

Stabilizing subgroup proficiency results to improve the identification of low-performing schools

A Publication of the National Center for Education Evaluation and Regional Assistance at IES



U.S. Department of Education

Miguel Cardona

Secretary

Institute of Education Sciences

Mark Schneider

Director

National Center for Education Evaluation and Regional Assistance

Matthew Soldner

Commissioner

Liz Eisner

Associate Commissioner

Heidi Gansen

Project Officer

Chris Boccanfuso

REL Branch Chief

The Institute of Education Sciences (IES) is the independent, nonpartisan statistics, research, and evaluation arm of the U.S. Department of Education. The IES mission is to provide scientific evidence on which to ground education practice and policy and to share this information in formats that are useful and accessible to educators, parents, policymakers, researchers, and the public.

We strive to make our products available in a variety of formats and in language that is appropriate for a variety of audiences. You, as our customer, are the best judge of our success in communicating information effectively. If you have any comments or suggestions about this or any other IES product or report, we would like to hear from you. Please direct your comments to ncee.feedback@ed.gov.

This report was prepared for the Institute of Education Sciences (IES) under Contract 91990022C0012 by Mathematica, Inc. The content of the publication does not necessarily reflect the views or policies of IES or the U.S. Department of Education nor does mention of trade names, commercial products, or organizations imply endorsement by the U.S. Government.

This is a provisional report that is being released by IES after expert peer review, but before being fully copy edited and formatted by the REL program editors. IES allows release of provisional reports when a partner in the education community needs access to the report's findings to inform a time-sensitive policy or programmatic decision. When the fully edited and formatted version of the study report is released by IES, IES will release a newsflash announcing that final version.

December 2022

This report is in the public domain. Although permission to reprint this publication is not necessary, it should be cited as:

Forrow, L., Starling, J., and Gill, B. (2022). *Stabilizing subgroup proficiency results to improve the identification of low-performing schools* (REL 2023-001). U.S. Department of Education, Institute of Education Sciences, National Center for Education Evaluation and Regional Assistance.

<https://ies.ed.gov/ncee/rel/Products/Publication/106926>

This report is available on the Institute of Education Sciences website at <https://ies.ed.gov/ncee/rel/>

Stabilizing subgroup proficiency results to improve the identification of low-performing schools

Lauren Forrow, Jennifer Starling, and Brian Gill

December 2022

The Every Student Succeeds Act (ESSA) requires states to identify schools with low-performing student subgroups for Targeted Support and Improvement (TSI) or Additional Targeted Support and Improvement (ATSI). Random differences between students' true abilities and their test scores, also called measurement error, reduce the statistical reliability of the performance measures used to identify schools for these categorizations. Measurement error introduces a risk that the identified schools are unlucky rather than truly low performing. Using data provided by the Pennsylvania Department of Education, the study team used Bayesian hierarchical modeling to improve the reliability of subgroup proficiency measures and demonstrate the approach's efficacy.

Why this study?

ESSA requires states to identify their lowest-performing schools for Comprehensive Support and Improvement (CSI) and to identify schools with low-performing student subgroups for TSI or ATSI. Identifying the schools that most need support hinges on accountability data that reliably measure school performance, as defined by the state. However, random differences between students' true abilities and their test scores—formally called measurement error—can obscure school performance, especially in small schools or subgroups, where random factors that affect a small number of students can have an outsized impact on the school's or subgroup's average score. As a result, accountability data for small schools and subgroups may not reliably reflect the performance of those schools and subgroups—small schools and subgroups could be identified for additional support primarily because of bad luck. States typically seek to reduce the risk of these errors by setting a minimum number of students that a subgroup must exceed to be included in accountability measures. However, a minimum subgroup size requirement may not fully solve the problem because data for the smallest included subgroups are still likely to be less reliable than data for larger subgroups. At the same time, such minimum subgroup size requirements come at the cost of barring some subgroups from contributing to a school's performance measures.

In this study, the Mid-Atlantic Regional Education Laboratory (REL Mid-Atlantic) investigates the potential of Bayesian hierarchical modeling to increase the robustness of the Pennsylvania Department of Education's (PDE's) accountability data to measurement error, thereby improving its accuracy. Specifically, the study team implemented a Bayesian modeling approach to stabilize subgroup proficiency rates used to identify Pennsylvania schools for TSI and ATSI. This modeling approach is referred to as Bayesian stabilization, or stabilization for short, throughout the rest of this report.

Proficiency rates are just one of six accountability indicators included in Pennsylvania's accountability system.¹ Although academic proficiency by no means provides a comprehensive picture of school performance, it is well suited to this pilot study because historical data are readily available and because there is a clear mechanism through which measurement error could affect a school's observed proficiency rate. Acknowledging that academic proficiency is just one dimension of school performance and could reflect factors outside the school's control, such as parental involvement or motivation, the REL Mid-Atlantic team considers this study a proof of

¹ The six indicators are academic proficiency, academic growth, progress toward fluency for English learner students, career readiness, regular attendance, and graduation rates.

concept that could inform Pennsylvania’s or other states’ decisions of to apply stabilization to proficiency or other accountability indicators in the future. For example, Pennsylvania or other states could choose to use accountability results based on stabilized indicators as a “safe harbor” that can move schools out of, but not into, ATSI.

The study goal was to determine whether Bayesian stabilization improves the reliability of the academic proficiency rates used in PDE’s accountability system, focusing on TSI and ATSI identification. TSI and ATSI walk a tightrope between reliability and inclusivity. On the one hand, performance data for small subgroups may largely reflect measurement error and should therefore not be used for accountability, lest schools be identified for additional support based on bad luck. On the other hand, excluding small subgroups from accountability for fear of assessing them based on unreliable data may bar them from receiving the support they need. A Bayesian stabilization approach could alleviate this tension by improving the precision and plausibility of proficiency rates for small subgroups, perhaps even for subgroups that are smaller than current minimum sample sizes used for accountability determinations.

If found to be effective, this approach could both improve the reliability of TSI and ATSI designations and expand the set of schools and subgroups included in the accountability system, allowing PDE to target the schools and students that most need additional support. Specifically, informed by the results of this study, the REL plans to collaborate with PDE to incorporate stabilization as a safe harbor alternative in its 2022 accountability calculations, addressing underlying measurement error in accountability indicators as well as the implications of accountability data collection disruptions in 2020 and 2021 due to the COVID-19 pandemic. The study’s findings could be relevant to states across the country, all of which face the same need to identify schools for TSI and ATSI, and could address the same tension between an inclusive approach to subgroups and measurement error when small numbers of students are involved.

Research questions

The study explored the usefulness of stabilization for school accountability through the following research questions:

1. How much do stabilized proficiency rates differ from unstabilized proficiency rates? Do these differences vary depending on the number of students in the subgroup?
2. How much does the Bayesian stabilization approach increase the reliability (long-run stability) of proficiency rates? Are these improvements sufficient for small subgroups to reduce the minimum threshold for a subgroup size from 20 students to 10?
3. How does incorporating stabilized proficiency rates into the ATSI determination process change the set of schools identified as eligible, as compared to a determination process implemented without stabilization?

Research question 1 gauged the impact of stabilization on proficiency rates, overall and by the number of tested students. This analysis determined whether stabilization made a difference to schools’ estimated proficiency rates, as context for further analysis. At the same time, it served as a check that the relationships between unstabilized and stabilized proficiency rates followed the pattern to be expected from theory and the literature, namely that stabilization would have a larger effect on smaller subgroups.

Research question 2 investigated how much stabilization increased statistical reliability and specifically how much stabilization improved reliability for small subgroups that are currently excluded from PDE’s accountability system because they do not meet the minimum sample size requirement, which Pennsylvania currently sets at 20 students. The minimum sample size requirement is intended both to preserve student privacy and to guarantee a minimum level of statistical reliability so that the accountability system focuses on schools and subgroups where enough information is available to assess performance. If stabilization achieves a comparable level of statistical

reliability for subgroups with 10–19 students as achieved without stabilization for subgroups with 20 or more students, the benefits of including small schools and subgroups in accountability may outweigh the statistical reasons to exclude them. Specifically, the current minimum subgroup size requirements eliminate small schools and subgroups from consideration for additional support that they may need. Including smaller schools and subgroups in accountability would make them eligible to receive this support.

Research question 3 explored how the effects of stabilizing proficiency rates affected which schools were identified for ATSI. This analysis assessed the possible impact of stabilization on accountability decisions, not its effect on the reliability of accountability data. Although it is not possible to assess the accuracy of accountability decisions, if stabilization improves the reliability of accountability measures, it is also likely to increase the accuracy of ATSI and TSI designations.

Box 1. Data sources, sample, methods, and limitations

Data and sample. The Pennsylvania Department of Education (PDE) provided school-level data for this study. The dataset contains one record for each combination of school and student subgroup for each school year from 2015/16 through 2018/19, for all schools—elementary, middle, and high schools—included in accountability in those years (2,678 schools). For each school, the analysis includes all tested students across courses and grade levels, in the following eight subgroups: racial categories (White students, Black students, Hispanic students, Asian students, multiracial students); economically disadvantaged students; students with disabilities; and English learner students. The Native American/Alaska Native and Hawaiian/Pacific Islander student groups are excluded because there were fewer than five schools with at least 10 students in those subgroups. Key variables include the percentage of students scoring at or above the state’s threshold for academic proficiency in each school-subgroup combination and the number of tested students in each school-subgroup combination in each year. For accountability purposes, Pennsylvania uses an average proficiency rate across math and English language arts (ELA); throughout the report, this average of math and ELA proficiency rates is called the school’s proficiency rate.

Methods. The Regional Education Laboratory (REL) team implemented two Bayesian statistical models to stabilize PDE’s academic proficiency rates. For model statements and technical details, see the technical appendix.

- *ATSI model:* The first model mirrors PDE’s Additional Targeted Support and Improvement (ATSI) designation process. This model stabilized average student proficiency rates across two academic years, 2016/17 and 2017/18. The model is cross-sectional, with one data point, the average proficiency rate across 2 years, per school-subgroup combination.
- *TSI model:* The second model mirrors PDE’s Targeted Support and Improvement (TSI) designation process. This model stabilized annual student proficiency rates across 4 academic years, 2015/16, 2016/17, 2017/18, and 2018/19. The model is longitudinal, with one data point per school-subgroup-year combination. To avoid redundancy with ATSI results, results for the TSI model appear in the appendix.

After fitting the models, the study team assessed the effect of stabilization on statistical reliability. In the absence of an error-free benchmark or student-level data with which to compute conventional reliability metrics, the study team approximated reliability by comparing the relationship between sample size and variation in proficiency rates between unstabilized and stabilized estimates. In less reliable estimates, small sample size is associated with more variation—for example, a wider range or larger standard deviation—whereas in more reliable estimates, the correlation between sample size and variation will be weaker. To gauge the effect of stabilization on ATSI designations, the study team compared the set of schools and subgroups identified for ATSI using PDE’s accountability rules with unstabilized versus stabilized proficiency rates.

Limitations. The primary limitation of this study is that, in the time frame required to conduct the study, it was not possible to obtain access to student-level data. Using school-level data constrained the set of models included in the analysis, which may understate the benefits of stabilization. Without student-level data, it is not feasible to calculate classical measures of statistical reliability, so the study team relied on visualizations and descriptive analysis to assess the extent to which stabilization improves the reliability of academic proficiency rates.

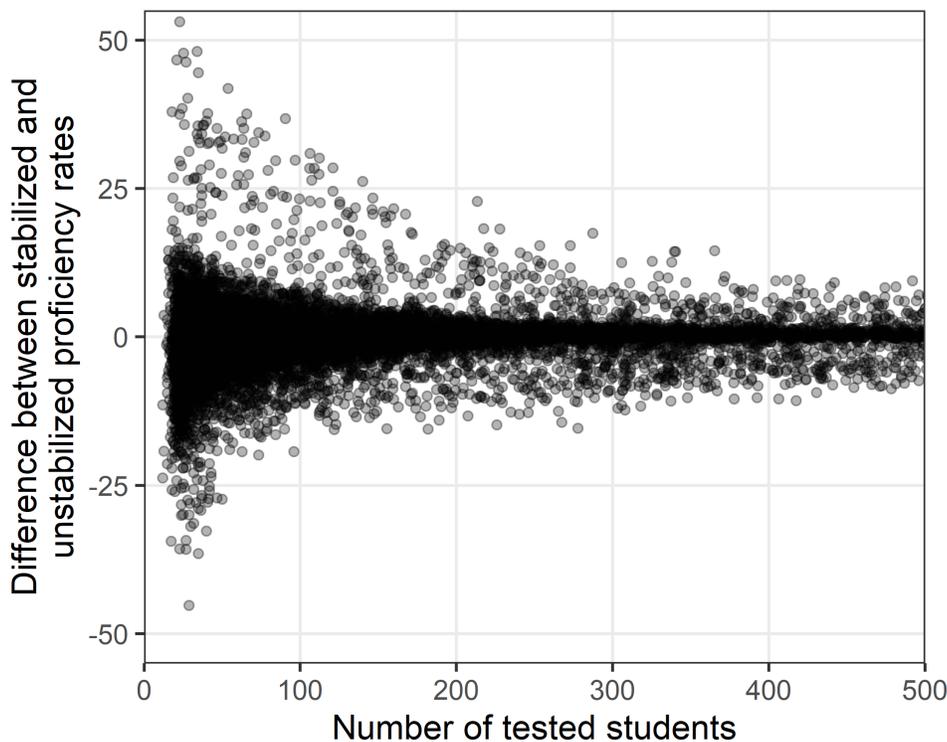
Findings

Stabilization had the greatest effect on small school-subgroup combinations

The average amount of stabilization—that is, the difference between stabilized and unstabilized estimates—for ATSI subgroup proficiency was 3.5 percentage points in either direction, with 75 percent of subgroups stabilized less than 4.5 percentage points; results for TSI were similar (see the technical appendix).

Stabilized proficiency rates differed from unstabilized rates particularly for small school-subgroup combinations. Figure 1 shows the difference between unstabilized and stabilized annual proficiency rates by sample size. The figure has a characteristic funnel shape, with a wider range of differences among smaller school-subgroup combinations (left side of figure) than among larger school-subgroup combinations (right side of figure). This pattern aligns with the expectation that stabilization is most influential for small school-subgroup combinations, where small sample size increases measurement error and correspondingly decreases reliability. Large school-subgroup combinations were minimally affected.

Figure 1. Stabilization was more influential for smaller subgroups (fewer than 100 tested students) than larger ones



Source: Regional Education Laboratory calculations using Pennsylvania Department of Education data.

Note: Each point in the figure represents a combination of school and subgroup for the combined 2016/17 and 2017/18 academic years. The horizontal axis represents the number of tested students in that school-subgroup, and the vertical axis represents the difference between the stabilized and unstabilized academic proficiency rates for that school-subgroup. The funnel shape of the points, with greater dispersion on the left of the figure than on the right, indicates that stabilization affects smaller schools more than larger ones, in line with theory.

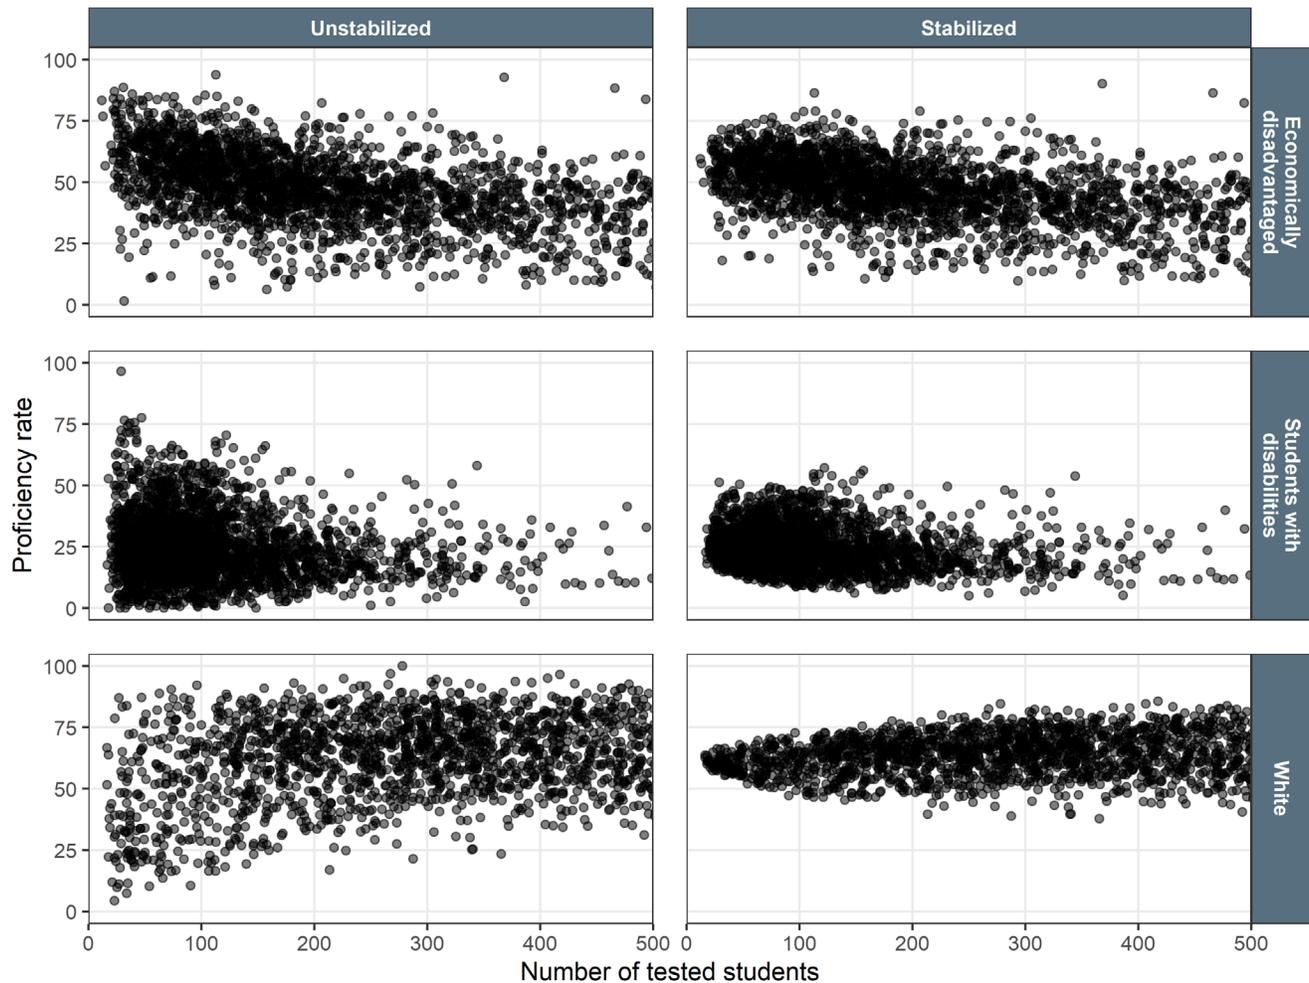
Stabilization moderated the relationship between subgroup size and variation in proficiency rates, indicating an improvement in statistical reliability

The unstabilized data also showed a characteristic funnel pattern; proficiency rates among smaller school-subgroup combinations varied more than proficiency rates among larger school-subgroup combinations. There is little reason to expect that variation in true academic proficiency would be greater for small groups of students

than for large groups of students, so this relationship is most likely to reflect measurement error—instability in the proficiency rate estimates for smaller school-subgroup combinations, due simply to their size.

To assess how much stabilization increased statistical reliability, the study team visually compared the relationship between variability and sample size in unstabilized and stabilized proficiency rates. Figure 2 depicts this relationship, separately for three of the eight subgroups used to make accountability decisions in Pennsylvania.

Figure 2. Compared to unstabilized data, stabilized data showed a weaker relationship between variability and sample size



Source: Regional Education Laboratory calculations using Pennsylvania Department of Education data.

Note: The rows represent three of the eight student subgroups included in the study. In each panel, the horizontal axis represents the number of tested students in a school-subgroup combination, and the vertical axis represents the proficiency rate in that school-subgroup combination. Each data point represents a two-year average of proficiency rates in a given school and subgroup for the combined 2016/17 and 2017/18 academic years. The left column shows unstabilized proficiency rates, while the right column represents stabilized proficiency rates.

Figure 2 shows the relationship between the number of tested students and the proficiency rate for the stabilization model mirroring data used in PDE’s ATSI accountability rules in three of the student subgroups included in the study. Each panel shows this relationship for unstabilized data (left column) and stabilized data (right column).

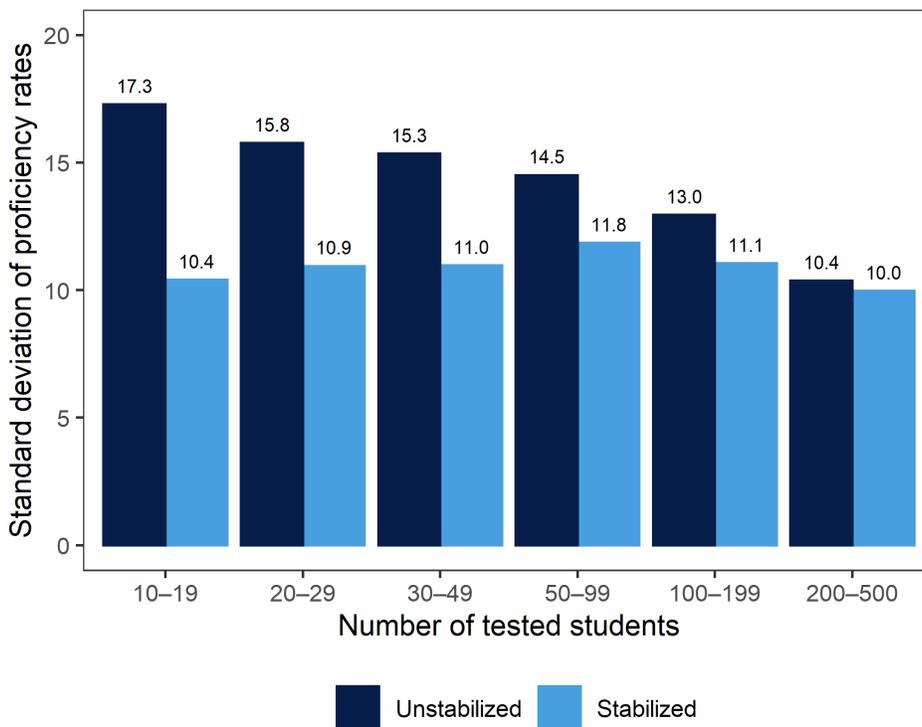
In all rows of figure 2, the unstabilized points (left column) follow the characteristic funnel shape, spanning a wider range on the vertical axis for smaller sample sizes than larger sample sizes, with smaller subgroups much more likely than larger subgroups to have extreme values due to measurement error. As a result of their greater

susceptibility to measurement error, smaller subgroups are more likely than larger subgroups to trigger ATSI identification just by bad luck.

In the stabilized points (right column) in figure 2, by contrast, a relationship between subgroup size and the range of estimated proficiency is less evident. Increased consistency in the estimates' variability across subgroup sizes suggests that stabilization has improved the statistical reliability of the proficiency rates.

To provide quantitative support for these relationships, the study team calculated the standard deviation of unstabilized and stabilized proficiency rates separately in each combination of subgroup and six sample size categories: school-subgroups with 10–19, 20–29, 30–49, 50–99, 100–199, and 200–500 tested students. Within a sample size category, the study team then took the median and interquartile range (IQR) of the subgroup-specific standard deviations, as a summary of the distribution across subgroups. Finally, these median standard deviations of unstabilized and stabilized proficiency rates were compared by sample size category (figure 3).

Figure 3. Stabilization substantially reduced the variability of proficiency rates for small subgroups, making the median standard deviation relatively constant across sample size categories



Source: Regional Education Laboratory calculations using Pennsylvania Department of Education data.

Note: Sample size categories are along the horizontal axis, and the median of subgroup-specific standard deviations in proficiency rates are on the vertical axis. Darker bars show the median of subgroup-specific standard deviations of unstabilized estimates for a certain sample size category, while lighter bars show the median subgroup-specific standard deviation of stabilized estimates for that sample size category. Data is for the combined 2016/17 and 2017/18 academic years.

In figure 3, the horizontal axis shows the sample size categories, from smallest to largest, and the vertical axis gives the median of the subgroup-specific standard deviations of proficiency rates for subgroups of that size. The bar color distinguishes between calculations based on unstabilized (darker) and stabilized (lighter) estimates.

The dark bars, representing the medians of subgroup-specific standard deviations of unstabilized proficiency rates, decrease with increasing sample size. This relationship reaffirms the pattern in figure 2, where proficiency rates for smaller subgroups were more widely dispersed than proficiency rates for larger subgroups. The lighter bars (stabilized), by contrast, are more similar across subgroups. Because the additional variation in smaller subgroups is likely to be measurement error, this result indicates that the stabilized rates are more reliable—less

prone to measurement error—than the unstabilized rates. Indeed, the consistency of the median standard deviations of stabilized proficiency rates across sample size categories suggests that with stabilization, the proficiency rates of smaller subgroups (10–19 students; median 10.4, IQR 5.5 to 12.4) are more reliable than the unstabilized proficiency rates of larger subgroups (20 or more students; median 15.8, IQR 14.0 to 18.3).² Thus, stabilization may make it possible to include smaller subgroups in accountability, without sacrificing statistical reliability.

Stabilization moved a modest number of subgroups above the proficiency cutoff for Additional Targeted Support and Improvement

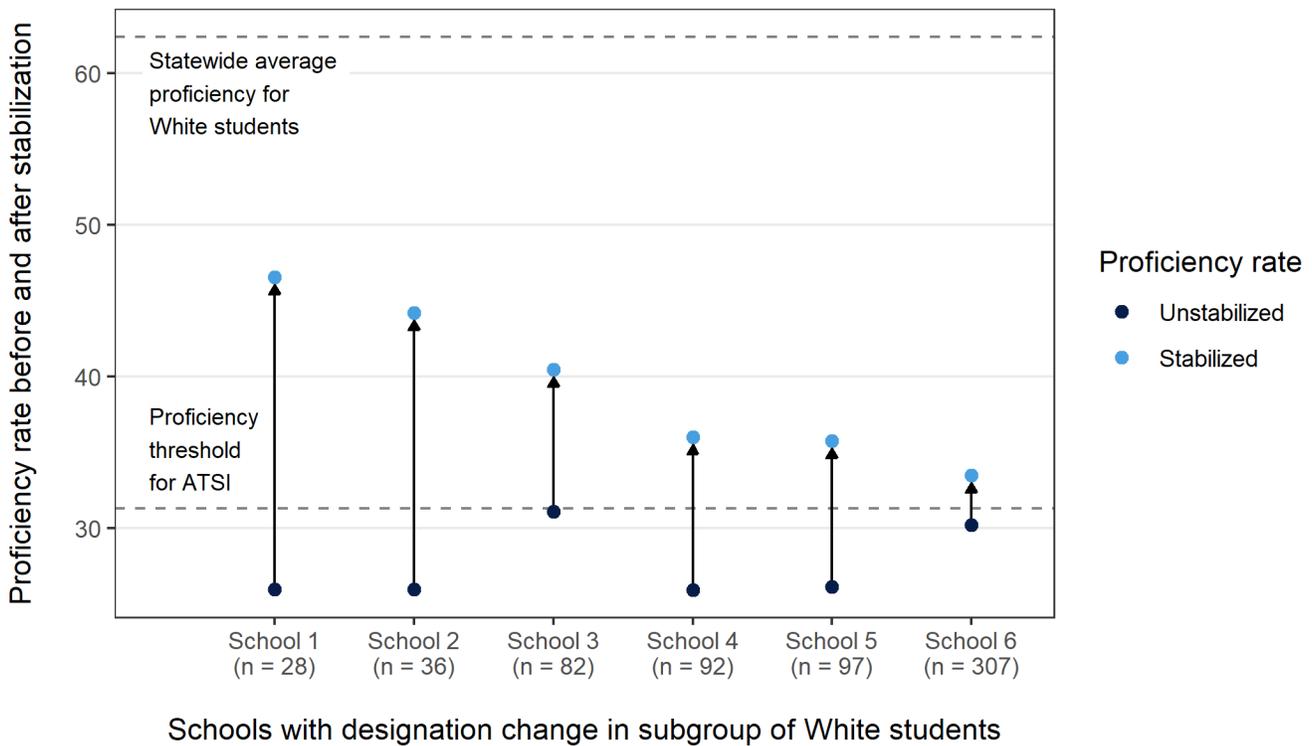
After concluding that stabilization increased the reliability of academic proficiency rates, the study team assessed the effect of including stabilized proficiency rates in PDE’s ATSI accountability rules instead of unstabilized proficiency rates. PDE applies a two-phase system of decision rules to identify schools for ATSI. In the first phase, PDE assesses whether schools are eligible for ATSI based on the combination of their academic proficiency and academic growth rates, relative to performance cutoffs that vary by the school’s academic proficiency rate. For example, a school with a proficiency rate between 20 and 30 percent must receive a negative academic growth rate to be eligible for ATSI, whereas a school with a proficiency rate between 10 and 19.9 percent is eligible for ATSI if its academic growth rate is below a designated positive threshold. In the second phase, PDE considers the eligible schools’ performance on the remaining accountability indicators, relative to performance cutoffs for those indicators that are determined in each year based on the distributions of the observed data.

The study team compared ATSI designations computed using PDE’s accountability rules with unstabilized rates to the same calculation performed using stabilized rates; the accountability data for the five remaining accountability indicators were not adjusted in any way.³ Stabilization did not substantially change ATSI accountability designations. Of the 2,678 schools in the analytic dataset (90 of which were designated as CSI, and so not eligible for ATSI designation), 193 were designated as ATSI by PDE. Of these 193 schools, 9 had a subgroup that switched from ATSI to non-ATSI designation under stabilization. In these nine schools, the subgroup switching from ATSI to non-ATSI designation under stabilization did not result in a change to the school’s overall ATSI designation because in each school at least one other subgroup retained ATSI designation after stabilization. In other words, even though stabilized results did not remove any school’s designation as ATSI, stabilization provided information allowing those schools to better focus on the subgroups for which true performance was most likely in the ATSI range.

² The wider interquartile range for stabilized than unstabilized standard deviations reflects variability in the amount of stabilization across subgroups.

³ In practice, PDE may use its discretion to adjudicate difficult cases. For this comparison, the accountability rules were applied with no discretionary adjustments.

Figure 4. Stabilization changed Additional Targeted Support and Improvement designation for the subgroup of White students in six schools; for these schools, stabilization toward the statewide average proficiency for White students pulled proficiency rates above the cutoff



Source: Regional Education Laboratory calculations using Pennsylvania Department of Education data.
 ATSI is Additional Targeted Support and Improvement.
 Note: Data is for the combined 2016/17 and 2017/18 academic years.

The nine ATSI-to-non-ATSI subgroup changes were observed among White (six schools, figure 4) and economically disadvantaged (three schools, not shown) subgroups. It is expected that stabilization would move the White and economically disadvantaged subgroups from ATSI to non-ATSI (rather than in the opposite direction) because stabilization pulls individual subgroups’ proficiency rates toward the mean proficiency rate in that subgroup,⁴ and in these subgroups the mean proficiency rate is above the ATSI proficiency cutoff.

In principle, stabilization could also reduce a subgroup score enough to move a school into ATSI status—for example, if the mean proficiency rate in the subgroup were below the proficiency cutoff. The study team did not examine that possibility here, for two reasons. First, this is not likely to occur frequently because statewide average proficiency rates are usually higher than the ATSI proficiency cutoff; low-scoring schools are therefore more likely to see stabilized scores move up than down. (The only subgroups in Pennsylvania with average proficiency rates below the ATSI proficiency cutoff are students with disabilities and English learner students.) Second, Pennsylvania is considering using stabilized scores only as a safe harbor that could remove a school from ATSI designation but not move a school into ATSI status.

⁴ The core assumption of the Bayesian stabilization models implemented in this study is that in the absence of information to the contrary, a specific school’s proficiency rate will be close to the overall average proficiency rate. Because the study team fit separate stabilization models for each subgroup, the assumption becomes that a specific school-subgroup’s proficiency rate is likely to be close to the average proficiency rate for that subgroup, across schools. As a result, the proficiency rates for smaller school-subgroup combinations are pulled toward their subgroup mean.

Limitations

This study's results suggest that stabilization has the potential to improve the statistical reliability of school accountability measures. However, the study's use of school-level, rather than student-level, data limits both the methods employed in the study and the conclusions that can be drawn from its results.

The use of school-level data constrains the set of models the study team tested and the metrics used to evaluate them. Without student-level data, the study team could not take advantage of repeated measures of students over time or of student-level variability. To avoid overstating the statistical precision of the results due to students who belong to more than one subgroup, the study team fit models separately for each subgroup, rather than capitalizing on relationships across subgroups. Nonetheless, the study team overstated precision to some degree in both the ATSI and TSI models, even within a subgroup, because it was not possible to adjust appropriately for correlation across students whose test scores were included in more than one academic year.

At the same time, without student-level data the study team could not calculate traditional measures of statistical reliability. Instead, the study team had to rely on visualizations and other descriptive assessments to gauge whether stabilization achieves the desired goals.

Finally, most of the subgroups with designation changes had proficiency rates close to the ATSI cutoff, with the stabilized proficiency just on the other side of the cutoff. For these borderline cases, binary decision rules for ATSI designation may overstate the study team's confidence in the available evidence about school or subgroup performance. Accounting for statistical uncertainty could allow for a more nuanced understanding. Future work that estimates uncertainty bounds, either around the stabilized proficiencies or around the accountability designations, could take full advantage of the properties of the Bayesian stabilization models explored in this study.

Implications

The results of this analysis indicated that stabilization could reduce the tension between reliability and inclusivity that characterizes accountability for small schools and subgroups. Stabilization improved the reliability of subgroups' proficiency rates; the relationship between subgroup sample size and variation in proficiency rates, which largely reflects measurement error, was weaker for stabilized than unstabilized proficiency rates.

Indeed, the results suggest that Bayesian stabilization could permit Pennsylvania to reduce its minimum subgroup size from 20 students to 10 while simultaneously reducing the likelihood of erroneously identifying a school for ATSI. In the context of Pennsylvania's ATSI calculations, stabilization improved reliability for small subgroups of 10–19 students enough that their distribution of stabilized scores had a similar, or even smaller, variance than the distribution of unstabilized scores of much larger subgroups of up to 200 students (see figure 3). This smaller variance implies that, with stabilization, proficiency rates for smaller subgroups may no longer be too unreliable or imprecise to use in accountability. Including smaller subgroups in accountability could be an important step toward ensuring that all students—even those in small schools or subgroups—receive additional support when they need it.

Despite increasing the reliability of the estimated proficiency rates, stabilization had little effect on subgroups' identification for ATSI. Nonetheless, stabilized results would provide better information to PDE and to schools about the subgroups with performance that is truly below the designated threshold. In about 5 percent of schools identified for ATSI in Pennsylvania in 2018, stabilization suggested that one subgroup that was below the ATSI proficiency cutoff should have scored above the cutoff. Although these changes did not affect whether the school as a whole would be identified for ATSI, they would permit the school to target improvement efforts more effectively.

These results will inform analyses of 2022 PDE data, which in turn may affect future ATSI and TSI designations. On the strength of these results, REL Mid-Atlantic hopes to collaborate with PDE to implement both models using data through 2022, with some modifications to account for the long-term effects of the COVID-19 pandemic.

Whether Bayesian stabilization improves the reliability and statistical precision of the estimated proficiency rates is just one of several important considerations for states or districts that contemplate using these methods. Simplicity and transparency are important qualities of a school accountability system, so that schools understand how they are assessed and deem the identifications credible. Despite better statistical properties, a system that relies on Bayesian stabilization is more complex and opaque than the existing system. Even so, improvements in the year-on-year consistency of accountability indicator scores, reflecting increased reliability, could gain stakeholders' trust. Individual states and districts must weigh these factors to determine whether the stabilization approach best meets the goals of their accountability system.