

NBER WORKING PAPER SERIES

RATIONAL IGNORANCE IN EDUCATION:
A FIELD EXPERIMENT IN STUDENT PLAGIARISM

Thomas S Dee
Brian A. Jacob

Working Paper 15672
<http://www.nber.org/papers/w15672>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2010

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2010 by Thomas S Dee and Brian A. Jacob. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Rational Ignorance in Education: A Field Experiment in Student Plagiarism
Thomas S Dee and Brian A. Jacob
NBER Working Paper No. 15672
January 2010
JEL No. I2,K4

ABSTRACT

Despite the concern that student plagiarism has become increasingly common, there is relatively little objective data on the prevalence or determinants of this illicit behavior. This study presents the results of a natural field experiment designed to address these questions. Over 1,200 papers were collected from the students in undergraduate courses at a selective post-secondary institution. Students in half of the participating courses were randomly assigned to a requirement that they complete an anti-plagiarism tutorial before submitting their papers. We found that assignment to the treatment group substantially reduced the likelihood of plagiarism, particularly among student with lower SAT scores who had the highest rates of plagiarism. A follow-up survey of participating students suggests that the intervention reduced plagiarism by increasing student knowledge rather than by increasing the perceived probabilities of detection and punishment. These results are consistent with a model of student behavior in which the decision to plagiarize reflects both a poor understanding of academic integrity and the perception that the probabilities of detection and severe punishment are low.

Thomas S Dee
Department of Economics
Swarthmore College
Swarthmore, PA 19081
and NBER
dee@swarthmore.edu

Brian A. Jacob
Gerald R. Ford School of Public Policy
University of Michigan
735 South State Street
Ann Arbor, MI 48109
and NBER
bajacob@umich.edu

1. Introduction

Several high-profile scandals involving plagiarism by prominent writers have recently captured the public's attention (McGrath 2007). Similarly, the dramatic technological progress of the last two decades (e.g., the widespread availability of full-text resources and "cut and paste" word processing) has fueled the impression that this illicit behavior has become both substantially easier and more common among students (e.g., Rimer 2003). Posner (2007) has recently argued for a nuanced understanding of plagiarism in which the gravity of an offense (and, correspondingly, the appropriate policy response) depends on the extent to which the misrepresentation imposes negative externalities (i.e., a "detrimental reliance"). By this criterion, plagiarism committed by students is a particularly serious offense. Students who plagiarize without punishment are likely to benefit at the expense of their peers with respect to valued outcomes such as grades and recommendations.¹

Unsurprisingly, plagiarism is strongly proscribed in the guidelines of most post-secondary institutions. However, when used in isolation, these stern institutional statutes may be misaligned with the behavioral context that influences the decision to plagiarize. Student plagiarism is a fairly unique type of illicit behavior in that it may often reflect a type of "rational ignorance."² The available evidence suggests that college students do not have a clear understanding of what constitutes plagiarism and how it can be avoided (e.g., Power 2009, Howard and Davies 2009). Furthermore, the incentives for students to educate themselves about plagiarism may be fairly weak given that college instructors tend to put little effort into detection and to prosecute cases of detected plagiarism informally and lightly (e.g., Schneider 1999, McCabe, Treviño, and Butterfield 2001).

¹ As a counterexample of plagiarism that is less problematic, Posner (2007) outlines how Shakespeare's Antony and Cleopatra borrows liberally but elegantly from Sir Thomas North's translation of Plutarch. Posner (2007) suggests that such borrowing is a pedigreed form of creativity that does not create material harm, concluding "if this is plagiarism, we need more plagiarism."

² Downs (1957) seminal public-choice research popularized the term "rational ignorance" in reference to voters who choose to obtain low levels of civic knowledge because of the non-trivial costs of gaining knowledge and the vanishingly small likelihood that their vote would actually be instrumental in decision-making. Ironically in the context of this study, Mueller (1989, page 205) notes that this idea was also present in the earlier writings on democracy by Joseph Schumpeter. An alternative example of an illicit behavior that may reflect an endogenously determined level of ignorance is tax evasion.

Despite the growing concerns about plagiarism in higher education, there is little objectively measured descriptive evidence on its prevalence and character.³ Furthermore, policy interventions designed to understand and reduce the prevalence of plagiarism have not been rigorously evaluated. This study presents the results of a “natural field experiment” (Harrison and List 2004) designed to provide evidence on these questions.

More specifically, we conducted a group-randomized trial in undergraduate social-science and humanities classes at a single, selective post-secondary institution during the fall 2007 semester. Each of the participating classes had a Blackboard web site that provided students with access to course materials. All of the students in the participating courses (i.e., 573 students who wrote a total of 1,256 papers) were unaware that they were participating in a research study but were required by their instructors to submit their writing assignments electronically through these Blackboard-based course web sites. We paired the 28 participating courses on baseline traits (e.g., the same instructor) and randomly assigned courses within these pairs (i.e., “block randomization”) to treatment and control conditions. The classroom-level treatment consisted of requiring students to complete a short but detailed and interactive Blackboard-based tutorial on understanding and avoiding plagiarism. Students in the treatment courses were encouraged to complete the treatment early in the semester but were *required* to do so before they were allowed to upload their writing assignments. We collected searchable, electronic files of the papers from all of the participating courses and analyzed these papers using proprietary plagiarism-detection software (i.e., Turnitin.com).

Our results indicate that plagiarism occurred in 3.3 percent of the papers from courses randomly assigned to the control condition. The type of plagiarism in these papers was almost exclusively of the “mosaic” variety (i.e., using sentences or sentence clauses without attribution). While the overall prevalence of plagiarism was low, we found that it was largely concentrated among males and students with lower SAT scores. The pattern with regard to SAT scores suggests that plagiarism may be substantially more

³ Some evidence suggests that student plagiarism is quite common. For example, in a recent survey fielded to over 60,000 students at 83 different campuses in the United States and Canada, over a third of students admitted to paraphrasing or copying a few sentences from a print or Internet source without attribution, sometime in the past year (McCabe 2005). However, the reliability of much of this evidence has been questioned because it is based on student self-reports of a behavior which they may not understand well (Brown and Emmett 2001, Power 2009).

common among the typical college student. Based on our data, one would predict a plagiarism rate of 17.7 percent among students at the national mean of SAT scores (i.e., 1017) and 31.4 percent among students at the 25th percentile of SAT scores (i.e., 850). These figures are roughly consistent with prior self-reports on the prevalence of plagiarism.

We find that random assignment to the web tutorial reduced instances of plagiarism by roughly 2 percentage points overall (i.e., a two-thirds reduction) and that this treatment effect was concentrated among students with lower SAT scores. The results of an ex-post survey and quiz completed by participating students suggest that the treatment was effective in large part because it increased student awareness about what constitutes plagiarism and knowledge of effective writing strategies. We find much weaker evidence that the intervention altered student perceptions about the likelihood of detection and/or sanctions associated with detection. Our results indicate that an easily replicable, scalable, and virtually costless educational intervention can be highly effective at reducing the prevalence of student plagiarism. Furthermore, because the treatment was particularly effective among students with lower SAT scores, these inferences may have good external validity for post-secondary institutions that are less selective than the institution in this study.

The remainder of the paper proceeds as follows. In section 2, we review the literature on student plagiarism. In section 3, we describe our field experiment in greater detail. Section 4 discusses our methodology, including how we handle several key analytical concerns. Section 5 presents our results, and section 6 concludes.

2. Literature Review

This section provides a selective overview of the multidisciplinary research on plagiarism by students.

2.1 Prevalence and Trends

Plagiarism by students is widely thought to be common, particularly with the diffusion of Internet access to full-text resources and word-processing software (Rimer 2003). Consistent with this view, a recent study by Baruchson-Arbib and Yaari (2004)

found that undergraduates are less likely to view plagiarism from resources that are available on-line as a form of academic dishonesty. However, the available data on plagiarism, which is based largely on student self-reports from surveys with distinctive design features, do not appear to provide clear evidence on the prevalence or trend in student plagiarism.

An influential early survey of students from 99 colleges and universities during the spring of 1963 (Bowers 1964) found that 28 percent of students indicated that they had plagiarized from published materials at least once since entering college. More recent surveys of college students suggest both higher and lower levels of plagiarism (e.g., McCabe and Treviño 1997, Scanlon and Neumann 2002, McCabe 2005).

A number of factors complicate any comparisons of the extant survey data. One is that the available data are not based on a consistently defined sampling frame of post-secondary institutions and often have quite low response rates. Differences in the design of the salient survey questions are another complicating factor (e.g., misconduct defined since entering college or over the past year). The conflicting data on the prevalence of academic misconduct may also reflect other survey artifacts. For example, Brown and Emmett (2001) examined data from multiple studies of student dishonesty and found that the number of student practices included in the study was related to the overall level of student cheating. Another more subtle complication suggested by McCabe, Treviño, and Butterfield (2001, page 221) is that comparisons of self-reports by students may understate the growth in plagiarism because of shifting norms about what actually constitutes an inappropriate use of source texts.

These concerns about the subjectivity of self reports and low response rates underscore the usefulness of this study's data, which are based on objective analyses of the papers submitted by the universe of students in the participating courses. At least two other studies have provided similar evidence on the prevalence of plagiarism, though on a smaller scale. Lau et al. (2005) collected the two papers written by 328 students in two sections of a psychology class and analyzed them using the plagiarism-detection service, Turnitin.com. This software generates a "similarity score", which indicates how much of a paper's text matches the text catalogued in their databases. They found that over 20

percent of papers had similarity scores above two percent.⁴ Bilic-Zulle, Azman, Frkovic, and Petrovecki (2008) used different plagiarism-detection software (WCOPYFIND) to identify the pre-intervention prevalence of plagiarism among 111 medical students in Croatia during 2000 and 2001. The writing assignment in question was based on known source articles and the software flagged text as plagiarized if six or more consecutive words matched the source text. Their baseline results indicated a considerable amount of plagiarism: 66 percent of student papers had a matched-text rate that exceeded 10 percent.

2.2 Individual and Contextual Determinants

The extant literature also suggests that a number of individual and contextual traits exert an important influence on the likelihood of committing plagiarism. For example, studies based on student self-reports generally indicate that younger students, males, those engaged in more extracurricular activities and those with weaker academic performance are more likely to engage in academic misconduct (e.g., McCabe and Treviño 1997, Genereux and McLeod 1995). Several studies also examine the relationship between psychological traits and academic misconduct. For example, Lau et al. (2005) find that a measure of psychopathy along with low verbal ability were strongly predictive of an objective measure of plagiarism.

In a recent synthesis of the broader psychological literature on academic misconduct, Murdock and Anderman (2006) emphasize the role of three motivational mechanisms. First, students whose academic motivation is extrinsic rather than intrinsic are more likely to cheat. Second, students are more likely to cheat when they have low expectations of their capacity to reach their academic goals through personal effort. Third, students who view the potential costs of cheating as minimal are more likely to cheat. Posner (2007, page 89) also stressed the role of self-interest noting that students "plagiarize to save time, to get better grades, or both; the effect on learning and evaluation is significant and punishment often and appropriately severe."

⁴ This may actually overstate the true prevalence of plagiarism. We use Turnitin.com to detect plagiarism in our study as well. However, our review of the similarity scores indicated that they contained a significant number of false positives. Therefore, we relied on multiple-rater review of the underlying "originality reports" to identify unambiguous instances of student plagiarism.

A review article by McCabe, Treviño, and Butterfield (2001) underscored the empirical relevance of related but broader contextual factors such as the perception of cheating by peers and the perceived severity of penalties. The literature on the apparent relevance of contextual factors lends credence to the consensus view that "cheating can be most effectively addressed at the institutional level" (McCabe, Treviño, and Butterfield 2001). Another complicating factor that supports this view is the somewhat indifferent role that may be played by faculty. College instructors do not generally view either educating students about avoiding plagiarism or policing instances of plagiarism as a primary responsibility. Faculty also tend to deal with academic misconduct they may uncover informally and in a manner that is perceived by students as fairly lenient (e.g., Schneider 1999, McCabe, Treviño, and Butterfield 2001).

2.3 Evaluation Studies

There is relatively little research that evaluates specific policies or interventions designed to reduce student plagiarism. One prominent institutional policy that has received some attention is an honor code. Basic cross-sectional comparisons suggest that honor codes are fairly effective. For example, McCabe, Treviño, and Butterfield (2001) report that, in a study of students at 31 selective colleges and universities during the 1995-96 academic year 43 percent of students at institutions without honor codes admitted to copying a sentence or two without attribution while at institutions with honor codes the corresponding prevalence was 32 percent.

Three other small-scale studies provide similarly encouraging evidence on the potential efficacy of targeted classroom and institutional policies. Bilic-Zulle, Azman, Frkovic, and Petroveckii (2008) compared the plagiarism rates among different cohorts of Croatian medical students and found that it dropped dramatically (i.e., from 66 percent to 11 percent) after the introduction of warnings about the penalties for plagiarism and making it known that plagiarism-detection software would be used. Second, a similar pre-post comparison based on a sample of students at San Diego State University suggests that completing a web tutorial raised student performance on a plagiarism quiz from 85.6 percent to 91.6 percent (Jackson 2006). Third, a lab-experimental study provides evidence consistent with the evidence from cross-sectional and time-series comparisons

about the importance of contextual factors. Specifically, in a lab study based on psychology students at a university in the United States, Rettinger and Kramer (2009) experimentally manipulated student exposure to vignettes about peer attitudes and behavior regarding academic misconduct and found that this manipulation influenced the stated likelihood that they or a protagonist would engage in such behavior.

However, the extant literature has a number of important shortcomings that the design of this study addresses directly. First, this study provides evidence on the prevalence and characteristics of student plagiarism using objective measures from a large number of actual student papers rather than potentially unreliable student self-reports. Second, this study provides evidence on the efficacy of a web-based tutorial that seeks to educate students about what constitutes plagiarism and effective research and writing strategies for avoiding it. This is a distinctive approach in that it does not rely directly on the perceived threats of detection but rather on enhancing human capital. Given the stylized evidence that students with low measures of prior achievement are more likely to engage in academic misconduct, this may be a particularly effective margin on which to improve student outcomes. Similarly, policy recommendations often underscore the need for a "hidden curriculum" that informs students about academic ethics (e.g., McCabe, Treviño, and Butterfield 2001). The tutorial evaluated in this study is one quite explicit way to operationalize such a curriculum. Furthermore, the format of this tutorial is such that it could be replicated and scaled up in different institutions at little cost.

Third and perhaps most important, by utilizing random assignment in a field setting, the evaluation evidence presented here is credibly robust to concerns about internal validity. The available inferences based on cross-institutional comparisons of student surveys could be seriously confounded by non-random student selection into schools with particular traits as well as by omitted variables. However, the random assignment procedures in this study should balance the potentially confounding unobservables across the treatment and control conditions and allow the effect of the tutorial to be reliably identified.

3. A Natural Field Experiment

The setting for our field research is a single, highly selective post-secondary institution in the United States. Specifically, we collected and analyzed electronic versions of anonymized student papers from 28 undergraduate social-science and humanities courses during the fall 2007 semester. The collection of student papers occurred largely through the Blackboard classroom-management web page for each course and the participating students were unaware of the study's existence. As part of the human-subject protocols for this research project, we do not identify the participating institution and all student papers were anonymized prior to analysis.

3.1 Study Recruitment and Randomization

We began the recruitment of courses by identifying all the social-science and humanities courses offered during the Fall 2007 semester. We excluded quantitative-methods courses, small-scale seminars, and research colloquia as well as independent study and thesis-related courses. We approached the instructors for 46 classes that had comparatively large enrollments and solicited their participation in a campus-wide study on student writing. The motivation for our emphasis on larger classes (i.e., typically 18 or more students) was both increased statistical power for our research effort and a potential increase in the external validity of our inferences for institutions that, on average, have larger class sizes than the participating institution. To complement the block-randomization strategy we describe below, we also recruited courses with somewhat smaller enrollments in situations where those courses were taught by the instructor of another recruited course.

The instructors were asked whether they would be willing to include their course in an IRB-approved, field-research project on the characteristics of student writing. They were told that participation would not involve any substantive change in their course. Participation would simply require using Blackboard's classroom-management software to collect student papers electronically and to provide students with information on their writing assignments. To encourage participation, the research team made it clear that they would design and manage this aspect of Blackboard as well as provide participating instructors with printed or electronic versions of all their submitted writing assignments.

The instructors for 9 of the 46 recruited courses provided no responses to recruitment queries. Four other courses had no valid writing assignment. The instructors for 5 additional courses refused participation.⁵

The remaining 28 courses were randomly assigned to treatment and control conditions. Courses in the control state merely had students use Blackboard to submit their writing assignments. In courses assigned to the treatment state, students also submitted their writing assignments through Blackboard. However, before they were allowed to do so, they also had to complete a Blackboard-based tutorial and quiz on plagiarism. This intervention is described in more detail below.

Our course-based randomization avoids the contamination that might have occurred if students within the same courses had been randomly assigned to the treatment. However, a potential drawback of randomizing over only 28 units is that the treatment and control courses might not be balanced with regard to observed and, more important, unobserved baseline traits. To reduce this possibility, we employed a simple block randomization strategy, pairing participating courses on baseline traits and then randomizing within those pairs.

In an ideal situation, we would be able to match each course to another course with a similar propensity for student plagiarism by using baseline traits that are highly predictive of the prevalence of plagiarism. Unfortunately, reliable baseline variables of this sort are unavailable in this context. However, in light of the prior evidence on the importance of contextual factors (McCabe, Treviño, and Butterfield 2001), we conjectured that the likelihood of plagiarism would be related to the many unobservables associated with particular instructors (e.g., types of writing assignment, writing support, and the apparent threat of detection) and with particular academic disciplines. Our review of the syllabi and writing assignments for the participating courses provided some confirmation for these priors (e.g., the presence of a plagiarism warning on the syllabus and the extent to which writing assignments involved a student's response to instructor-chosen source material as opposed to researching a topic through self-identified references).

⁵ One instructor provided no reason for refusing while a second instructor was uncomfortable with using Blackboard despite the facilitation by the research team. Three other instructors refused because they were uncomfortable with the "deception" of students, despite the data-security protocols.

Based on the available baseline information about the courses and their instructors, we paired courses prior to randomization in the following manner. First, for 12 of the participating courses, we were able to form pairs among courses taught by the same professor in the same department (and, in 6 of these cases, randomization was also within sections of the same course). Second, for 10 other courses, randomization occurred among courses taught in the same department. In cases where there were multiple courses from a given department, we paired courses that had similar writing assignments as indicated by the syllabi (e.g., research content of the assignments and the presence of a plagiarism warning). The remaining 6 courses were paired to another course in the same academic division (i.e., social sciences or humanities) using the same data on the character of the writing assignments.

3.2 The Treatment

In the courses assigned to the treatment, students were required to complete a Blackboard-based tutorial on understanding and avoiding plagiarism. The tutorial was adapted for Blackboard from resources available at the Plagiarism Resource Site (<https://ats.bates.edu/cbb/>) developed by staff at Colby, Bates and Bowdoin Colleges.⁶ The tutorial required students to click through 18 sequential screens with text that defined different forms of plagiarism. This tutorial also provided explicit examples of what constitutes plagiarism by showing side-by-side examples of source material along with examples of the correct and incorrect use of that material in a student paper. The tutorial also outlined effective strategies for avoiding plagiarism (e.g., not procrastinating and careful note-taking). At the end of this sequence of material, students completed a nine-question quiz consisting of several detailed and example-driven questions on plagiarism. Each response triggered detailed feedback on why that answer was either correct or incorrect before proceeding to the next question. Appendix A contains several illustrative screenshots of this tutorial.

⁶ We secured permission for the use of this material, which was also available for sharing and adaptation under a Creative Commons license. To avoid unintended irony, the tutorial clearly made an attribution to its source. Consistent with the license conditions, our use of this material was non-commercial and our adaption of this material is available for sharing upon request.

This intervention was deployed on the Blackboard sites of the treatment courses at the beginning of the third week of the semester (i.e., immediately after the date at which students could drop the course with no record of having been enrolled). The Blackboard sites made it clear that students would not be allowed to upload their completed writing assignments (i.e., the upload mechanism would not activate) until students had completed the tutorial. However, because of the role that early research and note-taking can play in unintentional plagiarism, the instructors within the participating treatment courses were encouraged to email students early in the semester about the need to complete the Blackboard tutorial. They were also provided with the names and email addresses of students who had not promptly completed the tutorial and encouraged to provide targeted follow-up reminders. As a result of this effort, there are no empirically meaningful distinctions between the effects associated with the intent-to-treat and the effect of the treatment-on-the-treated in this study. Over 97 percent of students in the treatment courses fully completed the tutorial while an additional 1.5 percent partially completed the tutorial.⁷

An interesting and important feature of this study is that the participating students were not aware that they were participating in a research study.⁸ However, as noted earlier, the participating instructors did know that their courses were involved in a writing study. In theory, the general awareness among instructors that their student papers were being externally evaluated in some way may have muddied the treatment contrast by encouraging all instructors to manage these assignments in a manner that reduced plagiarism. The existence of sizable treatment effects suggest that, if there were any effects associated with this general awareness, they were not empirically confounding. Several of the participating instructors were also clearly aware of the broad intent of the

⁷ Students who exited the tutorial before completing its entire sequence were allowed to upload their writing assignments. A small number of students may not have completed the tutorial at all because they dropped the course or because they submitted hard copy papers directly to the instructor. We collected all available hard copies for our analysis and assess the implications of study attrition for our key inferences.

⁸ Nonetheless, the use of electronic paper collection as opposed to printed copies could conceivably constitute a study-wide deterrent to plagiarism. In theory, this could compromise the external validity of our results for papers that are collected as printed copies. And the fairly low prevalence of plagiarism in our study suggests this caveat. If such an effect existed it could also muddy our treatment contrast, which would bias us towards finding no effect of the intervention. However, given the magnitude of the apparent treatment effects in this study, this seems less problematic. Furthermore, we suspect that the electronic submission of papers is an increasingly common mechanism.

research study (i.e., the focus on plagiarism). Four participating instructors asked to be aware of the study goals as a condition of participation and a fifth clearly inferred the study goals because he or she actively managed other components of their Blackboard sites. Fortunately, our block-randomization strategy implies that there is uniform treatment and control variation within these instructors. Three of these instructors each taught two of the participating courses so they were paired with other courses they taught prior to randomization. The remaining two instructors taught courses in the same department and these courses were also paired with each other. This pattern of pairing implies that any effects that might be associated with an awareness of the study's focus should again create an attenuation bias in the estimated treatment effects.

3.3 Identifying Plagiarism

We relied on the proprietary web service, Turnitin.com, to analyze the participating papers for plagiarism.⁹ For each submitted paper, Turnitin.com generates a "similarity score" that identifies the percentage of submitted text that matches their continually updated database of journal articles, newspapers, magazine articles, books, and web pages. An "originality report" also makes it possible to connect suspicious text to the potentially plagiarized source. At a pilot stage for this project, we compared the performance of this service to that of other available software and found that it was particularly discriminating both with respect to its extensive database and with regard to identifying plagiarized text that may have been lightly edited.

Settings for the originality reports allowed most quoted text and bibliographies to be ignored in generating similarity scores. Nonetheless, our review of the similarity scores from the papers in this field experiment indicated that a large share of the highest similarity scores reflected false positives. This occurred when the software failed to recognize correct citations of quoted material and when it flagged oddly formatted bibliographies as plagiarized text. For example, some high similarity scores occurred when a paper legitimately quoted other text but used margin offsets instead of quotation

⁹ As part of the human-subject protocols for this research, all of the collected papers were assigned random identifiers and anonymized (e.g., names removed from file name, paper titles and headers) prior to analysis.

marks. We also found that some high similarity scores were simply due to the accumulation of common word fragments used throughout a given paper.¹⁰

Given the pervasive amount of measurement error in the similarity scores generated by Turnitin.com, we adopted a straightforward rating strategy using multiple reviewers. We first reviewed the papers with high similarity scores (i.e., 15 or higher). Roughly half of these had clearly plagiarized content while the remaining papers appeared to be false positives exclusively. We then reviewed each paper with a similarity score between 11 and 15. Of these papers, roughly one third had plagiarized content. We then reviewed the papers with similarity scores between 8 and 10. Only 16 percent of these papers were judged to have plagiarized content. As we moved to (and through) the third strata, the probability of having identifiable plagiarism clearly dropped. Furthermore, within the lower stratum, the extent of plagiarism in papers with plagiarized content was substantially lower. Interestingly, our exhaustive review of these papers indicated that the plagiarism that did occur was predominately of the "mosaic" variety (e.g., copied sentences, sentence clauses and phrases).

Our analysis focuses on a binary dependent variable that indicates whether a paper had plagiarized content and was in the two rating strata defined by similarity scores of 11 or higher. This focus reflected a judgment that the plagiarism that occurred in the highest two strata was distinctly more consequential in scale. However, as robustness checks, we also present results based on binary indicators for more and less restrictive measures of plagiarism (i.e., papers identified as having plagiarized content with similarity scores of 15 or higher and 8 or higher, respectively).

3.4 Data Description

The 28 participating courses had collective enrollment of 697 students.¹¹ The writing assignments in these 28 courses and the corresponding course enrollments implied that there were 1,329 potential papers to be collected. Because attrition from the collection of papers is a potential threat to both the internal and external validity of our

¹⁰ In his recent book on plagiarism, Posner (2007, page 84) discusses Turnitin.com and notes it generates false positives because of indented quotations and the flagging of incidental phrases.

¹¹ However, because some unique students were enrolled in more than one participating course, there were 573 unique students in the study.

study, we made an aggressive effort, in cooperation with the participating instructors, to obtain physical copies of papers that were submitted as print outs rather than through the web-based upload mechanisms. More specifically, nearly 6 percent of the potential papers (i.e., 79 of 1,329) were obtained as printouts. Through the use of scanning and optical character recognition (OCR) software, we were able to convert these papers to searchable text and include them in our analysis.

Our final, analytical sample consisted of 1,259 papers, implying a fairly low attrition rate of 5.3 percent (i.e., 70/1,329). The attrition of these 70 papers was due in part to students who withdrew from courses or had taken a grade of incomplete ($n = 19$). The remaining papers ($n = 51$) were either not submitted or were submitted directly to the instructor as printed copies that we could not obtain. As we discuss below, the differences in attrition across treatment and control classrooms were small and statistically insignificant.

We were able to identify a number of student traits (e.g., race, gender, SAT scores) through access to the institution's administrative data (Table 1). The composite (math and verbal) SAT scores were imputed for those who only had ACT composite data using a concordance table available from the College Board. We were also able to identify other student traits (e.g., class status, pass/fail status) from the class enrollment data.¹² We also identified several class-level observables (e.g., class size, the presence of a plagiarism warning on the syllabus, the number of required papers for the course, the academic rank and gender of the instructor) that may be relevant determinants of student plagiarism.

4. Empirical Strategy

The randomized nature of the field experiment alleviates many of the common selection concerns associated studies of plagiarism, and allows for a straightforward analysis of the data. We estimate variants of the following OLS regression:

$$(1) \quad y_{ic} = \beta T_c + \Gamma X_i + \Pi C_c + \alpha_c + \varepsilon_{ic}$$

¹² A small number of potential papers ($n = 31$), only one of which was lost to attrition, were from students who were taking a course at the participating institution but were not enrolled there. Because of their unique enrollment status, some data (e.g., SAT scores) were unavailable for these students. However, their gender was accurately identified from their first names and other public sources.

where i denotes individuals and c denotes classrooms.¹³ T_c is a binary indicator for whether the class was in the treatment group, X_i is a vector of student characteristics and C_c is a vector of classroom characteristics. The α_c term is a classroom-specific error term that will be a concern in properly estimating the precision of our regression estimates, which we discuss in greater detail below.

In this section, we discuss three issues of particular concern in cluster-randomized trials such as this: 1) treatment-control balance, 2) sample attrition and 3) proper treatment of the clustered nature of our data for the purposes of statistical inference.

4.1 Treatment-Control Balance

In expectation, the randomization of classrooms to treatment and control conditions will ensure that all observable and unobservable characteristics of students and classrooms are balanced across the two groups. In small samples, however, it is possible for a specific realization of random assignment to result in poor balance. Our block-randomization strategy was explicitly designed to avoid this potential problem. Nonetheless, it is still important to explore the realized balance of baseline traits across the treatment and control conditions.

Table 1 presents summary statistics that speak to this concern. The first column shows sample means for the full sample. We can see that the sample contains very high-performing students (i.e., an average SAT score of 1407), who are typical of the participating institution. There are an equal proportion of males and females, and reasonable distribution of different race/ethnicity types, with 9 percent African-American, 11 percent Hispanic, 20 percent Asian and 54 percent Caucasian. Importantly, 21 percent of students do not indicate a race/ethnicity on school records. The majority of the sample is composed of freshman and sophomores (21 and 33 percent respectively). All freshmen take their fall courses pass/fail and roughly 6.8 percent of other students are taking pass/fail courses in our data. A bit less than 30 percent of classes included some warning about plagiarism in the syllabus. Female professors taught roughly one-third of the courses in our sample.

¹³ Here we abstract away from student-by-paper as the unit of observation.

Columns 2 and 3 present classroom-level summary statistics for the treatment and control groups. In none of these comparisons do we find that the treatment-control differences are statistically significant. This is striking because, in conducting these “multiple comparisons,” one might expect to sometimes reject some null hypotheses of no difference, even when the null hypotheses are true. Various procedures (i.e., Bonferroni and Benjamini-Hochberg) are designed to correct for the Type I errors in multiple comparisons. These corrections would only imply that the p-values in Table 1 are even larger.

However, multiple-comparison procedures that also allow for a joint error structure across these comparisons might create a more powerful test of the treatment-control differences in baseline covariates. As an alternative way to test whether we are likely to observe this distribution of covariates under the null hypothesis of random assignment, we conduct a permutation test analogous to Fischer’s exact test.¹⁴ To do so, we conduct a 1,000 replications in which we randomly assign treatment status to 14 of the 28 classrooms (keeping the covariates in the classrooms fixed as they are in reality). We then run a seemingly unrelated regression with each of the covariates as outcomes and a treatment indicator as the single predictor in each equation. We obtain the F-statistic from a joint test of the null hypothesis that the coefficients on the treatment indicator for all equations are equal to zero. We then ascertain at which point the true F-statistic we obtained in our sample would fall in the distribution of the 1,000 F-statistics we obtained from our permutations. The p-value from this exercise indicates that we cannot reject the null hypothesis of random assignment. We also conducted an identical permutation test, but instead of running a SUR, we estimate a single equation OLS model in which the treatment indicator is the dependent variable and all 20 of the covariates are predictors.

In summary, we cannot reject the null hypothesis of random assignment, and thus it appears that our covariates are reasonably well balanced. However, the results in Table

¹⁴ Given the small number of observations (i.e., 28) and the relatively large number of covariates we have (i.e., 20), many standard regression techniques that rely on asymptotic results do not work. For example, the SUR and OLS regression approaches described above failed miserably in simulation exercises. We generated a test data set with 28 observations and 20 covariates, which were drawn at random but were set to match the means and covariances of the 20 covariates in our actual data. We then randomly assigned treatment status to 14 of the 28 classrooms, and estimated the SUR and OLS models described above.

I suggest that there were some noticeable treatment-control differences in classroom-level traits. For example, treatment classrooms were substantially more likely to have a female instructor and less likely to be taught by a full professor. We examine the robustness of our impact estimates to these statistically insignificant differences through regression adjustments for classroom observables.

4.2 Sample Attrition

Even if randomization results in good balance across treatment and control groups, differential sample attrition between the conditions may still result in a biased estimate of the treatment effect. For example, if students in treatment classrooms were more likely to drop the class when they feel uncertain about their writing skill, the result may be that control classrooms have a disproportionate fraction of good writers who may have a lower propensity to plagiarize even in the absence of the treatment. This dynamic would lead our empirical strategy to underestimate any beneficial impact of the treatment.

To test for the presence of differential attrition, we estimate specifications similar to equation (1) where the outcome is a binary indicator for whether we have any outcome data for the student-paper observation. We estimate a variety of different variations on this basic specification, and in no case does assignment to the treatment group have a statistically significant or substantively important impact on attrition. Hence, sample attrition does not appear to be a concern with respect to the internal validity of our results. Furthermore, the low level of attrition also suggests that its implications for the external validity of this study are negligible.

4.3 Estimation and Statistical Inference

While our final analysis sample contains over 1,200 student-paper observations, the treatment was randomly assigned across only 28 classrooms. As others have pointed out, the nested structure of the data has important implications for accurately estimating the precision of the treatment effects and conducting statistical inference (Bertrand et al. 2004, Cameron et al. 2008, Donald and Lang 2007, Angrist and Pischke 2009). Specifically, statistical inference must take into account the within-group dependence in

the data. A common approach is to report cluster-robust standard errors that generalize the White (1980) heteroskedasticity-consistent estimates of OLS standard errors. Such cluster-robust standard errors provide consistent estimates as the number of clusters goes to infinity. In practice, however, many applied studies use samples with a small number of clusters.

Several recent papers demonstrate that cluster-robust standard errors may not be consistent when the number of clusters is as small (Cameron et al. 2008, Donald and Lang 2007). More importantly, the direction of the bias generally leads one to over-reject the null hypothesis. In the analysis below, we present a several alternative estimates suggested in the recent literature. There are two broad approaches we pursue.

The first strategy utilizes group-level data. In our case, this means that we will collapse our data to the classroom level and estimate specifications like the following using the 28 classroom-level observations:

$$(2) \quad y_c = \beta T_c + \Pi C_c + \varepsilon_c$$

where y_c is the rate of plagiarism in classroom c . For the purpose of inference, we calculate bias-adjusted robust standard errors to account for heteroskedasticity. The bias adjustment we use is HC2 (described in Angrist and Pischke (2009)) and is meant to adjust for the finite-sample bias of the commonly used "robust" (White 1980) standard errors.

A variant of this group-data approach is a two-step procedure that allows us to incorporate the student-level covariates we have in an effort to gain greater precision. In the first step, we estimate

$$(3) \quad y_{ic} = \Gamma X_i + \mu_c + \eta_{ic}$$

where μ_c provide estimates of the covariate-adjusted group effects, in our case the adjusted plagiarism rate in each classroom. In step two, we regress these adjusted group effects on a set of classroom-level variables, which can include our treatment indicator as well as other classroom covariates (and pair effects):

$$(4) \quad \hat{\mu}_c = \beta T_c + \Pi C_c + (v_c + (\hat{\mu}_c - \mu_c))$$

We show GLS estimates of equation (4) that use the inverse of $\text{var}(\hat{\mu}_c)$ from equation (3) as weights. We also report estimates of equation (4) that are unweighted and

a third set that are weighted based on the number of student-papers within each classroom. Following the suggestion of Donald and Lang (2007), we use the t-distribution with $C-K$ degrees of freedom (where C is the number of group-level observations and K is the number of regressors) to conduct inference on estimates from all group-level models.

Our second broad approach directly utilizes the micro (i.e., student-level) data. One of the virtues of using the micro data is that it facilitates identifying heterogeneity in treatment effects by student-level traits. However, our approach still needs to account for the nested structure of the data and the relatively small number of classrooms in our sample. To do so, we calculate and report bias-corrected clustered standard errors using the method proposed by Bell and McCaffrey (2002). This procedure, called bias-reduced linearization or BRL, is essentially a generalization of the HC2 correction for the case of clustering. In recognition of the within-group dependence and the small number of clusters, we conduct inference based on a t-distribution with $C-K$ degrees of freedom despite the fact that the estimation utilizes student-level observations.¹⁵

Furthermore, we also show results from the bootstrap-based approaches recommended in Cameron et al. (2008). These authors propose cluster bootstrap-t procedures to improve inference in cases with a small number of clusters. Bootstrap estimates of a t-statistic provide "asymptotic refinement" because the asymptotic distribution of the t-distribution does not depend on any unknown parameters (unlike regression coefficients, whose asymptotic distribution depends on the unknown residual variance). In a bootstrap-t procedure, one calculates a t-statistic for each bootstrap sample, and compares the t-statistic from the original sample to the distribution of t-statistics from the bootstrap replications. If the absolute value of the original t-statistic is above the 95th percentile of the absolute values from the bootstrap distribution, one rejects the null hypothesis at the 5 percent level. While this approach provides some efficiency gains, it does not yield standard errors, which might be of independent interest, for example to calculate a confidence interval. On the basis of Monte Carlo simulations,

¹⁵ Angrist and Pischke (2009) recommend using the maximum of robust and conventional standard errors for inference since robust standard errors can be subject to considerable sampling variance. In practice, the BRL standard errors are virtually identical to the conventional standard errors, both of which are larger than the standard cluster-robust standard errors. In discussing our results, we present the BRL standard errors, effectively adopting the conservative rule of thumb recommended by Angrist and Pischke (2009).

Cameron et al. (2008) recommend using a wild-cluster bootstrap rather than a simple block bootstrap. The wild-cluster bootstrap re-samples residuals while holding the regressors fixed. A key advantage of this approach is that it avoids bootstrap replications in which β or $\text{var}(\beta)$ are inestimable, as can happen more frequently with a small number of clusters when the treatment varies exclusively at the cluster level.¹⁶

5. Results

We begin by presenting in Table 2 some descriptive evidence on the relationship between various student and class characteristics and the prevalence of plagiarism. For this exercise, we use our main plagiarism indicator (i.e., papers with similarity scores of 11 or higher that were rated as plagiarism by multiple raters) and simply note that the results are qualitatively the same using the other measures. Column 1 presents estimates from a series of bivariate OLS regressions that model an indicator for plagiarism as a function of a single student or classroom characteristics. The standard errors for the student level regressors are clustered by student whereas the standard errors for the classroom regressors are clustered by class with a bias-reducing linearization (BRL) adjustment.¹⁷ Column 2 presents estimates from a single regression model in which all of the student and classroom characteristics shown are entered jointly as regressors. Hence, column 1 shows the unconditional relationship between a particular regressor and the outcome whereas column 2 shows the conditional relationship between the regressor and the outcome.

Several interesting patterns emerge. In column 1, we see that African-American and Asian students are more likely to plagiarize than other students while students with higher SAT scores are less likely to plagiarize. Indeed, the relationship between SAT score and plagiarism appears convex. The joint model in column 2 indicates that there is still a significant relationship between SAT score and plagiarism, even after controlling for other factors. The large positive coefficient on African-American students disappears, but the positive effect for Asian students remains significant.

¹⁶ We are very grateful to Doug Miller for providing STATA code we use to implement the wild-cluster bootstrap method used in Cameron et al. (2008).

¹⁷ More specifically, the standard errors reflect cluster adjustments and a "bias reducing linearization" (BRL) adjustment, which we discuss below.

The relationship between SAT score and plagiarism is quite strong. The bottom quintile of students at the school, who scored between 1,000 and 1,200 on the SAT, had plagiarism rates of nearly 14 percent. Using the estimates from column 2 and extrapolating out of sample, we would predict that students scoring at the national average on the SAT (a score of 1,017) would have rates of 17.7 percent and students at the 25th percentile of national SAT scores (i.e., 850) would have a plagiarism rate of 31.7 percent. These results are consistent with prior work indicating that lower-performing students are more apt to plagiarize (McCabe, Treviño, and Butterfield 2001). And perhaps not surprisingly, students taking classes pass/fail are significantly less likely to plagiarize than those taking the course for a letter grade.

As an initial view of the treatment effect, Figure 1 shows histograms of unadjusted plagiarism rates by classroom organized so that the class pairs are adjacent to each other. The asterisks below the horizontal axis indicate pairs comprised of one instructor teaching two classes. In 11 out of 14 pairs, the average plagiarism rate in the control classroom exceeds that in the treatment classroom. In 2 out of 3 of the other cases, there were no cases of plagiarism in either classroom. Hence, the plagiarism rate in the control classroom exceeded the rate in the paired treatment class in only 1 out of 14 pairs. Simple nonparametric tests based on these means suggest that the intervention substantially reduced the prevalence of plagiarism. For example, a Wilcoxon rank-sum test for matched-pair data rejects the null hypothesis of equality across the treatment and control conditions with a p-value of .008 for a two-sided test. An analogous test using adjusted classroom means derived from a regression that includes all of the student and classroom covariates shown in Table 2 also rejects the null of equality with a p-value of .048.

5.1 Baseline Estimates

Tables 3 and 4 present parametric estimates of the treatment effect to properly quantify the magnitude and determine the statistical precision of the effect. We start with specifications that use classroom-aggregate data as a conservative estimate of the precision of our treatment effect estimates. In Panel A of Table 3, the dependent variable is the unadjusted classroom mean plagiarism rate, which ranges from zero to .31. Column

1 shows the results from a model that includes no other covariates besides the treatment indicator, weights each classroom observation equally and does not make any adjustment to the standard errors. The resulting point estimate of -0.036 suggests that the treatment reduced plagiarism by roughly 3.6 percentage points, a very large effect given the classroom-level control mean of 5.6 percent (i.e., a decrease of 64 percent). With a standard error of $.024$, however, this point estimate is not statistically different than zero. However, after controlling for either class-level observables or pair fixed effects, this impact estimate becomes statistically significant. In particular, introducing pair fixed effects has virtually no effect on the impact estimate but increases its statistical precision appreciably. Furthermore, the robust and HC2 standard errors are quite similar to the conventional standard errors in the pair fixed-effects specification.

Panel B shows the key results from the two-step estimation procedure described above, including different sets of controls and with different weights (i.e., equation (4)). Our preferred specification in column 6 includes student characteristics and pair fixed effects, and weights the second-step regression with the inverse of the covariance matrix on the classroom fixed effects in step one. The resulting point estimate of $-.023$ is smaller but statistically significant, with a p-value of $.027$. This impact estimate implies a 41 percent reduction in plagiarism relative to the control group mean.

Table 4 presents the estimated treatment effect conditional on student covariates and the pair fixed effects. Interestingly, the BRL-adjusted standard errors shown in column 3 are nearly identical to the conventional standard errors shown in column 1, both of which are substantially larger than the typical cluster-robust standard errors shown in column 2. Indeed, the standard error in column 3 is virtually identical to the GLS estimate from Table 3, Panel B, column 6. Columns 4 and 5 present results from the standard block bootstrap-t and the wild cluster bootstrap-t suggested by Cameron et al. (2008), which should provide more efficient estimates. The 95% confidence intervals of the t-statistics come from 10,000 replications of the bootstrap. The p-value reported indicates the fraction of the 10,000 replications in which the t-statistic was larger, in absolute value, than -2.684 , the t-statistic from the full sample.

Taken together, these results suggest a treatment effect that is large in magnitude, statistically significant and robust to a variety of alternative strategies to calculating the

standard errors.¹⁸ Table 5 presents several additional robustness checks. Column 1 replicates the estimates from column 3 of Table 4 as a baseline. Columns 2 and 3 show that using more and less restrictive definitions of plagiarism (i.e., similarity scores 15 or above and 8 or above, respectively) do not materially change our results. Some students in our sample appear in multiple classes. As a result, roughly 11 percent of students were simultaneously in at least one treatment and control classroom. For these students, it is possible that exposure to the treatment in one class might have influenced behavior in other classes. For this reason, the specification in column 4 assigns treatment status to all observations of students who were in at least one treatment class. The results are virtually identical to the baseline. In column 5, we limit the sample to classroom pairs in which the instructor taught both the treatment and control classes. In column 6, we re-estimate the main specification dropping all observations from the four pairs in which one of the instructors in the pair was known to be broadly aware of the objective of the study.¹⁹ While the point estimates differ slightly across these specifications, they are not significantly different than the baseline results.

Columns 7-9 show results for logit and probit models as well as OLS. With the full set of baseline student and classroom controls, many observations drop from the non-linear models. For this reason, the specifications in columns 7-9 include a limited set of covariates.²⁰ In addition, we report conventional standard errors, which Table 4 indicates are virtually identical to the BRL-adjusted standard errors. The OLS estimate in column 7 is slightly large than the baseline estimate shown in column 1, as one would expect since we drop classroom pairs with no observed cases of plagiarism (and thus no potential treatment effect). More importantly, however, the average marginal effect from the probit (column 8) and logit (column 9) models yield very similar results to those in column 7.

¹⁸ In results not reported here but available upon request, we examined the impact of the intervention on student grades in subsequent courses. We found no statistically significant or substantively important impacts.

¹⁹ While we did not systematically inform instructors of the purpose of the study, several instructors insisted on knowing as a condition of participation and others inadvertently discovered the objective while working with the electronic submission system we used.

²⁰ These specifications drop observations with missing SAT score (no variation in outcome), drop four classrooms with no variation in outcome, drop the pass-fail indicator variable (no variation), drop the indicators for assistant, associate and visiting professor (almost no variation), and combine the Hispanic and other race indicators (because the other race indicator has no variation in outcome).

This suggests that our baseline OLS results are robust to the use of alternative specifications.

5.2 Treatment Heterogeneity

Prior literature as well as the analysis shown in Table 2 suggests that the prevalence of plagiarism varies systematically with student characteristics. It thus seems likely that the impact of any particular intervention may also vary across students. Table 6 shows treatment effects separately for several key subgroups. In columns 1 and 2, we see that the intervention had a similar effect on male and female college students.²¹ Columns 3-6 show the results separately by year in college. The point estimates are roughly equivalent for all but sophomores (for which the point estimate is essentially zero). More specifically, though the treatment effects for juniors and seniors are estimated with comparative precision, the differences in treatment effects by class are not statistically significant.

To explore the relationship between initial achievement/cognitive ability and the intervention, we estimate models that allow the treatment effect to vary with a student's SAT score. Figure 2 shows the treatment effects and control mean estimated by local linear regression using a triangular kernel with bandwidth of 200 SAT points. Confidence intervals use BRL corrected standard errors. The treatment effect model includes all covariates from the primary specification described above.²²

To begin, note that the mean plagiarism rate among control students is over 10 percent for students at the bottom of the SAT distribution. The rate declines steadily as SAT score rises, asymptoting to nearly zero at the upper end of the SAT distribution. More interestingly, we see that the intervention had a much larger impact on students at the bottom tail of the SAT distribution. Based on the estimates from this model, we would conclude that the intervention reduced the likelihood of plagiarism by roughly 10 percentage points among students with SAT scores below 1,200. While these students comprise just less than 10 percent of the students in our sample, the national average of

²¹ The treatment estimate for female students becomes smaller and statistically insignificant in models that control for classroom observables instead of pair fixed effects.

²² A model that interacts SAT score with linear and quadratic terms of the SAT score yields comparable results.

math and verbal SAT scores among all test-takers in 2007 was 1,017 and the 25th percentile was 850. While the external validity of any intervention trial is open to question, this treatment heterogeneity suggests that the intervention may have a large impact on the typical college student.

In theory, it is possible that the treatment effect may have varied with classroom or instructor characteristics. Unfortunately, with only 28 classrooms and 14 pairs, our ability to detect such differences is quite limited. The results shown in Table 7 suggest that impacts may have been larger in classes with female professors, professors below the rank of full professor and in classes that did not include a warning on the syllabus regarding plagiarism. However, none of these treatment effect differences are statistically significant at conventional levels.

5.3 Human Capital or Deterrence?

Our intervention was designed to reduce the prevalence of plagiarism by educating students about what constitutes plagiarism and providing them with effective strategies for avoidance. However, it may also (or even exclusively) be that this intervention reduced plagiarism simply by increasing the perceived likelihood that plagiarism would be detected and prosecuted. To assess the mediating mechanisms by which this intervention was effective, we fielded a web-based survey of the participating students approximately one month after the end of the semester and after the collection of writing assignments for this study had concluded. The survey contained 10 questions tapping student attitudes regarding the course and the instructor, along with three true-false questions assessing the student's knowledge of plagiarism (see Appendix B). The response rate was 51 percent and did not differ significantly across treatment and control groups.²³

Table 8 presents results from analyses that examine the impact of the treatment on survey responses. More specifically, Table 8 reports the estimated treatment effects from

²³ Because of the operating constraints implied by Blackboard's survey mechanism, student identifiers were not available for the student-level survey responses within participating courses. Blackboard's design also implied that one course pair could not participate in this follow-up survey. One instructor effectively merged the Blackboard site for two sections of the same course. This did not complicate the treatment, which could be viewed only by the treatment course students in this pair. However, because Blackboard's survey mechanism strips individual identifiers, it was not possible to separate treatment and control responses for this pair of courses.

OLS regressions that condition on student traits and pair fixed effects, with BRL-adjusted clustered standard errors. The results in the first row of Table 8 indicate that, though students in treatment courses were somewhat more likely to not complete the survey, this difference was not statistically significant. In other words, the response rate to the survey appears balanced across treatment and control conditions.

Perhaps most interestingly, the next row in Table 8 indicates that students in treatment classes were substantially more likely to correctly answer all the three quiz items, which assessed student understanding of plagiarism. This is not surprising insofar as these items were based directly on the information provided in the on-line tutorial. The fact that roughly 87 percent of control students answered all three items correctly indicates that many students were aware of much of the information contained in the tutorials. However, the fact that virtually 100 percent of students in treatment classes answered all three items correctly confirms that the intervention provided information to a non-trivial fraction of students and that these students retained such information for at least one semester.

The remaining rows in Table 8 identify the treatment-control differences for the other survey responses where responses were on a scale of 1 to 5 with 1 indicating strong disagreement with the statement and 5 reflecting strong agreement.²⁴ Interestingly, the data from question 6 indicate that students in the treatment courses were significantly more likely to agree that they had a good understanding of plagiarism, a finding consistent with the quiz results and the educational intent of the tutorial. In contrast, the results to questions 8, 9, and 10 suggest that the intervention did not have a statistically significant deterrent effect. That is, respondents in treatment courses were not significantly more likely to think that a professor would detect plagiarism, respond to it in some way or report it to a judiciary committee. As a composite measure of student perceptions regarding the likelihood of detection and sanction for plagiarism, we also calculated the average of survey questions 8, 9 and 10. In the last row of Table 8, we see that assignment to the treatment condition is not significantly related to this composite

²⁴ In results available upon request, we confirm that we obtain comparable results if we use binary outcomes reflecting the top two categories (i.e., agree or strongly agree) instead of the continuous level of agreement measure.

measure.²⁵ While not definitive, these results suggest that the primary mediating mechanism for the intervention was education rather than deterrence.

6. Discussion

Rapid technological advances (e.g., access to full-text resources and cut-and-paste word processing) have contributed to the widely held impression that plagiarism by students has dramatically increased on college campuses. Some commentators have suggested that instructors and institutions should respond to this academic misconduct by increasing the penalties for plagiarism as well as the probability that an infraction is detected (e.g., Galles et al. 2003). In contrast, other observers have expressed unease about fostering increasingly adversarial classroom environments. However, despite the degree of concern and the corresponding calls for reform, there has been surprisingly little credible evidence on how much student plagiarism actually occurs and on the policy determinants of this illicit behavior.

The results of the unusual field experiment discussed here suggest that plagiarism is not all that common at selective institutions but may be substantially more common among college students with average SAT scores. This study also presented experimental evidence that a fairly brief but detailed educational tutorial can be highly effective at reducing the prevalence of plagiarism. An ex-post survey of the participants in this experiment suggested that this tutorial was effective by increasing student knowledge about plagiarism rather than by increasing the perceived probabilities of detection and punishment.

These results are consistent with a model of student behavior in which plagiarism persists in equilibrium because of rational ignorance. That is, the available evidence clearly suggests that students have a generally poor understanding of what constitutes plagiarism. However, they may also have weak incentives to educate themselves about this illicit behavior. College instructors often view policing plagiarism and teaching

²⁵ It should be noted that a specification that includes classroom covariates in addition instead of pair fixed effects yields a point estimate of roughly .13, which represents a moderate size effect (i.e., roughly .25 of the standard deviation of the measure) that is statistically different than zero. This is one of the few instances in which the inclusion of classroom covariates instead of (or in addition to) pair fixed effects leads to any substantive change in our estimates. However, the estimate of .13 is not significantly different than the estimate of .06 shown in Table 8 here. Our read of these results is that the present study does present compelling evidence in either direction with regard to potential deterrent effects.

students about it as outside their responsibilities. Furthermore, when plagiarism is detected, the penalties are often informal and fairly light.

If this characterization of student plagiarism is accurate, it suggests that academic integrity at colleges and universities resembles a “tragedy of the commons” in that the individual incentives of both students and their instructors are not well aligned to support collectively advantageous institutional norms. However, our evaluation results also suggest that educationally themed interventions can meaningfully address this problem. In particular, the web-based tutorial evaluated in this study constituted a cost-effective and easily scalable approach to supporting a “hidden curriculum” of academic integrity and did so without necessarily relying on meaningful effort from instructors.

References

- Angrist, Joshua D. and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press, 2009.
- Baruchson-Arbib, Shifra and Eti Yaari. “Printed versus Internet Plagiarism: A Study of Students’ Perception,” *International Journal of Information Ethics* 1, 2004, pages 1-7.
- Bell, Robert M., and Daniel F. McCaffrey, “Bias Reduction in Standard Errors for Linear Regression with Multi-Stage Samples,” *Survey Methodology* 28(2), 2002, pages 169–179.
- Bertrand, M., Duflo, E., Mullainathan, S. How Much Should We Trust Differences-in-Differences Estimates? *Quarterly Journal of Economics* 2004; 119(1), 249-75.
- Bowers, William J. *Student Dishonesty and its Control in Colleges* (New York: Bureau of Applied Social Research, Columbia University, 1964), 155.
- Bilic-Zulle, Lidija, Josip Azman, Vedran Frkovic, Mladen Petrovecki. "Is There an Effective Approach to Detering Students from Plagiarizing?" *Science and Engineering Ethics* 14, 2008, 139-147.
- Brown, B.S. and D. Emmett. “Explaining variations in the level of academic dishonesty in studies of college students: Some new evidence,” *College Student Journal* 35(4), 2001.
- Cameron, A. Colin, Jonah. B. Gelbach and Doug Miller. “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, Vol 90, No 3, 2008, pages 414-427.
- Donald, Stephen G., and Kevin Lang, “Inference with Difference-in-Differences and Other Panel Data,” *Review of Economics and Statistics* 89(2), 2007, pages 221–233.
- Downs, Anthony. *An Economic Theory of Democracy*. New York: Harper, 1957.

- Galles, Gary, Philip E. Graves, Robert L. Sexton and Surrey M. Walton. "Monitoring Costs and Tolerance Levels for Classroom Cheating," *American Journal of Economics and Sociology*, Vol. 62, No. 4 (Oct., 2003), pp. 713-719
- Genereux, R.L. and B.L. McLeod. "Circumstances surrounding cheating: A questionnaire study of college students," *Research in Higher Education* 36, 1995, pages 687-704.
- Harrison, Glenn W. and John A. List. "Field Experiments," *Journal of Economic Literature* Vol XLII, December 2004, pages 1009-1055.
- Howard, Rebecca Moore and Laura J. Davies. "Plagiarism in the Internet Age," *Educational Leadership*, March 2009, 64-67.
- Jackson, Pamela. "Plagiarism Instruction Online: Assessing Undergraduate Students' Ability to Avoid Plagiarism," *College & Research Libraries* 67(5), 2006, 418-428.
- Lau, Katherine S.L., Craig Nathanson, Kevin M. Williams, Bryce Westlake, and Delroy L. Paulhus. "Investigating Academic Dishonesty with Concrete Measures," Department of Psychology, University of British Columbia, Poster presented at the 17th annual meeting of the American Psychological Society, Los Angeles, 2005
- McCabe, Donald L. "Cheating Among College and University Students: A North American Perspective," *International Journal for Educational Integrity* (2005) Volume 1, No. 1.
- McCabe, Donald L. and Linda Klebe Treviño. "Individual and Contextual Influences on Academic Dishonesty: A Multicampus Investigation," *Research in Higher Education* 38(3), 1997, pages 379-396.
- McCabe, Donald L., Linda Klebe Treviño, and Kenneth D. Butterfield. "Cheating in Academic Institutions: A Decade of Research," *Ethics and Behavior* 11(3), 2001, 219-232.
- McGrath, Charles. "Plagiarism: Everybody Into the Pool," *The New York Times*, January 7, 2007.
- Mueller, Dennis C. *Public Choice II*, Cambridge University Press, 1989.
- Murdock, T.B. and E. Anderman, E. "Motivational perspectives on student cheating: Current status and future directions," *Educational Psychologist* 41, 2006, pages 121-145.
- Posner, Richard A. *The Little Book of Plagiarism*. Random House, Inc. 2007.
- Power, Lori G. "University Students' Perceptions of Plagiarism," *Journal of Higher Education* 80(6), November/December 2009, pages 643-662.
- Rettinger, David A. and Yair Kramer. "Situational and Personal Causes of Student Cheating," *Research in Higher Education* 50, 2009: 293-313.
- Rimer, Sara. "A Campus Fad That's Being Copied: Internet Plagiarism," *The New York Times*, September 3, 2003.
- Scanlon, Patrick M., & Neumann, David R. "Internet plagiarism among college students," *Journal of College Student Development* 43, 2002, pages 374-385.
- Schneider, Alison, "Why Professors Don't Do More to Stop Students Who Cheat," *Chronicle of Higher Education*, 22 January 1999, A9.

Appendix A – Selected Tutorial Screens

Introduction



Plagiarism Tutorial



Avoiding Plagiarism

Note that plagiarism - whether intentional or unintentional - is a serious form of academic misconduct. This brief tutorial will help you understand the different forms of plagiarism as well as how to avoid them.

The tutorial concludes with a short 9-question quiz that tests your knowledge of plagiarism. At the conclusion of the quiz, you will be able to see which questions you answered correctly along with detailed feedback on each question.

The materials in this tutorial have been taken, largely verbatim, from resources developed by Bates, Bowdoin, and Colby Colleges and adapted for use in Blackboard.

Page 1

1. What is plagiarism?

[Contents](#) [Close Window](#)

Page 1 of 18

"To plagiarize" comes from the Latin word "plagiare" which means, "to kidnap." There are many ways to "kidnap" or steal ideas, both intentional and unintentional. As a member of an academic community that takes the sharing of ideas and information very seriously, it is important to avoid even the suspicion of plagiarism. To that end, it is your responsibility to learn how to cite your sources. It is also important to remember that understanding your materials is paramount to writing a good paper, and that plagiarizing reveals a lack of confidence in your own understanding. If you are ever tempted to kidnap someone else's words or ideas – think again – and go to your professor for help.

There are different types and degrees of plagiarism. The four most common are described here and illustrated with examples.

Page 10

10. Mosaic Plagiarism - Example 3

[Contents](#) [Close Window](#)

A third example of mosaic plagiarism in this example of student writing...

Student Writer B

"Only two years later, all these friendly Sioux were suddenly plunged into new conditions, including starvation, martial law on all their reservations, and constant urging by their friends and relations to join in warfare against the treacherous government that had kept faith with neither friend nor foe."

Source

[Clear source text of highlighting.](#)

"Contrast the condition into which all these friendly Indians are suddenly plunged now, with their condition only two years previous: martial law now in force on all their reservations; themselves in danger of starvation, and constantly exposed to the influence of emissaries from their friends and relations, urging them to join in fighting this treacherous government that had kept faith with nobody--neither with friend nor with foe" (Jackson 178).

Example Five: Writer B has borrowed a phrase from Jackson without acknowledging the source or changes to the original text, including substituting or inserting words without indicating he/she has done so.

17. Error 4 - Poor Note-taking

[Contents](#)
[Close Window](#)


Page 17 of 18



Poor Note-taking

Inexperienced students often forget to put quotation marks around notes taken directly from text, or find that their notes are disorganized. As a result, they cannot tell which notes came from which source when they are in the stages of writing up their assignment.

Tutorial Quiz Question 3 with interactive feedback

Question

You are writing a research paper on the history of public education in the United States. You have cut and pasted a lot of information from articles you found on web sites and databases into a Word file on your computer. While writing your essay, you find yourself patching together pieces from different sources, and you have occasionally lost track of which ideas were your own and which were from various articles and websites. You could go back to the original sources but the prospect is daunting. Fortunately, if your professor queries your sources, you can say that you didn't intentionally plagiarize, and this will result in a lesser punishment.

Answer True
 False

Correct Feedback Right! As a general rule, unintentional plagiarism is still intellectual theft and bad note-taking skills are not a mitigating circumstance when punishment is meted out. Recently, there have been well-publicized cases of famous authors whose poor note-taking skills led them to plagiarize. They have had to suffer public humiliation and severe blows to their professional reputations.

Here are some tips for avoiding unintentional plagiarism:

If you take notes on the computer rather than on paper, create a special folder for citation information. In fact, it would be a good idea to create a number of folders: one for your paper; another for sources, with individual files for each and every source; and another folder for the notes you take from each source. Maintain all the information for the bibliography as you go – it'll save time and effort later.

When taking notes, identify your source. Put quotation marks around direct quotes and double check to make sure you've duplicated every punctuation mark. Avoid using the author's language when paraphrasing or summarizing information – unless, of course, you quote verbatim from the original. Here's a tip for keeping your ideas separate from those in your sources; you can either identify each idea as your own, that is, cite yourself, or put your ideas in a different font, case, or color on the screen. Another good idea is to print out your sources whenever possible, even when you have a file-version on your computer. Working from the paper sources will allow you to check quotations for accuracy.

Incorrect Feedback Incorrect. As a general rule, unintentional plagiarism is still intellectual theft and bad note-taking skills are not a mitigating circumstance when punishment is meted out. Recently, there have been well-publicized cases of famous authors whose poor note-taking skills led them to plagiarize. They have had to suffer public humiliation and severe blows to their professional reputations.

Here are some tips for avoiding unintentional plagiarism:

If you take notes on the computer rather than on paper, create a special folder for citation information. In fact, it would be a good idea to create a number of folders: one for your paper; another for sources, with individual files for each and every source; and another folder for the notes you take from each source. Maintain all the information for the bibliography as you go – it'll save time and effort later.

When taking notes, identify your source. Put quotation marks around direct quotes and double check to make sure you've duplicated every punctuation mark. Avoid using the author's language when paraphrasing or summarizing information – unless, of course, you quote verbatim from the original. Here's a tip for keeping your ideas separate from those in your sources; you can either identify each idea as your own, that is, cite yourself, or put your ideas in a different font, case, or color on the screen. Another good idea is to print out your sources whenever possible, even when you have a file-version on your computer. Working from the paper sources will allow you to check quotations for accuracy.

Appendix B – Follow-up Survey

Questions 1 through 10, which are listed below had five possible responses: Strongly Agree, Agree, Neither Agree nor Disagree, Disagree, Strongly Disagree

Question 1 - Overall, I enjoyed this class.

Question 2 - I found this class to be fairly difficult academically.

Question 3 - I found the writing assignment(s) for this class somewhat stressful.

Question 4 - I tended to get an early start, rather than procrastinate, on writing assignments for this class.

Question 5 - When working on the writing assignments for this class, I paid particular attention to avoiding plagiarism.

Question 6 - I have a good understanding of what constitutes plagiarism in academic writing.

Question 7 - I know how to avoid plagiarism in my writing assignments.

Question 8 - If my writing assignments for this class contained any plagiarism, this instructor would detect it.

Question 9 - If this instructor felt that one of my writing assignments contained any plagiarism, he or she would ignore it.

Question 10 - If this instructor felt that one of my writing assignments contained any plagiarism, he or she would report it to the [*institutional judiciary authority*].

Questions 11, 12, and 13 were true/false questions.

Question 11 - Suppose you are writing a research paper. You have cut and pasted a lot of information from articles you found on web sites and databases into a Word file on your computer. While writing your essay, you find yourself patching together pieces from different sources, and you have occasionally lost track of which ideas were your own and which were from various articles and websites. You could go back to the original sources but the prospect is daunting. Fortunately, if your professor queries your sources, you can legitimately claim that you didn't plagiarize because it wasn't intentional.

Question 12 - Suppose it would be quite easy for you to re-tool whole sections of a paper you have written for a previous to satisfy the requirements of another course you are currently taking. It is acceptable practice to re-submit this edited paper – without checking with either professor -- because you are writing a paper for a different professor and a different course.

Question 13 - Plagiarism is not limited to taking something from a book; it also includes stealing ideas from a movie, a professor's lecture, or from an interview on a radio news program.

Table 1: Summary Statistics

	Student		Classroom Aggregate Data		P-value
	Level Data Full Sample	Control Mean	Treatment	Diff: T-C	
Fraction Female	(1) 0.561	(2) 0.499	(3) 0.558	(4) 0.06	(5) 0.30
Fraction Black	0.086	0.090	0.088	0.00	0.93
Fraction Hispanic	0.116	0.108	0.134	0.03	0.45
Fraction Asian	0.180	0.201	0.216	0.02	0.76
Fraction White	0.563	0.537	0.515	-0.02	0.70
Fraction Other Race	0.055	0.063	0.047	-0.02	0.49
Average SAT Score	1407	1404	1409	4.52	0.66
Fraction papers submitted as hard copy	0.059	0.087	0.058	-0.03	0.71
Fraction Seniors	0.213	0.181	0.226	0.05	0.50
Fraction Juniors	0.251	0.196	0.268	0.07	0.16
Fraction Sophomores	0.331	0.284	0.336	0.05	0.28
Fraction Freshmen	0.205	0.340	0.171	-0.17	0.15
Fraction taking the course as Pass/Fail	0.068	0.055	0.081	0.03	0.28
Fraction taking class for a letter grade	0.732	0.607	0.749	0.14	0.23
Class size	27.2	25.9	23.9	-2.07	0.47
Number papers required for the class	2.5	2.1	1.9	-0.21	0.59
Warning about plagiarism on the syllabus	0.301	0.286	0.429	0.14	0.46
Social science course	0.799	0.857	0.857	0.00	1.00
Assistant professor	0.231	0.214	0.286	0.07	0.68
Associate professor	0.221	0.214	0.214	0.00	1.00
Full professor	0.271	0.357	0.214	-0.14	0.42
Visiting professor	0.277	0.214	0.286	0.07	0.68
Female professor	0.327	0.286	0.571	0.29	0.14
Number of observations	1259	14	14		
P-value from SUR with all covariates as the outcome variables and the treatment indicator as the single predictor	0.509				
P-value from OLS regression of treatment indicator on all covariates	0.258				

Notes: * = significant at the 10% level; ** = significant at the 5% level.

Table 2: Relationship between Plagiarism and Student and Class Characteristics

	Bivariate Regressions (1)	Multivariate Regression (2)
Female	0.012 (0.009)	0.011 (0.007)
Black	0.054** (0.027)	-0.005 (0.026)
Hispanic	0.018 (0.017)	-0.010 (0.020)
Asian	0.036** (0.017)	0.032** (0.014)
Other race	-0.006 (0.004)	-0.011** (0.005)
Missing race	0.019* (0.010)	0.019** (0.009)
SAT score	-0.227* (0.127)	-0.255* (0.149)
SAT score squared	0.008* (0.005)	0.009 (0.005)
Missing SAT score	-1.695* (0.890)	-1.890* (1.053)
Hard copy paper	0.045 (0.030)	0.046 (0.034)
Freshman	-0.005 (0.012)	-0.003 (0.010)
Sophomore	0.007 (0.012)	0.004 (0.012)
Junior	0.024 (0.017)	0.024 (0.016)
Missing class year	-0.017* (0.010)	-0.024** (0.009)
Pass-Fail	-0.026** (0.007)	-0.015** (0.007)
Missing grade designation	-0.021* (0.012)	0.001 (0.035)
Class size	-0.000 (0.001)	0.003 (0.004)
Number of papers	-0.008 (0.010)	0.002 (0.029)
Warning about plagiarism in syllabus	-0.004 (0.019)	-0.035 (0.037)
Female professor	-0.012** (0.003)	0.031 (0.067)
Assistant professor	-0.034 (0.044)	-0.058 (0.066)
Associate professor	-0.038 (0.043)	-0.095 (0.145)
Visiting professor	-0.039 (0.041)	-0.071 (0.101)
Number of students	1259	
Mean of dep. var.	0.024	

Notes: Standard errors on student-level covariates are clustered by student and standard errors on classroom level covariates are clustered by classroom with BRL-cluster adjustment.

All models contain pairid indicators.

*=significant at the 10% level; ** = significant at the 5% level.

Table 3: Group Data Estimates of the Effect of Treatment Status on Plagiarism

Panel A: Dependent Variable = Unadjusted Classroom Mean Rate of Plagiarism						
	(1)	(2)	(3)	(4)	(5)	(6)
coeff	-0.036	-0.049**	-0.036**	-0.043	-0.036**	-0.036**
s.e.	(0.024)	(0.023)	(0.014)	(0.027)	(0.014)	(0.014)
p-value (t dist)	0.148	0.043	0.025	0.165	0.025	0.025
N	28	28	28	28	28	28
Df	26	19	13	6	13	13
Student covariates	No	No	No	No	No	No
Class covariates	No	Yes	No	Yes	No	No
Pair Fixed Effects	No	No	Yes	Yes	Yes	Yes
s.e.	regular	regular	regular	regular	robust	hc2
Panel B: Dependent Variable = Adjusted Classroom Mean Rate of Plagiarism (adjusted for student characteristics)						
	(1)	(2)	(3)	(4)	(5)	(6)
coeff	-0.031	-0.038	-0.031**	-0.034	-0.022**	-0.023**
s.e.	(0.024)	(0.024)	(0.013)	(0.027)	(0.010)	(0.009)
p-value (t dist)	0.206	.113	0.038	0.254	0.050	.027
N	28	28	28	28	28	28
Df	26	18	13	6	13	13
Student covariates	Yes	Yes	Yes	Yes	Yes	Yes
Class covariates	No	Yes	No	Yes	No	No
Pair Fixed Effects	No	No	Yes	Yes	Yes	Yes
Weighting	Identity	Identity	Identity	Identity	Number of student-papers	Inverse covariance matrix from step one

Notes: *=significant at the 10% level; ** = significant at the 5% level.

Table 4: Micro Data Estimates of the Effect of Treatment Status on Plagiarism

	(1)	(2)	(3)	Block Bootstrap-t (4)	Wild Block Bootstrap-t (5)
Treatment Effect	-0.023**	-0.023**	-0.023**	-0.023**	-0.023**
Standard error	(0.009)	(0.006)	(0.009)	(0.009)	(0.009)
P-value (from the t distribution)	0.026	0.001	0.018	0.018	0.018
T-statistic (from full sample estimation)		-3.973	-2.684	-2.684	-2.684
95% percent CI of T-statistic from bootstrap				[-1.460, 2.104]	[-2.960, -1.286]
p-value of T-statistic differs from the bootstrap				0.0104	0.0872
s.e. correction	None	Cluster-robust	Cluster-robust with BRL correction	Cluster-robust with BRL correction	Cluster-robust with BRL correction

Notes: All models include student covariates and pair fixed effects.
N=1259 and cluster df=14 for all models.
*=significant at the 10% level; ** = significant at the 5% level.

Table 5: Robustness Tests

	Baseline	Less restrictive definition of plagiarism	More restrictive definition of plagiarism	Assign T to student	Classroom pairs with same professor	Excluding pairs in which both instructors likely knew the purpose of the study	OLS	Probit	Logit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
coeff	-0.023**	-0.025**	-0.020**	-0.020**	-0.011	-0.028**	-0.028**	-0.023**	-0.024**
s.e.	(0.009)	(0.011)	(0.007)	(0.007)	(0.013)	(0.013)	(0.010)	(0.011)	(0.009)
p-value (t dist)	0.018	0.031	.013	0.012	0.413	0.061	0.017	0.048	0.026
N	1259	1259	1259	1259	438	875	1091	1091	1091
Df	1221	1228	1221	1228	417	848	1070	1106	1106
Cluster df	7	14	7	14	6	10	12	12	12
Control mean	0.033	0.040	0.021	0.033	0.032	0.036	0.033	0.033	0.033

Notes: All estimates in columns 1-6 come from specifications using the micro data with BRL-adjusted s.e.

Column 7 presents OLS estimates with no standard error corrections.

Columns 8 and 9 present average marginal effects with standard errors calculated using the delta method.

All models include student covariates and pair fixed effects.

*=significant at the 10% level; ** = significant at the 5% level.

Table 6: Robustness by Student Characteristics

	Gender		Class				
	Female (1)	Male (2)	Freshman (8)	Sophomore (9)	Junior (10)	Senior (11)	
coeff	-0.023** (0.011)	-0.026** (0.012)	-0.039 (0.036)	0.001 (0.015)	-0.042** (0.017)	-0.028* (0.016)	
s.e.							
p-value (t dist)	0.059	0.054	0.313	0.952	0.026	0.111	
N	713	546	260	408	302	259	
Cluster df	14	14	7	14	12	9	
Control mean	0.036	0.029	0.056	0.025	0.029	0.024	

Notes: Each column represents a separate regression.

All models include student covariates and pair fixed effects.

Standard errors shown in parenthesis are BRL-adjusted clustered standard errors.

*=significant at the 10% level; ** = significant at the 5% level.

Table 7: Robustness by Class Characteristics

	Female Professors (1)	Male Professors (2)	Full Professors (3)	Not Full Professors (4)	Syllabus Warning (5)	No Syllabus Warning (6)
coeff	-0.050**	-0.028	-0.013	-0.024**	-0.011	-0.023*
s.e.	(0.015)	(0.020)	(0.039)	(0.007)	(0.019)	(0.013)
p-value (t dist)	0.031	0.209	0.771	0.010	0.592	0.122
N	414	845	350	909	388	871
Cluster df	4	6	3	9	3	7
Control mean	0.090	0.024	0.038	0.030	0.023	0.035

Notes: Each column represents a separate regression.

All models include student covariates and pair fixed effects.

Standard errors shown in parenthesis are BRL-adjusted clustered standard errors.

*=significant at the 10% level; ** = significant at the 5% level.

Table 8: Impacts on Student Attitudes and Perceptions

	Control Mean	Diff: T-C
	(s.d.)	(s.e.)
	(1)	(2)
Answered all 3 quiz items correctly	0.866	0.157** (0.070)
Question 1 - Overall, I enjoyed this class.	3.766 (1.035)	0.301 (0.236)
Question 2 - I found this class to be fairly difficult academically.	3.416 (0.938)	-0.030 (0.122)
Question 3 - I found the writing assignment(s) for this class somewhat stressful.	3.237 (0.998)	-0.176 (0.129)
Question 4 - I tended to get an early start, rather than procrastinate, on writing assignments for this class.	2.968 (1.208)	0.050 (0.161)
Question 5 - When working on the writing assignments for this class, I paid particular attention to avoiding plagiarism.	4.000 (0.829)	0.006 (0.136)
Question 6 - I have a good understanding of what constitutes plagiarism in academic writing.	4.363 (0.634)	0.075** (0.032)
Question 7 - I know how to avoid plagiarism in my writing assignments.	4.356 (0.552)	0.034 (0.336)
Question 8 - If my writing assignments for this class contained any plagiarism, this instructor would detect it.	4.021 (0.776)	0.076 (0.196)
Question 9 - If this instructor felt that one of my writing assignments contained any plagiarism, he or she would ignore it.	1.598 (0.770)	-0.112 (0.276)
Question 10 - If this instructor felt that one of my writing assignments contained any plagiarism, he or she would report it to the College Judiciary Committee.	3.811 (0.814)	0.007 (0.102)
Question 8/9/10 mean, 5=strongly agree that instructor will notice/address plagiarism (question 9 reverse coded).	4.077 (0.608)	0.061 (0.407)

Notes: N=369.

Each row reflects a separate OLS regression in which the outcome is the student response to a particular survey question, coded 1 to 5 as described in the text.

BRL-cluster adjusted s.e. are shown in parentheses.

All regressions include pair fixed effects in addition to the treatment indicator.

*=significant at the 10% level; ** = significant at the 5% level.

Figure 1: Unadjusted plagiarism rates in treatment and control classes

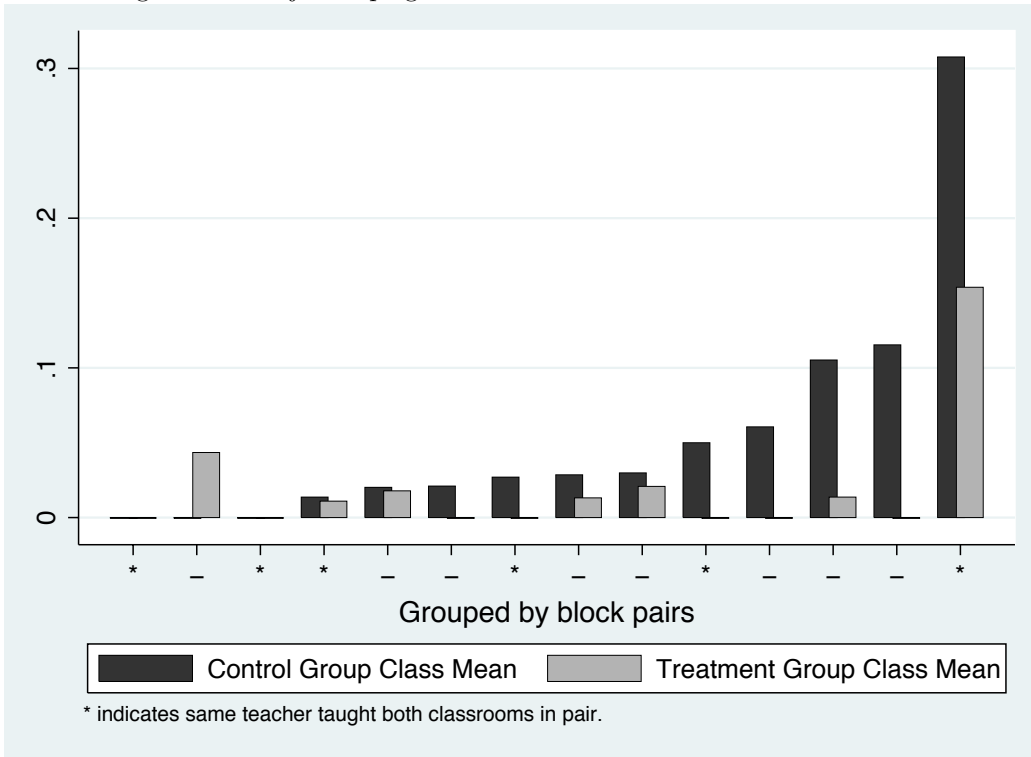


Figure 2: The impact of the web-based tutorial on plagiarism, by student SAT score

