

Familiarity Does *Not* Breed Contempt: Generosity, Discrimination and Diversity in Delhi Schools

Gautam Rao*

October 16, 2014

Abstract

I exploit a natural experiment in India to identify how economic diversity in schools affects wealthy students. A sudden policy change in 2007 forced many private schools in Delhi to meet a quota of poor children in admissions. This led to a sharp increase in the presence of poor children in new cohorts in those schools, but not in older cohorts or in other schools. Using this variation, I study impacts on three broad classes of outcomes: (i) prosocial behavior; (ii) social interactions and discrimination, and (iii) academic outcomes. My first finding is that having poor classmates makes wealthy students more prosocial and generous. They are more likely to volunteer for a charity at school, more generous towards both rich and poor students in dictator games, and exhibit more egalitarian preferences. Second, wealthy students discriminate less against poor children, as measured in a team-selection field experiment. Closely related, having poor classmates increases their willingness to socially interact with poor children outside school. I additionally exploit idiosyncratic assignment to study groups within classrooms to show that these effects on social behaviors are largely driven by personal interactions between wealthy and poor students, rather than by changes in teacher behavior or curriculum. In contrast to the strong impacts on social outcomes, I find that poor students have mixed effects on the academic achievement of their wealthy classmates. In particular, I find marginally significant negative effects on test scores in English, but no effects on Hindi, or Math. I also detect no peer effects on disruptive behavior in the classroom, but do find substantial increases in the use of inappropriate language at school.

*Microsoft Research and Department of Economics, Harvard University. E-mail: grao@fas.harvard.edu. The latest version of this paper is available at <https://sites.google.com/site/graoeconomics/>. I thank my extraordinary advisors: Stefano DellaVigna, Edward Miguel, Matthew Rabin, Frederico Finan and Ernesto Dal Bo for their encouragement, counsel and patience over the course of this project. This paper also benefited from helpful comments from Lorenzo Casaburi, Jacqueline Doremus, Greg Duncan, Willa Friedman, Matt Gentzkow, Paul Gertler, Jonas Hjort, Simon Jaeger, David Laibson, Steven Levitt, Ulrike Malmendier, Jesse Shapiro, Richard Thaler, Betty Sadoulet, and many others, including audiences at UC Berkeley, U Chicago, Stanford, PACDEV and the National Academy of Education. Tarunima Sen and Dheeraj Gupta provided excellent research assistance. Funding for this project was generously provided by the Spencer Foundation, the National Academy of Education, the Center for Equitable Growth, the Program in Psychological Economics at UC Berkeley, the Levin Family Fellowship, the Center for Evaluation of Global Action (CEGA) and the UC Berkeley Summer Research Grant.

1 Introduction

Schools are de facto segregated across social and economic lines in many countries. Much research has examined the effects of such segregation on learning outcomes.¹ But desegregation and affirmative action efforts have historically been motivated not only by equity concerns, but also by the argument that diversity in schools benefits society by positively influencing inter-group attitudes and social behavior (Schofield 1996). Yet empirical evidence on such effects is exceedingly scarce. More generally, little is known about how social preferences and behaviors are shaped, and whether they can be influenced by policy. This question is of particular interest in diverse and polarized societies, where the costs of social divisions are well documented.²

I focus on a particular dimension of diversity - economic status - and seek to answer the following question: What effect do peers from poor households have on students from relatively wealthy families? I assemble a data set of about two thousand students in fourteen schools in Delhi, India and use a combination of field and lab experiments, tests of learning and administrative data to measure the following outcomes: (i) generosity and prosocial behavior; (ii) tastes for socially interacting with or discriminating against the poor; and (iii) learning and classroom behavior.

My first econometric strategy exploits the plausibly exogenous staggered timing of a policy change that required elite private schools to offer places to poor students. This causes a sharp discontinuity across cohorts in the presence of poor students. In most schools, cohorts beginning schooling in 2007 or later have many poor students, while older cohorts are comprised exclusively of rich students. However, a small control group (about 4%) of elite private schools are entirely exempt from the policy for historical reasons, while another handful (6%) of schools complied with the policy a year late - in 2008 instead of 2007. I can therefore identify the effect of the presence of poor students (the “treatment”)

¹Buchmann and Hannum (2001) and Karsten (2010) report measures of educational segregation or stratification in a number of countries. Hattie (2002) and Van Ewijk and Slegers (2009) provide meta-analyses of the effects of segregation on inequalities in learning.

²A large empirical and theoretical literature links greater social diversity, inequality and polarization with conflict (Esteban and Ray 2011), collective action problems (Bardhan et al. 2007), low levels of public good provision (Miguel and Gugerty 2005), political instability (Alesina and Perotti 1996) and diminished economic growth (Easterly and Levine 1997).

by comparing both within schools (comparing treated and untreated cohorts) and within cohorts (comparing treated and untreated schools) using a difference-in-differences regression model. This approach identifies the average effect on wealthy students of adding poor children to their classroom - an important estimate for policy.

The second econometric strategy isolates the role of personal interactions between rich and poor students by exploiting idiosyncratic variation in peer groups *within* the classroom. Some schools in the sample use alphabetic order of first name to assign students to group-work and study partners. In these schools, the number of poor children with names similar to a given rich student provides plausibly exogenous variation in personal interactions with a poor student. This allows me to distinguish between changes occurring due to personal interactions between students, and the effects of other possible changes at the classroom level, say in teacher behavior or curriculum.

My first finding is that having poor classmates makes students more prosocial, as measured by their history of volunteering for charitable causes at school. The schools in my sample occasionally offer students opportunities to volunteer or fundraise for affiliated charities. One such activity involves attending school on two weekend afternoons to help fundraise for a charity for disadvantaged children. I collect attendance records from such events and find that having poor classmates increases the share of volunteers by 10 percentage points (se 2.5%) on a base of 24%, while having a poor study partner increases volunteering by an imprecisely estimated 13 percentage points (se 9%).

To complement the field measure of prosocial behavior, I invite students to participate in a set of dictator games in the lab. Their incentivized choices in the games show that having poor classmates makes them more generous and egalitarian. Treatment students share 45% (se 7%) or about 0.45 standard deviations more than control students when offered the chance to share money with an anonymous poor recipient at another school. But importantly, they are also 27% (se 5%) more generous when paired with other *rich* students. These effects are driven largely by increases in the share of students choosing a 50-50 split of the endowment. Consistent with this, I find increases in separate experimental measures of egalitarian preferences. Thus, exposure to poor students does not just make

students more charitable towards the poor. Instead, it affects generosity and notions of fairness more generally.

The second finding is that economically diverse classrooms cause wealthy students to discriminate less against other poor children outside school. I measure discrimination using a field experiment in which participants select teammates for a relay race. By having participants choose between more-athletic poor students and less-athletic rich students, I create a tradeoff between ability and social similarity. Ability was revealed in a first stage using individual sprints, and the reward offered for winning the relay race was randomly varied across students. This provides exogenous variation in the price of discrimination. I find that when the stakes are high - Rs. 500 (\$10), about a month's pocket money for the older students - only 6% of wealthy students discriminate by choosing a slower rich student over a faster poor student. As the stakes decrease, however, I observe much more discrimination. In the lowest stakes condition (Rs. 50), almost a third of students discriminated against the poor. But past exposure to poor students reduces discrimination by 12 percentage points. I structurally estimate a simple model of taste-based discrimination and find that wealthy students dislike socially interacting with a poor teammate relative to rich one by an average of Rs. 34, about two days worth of pocket money. Having poor classmates reduces this distaste by 30%.³

To shed light on the observed reduction in discrimination, I conduct a separate experiment to directly measure tastes for social interactions. Preferences for interacting with people from other social groups provide a natural foundation for taste-based discrimination. To measure such preferences, I invite students to attend a play date at a school for poor students, and elicit incentivized measures of their willingness to accept. I find that having poor classmates makes students more willing to attend the play dates with poor children. In particular, it reduced the average size of the incentive they required to attend the play date by 19% (se 3%). Having a poor study partner affects "willingness to play" by a similar

³The observed behavior is inconsistent with a simple model of statistical discrimination. When students are asked which of the prospective teammates is most likely to win the relay race, they invariably (98%) select the fastest student in the ability revelation round, similarly across treated and control classrooms. This implies that a substantial number of students chose a rich teammate despite believing that he is *less* likely to help them win the race. This suggests that taste-based discrimination dominates in this setting.

amount.

Having established the effects of having poor classmates on social preferences and behaviors, I turn attention to impacts on learning and classroom discipline. A traditional concern with integrating disadvantaged students into elite schools is the potential for negative peer effects on academic achievement. To evaluate this concern, I conduct tests of learning in English, Hindi and Math, and collect teacher reports on classroom behavior. I detect marginally significant but meaningful decreases in wealthy students' English language scores, but find no effects on Hindi or Math scores, or on a combined index over all subjects. This pattern of findings is consistent with the measured achievement gap between poor and wealthy students, which is largest in English – perhaps because wealthy students are more likely to speak English at home. And while teachers do report higher rates of disciplinary infractions by wealthy students in treated classrooms, the increase comprises entirely of the use of inappropriate language (that is, swearing), as opposed to disruptive or violent behavior. My third finding is thus of mixed but arguably modest effects on academic achievement and discipline.

For each of the outcomes above, I compare the effects of the two types of variation: across-classroom variation in the presence of poor students, and within-classroom variation driven by assignment to study groups. This sheds light on mechanisms underlying the results by teasing apart the effect of direct personal interactions from the impact of other changes such as those in teacher behavior or curriculum. I find that personal interactions are an important driver of the overall effects. For example, having a poor study partner alone explains 70% of the increase in “willingness to play” with a poor child, and 38% of the increase in generosity towards the poor. This likely underestimates the importance of personal interactions, since students surely also interact with other poor classmates outside their study groups.

This paper relates to four bodies of work in economics. First, a recent literature studies whether interaction reduces inter-group prejudice. Most closely related are Boisjoly et al. (2007) and Burns et al. (2013), who find that being randomly assigned a roommate of a

different race at college increases inter-racial social interactions in later years.⁴ Second, this paper relates to research on the effects of desegregation and (more generally) peer effects in education. Evidence on peer effects in learning is mixed, with impacts on non-academic outcomes such as church attendance and drug and alcohol use a more robust finding (see Sacerdote 2011 for a review). Consistent with this, I find substantial effects on prosocial behavior and discrimination, but mixed and overall modest effects on test scores. A third connection is to the growing literature on how distributional social preferences are shaped, for example by exposure to violent conflict (Voors et al. 2012) and the ideology of one's college professors (Fisman et al 2009). I add to this literature by showing that peers at school can also shape social preferences. Finally, this paper relates to research on the economics of discrimination. I contribute to this literature by showing evidence of taste-based discrimination in an experiment (albeit in a non-market setting amongst students), and more importantly by showing that past exposure to out-group members causally reduces such discrimination.⁵

My findings are also of relevance to policy makers: the policy I study will shortly be extended throughout India under the Right To Education Act, with consequences for over 300 million school-age children. This policy is controversial, with legal battles over its legitimacy reaching India's Supreme Court. Opponents have prominently argued that any gain for poor children will come at a substantial cost to the existing clientele of private schools. Proponents have responded that diversity will benefit even rich students by providing them with "a clearer idea of the world".⁶ While we must be cautious in extrapolating from elite private schools in Delhi to the rest of India, my findings provide some support for each side of the debate. A radical increase in diversity in the classroom did have modest negative impacts on the academic achievement and behavior of advantaged students. But it also

⁴Outside of economics, a long tradition of related work in social psychology following Allport (1954) generally documents a negative correlation between inter-group contact and prejudice, but suffers from both selection problems and a reliance on stated attitudes rather than observed behavior as outcome measures (Pettigrew and Tropp 2006).

⁵Scholars have investigated the extent to which discrimination exists and matters in the labor market using audit studies (Bertrand and Mullainathan 2004), quasi-experiments (Goldin and Rouse 2000) and experiments (List 2004, Mobius and Rosenblat 2006).

⁶The Indian Express, April 13 2013. <http://www.indianexpress.com/news/learning-curve/936084/>

made them substantially more generous and prosocial, more willing to socially interact with poor children, and less likely to discriminate against them. A full accounting of the effects of economic diversity in schools on privileged students should consider all these effects.

The rest of this paper is organized as follows. In Section 2, I describe the policy change underlying the natural experiment. Section 3 discusses the two econometric strategies and addresses possible challenges to identification. Section 4 reports impacts on the first class of outcomes, prosocial behavior and generosity. Section 5 describes the experiments and results relating to discrimination and social interaction. Section 6 reports effects on learning and discipline. Section 7 summarizes the results and discusses shortcomings and avenues for future research.

2 Background and Policy Experiment

In this section, I describe a policy change which forced most elite private schools in Delhi to offer places to poor children, thus integrating poor and wealthy students in the same classrooms. I briefly describe how the timing of the policy change varied across schools, as well as key features of the selection process for both poor and wealthy students. In particular, poor students are selected using randomized lotteries, while wealthy applicants are selected using a transparent scoring system, which does not allow the use of baseline test scores or ability measures.

Delhi - like most cities in India - has a highly stratified school system. Public schools and a growing number of low-fee private schools serve the large population of urban poor. Relatively expensive ‘elite’ private schools cater to students from wealthy households.⁷ These types of schools differ widely in affordability, school inputs and acceptance rates. Public schools are free, and students are typically guaranteed admission to at least one public school in their neighborhood. In contrast, elite private schools as I define them charge tuition fees in excess of Rs. 2000 per month (approximately \$40, 25% of median monthly household consumption in 2010), and are vastly over-subscribed. Private schools in my sample report

⁷A loosely defined middle class typically sends its children to private schools intermediate in their price and exclusivity to public and elite private schools.

average acceptance rates of 11%, and monthly fees of up to Rs. 10,000.⁸

Policy Change. Many private schools in Delhi – including over 90% of the approximately 200 elite private schools – exist on land leased from the state in perpetuity at highly subsidized rates. A previously unenforced part of the lease agreement required such schools to make efforts to serve “weaker sections” of society. In 2007, prompted by the Delhi High Court, the Government of Delhi began to enforce this requirement. It issued an order requiring 395 private schools to reserve 20% of their seats for students from households earning under Rs. 100,000 a year (approximately \$2000). Schools were not permitted to charge the poor students any fees; instead, the government partially compensated the schools. Decades after most of these private schools were founded, the policy change forced open their doors to many relatively poor children.

Two features of the policy change are particularly important for my analysis: (i) schools were not permitted to track the students by ability or socioeconomic status. Instead, they were required to integrate the poor students into the same classrooms as the rich, and (ii) the policy only applied to new admissions, which occur almost exclusively in the schools’ starting grades (usually preschool). Thus, the policy did not change the composition of cohorts that began schooling before 2007.

Variation in timing. I divide elite private schools into three categories based on their response to the policy change. (i) *Treatment schools* were subject to the policy, and complied with it in the very first year. In these schools, cohorts admitted in 2007 or later have many poor students, while older cohorts comprise exclusively of wealthy students. This group includes about 90% of all elite private schools. (ii) *Delayed treatment schools* were also subject to the policy, but failed to comply in the first year - either because they expected the policy to be overturned in court, or because they felt the order was issued too late for them to modify their admission procedures. These schools complied with the policy a year later, in 2008, following a court ruling upholding the policy. This group comprises about 6% of all elite private schools. *Control schools* are the 4% of elite private schools which were not

⁸Parents of the wealthy students in the elite private schools I study apply to 8.8 schools on average and are offered admission to 1.8 of them. An article in the Indian Express memorably lamented that gaining admission to preschool in an elite private school in Delhi is harder than getting into Harvard.

subject to the policy at all, typically because they were built on land belonging to private charitable trusts or the federal government instead of the state government. In control schools, therefore, all cohorts comprise exclusively of rich students. The important point, discussed in detail in the next section, is that while schools are not randomly assigned to treatment, delayed treatment and control status, variation in the presence of poor children exists both within schools (across cohorts) and within cohorts (across schools).

Selection of Poor Students. If the seats for poor children are over-subscribed, schools are required to conduct a lottery to select beneficiaries. Conditional on applying to a given school, poor students are thus randomly selected for admission. While applications are free, they do involve the time costs of filling out and submitting the application form, and of obtaining documentation of income. Within the universe of eligible households, applicants are thus likely to be positively selected on their parents' preferences for their education, knowledge of the program, and their ability to complete the necessary paperwork. Since the children themselves are between 3 and 4 years of age when applying, it is less likely that their own preferences are reflected in the decision to apply.

The key point for this paper is that while the poor students may not be a representative sample of poor children in Delhi, they are without doubt from a very different economic class than the typical wealthy student in an elite private school. Figure 1 shows that the income cutoff of Rs. 100,000 per year is around the 45th percentile of the household income distribution, and the average poor student in my sample is from the 25th percentile. In contrast, the typical rich student in the sample is from well above the 95th percentile of the consumption distribution. In the US, a corresponding policy change would see students from households making \$23,000 a year attend the same schools as those making \$200,000 a year.

Selection of wealthy students. The admissions criteria used by elite private schools to select wealthy (fee-paying) students are strictly regulated by the government, and publicly declared by the schools themselves. Schools rank applicants using a point system, with the greatest weight placed on distance to the applicant's home and whether an older sibling is already enrolled in the school. Other factors include a parent interview, whether parents are

alumni, and gender (a slight preference is given to girls). Importantly, schools are prohibited from interviewing or testing students before making admissions decisions. Thus, it is difficult for schools to screen applicants on ability.⁹ The overwhelming majority of admissions to elite private schools occur in preschool, which is the usual starting grade. New students are typically only admitted to higher grades when vacancies are created by transfers, which are rare: 1.7% per year in my sample.

3 Econometric Strategies

Between February 2012 and September 2013, I conducted field and lab experiments, and gathered test scores and administrative data on 2032 students in 14 elite private schools in Delhi. The sample consists of 9 treatment schools, 2 delayed treatment schools and 3 control schools, recruited as part of a larger project studying a variety of learning and behavioral outcomes.¹⁰ Within each school, I constructed a representative sample of wealthy (that is, fee-paying) students in the four cohorts who began preschool between 2005 and 2008. Given the timing of the policy change, these students were in grades 2 (cohort of 2008) through 5 (cohort of 2005) at the time of data collection.

Using this data, I exploit two types of variation to identify the effects of poor students on their rich classmates: whether or not poor students are present in a particular cohort and school, and idiosyncratic variation in interactions with poor students within the classroom.

3.1 Variation within schools and cohorts

The first approach identifies the average effect of having poor students in one’s classroom. Recall that in treatment schools, wealthy students in grades 2 and 3 are “treated” with poor classmates, while grades 4 and 5 have no poor students. In delayed treatment schools, only

⁹Schools may, of course, use parent interviews to judge the ability of applicants. But parents cannot easily provide schools with credible information about student ability in the interviews, since the child is typically under 4 years of age and has no prior schooling.

¹⁰I contacted a total of 16 schools. Two of these schools (one control and one treatment school) declined to participate. The 16 schools were selected partly for convenience, but also to cover all 12 education districts of the Delhi Directorate of Education, while oversampling control and delayed treatment schools and satisfying my criteria for being elite schools (monthly fees exceeding Rs. 2000). All schools were guaranteed complete anonymity in exchange for participating.

grade 2 is treated, while grades 3-5 are untreated. And in control schools, grades 2-5 are all untreated. Restricting the sample to rich students, I estimate the following difference-in-differences specification by OLS:

$$Y_{igs} = \alpha + \delta_s + \phi_g + \beta \text{TreatedClassroom}_{gs} + \gamma X_{igs} + \varepsilon_{igs} \quad (1)$$

where Y_{igs} denotes outcome Y for student i in grade g in school s ; X is a vector of controls, δ_s are school fixed effects, ϕ_g are grade or cohort fixed effects and ε_{igs} is a student specific error term. $\text{TreatedClassroom}_{gs}$ is the treatment indicator; it equals one if grade g in school s contains poor students, and is zero otherwise. β is thus the average effect of having a poor classmate, and is the key parameter to be estimated. I cluster standard errors at the grade-by-school level, since this is the unit of treatment. With 14 schools and 4 grades (2 through 5) per school in the sample, this results in 56 clusters. For robustness, I also cluster standard errors at the school level. Given the small number of schools ($k = 14$), I use the wild-t cluster bootstrap method of Cameron, Gelbach and Miller (2008).

Note that average differences in outcomes across schools are permitted; they are controlled for by the school fixed effects. Thus, I do not assume that treatment, delayed treatment and control schools would have the same average outcomes without treatment. Similarly, average differences across cohorts (or grades) are controlled for using cohort fixed effects. This is important, given the possibility of age effects in social behaviors and preferences, as shown by Fehr et al. (2008) and Almas et al. (2012). I only utilize variation within schools (comparing students in different cohorts) and within cohorts (comparing students in different schools).

The identifying assumption is that in the absence of treatment, the *gaps* in outcomes across the different types of schools would be the same across treated and untreated grades. This would be violated if, for example, even in the absence of the policy, treatment schools had (say) better teachers than control schools in grades 2 and 3, but not in grades 4 and 5.

Challenges to Identification. This identification strategy faces the following potential challenges, each of which I briefly address below. (i) Wealthy students may select into control

schools based on their affinity for poor children. (ii) Treatment and delayed treatment schools have fewer seats for wealthy students after the policy change, which might increase the average ability of admitted students. (iii) There may be spillovers between treated and untreated grades within treated schools, and (iv) The policy may cause an increase in class size, which could directly affect outcomes.

The concern most relevant to estimating effects on social outcomes is that students might sort across the different types of schools based on their affinity for poor children.¹¹ In practice, this mechanism is of limited concern for the following reasons. First, it is difficult for parents to be picky, since acceptance rates at elite schools are low (about 10%) and less than 5% of such schools are control schools. Transfers between elite schools are also rare; control schools report very few open spaces in grades other than preschool each year.¹² Second, as a robustness check, I can restrict attention to students who had older siblings already enrolled in the same school. These students are likely to be less selected, both because parents might prefer to have both children in the same school, and because younger siblings of a current student are much more likely to be offered admission to the school. I show that none of the main results substantially change when restricting the sample in this way. Finally, the second identification strategy I describe below is entirely exempt from this concern, since it does not rely on variation across schools.

The main concern with estimating effects on academic outcomes is that the policy change may force treatment schools to become more selective when admitting wealthy students. And indeed, while the share of poor students in the incoming cohorts is around 18%, total cohort size only increases by 5%.¹³ This implies that fewer wealthy students are accepted into treated private schools after the policy change. If schools select students based on academic ability, this would mechanically raise the average quality of admitted wealthy students, and bias my estimate of the effect on learning outcomes. I can deal with this

¹¹For example, a parent who particularly dislikes the thought of his son sitting next to a poor child might try extra hard to have him enroll in one of the few control schools. Or students who find that they particularly dislike their poor classmates might transfer to a control school in later years.

¹²Additionally, I find that the number of applications to control schools relative to treatment schools does not increase after the policy change, suggesting that the policy change did not increase overall demand for the control schools amongst wealthy parents.

¹³The target of 20% reservation was not always met in the early years of the program.

concern as above - by restricting attention to the less-selected younger siblings of previously enrolled students, and by relying on the second identification strategy. However, it is also worth emphasizing that the schools are prohibited from testing or interviewing prospective students in starting grades. Since preschool applicants are between 3 and 4 years old, schools also have no prior test scores available while making their decisions. Thus, it is difficult for schools to screen applicants based on ability.

Spillovers between grades are likely minimal, since students spend over 85% of the school day exclusively with their assigned classmates, and little time interacting with students in other classrooms of the same grade, let alone students in other grades. To the extent that any spillovers do exist, they would bias against finding effects. And finally, class sizes increase by only 5% after the policy change. It is therefore unlikely that changes in class size could be important drivers of any effects.

The econometric strategy described above identifies the overall effect on wealthy students of having poor students integrated into their classrooms. This effect would be one important input to any evaluation of the costs and benefits of such programs. However, it tells us little about the mechanisms underlying any effects. In particular, it does not separate the effect of increased personal interactions between rich and poor students from other plausible classroom-level changes such as teachers spending more time teaching students about inequality and poverty.

3.2 Idiosyncratic variation within classrooms

The second approach uses membership in the same “study groups” as a proxy for personal interactions between students. Students in my sample spend an average of an hour a day engaged in learning activities in small groups of 2-4 students. Examples of such activities include collaborative craft projects, and working on math problems or reading comprehension. I collect data on study group membership in each school, and determine whether each student i has *any* poor children in his study group. I denote this binary measure by $HasPoorStudyPartners_i$.

In 8 of the 14 schools, students are assigned to study groups by alphabetic order of first

name (*SchoolUsesAlphaRule* = 1). In the remaining schools, groups are either frequently reshuffled by teachers, or no systematic assignment procedure is used (*SchoolUsesAlphaRule* = 0). I obtain class rosters, and sort them alphabetically to compute whether each student i is immediately followed or preceded by a poor student. I denote this measure by *HasPoorAlphabeticNeighbor_i*. I then estimate the following regression by two-stage least squares:

$$Y_{icgs} = \alpha + \nu_{cgs} + \beta_1 \text{HasPoorStudyPartners}_i + \gamma X_i + \varepsilon_{igs} \quad (2)$$

where Y_{icgs} denotes outcome Y for student i in classroom c in grade g in school s , ν_{cgs} is a classroom fixed effect, and $\text{HasPoorStudyPartners}_i$ is instrumented for using *SchoolUsesAlphaRule_s* * *HasPoorAlphabeticNeighbor_i*.

This identification strategy isolates the effect of personal interactions between rich and poor students. Identification comes entirely from within treated classrooms in treatment and delayed treatment classrooms, and average differences across classrooms (or schools and cohorts) are controlled for using classroom fixed effects. Thus, this strategy is not subject to concerns about the sorting of wealthy students across different types of schools, or changes in class size or teacher behavior.

Note that this approach does not require that poor and rich students have a similar alphabetic distribution of names. Nor do I assume that a rich student's first name has no direct effect on his outcomes. Instead, I only utilize variation in study groups predicted by the differential effect of alphabetic ordering in schools which use versus do not use such ordering to assign study groups. In the most aggressive specification, the individual level controls include first letter of first name, and thus directly control for any average differences across alphabetic order of names.

Figure 2 graphically reports the first stage of this regression. It shows that in the schools which report using alphabetic order to assign study groups, having a name alphabetically adjacent to at least one poor student substantially increases the probability of having at least one poor study partner, from about 40% to 90%. In contrast, alphabetic adjacency

has no effect in other schools. Table 1 provides the first stage regression, and reports that the instrument is strong, with an F -statistic of over 40.

4 Generosity and Prosocial Behavior

Common sense and empirical evidence suggest that human beings care about others, and about fairness. Economists have argued for the importance of such “social preferences” in domains including charitable donations (Andreoni 1998), support for redistribution (Alesina and Glaeser 2005), voter turnout (Edlin et al. 2007), and labor markets (Akerlof 1984, Bandiera et al. 2005). Researchers have measured social preferences in the field using behaviors like charitable giving and public goods provision (DellaVigna 2009), and in the lab using dictator games (Kahneman et al. 1986), where the participant (the “dictator”) typically decides how to split a pot of money between himself and an anonymous recipient.¹⁴

Recently, scholars have begun to investigate how social preferences are shaped by life experiences, including education (Jakiela et al. 2010), the ideology of one’s college professors (Fisman et al. 2012), political violence (De Voors et al. 2012) and macroeconomic conditions (Fisman et al. 2013). Researchers have also begun to trace the emergence of social preferences in children, where egalitarian preferences are seen to emerge around age 4-8 (Fehr et al. 2008), while more sophisticated notions of fairness emerge in adolescence (Almas et al. 2010).

In this section, I estimate how having poor classmates affects the prosocial behavior of wealthy students. I measure such behavior in two ways: in the field using administrative data on volunteering for charities, and in the lab using dictator games. I first find that wealthy students are substantially more likely to volunteer for a charity if they have poor classmates in school. Next, I show that the increase in prosocial behavior is also visible in the lab. Having poor classmates makes wealthy students more generous in dictator games. This

¹⁴More sophisticated versions of dictator games might vary the identity of the recipient (Hoffman et al. 1996) or vary the exchange rate at which money can be transferred between the dictator’s endowment and the recipient (Andreoni and Miller 2002). Choices made in such lab games have been shown to predict real-world behavior such as charitable donations (Benz and Meier 2008), loan repayment (Karlan 2005) and voting behavior (Finan and Schechter 2012).

increased generosity is partly driven by the wealthy students displaying more egalitarian preferences over monetary payoffs.

4.1 Prosocial Behavior - Volunteering for Charity

I begin by studying prosocial behavior in a setting familiar to students in elite private schools in Delhi. All the schools in my sample provide students with occasional opportunities to volunteer for charities. One such activity common across the schools involves spending two weekend afternoons in school to help fundraise for a charity serving disadvantaged children. The task itself is to help make and package greeting cards, which are then sold to raise money for the charity. Participation in these events is strictly voluntary; only 28% of students choose to attend. Volunteering activities thus serve as a natural real-world measure of prosocial behavior.

I collect administrative data on attendance at these volunteering events, and apply both econometric strategies described in the previous section to identify the effects of poor students on their wealthy classmates. Figure 3a graphically depicts the difference-in-differences strategy, plotting the share of students volunteering by grade and school type. The graph shows that wealthy students in grades 4 and 5 - which have no poor students - have similar volunteering rates across the three types of schools (control, treatment and delayed treatment). This suggests that the control schools are not especially bad at teaching prosocial behavior; before the policy change, all the schools had similarly prosocial students. However, wealthy students in treatment schools volunteer substantially more in grades 2 and 3 - precisely the grades which contain poor students. The same pattern is evident for delayed treatment schools, which are only treated in grade 2. This pattern of volunteering behavior suggests that it is having poor classmates that causes the increase in wealthy students' prosocial behavior.

Figure 3b shows the effect of having a poor study partner, and conveys the essence of the instrumental variable strategy. It shows the share of volunteers, separately by whether or not the wealthy student's name is alphabetically adjacent to at least one poor student in his class roster. The graph shows that wealthy students with names close to a poor student

are more likely to volunteer for the charity - but only in those schools which report using alphabetic order to assign study groups. This result suggests that it is having a poor student in one's study group - and therefore personally interacting with a poor student - that causes an increase in prosocial behavior.

The regression results in Table 2 confirm these findings, and attach a precise magnitude to the effects. Column 1 reports the main difference-in-differences estimate and shows that having poor classmates increases volunteering by 10 percentage points (se 2.4), an increase of 43% or 0.25 standard deviations over the volunteering rate in control classrooms. The effect remains highly significant ($p < 0.01$) when standard errors are clustered at the school level using the wild cluster bootstrap method of Cameron et al. (2008). Column 2 reports the same specification estimated on the restricted sample of students who have older siblings in the same school. The results are similar and not statistically different: an increase in volunteering of 7 percentage points (se 3.0). Column 3 reports the instrumental variable estimate of the effect of having a poor study partner. It shows that having at least one poor study partner causally increases volunteering by 13 percentage points (se 8.6), an imprecisely estimated increase of 53%.

4.2 Dictator Games

To complement the field measure of prosocial behavior, and to better understand any changes in social preferences, I invite students to play dictator games in a lab setting. In particular, I use two dictator games to measure their generosity towards both poor and rich recipients. In the first game, students are knowingly paired with an anonymous poor student at a school catering solely to poor children. In the second game, instead, students are paired with an anonymous rich student at a control school catering exclusively to rich students.

Protocol. Within each school, I randomly selected about 30 students from each of grades 2 through 5. I invited these students to experimental sessions conducted in a separate room in their school during a regular school day. Each student was provided with two numbered and colored envelopes, corresponding to the two dictator games. Each envelope contained a decision sheet and a photograph and brief description of the recipient's school. The

only difference between the two games was the identity of the recipient. In one game, the recipient was an anonymous student in a school catering to disadvantaged children. In the other, the recipient was an anonymous student in an elite private school which caters to wealthy students. In each case, the experimenter read the description aloud to the students, and also asked them to carefully look at the materials themselves. Students were informed that they would play two such games, and that one of the two games would be selected for implementation through a coin toss.

The decision sheet contained a table with rows corresponding to the possible splits of the endowment. The student's allocation was in the left column, and the recipient's endowment in the right column. The allocations summed to 10 rupees, with only integer values permitted. Students were asked to circle their desired allocation. As a check, they also filled in blanks stating how much they would receive, and how much the recipient would receive. After the game was verbally described to the students, the experimenter answered any questions they had. The game proceeded only once the students appeared to understand the procedure very well, and were able to correctly state how much they would receive by circling each row. Students then played the two games, in an order randomly varied by experimental session. After subjects made each decision, they placed the decision sheet and descriptive materials back in the envelope, closed it and returned it to the experimenter. Then, as promised, a coin was flipped to determine which choice would be implemented.

Students next completed a brief survey to debrief, and a set of additional games and questions for a related research project. At the end of this period, sealed envelopes were returned to students containing their chosen payoff. Students then had the option to use their payoff (and any other money they may have had) to purchase candy from a small store set up by the experimenter.

Poor Recipient. I find that having poor classmates and interacting with them in study groups makes wealthy students substantially more generous towards poor recipients. Fig 4 shows the results graphically, while Table 3 provides numerical estimates. Having poor classmates increases the average amount shared with a poor recipient by 12 percentage points (se 1.9), an increase of 45% or 0.45 standard deviations over the average giving in

control classrooms. The results are very similar for the reduced sample of younger siblings (Column 2). The instrumental variable estimates of Column 3 show that having at least one poor study partner partner causally increases giving by 7 percentage points (se 3.1), an increase of 22%.

Rich Recipient. Figure 5 plots the corresponding results for the amounts shared with rich recipients. They show a very similar pattern to the results for poor recipients, albeit with slightly smaller effect sizes. Table 4 reports that having poor classmates increases giving to wealthy recipients by 27% (se 5%), while having a poor study partner increases giving by a less precisely estimated 42% (se 23%).

Egalitarian Preferences. Digging deeper, Figure 6 plots the distribution of giving in the two games, separately for students in treated and untreated classrooms. The figures show a distinct increase in the probability of sharing exactly 50% of the endowment with the recipient. This raises the intriguing possibility that exposure to poor classmates makes wealthy students more egalitarian.

To explore this possibility further, I use a set of three simple dictator games designed to identify whether subjects dislike unequal allocations.¹⁵ Each game poses dictators with a binary choice between more and less equal distributions of payoffs. The less equal option provides a higher personal payoff (in the “equality game”) or a higher sum of payoffs for the two recipients (in the two “disinterested” dictator games). The payoffs in the games are listed in the table below.¹⁶

	More equal option	Less equal option
Equality Game	Dictator=5, Recip =5	Dictator=6, Recip =1
Disinterested Game 1	Recip A = 4, Recip B = 4	Recip A = 8, Recip B = 3
Disinterested Game 2	Recip A = 4, Recip B = 4	Recip A = 12, Recip B = 0

Table 5 reports that students with poor classmates are consistently more likely to pick the more equal outcome. Column 1 shows that treated students are 9 percentage points more likely to reduce their own payoff by choosing (5,5) over (6,1) in the equality game, compared to a base of 55% in the control group. And when choosing allocations for two

¹⁵These games are adapted from Charness and Rabin (2002).

¹⁶Note that the games themselves were presented without labels, and the order of the two options was randomized.

anonymous recipients (holding their own payoff fixed) in the disinterested dictator games, they are 12 percentage points more likely to choose (4,4) over (8,3) and 14 percentage points more likely to pick (4,4) over (12,0).

Considering the set of dictator game results together, I conclude that having poor classmates does not simply make students more charitable towards the poor. Instead, it makes them more generous overall, and in particular makes them exhibit more egalitarian preferences over monetary payoffs. This is an important point, suggesting that interacting with poor children in school does not just make students more favorably inclined towards the poor. Rather, it changes more fundamental preferences regarding fairness.

The dictator game measures were entirely independent of the field observations of volunteering behavior described previously. Putting together the findings of increased generosity in the lab and increased volunteering in the field thus substantially strengthens my conclusion that being exposed to poor children in school makes wealthy students more prosocial.

5 Social Interactions and Discrimination

Discrimination is a pervasive and important phenomenon in labor markets (Goldin and Rouse 2000, Bertrand and Mullainathan 2004), law enforcement (Persico 2009), residential location choice (Becker and Murphy 2009) and other contexts. Theories of discrimination are of two main types: taste-based discrimination, reflecting an innate animosity towards individuals from a particular group (Becker 1957), and statistical discrimination, which results from imperfect information about productivity or ability (Phelps 1972, Arrow 1973, Aigner and Cain 1977).

Tastes for social interactions provide a natural foundation for taste-based discrimination. But social interaction models also explain features of residential patterns (Schelling 1971), collective action (Granovetter, 1978), job search (Beaman and Magruder 2012) and the marriage market. Changes in willingness to interact with members of other social groups are therefore a potentially important outcome of diversity in schools. Indeed, theory suggests that even small changes in these tastes can lead to large differences in aggregate outcomes,

since social interaction models often feature multiple equilibria (Card et al. 2008).

In this section, I estimate how having poor classmates in school affects rich students' willingness to socially interact and work with other poor children in teams, or conversely to discriminate against them. I design two novel experiments to measure these outcomes. The first is a team selection field experiment designed to estimate taste-based discrimination using exogenous variation in the price of discrimination. The second experiment elicits students "willingness to play" - the cost they attach to attending a play date with poor children.

5.1 Team-Selection Field Experiment

Design. The main idea of the team-selection experiment is to create a tradeoff for wealthy students between choosing a high ability teammate (and thus increasing their own expected payoff) versus choosing a lower-ability teammate with whom they would prefer to socialize. The team task I used in the experiment was a relay race, a task which was familiar to all the students, and in which ability is easily revealed through times in individual sprints. In addition to running the relay race together, participants were required to spend time socializing with their teammates.

The experiment was conducted on the sidelines of a sports meet featuring athletes from two elite private schools - one a treatment school, and the other a control school. The participants in the experiment were not the athletes themselves, but were instead drawn from the large contingents of students who were present to support their teams. Note that attendance in this supporting role was compulsory for students in both schools; the attendees were not a selected set of cheerleaders. In addition to students from the two elite private schools, I invited selected students from a public school catering to relatively poor students to participate in the experiment. These students were selected for having a particular interest in athletics.

The experiment proceeded in four stages.

Randomization. First, students were randomized to different sessions (separately by gender) with varying stakes for winning the subsequent relay race - either Rs. 500, Rs. 200

or Rs. 50 per teammate for winning the race. 500 rupees are approximately a month's pocket money for the oldest students in the sample, so the stakes are substantial. Within each session, students were asked to mix and introduce themselves to each other for about fifteen minutes. This ensured that students were able to accurately identify the difference in the social groups that the various participants belonged to. School uniforms made group membership salient, and debriefing suggested that students were quickly able to identify that the students from the public school were relatively poor, while the students from the private schools were wealthy. At the end of this phase, the following three stages were described to the students, and the experiment proceeded.

Ability Revelation and Team Selection. Students watched a series of one-on-one sprints, designed to reveal each runner's ability. In most cases, one runner was from the public school, while the other was from one of the private schools. However, some pairs included two students from private schools, or two from the public school. After each sprint, the rank (first or second) and times of the two runners were announced.

Decision Stage. After each such ability revelation sprint, students privately chose on a worksheet which of the two students they would like to have in their two-person team for a relay race. After the sprints were complete, six students were picked at random to participate in the relay race, and one of their choices was randomly selected for implementation.

Relay Race and Socializing with Teammate. The relay races were conducted and rewards were distributed as promised. After the rewards were distributed, students were required to spend two hours socializing with their teammate. They were provided with board games and could also use playground equipment. However, they were not permitted to play in larger groups. This part of the experiment was described to the participants in advance, so they knew that their interactions with their selected teammate would exceed the few minutes spent on the relay race itself.

Reduced Form Results. The first reduced form finding is significant discrimination against the poor on average. I classify a wealthy student as having discriminated against the poor if he or she chooses a lower ability (i.e. slower) rich student from another school

over a higher ability poor student from the public school.¹⁷ Averaging over the different reward conditions, participants discriminate 19% of the time. These are not just mistakes, since the symmetric mistake of “discriminating” against a rich student occurs only 3% of the time. And when participants are choosing between two runners from the same (other) school, they pick the slower runner only 2% of the time. Thus, only poor students competing against rich students are systematically discriminated against.

The second finding is that discrimination decreases as the stakes increase. In the control school, 35% of choices exhibit discrimination against the poor in the Rs. 50 condition, but this falls to 27% when the reward is Rs. 200, and only 5% in the highest stakes condition of Rs. 500. This result is shown by the solid line in Figure 7, which I interpret as a demand curve for discrimination.

The third and most important finding is a reduction in discrimination from having poor classmates and study partners. Figure 7 shows that for each level of stakes, wealthy students with poor classmates are less likely to discriminate against the poor. In addition, the slope of the demand curve for discrimination is higher for students with poor classmates. Figure 8(a) depicts the difference-in-differences estimates graphically by plotting rates of discrimination by school and grade. In the treatment school, discrimination is substantially lower than in the control school in the treated grades 2 and 3, but not in grades 4 and 5. Figure 8(b) instead depicts the reduced form of the IV strategy, plotting rates of discrimination by whether the student has a name alphabetically adjacent to a poor students. Consistent with the difference-in-differences result, the figure shows that students with names close to a poor student (and therefore a higher likelihood of having a poor study partner) discriminate less.

Regression estimates are reported in Table 6. Column 1 shows that having a poor classmate reduces discrimination by 12 percentage points (se 5).¹⁸ This effect is comparable to the 11 percentage point reduction in discrimination caused by increasing the stakes from

¹⁷I do not consider a choice to be discriminatory if it involves choosing ones own schoolmate over a higher-ability poor student, since participants may prefer to partner with children they already know.

¹⁸Since the discrimination experiment has wealthy students from only two schools, I do not attempt to cluster standard errors at the school level. Instead, I report unclustered standard errors and, as a robustness check, cluster at the school-by-grade level (8 clusters) using the wild cluster bootstrap method.

Rs. 50 to Rs. 200 (an increase of about \$3). Column 2 shows that having poor classmates has the biggest effect on discrimination in the lowest stakes condition (a 25 percentage point reduction). Column 3 reports the IV result that having a poor study partner reduces discrimination by 14.7 percentage points (se 8.8).¹⁹

The observed behavior is more consistent with taste-based discrimination rather than statistical discrimination. When a separate sample of students is asked which of the two runners is more likely to be in the winning relay race, 98% pick the faster student. This implies that many students prefer a wealthy teammate even though they believe he makes them *less* likely to win, a fact inconsistent with a simple model of statistical discrimination. This is not surprising, since the experiment was designed with the intention of measuring taste-based discrimination. The clear signals of ability provided by the sprints additionally make statistical discrimination unlikely. And the fact that participants are forced to actually spend time socializing with their teammates - as is often the case when hiring colleagues or employees - provides a natural setting for taste-based discrimination.

Since the experiment does not simulate a market with a wage offered to teammates, I cannot directly test the classic prediction of taste-based discrimination that only the marginal employer's tastes matter. However, the fact that about 35% of control students discriminate in the low stakes condition suggests that at least a third of wealthy students dislike having a poor teammate (relative to a wealthy teammate).

Model and Structural Estimation. The reduced form results provide evidence of a reduction in discrimination. But they do not inform us of the precise magnitude of the distaste that wealthy students have for partnering and socializing with a poor child, nor how much this is changed by having poor classmates. In order to estimate these quantities, I structurally estimate a simple model of discrimination.

Model. Suppose the decision-maker has expected utility:

$$U_t = p_t M + S_t \tag{3}$$

¹⁹The treatment school uses alphabetic order to assign study partners. Since the sample for this experiment does not include any other treatment schools which do not use such a rule, I directly use alphabetic adjacency to a poor student as the instrument for having a poor student in one's study group.

where p_t is the probability of winning the race with teammate t , M is the monetary reward for winning the race and S_t is the utility from socially interacting with teammate t . I assume that teammates are of two types, $t \in \{R, P\}$, where R denotes a rich student and P denotes a poor student.

Then, she chooses the rich teammate if

$$p_R M + S_R > p_P M + S_P$$

$$\Leftrightarrow S_R - S_P > (p_P - p_R) M$$

In the absence of a particular distaste for having a poor teammate, $S_P = S_R$. And in the absence of statistical discrimination, rich and poor students with the same performance in the sprint would be perceived to be equally able, $p_P = p_R$. Define $D_{poor} \equiv S_R - S_P$ as the distaste for interacting with a poor student (relative to a rich student), and $\delta_{poor} \equiv p_P - p_R$ as the perceived *increase* in probability of winning from having a poor teammate, provided the poor student won the ability-revelation sprint. Then, the decision-maker discriminates against a poor student if:

$$D_{poor} > \delta_{poor} M \tag{4}$$

Similarly, in the case where the rich student wins the sprint, we can define the increase in probability of winning from choosing the rich teammate, δ_{rich} .

In order to estimate the model, I impose the following distributional assumption: D_{poor} is distributed normally with mean μ_D^T and standard deviation σ_D^T , separately for students from treated classrooms ($T = 1$) and untreated classrooms ($T = 0$). Consistent with the fact that 98% of students state that the winner of the ability sprint is more likely to win the relay race (regardless of whether the winner was rich or poor), I additionally impose the assumption of no statistical discrimination, $\delta_{poor} = \delta_{rich} \equiv \delta$.

Then, the parameters to be estimated are: (i) μ_D^1 and μ_D^0 , the average distaste for

having a poor teammate amongst students with and without poor classmates, respectively; (ii) σ_D^1 and σ_D^0 , the standard deviations of the distribution of distaste; (iii) δ , the increase in probability of winning from choosing the teammate who won the ability-revelation sprint.

I estimate these parameters using a classical minimum distance estimator. Specifically, the estimator solves $\text{Min}_\theta (m(\theta) - \hat{m})' W (m(\theta) - \hat{m})$, where $\hat{m}$ is a vector of the empirical moments and $m(\theta)$ is the vector of theoretically predicted moment for parameters θ . The weighting matrix W is the diagonalized inverse of the variance of each moment; more precisely estimated moments receive greater weight in the estimation.

The moments for the estimation are the following: (i) The probability of discriminating against a higher-ability poor student, separately by stakes $M \in \{50, 200, 500\}$ and treatment status $T \in \{0, 1\}$, and (ii) The probability of discriminating against a higher-ability rich student, by stakes M and treatment status T . The empirical moments are simply shares of students observed to discriminate in each condition, estimated by an uncontrolled OLS regression.

Identification. All 5 parameters are jointly identified using the 12 moments. The intuition for the identification is straightforward. Conditional on δ , the exogenous variation in the stakes M pins down the mean μ_D and standard deviation σ_D of the distribution of distaste D . Conditional on the distribution of D , the perceived increase in probability of winning from choosing a high-ability teammate (δ), is identified from comparing the probabilities of choosing (say) the poor student when he wins versus loses the ability-revelation sprint.

Estimates. The lower panel in Table 7 reports the empirical and fitted values of the moments. The model overall does a good job of fitting the moments, with the exception of slightly over-predicting discrimination against the poor in the lowest stakes condition. Table 7 also reports the structural estimates of the parameters. The perceived increase in probability of winning from choosing a high-ability teammate is imprecisely estimated, $\delta = 0.08$ (se 0.1). Students without poor classmates are estimated to feel an average distaste for having a poor teammate of $\mu_D^0 = \text{Rs. } 37$ (se Rs. 4.4), with a standard deviation $\sigma_D^0 = \text{Rs. } 6$ (se Rs. 1.9). In contrast, treated students are estimated to have a substantially

lower distaste of $\mu_D^1 = \text{Rs. } 26$ (se Rs. 4.8) and a similar standard deviation, $\sigma_D^1 = \text{Rs. } 5$ (se Rs. 2.1). The difference in average distaste of Rs. 11 is significant at the 10% level, and constitutes a 30% reduction relative to students without poor classmates.

5.2 Willingness to Play

To shed more light on the observed reduction in discrimination in the team-selection experiment, I directly test wealthy students' tastes for socially interacting with poor children. I do so by inviting them to play dates at neighborhood schools for poor children. The play dates were motivated as an opportunity to make new friends, and involved two hours of games and playground activities. In order to measure tastes, I elicited incentivized measures of wealthy students' willingness to accept to attend these play dates. I find that having poor classmates in school makes wealthy students substantially more willing to play with other poor children.

Protocol. First, students were informed in school about the play dates. The play date was presented to them as an opportunity to make new friends in their neighborhood. The host school was named and described, and the experimenter showed the students a photograph of the school. The play dates all occurred on a weekend morning, and the students were informed about them approximately two weeks before the play date.

After answering students questions about the planned play dates, I elicited their willingness to accept - the payment they required - to attend the play date.²⁰ I employed a simple Becker-DeGroot-Marschak mechanism, where students were presented with a decision sheet showing possible levels of payments for attending the play date. For each such price level, they were asked to indicate whether they would like to attend the play date. Then, a numbered ball was drawn from a bag, and their decision corresponding to that price was implemented. In particular, if they had indicated they would like to attend for the drawn price, their name was written down on a list, and they were provided an invitation form to take

²⁰Pilot work revealed that students generally find the play dates unattractive - nearly all students expressed a negative willingness to pay. This is unsurprising, given that the play dates were held on weekends, which are surely a precious part of a child's week. The opportunity cost of attending anecdotally included watching television and playing with existing friends.

home to their parents.²¹ The entire procedure was first explained to the students verbally, and then they played three practice rounds, at which point they appeared to understand the decision well.

Results. The key finding is that students become more willing to socialize with poor children if they already have poor classmates. Figure 9 shows the results of the two identification strategies graphically. Panel (a) plots average willingness to accept by school type (control, treatment and delayed treatment) and grade. For treatment schools, willingness to accept is lower than control schools only in the treated grades 2 and 3, but not in the untreated grades 4 and 5. A similar pattern is visible for delayed treatment schools, in which only grade 3 is treated. A lower willingness to accept indicates greater willingness to socially interact with poor children. Figure 10 plots the resulting supply curves for attending the play date, separately for students with and without poor classmates.

Figure 9(b) depicts the IV strategy graphically, by plotting average willingness to accept by whether the wealthy student has a name alphabetically adjacent to any poor children. We see that having an alphabetic neighbor only increases willingness to play in schools which use the alphabetic assignment rule. This suggests that personally interacting with a poor student through assignment to a common study group reduces a wealthy students' distaste for interacting with other poor children.

Finally, Table 8 reports numerical estimates of the effects, using the specifications discussed in Section 3. I find that having poor classmates decreases willingness to accept (i.e. increases willingness to play) by Rs. 7 (se 1.1) on a base of Rs. 37, a decrease of 19%. The effect is highly significant ($p < 0.01$) even when clustering standard errors at the school level, and the result is similar in the restricted sample of younger siblings. Having a poor study partner reduces willingness to accept by 24% (se 9%).

In contrast, I find no effects on willingness to attend play dates with rich students. In August 2013, I conducted a parallel experiment in a smaller sample as a placebo test. Students now had the opportunity to spend two hours playing with other wealthy students

²¹Parents, of course, had the ability to veto their childrens' choice to attend the play date, and did so in about 35% of cases. Since I wish to isolate the child's tastes rather than the parents, I use the elicited willingness to pay (or accept) as the outcome measure. Using actual attendance of play dates as an alternative outcome, I find similar but muted effects.

from a control private school. While the estimates are less precise due to a smaller sample size, Appendix Table 1 reports no average effect on willingness to play with rich students.

6 Academic Outcomes

One concern with integrating disadvantaged students into elite schools is that wealthy students' academic outcomes may suffer as a result. This concern is motivated by the large literature studying peer effects in education, which has sometimes found substantial peer effects (Hoxby 2000, Hanushek et al. 2003) and at other times no evidence that peers affect test scores (Angrist and Lang 2004, Imberman et al. 2009). Classroom disruptions by poorly disciplined students have been proposed to be a key mechanism underlying any negative effects (Lazear 2001, Lavy and Schlosser 2011, Figlio 2007). Indeed, principals in the schools I studied reported being particularly concerned about classroom disruptions and learning. In this section, I therefore turn attention to estimating the impact of poor students on the learning and classroom discipline of their wealthy peers.

6.1 Learning

To measure effects on learning, I conduct simple tests of learning in English, Hindi and Math. With the assistance of teachers in a non-sample school, I first assembled a master list of questions from standard textbooks for grades 1 through 7. Students in each grade were asked to answer a set of questions considered appropriate for their grade, and a smaller set of questions at lower and higher grade levels. The test was designed to be quick and easy to implement, and therefore provides a somewhat coarse measure of learning. Nonetheless, it provides comparable test scores across different schools in the absence of any existing system of standardized testing in primary schools. I normalize the test score in order to provide standardized effect sizes.

I find that poor students do worse than rich students on average, but with substantial heterogeneity. Poor students score 0.32 standard deviations (s.d.) worse than wealthy students in English, 0.12 s.d. worse in Hindi and 0.24 s.d. worse in Math. The lower

average learning levels of poor students make the possibility of negative peer effects very real. But the variance in poor students' test scores is similar to that of wealthy students; there is thus plenty of overlap in the distributions of academic achievement. For example, poor students have weakly higher scores than 40% of their wealthy study partners even in English.

Table 9 reports regression estimates of the effects of poor students on their wealthy classmates' test scores. The first two columns show a small and insignificant effect on an equally weighted average of standardized scores in the individual subjects. I also consider effects on each subject in turn. Most importantly, I estimate a 0.09 standard deviation reduction in average test scores in English in treated classrooms, significant at the 10% level. The coefficient on the IV regression of English scores on having a poor study partner is also negative, although quite imprecisely estimated. In contrast, I find no effects of poor classmates on wealthy students' test scores in Hindi or Math. The online appendix reports similar effects for the likely less-selected subsample of younger siblings.

Considering the results for the different subjects together, the overall pattern is one of mixed but arguably modest effects on academic achievement. The only negative effect is on English scores. This is consistent with English being the subject with the largest achievement gap between rich and poor students, perhaps due to the fact that poor students almost exclusively report speaking only Hindi at home. But substantial learning gaps also exist in Math and (to a lesser extent) Hindi, and yet I detect no negative peer effects in those subjects. These latter non-effects are consistent with those of Muralidharan and Sundararaman (2013), who find no effects on the achievement of existing students in private schools in rural India when initially lower-achieving voucher-recipients enter their schools.²² An additional mechanism, supported by anecdotal reports from teachers, is that the presence of poor students causes conversations between students to shift more from English to Hindi, which might well reduce wealthy students' fluency in English. However, I find no evidence of a significant increase in Hindi test scores.

²²But note that Muralidharan and Sundararaman (2013) study a quite different setting. I study the effects on quite wealthy students attending elite private schools in urban Delhi, while they study the effects on students attending relatively modest private schools in rural Andhra Pradesh, where the social and economic disparity between the existing and incoming students is likely much smaller.

6.2 Discipline

To measure classroom discipline, I ask teachers to report whether each student has been cited for any disciplinary infractions in the past six months²³. I find that 22% of wealthy students have been cited for the use of inappropriate language (that is, swearing) in school, but only about 6% are cited for disruptive or violent behavior. Poor students are no more likely than rich students to be disruptive in class, but they are 12 percentage points more likely to be reported for using offensive language.

Table 10 reports regression estimates of the effects of poor students on disciplinary infractions by their wealthy classmates. The results suggest that having poor classmates increases the share of wealthy students reported for using inappropriate language by 7.5 percentage points (se 3.7). Having a poor study partner causes an even larger increase of 10 percentage points (se 6), an increase of about 45%. In contrast, I find precisely estimated zero effects on the likelihood of being cited for disruptive or violent behavior.

The finding that poor students do not make their wealthy classmates more disruptive – and indeed are no more disruptive than wealthy students themselves – is consistent with the absence of negative peer effects on Hindi and Math scores. In the context I study, concerns about diversity affecting test scores through indiscipline appear to be unwarranted. In contrast, the effects on inappropriate language use are substantial, although likely without implication for learning and future academic achievement.

7 Conclusion

In this paper, I exploit a natural experiment in education policy in India to estimate how greater economic diversity in classrooms affects wealthy students. I assemble a variety of evidence from field and lab experiments, administrative data and tests of learning to reach three main findings. The first finding is that having poor classmates makes wealthy students more prosocial and concerned about equality, and thus more generous towards others. The second finding is that wealthy students become more willing to socially interact

²³This question was asked at the end of the academic year

with poor children outside school, and thus exhibit less taste-based discrimination against the poor. The third finding is of mixed but overall modest impacts on academic outcomes, with negative effects on English language learning but no effect on Hindi or Math. Thus, my overall conclusion is that increased diversity in the classroom led to large and arguably positive impacts on social preferences and behaviors, at the cost of negative but modest impacts on academic outcomes.

One implication of these findings is that school policies involving affirmative action, desegregation and tracking should be evaluated not only on learning outcomes - which are of unarguable importance - but also on other economically important outcomes related to social behaviors. More generally, my findings support the view that increased interactions across social groups, perhaps especially in childhood, can improve inter-group behaviors. This finding has implications for diverse and polarized societies, where the costs of social divisions are thought to be high. Finally, my findings are of relevance to the planned expansion of this policy across India over the next few years.

One limitation of this paper is that due to the recency of the policy experiment, it does not study important long-term outcomes and behaviors such as political beliefs, social interactions as adults, and marriage market choices. Another limitation is the very particular nature of the sample - wealthy students in elite private schools in Delhi. Finally, this paper does not examine the effects on poor students of attending elite private schools, which might have profound consequences for their academic achievement and life outcomes. In ongoing and future work, I hope to address these shortcomings, as well as examine other aspects of the impacts of diversity in schools.

References

- [1] AIGNER, D. J., AND CAIN, G. G. Statistical theories of discrimination in labor markets. *Industrial and Labor relations review* 30, 2 (1977), 175–187.
- [2] AKERLOF, G. A. Gift exchange and efficiency-wage theory: Four views. *The American Economic Review* 74, 2 (1984), 79–83.
- [3] ALESINA, A., AND GLAESER, E. Fighting poverty in the us and europe: A world of difference. *OUP Catalogue* (2005).
- [4] ALESINA, A., AND PEROTTI, R. Income distribution, political instability, and investment. *European Economic Review* 40, 6 (1996), 1203–1228.
- [5] ALLPORT, G. W. *The nature of prejudice*. Addison-Wesley, 1954.
- [6] ALMÅS, I., CAPPELEN, A. W., SØRENSEN, E. Ø., AND TUNGODDEN, B. Fairness and the development of inequality acceptance. *Science* 328, 5982 (2010), 1176–1178.
- [7] ANDREONI, J. Toward a theory of charitable fund-raising. *Journal of Political Economy* 106, 6 (1998), 1186–1213.
- [8] ANDREONI, J., AND MILLER, J. Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica* 70, 2 (2002), 737–753.
- [9] ANGRIST, J. D., AND LANG, K. Does school integration generate peer effects? evidence from boston’s metco program. *The American Economic Review* 94, 5 (2004), 1613–1634.
- [10] ARROW, K. The theory of discrimination. In *Discrimination in Labor Markets*, O. Ashenfelter and A. Rees, Eds. Princeton University Press, 1973.
- [11] BANDIERA, O., BARANKAY, I., AND RASUL, I. Social preferences and the response to incentives: Evidence from personnel data. *The Quarterly Journal of Economics* 120, 3 (2005), 917–962.
- [12] BARDHAN, P., GHATAK, M., AND KARAIVANOV, A. Wealth inequality and collective action. *Journal of Public Economics* 91, 9 (2007), 1843–1874.
- [13] BECKER, G. S. *The economics of discrimination*. University of Chicago press, 1957.


- [14] BECKER, G. S., AND MURPHY, K. M. *Social economics: Market behavior in a social environment*. Harvard University Press, 2009.
- [15] BENZ, M., AND MEIER, S. Do people behave in experiments as in the field? evidence from donations. *Experimental Economics* 11, 3 (2008), 268–281.
- [16] BERTRAND, M., AND MULLAINATHAN, S. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *The American Economic Review* 94, 4 (2004), 991–1013.
- [17] BOISJOLY, J., DUNCAN, G. J., KREMER, M., LEVY, D. M., AND ECCLES, J. Empathy or antipathy? the impact of diversity. *The American Economic Review* 96, 5 (2006), 1890–1905.
- [18] BUCHMANN, C., AND HANNUM, E. Education and stratification in developing countries: A review of theories and research. *Annual review of sociology* (2001), 77–102.
- [19] BURNS, J., CORNO, L., AND LA FERRARA, E. Does interaction affect racial prejudice and cooperation? evidence from randomly assigned peers in south africa. *Unpublished working paper* (2013).
- [20] CAMERON, A. C., GELBACH, J. B., AND MILLER, D. L. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics* 90, 3 (2008), 414–427.
- [21] CARD, D., MAS, A., AND ROTHSTEIN, J. Tipping and the dynamics of segregation. *The Quarterly Journal of Economics* 123, 1 (2008), 177–218.
- [22] CHARLES, K. K., AND GURYAN, J. Studying discrimination: Fundamental challenges and recent progress. Tech. rep., National Bureau of Economic Research, 2011.
- [23] CHARNESS, G., AND RABIN, M. Understanding social preferences with simple tests. *The Quarterly Journal of Economics* 117, 3 (2002), 817–869.
- [24] EASTERLY, W., AND LEVINE, R. Africa’s growth tragedy: policies and ethnic divisions. *The Quarterly Journal of Economics* 112, 4 (1997), 1203–1250.

- [25] EDLIN, A. S., GELMAN, A., AND KAPLAN, N. Voting as a rational choice: Why and how people vote to improve the well-being of others. *Rationality and society* 1 (2007).
- [26] ESTEBAN, J., AND RAY, D. Linking conflict to inequality and polarization. *The American Economic Review* 101, 4 (2011), 1345–1374.
- [27] FEHR, E., BERNHARD, H., AND ROCKENBACH, B. Egalitarianism in young children. *Nature* 454, 7208 (2008), 1079–1083.
- [28] FIGLIO, D. N. Boys named sue: Disruptive children and their peers. *Education* 2, 4 (2007), 376–394.
- [29] FINAN, F., AND SCHECHTER, L. Vote-buying and reciprocity. *Econometrica* 80, 2 (2012), 863–881.
- [30] FISMAN, R., JAKIELA, P., AND KARIV, S. How did the great recession impact social preferences? *Unpublished working paper* (2012).
- [31] FISMAN, R., KARIV, S., AND MARKOVITS, D. Exposure to ideology and distributional preferences. *Unpublished paper* (2009).
- [32] GOLDIN, C., AND ROUSE, C. Orchestrating impartiality: The impact of. *The American Economic Review* 90, 4 (2000), 715–741.
- [33] GRANOVETTER, M. Threshold models of collective behavior. *American journal of sociology* (1978), 1420–1443.
- [34] HANUSHEK, E. A., KAIN, J. F., MARKMAN, J. M., AND RIVKIN, S. G. Does peer ability affect student achievement? *Journal of Applied Econometrics* 18, 5 (2003), 527–544.
- [35] HATTIE, J. A. Classroom composition and peer effects. *International Journal of Educational Research* 37, 5 (2002), 449–481.
- [36] HOFFMAN, E., MCCABE, K., AND SMITH, V. L. Social distance and other-regarding behavior in dictator games. *The American Economic Review* 86, 3 (1996), 653–660.
- [37] HOXBY, C. Peer effects in the classroom: Learning from gender and race variation. Tech. rep., National Bureau of Economic Research, 2000.

- [38] IMBERMAN, S., KUGLER, A. D., AND SACERDOTE, B. Katrina’s children: evidence on the structure of peer effects from hurricane evacuees. Tech. rep., National Bureau of Economic Research, 2009.
- [39] JAKIELA, P., MIGUEL, E., AND TE VELDE, V. L. You’ve earned it: Combining field and lab experiments to estimate the impact of human capital on social preferences. Tech. rep., National Bureau of Economic Research, 2010.
- [40] KARLAN, D. S. Using experimental economics to measure social capital and predict financial decisions. *American Economic Review* (2005), 1688–1699.
- [41] KARSTEN, S. School segregation. In *Equal Opportunities?: The Labour Market Integration of the Children of Immigrants*. OECD Publishing, 2010.
- [42] LAVY, V., AND SCHLOSSER, A. Mechanisms and impacts of gender peer effects at school. *American Economic Journal: Applied Economics* 3, 2 (2011), 1–33.
- [43] LAZEAR, E. P. Educational production. *The Quarterly Journal of Economics* 116, 3 (2001), 777–803.
- [44] LIST, J. A. The nature and extent of discrimination in the marketplace: Evidence from the field. *The Quarterly Journal of Economics* 119, 1 (2004), 49–89.
- [45] MIGUEL, E., AND GUGERTY, M. K. Ethnic diversity, social sanctions, and public goods in kenya. *Journal of Public Economics* 89, 11 (2005), 2325–2368.
- [46] MOBIUS, M. M., AND ROSENBLAT, T. S. Why beauty matters. *The American Economic Review* (2006), 222–235.
- [47] MURALIDHARAN, K., AND SUNDARARAMAN, V. The aggregate effect of school choice: Evidence from a two-stage experiment in india. 2013.
- [48] PERSICO, N. Racial profiling? detecting bias using statistical evidence. *Annu. Rev. Econ.* 1, 1 (2009), 229–254.
- [49] PETTIGREW, T. F., AND TROPP, L. R. A meta-analytic test of intergroup contact theory. *Journal of personality and social psychology* 90, 5 (2006), 751.


- [50] PHELPS, E. S. The statistical theory of racism and sexism. *The american economic review* 62, 4 (1972), 659–661.
- [51] SACERDOTE, B. Peer effects in education: How might they work, how big are they and how much do we know thus far? *Handbook of the Economics of Education* 3 (2011), 249–277.
- [52] SCHELLING, T. C. Dynamic models of segregation. *Journal of mathematical sociology* 1, 2 (1971), 143–186.
- [53] SCHOFIELD, J. W. School desegregation and intergroup relations: A review of the literature. *Review of research in education* 17 (1991), 335–409.
- [54] VAN EWIJK, R., AND SLEEGERS, P. The effect of peer socioeconomic status on student achievement: A meta-analysis. *Educational Research Review* 5, 2 (2010), 134–150.
- [55] VOORS, M. J., NILLESEN, E. E., VERWIMP, P., BULTE, E. H., LENSINK, R., AND VAN SOEST, D. P. Violent conflict and behavior: a field experiment in burundi. *The American Economic Review* 102, 2 (2012), 941–964.

Figure 1. Program eligibility and the household income distribution in Delhi





Note: This graph is based on the household consumption distribution reported in NSS-2010, with consumption amount converted to income levels using the ratio of household income to household consumption for urban Indian households reported in IHDS-2005.

Figure 2. First Stage of Instrumental Variable



Note: Having a name alphabetically adjacent to a poor student predicts having a poor student in one's study group, but only in the schools which explicitly use alphabetic ordering



Figure 3. Volunteering for Charity



Note: The top panel plots the share of wealthy students who participate in voluntary charitable fundraising activities in school, separately by type of school. Error bars plot 95% confidence intervals (unclustered).

The bottom panel plots share volunteering by whether the subject has a name alphabetically adjacent to any poor students, separately by whether schools use alphabetic order to assign study groups.



Figure 4. Giving to Poor Recipients in Dictator Game.



Note: The top panel plots the share given to a poor recipient against grade, separately by type of school. 95% confidence intervals (unclustered) for the mean are included. The figure shows that giving is higher in treatment and delayed treatment schools only in the treated grades.

The bottom panel plots the share of the endowment given to the poor recipient by whether the subject has a name alphabetically adjacent to any poor students, separately by whether schools use alphabetic order to assign study groups. This figure thus graphically depicts the reduced form regression of generosity to the poor on the excluded instrument.



Figure 5. Giving to Wealthy Recipients in Dictator Game.



Note: The top panel plots the share given to a wealthy recipient against grade, separately by type of school. 95% confidence intervals (unclustered) for the mean are included.


The bottom panel plots the share of the endowment given to the poor recipient by whether the subject has a name alphabetically adjacent to any poor students, separately by whether schools use alphabetic order to assign study groups.

Figure 6. Distribution of Giving in Treated and Untreated Classrooms





Note: The top panel shows the distribution of giving by wealthy students to poor recipients, separately for whether they have poor classmates (red bars) or not (blue bars). The bottom panel shows the same results for giving to wealthy recipients instead of poor recipients.

Figure 7. Demand Curve for Discrimination



Notes: This graph plots the share of wealthy students who discriminate against the poor (on the y axis) by the stakes of the decision, separately by whether the student has poor classmates (dotted red line) or not (solid green line). A student is classified as having discriminated against the poor if he chooses a lower-ability rich student over a higher-ability poor student.



Figure 8. Discrimination Against the Poor



Notes: The top panel plots the share of wealthy students who discriminate against the poor (on the y axis) by grade (on the axis), separately by school type. The control school is represented by the solid green line, while the treatment school is represented by the dotted red line. Error bars plot 95% confidence intervals (unclustered).

The bottom panel plots discrimination rates by whether the participant has a name alphabetically adjacent to any poor students, only for the treatment school.


Figure 9. Willingness to Play



Note: The top panel plots wealthy students' average Willingness To Accept a play date with poor children, separately by type of school. Error bars plot 95% confidence intervals (unclustered).

The bottom panel plots "willingness to play" by whether the subject has a name alphabetically adjacent to any poor students, separately by whether schools use alphabetic order to assign study groups.

Figure 10. Supply Curve for Attending Play Date



Note: This graph plots the share of wealthy students willing to attend the play date with poor children (on the y axis) for each given level of payment for attending the play date (on the x axis). The solid green line represents this supply curve for wealthy students without poor classmates, while the dotted red line represents wealthy students who do have poor classmates.