (The fact that the past presence dummy in A1 was positive is unsurprising given the positive correlation between the use of any plan and pass rates.)

Table 5 allows us to interpret these coefficients by predicting the effects of past presence and past perceived achievement effects of the two largest plans in the state, AVID (Panel A) and Accelerated Schools (Panel B), on plan adoption. For both plans

(see column A1 in Tables 5A and 5B), past usage of a plan within a district more than doubles the probability of future adoption for a given grant. Column A2 shows that pass rates one standard deviation above average (among CSR plans) further increase probability of choosing both AVID (from 41 to 56 percent conditional on past use in the district) and Accelerated Schools (from 29 to 42 percent) again.

Because Texas was adjusting to a new test (the TAAS) and a new accountability system in this period, scores on the state test were generally improving (although scores on a low-stakes national test, the National Assessment of Educational Progress, were not). We therefore standardize pass rates within each year and in specification A3 of Table 4, use data on changes in standardized pass rates. Reassuringly, these are statistically significant and positive as well—it is not simply the case that plans adopted in years with greater overall raw test score gains are more likely to be adopted in the future. In column A4 we present results from a specification using levels of pass rates rather than changes, and cannot rule out the possibility that administrators base their choices on observed *levels* alone.

In specification A5, we investigate the possibility that your peers using a plan in the past gives you information more closely approximating the true quality of the plan—better information than simply the change in test scores experienced in that limited set of peer schools. This model is observationally equivalent to one in which administrators are aware of particular plans because they have operated nearby, and then go to two distinct data sources (on Texas achievement levels and federal CSR grants), merge them, and estimate a regression using statewide data, including control variables. In other words, we add a variable interacting peers' past use with what we calculate as a regression-

adjusted correlation between that plan and test scores based on data from the entire state. Despite the fact that these correlations are based on more observations and remove potentially important confounders, we find no support for this model. (In results not shown, less sophisticated adjustments, using statewide data without control variables, were also uninformative.)

In Panels B and C of Table 4, we allow schools to learn from peers from outside of their own districts. Because we expect information sharing within districts to be stronger, we estimate the impact of experience in neighboring districts separately, while continuing to control for past experience (plan presence and perceived efficacy) within the school's own district. Based on the high correlation between raw changes, standardized changes, and levels of pass rates, we present results going forward for changes in raw pass rates only, as this seems most likely in a relatively but not completely statistically naïve population and the different concepts yield similar results regardless. Panel B shows results for neighbors defined geographically, using the 50-mile radius definition, and administrative, based on assignment to regional Education Service Center. As within district, the past presence of a plan in a neighboring district is statistically significantly positively associated with the choice of the plan, but we again do not infer any causal impact from this relationship, as it may capture unobservable characteristics associated with plan choice and correlated within peer groups. We find that schools are more likely to choose plans which they perceive to have served their geographic and administrative neighbors well in terms of test score gains. In future work, we plan to test the relative importance of these effects by considering the role of common ESC membership, controlling for geographic distance. Table 5 again interprets these

estimated logit coefficients. Panel C of Table 4 shows that the positive impact of test score increases at geographically or administratively “neighboring” schools is not present when we adapt the “neighbor” concept for the unique circumstances of districts that are unusually large or have high shares of English-language learners.

In results not shown, we see no support in any of the neighbor concepts in Panels B or C for a model in which a peer’s use of a plan facilitates further inference based on statewide data and statistical adjustments for observable characteristics and time trends. Finding no relationship between statewide achievement estimates of locally salient plans on plan choice is consistent with two very different models of decision making or two coexisting but very different types of school administrators: administrators may not have the capacity to access school-level micro-data from two sources, merge them together, and analyze them; alternatively (at least some) administrators with the capacity to conduct such research themselves or request it from another source do not think the exercise would be informative. Our data do not allow us to distinguish between these two scenarios, but we wish to emphasize important differences in policy implications between them nonetheless. If administrators simply lack low-cost access to data and analysis, there is a clear role for state and/or federal intervention in providing more accessible data and research findings to schools choosing providers. If, however, school administrators—like researchers—doubt that raw changes in pass rates capture the causal impact of plan adoption, the policy “fix” is much more complex. As one teacher said, “If I saw that a CSR plan were associated with a big increase in test scores at a given school, I would think that reflects more on the school’s capacity to choose and implement the plan than on the plan itself,” (personal communication, Clark, 2010). While it appears

administrators do use information available from their likely peers, information is then limited to those plans implemented by one's own peers.

Overall, the results of our empirical estimation suggest that schools do consider local information when choosing their CSR providers. We cannot rule out, however, that schools have access to information from outside their own peer group but strongly discount its relevance, perhaps due to the importance of school-provider match quality unobservable to the researcher.

VI. Conclusion

The question of how schools choose contractors is more salient now than ever due to an exponential increase in federal funds potentially available to education contractors, most notably \$4.35 billion in Race to the Top funds and \$3.5 billion in Title I School Improvement Grants (SIGs). This rapid expansion has created unprecedented demand for educational contractors. Rudy Crew, the former New York City schools chancellor and now a contractor paid from SIG funds, characterized the supply response to SIGs as “like the aftermath of the Civil War, with all the carpetbaggers and charlatans,” (Dillon 2010). Unsurprisingly, some firms previously receiving CSR grants are now also marketing themselves as school turnaround and transformation (specific options available to schools with SIGs) specialists.²¹

More broadly, our findings relate to a topical set of questions about how policymakers can help agencies—or households or firms—optimize through information and regulation of choice sets. We find that not only do peers' choices matter, but that

²¹ See, for example, Success for All <<http://www.successforall.net/Turnaround/turnaround.html>>.

their specific experiences matter as well, highlighting the importance of using outcome data instead of simply observing peers' prior choices. That schools appear to not obtain (or disregard) the vast majority of data, even when it is available publicly and centrally through state and federal databases, suggests a potentially productive role for government in making information more readily available and interpretable to schools (again, with the caveat that schools may be disregarding this information because it in fact is not predictive of what their own experience would be with a given plan). Research to date has focused on making information more readily accessible to consumers and households. For example, Hastings and Weinstein (2008) find experimental support for parental responsiveness to highly accessible information on school quality that is handed out to parents in one to three page hard copy documents, rather than available online for those who seek it out, and a growing research literature and policy debate emphasize the importance of communications and framing in personal financial decision-making. Our research suggests that public agencies, like individuals, potentially could benefit from informational interventions. These interventions naturally would be most useful in cases with sufficient data and variation to yield reliable inferences. With research funding from the Department of Education now focused on credibly identified variation, the research base on particular programs and providers may grow sufficiently to allow existing language limiting the use of federal funds to educational interventions supported by "scientifically-based research" to become meaningfully enforced.

References

- Berends, Mark, Sheila N. Kirby, Scott Naftel, and Christopher McKelvey. (2001). *Implementation and performance in New American Schools: three years into scale-up*. Washington, DC: RAND.
- Blank, Rebecca M. 2000. "When Can Public Policy Makers Rely on Private Markets? The Effective Provision of Social Services." *The Economic Journal* 110(462), C34-C49.
- Borman, Geoffrey D., Gina M. Hewes, Laura T. Overman, and Shelly Brown (2003). Comprehensive school reform and achievement: a meta-analysis. *Review of Educational Research* 73(2), 125-230.
- Burch, Patricia E. (2006). The new educational privatization: educational contracting and high stakes accountability. *Teachers College Record* 108 (12), 2582–2610.
- Conley, Timothy G. and Christopher R. Udry. (2010). Learning about a new technology: pineapple in Ghana. *American Economic Review* 100(1), 35-69.
- Cullen, Julie Berry and Randall Reback. (2006). Tinkering toward accolades: school gaming under a performance accountability system. In: T. Gronberg and D. Jansen, eds. *Improving School Accountability: Check-Ups or Choice, Advances in Applied Microeconomics* 14 (Amsterdam: Elsevier Science), pp.1-34.
- Dranove, David and Ginger Z. Jin. (Forthcoming). Quality Disclosure and Certification: Theory and Practice. *Journal of Economic Literature*.
- Datnow, Amanda. (2000). Power and politics in the adoption of school reform models. *Educational Evaluation and Policy Analysis* 22(4), 357-374.
- Department of Education. (2007). About CSR. <http://www.ed.gov/programs/compreform/2pager.html>, accessed November 29, 2007.
- Dillon, Sam. (2010). "Inexperienced companies chase school funds." *New York Times*. August 10, 2010, A11.
- Dillon, Sam (2010). "U.S. awards millions for school innovations." *New York Times*. August 6, 2010.
- Hastings, Justine S. and Jeffrey M. Weinstein (2008). Information, School Choice, and Academic Achievement: Evidence from Two Experiments. *Quarterly Journal of Economics* 123 (4), 1373-1414.

- Lang, Kevin (2010). Measurement matters: perspectives on education policy from an economist and school board member. *Journal of Economic Perspectives* 24, 167-182.
- Levin, Jonathan and Steven Tadelis. (2010). Contracting for Government Services: Evidence from U.S. Cities. *Journal of Industrial Economics* 58(3), 507-541.
- Manski, Charles F. (2000). Economics of social interactions. *Journal of Economic Perspectives* 14(3), 115-136.
- Mason, Bryce (2005). Achievement Effects of Five Comprehensive School Reform Designs Implemented in Los Angeles Unified School District. *Pardee RAND Graduate School dissertation*.
- Munshi, Kaivan (2004). Social learning in a heterogeneous population. *Journal of Development Economics* 73, 185-213.
- Poterba, James. 1996. "Government Intervention in the Markets for Education and Health Care," in *Individual and Social Responsibility*, Victor Fuchs, editor (Cambridge, MA: National Bureau of Economic Research), 277-307.
- Sorenson, Alan T. (2006). Social learning and health plan choice. *RAND Journal of Economics* 37 (4), 929-945.
- Vernez, Georges, Rita Karam, Louis T. Mariano, and Christine DeMartini. (2006). Evaluating Comprehensive School Reform models at scale: Focus on implementation. RAND Corporation Monograph.

Table 1: Comparing CSR and non-CSR elementary schools in Texas

Panel A: Means and standard deviations of selected variables (spring 1998)

Variable	CSR schools	Non-CSR schools
Percent White	29.4 (28.4)	51.4 (31.4)
Percent Hispanic	48.8 (34.0)	33.3 (30.5)
Percent Black	20.0 (24.4)	13.2 (19.3)
Percent Asian	1.6 (3.0)	1.8 (3.7)
Percent Limited English Proficiency	18.5 (20.0)	11.3 (16.8)
Percent Special Education	12.5 (4.7)	12.9 (4.9)
School enrollment (in thousands)	0.867 (0.633)	0.599 (0.439)
Number of schools in the district	51.2 (72.5)	35.5 (60.7)
Large district dummy*	0.347 (0.477)	0.220 (0.414)
Observations	372	5565

* Large districts are defined as districts with at least 40,000 students. There are 15 such districts.

Notes: Table shows variable means. Standard deviations in parentheses.

Sources: CSR Awards Database, NCES Common Core

Panel B: Distribution of pass rates (spring 1998)

Percentile of passrate	CSR-schools	Non-CSR schools
5	38.2	44.5
25	51.4	61.9
50	60.2	71.5
75	69.9	79.9
95	81.0	90.4
Mean	60.3	70.1

Source: CSR Awards Database, Texas Education Agency Academic Excellence Indicator System.

Table 2A: Plan choices by wave

CSR Plan	Year of Grant				Market Share	AIR rating of elementary CSR plans
	2000	2002	2005	Total		
Accelerated Schools	30	29	26	85	17.1%	moderate
AVID	25	18	24	67	13.5%	
Co-nect	4	10	3	17	3.4%	limited
Success for All	5	9	2	16	3.2%	moderately strong
High Schools That Work	0	5	9	14	2.8%	
Lightspan	6	8	0	14	2.8%	
HOSTS	0	11	2	13	2.6%	
Coalition of Essential Schools	4	7	0	11	2.2%	zero
Direct Instruction	9	1	0	10	2.0%	moderately strong
Modern Red Schoolhouse	2	1	5	8	1.6%	limited
Roots and Wings	7	1	0	8	1.6%	
Creating Independent Student-owned Strategies	6	1	0	7	1.4%	
Literacy Collaborative -- Ohio State University	4	3	0	7	1.4%	limited
Breaking Ranks	0	3	3	6	1.2%	
Effective Schools	3	3	0	6	1.2%	
Integrated Thematic Instruction	0	6	0	6	1.2%	limited
Reading Recovery	0	6	0	6	1.2%	
Renaissance Learning	2	1	3	6	1.2%	
Success-in-the-Making	0	5	1	6	1.2%	
Target Teach	0	6	0	6	1.2%	
Capturing Kids Hearts	0	3	2	5	1.0%	
other	29	64	80	173	34.7%	
Total	137	201	160	498	100.0%	

Sources: CSR Awards Database, AIR.

Note: zero rating means no evidence "was of sufficient quality to be counted as reliable evidence."

Table 2B: Plan choices by grade level

CSR Plan	Grade Level			Total
	Elementary	Middle	High	
Accelerated Schools	45	17	17	79
AVID	0	28	38	66
Co-nect	6	7	4	17
Success for All	12	3	1	16
Lightspan	11	3	0	14
HOSTS	9	2	2	13
High Schools That Work	0	1	11	12
Coalition of Essential Schools	5	4	2	11
Direct Instruction	2	7	1	10
Modern Red Schoolhouse	3	2	3	8
Roots and Wings	7	1	0	8
Creating Independent Student-owned Strategies	3	3	1	7
Literacy Collaborative -- Ohio State University	7	0	0	7
Breaking Ranks	0	0	6	6
Effective Schools	3	2	1	6
Reading Recovery	6	0	0	6
Success-in-the-Making	4	1	1	6
Target Teach	3	2	1	6
Capturing Kids Hearts	0	1	4	5
Integrated Thematic Instruction	5	0	0	5
Renaissance Learning	3	1	1	5
other	76	37	51	164
Total	210	122	145	477

Note: Number of observations is lower in this panel because grade level is missing for some schools.
Source: CSR Awards Database.

Table 3: Prior Experience with CSR Plans
(Level of observation is school-plan combination; 21 plans)

Panel A: Choice by Prior Experience

Neighbor concept	Geographic (50-mile radius)		Administrative (Educ. Svc. Ctr.)		Demographic (large LEAs)		Demographic (high LEP LEAs)	
Variable	Share of all school-possible observations	Percent for which choice=1	Share of all school-possible observations	Percent for which choice=1	Share of all school-possible observations	Percent for which choice=1	Share of all school-possible observations	Percent for which choice=1
Schools with no prior experience with this plan in own or neighboring district	88.27	79.48	96.24	79.68	64.78	64.19	61.21	58.15
Schools with prior experience with this plan in own district, but not in neighboring district	2.09	3.82	2.41	5.03	0.85	2.41	1.41	4.83
Schools with prior experience with this plan in neighboring district, but not in own district	8.94	11.87	0.97	11.67	32.43	27.16	36.00	33.20
Schools with prior experience with this plan in own and neighboring district	0.70	4.83	0.39	3.62	1.94	6.24	1.38	3.82

Notes: Panel shows variable means. Standard deviations in parentheses. Choice equals 1 if the school chooses the plan, 0 otherwise.

Table 3: Prior Experience with CSR Plans
(Level of observation is school-plan combination; 21 plans)

Panel B: Perceived Effectiveness of CSR Plan, conditional on having prior experience in own/neighbor district

Variable	All	Choice=1
<i>Schools that have prior experience with this plan in own district</i>		
Change in raw pass rates	1.21	3.08
	(2.49)	(2.72)
<i>Schools that have prior experience with this plan in neighboring district</i>		
<i>Neighbor concept: geographic, 50-mile radius</i>		
Change in raw pass rates	0.28	1.62
	(1.13)	(2.30)
<i>Neighbor concept: administrative; districts in common Education Service Center</i>		
Change in raw pass rates	0.33	0.40
	(1.54)	(1.25)
<i>Neighbor concept: demographic; large (>40,000 students) districts</i>		
Change in raw pass rates	1.62	1.91
	(3.12)	(2.25)
<i>Neighbor concept: demographic; high share LEP districts</i>		
Change in raw pass rates	2.31	2.86
	(5.82)	(6.77)

Notes: Panel shows variable means. Standard deviations in parentheses. Choice equals 1 if the school chooses the plan, 0 otherwise.

Table 4: Perceived achievement effects and plan choice

Panel A: Effect of perceived achievement effects in own district

Perceived effect measured by...	(A1)	(A2)	(A3)	(A4)	(A5)
Past presence of plan in own district	1.246**	0.790+	1.076*	-0.576	1.261+
	(0.401)	(0.474)	(0.452)	(0.677)	(0.716)
Change in raw pass rates (own district)		0.243*			
		(0.119)			
Change in pass rates standardized by year (own district)			1.471*		
			(0.712)		
Level of raw pass rate (own district)				0.045**	
				(0.009)	
Past presence in own district X statewide predicted value-added					-0.093
					(3.315)
Pseudo-R2	0.2678	0.2726	0.2692	0.2867	0.2678
Observations	444	444	444	444	444

Panel B: Effect of perceived achievement effects in own districts and geographic/administrative neighbors

	(B1)	(B2)	(B3)	(B4)
<i>Neighbor concept:</i>				
	<i>Geographic neighb.</i>		<i>Admin.: Education</i>	
	<i>(50 miles)</i>		<i>Service Centers</i>	
Past presence of plan in own district	1.140**	0.676	0.886	0.354
	(0.378)	(0.470)	(0.603)	(0.559)
Change in raw pass rates (own district)		0.211+		0.261**
		(0.116)		(0.086)
Past presence of plan in neighb. district	1.907**	1.609**	3.128**	3.011**
	(0.400)	(0.436)	(0.422)	(0.404)
Change in raw pass rates (neighb. district)		0.451**		0.481*
		(0.087)		(0.188)
Pseudo-R2	0.2885	0.3033	0.3294	0.3373
Observations	444	444	444	444

Panel C: Effect of perceived achievement effects in own district and demographic neighbors

	(C1)	(C2)	(C3)	(C4)
<i>Neighbor concept:</i>				
	<i>Large Districts</i>		<i>High Share LEP</i>	
			<i>Districts</i>	
Past presence of plan in own district	1.781**	1.162+	1.171*	0.784
	(0.596)	(0.704)	(0.543)	(0.612)
Change in raw pass rates (own district)		0.228*		0.210+
		(0.089)		(0.115)
Past presence of plan in neighb. district	-1.585*	-1.738*	-0.859	-0.791
	(0.647)	(0.831)	(0.575)	(0.558)
Change in raw pass rates (neighb. district)		0.075		-0.036
		(0.163)		(0.025)
Pseudo-R2	0.2040	0.2140	0.3010	0.3084
Observations	162	162	153	153

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

Notes: Standard errors are clustered at district level. Each specification includes plan dummies and the school's initial pass rate interacted with plan dummies.

Specifications C1-C3 only include schools in the 15 largest districts.

Sources: CSR Awards Database, Common Core of Data and Texas Education Agency Academic Excellence Indicator System.

Table 5: Predicted choice probabilities based on estimates from Table 4

Panel A: Predicted choice probabilities for AVID

Specification	(A1)	(A2) ¹	(B2)	(B4)
Neighbor concept	None	None	Geographic (50 miles)	Administrative (ESC)
No prior presence of plan in own (or neighboring) district ²	0.194	0.193	0.147	0.148
Plan was used previously in own district	0.456			
Plan was used previously in own district and passrate change was equal to mean		0.411	0.302	0.249
Plan was used previously in own district and passrate change was equal to mean plus one std dev		0.557	0.419	0.386
Plan was used previously in own and neighboring districts and passrate changes were equal to the mean			0.658	0.854
Plan was used previously in own and neighboring districts and passrate changes were equal to the mean plus one std dev			0.761	0.924

Panel B: Predicted choice probabilities for Accelerated Schools

Specification	(A1)	(A2) ¹	(B2)	(B4)
Neighbor concept	None	None	Geographic (50 miles)	Administrative (ESC)
No prior presence of plan in own (or neighboring) district ²	0.125	0.121	0.100	0.090
Plan was used previously in own district	0.333			
Plan was used previously in own district and passrate change was equal to mean		0.285	0.218	0.159
Plan was used previously in own district and passrate change was equal to mean plus one std dev		0.418	0.318	0.263
Plan was used previously in own and neighboring districts and passrate changes were equal to the mean			0.553	0.769
Plan was used previously in own and neighboring districts and passrate changes were equal to the mean plus one std dev			0.672	0.873

¹ Specification A1 does not include changes in raw pass rates; A2 does.

² No neighboring districts for specifications A1 and A2.

Appendix Table 1
CSR Plan Adoption within Large Districts over Time

District and spring of school year plans first implemented	Number of CSR schools, by year and level	School level	Plan adopted
Austin ISD			
1998	3	high	AVID
	4	middle	AVID
2001			
2004	1	elementary	Accelerated Schools
	1	elementary	Cognitively Guided Instruction
	1	middle	AVID
	3	high	AVID
Cypress-Fairbanks			
1998	3	elementary	Literacy Collaborative--Ohio State U.
	2	middle	AVID
	1	high	AVID
	1	high	locally-developed plan
2001	1	elementary	locally-developed plan
	1	middle	AVID
	1	high	AVID
Dallas ISD			
2001	2	elementary	Success for All
2004	3	elementary	3-Tier Intervention Model
	3	middle	AVID
El Paso ISD			
1998	3	elementary	Success for All
			El Paso Collaborative for Academic Excellence
	2	elementary	El Paso Collaborative for Academic Excellence
	1	middle	Cooperative Integrated Reading and Comprehension
	1	middle	Community for Learning
	1	high	El Paso Collaborative for Academic Excellence
2001	1	elementary	Excellence
2004	2	high	missing
Fort Worth ISD			
1998	8	middle	Direct Instruction
	1	middle	AVID
	2	high	AVID
	1	high	Direct Instruction
2001	1	elementary	Direct Instruction
	2	middle	AVID
2004	3	middle	AVID
	3	high	AVID

Sources: CSR Awards Database; authors' Public Information Request (available upon request).

Appendix Table 1 (continued)
CSR Plan Adoption within Large Districts over Time

District and spring of school year plans first implemented	Number of CSR schools, by year and level	School level	Plan adopted
Houston ISD			
1998	2	elementary	Accelerated Schools
	2	elementary	Core Knowledge
	1	elementary	Different Ways of Knowing
	1	middle	Accelerated Schools
	1	middle	Coalition of Essential Schools
	1	middle	Co-nect
	1	middle	Learner-Centered Framework
	1	high	Co-nect
2001	1	elementary	Consistent Mangement and Cooperative Discipline
	1	elementary	Core Knowledge
	1	elementary	Houston Annenberg Challenge
	2	middle	Co-nect
	1	middle	First Things First
	2	high	Co-nect
	1	high	Houston Annenberg Challenge
	1	high	Coalition of Essential Schools
2004	4	elementary	Co-nect
	2	middle	Accelerated Schools
	1	middle	Success for All
	1	middle	School Development Program
	1	high	Project CLEAR
	1	high	Houston Annenberg Challenge
	1	high	Accelerated Schools
	2	high	High Schools That Work
1	high	Co-nect	
San Antonio ISD			
1998	6	elementary	Roots and Wings
	2	elementary	Success for All
	1	elementary	Co-nect
	1	elementary	locally developed
	1	middle	Modern Red Schoolhouse
	2	middle	Capacity Building Model
	1	middle	locally developed
	1	elementary	HOSTS
2001	1	elementary	Quantum Learning
	1	elementary	Project Reach
2004	1	elementary	Project Reach
	2	high	Capturing Kids Hearts
	1	high	AVID
	1	high	missing

Sources: CSR Awards Database; authors' Public Information Request (available upon request).

Appendix Table 2: Value-added regressions, 1994-2002

Dependent variable: Standardized pass rate

Exempt rate	1.656** (0.097)
Percent white	0.019** (0.001)
Percent LEP	-0.021** (0.001)
Initial pass rate	0.037** (0.009)
Enrollment	-0.440** (0.049)
Enrollment squared	0.058** (0.014)

Observations	52352
R-squared	0.771

Note: Estimation includes plan and year fixed effects

Standard errors in parentheses

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