Racial Disparities in School Discipline: Racism or Rational Choice?

Josh Kinsler, Duke University

April 9, 2005
Preliminary Draft - Incomplete

Abstract

Using a cross-section of disciplinary data from North Carolina, I show that differential suspension rates between white and black students can be attributed to rational behavior on the part of school administrators. Racial bias is typically cited as the primary factor in explaining the punishment gap; however, within school analysis suggests racial bias is responsible for little of the overall gap in punishment. Rather, the aggregate disparities in school punishment result from varying school disciplinary levels. We present a theoretical model of the punishment decision which suggests schools with high levels of misbehaving children will use stricter punishment. To test our theoretical results, we examine the deterrent effects of school discipline for both offending and non-offending students. Both groups respond to the threat of discipline. However, non-offenders with observable characteristics associated with misbehavior respond significantly stronger to the threat of discipline. This behavioral response to punishment results in high risk schools utilizing severe discipline. Because black students are more likely to attend a high risk school a racial gap in punishment persists.

1 Introduction

During the 2000-2001 school year in the state of North Carolina, over 160,000 students were cited for violating school policy. As a result, the offending students lost over 800,000 school days to suspension. Black students were impacted most severely, accounting for 55.6% of the total days suspended, while constituting only 31% of the entire student population. The

racial disparity in school discipline is disconcerting, particularly when research suggests strong correlations between suspensions, dropout, and criminal activity.

The disproportionate representation of minority students in the use of suspension is well documented in the education literature. The Children's Defense Fund (1975) found that black students were suspended at twice the rate of any other ethnic group, and at the secondary school level were three times as likely to be suspended as white students. More recent studies have continued to document the over representation of blacks receiving out of school suspensions (Costenbader & Markson, 1994; Skiba et al, 1997). A number of causes for the disparity in suspension have been suggested, including administrators who are biased against black or low income students, greater behavioral problems among black students, and a lack of understanding of black youth behavior (Skiba, 2002; Townsend, 2000). Although a number of potential explanations for the differences in punishment have been suggested, the overarching theme in the literature is black students are unfairly singled out for harsher disciplinary measures in school.

This conclusion has not gone unnoticed by the general public. In just the past few years, school boards all over North Carolina have been grappling with the issue of student discipline and its disproportionate effect on black students. In Union County, the school board attempted to ban corporal punishment from its schools, largely because black students are the primary recipients of the punishment (AP, 3/9/05). In the past five years, out of school suspension rates dropped in Guilford County by 4.5% overall, but have risen 28% for black students. One school board member suggested institutional racism is to blame (News & Record, 12/04/04). Many other articles from across the state echo the general sentiments in these two examples. However, in almost every article the disciplinary data provided is aggregated at the district level rather than the the school level. District level disparities in punishment may result from racial bias on the part of individual schools, but it may also reflect varying disciplinary strategies across schools.

The education literature illustrates racial disparities in a similar fashion, using data aggregated at the district, state, or national level. Little attention has been paid to the variation in disciplinary levels across schools with a state or district. One exception to this is Mendez et. al (2002), who use data from a large school district in Florida to show that schools with the highest suspension rates also tend to serve minority children from low socioeconomic back-

grounds. This pattern of discipline results in higher levels of suspension at the district level for minority and less affluent children. Using only district level data to conclude racial bias can be misleading. Students of various races may be treated identically within a school, with some schools simply choosing harsher disciplinary measures. If more black students attend schools using severe punishment strategies, aggregate disparities in punishment will exist.

In order to differentiate between racial bias and variation in school policy with regards to discipline, examination of within school discipline is necessary. Racial bias in the administration of suspension would suggest that all else equal (severity of misconduct, student history, etc.) a black student either receives an out of school suspension when a white student does not, or receives a longer sentence than the corresponding white student would. Because punishment is determined at the school level, racial bias can only exist at the school level. Disparities in student suspension at the aggregate level may result from bias by individual schools, but is not in itself evidence of racial bias.

The purpose of this study is twofold. First, we want to determine whether racial variation in school discipline is a result of bias a byproduct of varying school behavior. Second, if the latter is true, then why do schools serving minority and low socioeconomic students use harsher disciplinary practices? To answer the second question We develop a simple theoretical model of school discipline incorporating the deterrent effects of punishment. Offending and non-offending students observe the school's disciplinary level and alter their behavior accordingly. The level of discipline in a school varies according to the type of student in the school. Schools serving students who respond strongly to punishment will choose harsher disciplinary measures.

Using a unique disciplinary data set from the state of North Carolina, we examine within school variation in discipline and test our theoretical model of principal behavior. The data include information on every student infraction for the 2000-2001 school year, capturing information such as the type of behavior, the school's response, and demographic variables at both the student and school level. Two disciplinary outcomes are considered: whether a student receives an out of school suspension, and if so, the length of the sentence. The data allow us to determine which individual factors affect the school's choice of punishment and the relative importance of these factors.

Our findings suggest that aggregate disparities in punishment result from varying discipli-

nary levels across schools. Absent school fixed effects, simple regression analysis indicates that race plays an important role in school discipline. For example, a black high school student receives a suspension that is half a day longer on average. However, within school analysis suggests that race, along with most other observable individual characteristics, plays little role in the choice of punishment. The most important individual characteristics are the number of infractions the student has committed prior to the current offense and the level of student commitment to the school, evidenced by participation in extra-curricular activities or standardized tests.

Schools with high percentages of black, low income, and low scoring students tend to impose harsher penalties on offenders. We show that this behavior is consistent with achievement maximizing behavior by principals. Out of school suspensions increase total achievement by reducing student infractions through both a specific and general deterrent effect. Our results indicate that the specific deterrent effect, or the impact of punishment on the punished, is significant. The longer a student is suspended, the less likely a future offense becomes. However, certain types of offenders do not respond as well to punishment. At each level of discipline, black, low income, and low performing offenders are more likely to recidivate, suggesting this group receives less discipline. Counteracting the specific deterrent effect is the fact that black, low income, and low performing non-offenders respond very strongly to the threat of suspension. This general deterrent effect appears to explain why schools with large populations of students at risk for behavioral problems use harsher disciplinary practices.

The remainder of the paper is organized as follows. Section 2 provides further background information on school discipline and how the issues considered here relate to prior research. Section 3 provides a brief description of the data, and racial bias within schools is addressed in Section 4. Section 5 examines why variation in punishment exists across schools from both a theoretical and empirical perspective. Section 6 concludes.

2 Background

The consensus result in the education literature has been that black students receive harsher penalties even when controlling for socioeconomic status and infraction severity. Two examples of such research are McFadden et al. (1992) and Gregory (1995). McFadden et al. examine

the difference in student punishment once a referral has been made in a large school district in Florida. Their results indicate that although black students do not commit more severe offenses on average, they do receive harsher penalties. Using data aggregated at a national level, Gregory (1995) finds that black students, particularly males, are significantly more likely to receive corporal punishment. Variation in infraction severity is ruled out as a potential cause since the disparity in suspensions between white and black students is significantly smaller. Gregory (1995) hypothesizes that the disparity in corporal punishment is the result of discrimination against black males.

On the surface, disciplinary data from North Carolina appear to support the claims of racial bias in student punishment. Black students are suspended from school at a rate 17% greater than white students, and receive suspension lengths that are 7% longer on average. Conditioning on infraction type and other student characteristics only exacerbates these differences. Similar to the analysis above, these differences do not account for any variation across schools in disciplinary policy. Figure 1, which plots the average suspension length against the proportion of the student body that is black, illustrates why this may be a concern. At least unconditionally, it appears that schools serving large proportions of black students tend to suspend students for longer periods of time.

Schools serving high proportions of black students are also more likely to struggle academically and have larger percentages of students dropping out. Achievement scores in North Carolina from 2001 indicate black students test at grade level only 54% of the time as compared to 84% of white students. 7.2% of black students dropped out of high school in 2001, while only 5.3% of white students dropped out. There is little doubt these outcomes are correlated. In fact, suspensions may exacerbate the achievement and dropout gaps. Students who are out of school lose learning time, and are more likely not to return. Approximately 6% of high school dropouts in 2004 were either serving a suspension and did not return, or cited disciplinary problems as the primary reason for dropout.

Much effort has been spent examining the role schools play both in perpetuating and alleviating differences in achievement, drop out rates, and juvenile delinquency. At higher levels of education, affirmative action is a tool schools use to level the playing field between white and minority applicants. In an effort to minimize differences in the quality of public education between communities, states have tried to equate the amount of resources available

to each school district. The aim of such programs is to reduce the educational gaps that exist across race and socioeconomic groups. Despite these efforts, a distinct achievement gap still remains between black and white students (Fryer & Levitt, 2004; Ferguson, 1998). A number of theories have been espoused to explain this gap, including variation in school quality (Currie & Thomas, 1995), different environmental factors (Crane, 1994), and the discounting of academic achievement by black students (Cook & Ludwig, 1998). Most closely related to the work at hand is the suggestion that racial bias in teacher expectations might exacerbate the observed achievement gap (Ferguson, 1998). Using student names as instruments for race, Figlio (2004) demonstrates how negative teacher expectations can reduce test scores and decrease the likelihood of being labeled as gifted. In the case of student punishment, principal expectations about the true criminal behavior of the offender may vary by race. This may lead to observable differences in patterns of school discipline.

Economists have weighed in on the issue of racial differences in arenas outside of education. A particular focus recently has been on the use of racial profiling by police. Knowles et al. (2001) develop a test for racial discrimination based on the success rates of police searches. The research question they pose is very similar to the one considered here: Is the fact that police officers search black drivers more often indicative of racial bias or simply optimal behavior given varying search success rates by race? The results of the study suggest that police are not biased in their search behavior since success rates are equal across racial groups. Antonovics and Knight (2004) extend the basic statistical discrimination model to allow for heterogeneity in officer preferences. They hypothesize that if racial bias does not play a role in police searches, then the race of the police officer should not matter. The results indicate that police officers in Boston do discriminate by race when searching drivers.

How a criminal is apprehended, either through good detective work or racial profiling, is irrelevant for punishment if the individual is guilty. Once an offender is convicted, the appropriate sentence is determined by a judge or jury. There is nice parallel between the role of a judge in a criminal court and the role of a principal when disciplining a child. The principal must consider the offender's record and likelihood of repeating an offense, just as a judge might when sentencing a criminal. Although no micro-analysis has been completed on principal decision making, many studies have examined judge's decisions in criminal cases. Mustard (2001) examines over 75,000 federal offenders sentenced under the Sentencing Reform

Act of 1984. After controlling for a host of demographic and criminological variables, he finds that black offenders receive substantially longer sentences. Everett and Wojtkiewicz (2002) and Steffensmeier and Demuth (2000) also find that white defendants consistently receive less severe sentences then both black and hispanic defendants. A related study on bail settings by Ayres and Waldfogel (1994) finds that judges discriminate against black and male defendants by setting their bail at higher rates than similar white defendants. The authors use the pricing behavior of bail bondsmen as a market test of discrimination.

A judge's job is to enforce a set of laws in a fair and equitable manner, rather than writing the laws themselves, which is typically the job of elected representatives. Principals, to a certain extent, create and enforce the rules in a school. Behavior deemed acceptable in one school may be harshly punished in another. Consideration must be given to the impact these rules have on the behavior of the students and the output of the school. A number of researchers have examined how juveniles respond to the threat of punishment. Using state level data, Levitt (1998) shows how juvenile offenders change their criminal behavior upon arriving at adulthood. This change in behavior is a response to an increase in the severity of punishment when a offender ceases to be a minor. Mocan and Rees (1999) use micro data to show that juveniles do respond to economic incentives when making criminal decisions. Children who come from poor families and communities are more likely to commit crimes, but are less likely to do so if the perceived punishment is high. The idea that students respond rationally to expected punishment levels will be an integral part of our model of school discipline.

3 School Infraction Data

The state of North Carolina consists of approximately 120 school districts and over 2000 public schools. More than 1 million students were enrolled in the 2000-2001 school year. Beginning in the 2000-2001 academic year, the state of North Carolina instructed schools to record and catalog all student infractions resulting in an out of school suspension[1]. Information about the

[1]Many schools in the sample chose to record infractions not resulting in out of school suspensions. These schools allow us to examine the principal's decision to remove a student from school. When analyzing the decision to suspend a student, we focus on schools for which out of school suspensions constitute less than 75% of reportable offenses. We assume that for these schools, all infractions brought to the principal's attention are

type of infraction, student involved, and school response is included in the record. A limited set of demographic variables is also available, along with the school district the offending student is enrolled in.

The infraction data set for the 2000-2001 school year does not include the school of the offender. In order to examine punishment within schools, we must link individual offenders to testing data that captures the school the student attends, as well as further individual information. North Carolina requires all students in grades three through eight participate in end-of-grade exams. Thus all students committing infractions between the grades of three and eight should also exist in the testing data. Linking the two data sets requires a unique student identifier, typically the student's social security number. If this information is not captured accurately when a student commits an infraction we will not be able to link this student back to a school. The match rates for individual offenders are 87.4% and 84.6% for primary and middle schools respectively. The majority of the non-matches likely come from a handful of schools misreporting infraction data[2]. For example, it is unreasonable to believe that a large middle school reports no infractions, yet this is what we find for a number of schools after the match is complete. Offenders who do not match tend to commit slightly more serious offenses, but look similar across other observable characteristics.

Matching students at the secondary school level is slightly less successful. Beginning in 9th grade, testing is no longer mandatory, making it difficult to identify offenders in the testing data. Also, once students reach the age of sixteen they are able to dropout of school. Since suspensions are often a precursor to dropout and children who misbehave are likely to be struggling academically, it is no surprise that the match rate, 68.4%, is much lower for secondary schools. Offenders not matching at the high school level tend to commit more serious infractions, but again look similar along other observable attributes.

Summary statistics for the variables used in our primary specifications are listed in Table 1. Students are separated into offender and non-offender groups by race. We will focus on only two categories of race, white and black, since these groups commit over 93% of state-wide

being recorded.

[2]A number of schools report no infractions, automatically eliminating them from the sample. These tend to be primary schools, but there are middle and secondary schools falling in this category as well. Schools consistently leaving the disciplinary action field blank are dropped.

infractions. Despite making up only 26% of the non-offender group, 45.8% of offenders are black. Students with behavioral issues tend to be older than the average student and come from lower income and less educated households. Overall, the average offender commits more than two infractions per year, with a maximum of over 30 incidents. Black offenders on average commit the most infractions, 2.45 offenses per offender. 67.5% of all offenders are male. White students, both offenders and non-offenders, are less likely to be on free or reduced price lunch, have more educated parents, and tend to be held back less often then black students.

The two outcome variables when analyzing a principal's choice of punishment are whether a student is removed from school, and if so, for how long. Table 1 indicates that black students are removed from school 66.2% of the time, while only 53.7% of white students are removed. The average length of suspension appears very similar for white and black students, averaging 2.91 and 2.95 days respectively. Without any controls, the evidence suggests black students are suspended more often, but not for significantly greater lengths of time than white students.

When a student is written up, the incident is categorized as one of 26 offense types. The entire list of offenses along with the average suspension length is provided in Table 2. The majority of infractions fall into three categories: rule violations, disruptive behavior, and fighting. Severe violations likely to require assistance from law enforcement are dropped from the sample. Examples of such violations are homicide, distribution of a controlled substance, possession of a firearm, or rape. These offenses constitute less then .1% of all infractions. Figure 2 depicts the distribution of offenses by race, where the offenses have been aggregated into nine categories. White students are written up more often for truancy, drug offenses, and minor violations[3], while black students commit higher levels of fighting and disruptive behavior violations.

The core data set includes 75.7% of elementary schools, 85.6% of middle schools, and 91.1% of high schools[4]. Table 3 breaks down the infraction data by school type[5]. Primary

[3]Minor violations consist of rule violations, health immunization infractions, and infractions categorized as other.

[4]We only consider regular public schools. Charter, magnet, vocational, and special education schools are dropped from the sample.

[5]Primary Schools have a low grade between PK-3 and a high grade between PK-08. Middle schools have a low grade between 4-7 and a high grade between 4-9. Secondary schools have a low grade between 7-12 and a high grade of 12 only.

schools produce the lowest number of infractions, and face a different distribution of offenses as compared to both middle and secondary schools. In particular, truancy and drug offenses are much rarer in elementary schools. The categorization of offenses is also likely to vary by school level. What might be considered a fight in primary school may be be considered rowdiness at the middle school level. High school students are old enough to work and drop out of school, while most primary school kids are too young to be left home alone. These differences likely lead to varying punishment behavior across schooling level, as evidenced by the percent of students receiving suspension and the average length of suspension. In much of the analysis to follow, we examine each school type independently to remove any variation in punishment attributable to school level differences.

4 Testing for Racial Bias in School Discipline

Simple analysis of the disciplinary data suggests black students are more likely to receive an out of school suspension, and receive slightly longer suspensions when removed from school. To determine if racial bias is at the root of these disparities in school discipline, we examine punishment within schools. If race plays no role in the disciplinary decision within schools, offender race coefficients should not differ from zero.

Two primary models are considered. The first, a simple logit model, measures the likelihood that individual i attending school j receives an out of school suspension for committing infraction k.

$$\text{Out of School}_{ijk} = \beta_1 X_i + \beta_2 Y_i + s_k + s_j$$

X_i contains individual demographic variables such as race and age, while Y_i captures the behavioral history of the student during the school year. The severity of the infraction is captured by s_k, while the overall punitiveness of a school is reflected in s_j. Evidence of racial bias will be reflected in the coefficients on the individual race characteristics. The second model is identical to the first except for the outcome measure. In this case we measure the number of days a student is suspended for, conditional on receiving an out of school punishment. Both linear and truncated binomial regression models are fit to the data. Because the results are generally consistent between the two models, we only display results from the linear specification.

4.1 Removing a Student from School

Table 4 lists results for each school type, with and without school fixed effects, when the decision is whether or not to remove a student from school. The sample of offenders is limited to those schools reporting all infractions[6]. Although not listed in the table, controls for the type of infraction and timing of the infraction are included.

The black coefficient is positive and significant across all school types and samples when school fixed effects are not included in the model. The overall disparity in school discipline noted earlier does not dissipate once we control for offender history, socioeconomic variables, and offense type. For example, the model predicts that an average white male attending middle school who is caught fighting receives an out of school suspension with a probability of .61. A black student with the same observable characteristics caught fighting has a predicted probability of suspension of .70. Conditioning on infraction type, behavioral history, and other observable factors does not mitigate the suspension gap between black and white offenders.

The impact of race is eliminated at both the primary and middle school levels once school fixed effects are included in the regressions. The black coefficient is reduced significantly at each level, and actually changes sign for middle schools. For secondary schools, the black coefficient remains positive and statistically significant, however, its economic significance is limited. For example, the predicted probability of suspension for disruptive behavior for black and white students is .28 and .26 respectively. These results suggest that within schools principals do not consider race when deciding whether or not to remove a student from school.

A number of other individual controls are included in both the cross-school and within-school regressions to ensure that the race coefficients do not suffer from extensive omitted variable bias. The results suggest that male students, repeat offenders, and students who are old for their grade are consistently more likely to receive an out of school suspension. Offenders with relatively higher educated parents who also exhibit some commitment to the school (either by taking end of grade exams or participating in school activities) are less likely to be removed from school at higher grade levels.

Despite the statistical significance of many of the individual characteristics, the probability of receiving an out of school suspension is driven primarily by the type of misconduct or overall

[6]Schools who suspend students more than 75% of the time are not included. Results are robust to variations in this value.

school punitiveness. As one might expect, violent student offenses and offenses involving drug use or weapon possession lead to an out of school suspension in the majority of cases. For a relatively average middle school, the probability of receiving an out of school suspension rises from .12 to .62 to .97 as offense severity increases from truancy to fighting to weapon possession. The probability of receiving an out of school suspension varies as much across schools as it does across infraction categories. Using fighting as the base infraction type, the predicted probability of receiving an out of school suspension ranges from as low as .1 to above .9, with a standard deviation of .2. As a result of the cross-school variation in discipline and the non-uniform distribution of race across schools, black students tend to receive more out of school suspensions than do white students.

4.2 Length of Student Suspension

Results from Table 1 indicate that on average, black and white students are suspended from school for similar lengths of time. This result does not control for any other individual characteristics, such as behavioral history or infraction type. As Table 5 illustrates, inclusion of these individual characteristics exacerbates the difference in suspension length between white and black students across schools. At the middle and secondary school levels black students receive significantly longer sentences conditional on infraction type and other individual characteristics. The racial disparity appears particularly strong for secondary schools where the model predicts that on average a black student receives almost half a day longer for each suspension.

Examining suspension lengths within schools yields very different results. The black coefficient is insignificant for middle schools, and is drastically reduced for both middle and high schools. This suggests that the disparity in punishment is primarily a result of varying punishment levels across schools. There does appear to be some difference in treatment within high schools for black and white students. Racial bias is one possible explanation. Another plausible alternative is unobserved differences in behavioral history. We only observe infractions for the current year, but students may have a long history of behavioral trouble. Because black students are more likely to commit an offense, they are more likely to have an extended behavioral history. We will investigate this issue further in the next section.

Similar to the removal decision, the length of an individual sentence varies greatly with

the type of infraction and the school to which the student belongs. Increasing the severity of the infraction from truancy to fighting results in an extra day and a half of out of school suspension. A weapons possession infraction increases the predicted suspension length almost six days as compared to a truancy violation. The predicted penalty for fighting varies across schools from a minimum suspension of one day, to as long as two weeks. The only individual characteristic consistently affecting the disciplinary decision is behavioral history during the current academic year.

The results thus far indicate that racial bias within schools plays a small role in generating aggregate differences in disciplinary outcomes. Before moving on to cross school variation in punishment, we want to validate our results with some further analysis. The next few sections delve further into the examination of racial bias within schools to ensure that our results thus far are robust.

4.3 Robustness Checks

4.3.1 Distribution of Bias Across Schools

One criticism of the within-school model of racial bias is the restriction of equal bias across schools. With only one set of race coefficients, we are actually capturing the average bias within schools. A reasonable theory would suggest that racial bias differs within each school both in magnitude and direction. For example, one school may exhibit no differences in punishment across race, a second school may suspend black students at significantly higher rates, while a third school disciplines white students more severely. If there are enough schools falling into the last two categories, then we may find no racial disparities in punishment when significant bias exists in both directions.

To address this concern the within-schools model is extended by interacting each school fixed effect with a white and black dummy variable[7]. For the decision of whether to remove a student from school or not, the resulting estimating equation is given by

$$\text{Out of School}_{ijk} = \beta_1 * X_i + \beta_2 * Y_i + s_k + s_j * (WHITE) + s_j * (BLACK).$$

This allows each school to vary its overall level of punitiveness by race. Significant differences

[7]Only schools with more than 5 incidents for each racial group are considered.

in the white and black fixed effects for a particular school would suggest that a school is racially biased.

The differences $s_j * (BLACK) - s_j * (WHITE)$ have a straightforward interpretation in the length of suspension model. They measure in days the difference in the average suspension between white and black students within a particular school. For the model of the removal decision, the difference in the coefficients has no measurable character without considering the other parameters of the model. We can say that if this difference is positive, black students are more likely to receive an out of school suspension all else equal.

Generally, there are very few schools treating black and white students differently. Out of the 270 schools included in the suspension decision analysis, only 29 have significantly larger black fixed effects at the 10% level. This number is less dramatic when we consider that 21 schools have white fixed effects that were significantly larger at a 10% level. Thus, not only do less than 20% of schools exhibit any bias between whites and blacks in out of school suspensions, but the disparities that do exist are relatively equal. Length of student suspension varies even less across race within each school. 33 out of 641 schools have a black fixed effect that is significantly larger than the corresponding white fixed effect, while only 8 schools exhibit the opposite. In the case of suspension length, it appears that wherever punishment differences exist, black students tend to suffer. However, only 5% of the schools in our analysis exhibit this behavior.

Despite the limited number of schools exhibiting inconsistent punishment behavior, we might be concerned if these schools are similar along other observable characteristics. This would suggest that bias is not randomly distributed among schools and principals, but rather the effect of particular community attributes. To investigate this issue, a standard logit model is estimated at the school level where the dependent variable is an indicator for bias against black students. Bias in both the suspension and duration decisions is considered. School attributes such as the percentage of students on free or reduced price lunch and pupil-teacher ratios are included along with district-level characteristics, such as the number of violent offenses in the district last year. Results are listed in Table 6. Across both regressions, only one attribute, percentage of students on free or reduced price lunch, appears to significantly impact the likelihood that a school disciplines black students more severely than white students.

Allowing for heterogeneity across schools does not alter the basic conclusion that racial

bias alone cannot explain observed punishment disparities. Very few schools treat white and black students differently when choosing discipline. The small number of schools exhibiting bias against black students in length of suspension tend to have more affluent student bodies.

4.3.2 First Offenses

Racial differences in punishment for first time offenders may occur because principals have varying beliefs about true criminal tendencies. This would result in within and across school differences in discipline outcomes attributable to racial bias. If a student has committed previous infractions, a principal is better able to gauge the threat a particular student poses. This will reduce the effect of any inherent biases on the part of a principal. In a similar vein, principals anticipate how a student will respond to punishment. Varying beliefs across race about the benefits of punishment would also generate differences in discipline across race. Again, a history of behavioral problems for a particular student provides the principal better information about the expected response to punishment, minimizing the specific role of race.

Student behavior prior to 2000-2001 is unobservable. This might exacerbate any punishment disparities since black students are more likely to have a history of misbehavior. Although each school year starts with a clean slate, a student may have built a reputation for poor behavior among teachers and administrators within a school. Inconsistent punishment across racial groups for first offenses in this case would not indicate discrimination at all, rather bias in our coefficients caused by a correlation between race and some unobserved individual characteristics.

To examine these issues, we re-run models for the school removal and duration decisions using first offenses only. Table 7 lists results by school type and suspension margin. Each regression includes school fixed effects. Based on our earlier results, we would expect the black coefficient to be marginally significant at most. Compared to the results using all offenses, the magnitude of the black coefficient is larger in every case. Notice that the black coefficient remains positive and significant at a 10% level in primary and middle schools with regards to the removal decision, and in all school levels for the length of punishment decision. These results suggest that racial disparities in punishment are exaggerated for first offenses, however, it does not help us determine whether racial bias is a causal factor.

In order to determine if racial bias is the root cause of these punishment differences, we

need to minimize the bias created by unobserved student misbehavior. A student often creates a reputation for bad behavior within a school. When a student changes schools, this behavior is less likely to travel with the student, particularly if the offenses to date have been relatively minor. This idea leads us to examine first offenses for students who are new to a particular school. In this case, new is defined as a student who is enrolled in the lowest grade within a school. The majority of students falling into this category are in either 6th grade or 9th grade. Because a new student in most primary schools is in either KG or 1st grade, primary schools are not considered in this analysis.

Table 8 lists the results of estimating a logit specification for the suspension decision, and a linear model for length of suspension for all middle and secondary school students qualifying as new students. The black coefficient is no longer significant for any specification including school fixed effects and is significantly reduced in all cases. Matriculating to a new school appears to provide a fresh start for the majority of students. When a student in a new school misbehaves for the first time, race is not an integral factor when choosing discipline. These results strengthen our earlier findings suggesting no racial bias in the administration of punishment. Because black students are more likely to have committed an offense in the past, we would expect the black coefficient to be upward biased if we do not control for the complete behavioral history. Although our initial analysis suffers from this upward bias, race remains an insignificant factor in the punishment decision.

Aggregate data for the state of North Carolina during the 2000-2001 academic year indicate that black and white students receive different treatment. However, when we examine disciplinary decisions within schools, it appears that little to no racial difference in discipline persist. In order to explain the racial gap in disciplinary outcomes we need to examine why schools choose varying punishment strategies.

5 Achievement Maximizing Schools and the Discipline Decision

5.1 A Model of School Discipline

Rational choice theory was first applied to criminal behavior by Becker (1964). In this seminal article, Becker provides a framework by which individuals commit crimes if in expectation the benefits outweigh the costs. The cost to committing a crime includes the likelihood of being caught, the subsequent probability of being convicted, and the expected fine or length of incarceration if found guilty. This model has been used extensively to examine individual decision-making with respect to crime. Slightly less attention has been given to the other agent in Becker's model, the governing body. This entity is also assumed to act in a rational manner with regard to punishment and apprehension. Sentence severity and police expenditures are set to maximize the deterrent effect at least cost. Determining the optimal disciplinary level in a school is a basic extension of this framework.

The problem facing a principal is to choose a disciplinary level for the entire school that will maximize educational output at minimum cost[8]. A principal's choice of discipline in this framework should not be thought of as pertaining to any individual student. Rather the principal is setting the tone for all offenders in the school. The choice variable in the model is the length of suspension any offending student will receive[9]. Suspension is similar to imprisonment in that an offender is completely isolated. Just as a prisoner who is in jail cannot commit another crime, a student who is at home cannot disrupt the educational process. Being at home also implies that the offending student loses valuable classroom time.

Before addressing the discipline decision of the principal, we must first specify how suspension length impacts student achievement and behavior. Human capital production for student

[8]We do not focus on the safety of the students or school personnel since the vast majority of infractions do not put anyone in danger. As a result of the intense media coverage, school shootings and other violent incidents often appear to occur more than they actually do.

[9]Lazear(2001) models discipline as an instantaneous reaction to misbehavior. Students are reprimanded in class and do not lose learning time as a result. Thus the costs of discipline are administrative rather than academic.

i in school j takes a rather simple form,

$$\text{Ach}_{ij} = \text{F}(\alpha_i, d, k_j)$$

where α_i measures ability, d is the number of days a student is suspended, and k_j is the total number of incidents in school j[10]. Achievement is strictly increasing in α_i, and weakly decreasing in d and k_j. Longer suspensions negatively impact achievement through lost learning opportunities, and increased detachment from the learning process. Achievement is reduced by the overall disruptiveness in a school since productive learning time is lost dealing with troubled students.

Student actions are determined by the utility received from good and bad behavior. Once a decision has been made, a misbehaving student suffers the disciplinary consequences, and increases human capital according to the achievement equation. Students have perfect information on the level of discipline a principal will impose[11], and are caught with probability 1. For simplicity we assume there is only one type of misbehavior[12]. Utility for good behavior, $\text{U}_\text{g}(t_i)$, is a function of student type only, while utility from bad behavior, $\text{U}_\text{b}(d, t_i)$, is negatively impacted by the length of suspension in the school. The longer the suspension, the less attractive misbehavior becomes, $\frac{\partial \text{U}_\text{b}}{\partial d} < 0$. High types receive larger benefits from misbehavior and smaller benefits from good behavior, resulting in $\frac{\partial \text{U}_\text{b}}{\partial t} > 0$ and $\frac{\partial \text{U}_\text{g}}{\partial t} < 0$. A student will choose to misbehave if

$$\text{U}_\text{b}(d, t_i) + \epsilon_{\text{b}i} > \text{U}_\text{g}(t_i) + \epsilon_{\text{g}i}.$$

The probability student i misbehaves is simply a function of the disciplinary level at the school

[10] We can either think of j as a class or school identifier. For elementary schools it is appropriate to think of j as a classroom identifier since students spend the entire day in one class. Disruptions effect only the students within the classroom. For middle or secondary students j signifies a school level variable. Thus the total number of disruptions in a school impacts achievement even though it might not occur in all classes. Grogger(1997) shows how the overall level of violence in the community surrounding a school impacts achievement. Violence inside the school will likely have an even stronger effect.

[11] Pamphlets discussing unacceptable behavior and the potential consequences are typically handed out at the start of the school year.

[12] Relaxing this assumption does not change the results of the model. In the case of multiple infraction types, the principal would choose a discipline level for each infraction type. The length of suspension would vary with the disruptiveness of the infraction.

and student type, as illustrated by

$$\mathrm{p}_i = \mathrm{p}(d, t_i).$$

The assumptions about the returns to behavior for high type students suggests $\frac{\partial \mathrm{p}}{\partial t} > 0$, or that high type students are more likely to commit an offense. Because the returns to misbehavior are decreasing in suspension length, the likelihood of an offense also decreases in suspension length. Thus the rational response of a student body to harsh disciplinary measures is to reduce misbehavior. This simple conclusion has never been tested empirically and remains a controversial issue in the debate on school discipline. In the following empirical section we show that students do respond to the threat of punishment.

From the school's perspective, the use of discipline is necessary for two reasons: to align incentives, and to force students to internalize the public costs of misbehavior. While misbehavior yields a private gain for the offending student, learning costs are borne by the entire class or school. In order to maximize achievement, the individual costs of discipline must be raised to reflect the social costs of misbehavior. Principals use discipline judiciously however, since the more time a student spends out of school, the more individual achievement suffers. The principal's maximization problem is given by

$$\max_d \ \mathrm{Ach}_j = \sum_{i=1}^{n} \Big[\mathrm{p}_i \mathrm{F}(\alpha_i, d, k_j + 1) + (1 - \mathrm{p}_i)\mathrm{F}(\alpha_i, 0, k_j) - \mathrm{p}_i \mathrm{A}(d) \Big]$$

The term inside the summation is the expected achievement for student i, conditional on student i's individual characteristics and discipline, d, minus the expected administrative costs of punishment. Administrative costs, which are increasing in d, reflect the time spent dealing with irate parents. For illustrative purposes, assume each school contains one type of student, such that $t_i = t \ \forall \ i$[13]. Utilizing the fact that in equilibrium the expected number of total offenses in school j is simply the sum of the individual probabilities, the principal's problem can be written as

$$\max_d \ \mathrm{Ach}_j = \sum_{i=1}^{n} \Big[\mathrm{p}_i \mathrm{F}(\alpha_i, d, \sum_{\substack{j=1 \\ j \neq i}}^{n-1} \mathrm{p}_j + 1) + (1 - \mathrm{p}_i)\mathrm{F}(\alpha_i, 0, \sum_{\substack{j=1 \\ j \neq i}}^{n-1} \mathrm{p}_j) - \mathrm{p}_i \mathrm{A}(d) \Big].$$

The principal chooses d^* conditional on the type of student body, the distribution of abilities, and the administrative costs associated with discipline.

[13] Assuming multiple types within a school results in the same general conclusions.

With some slight re-arranging, the first order condition from the principal's maximization problem is given by

$$\frac{\partial \text{Ach}}{\partial d} = \sum_{i=1}^{n} \Bigg[\underbrace{\frac{\partial \text{p}_i}{\partial d}\Big(\text{F}(\alpha_i, d, \sum_{\substack{j=1 \\ j\neq i}}^{n-1} \text{p}_j + 1) - \text{F}(\alpha_i, 0, \sum_{\substack{j=1 \\ j\neq i}}^{n-1} \text{p}_j)\Big)}_{\text{specific deterrent effect}}$$

$$+ \underbrace{\text{p}_i \frac{\partial \text{F}}{\partial d}}_{\text{incapacitation cost}} + \underbrace{\frac{\partial \text{F}}{\partial k} \sum_{\substack{j=1 \\ j\neq i}}^{n-1} \frac{\partial \text{p}_j}{\partial d}}_{\text{general deterrent effect}} - \Big(\underbrace{\frac{\partial \text{p}_i}{\partial d}\text{A}(d) + \text{p}_i \frac{\partial \text{A}}{\partial d}}_{\text{administrative costs}} \Big) \Bigg] = 0.$$

The first term inside the summation is positive, and represents the deterrent effect of punishment for individual i. Increasing discipline reduces student i's probability of an offense, resulting in an achievement gain proportional to the difference in achievement at $d = d^*$ and $d = 0$. Achievement costs of discipline for student i are captured by the second term inside the summation. Longer suspensions reduce student i's achievement with probability p_i. The achievement costs of discipline can be thought of as the incapacitation costs of suspension. The third term inside the summation is positive, and reflects the general deterrent benefits of suspension for student i. High levels of discipline keep the other students in the school in line, reducing the total number of distractions experienced by student i. The last term captures the total administrative costs of punishment.

The principal's first order condition implicitly defines d^* as a function of t, the type of student in school j. By implicitly differentiating this equation, we can determine how the optimal suspension length varies with student type, as defined by the following equation,

$$\frac{\partial d^*}{\partial t} = -\frac{\frac{\partial(\frac{\partial \text{Ach}}{\partial d})}{\partial t}}{\frac{\partial^2 \text{Ach}}{\partial d^2}}.$$

The denominator in the above expression characterizes the slope of the school's achievement function. As suspension length increases, the academic benefits diminish while parental appeals and scorn from the local community rise at an increasing rate. As a result, the denominator in the above expression is negative, leaving the sign of $\frac{\partial d^*}{\partial t}$ to be determined by the

derivative of the principal's first order condition with respect to t.

$$\frac{\partial(\frac{\partial \mathrm{Ach}}{\partial d})}{\partial t} = \sum_{i=1}^{n} \left[\frac{\partial^2 \mathrm{p}_i}{\partial d \partial t} \left(\mathrm{F}(\alpha_i, d, \sum_{\substack{j=1 \\ j \neq i}}^{n-1} \mathrm{p}_j + 1) - \mathrm{F}(\alpha_i, 0, \sum_{\substack{j=1 \\ j \neq i}}^{n-1} \mathrm{p}_j) \right) \right.$$

$$\left. + \frac{\partial \mathrm{p}_i}{\partial t} \frac{\partial \mathrm{F}}{\partial d} + \frac{\partial \mathrm{F}}{\partial k} \sum_{\substack{j=1 \\ j \neq i}}^{n-1} \frac{\partial^2 \mathrm{p}_j}{\partial d \partial t} - \left(\frac{\partial^2 \mathrm{p}_i}{\partial d \partial t} \mathrm{A}(d) + \frac{\partial \mathrm{p}_i}{\partial t} \frac{\partial \mathrm{A}}{\partial d} \right) \right] = 0.$$

Only two terms in the above expression can be signed based on the assumptions already imposed. $\frac{\partial \mathrm{p}_i}{\partial t} \frac{\partial \mathrm{F}}{\partial d}$ and $-\frac{\partial \mathrm{p}_i}{\partial t} \frac{\partial \mathrm{A}}{\partial d}$ are both negative. These terms capture the increased achievement and administrative costs associated with an increase in the probability of an offense. If the response to discipline does not vary by type , $\frac{\partial^2 \mathrm{p}}{\partial d \partial t} = 0$, these costs result in less severe punishment for poorly behaved students. If we assume different types of students respond differently to punishment, then the sign of the above expression will depend on which type of student responds more favorably to punishment.

In his model of school discipline, Lazear(2001) assumes well behaved children, or lower types in our scenario, are more likely to respond to punishment. Under this assumption, Lazear shows that punishment should always be more severe for better behaved students. In our model, assuming lower types respond increasingly well to punishment, or $\frac{\partial^2 \mathrm{p}}{\partial d \partial t} < 0$, yields the same conclusion. Under this assumption $\frac{\partial d^*}{\partial t}$ is unambiguously negative. This theoretical solution, however, is not supported by data from North Carolina or previous education research. Past school discipline studies, as well as the evidence presented here, suggest black students not only commit more infractions, but receive harsher discipline. If the model were correct, it would imply that black students receive less discipline since they are prone to behavioral problems. If we assume $\frac{\partial^2 \mathrm{p}}{\partial d \partial t} > 0$, or that poorly behaved students respond more favorably to discipline, $\frac{\partial d^*}{\partial t}$ may be greater than zero if the deterrent effects are strong enough.

An alternative assumption is that the response to discipline is not strictly increasing or decreasing in type. Students who reap large rewards from misbehavior are likely to commit an offense regardless of the discipline threat. On the other hand, even a small amount of discipline may impose huge costs for some students. For these students a small change in punishment will not alter behavior since the probability they misbehave remains close to zero. This theory of behavior suggests that $\frac{\partial^2 \mathrm{p}}{\partial d \partial t} < 0$ for $t > t*$ and $\frac{\partial^2 \mathrm{p}}{\partial d \partial t} < 0$ otherwise. Figure 1 illustrates this concept graphically.

Every school can be summarized by one point on the line in Figure 1. This is a result of

assuming that schools serve only one type of student. In reality, schools serve students spanning the distribution of types illustrated in Figure 1. Incorporating this type of heterogeneity does not change the basic tenor of the results. As a simple example, consider a school serving two types of students. The principal's maximization problem remains the same except now total achievement is summed from two types of students. The optimal level of discipline will depend not only on the two types of students, but their proportions in the student population. Similar to the case of one student type, the optimal level of discipline varies according to the response by each type to discipline. If students likely to get in trouble respond to discipline more than well behaved students, then increasing the number of poorly behaved students will increase the optimal discipline level.

Without specifying precisely how the response to discipline varies by student type, we cannot reach a definitive conclusion about how punishment varies with student type. However, the model does suggest some intuitive results. Consider a school with extremely well behaved students. On the off chance a student does misbehave, the benefits of punishment appear to be limited. The specific deterrent effect may be large, but the general deterrent effects are likely to be small. At schools with significant behavioral problems, the specific deterrent effects may be relatively small, but the general deterrent effects may be significant. If the general deterrent effects are strong enough, punishment may be more severe in this environment. In order to make such judgements, a principal must have information on the distribution of types in the class.

Student type is unobservable, both in the case of a homogenous school and heterogenous school. However, principals have many sources from which to draw when making inferences about the student body as a whole and individually. A history of violence in the school and community may suggest what the overall mood will be within a school. Each individual student has a history of behavior, commitment to academic achievement, and commitment to the school. Principals aggregate this information to better understand what type of students the school serves. If race happens to be correlated with community or individual factors that put students at high risk for behavior problems, than punishment will differ significantly across racial groups. This does not result from pernicious bias on the part of school administrators, rather it reflects the principals optimal choice of punishment given the expected behavior of the student body.

5.2 Empirical Analysis of Student Responses to Discipline

The theoretical model of school discipline illustrates that the overall disciplinary level in a school will vary according to the type of student served. For most public schools in North Carolina the composition of the student body is determined by the types of families living in the local community[14]. Any variation in socioeconomic status or racial composition across communities will be reflected by variation in student composition across schools. The fact that school disciplinary procedures vary significantly is not surprising when we consider how different students look at particular schools.

Table 9 presents summary statistics at the school level for some key variables. In some primary and middle schools almost every child receives reduced price lunch, while in other schools almost no child receives aid. Because of inconsistencies in the free and reduced price lunch data at the high school level, we use median income in the area defined by the 5-digit zip code where the school is located to assess economic variation across schools. Although the variation in income is not as stark, it still suggests significant socioeconomic differences across schools. Results for each of the other variables under consideration are quite similar. Race, academic achievement, and the amount of serious violence vary significantly across schools in North Carolina. This is important since these variables are highly correlated with the probability of committing an offense. As Table 1 suggests, offenders are more likely to be low-income, black, and come from less educated families.

Variation among these key predictor variables suggests that certain schools will have higher at-risk student populations. Behavior differences across schools would be muted if variation in these predictor variables is random. However, as Table 10 indicates, schools with a high percentage of students on reduced price lunch are also more likely to be predominantly black, struggling academically, and located in a neighborhood with high levels of crime[15]. The confluence of these characteristics in one school indicates a student population at extreme risk for behavioral problems.

Variation in student behavior will still exist within schools. There will always be good students in what appears to be an at-risk school and vice versa. However, when setting the disciplinary tone, the principal must consider the overall behavioral risk of the school. This

[14]One of the largest districts in the state was still bussing students in the 2000-2001 school year.

[15]The correlation matrices for primary and secondary schools look very similar

type of optimization will result in varying punishment levels across schools for what appear to be similar offenses and offenders. Our analysis of within-school punishment reflected this strategy, with individual characteristics playing only a small role in the discipline decision. In order to better understand the punishment choice, we need to focus on the characteristics of the school and community instead of the individual.

5.2.1 Impact of School and Community Characteristics on Punishment

When examining the punishment decision in previous sections, we used school fixed effects to soak up any variation across schools in disciplinary levels. Replacing these fixed effects with school level variables allows us to determine which types of schools tend to use harsh disciplinary practices. Table 11 presents the results for each outcome measure and school type. Each regression includes the same individual and time controls present in Tables 4 and 5. Only the coefficients for the school variables are listed.

At each school level and student outcome, increasing the proportion of students in a school who are black increases the severity of the expected punishment. The percent black coefficient is positive in every case and significant at a 10% level in all but one. This is not surprising since our earlier cross-school results suggest that black students receive more punishment. The coefficients on the other variables in the regressions show mixed results, likely a result of the high correlation between percent black and the other controls. Conditional on other characteristics, the results do seem to suggest that higher socioeconomic status(SES) students receive harsher penalties. Similarly, increasing the education level in the community, all else equal, tends to result in more severe punishment. The number of violent offenses in the district during the previous year has a positive impact on sentence severity in all but one case.

Students who misbehave in at-risk schools, defined by the proportion on reduced price lunch and the proportion black, appear to receive the harshest punishment. In the context of our model, this suggests $\frac{\partial d}{\partial t} > 0$, or that punishment is increasing in type. To determine if this pattern of punishment is consistent with our model of profit maximizing behavior, we examine the deterrent and incapacitation effects of school discipline.

5.2.2 Specific Deterrence

Specific deterrence refers to the impact punishment, or the threat of punishment, has on the individual offender. In the context of school discipline, this is equivalent to the effect of suspension on the likelihood of recidivism. This is slightly different from the conventional definition since we are assuming the probability of apprehension is 1. Given the level of adult supervision in a school, this does not seem unreasonable. Our theoretical model assumes $\frac{\partial p}{\partial d} > 0$, or that the probability of misbehavior is reduced by the anticipated level of discipline. Greater levels of discipline may increase discipline at home, ridicule from classmates, or result in reduced school privileges, such as the opportunity to play sports. Students who misbehave and receive an out of school suspension will be better able to anticipate these costs and alter their behavior optimally.

Critics of out of school suspensions often claim that student behavior is not altered by discipline. No empirical evidence to date has been able to refute or support this argument. To determine if punishment has any impact on offending students, we examine how suspension length effects the probability of a future out of school suspension. The estimating equation is given by

$$\begin{aligned}\text{Offense}(n+1|n,i,j,k) = \quad & \beta_1 * (\text{days suspended for offense } n) \\ & + \beta_2 * (\text{days suspended for offense } n)^2 \\ & + \delta * (\%\text{ of students with at least } n \text{ offenses})_j \\ & + \gamma * X_i + s_k + week_i\end{aligned}$$

where $\text{Offense}(n+1|n,i,j,k)$ is 1 if offender i in school j commits offense $n+1$ conditional on having committed n offenses already, the last of which was type k. The right hand side includes controls for individual characteristics, X_i, overall behavior in school j, infraction type, s_k, and the week individual i will return to school. The coefficients of interest, β, measures the impact previous suspension length has on future behavior. By including controls for the week the student returns to school, β will measure only the deterrent effect of punishment. Without the week of return controls, β would capture the sum of the deterrent and incapacitation effects.

The analysis is limited to students who receive suspensions for twenty days or less, which

includes over 99% of all suspensions. A suspension for more than four weeks indicates either a severe incident or uncontrollable child. In either case, an intervention independent of school is likely, biasing any impact of the suspension itself. Students whose suspensions extend beyond the end of the school year are also dropped.

Each column of Table 12 lists the results of a logistic regression where the outcome variable equals 1 if the student commits offense $n + 1$ conditional on committing offense n. The coefficient on days suspended is negative and significant at each offense level[16]. The coefficient for the square of the days suspended is positive and significant in all cases but one. This indicates decreasing behavioral returns as the number of days a student is suspended increases. In fact, as the number of days suspended approaches the highest levels in the sample, an extra day of suspension results in slightly higher probabilities of recidivism. This likely results from the highly selected set of students receiving long term suspensions.

To put the results of the model in perspective, consider a black, 12 year old student suspended for disruptive behavior in a school where 20% of the students receive at least one out of school suspension. Assuming that the return date for the student is the same regardless of suspension length, the predicted probability of a second infraction decreases from .70 to .64 to .60 as the length of suspension increases from 1 to 5 to 10 days respectively. When we consider that a student suspended for 5 or 10 days will return to class one or two weeks later, the probability of a second infraction decreases from .70 to .62 to .57. The deterrent effect of punishment comprises almost 75% of the reduction in the probability of committing a second offense. Moving to a policy that requires mandatory 10 day suspensions for a first offense in a given year would reduce the number of 2nd infractions by more than 8000 across the state. This probably understates the effect since students would likely respond to the policy by not getting in trouble to start. The remaining coefficients correspond to our previous findings. Students who come from less educated families, and are black or low income pose greater risks of recidivism. Not shown in the table are the dummies capturing the week the student returns

[16]Numerous variations of the model were estimated to ensure the robustness of the basic results. In each case the results coincided with those listed in Table 12. Variations on the model included using the number of school days remaining after serving a suspension rather than week of return, limiting the sample to students below 9th grade, and limiting the sample to only those students taking end of grade exams. Splitting the sample into middle and primary school students also had no appreciable impact.

to class. Not surprisingly, the later a student returns to school the lower the probability of a next offense.

The results in Table 12 support our assumption that behavior responds to punishment in a positive way. Yet it does not provide insight into the question of why at-risk schools punish students more severely. To address this question we repeat the analysis from Table 12, allowing the deterrent effect to vary according to race and socioeconomic status. The coefficients of interest, listed in Table 13, indicate that black, low income, and academically struggling students are less responsive to punishment at each infraction level. Optimal school behavior would suggest that at risk students receive less discipline if they are less responsive. However, this is the exact opposite of what we observe in the data.

This result could be rationalized by considering the costs of punishment from the school's perspective. If the academic and administrative costs of punishment are negligible, then each school is simply aiming to minimize the number of infractions. As a result, at-risk students receive longer sentences in order to achieve the same level of behavior. On the other hand, if punishment is costly, particularly for high types, then punishment might be higher for at-risk students. Neither of these arguments is terribly convincing since punishment certainly has costs associated with it, and any differences in costs are likely to be small. Thus in order for the observed behavior to be consistent with our theoretical model, some other motivating factor for punishment must exist.

5.2.3 General Deterrence

One concern with the analysis of specific deterrence is the exclusion of well-behaved children. Punishment acts as a deterrent not only for the offending student, but also for the rest of the student body. This type of deterrence, known as general deterrence, is often used to justify harsh criminal sentences such as the death penalty or the three strikes laws in California. Although not all students observe every infraction, word spreads quickly around the school yard when a student is suspended. Thus even students who have not been in trouble form expectations about what the punishment will be if they misbehave.

Measuring the general deterrent effects with only one year of data is a challenge. Without data on the punitiveness of schools in previous years, no measure of a student's expectation of punishment is available. To side step this problem, we split the academic year into ap-

proximately two halves. Using disciplinary choices in the first half of the year, a measure of expected punishment is constructed. In practice, we use the median suspension length as the expected punishment. Although this does not control for other mitigating factors, such as severity of infraction or student history, this information is not likely known by the majority of the student body.

Expected punishment is then used as a regressor in an equation of student behavior. Only students who did not receive an out of school punishment in the first half of the year are considered. Offending students are more likely to respond to their own punishment than to the median punishment level in the school. The estimating equation is given by

$$\text{Offense in Second Half}_{ij} = \gamma * (\text{E}[d_j]) + \beta * X_i + \delta * (\text{P}_j) + \alpha * (\text{Other Offenders}_i).$$

The probability of committing any offense in the second half of the school year is a function of the expected discipline in school j, the probability of committing an infraction in school j in the first half of the year, individual characteristics, X_i, and the number of offenders a student is exposed to. The number of other offenders is defined at the classroom level in primary schools and at the grade level in middle school. Secondary school students are not included in the analysis since the students who do not match in high school typically get longer suspensions. Thus our measure of expected discipline would be significantly biased. The likelihood of committing an offense in the first half of the year is included to control for the overall behavior type in the school.

Results of the above model, presented in Table 14, indicate that the behavior of non-offending students is affected by the level of discipline in a school. The coefficient on average days suspended is negative and significant in primary and middle school. The predicted probability of committing any offense in the second half of the year for a fourth grader exposed to two other offenders decreases from .075 to .045 if the the median suspension level increases from one to five days. The predicted probability of an offense for a seventh grader is reduced from .14 to .07 for the same increase in median suspension levels. For a class of 250 seventh graders, this reduces the expected number of first time offenders during the second half of the school year from thirty-six to eighteen.

To test the validity of these results we substitute two other measures of expected punishment: average and maximum suspension lengths within a school. The maximum suspension

length during the first half of the school year does not significantly impact behavior. This suggests that students are able to discern extreme behavior from rather mundane school disruptions. Using the mean suspension level as a measure of expected discipline yields results similar to those in Table 14. However, the magnitude of the coefficient on expected discipline is significantly reduced. This is likely due to the increased noise present in the mean suspension level.

To determine if general deterrent effects can explain the difference in punishment between schools, we allow the impact of expected discipline to vary by student type. Table 15 lists the results using different proxies for type, including race, reduced price lunch, and academic achievement. In almost every regression, the group more likely to misbehave is also more likely to respond to expected discipline. The one case this is not true is for black students in elementary schools. This result coincides with the earlier results in Table 5 suggesting that elementary schools do not suspend black students for longer periods of time.

The fact that at risk students are more responsive to expected punishment is not surprising if we reconsider Figure 1. Figure 1 indicates the lowest type students will behave regardless of punishment. Schools containing majority low type students receive no general deterrent benefit when disciplining students. The marginal benefit from discipline arises only from the specific deterrent effect. Contrast this with a school that has a large population of at-risk students. The marginal benefit is the sum of the specific deterrent effect, which may be small, and the general deterrent effect for all the other students in the school. The marginal benefit to suspension in going to be much higher in schools with high populations of at-risk students.

Achievement maximizing behavior suggests stricter disciplinary procedures for schools with at-risk student bodies. The benefit to discipline is greater in such schools since punishment positively impacts both the offender and the other students in the school. This behavior results in aggregate racial disparities in punishment since black students often attend schools in high crime, low SES communities. While stricter punishment may negatively impact the offending student, it also has positive impacts for the students who are not getting in trouble. These students are less likely to misbehave and begin the downward spiral towards poor academic outcomes and future dropout.

6 Conclusion

The disproportionate impact of school discipline on black youth continues to be a controversial topic in North Carolina. In just the past few months, one school district attempted to ban corporal punishment because of the high number of black students being affected. Other districts have changed suspension policies significantly to reduce the amount of days minority students are out of class. The results presented here suggest that measures reducing punishment in schools with at risk students could lead to greater disruption in the future.

Using statewide disciplinary data from the state of North Carolina, we show that black students are more likely to receive out of school suspensions and receive longer sentences when punished. This results not from any pernicious bias on the part of schools, but is due to varying disciplinary policy across schools. Schools pursue differing disciplinary policy as a function of the type of student in the school.

Public schools serving majority black students also tend to serve high proportions of low income and low performing students. Each of these factors suggest a student body that is prone to high levels of misbehavior. Principals use harsher punishments in such schools to discourage students from misbehaving. The expected level of punishment facing a student is shown to have significant specific and general deterrent effects. At risk offenders, defined by race, socioeconomic status, and academic achievement, respond less positively to punishment than their peers. The same type of student who has yet to commit an infraction responds more positively to the threat of punishment. This pattern of student behavior suggests that schools with high proportions of at risk students should use more severe punishment.

Table 1: Summary Statistics by Race

	Non-Offenders			Offenders		
	All	White	Black	All	White	Black
Total	580,233	375,646	152,342	79,845	37,946	36,553
% Male	48.35	48.99	46.55	67.48	71.47	62.88
% Parents HS Grade	37.30	32.38	49.30	45.45	40.88	50.88
% Parents 2-yr coll grad	22.60	23.57	23.10	21.90	24.56	20.13
% Parents 4-yr coll grad	31.10	38.02	17.71	16.33	20.43	12.77
Age	12.87	12.99	12.7	13.34	13.67	13.13
% Reduced Price Lunch [a]	41.21	23.64	71.82	65.38	44.22	81.14
% Old for Grade	19.28	16.35	24.13	36.58	33.20	39.24
% Primary School	35.65	33.92	37.47	14.39	10.22	17.24
% Middle School	29.52	28.96	30.99	45.12	41.7	48.45
% Secondary School	34.83	37.12	31.54	40.49	48.09	34.30
Total Infractions				184,557	83,613	89,666
Avg Infractions per Offender				2.31	2.20	2.45
% Out of School Suspension				60.33	53.71	66.16
Avg Suspension Length in Days				2.96	2.91	2.95

[a]These numbers are calculated using only primary and middle school students. Secondary schools do not capture this information.

Table 2: Infraction Distribution

Infraction Category	Count	Proportion of Total	Average Suspension Length
Property Damage	1.045	0.57	2.8
Theft	2,980	1.62	2.81
Truancy	10,309	5.60	0.63
Undisciplined-Rowdy	41,824	22.73	1.34
Aggressive Behavior-Fighting	37,140	20.18	2.9
Substance Abuse	1,967	1.07	3.06
Health Immunization	171	0.09	2.27
Rule Violation	68,752	37.37	1.01
Assault with a Weapon	57	0.03	9.19
Assault Resulting in Serious Personal Injury	361	0.20	9.69
Assault on a School Official	419	0.23	7.24
Possession of a Controlled Substance	1,863	1.01	8.98
Distributing a Controlled Substance	317	0.17	10.19
Possession of a Weapon	1,243	0.68	7.93
Possession of a Potentially Harmful Object	360	0.20	3.51
Sexual Assault	124	0.07	4.98
Sexual Offense	736	0.40	3.51
Taking Indecent Liberties with a Minor	47	0.03	1.76
Serious Threat to Self or Others	230	0.13	5.64
Other	13,929	7.57	1.94

Table 3: School Characteristics

	Primary	Middle	Secondary
Total Schools	910	368	288
Average Size	520	667	1053
% Match	87.4	84.6	67.1
Total Offenders	14,796	37,136	37,468
Total Infractions	26,565	95,621	87,285
Average Infractions per Offender	1.80	2.57	2.33
% Schools Reporting Only OSS	69.60	40.18	42.73
% OSS	72.86	55.3	53.41
Average Suspension Length	2.12	3.45	3.51
Average Age	10.09	12.53	15.63
% White	37.87	43.27	53.11
% Black	55.15	49.50	40.15
% Violent	1.77	0.94	0.50
% Dangerous Weapon	2.64	1.24	1.04
% Sex Offense	0.78	0.76	0.34
% Drug Offense	1.09	1.91	4.94
% Fighting	41.74	33.60	15.93
% Property Related	4.04	2.90	2.14
% Rowdy	24.72	26.97	17.04
% Truancy	0.41	1.83	7.97
% Minor	22.81	29.84	50.09

Table 4: Predicting Out of School Suspensions

Logit model, dependent variable indicate an out of school suspension						
	Cross-School Variation			Within-School Variation		
	Primary	Middle	Secondary	Primary	Middle	Secondary
Male	0.1025	0.0556*	0.0763*	0.0098	0.0090	0.0446
Age	0.0164	0.0966*	-0.0150	0.0343	-0.0121	-0.1071*
Black	0.2124*	0.3794*	0.0491**	0.1191	-0.0709*	0.0917*
Total Infractions to Date	0.0279**	0.0690*	0.0534*	0.1273*	0.1176*	0.1078*
First Offense This Year	-0.4130*	-0.2477*	-0.3173*	-0.4735*	-0.3061*	-0.4566*
Reduced Price Lunch	0.1005	-0.1212*		0.0700	0.0696*	
Parents did not grad. HS	0.1300*	0.1090*	0.0540	0.1759*	0.0879*	0.0493
Parents grad. 2-yr coll.	0.1469	0.0559*	-0.0885*	-0.0183	-0.1447*	-0.0542**
Parents grad. 4-yr coll.	0.1317	-0.1667*	-0.2132*	0.1771	-0.1149*	-0.1680*
Old for grade, 1yr	0.1069	-0.1183*	0.1666*	0.1046	0.1190*	0.1277*
Old for grade, >1yr	0.2190	0.0334	0.3805*	0.2657**	0.2614*	0.3008*
Took EOG test last yr		-0.2203*			-0.3516*	
Age≥16			-0.0514			0.1728*
Participates in Sports			-0.1819*			-0.1507*
Participates in Other Activities			-0.0880*			-0.1008*
Sample Size	7,663	47,953	40,446	7,585	47,953	40,446
School FE	N	N	N	Y	Y	Y
Infraction Controls	Y	Y	Y	Y	Y	Y
Time Controls	Y	Y	Y	Y	Y	Y

Table 5: Suspension Length Conditional on Out of School Suspension

Linear model, number of days suspended is dependent variable						
	Cross-School Variation			Within-School Variation		
	Primary	Middle	Secondary	Primary	Middle	Secondary
Male	0.0537	-0.0402	0.1213	0.0443	-0.0436	0.0009
Age	0.1165*	0.1914*	-0.1090	0.1075*	0.1674*	0.0493**
Black	0.0730	0.2631*	0.4507*	0.1018	0.0891	0.2056*
Total Infractions to Date	0.0989*	0.0172	0.1576*	0.1529*	0.1081*	0.2492*
First Offense This Year	-0.1588*	-0.3095*	-0.1673	-0.1676*	-0.2713*	-0.3284*
Reduced Price Lunch	-0.0706	-0.0922		-0.0954	-0.1651*	
Parents did not grad. HS	0.0076	0.0478	-0.2940*	0.0328	0.0555	-0.2201**
Parents grad. 2-yr coll.	0.0412	-0.0463	-0.2408*	0.0308	-0.0630	-0.1987*
Parents grad. 4-yr coll.	0.1473	0.0807	-0.0908	0.3248*	-0.0383	-0.1258
Old for grade, 1yr	-0.0507	-0.0392	-0.0722	-0.0634	-0.0541	-0.1647
Old for grade, >1yr	-0.0082	-0.1143	0.1399	-0.1069	-0.1029	-0.0021
Took EOG test last yr		0.0734			0.1573*	
Age≥16			0.2801*			0.0302
Participates in Sports			-0.2323*			-0.1453**
Participates in Other Activities			-0.0039			0.0288
Sample Size	16,290	45,450	40,204	16,290	45,450	40,204
School FE	N	N	N	Y	Y	Y
Infraction Controls	Y	Y	Y	Y	Y	Y
Time Controls	Y	Y	Y	Y	Y	Y

Table 6: Predicting School Level Bias Using Race Specific School FE

	Bias in Out of School Suspension	Bias in Suspension Length
% Reduced Price Lunch	0.9529	-3.393*
% Black	-0.8935	0.6767
Pupil-Teacher Ratio	0.0563	-0.0626
Mid-Size City	-0.4169	0.7549
Fringe of Mid-Size City		-0.3061
Small Town	0.7399	-0.8138
Rural, Outside MSA	-0.1700	0.4647
Rural, Inside MSA	0.04405	0.0429
Violent Offense Rate in 2000	0.0794	-0.0676
School Personnel Average Experience Level	0.0401	-0.0023
Constant	-4.4645*	-0.5709
Total Schools	261	611

Table 7: Discipline for First Offenses in Current Academic Year

	Out of School Suspension- Logit Model			Suspension Length - Linear Model		
	Primary	Middle	Secondary	Primary	Middle	Secondary
Male	-0.0339	0.0543	-0.0076	0.1273**	-0.0417	-0.0202
Age	-0.0308	0.1055*	-0.0184	0.1262*	0.2018*	0.1075*
Black	0.5198*	-0.0661	0.1318*	0.1531**	0.2479*	0.1939*
Reduced Price Lunch	-0.0877	0.1445*		-0.0666	-0.2532*	
Parents did not grad. HS	0.04030*	0.0545	0.0583	-0.0118	-0.0817	-0.1295
Parents grad. 2-yr coll.	-0.1128	-0.1673*	-0.1371*	-0.1235	-0.1771**	-0.0900
Parents grad. 4-yr coll.	0.2289	-0.1494*	-0.2958*	0.1795	-0.1456	-0.1695
Old for grade, 1yr	-0.0065	0.0167	0.1031**	-0.1553*	-0.1087	-0.1453
Old for grade, >1yr	-0.0065	0.0684	0.3033*	-0.1512	-0.2271	0.3018**
Took EOG test last yr		-0.4242*			-0.0578	
Age≥16			0.0149			-0.2355*
Participates in Sports			-0.0567			-0.1325
Participates in Other Activities			0.0057			0.0618
Sample Size	3,130	15,703	15,259	8,873	18,384	18,443
School FE	Y	Y	Y	Y	Y	Y
Infraction Controls	Y	Y	Y	Y	Y	Y
Time Controls	Y	Y	Y	Y	Y	Y

Table 8: Discipline for First Offenses in a New School

	Out of School Suspension- Logit Model		Suspension Length - Linear Model	
	Middle	Secondary	Middle	Secondary
Male	0.0910	-0.0810	0.0761	-0.1237
Age	0.0793	-0.0464	0.3799*	0.2717*
Black	-0.10623	0.0536	0.0629	0.1751
Reduced Price Lunch	0.3030*		0.0112	
Parents did not grad. HS	-0.0409	0.0520	0.0057	-0.2996
Parents grad. 2-yr coll.	-0.1072	-0.1281	-0.0557	0.0970
Parents grad. 4-yr coll.	-0.2483	-0.3261*	-0.1683	0.0258
Old for grade, 1yr	0.02327	0.1202	-0.2658*	-0.3903*
Old for grade, >1yr	0.2366	0.4444**	-0.7580*	-0.2891
Took EOG test last yr	-0.4529*		0.1867*	
Age≥16		0.0634		0.1500
Participates in Sports		-0.0405		-0.2932*
Participates in Other Activities		-0.0248		-0.0579
Sample Size	4,876	5,655	6,082	7,963
School FE	Y	Y	Y	Y
Infraction Controls	Y	Y	Y	Y
Time Controls	Y	Y	Y	Y

Table 9: Variation in Student Population Across Schools

	Mean	Standard Dev.	Minimum	Maximum
Primary - 855 Schools				
Proportion Black	0.310	0.245	0.00	0.998
Proportion Reduced Price Lunch	0.501	0.224	0.008	0.996
EOG Composite Score	74.64	9.70	39.10	94.90
Violent Offense Rate	5.82	2.05	1.11	16.54
Middle - 338 Schools				
Proportion Black	0.328	0.234	0.000	0.989
Proportion Reduced Price Lunch	0.436	0.198	0.018	0.995
EOG Composite Score	76.57	9.12	44.10	93.90
Violent Offense Rate	5.80	2.13	1.11	16.54
Secondary - 278 Schools				
Proportion Black	0.273	0.210	0.000	0.987
Avg. Median Income [a]	37,622	9,563	22,204	76,795
EOG Composite Score	61.30	11.03	23.70	86.10
Violent Offense Rate	5.84	2.16	1.11	16.54

[a]Income data is from the 2000 U.S. census, aggregated at the 5-digit zip code area where the school is located.

Table 10: Correlation in Student Population Variables for Middle Schools

	% Red. Lunch	% Black	EOG Score	Violent Inf. Rate	Arrest Rate	Avg. Experience	Per-P
% Red. Lunch	1						
% Black	0.695	1					
EOG Score	-0.722	-0.671	1				
Violent Inf. Rate	0.222	0.085	-0.247	1			
Arrest Rate	0.212	0.219	-0.174	0.083	1		
Avg. Experience	-0.028	-0.172	0.191	0.111	0.037	1	
Per-Pupil Exp.	0.202	0.322	-0.144	0.024	-0.036	0.024	

Table 11: Impact of School and Community Characteristics on Punishment

	Out of School Suspension- MFX from Logit Model			Suspension Length - Linear M		
	Primary	Middle	Secondary	Primary	Middle	Second
% Reduced Lunch	-0.4637*	-0.0468	-0.0629	1.1549*	1.0215*	-0.00
% Black	0.3594*	0.4644*	0.1583*	0.4009*	0.1183	0.762
Violent Offense Rate in 2000	0.0292*	0.0462*	0.0031*	0.0171	-0.0548*	0.014
Average Experience Level in School	0.0032	-0.0165*	-0.0191*	0.0052	-0.0032	0.028
Pupil-Teacher Ratio	-0.0150*	0.0135*	0.0005	0.0394*	0.0457*	-0.027
EOG/EOC Composite Score	0.0003	0.0054*	-0.0053*	0.0106*	-0.0117	0.028
% BS Degree	0.2104*	0.3918*	-0.2763*	-0.0918	1.9942*	-0.30
County Arrest Rate	1.4270*	0.4579*	-0.2370*	1.8181*	-3.6238*	1.07
Sample Size	6,960	43,617	36,689	16,502	41,675	37,5
Individual Controls	Y	Y	Y	Y	Y	Y
Infraction Controls	Y	Y	Y	Y	Y	Y
Time Controls	Y	Y	Y	Y	Y	Y

Table 12: Behavioral Effects of Suspension - Specific Deterrence

	P(n=2 | n=1)	P(n=3 | n=2)	P(n=4 | n=3)	P(n=5 | n=4)
Days Suspended	-0.0883*	-0.1038*	-0.0940*	-0.1608*
Days Suspended^2	0.0038*	0.0059*	0.0033	0.0105*
% of School Committing n infractions	0.3077*	1.1796*	1.5712*	2.4968*
Female	-0.4162*	-0.3231*	-0.1805*	-0.2772*
Black	0.3229*	0.1639*	0.0812	0.1776*
Reduced Price Lunch	0.2142*	0.1295*	0.1126**	0.1090
Parents did not grad. HS	0.2641*	0.0904*	-0.0075	0.1587**
Parents grad. 2-yr coll.	-0.1858*	-0.2007*	-0.1447*	-0.0985
Parents grad. 4-yr coll.	-0.3425*	-0.2945*	-0.0416	0.0636
Old for grade	0.1443*	0.0701	0.1631*	0.0318
Age	0.1177*	0.0712*	0.0632*	0.0879*
Sample Size	46,365	18,920	9,932	5,627
Infraction Controls	Y	Y	Y	Y
Week of Return Controls	Y	Y	Y	Y

Table 13: Behavioral Effects of Suspension - Specific Deterrence by Student Type

	P(n=2 | n=1)	P(n=3 | n=2)	P(n=4 | n=3)	P(n=5 | n=4)
Days Susp. x Black	-0.0993*	-0.0989*	-0.0365	-0.1392*
Days Susp. x White	-0.0831*	-0.1223*	-0.1833*	-0.2319*
Days Susp. x Reduced	-0.0852*	-0.1010*	-0.0632	-0.1648*
Days Susp. x No Reduced	-0.1034*	-0.1149*	-0.1253*	-0.2110*
Days Susp. x Below Avg. Student	-0.0454	-0.1049*	-0.0444	-0.0526
Days Susp. x Avg. Student	-0.1113*	-0.1213*	-0.0826	-0.2153*
Days Susp. x Above Avg. Student	-0.1175*	-0.0986	-0.1887	-0.3373*
Infraction Controls	Y	Y	Y	Y
Week of Return Controls	Y	Y	Y	Y
Individual Char. Controls	Y	Y	Y	Y

Table 14: Behavioral Effects of Suspension - General Deterrence

	Primary	Middle
Median Suspension Length	-0.1327*	-0.1882*
Other Offenders	0.1445*	0.0087*
Proportion of Students Offending in 1st half	11.0998*	4.6894*
Female	-1.2408*	-0.7705*
Black	0.7873*	0.5867*
Reduced Price Lunch	0.5166*	0.4320*
Parents did not grad. HS	0.4215*	0.4572*
Parents grad. 2-yr coll.	-0.2642*	-0.2036*
Parents grad. 4-yr coll.	-0.7454*	-0.6674*
Old for grade, 1yr	0.3638*	0.4511*
Old for grade, ¿1yr	0.5114*	0.7262*
Constant	-4.0172*	-2.8394*
Sample Size	182,275	196,729
Grade Controls	Y	Y

Table 15: Behavioral Effects of Suspension - General Deterrence by Student Type

	Primary	Middle
Median Suspension Length x Black	-0.1229*	-0.2366*
Median Suspension Length x White	-0.1481*	-0.1398*
Median Suspension Length x Reduced	-0.1626*	-0.2056*
Median Suspension Length x No Reduced	-0.0686*	-0.1650*
Median Suspension Length x Below Avg. Student	-0.1167*	-0.2165*
Median Suspension Length x Avg. Student	-0.1163*	-0.1838*
Median Suspension Length x Above Avg. Student	-0.0466	-0.1661*
Individual Char. Controls	Y	Y
Grade Controls	Y	Y

Figure 1:

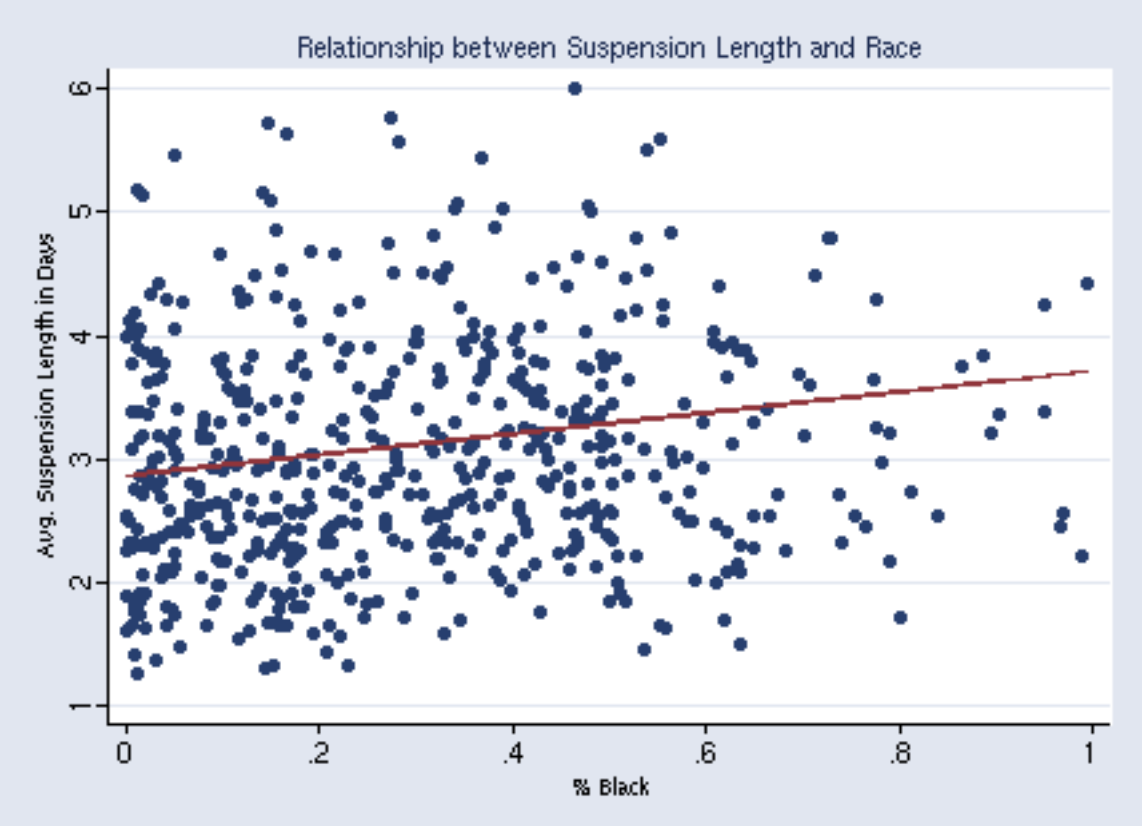

Figure 2:

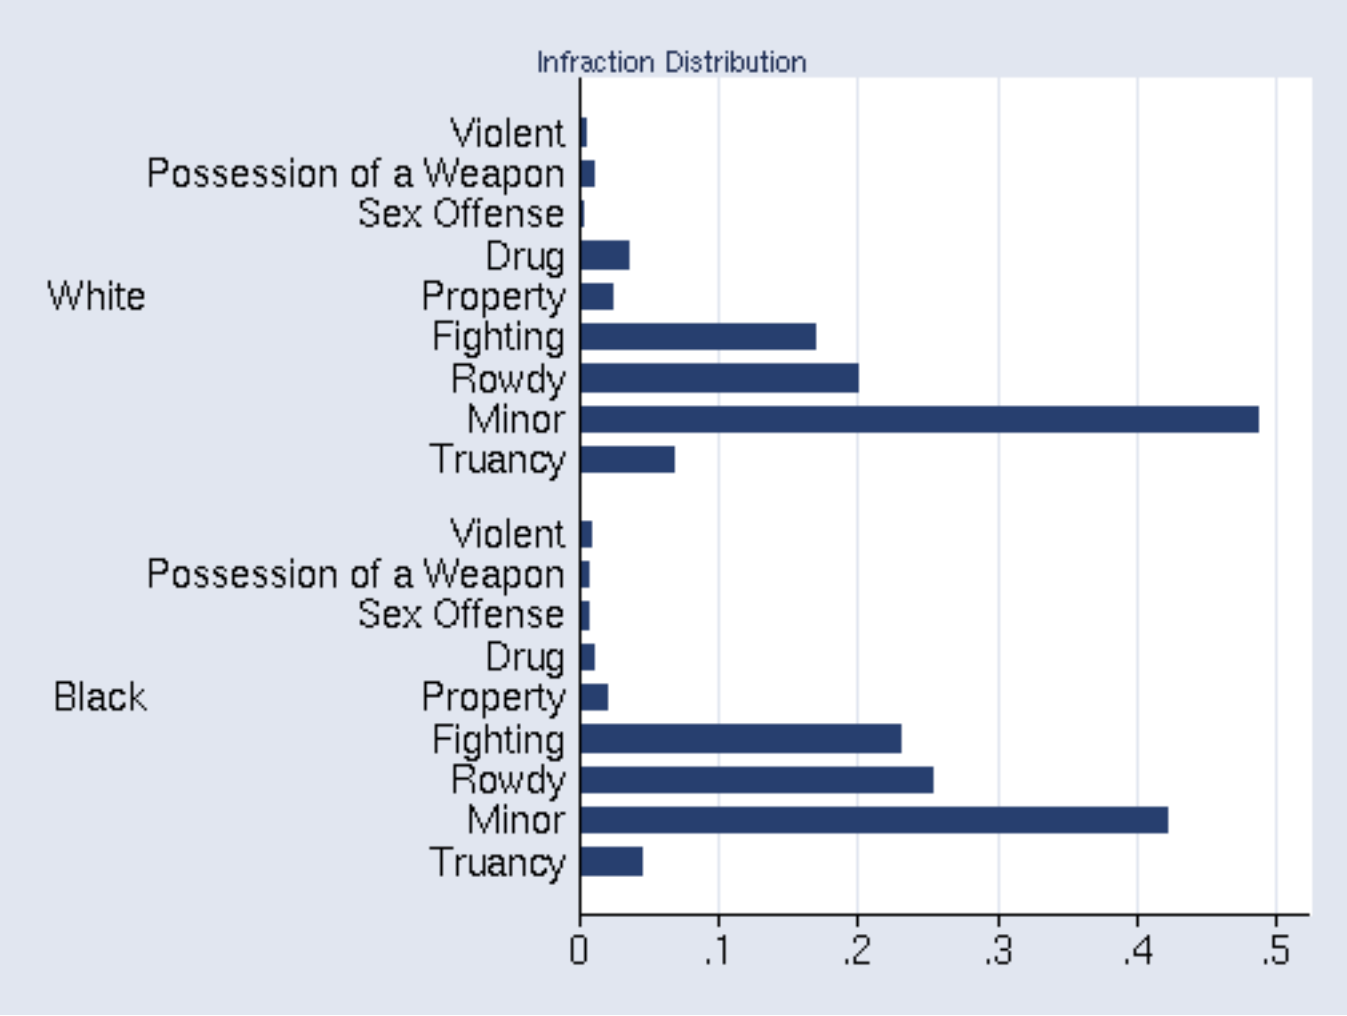
Infraction Distribution
White
Violent
Possession of a Weapon
Sex Offense
Drug
Property
Fighting
Rowdy
Minor
Truancy
Black
Violent
Possession of a Weapon
Sex Offense
Drug
Property
Fighting
Rowdy
Minor
Truancy
0
.1
.2
.3
.4
.5

Figure 3: