

New York City School Survey 2008-2010: Assessing the Reliability and Validity of a Progress Report Measure

To what extent do you feel supported by your principal?
 Most students in my school treat teachers with respect.
 School leaders communicate a clear vision for this school. My school has high expectations for all students.
My child is safe at school.
 The school contacts me about my child's achievements and successes.
I am safe in my classes.
 My school is kept clean. Most students in my school just look out for themselves.
 Teachers in my school trust each other.
 My child is learning what he or she needs to know to succeed in later grades.
Discipline in my school is fair. **My teachers inspire me to learn.**
 My child's teachers give helpful comments on classwork and tests.
 I have sufficient materials to teach my classes.
I feel welcome in my child's school.

Lori Nathanson
 Rachel Cole
 James J. Kemple
 Jessica Lent
 Meghan McCormick
 Micha Segeritz

The **Research Alliance** for
New York City Schools

New York City School Survey 2008-2010:

Assessing the Reliability and Validity of a Progress Report Measure

Lori Nathanson

Rachel Cole

James Kemple

Jessica Lent

Meghan McCormick

Micha Segeritz

June 2013

ACKNOWLEDGMENTS

Many New York City Department of Education staff played vital roles over the course of this project, particularly Adina Lopatin, Simone D’Souza, Jessica Richman Smith, and Sonali Murarka. We’d also like to recognize former DOE staff, Jennifer Bell-Ellwanger and Martin Kurzweil, who shaped this work in the early days of the partnership. Of course, Lauren Sypek deserves a special thank you, not only for drafting the policymaker perspective presented in the accompanying brief, *Strengthening Assessments of School Climate*, but for her commitment to improving the School Survey through collaboration with many stakeholders, including the Research Alliance.

Thank you to all of the Research Alliance staff who contributed this work. Chelsea Farley edited the report and brief and connected the Research Alliance’s work on the NYC School Survey with similar efforts taking place in other school districts. Shifra Goldenberg copyedited and formatted the brief and this technical report. Her cover design highlights the content of the surveys that are at the heart of this work.

CONTENTS

Acknowledgments	3
1. Survey Development, Context, and Goals.....	5
2. Preliminary Analysis of the NYC DOE School Survey	11
3. Improving the School Survey.....	28
4. Summary and Conclusions.....	37
Endnotes	40
References	41

CHAPTER 1: SURVEY DEVELOPMENT, CONTEXT, AND GOALS

Each spring, the New York City Department of Education (DOE) invites all public school students in grades 6 through 12 as well as parents and teachers throughout the City to complete the School Survey. In 2012, 476,567 parents, 428,327 students, and 62,115 teachers completed the NYC School Survey, which ranks among the largest of any kind ever conducted nationally.¹ Survey results provide insight into a school's learning environment through questions that collect information on perceptions of four broad reporting categories: (1) Academic Expectations; (2) Communication; (3) Engagement; and (4) Safety & Respect. School Survey results are also included in the calculation of each school's Progress Report grade (the exact contribution to the Progress Report is dependent on school type). These school Progress Report grades are used by the DOE to track a variety of factors related to schools' quality and progress over time.

The Research Alliance for New York City Schools examined DOE School Survey data from 2008-2010 to better understand the richness and complexities of the information elicited by the Survey from parents, students, and teachers. This report provides background information on the development of the NYC School Surveys during this time period and assesses the usefulness and appropriateness of measures derived from the survey and included in Progress Report grades. To do this, we first provide context about the School Survey's multiple purposes. Next, we outline the survey development process, and give information about similar large-scale survey efforts that informed NYC's survey measures and administration. We then present a series of statistical tests used to examine whether the School Survey is useful and appropriate for describing the school learning environment, particularly whether it contributes meaningful information to Progress Report grades, and whether it identifies components of the learning environment that schools can target for improvement. Finally, the report outlines steps for improving the School Survey as a measure of the school environment, while also maintaining continuity in items and remaining a stable measure for School Progress Reports.

Goals for the New York City DOE School Surveys, 2008-2010

Because the DOE has identified multiple purposes for the School Survey, and its use differs across groups of key stakeholders, the School Survey instrument is broad and

complex. Identifying the DOE's key goals for the School Survey is important for understanding its design and implementation. The first goal is to provide actionable information to schools to help them better understand their strengths and weaknesses and target areas for improvement. Second, the School Survey provides an opportunity for community engagement and feedback, by giving all parents, students in grades 6-12, and teachers the opportunity to participate. Third, the survey provides evidence that enables researchers to link malleable characteristics of the school learning environment to overall school effectiveness. Finally, the School Survey adds information about perceptions of the school learning environment to schools' Progress Reports (which are otherwise largely based on measures of academic achievement).

To achieve the DOE's first goal of providing useable, actionable information to school leaders about key stakeholders' perceptions of the learning environment, the DOE produces a report for each school with a summary of survey results. These reports provide detailed, item-level information to schools on a range of indicators that represent dimensions of the school learning environment. Respondent groups' aggregate reports of each item are presented using frequencies (e.g., the percent of respondents who selected "strongly agree", "agree", "disagree," or "strongly disagree") for each survey item, organized by respondent group (parents, students, teachers) and reporting category (Academic Expectations, Communication, Engagement, and Safety & Respect).

The act of survey administration itself addresses the second goal of providing an opportunity for community engagement and feedback with the survey participants. Schools work hard to connect with all parents, students, and teachers during the administration window. The DOE makes an effort to ensure respondents' confidentiality by using a third party survey vendor to distribute and collect surveys and process survey results. Different respondents' perspectives on survey items are understood as distinct and important, given their varied experience with the school learning environment. Taken together, the voices of the three reporter groups provide a comprehensive look at how each reporting category looks in a given school, and across the district. Community engagement is thus a product of participation in the survey itself as well as interest in the resulting school and citywide reports.

Nuanced reporting also serves the third goal of the survey—to enhance the quality of evidence linking malleable characteristics of the school learning environment and overall school effectiveness. Information from key stakeholders about their perceptions of the school environment is an important step toward understanding how characteristics of school climate—which are not necessarily measured by student achievement data—may link to students’ educational attainment, and the organizational health of the schools themselves. Applying the School Survey to the larger context of improving schools allows the DOE to build quality evidence linking characteristics of the learning environment to overall school effectiveness. This is likely a good first step toward identifying school characteristics that can be targeted for improvement.

The fourth goal of the DOE School Survey, and the subsequent focus of this report, is its inclusion in schools’ Progress Reports as the School Learning Environment score. School Progress Reports incorporate four School Survey reporting category scores (Academic Expectations, Communication, Engagement, and Safety & Respect). These scores, along with attendance, comprise the School Environment score, one of the measures incorporated in the calculation of each school’s overall Progress Report grade. Using the School Survey in this manner attaches considerable weight to the survey and may pose challenges for interpreting some results. For example, one criticism argues that because respondents know that survey results count for their school’s Progress Report, they may be inclined to answer more positively than they would if the results were not used for accountability. Such a situation puts the utility of the School Survey into question. If all scores are skewed positively, it may be difficult to identify key components of the school learning environment to target for improvement.

However, even given the possibility of response bias (which is an issue in all social science survey-based research), it is unlikely that accountability alone is driving survey results. First, there are three key stakeholders reporting on the survey, and each group is comprised of diverse individuals with differing perspectives on the school learning environment. Each type of reporter will likely provide unique information (which we used a statistical test to determine, as discussed later in this report). Although some accountability issues may bias these reports, it is unlikely that this would be the case for all three reporters, and that the extent and nature of the bias would be similar across the groups. Moreover, it is important to look at the variation within and across schools on the four reporting categories. Although it

appears that most responses are skewed positively, there is substantial variation in measures (also discussed later in this report). This finding suggests that, at least between reporter groups within schools, there is a sizeable group of individuals who differ in their perceptions of the school's learning environment. Such variation suggests that accountability is not the factor driving survey item responses across all respondents. The current report investigates this variation by using a series of statistical analyses to determine whether the survey results provide helpful information to schools and policymakers, over and above complications posed by reporter bias.

School Survey Development

In 2006, the DOE issued an RFP and ultimately contracted KPMG to begin the survey design process. The survey development team reviewed a number of extant surveys of school climate and environment, and collected feedback from a range of key stakeholders in the NYC public education system. Parents, students, teachers, parent coordinators, principals, and representatives from organizations like the United Federation of Teachers and the Citywide Council on High Schools contributed ideas about survey content and administration. Based on these initial comments, conversations, and working groups, survey developers identified four reporting categories that describe the school learning environment: (1) Academic Expectations; (2) Communication; (3) Engagement; and (4) Safety & Respect. These categories are the same four that exist today.

Next, the DOE conducted a research review that included education researchers' advice on refining and testing the surveys. During this process, multiple stakeholders, including researchers, community-based organizations, educators, parents, and internal DOE staff provided feedback on the types of questions that the survey should include, as well as the question phrasing. Because there was specific interest in using the survey results in the School Progress Reports, developers focused on collecting information about aspects of the learning environment that schools could control and could focus on improving after receiving their survey results. DOE then tested the surveys in the field with parents, students, and teachers, and further modified them based on feedback. To ensure parent participation, the DOE also sought out advice from community-based organizations about methods for survey administration before beginning data collection in 2007.

The parent, student, and teacher versions of the School Survey all underwent another major round of revisions, including changes to individual items, after the first year of administration. The survey stayed fairly stable in terms of overall content and individual questions from its inception through the 2010-2011 school year, when the DOE began implementing some of the Research Alliance's recommendations.

Broader Context for the School Survey

While the School Survey is unique in terms of its scale, there are a number of examples of survey administration efforts that help to provide background on the use and utility of similar survey-based research efforts. Like the School Survey, these other surveys have broad, multi-faceted goals, have been administered in and used for a range of purposes, and have helped to inform development of the School Survey. For example, the University of Chicago Consortium on Chicago School Research (CCSR) and the National Center for Education Statistics (NCES) have a long history of education-focused survey efforts that address multiple goals. These organizations' survey work provides context for one of the broader, underdeveloped goals for the School Survey—building a body of evidence about malleable school characteristics related to overall school effectiveness.

The NYC School Survey is most like the CCSR survey effort, which uses teacher and student perceptions to assess the school environment and produce school-level reports. Over the course of 20 years of work with teacher and student surveys, CCSR has developed a number of school environment constructs and used this research to present findings to schools. CCSR's biannual teacher survey is used to create school-level reports. By using survey results to develop confidential school reports, CCSR was able to not only build trust and rapport with schools but to develop the survey in ways that provided schools with information that was useful and appropriate for improvement. Survey results also provided the foundation for Bryk and colleagues' 2010 book, *Organizing Schools for Improvement*, which identified five essential supports for successful schools. Using this framework, CCSR was able to link characteristics of the learning environment to indicators of school quality, such as improving achievement trajectories. CCSR's work combines the practical goal of the survey—to provide actionable information to schools—with broader research implications. The NYC School Survey holds the same potential as the CCSR surveys; it is a source of information for individual schools and can be

leveraged for research purposes. However, the CCSR surveys take place solely in a research context and have not yet been used for accountability purposes.

Similarly, NCES has multiple surveys and programs that build publically available, nationally representative data sources that play an important role in education research. The NCES program studies cover a range of topics and cohorts, including the National Longitudinal Study of the High School Class of 1972 (NLS-72), High School and Beyond (HS&B), the National Education Longitudinal Study of 1988 (NELS:88), the Education Longitudinal Study of 2002 (ELS:2002), and the High School Longitudinal Study of 2009 (HSL:09).² These studies result in nationally representative data sources. Although these surveys cover a wide range of topics, samples, and school types, they do not provide information to or about individual schools. Researchers use NCES survey data to build a research base of parent, student, and teacher perspectives on topics such as student locus of control or parent expectations, similar topics to those on the NYC School Survey.

CHAPTER 2: PRELIMINARY ANALYSIS OF THE NYC DOE SCHOOL SURVEY

Like the CCSR and NCES surveys, the NYC School Survey is a rich source of data for research. Unlike those other survey efforts, however, less is known about the measures used for the School Survey. The ultimate goal of this report is to assess the usefulness and appropriateness of the measures derived from the NYC School Surveys from 2008-2010 for use in the Progress Reports. To do this, we provide an overview of key measurement properties and descriptive statistics used for the preliminary analysis. First, we explain how we calculated School Survey scores and how those differed from the the DOE's original calculations. Second, we provide response rates to show who participated in the School Survey from 2008-2010. Next, we examine the reliability and construct validity of the reporting categories, as defined by the DOE. Finally, we examine the extent to which the School Survey relates to other Progress Report scores and to student test scores in general.

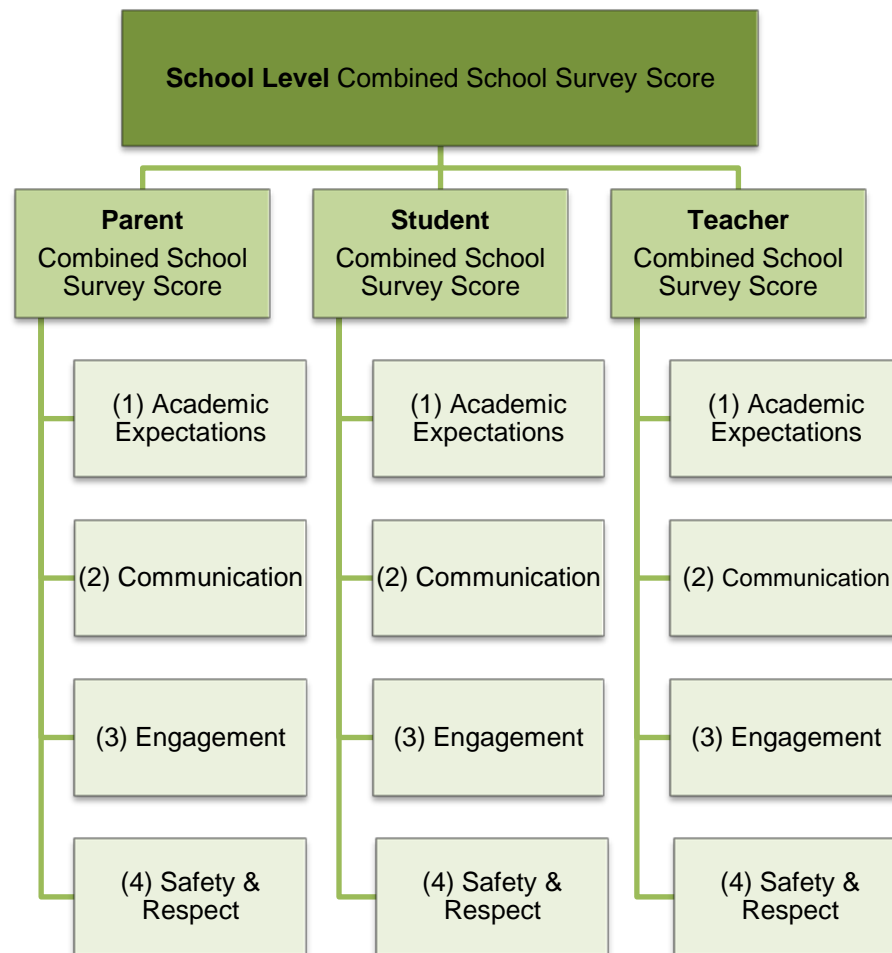
Results presented in this report focus on three major school types—elementary (ES), middle (MS), and high (HS) schools. Appendix Table A-1 shows the number of schools by type for 2008-2010. In 2010, 43 percent of all NYC public schools were classified as grade K-5/6, 17 percent as grade 6-8, and 25 percent as grade 9-12.³ Results for additional grade configurations (K-8, 10 percent NYC public schools; and K-12, 5 percent of NYC public schools) are included in Appendix A.⁴

How Can We Measure Components of the School Learning Environment? Constructing School Survey Scores.

As noted earlier, initial survey development efforts led to the creation of individual items on the School Survey covering four reporting categories: Academic Expectations, Communication, Engagement, and Safety & Respect. Each reporting category is scored separately by respondent group (parents, students, teachers) and includes a large number of individual items (ranging from 5 to 17 items, depending on respondent type) per category. By analyzing data from each respondent group separately, we can ascertain whether they have different perceptions of the learning environment and can provide unique information. Appendix B briefly describes how we calculated school-level scores for reporting categories and provides additional details, including tables listing survey items by respondent.

For each school, we constructed four reporting category scores for each respondent group, resulting in 12 scores per school (represented by the bottom four rows of boxes in Figure 1 below).⁵ We calculated these school-level scores by combining items within the four reporting categories at the individual level and then aggregating the category scores to the school level by taking the mean of the scores of all the individual reporters within the school. These calculations resulted in one school-level score for each respondent group in each reporting category (again, the green boxes). We further reduced the data by combining all four reporting categories into a single combined School Survey score for each respondent group, resulting in three scores per school (represented by the second row of boxes in Figure 1). Ultimately, we combined these scores when calculating the total School Environment Score (represented by the top box in Figure 1). For the majority of this report, however, we examine scores separately for each respondent group.

Figure 1: Constructing School Survey Scores



Our calculation of the School Environment Score differed from the DOE's original calculation in that we did not combine category scores across respondents. However, when validating our calculation against the DOE's, we found very strong correlations between the Research Alliance School Survey scores and those in the DOE's publicly available data files (e.g., average correlations between Research Alliance and DOE scores for parents, students, and teachers were $r = .96$, $r = .98$, and $r = .99$, respectively). This correlational evidence supports our hypothesis that the scores we calculated were equivalent to the DOE scores. Therefore, analyses linking these scores to outcomes (e.g., achievement, retention, grade completion) can likely be replicated with the reporting category scores calculated by the DOE. Appendix Tables A-2 to A-4 show correlations between comparable scores for 2008-2010.

How Many Parents, Students, and Teachers Participated in the Survey? Analyzing School Survey Response Rates.

Evaluating response rates is an important first step in analyzing the School Survey data because response rates indicate whether or not results are representative of the whole population of NYC parents, students, and teachers. Higher response rates reflect broader participation among key stakeholders. Lower response rates suggest that the survey results may not be representative of the entire group of intended participants.

A common approach to survey research is to sample from the total population, randomly selecting a subgroup of survey respondents from the total pool of respondents. However, the NYC DOE uses a census approach when administering the School Survey and asks *all* parents, students in grades 6-12, and teachers to complete the survey. By taking this approach, the DOE fulfills its goal of engaging the community and soliciting feedback from all constituents. NYC public schools employ over 80,000 teachers and serve more than 1.1 million students (and by extension, their parents). As such, a survey of this scale is second in size only to the US census. Although labor- and resource-intensive, the ambitious scale means that the survey represents a range of perspectives about the nation's largest school district.

Figure 2 below displays response rates by respondent group from 2008 to 2010. (Figure 2 includes response rates by respondent for all school types for 2008-2010). Overall, there was a positive trend in response rates over time, with the survey generating the highest response rates in 2010. However, response rates varied by respondent group and school type. Figure 3 on the next page shows response rates for each respondent group by school type for 2010. (Patterns were similar for 2008 and 2009). In 2010, parents had relatively lower participation than students or teachers. Lower response rates were especially pronounced for middle and high school parents (50 and 30 percent, respectively). One reason for this difference may be that schools have greater access to students and teachers who come to school every day, compared with parents who may not have direct daily interactions with schools. The lower parent response rates for the School Survey suggest that the respondent group of parents may not be representative of the entire parent population. However, it is important to consider that, historically, response rates for parent surveys in large school districts have been low (an estimated 30 percent for similar district-sponsored surveys). By comparison, the parent response rate in NYC is high.⁶

Figure 2: Parent, Student, and Teacher Response Rates, 2008-2010





Figure 3: Parent, Student, and Teacher Response Rates by School Type, 2010



Note: Elementary school students do not complete the school survey.

Response rates improved over time for all respondents. For students and teachers, high response rates provide confidence that respondents represented the broader population. Parent response rates, although improving over time, highlight an area for improvement.

Do Survey Items Reliably Reflect Each Reporting Category? Examining the Internal Consistency of School Survey Reporting Categories.

The DOE created the four reporting categories for the survey based on the premise that the individual items in each category measured the same, homogeneous construct. It is, however, important to test this assumption empirically. Cronbach's alpha (α) is a coefficient of reliability commonly used to assess how closely a set of items is related as a group. Both the strength of correlation between items and the number of items contribute to Cronbach's alpha. If all items in the reporting category measure the same characteristics, the correlations between items should be strong. Although the number of items does not necessarily relate to the strength of correlations, increasing the number of items can inflate Cronbach's alpha. Alphas between .70-.95 indicate that the items within a reporting category measure the

same construct and have good internal consistency. Alphas above .95 suggest redundancy among items and indicate that it is possible to delete individual items without losing any accuracy in measuring the larger reporting category.

First, we calculated a Cronbach's alpha to evaluate internal consistency for each of the four reporting categories by respondent type. As such, we calculated 12 total alphas—four for each respondent group. Then, we calculated a Cronbach's alpha to assess the internal consistency by respondent type (aggregated all the individual School Survey survey items for parents, students, and teachers respectively)—three alphas in all.⁷

In general, items within each reporting category successfully measured the same characteristics. When all items were combined across categories into one score, some items appeared to be redundant.

Most alphas that we calculated in these analyses were higher than the .70 threshold denoting high internal consistency, and exhibited little variation across years and school types (Table 1 below shows alphas for 2010 as an example, and Appendix Tables A-6 and A-7 show alphas for 2008 and 2009). All alphas for the parent and teacher reporting categories in 2010 were above the .70 threshold. However, the student-reported Communication category had lower than desired alphas (below .70) in all years. One possible explanation for this finding is that the six items comprising the student Communication category used three different rating scales, including three items that were originally rated on a five-point, as opposed to a four-point scale used to administer most of the remaining survey items. These inconsistencies in administration may account for the consistently low alphas in this category. Two additional categories had alphas below .70 for specific respondents in single years: student-reported Academic Expectations in 2008 and teacher-reported Engagement in 2009. Aside from these anomalies, alphas were sufficiently large to support grouping items by the reporting categories established by the DOE.

Table 1 on the next page shows that when all items were combined for each respondent group, alphas were between .93 and .98. These high alphas indicated that all items measured the same reporting category to some extent, and in fact, the majority of alphas for all items combined were equal to or above .95, indicating redundancy among items. In the next set of analyses we continue to examine the extent to which the four category scores identified separate reporting categories or measured the school environment as a whole.

Table 1: Cronbach's Alpha by School Type and Respondent, 2010

	ES	MS	HS	K-8	K-12	All
Parent						
Academic Expectations (5)	0.87	0.87	0.88	0.87	0.87	0.87
Communication (10)	0.89	0.89	0.90	0.89	0.90	0.90
Engagement (7)	0.84	0.84	0.85	0.84	0.85	0.85
Safety & Respect (11)	0.84	0.88	0.92	0.87	0.89	0.89
Combined (33)	0.96	0.96	0.97	0.96	0.96	0.96
Student						
Academic Expectations (9)	--	0.73	0.79	0.72	0.75	0.76
Communication (6)	--	0.61	0.66	0.61	0.62	0.64
Engagement (10)	--	0.83	0.85	0.82	0.83	0.84
Safety & Respect (20)	--	0.91	0.92	0.90	0.91	0.91
Combined (45)	--	0.94	0.95	0.93	0.94	0.94
Teacher						
Academic Expectations (15)	0.95	0.95	0.95	0.95	0.94	0.95
Communication (10)	0.89	0.89	0.89	0.88	0.89	0.89
Engagement (12)	0.93	0.92	0.92	0.92	0.91	0.92
Safety & Respect (17)	0.91	0.92	0.91	0.92	0.90	0.91
Combined (54)	0.97	0.98	0.97	0.98	0.97	0.97

Note: The number of items follows the reporting category for each respondent group. Alphas below the .70 threshold are shaded.

Do Category Scores Measure Four Different Reporting Categories or One School Environment Construct (the Overall School Learning Environment)? A Multi-Method Construct Validity Analysis.

When combining individual items to describe larger categories, it is important to determine whether the reporting categories themselves represent distinct concepts. For example, we want to know whether Safety & Respect as a construct is statistically distinct from Engagement, or if the two reporting categories are actually capturing a larger perception of the school learning environment as a whole. Such information will be helpful for using survey results to make changes in schools.

With this in mind, we used two methods to explore the associations between different reporting categories: bivariate correlations and confirmatory factor analysis. Results from these two analytic methods helped us to understand the underlying structure of the School Survey items.

- **Bivariate correlations**

It is important to examine correlations to determine the extent to which the reporting categories successfully measure different categories. For example, a high correlation indicates that a respondent's score on Academic Expectations is fairly similar to their score for Safety & Respect. We expected there to be moderate correlations (e.g., $r = .30-.69$) between reporting categories, which would suggest some association between categories but also some differences. We used the following stringent guidelines to evaluate strength of correlations: correlations .70 or higher are strong; .30-.69 are moderate; .29 or lower are weak.⁸

Appendix Tables A-4 to A-6 show that the majority of correlations between reporting categories within a respondent group were strong (above .70). In 2010, all correlations between category scores were strong. There were a few lower, but still moderate correlations in 2008 and 2009. For example, Safety & Respect scores had moderate (as opposed to strong) correlations with other categories across a variety of school types in 2008 and 2009. Overall, correlations were consistently strong across years and categories, meaning that category scores overlapped and were strongly related to one another. This finding suggests that the reporting categories were likely capturing similar information from respondents.

- **Confirmatory factor analysis**

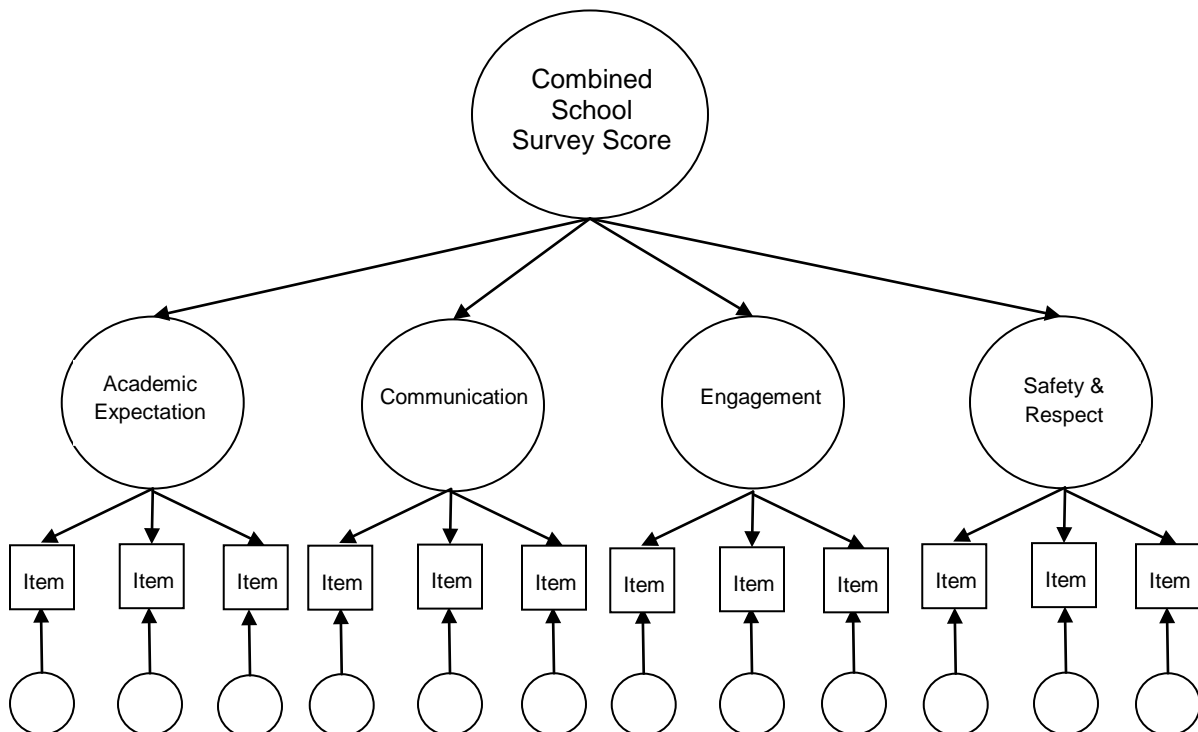
A bivariate correlation only represents the concurrent association between two variables. A number of additional analyses can be used to determine whether shared variance between individual items describes higher order reporting categories and a global score of the school learning environment. The DOE uses School Survey results to calculate reporting category scores that are provided to schools and used for accountability purposes. As such, we decided to use a second-order confirmatory factor analytic model to test whether individual items successfully map onto reporting categories, and then test whether those reporting categories accurately represent an overall school score.

After engaging key stakeholders and leading education researchers about how best to capture important information about a school's learning environment, the DOE hypothesized that four factors matter for school climate—Academic Expectations, Engagement, Communication, and Safety & Respect. Because one cannot directly measure these latent factors, the DOE worked with education researchers to

developed survey items they hypothesized were related to the factors based on theory and practical applications. For example, they hypothesized that the item “I need to work hard to get good grades at my school,” was related to Academic Expectations. If all of the survey items for Academic Expectations measure the factor well, confirmatory factor analysis (CFA) results would show that these items were strongly associated with one another and the individual items would demonstrate a high degree of shared variance in representing the reporting category. Thus, a CFA can be used to determine whether the data measure four distinct factors (that is, the four reporting categories) or conversely, if the current factor structure does not fit the data well.⁹

Figure 4 below shows the theoretical model we tested with specific relationships between items, reporting categories, and the combined School Survey score. We tested a second-order factor model where the first-order factors were the four reporting categories and these factors loaded onto the combined School Survey score factor. A key benefit of CFA is that the analysis produces model fit statistics that are useful for assessing how well a single theory-specified model fits the data. We used industry standard values to assess model fit for three independent models (e.g., students, teachers, parents).¹⁰

Figure 4: Theorized Factor Structure for School Survey Items



• **Confirmatory Factor Analysis Results**

We ran CFA models for parents, students, and teachers based on 2010 survey data (see Appendix Table A-11 for the resulting fit statistics for each model). None of the results met the recommended goodness-of-fit threshold. Poor model fit indicated that grouping the items into four factors and then one overall score was not an appropriate way to represent the data. By allowing for an alternative model fit, we were able to identify reasons why the original model did not fit the data. We subsequently changed the model to generate a better fit.

For all respondent groups, alternative models suggested that fit could be improved by allowing most items to load onto multiple factors. In other words, although items should theoretically load onto only one reporting category, we discovered that if we allowed items to load onto more than one category, the model fit would improve. When we did so, we found that model fit could be improved by allowing items to load directly onto the second-order combined School Survey factor, without assigning them to a single reporting category.

Both model fit statistics and modification indices suggested that the theorized model does not fit the data well and a different, more streamlined model may be more appropriate. Strong associations between category scores and a lack of four-factor model fit both support combining all four categories into a single score for each respondent group.

Do Parents, Students, and Teachers Have the Same Perceptions of the School Environment? Examining Correlations Between Respondent Groups.

We know that parents, students, and teachers all participate in the School Survey in large numbers each year. In order to examine whether these key stakeholders perceive the school environment in the same way, we calculated a single combined School Survey score from the four reporting categories, resulting in three summary scores per school—parent, student, and teacher combined School Survey scores. We then calculated bivariate correlations (r) between different respondent group scores and used the same guidelines from the previous section to interpret the strength of correlations (correlations .70 or higher are strong; .30-.69 are moderate; .29 or lower are weak).

Correlations between these combined School Survey scores showed some association between different respondent group perceptions. However, correlations were not as consistently strong as the correlations we calculated between reporting categories *within* respondent groups (that is, between student Academic Expectations and student Communication). Correlations between parents and students were strong in 2008 and moderate in more recent years. Correlations between teachers and each of the other respondent groups were low to moderate (Appendix Table A-12). These correlations do not provide strong evidence for combining all participants' scores into a single school-level score. Instead, the low correlations indicate that because parents, students, and teachers have unique interactions with schools, it is important to consider each of these perspectives separately. Doing so may provide a more complete, multidimensional picture of a school. We continue to examine the School Survey separately for each respondent group in the following sections.

Parent and student perceptions are more similar to one another than they are to teacher perceptions. Still, parents, students, and teachers each have a unique perspective of the school environment.

What Does the Average School Look Like? Examining Mean Scores by Reporting Categories and the Combined School Survey Score.

To calculate school scores for each reporting category, we aggregated scores from each respondent group for each reporting category to obtain 12 scores for each school (8 for elementary schools, where students did not participate in the survey). In addition, we combined the four reporting categories to calculate a combined School Survey score for each respondent group, which provides a concise summary we can use to analyze trends. Tables A-13 to A-15 display the means and standard deviations for each reporting category and combined School Survey score by respondent group and school type for 2008-2010. Higher scores indicate more positive perceptions of the reporting category (based on a 1 to 4 rating scale). School means provide a sense of average levels of each reporting category and combined School Survey scores by school type. The related standard deviation gives a sense of the variation of the school-level scores in relation to the overall mean.


In Figure 5 on the next page, we use the combined School Survey score to examine trends in reporting categories by respondent, school type, and year. First, we identified a restricted range, or tendency for respondents to respond positively to all

items in the 2010 data. Means for combined School Survey scores ranged from 2.86 (HS students) to 3.44 (ES parents). In addition, parents had more positive perceptions than students or teachers; the mean combined School Survey score for elementary school parents (3.44) was the highest of any respondent or school type. Students had the least positive perceptions of the three respondent groups; high school students in 2010 had the lowest mean for any respondent group (2.86).

On average, parents, students, and teachers all have generally positive perceptions of their schools. Parents had more positive perceptions than students or teachers. Respondents generally had more positive perceptions about elementary schools than middle or high schools.

Next, respondents had more positive perceptions about elementary schools (represented by darkest green in the figure), than middle schools (medium green), and high schools (light green). School type differences may be associated with respondent types because the largest proportion of parent participants (who had the most positive perceptions) report about elementary schools whereas student reporters (who had the least positive perceptions) were concentrated in middle and high schools. Overall, means for all respondent groups reflected positive perceptions of schools, regardless of year or school type.

Figure 5: Mean Combined School Survey Scores by School Type and Respondent, 2010



Note: Elementary school students do not complete the school survey.


Can the School Survey Be Used to Distinguish Between Schools? Examining Between- Versus Within-School Variation.

To help determine whether the School Survey is a reliable measure for differentiating between schools, we assessed the amount of variation between schools (as opposed to individuals within schools). For example, in theory, School Survey scores for different reporting categories will help us understand systematic differences between schools with poor Communication and schools with excellent Communication. This analysis is particularly relevant to the use of School Survey scores for Progress Reports, which use School Survey scores to rank and distinguish between schools.

The variance in School Survey scores can be divided into two parts: the variance attributed to systematic differences *between* schools and the variance attributed to individual respondents *within* schools. Again, in an ideal situation, School Survey scores help us understand systematic differences between schools with poor Communication and schools with excellent Communication. However, like many school-level characteristics such as test scores, we found that School Survey scores accounted for more variance within schools than between schools. In other words, there was more information distinguishing individuals within a school than information that helped differentiate one school from another.

By evaluating mean scores, we learned that all three respondent groups shared generally positive perceptions of the school environment. Figure 6 on the next page helps to illustrate the restricted range, or amount of variation in school-level scores, by showing means for schools in the 10th to 90th percentile of combined school survey scores. By plotting scores for these two extreme groups, the figure shows the limited variation between the most “extreme” schools. The mean scores based on parents and students showed the least variation; there was very little difference between the top and bottom groups. For teacher report scores, the range between the top and bottom schools was broader than for parents and students, but still was restricted to the positive range. This restricted range limits interpretation; although there may be statistically significant differences between a top school and a bottom school, the practical significance between 3.1 and 3.4 is limited.

Figure 6: Mean Combined School Survey Scores for 10th to 90th Percentile Schools by School Type and Respondent, 2010



Note: Elementary school students do not complete the School Survey.

On average, there was less variance in reporting category scores between schools for parents and students than there was for teachers (in 2010, 8 percent, 12 percent, and 24 percent, respectively). Combined School Survey scores showed more between-school variation for all respondent groups (Appendix Tables A-8 to A-10), but parent and student combined School Survey scores still showed less between-school variance than teacher scores. While teachers within a school may have different opinions from one another, collectively, teachers provide information that helped us distinguish between schools. Variation between schools based on parent and student scores ranged from 4 percent to 19 percent in 2010, depending on the reporting category and school type. These scores did not consistently distinguish between schools and therefore were not reliable scores for examining differences between schools.

Teachers' combined School Survey scores have potential to distinguish between school learning environments. Parent and student scores did not consistently distinguish between schools.

Does the School Survey Relate to Other School Characteristics? Examining Associations Between School Survey Scores and Indicators of School Quality.

Thus far, we have analyzed the School Survey from many angles—across years, categories, and respondents—using a variety of statistical analyses. It is also important to understand whether School Survey results are associated with other school characteristics, namely state test scores and other indicators of achievement. We used regressions to model concurrent relations between School Survey scores and school characteristics that are indicators of school quality, controlling for a number of non-malleable background characteristics of schools (for example, percent English Language Learners) that are theoretically related to School Survey scores and outcomes. Controlling for these structural characteristics helps us better understand the associations between School Survey scores and outcome measures, beyond background factors.

For each regression model, School Survey scores predicted test scores. In other words, controlling for school-level background characteristics, schools with higher School Survey scores were more likely to have higher English Language Arts (ELA) and math state test scores for elementary and middle schools. High schools also had a larger proportion of students who were “on track” to graduate (based on credits and Regents accumulated), relative to schools with lower School Survey scores.¹¹ For ease of interpretation we standardized the Survey Scores, calculating z scores with a mean score of 0 and a standard deviation of 1. Because the original scores were on different scales (e.g., some questions were on a five-point scale, others were on a four-point scale), standardizing the measures allows us to compare the magnitude of scores directly. Thus, the reported regression coefficients (beta) represent the change in test scores (or proportion of on-track students) for one standard deviation in School Survey scores.

Overall, there were significant positive associations between School Survey scores and test scores across respondent groups and years (Tables A-11 to A-14). However, even though the regression coefficients were statistically significant for most of the models, the practical relevance of these coefficients is unclear. For example, consider the association between the elementary school teachers’ School Survey scores in 2010 as a predictor of ELA test scores: If a school’s teacher School Survey score increased by one standard deviation—a relatively large change on the survey scale—the average ELA test score of this school would increase by .71 ELA

scale score points (controlling for school characteristics including previous test scores). For middle schools, the association between School Survey scores and test scores appeared to weaken from 2008 to 2010; in fact, teacher School Survey scores were significant predictors of test scores in 2008 and 2009, but not in 2010. Both limited practical relevance and changes over time in the association between School Survey scores and test scores pose problems for interpreting these correlations.

Similarly, associations between School Survey scores and on-track indicators for high schools were inconsistent and therefore difficult to interpret. Regression coefficients were close to and sometimes not significantly different from zero. Even so, because the outcome variable was the percentage of on-track students, it is possible to interpret this finding as practically meaningful because a change of one standard deviation in School Survey scores predicted a 1 percent change in the percent of on-track students. In this respect, School Survey scores from each respondent group predicted 2 to 3 percentage point changes in percentage of on-track students in 2010. Again, year-to-year changes in the association between School Survey scores and percent on-track students do pose problems for the interpretation of these coefficients. However, given that this finding translates into a significant number of students within a school being on track to graduate from high school, the change predicted by the School Survey score is practically meaningful.

Large changes in School Survey scores were associated with relatively small changes in test scores, after controlling for background school characteristics. Inconsistent changes over time add to the difficulty of summarizing how the School Survey relates to other school characteristics.

Summary

The preceding chapter examined the School Survey from many angles to understand whether it is useful and appropriate a component of the school Progress Report. We replicated the DOE's four reporting categories, recalculating them on a consistent four-point scale for each reporter type (parents, students, and teachers) and evaluated their ability to empirically measure different aspects of a school's environment. Our interpretation of the results is that the School Survey captured an

overall construct of learning environment rather than four distinct reporting categories. Reliability statistics indicated that reporting category scores are reliable, but when we combined the categories into an overall school score, we discovered that some survey items may be redundant. Further, scores reflected a restricted range; the School Survey assessed a range of positive perceptions but did not do as good a job measuring negative perceptions about the school environment.

Although the reporting categories were empirically similar to one another (e.g., there were strong correlations between reporting categories), the three respondent groups were distinct from one another. The School Survey tapped into parents', students', and teachers' unique perspectives about school environment. A single combined School Survey score from each respondent group at each school provided varied information about schools' learning environments. Furthermore, teachers stood out as the reporters whose ratings provided the information that best distinguished between schools. These results suggest that School Survey scores can be combined in ways that take into account the unique advantages of each respondent group. For example, it makes sense to weight teacher scores more heavily when calculating School Survey scores that distinguish between schools.

Finally, School Survey scores were significantly associated with other policy-relevant school characteristics, although these associations were inconsistent over time. In addition, relatively large changes in School Survey scores were associated with relatively small changes in elementary and middle school test scores. For high school outcomes, changes in School Survey scores were associated with small but meaningful changes in the percent of on-track students.

CHAPTER 3: IMPROVING THE SCHOOL SURVEY

Findings from the previous chapter build a strong case for revising the School Survey. In this chapter, we explore ways to improve the School Survey as a measure of the school learning environment. We focus on the possibility of using fewer items to calculate school environment scores that are comparable to the ones currently used for school Progress Reports. Because we recognize the advantages of stable measures for accountability reporting over time, we have balanced suggestions for revisions with the need to preserve the comparability of the School Survey's contribution to Progress Reports.

Can We Measure the School Environment More Efficiently with a Subset of Survey Items?

Increasing efficiency (measuring the same construct using fewer items) is an approach for improving the School Survey as a measure of the school environment. Using fewer items to produce equivalent scores would free up space for adding measures that could capture other aspects of the school environment. Because we found that many of the items on the School Survey are collecting duplicative information, it is possible to significantly reduce the number of individual survey questions without significantly changing the overall School Survey scores used on the Progress Report. At the same time, the overall length of the survey would not increase.

We ran two sets of analyses to explore whether a subset of School Survey items could measure the school learning environment as well as the full set of items. In the first set of analyses, we gauged the possibility of reducing the number of survey items while preserving schools' combined School Survey score for each respondent group. It is intuitive that deleting just one of many items only minimally influences a school's score. However, it is not immediately apparent how much a school's score would change if a larger number of items were excluded. Therefore, we conducted an analysis that simulated how much scores changed as a function of including progressively fewer survey items in that score. This simulation, called a Monte-Carlo experiment, randomly selected items for deletion.


We hypothesized that an intentional item selection process using criteria aimed at maintaining score comparability would further increase stability in mean scores. Thus, in the second set of analyses, we purposefully selected specific items to create

several subsets and then examined the differences between scores based on these item subsets and scores based on the full set of items. We examined mean scores at different levels of aggregation (for example, combining across respondents or categories) and correlations between scores based on the full set and subsets of items.

Monte-Carlo Experiment

Monte-Carlo analyses helped to determine the ideal number of items to include in a reporting category, ultimately resulting in school-level mean scores comparable to scores based on the original full set of items. For each iteration in the experiment, we recorded the difference between the original score and the reduced-item score. The random deletion process and score calculation was repeated 700 times for each number of deleted items (i.e., 700 times deleting one item, 700 times deleting two items, etc.). The mean differences between the full set and subset scores for all 700 iterations are plotted in Figure 7 below.

Figure 7: Mean Difference from Original School Survey Scores Based on Monte-Carlo Experiment



The mean differences between scores based on the full set and subsets of items are a good measure to judge how many items we could delete from the School Survey while still preserving comparability with the original score. However, the threshold for how large a difference is acceptable depends on the context and the overall variability of the measure, or the range in scores of the measure. Our Monte-Carlo results (Figure 7) clearly indicated that deleting up to 50 percent of items changed the student School Survey score by an average of .05 points (on a four-point scale). Parent and teacher scores based on the 50 percent item subset changed even less than the 50 percent subset for students. The Monte-Carlo experiments also indicated that deleting 25 percent of the items changed scores by less than .03 points across respondent groups. Based on these results, we concluded that we can calculate scores that were comparable to those based on the full set of items by using subsets that retained 50-75 percent of items (for each respondent group).

Item Subsets

The Monte-Carlo experiment was based on random item deletions. We hypothesized that intentional item deletion would result in a subset of items that would perform even better than Monte-Carlo estimates predicted. For the following set of analyses, we developed criteria for excluding items that (a) preserved school-level scores and (b) preserved the associations between the four reporting categories.¹² The Monte-Carlo experiments helped us estimate an optimal proportion of items to delete while maintaining score comparability. We used these criteria to create three subsets of items—retaining 75 percent of items, 50 percent of items, or a “balanced” subset retaining four items per reporting category (see selection process below)—for each respondent group. As in the Monte-Carlo experiment, we compared the absolute difference between scores based on item subsets and the full set to judge the level of consistency between survey scores. The absolute mean differences provided insight into how much scores would shift if fewer items were included. We also correlated scores based on item subsets and the full set of items to evaluate whether relative rank of schools was comparable.

Selecting Items for Subsets

To ensure each reporting category was still adequately represented in item subsets, we selected items for deletion with the goal of equalizing the number of remaining items within each category. Although the number of items in each category varied greatly in the original survey (for example, 5 Academic Expectations and 11 Safety

& Respect items for parents), each of the item subsets moved a step closer toward a balanced number of items within each category.

To identify specific items for deletion, we ranked all items in each reporting category based on two criteria: (a) how much the school level mean of a reporting category would change if the item was deleted and (b) how much this change varied across schools. We deleted the items in each category that would have the least impact on the original scores based on these two criteria. Figure 8 on the next page shows the number of items by reporting category in the original survey and in each of the subsets. Tables A-14 to A-16 show which specific items were included in each subset for parents, students, and teachers, respectively.

We created one subset that retained 75 percent of all items (Subset 1), one that retained 50 percent of all items (Subset 2), and a third “balanced” subset (Subset 3).

In Subset 3, we retained an equal number of items in each reporting category—four items for each respondent group in each reporting category. For parents, this subset differs from Subset 2 by only one item, but for the teacher survey, which initially had a larger item pool, the balanced subset only retains 30 percent of original items. For all categories and reporting groups combined, Subset 3 contained 48 items (36 percent retained from the original item pool).

Comparing Scores Based on Full Set and Subsets of School Survey Items

Maintaining score comparability while reducing the number of items is particularly important because School Survey scores combined across categories and respondents have a high stakes application as part of the DOE’s annual Progress Reports. Thus, our discussion of item subset score comparability focuses on the single school-level score that is equivalent to the School Survey contribution to the Combined School Score.

The DOE uses the School Survey to calculate the School Environment score on the Progress Report grade, which thus incorporates parent, student, and teacher scores for each of the four reporting categories. Each of the four category scores are

Figure 8: Number of School Survey Items by Reporting Category and Respondent for Full Set and Reduced-Item Subsets

	All survey items	Subset 1: 75% retained	Subset 2: 50% retained	Subset 3: "balanced"
Academic Expectations	Parents 5	5	4	4
	Students 9	9	6	4
	Teachers 15	10	7	4
Communication	Parents 10	7	4	4
	Students 6	6	5	4
	Teachers 10	10	6	4
Engagement	Parents 7	6	4	4
	Students 10	9	6	4
	Teachers 12	10	7	4
Safety & Respect	Parents 11	7	5	4
	Students 20	10	6	4
	Teachers 17	11	7	4
Total items	Parents 33	25	17	16
	Students 45	34	23	16
	Teachers 54	41	27	16

weighted equally for the Combined School Score, and therefore, for each subset, we averaged the four categories to form a single school-level score that incorporates all respondent groups and all categories. For additional information about comparability, we also disaggregated and examined scores by category and respondent group.


The subset scores that combined survey results across categories and respondents groups included more items than scores calculated separately by category or respondent group. This suggests that the subset scores that combined categories and respondent groups should be more stable and less sensitive when a greater proportion of items were deleted. Appendix Table A-25, which presents the absolute differences¹³ between scores based on the full set and subsets of items by school type, shows exactly this pattern. The shaded rows show differences for scores combined across reporting categories but calculated separately for each

respondent group (three scores per school). The shaded column shows differences for scores collapsed across respondents but calculated separately for each category (four scores).

Across school types, the absolute differences for the single school-level score were very small for the first two subsets (.01 or .02) and reached .03 for the balanced subset, still a reasonably small mean difference on a four-point scale (see Figure 9 below). For Subset 3, the school-level score based on 48 items maintained comparability to school-level scores based on the full set of items, with a mean absolute difference of .03 across all schools (.02, .03, .05 for elementary, middle, and high schools, respectively, as presented in shaded and bolded cells in Appendix Table A-25). The mean absolute differences between scores based on the full set of items and Subsets 1 and 2, were even smaller—.01 and .02, respectively.

When we disaggregated the single score by respondent group to create three school-level scores, the differences between the full set and Subsets 1 and 2 were similar to single score differences (all below .04; see Figure 9 below). For Subset 3, teacher scores differed from full set scores by .07. Still, a mean difference of .07 on a four-point scale was relatively minor considering the number of items was reduced from 54 to 16.

Figure 9: Absolute Differences Between Full Set and Subsets of Items: Respondent-Group Scores, Combined Across Reporting Categories




Note: The line at the 0.00 mark on the y-axis represents the mean scores for the full set of items.

When we further disaggregated the single score by response category (while still combined across respondent groups) to create four scores for each school, the differences between the full set and Subset 1 were all below .03, (see Figure 10 below). However, differences for Subsets 2 and 3 were slightly greater. Communication scores differed from full set scores by .07 for subset 2 and .14 for Subset 3. Reporting category scores based on Subset 3 appeared less stable than scores by respondents detailed in the previous section (Figure 9).


Finally, we disaggregated scores to create 12 school-level scores (four category scores for each of the three respondent groups) and compared scores based on the full set and subsets. With even fewer items per category after respondents were evaluated separately, we expected larger differences between subset and full set scores. Figure 11 on the next page shows absolute differences for one category, Communications, as an example. Figure 11 illustrates the most extreme differences found in the Communication category: absolute differences between scores based on subset 3 and the full set of items were as high as .31 for teachers. In general, teacher category scores based on Subsets 2 and 3 were less stable than scores for other respondent groups.

Figure 10: Absolute Differences Between Full Sets and Subsets of Items: Reporting Category Scores, Combined Across Respondent Groups



Note: The line at the 0.00 mark on the y-axis represents the mean scores for the full set of items.

Figure 11: Absolute Differences Between Full Set and Subsets of Items: Reporting Category Scores Separated by Respondent: Communication Category



Note: The line at the 0.00 mark on the y-axis represents the mean scores for the full set of items.

For the majority of comparisons described in this section, scores remained relatively stable. In fact, the single school-level score most closely related to the School Environment score and remained comparable to scores based on the full set of items for all subsets.

We also wanted to establish that the relative ranks of each school based on mean scores did not change substantially. For example, a school that had a high mean score relative to other schools based on the full set of items should also have a comparably high score based on a subset of items. We used correlation analysis to compare the consistency of relative rank for scores based on full sets and subsets of items. In this case, strong correlations between two scores indicated that scores based on subsets of items consistently reproduced the relative rank of scores based on the full set.

Appendix Table A-26 illustrates the high level of stability of the relative rank of scores based on the full set and subsets of items. For all school types and respondent groups, correlations were near perfect. In general, correlation coefficients can range from -1 to 1, where -1 and 1 indicate a perfect relationship and 0 no relationship at

all. Correlations between scores based on the full set of items and subsets were strong. Only two correlations were under .91, but the correlations were still strong: for Subset 3, .84 for teacher Safety & Respect and .88 for teacher Communication.

Summary

These findings suggest that it is possible to maintain comparability of the survey scores for all school types and respondent groups using fewer items. Again, we found that the single school-level score, which combined all respondents and categories, could be reproduced with as few as 36 percent of the total items currently used (Subset 3). We found that the four reporting categories could be reproduced with as few as 50 percent of items, albeit with slightly less stability than scores based on the full set of items. The most conservative subset of items, scores based on the 75 percent subset, showed very high comparability with scores based on the full set of items for all levels of disaggregation, including for all 12 reporting category scores by respondent group.

CHAPTER 4: SUMMARY

Since 2007, the DOE has conducted annual surveys of all parents, students from grades 6-12, and teachers. The surveys represent a unique opportunity to gain perspective on a wide range of issues from the district's primary stakeholders. The DOE incorporates these perspectives into the annual Progress Reports that assess the progress, performance, and learning environment of each school in the district, giving the survey high stakes.

The Research Alliance for New York City Schools is conducting an ongoing assessment of the reliability and utility of the survey-based measures that are incorporated into the school Progress Reports. Our goal is to inform discussions of potential improvements and extensions of this unique source of information about the City's schools. Following is a summary of the key findings that have arisen thus far.

Response Rates

Robust response rates for students and teachers demonstrate widespread participation, a key strength of the NYC School Survey. Response rates among students and teachers increased steadily over time and reached 78 percent and 83 percent, respectively, in 2010. These high response rates offer confidence that survey results reflect the opinions of the broader population. Parent response rates did not approach the same levels as student and teacher response rates, and thus the representativeness of parent survey results is more in question. However, it is important to consider that, historically, response rates for parent surveys in large school districts have been low (an estimated 30 percent for similar district-sponsored surveys¹⁴). By comparison, the parent response rate in NYC is high. The district has made it a priority to increase parent response rates, which have risen over time and surpassed 50 percent since 2010. These positive trends in parent response rates are encouraging.

Measures of the School Learning Environment

The existing School Survey items provide statistically reliable indicators of the four reporting categories currently incorporated into the annual school Progress Reports (Academic Expectations, Communication, Engagement, and Safety & Respect). In

other words, responses to the survey items used to create each of these measures were highly correlated with one another.

The four reporting category measures were statistically indistinguishable from each other. For the purposes of the school Progress Reports, this means that the survey items used to create the individual measures can be combined into a single measure of school environment without significantly diminishing the current breadth of information about schools. This single measure can be constructed reliably using about half of the items currently being used to construct the four individual measures.

Individually and combined, the 2008-2010 versions of the four reporting category scores were not strong measures for distinguishing between schools. In other words, the survey provides more information about differences between individuals within a school, and less information about how that school differs from other schools. In addition, School Survey scores were not powerful predictors of other performance indicators. As a result of these limitations, there is significant room for improving the School Survey's ability to identify relative strengths and weaknesses of schools, predictive of school level achievement and completion.

The measures derived from teacher, student and parent surveys provide different assessments of the school's learning environment characteristics. These differences are masked by the scores included in the annual school Progress Reports, which combine measures across the three respondent groups for each of the four categories.

Measures derived from the teacher survey offer the greatest potential to identify systematic differences between schools' learning environments. Although there were large differences in perceptions of the learning environment among teachers in the same school, teacher responses as a group provided relatively strong information about differences in learning environments across schools. Current measures derived from student and parent surveys, however, were weaker instruments for identifying distinctions between schools.

Recommendations

The universal administration of the New York City School Survey and its corresponding high response rates present a unique opportunity to accomplish the

measurement and accountability goals that DOE and other school systems increasingly see as priorities. These goals include giving parents, students, and teachers a voice in assessing the quality of their schools; providing formative information that can be incorporated into school improvement efforts; and constructing climate measures for which schools may be held accountable.

Findings from our analyses suggest that much of this potential remains untapped. The following are recommendations intended to increase the utility of the school surveys, while preserving the continuity of measures that are currently in use:

- Eliminate redundant items that are used to create the learning environment measures while preserving the reliability of those measures.
- Reduce the time required to complete the school surveys and improve the reliability of measures by creating more consistent question formats and response categories.
- Choose different and separate measures for the parent, student, and teacher surveys to capitalize on their distinctive perspectives on their schools. Limit or eliminate measures that combine responses from parent, student, and teacher surveys. For example, although teachers may be the best reporters of Academic Expectations, parents and students may have unique and important views on Safety & Respect.
- Incorporate new school environment measures that are more likely to distinguish between schools and are associated with other school performance indicators. For example, ask teachers to report about their principal's instructional leadership.
- Incorporate measures that more effectively gauge parent satisfaction and engagement with their child's school. For example, items like "I would recommend this school to other parents" tap into parents' overall satisfaction.

By building on this formative research and further refining the School Survey to address the issues raised in this report, the DOE will be better able to understand how schools differ from one another, and identify the features of the learning environment that should be targeted for improvement. It is also important to note that the development of these instruments is ongoing and future measurement will continue to inform these efforts.

Notes

¹ Lopatin (2012).

² Additional information on national education surveys through NCES is located at <http://nces.ed.gov/surveys>.

³ The relative percentage of each school type is consistent from 2008-2010 within one percent. About 50 of the Grade K-5/6 schools had 10 or more student responses, but we excluded all student responses for this school type to preserve parity among elementary schools.

⁴ Note that only students in grade 6 and above participate in the School Survey, and therefore, the student scores for other K-8 and K-12 schools only include student perspectives for students in grade 6 or higher.

⁵ There were 12 scores for middle and high schools (four for each of the three respondent groups), but only eight in elementary schools where students did not participate in the survey.

⁶ We drew several comparisons to place NYC School Survey parent response rates in context. Austin (2011) estimated that typical parent response rates for district-sponsored paper surveys is 30 percent and noted that NYC made a heavy investment in public relations outreach to get parent response rates to 45 percent in the 2009 School Survey administration. Other examples suggest even lower rates: A Cincinnati Public Schools news release reported an 8 percent parent response rate in 2012. And, based on a 2008 report about a Los Angeles Unified School District parent survey, we estimated that 22 percent of the 15,000 parents who received a survey responded (the response rate was not directly reported in the paper, and we based on estimate on Tables 1 and 2, which reported the number of responses to two specific survey items.)

⁷ By combining all items the alpha may be inflated as the number of items increases.

⁸ Cohen (1988) provided benchmarks for operational definitions of weak, moderate,

and strong correlations (.10, .30, and .50, respectively). Cohen notes that these benchmarks are not universal, and different benchmarks may be appropriate for different contexts and expectations for the strength of association between variables. For example, a .70 correlation means that .49 of the variance in one parent category score could be explained by another parent category score.

⁹ In 2007, the DOE commissioned CFA and adjusted the four reporting categories based on a combination of these analyses and theoretical considerations. (Lauren Sypek, personal communication, 2013).

¹⁰ Schreiber and colleagues (2006) recommended the following fit indices to assess a single model and Hu and Bentler (1999) suggested the associated cutoffs for judging good fit: TLI > .95, CFI > .95, and RMSEA < .06.

¹¹ Background characteristics include ethnic/racial composition, proportion of free and reduced price lunch students, proportion of ELL students, and previous year's test scores.

¹² By using this approach, we maximized comparability but potentially deleted items with more important content or better measurement properties. For example, the survey currently has limited ability to differentiate between schools, and it is possible that this approach resulted in a comparable measure that has even less ability to differentiate between schools.

¹³ "Differences" refers to the absolute difference between scores based on the full set of items and scores based on a subset of items. When we use differences in conjunction with specific subsets, we are referring to differences between subset scores and scores based on the full set of items.

References

- Austin, G. (2011).** "Using a Parent Survey to Improve Parent Involvement and School Climate." National Conference for the Office of Safe and Drug-Free Schools, Washington, DC.
- Cohen, J. (1988).** *Statistical Power Analysis for the Behavioral Sciences*, 2nd Edition. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.
- Fernández, A. & Gilstrap, S. (2008).** *The Parent Perspective: A Survey of Belief Systems and Involvement Behavior in the Los Angeles Unified School District.* Retrieved 2/15/13 from http://notebook.lausd.net/pls/ptl/docs/page/ca_lausd/fldr_organizations/fldr_plcy_res_dev/par_division_main/research_unit/publications/reports/parent%20survey%20report-r%26p.pdf.
- Hu, L. & Bentler, P.M. (1999).** "Cutoff Criteria for Fit Indexes in Covariance Structure Analysis: Conventional Criteria Versus new Alternatives," *Structural Equation Modeling*, 6(1), 1-55.
- Lopatin, A. (2012).** Presentation at the biannual meeting of the Society for Research on Educational Effectiveness, Washington, DC.
- Rockoff, J. & Speroni, C. (2008).** "Reliability, Consistency, and Validity of the NYC DOE Environmental Surveys: A Preliminary Analysis." Working Paper. New York: Columbia University.
- Schreiber, J.B., Stage, F.K., King, J., Nora, A., & Barlow, E.A. (2006).** "Reporting Structural Equation Modeling and Confirmatory Factor Analysis Results: a Review." *Journal of Educational Research*, 99(6), 323-337.
- Walsh, J. (2012).** "CPS Parent, Student Satisfaction Levels High." Retrieved 2/15/13 from http://www.cps-k12.org/sites/www.cps-k12.org/files/newsreleases/Sept12ParentTeacher_0.pdf.

The **Research Alliance** for
New York City Schools

285 Mercer Street, 3rd Floor | New York, New York 10003-9502
212 992 7697 | 212 995 4910 fax
research.alliance@nyu.edu | www.steinhardt.nyu.edu/research_alliance

The Research Alliance for
New York City Schools conducts
rigorous studies on topics that
matter to the city's public schools.
We strive to advance equity and
excellence in education by
providing non-partisan evidence
about policies and practices that
promote students' development
and academic success.