Measuring Discipline Disparities

James P. Scanlan

(Statement Prepared for U.S. Commission on Civil Rights Briefing "The School to Prison Pipeline: The Intersection of Students of Color with Disabilities" (Dec. 8, 2017)

Federal government policy regarding racial differences in school discipline outcomes has been consistently based on the belief that relaxing discipline standards and otherwise reducing adverse discipline outcomes will tend to reduce (a) relative (percentage) racial differences in rates of experiencing the outcomes and (b) the proportions African Americans and other racial minorities make up of persons experiencing the outcomes. In fact, exactly the opposite is the case.

By way of clarification, if the minority suspension rate is 15% and the white rate is 5%, the ratio of the minority rate to the white rate would be 3.0. That is, the minority rate is 200% greater than the white rate. The 200% figure is the relative, or percentage, difference. In the same situation, assuming minorities are 20% of students, they would be 43% of suspended students.

Federal policy has been based on the belief that activities that generally reduce suspensions (like Positive Behavioral Interventions & Support (PBIS) programs) will tend to reduce the 3.0 ratio and the 43% proportion figures. In fact, such activities will tend to increase those figures.

Test Score Illustration

Table 1 provides a simple illustration of why this is the case. The table is based on hypothetical test scores of higher- and lower- scoring groups (which are denominated AG for advantaged group and DG for disadvantaged group).

The first row of the table shows the pass rates for the two groups at a particular cutoff. The pass rates are 80% for AG and 63% for DG. Thus, AG's pass rate is 1.27 times (27% greater than) DG's pass rate.[1]

[1] While I commonly refer to patterns of relative differences in this statement, the table actually presents rate ratios (also termed risk ratios or relative risks). The relative difference is the rate ratio minus 1 where the rate ratio is above 1 and 1 minus the rate ratio where the rate ratio is below one. In the former case, the larger the rate ratio, the larger the relative difference; in the latter case, the smaller the rate ratio, the larger the relative difference. It is more common to employ the disadvantaged group's rate as the numerator for the favorable as well as the adverse outcome, which is the approach as to favorable outcomes of the "four-fifths" or "80 percent" rule for identifying disparate impact under the Uniform Guideline for Employee Selection Procedures. I have sometimes employed this approach, as in "Can We Actually Measure Health Disparities?," *Chance* (Spring 2006) (http://www.jpscanlan.com/images/Can_We_Actually_Measure_Health_Disparities.pdf). More recently, however, I have usually used the larger figure as the numerator for both rate ratios, in which case, as to both favorable and adverse outcomes, the larger the rate ratio, the larger the relative difference. Choice of numerator in the rate ratio, however, has no bearing the patterns described here whereby measures tend to be affected by the prevalence of an outcome.

Table 1. Illustration of effect of lowering test cutoff on relative difference between pass rates of advantaged group (AG) and disadvantaged group (DG)

Cutoff	AG Pass Rate	DG Pass Rate	AG/DG Pass Ratio
1 High	80%	63%	1.27
2 Low	95%	87%	1.09

The second row shows what would happen if the cutoff is lowered to the point where AG's pass rate is 95%. Assuming normal test score distributions, DG's pass rate would be about 87%. With the lower cutoff AG's pass rate would be only 1.09 times (9% greater than) DG's pass rate. The fact that lowering a cutoff tends to reduce relative differences in pass rates is the reason why lowering a test cutoff is universally regarded as reducing the disparate impact of tests on which some groups outperform others.

At this point it may seem that I have contradicted my point at the outset. But, whereas lowering a cutoff tends to reduce relative differences in pass rates, it tends to increase relative differences in failure rates. This pattern is illustrated in Table 2. The table presents the same information as Table 1, but with the failure rates of the two groups added, along with the ratio of DG's failure rate to AG's failure rate (in the final column). The column with the rate ratios for test passage is highlighted in blue and the column with the rate ratios for test failure is highlighted in red.

Table 2. Illustration of effect of lowering test cutoff on (a) relative difference between pass rates and (b) relative difference between failure rates of advantaged group (AG) and disadvantaged group (DG)

Cutoff	AG Pass Rate	DG Pass Rate	AG Fail Rate	DG Fail Rate	AG/DG Pass Ratio	DG/AG Fail Ratio
1 High	80%	63%	20%	37%	1.27	1.85
2 Low	95%	87%	5%	13%	1.09	2.60

The final (red highlighted) column shows that with the initial cutoff DG's failure rate was only 1.85 times (85% greater than) AG's pass rate. With the lower cutoff, DG's failure rate is 2.60 times (160% greater than) AG's failure rate.

That is, as the prevalence of test passage and test failure generally changed as a result of lowering the cutoff, the relative difference in the increasing side of the dichotomy (test passage) decreased and the relative difference in the decreasing side of the dichotomy (test failure) increased.

As suggested at the outset, appraisals of discipline disparities issue sometimes focus on the proportions racial minorities make up of persons disciplined (compared with the proportions such groups make up of students). Patterns of changes in the proportions groups make up of persons experiencing either of the two outcomes as the prevalence of the outcomes changes are corollaries to the patterns shown in Table 2.

Table 3 is the same as Table 2, but with two more columns added on the right. These columns show the proportions DG makes up of persons who pass the test (highlighted in blue) and persons who fail the test (highlighted in red) in circumstances where DG makes up 50% of persons who take the test.

Table 3. Illustration of effect of lowering test cutoff on (a) relative difference between pass rates and (b) relative difference between failure rates of advantaged group (AG) and disadvantaged group (DG) and proportion DG makes up of (c) persons who pass the test and (d) persons who fail the test (where DG makes up 50% of test takers)

Cutoff	AG Pass Rate	DG Pass Rate	AG Fail Rate	DG Fail Rate	AG/DG Pass Ratio	DG/AG Fail Ratio	DG Prop of Pass	DG Prop of Fail
1 High	80%	63%	20%	37%	1.27	1.85	44%	65%
2 Low	95%	87%	5%	13%	1.09	2.60	48%	72%

The penultimate column shows that lowering the cutoff causes the proportion DG makes up of persons who pass the test to increase from 44% to 48%. That would reduce the difference between the proportion DG makes up of persons who take the test and the proportion it makes up of persons who pass the test.

But the final column shows that lowering the cutoff also increased the proportion DG makes up of persons who fail the test, from 65% to 72%. That would increase the difference between the proportion DG makes up of persons who take the test and the proportion DG makes up of persons who fail the test.

These patterns are not peculiar to test score data or the numbers I used to illustrate them. Rather, changing the frequencies of virtually any outcome and its opposite tends to cause the relative difference in the increasing outcome to decrease and the relative difference in the decreasing outcome to increase (with related effects on the proportions groups more susceptible to the outcomes make up of persons who experience the increasing outcome and the decreasing outcome).

This will not invariably happen with the consistency that will be observed with hypothetical test score data. For many factors are at work. But it will typically happen, especially when the changes in the prevalence of an outcome are substantial. In the school discipline context in particular, generally reducing discipline rates, while tending to reduce relative racial differences in rates of avoiding discipline (analogous to test passage), will tend to increase relative racial differences in rates of being disciplined (analogous to test failure). And in fact that is being observed all across the country as school districts have been generally reducing discipline rates while mistakenly believing that doing so should reduce relative racial differences in discipline rates (or the proportions racial minorities make up of student who are disciplined).[2]

[2] See page 8 of my July 17, 2017 letter to the Departments of Education, Health and Human Services, and Justice. http://www.jpscanlan.com/images/Letter_to_Departments_of_Education,_HHS,_and_Justice_July_17,_2017_.pdf

It is important to recognize that the situation is not one where the government has reasoned that, while the above-described patterns will be found in test score data, there are reasons why the patterns will not ordinarily be found in other situations. Rather, despite dealing with issues about demographic differences in test outcomes for half a century, the government has failed even to understand that lowering a test cutoff tends to increase relative differences in failure rates.

It is also important to understand that an increase in the relative difference in the adverse outcome does not mean that a disparity has increased in some meaningful sense any more than the reduction in the relative difference in the favorable outcome means that a disparity has decreased in a meaningful sense. Rather, the problem is that neither relative difference is a useful indicator of the strength of the forces causing the outcome rates of two groups to differ (or, as we might otherwise put it, the size of the difference in the circumstances of two groups reflected by their outcome rates). That is quite important to recognize as we endeavor to understand the causes of disparities and determine whether they are growing larger or smaller over time or are larger in one setting than another.

Still focusing on either Table 2 or Table 3 (though the former is somewhat simpler), one may think of the pass and fail rates as reflecting any favorable and adverse outcome rates that result from decisions of individual decision-makers. In the school discipline context, consider the failure rates as if they are the suspension rates of minorities and whites and the pass rates as if they are the groups' rates of rates of avoiding suspension. To the extent that bias on the part or decision-makers contributes to the differences between rates, any actions that reduce that bias will tend to reduce all measures of racial differences between favorable or adverse outcomes.

At the same time, however, simple reductions in adverse discipline outcomes, such as those resulting from PBIS programs, will tend to change the measures of difference in the manner reflected in the tables. Thus, in consequence of general reductions in discipline rates, a school district that substantially reduces suspension rates will tend to show a pattern of changing measures of differences in outcome rates akin to that found in movement from the first row to the second row of the two tables.

In circumstances where decision-makers, including teachers and administrators, are being encouraged to generally reduce suspension rates, all other things being equal, the results for decision-makers who do not try very hard to reduce suspension rates will tend to look more like the first row than the second row. The results for decision-makers who try very hard to reduce suspension rates will tend to look more like the second row than the first row.

Thus, consider a situation where the two rows reflect the results of actions of two different decision-makers and an effort is made to determine which decision-maker is more likely to have made racially biased decisions. One would reach opposite conclusions depending on whether one examined relative differences in the favorable outcome or relative differences in the adverse outcome. In fact, however, there is no rational basis for distinguishing between the two rows with regard to the question of which is more likely to reflect the results of biased decisions.

It should be evident that it is essential for school administrators endeavoring to address discipline disparities issues, and those monitoring those efforts and otherwise attempting to ensure equal

treatment for all groups, to understand these patterns. Yet the situation is not simply that virtually no one involved in such efforts understands these patterns; rather, virtually everyone involved in such efforts proceeds on a belief about the effects of generally reducing discipline rates on the measures most commonly employed in quantifying racial and other demographic disparities that is the opposite of reality.

Illustration of the Effects of Substituting a Reprimand for What Would Otherwise Be a First Suspension on Proportions More Susceptible Groups Make up of Persons Suspended

Data made available in Department of Education reports provide other simple illustrations of the effects of generally reducing adverse discipline outcomes rates on measures of racial or other demographic differences in discipline outcomes.

Tables 4 and 5 are based on data from a March 21, 2014 Department of Education report titled "Data Snapshot: School Discipline."[3] The data in the report enable one to determine the proportions demographic groups make up of K-12 and preschool students who are suspended (a) one or more times and (b) two or more times.[4]

Table 4. Illustration of effect of giving all students a reprimand instead of their first suspension on proportion African Americans make up of K-12 and preschool students receiving one or more suspensions

Setting	Number of Suspensions	AA Proportion of Students Experiencing the Outcome
K-12	One or more	37%
K-12	Two or more	43%
Preschool	One or more	44%
Preschool	Two or more	48%

Table 4 provides that information with regard to African American students in K-12 and preschool. The first row of the first set of two rows shows the proportion African Americans make up of K-12 students suspended one or more times (37%) and the second of those rows shows the proportion they make up of K-12 students suspended two or more times (43%). Suppose, then, that in every situation that otherwise would have resulted in a first suspension, the students were given a reprimand rather than a suspension. In such case, the figure in the second row would tend to become the figure for one or more suspensions. Thus, the 37% figure for the proportion African Americans make up of K-12 students suspended one or more times would tend to rise to 43%.

[3] https://www2.ed.gov/about/offices/list/ocr/docs/crdc-discipline-snapshot.pdf

[4] The document provided information on the proportions demographic groups made up of K-12 and preschool students suspended one time and suspended multiple times. From the information provided in the report, one can then determine the proportions the groups made up of persons suspended (a) one or more times and (b) two or more times.

The second two rows of the table provide a similar illustration for preschool. In this setting, giving students a reprimand instead of their first suspension would tend to cause the proportion African Americans make up of students suspended one or more times to increase from 44% to 48%.

Table 5 presents the same type of information for boys, who commonly have higher suspension rates than girls and thus commonly make up a larger proportion of suspended students than the approximately 50% that they make up of all students. Here, too, the Department of Education data show that in both K-12 and preschool, giving students a reprimand rather than what would otherwise be their first suspension would tend to increase the proportion boys (the group more susceptible to suspension) make up of students suspended one or more times.

Table 5. Illustration of effect of giving all persons a reprimand instead of their first suspension on proportion boys make up of K-12 and preschool students receiving one or more suspensions

Setting	Number of Suspensions	Male Proportion of Students Experiencing the Outcome
K-12	One or more	70%
K-12	Two or more	72%
Preschool	One or more	80%
Preschool	Two or more	82%

Illustration of Effects of the Prevalence of Adverse Discipline Outcomes in Different Settings on Measures of Racial Disparity in Those Settings

I often describe the statistical pattern at work in the discipline context (and essentially every other context where disparities are quantified in terms of relative differences or measures that are functions of relative differences) as that whereby the rarer and outcome, the greater tends to be the relative difference in experiencing it and the smaller tends to be the relative differences in avoiding it. One important, though universally misunderstood, manifestation of that pattern is that in settings (or among subpopulations) where adverse discipline outcomes are comparatively uncommon, relative racial differences in rates of experiencing those outcomes will tend to be larger, while relative differences in the corresponding favorable outcome will tend to be smaller, than in settings where the outcomes are comparatively common.

Tables 6 and 7 are based on data from the Massachusetts and Loudoun County, Virginia. Both are areas where policymakers or others have expressed concern that, though the areas have comparatively low suspension rates, relative racial differences or other measures of racial differences in suspensions are comparatively high.

The two tables may be compared to Table 2 above (save that they do not show the rates at which the two groups avoid suspension, the equivalent of test passage) with columns reordered to be more consistent with the way the issues are commonly discussed (and with the same color-coding for the rate ratios for the adverse and favorable outcomes). But I have added an additional column at the end termed EES, for estimated effect size. This column presents a measure of the strength of the forces causing outcome rates of two groups to differ that is theoretically unaffected by the prevalence of an outcome. I describe it (and its strength and weaknesses) in my "Race and Mortality Revisited," *Society* (July/Aug. 2014)[5] and various other places.

Table 6: Out-of-school suspension rates for African American and white students in Massachusetts and nationally in 2012-2013, with measures of difference

Area	AA Rate	White Rate	AA/White Ratio-Susp	White/AA Ratio - No Susp	EES
Massachusetts	10.0%	2.7%	3.70	1.08	0.65
National	16.4%	4.6%	3.57	1.14	0.71

Table 6 shows the common patterns whereby the setting with comparatively low suspension rates (Massachusetts compared with national figures) shows larger relative differences in suspension rates, but smaller relative differences in rates of avoiding suspension, than are found nationally. The EES figures – .65 in Massachusetts and .71 nationally – indicate that the forces causing suspension rates of African American and white students to differ (whatever those forces may be) are weaker in Massachusetts than nationally.[6]

Table 7 presents similar information from schools in Loudoun, County Virginia (an affluent suburb of Washington, DC), where suspension rates are very low. In this case, the concern about large racial disparities was prompted by the comparatively high ratio of the proportion African Americans made up of suspended students to the proportion they made up of students.[7]

[5] http://jpscanlan.com/images/Race_and_Mortality_Revisited.pdf

[6] These data and similar data relating to students with disabilities are discussed more fully in my November 12, 2017 letter to the Boston Lawyers' Committee for Civil Rights and Economic Justice. http://jpscanlan.com/images/Letter_to_Boston_Lawyers_Committee_Nov._12,_2015_.pdf

[7] That areas with low African American representation among students tend to have higher such ratios than other areas even when the areas have same suspension rates for African American students and for other students is among a number of reasons beyond the statistical patterns addressed here that comparisons of the proportion a group makes up of persons potentially experiencing an outcome and the proportion the group makes up of persons actually experiencing the outcome cannot effectively quantify the forces causing outcome rates of advantaged and disadvantaged groups to differ. See references in the succeeding note. See also the IDEA Data Center Disproportionality Guide subpage of the Discipline Disparities page of jpscanlan.com. http://jpscanlan.com/disciplinedisparities/ideadatacenterguide.html

The ratio African American suspension rate to the white suspension rate is actually slightly lower in Loudoun County than nationally, while the relative difference in rates of avoiding suspension is much lower in Loudoun County than nationally. The EES figures – .55 in Loudoun County and .71 nationally – indicate that the forces causing suspension rates of African American and white students to differ are considerably weaker in Loudoun County than nationally.[8]

Table 7: Out-of-school suspension rates for African American and white students in Loudon County (VA) Public Schools and nationally in 2012-2013, with measures of difference

Area	AA Rate	White Rate	AA/White Ratio-Susp	White/AA Ratio - No Susp	EES
LCPS	4.65%	1.3%	3.54	1.04	0.55
National	16.4%	4.6%	3.57	1.14	0.71

Neither Massachusetts nor Loudoun County has any idea that to the extent that racial disparities in school discipline can be effectively measured, their disparities are smaller, not larger, than nationally. Nor do they have any idea that the actions to generally reduce discipline rates that they see as means of reducing the measures of racial disparity that are causing them concern will in fact tend to increase those measures.

Table 8, which is based on Table 8 of the aforementioned "Race and Mortality Revisited," is similar to Tables 6 and 7. But rather than comparing figures from a particular geographic area with national figures, Table 8 compares figures in preschool (where suspensions are comparatively rare) with figures from K12 (where suspensions are much more common). The table presents figures on multiple suspensions, which is the outcome respecting which racial disparities received the greatest attention when racial disparities in preschool suspensions first received substantial attention in 2014.[9]

[8] These data are discussed more fully in the Loudoun County (VA) Disparities subpage of the Discipline Disparities page of jpscanlan.com (http://jpscanlan.com/disciplinedisparities/loudounctydisparities.html). That subpage also discusses data showing that between the 2009-2010 and the 2013-2014 school years general reductions in suspension rates were accompanied by an increase in the relative differences between African American and white suspension rates and a decrease in the relative difference between African American and white rates of avoiding suspension, with negligible change in the EES. See also my September 5, 2017 letter explaining this issue to the Loudoun County School Board.
http://jpscanlan.com/images/Letter_to_Loudoun_County_Public_Schools_Sept._5,_2017_.pdf

[9] The facts receiving special attention in coverage of the issue were that African Americans were 18% of preschool children but 48% of preschool students receiving multiple suspensions. The figures in Table 8 are the suspension rates that can be derived from data in the previously mentioned Department of Education March 2014 document "Data Snapshot: School Discipline." The 18% and 48% figures were also highlighted in a March 21, 2014 Department of Education report titled "Data Snapshot: Early Childhood Education."
https://www2.ed.gov/about/offices/list/ocr/docs/crdc-early-learning-snapshot.pdf

Table 8. African American and white rates of multiple suspensions in preschool and K-12, with measures of difference

Level	AA Mult Susp Rate	White Mult Susp Rate	AA/Wh Ratio Mult Susp	Wh/AA Ratio No Mult Susp	EES
Preschool	0.67%	0.15%	4.41	1.01	.49
K12	6.72%	2.23%	3.01	1.05	.51

As will commonly be observed, Table 8 shows that in the setting where suspensions are less common (preschool) relative differences in multiple suspension rates are greater, while relative differences in rates of avoiding multiple suspensions are smaller, than in the setting where suspensions are more common (K-12). In this case, however, the EES figures are very similar suggesting that, whatever the forces causing African American and white suspension rates to differ, they are of approximately the same strength in the two settings.

Table 9 is based on data from a 2012 Department of Education report titled "Helping to Ensure Equal Access to Education: Report to the President and Secretary."[10] Data were provided only on the proportion African Americans make of students and expelled students overall and in zero tolerance schools. The actual expulsions rates were not available. But based on the data available, one can present those two proportions in each setting and derive therefrom the relative difference between the African American rate and the rate for all other students.

Table 9: Proportions African Americans make up of students and expelled students overall and in schools with zero tolerance policies, with ratio of the African American expulsion rate to the white expulsion rate

Setting	AA Proportion of Students	AA Proportion of Expulsions	AA/Non-AA Expulsion Ratio
Overall	18%	39%	2.91
Zero Tolerance Schools	19%	33%	2.10

In accordance with the pattern described above, the ratio of the African American expulsion rate to the expulsion rate of other students was higher where expulsions were presumably less common (overall) than in the setting where expulsions were presumably more common (zero tolerance schools). (I do not present an EES figure because one needs the actual expulsion rates to derive such figure.) There nevertheless continues to be a near universal belief that zero tolerance policies lead to larger relative racial differences in adverse disciplines outcomes (and larger African American proportions or persons experiencing those outcomes) than more lenient policies.

An understanding of these patterns is also essential to drawing sound inferences about processes based on the comparative size of disparities. Relative racial differences in suspension rates are commonly greater, while relative differences in rates of avoiding suspension are commonly smaller, among girls (where suspensions are less common) than among boys (where suspensions

[10] http://www2.ed.gov/about/reports/annual/ocr/report-to-president-2009-12.pdf

are more common). Correspondingly, relative gender differences in suspension are commonly greater, while relative gender differences in rates of avoiding suspension are commonly smaller, among whites (where suspensions are less common) than among African Americans (where suspensions are more common). See the Discipline Disparities page of jpscanlan.com.[11]

Similarly, relative racial differences in suspensions will commonly be greater, while relative differences in avoiding suspensions will commonly be smaller, among students without disabilities (where suspensions are less common) than among students with disabilities (where suspensions are more common). Correspondingly, relative differences between the suspension rates of students with and without disabilities will commonly be greater, while relative differences between rates at which such groups avoid suspension will commonly be smaller, among whites (where suspensions are less common) than among African Americans (where suspensions are more common).

On cannot draw inferences about processes on the basis that one of these disparities is larger than another, or otherwise usefully hypothesize about why any disparity is larger than another, without understanding the above-described and other patterns by which measures tend to be affected by the prevalence of an outcome.

Conclusion

The failure to understand the ways the prevalence of an outcome affects relative differences in rates of experiencing an outcome and relative differences in rates of avoiding the outcome is but part of a larger failure of the government (and the social science and statistical communities) to understand the ways standard measures of differences between outcome rates of advantaged and disadvantaged group tend to be affected by the prevalence of an outcome. For more extensive treatment of that issue with regard to all analyses of demographic differences in outcome rates in the law and the social and medical sciences, see the aforementioned "Race and Mortality Revisited," my November 14, 2016 Comments for Commission on Evidence-Based Policymaking,[12] and my October 8, 2015 letter to the American Statistical Association.[13] With regard to the way the larger failure has undermined Department of Education analyses of demographic differences regarding student outcomes apart from discipline, see my "Innumeracy at the Department of Education and the Congressional Committees Overseeing It," Federalist Society Blog (Aug. 24, 2017).[14] See also the July 17, 2017 letter to the Departments of Education, Health and Human Services, and Justice mentioned in note 2 *supra*, which, in addition to advising the agencies of their obligations to correct prior guidance to school administrators as to the likely effects of generally reducing discipline rates on measures of discipline disparities, suggests that the agencies halt all funding of research into demographic

[11] http://jpscanlan.com/disciplinedisparities.html

[12] https://www.regulations.gov/document?D=USBC-2016-0003-0135

[13] http://jpscanlan.com/images/Letter_to_American_Statistical_Association_Oct._8,_2015_.pdf

[14] http://www.fed-soc.org/blog/detail/innumeracy-at-the-department-of-education-and-the-congressional-committees-overseeing-it

differences that fails to consider implications of the ways the measures employed tend to be affected by the prevalence of an outcome.

But the mistaken belief that generally reducing an adverse outcome should tend to reduce, rather than increase, relative differences in rates of experiencing the outcome (and the proportions groups more susceptible to the outcome make up of persons experiencing it) – which informs federal civil rights policies regarding criminal justice, lending, employment, and voter qualification, as well as school discipline – is an extreme example of the larger failure of understanding. And it has pernicious consequences. These include the many anomalies where by complying with government encouragements to relax standards and otherwise reduce adverse outcomes, entities covered by civil rights law increase the chances that the government will accuse them of discrimination. Similar anomalies exist in situations where individual actors who comply with their employers' instruction to reduce adverse outcomes increase the chances that their employees will accuse them of discrimination. Further, in contexts where actions that are supposed to be reducing measures of racial disparity are followed by increases in those measures, observers will conclude that the forces causing outcome rates to differ must be growing stronger, often prompting increasing distrust in the fairness of systems.

Such conclusions will not have a sound statistical basis. But so far very few people understand that.