

Introduction and Literature Review

Student plagiarism has been a longstanding problem in academic settings, particularly among business students¹. While the incidence of plagiarism is hard to measure², the perceived severity of it has induced many teachers to change how they teach in order to discourage, avoid, or penalize student cheating. Some such changes may seriously harm the educational enterprise.

¹ See Callahan, David, 2004, Chapter 7 discusses student plagiarism in depth. Callahan is concerned with cheating in the United States. Concerns with plagiarism also are evident for Australia (Devlin, 2006), New Zealand (Sheridan, Alany and Brake, 2005), South Africa (Ellery, 2008), and the United Kingdom (Sciafe, 2007, Tennant and Duggan, 2008). I know of no reason to suppose that it is not universal. Concerning business students, see Mangan, (2006, 2009). Nor are concerns about plagiarism confined to students, see for example, Hoover, 2006, and Enders and Hoover, 2006, who argue that current sanctions against plagiarism among professional economists are not sufficient to deter rational but unethical actors. Godlee, 2007, illustrates the same point for the medical literature.

² See Cross, Michael, who notes: "There is no evidence that plagiarism is becoming more prevalent in research. But there is no doubt that plagiarism happens,..." On the other hand, Goffe and Sosin cite the Center for Academic Integrity to the effect that: "the use of the internet is increasing the plagiarism rate because of the ease of cutting and pasting unattributed material into writing assignments..." p.285.

For example, some teachers choose not to assign papers at all for fear that they would be unable to reliably detect plagiarism.

We explore less extreme approaches. For the last several years, we have been using Turnitin, primarily to detect questionable paper submissions, submitted in partial satisfaction of the intermediate finance course at George Washington University³. More recently, the University has made another similar product, SafeAssign, available through the course management software suite, Blackboard. These two software devices, Turnitin, and SafeAssign can both be used to detect questionable submissions.⁴

While there is little in the Finance literature, this is not the first comparison between Turnitin, SafeAssign or other software systems for detecting plagiarism. Most such studies relied primarily on surveys. Whittle and Murdoch-Eaton (2008) found that students in a writing course valued one-time access to Turnitin for drafts prior to submitting term papers. Also, Kaner and Fiedler (2008) compared Turnitin and SafeAssign's earlier incarnation "MyDropBox" with a selection of published papers from the engineering literature. They were distressed to find that only 3 (Turnitin), or 4 (MyDropBox) of the 13 published papers were detected as plagiarised⁵, and suggest that these packages are being oversold. Such incomplete detection seems puzzling if the original sources were actually in the software databases. It seems more likely that working papers, secondary citations and the like account for these mixed results. Evidently, the original sources of the published articles were not part of either software packages' databases. This may be somewhat of a concern for finance as well. Potential users may want to periodically check which journals are included in current versions of the software system databases.

Sciafe (2007) and Bull, et al. (2008), conducted much more extensive studies for the British Joint Information Systems Committee (JISC). Sciafe conducted an extensive search of available plagiarism detection software packages, identifying eleven of them for in-depth study. Based mainly on survey responses of the companies, Sciafe concluded that Turnitin was the best all-around plagiarism detection software suite, while a Swedish product "Urkund" and also "MyDropBox/SafeAssign" were rated second and third, respectively. (p. 29). Based on tests with a dozen "test files," including some purchased papers, MyDropBox, Turnitin and Urkund again scored well (p. 32) Bull, et al. similarly evaluated 5 software packages including Turnitin (but not SafeAssign), based on surveys of academic users, and eleven "test files." Turnitin, and two other products, "CopyCatch" and "Eve2 2.3" were all rated highly.

It is generally understood that writing assignments help students digest scholarly

³ We have also used Turnitin to screen papers from three other courses. Two were undergraduate courses that were broadly similar to the main group, as were the corresponding summary data. They are included with the main data. The third was an MBA course. MBA papers appear to be a bit different, with higher citation rates, and were set aside where appropriate.

⁴It is important to note that detection of plagiarism after the fact is not the only possible use of these software packages. The vendors of both software packages make clear that they can be used to teach proper citation practices as well, and some schools appear to be using them in that way. For example, see Sheridan, Alany and Brake (2005). On the other hand, Rebecca Attwood of the Times Higher Education Supplement (2008) notes that some British universities are limiting access to Turnitin because some students are attempting to "beat the system" with repeated submissions from which they progressively delete offending text.

⁵See Kaner and Fiedler, 2008, "A Cautionary note on Checking Software Engineering Papers for Plagiarism".

material. As noted above, in our experience, some finance and other business teachers shy away from writing assignments because of concern about plagiarism. A variety of reasons have been offered for this. In some cases, they are unsure of their ability to detect plagiarism, and so chose to not assign papers at all. In other cases, they hesitate to pursue cases of possible plagiarism because of the administrative burden involved⁶, or because of uncertainty about where to “draw the line”. Some colleagues have suggested that plagiarism is the fault of the instructor for not effectively communicating ground rules for paraphrasing techniques, citation practices, and the like. We suspect few of us would relish having our teaching ability or techniques questioned as a result of bringing cases of apparent plagiarism to the attention of our colleagues. One of our dear colleagues has even suggested that cheating ought to be defined as **not** helping fellow students.

Nor is the use of anti-plagiarism software beyond criticism. Purdy argues that “the services we sometimes use to ensure the integrity of students’ texts can themselves be of questionable integrity - largely through the design of their archives”. If we understand his argument, his concern is about copyright issues. In March 2007, after an extended controversy, Turnitin, and McLean High School in Fairfax county, Virginia were jointly sued because the school required students to submit papers to Turnitin⁷. The students objected on copyright grounds. Turnitin argued that they were protected by the fair use exception for copyrights⁸. Turnitin appears to have won that case⁹. Even so, it is understandable that some finance teachers might prefer to avoid the whole problem by relying mainly on in-class exams and homework for grading.¹⁰ Of course, homework is also vulnerable to plagiarism. We fear this reluctance to confront plagiarism seriously impedes student learning, and also harms the teacher’s ability to evaluate that learning.

At this university, charges of academic integrity violations can be handled in several ways. First, the individual professor could simply adjust grades to reflect their suspicions. Of course, our students are intensely curious about grading, and readily complain about what they perceive as unfair grading. Second, the professor might confront the student with their evidence or suspicions, and then sanction them within the context of the course. Neither of these approaches would necessarily leave a record, except with the individuals directly involved. Third, the professor can notify the university’s Office of Academic Integrity (OAI). If the apparent infraction is a first time (in the eyes of the Office of Academic Integrity), the professor can then suggest an appropriate sanction. The student is notified, and if they accept the charge and the sanction, their name is entered into a list kept by the office of academic integrity. It is our understanding that the list is destroyed if and when students graduate, provided there are no further charges. Fourth, if the student has already come to the attention of the office of academic

⁶ Some schools are beginning to address this issue as well. See Lipka, 2009, “Colleges sharpen tactics for Resolving Academic Integrity Cases”.

⁷See MacMillan, Douglas, 2007.

⁸See Read and Brock, Chronicle of Higher Education, 2008.

⁹See Young, Chronicle of Higher Education, 2008.

¹⁰The assumption may be that short, easy homework assignments are less likely to be plagiarized than longer, harder papers. The cost of the maximum potential sanction for plagiarising (suspension or expulsion) is so much greater than the benefit for small assignments. On the other hand, these sanctions may not be credible for small assignments. We know of no other reason why homework would not be subject to plagiarism.

integrity, there is a requirement of a formal hearing. A hearing panel of at least 2 members of the faculty, and two students would be convened to determine the accuracy of the charge, and that panel may suggest other sanctions. Finally, if a first-time student disputes the charge or the proposed sanction, they can ask for a hearing as well. Once a hearing panel is brought in, they can override the professor's suggested sanctions, making them more or less severe.

We have required papers of our students for more than a decade. We have had to come to grips with plagiarism. Coming to grips presumably implies looking for plagiarism once we have reason to suspect it, sensibly interpreting what we find, and developing an appropriate response.

In what follows, we describe our initial experience with Turnitin, and later SafeAssign, and how Turnitin and SafeAssign compare in evaluating papers submitted for our courses. Among our preliminary conclusions, we find that the two software packages are generally similar, but that the results depend importantly on the databases associated with them. Thus, depending on the context, the proprietary database that Turnitin has built-up over the years may confer a significant advantage.

The Data

In all, we have evaluated 1192 student papers for plagiarism with Turnitin and SafeAssign, including 123 in the initial test ba

Turnitin was with other papers I had already posted to Turnitin.

These results were presented to department faculty, and after extended discussion, a procedure was developed. For subsequent semesters, we warned students; and we tested them. In the syllabus, we explicitly reserved the right to compare student papers to work submitted at other times and places to ensure academic integrity. Also, we routinely posted all subsequent papers to Turnitin.¹¹

Concerns about student writing were also reflected at the University level. Beginning as early as 2004, the university initiated its “University Writing Program”, intended to upgrade the writing preparation our students received. The program became fully operational in 2006. In this program, students were required to take newly intensified writing courses, followed by required “writing in the disciplines” courses, where the latter were intended to reinforce skills developed in the general course work with writing more closely related to their chosen field of study. While my Intermediate Finance course was never classified as a “writing in the disciplines” course, it was increasingly clear that my focus on writing integrity was in keeping with the spirit of the University Writing Program.

In early 2008, SafeAssign was added to Blackboard. That Summer, we scanned most of the previously accumulated papers with SafeAssign, and then added them to the institutional database within SafeAssign. For Fall 2008 and Spring 2009, both Turnitin and SafeAssign were used to scan student papers, more or less simultaneously. Papers from Spring 2008 were also processed at nearly the same time. Thus, in our empirical work, we will focus on both the longer SafeAssign and Turnitin time series created, as well as the current properties of these two software packages for Spring 2008, Fall 2008 and Spring 2009¹².

The data naturally break up into two sections. All papers were tested using Turnitin,

¹¹ Beginning In Fall 2004, we inserted the following into our Syllabi: **“For research purposes, and also to encourage academic integrity, I reserve the right to compare any papers submitted to other work submitted at other times or places.”** Over time, this was also reinforced with progressively more explicit verbal warnings.

By Spring 2007, the text was strengthened to: **“For research purposes, and to encourage academic integrity, I reserve the right to compare papers submitted to those submitted at other times or places. If detected, plagiarism will be punished.”**

By Fall 2008, this was changed to: **For research purposes, and to encourage academic integrity, I reserve the right to compare papers submitted here to those submitted at other times or places. If detected, plagiarism will be punished.**

This year, we will go farther. The University and Blackboard have added “SAFEASSIGN to the suite of features in Blackboard. Safeassign is a type of “plagiarism prevention” software. Papers run thru it are linked to an assessment of how much of the paper is identifiably from other sources. If these sources are properly quoted and cited, that’s fine. If not, there may be a problem. While it is important to note that there may be innocent reasons for any such findings, it is best to take care to ensure that your writing is your own, and credit is given where credit is due.

¹²For both Turnitin and SafeAssign, papers are uploaded to the software system, which then generates what Turnitin calls an “originality report”. These reports indicate the degree of textual overlap between the paper and other known sources, in detail. For Turnitin, the sources are then ranked by percentage overlaps. SafeAssign does not provide such a percentage ranking, but does list apparent outside sources. In both cases, the overlapping phrases are specifically identified and marked to link them to the outside sources identified. Also, with Turnitin, papers within a given class can be sorted by name, time posted, percent overlap, or other variables.

generally at the end of the semester. Papers submitted before Fall 2008 were tested against SafeAssign when it became available, in Summer 2008, and then added to the associated in-house database at that time. So far we have matched up reports for 902 of 1192 papers across the two systems.¹³ For those, Turnitin averaged 24.0% (median 19.0%) overlap detected, with a standard deviation of 20.3%. SafeAssign averaged 19.0% (median 15.0%) overlap detected, with a standard deviation of 17.9%. However, for the earlier papers, the direct comparison is not fully appropriate since the two databases used by Turnitin and SafeAssign continue to grow over time. While we continue to try to match up papers so that our results should be regarded as preliminary, it appears that Turnitin detects somewhat more overlap than SafeAssign.

While the SafeAssign data were compiled mostly during Summer 2008 (and will be discussed in a later paper), we can consider the Turnitin results, which represent a more conventional time series (table 3). The percent overlap rises gradually, punctuated by minor spikes at Fall 2004 and Spring 2008, and a major spike at Spring 2006. These are also the three semesters that produced the preponderance of the academic integrity charges. This gradual rise might be due to any of several changes - each major part of the database grew over time, there may have been improvements in the search algorithm, and the students might have begun to copy more from detectable sources. To get a better sense of which of these changes are most important, we will need to look in more detail.

Consider the most recent three semesters. Papers submitted for Spring 2008 were tested with Turnitin near the end of the semester, and with SafeAssign a few months later, which would slightly favor SafeAssign. Papers submitted for Fall 2008 and Spring 2009 were tested with both software packages, generally near- or after- the end of the semester. Testing with SafeAssign was automated via Blackboard as the students posted the papers during the semester. Testing with Turnitin was done manually at the end of the semester. For these latter papers, any timing difference between when they were submitted would favor Turnitin by at most 6 or 7 weeks. On balance then, allowing for small delays, these last three semesters of papers were submitted at roughly the same time. Table 4 presents disaggregated results for Turnitin for the last three semesters. (Corresponding data for SafeAssign will be forthcoming.)

The databases used by these two software packages differed somewhat. During the early part of this period, Turnitin used a combination of Cengage Academic One Journals (as well as some other journal collections), Internet web sites, and student submissions submitted over the life of Turnitin, since 1997¹⁴. During 2006, Turnitin switched from Proquest to Cengage.¹⁵ More recently, according to Stowers and Hummel (2011), Turnitin is again using Proquest. SafeAssign may allow other configurations, but at this university, it screens papers mainly against Proquest journals and other sources, Internet web sites, and prior submissions from within the university, but not papers submitted by students at other schools.¹⁶ In both cases, each component of the databases have continued to grow over time.

¹³Most student papers are easy to match. However, for the remainder, because of various file name changes, this is surprisingly time consuming.

¹⁴Turnitin Help Desk, July 2009.

¹⁵ Turnitin sales rep., September 2009. I am still seeking information to clarify exactly when the switch to Cengage from Proquest was made.

¹⁶SafeAssign technical support, July 2009

A Provisional Model

One straightforward approach would be to model overall percent overlap as a function of other variables. For example:

$$1) PO = f(POJ, POW, POS, ALG, ST, T)^{17}$$

2)

where PO=Overall percent overlap, POJ is the percent overlap with journal sources, POW is the percent overlap with web site sources, POS is the percent overlap with student paper sources, ALG is the search algorithm, ST is student type, and T is time submitted. Unfortunately, while this variable is useful for an overview, it is not clear how to proxy or otherwise tease the independent variables apart.

In our empirical work (which is ongoing), we will disaggregate the data to focus on six components of PO. Thus:

$$2) POJ = POJC + POJN$$

$$3) POW = POWC + POWN$$

$$4) POS = POSC + POSN$$

where the “C” and “N” end letter indicates the overlap is properly cited (C), or not (N). In our modeling, this allows us to focus on simpler relations such as:

$$5) POJC = f(\text{Databases included (DI), Search Algorithm (SA), Student Type (ST), and Time Submitted (T)})$$

and,

$$6) POJN = f(\text{Databases included (DI), Search Algorithm (SA), Student Type (ST), and Time Submitted (T)})$$

where DI is Databases included, SA is Search Algorithm, ST is student type, and T is Time Submitted. Here, the effect of SA should be constant and captured in that term, while the effect of DI would be expected to rise over time since the databases are growing, with a possible discontinuity when Turnitin shifted databases. A trend variable for time submitted might reasonably proxy database growth. A dummy variable for the discontinuity would also be reasonable. Student type is limited by available data and privacy concerns, but these could include final exam grades (excluding the paper) and gender. Presumably the better students cite more sources properly. It would be interesting to see if better students fail to cite sources more or less than weaker students. We can model the other variables, POWC, POWN, POSC and POSN, similarly.

Preliminary Results and Discussion

¹⁷Note that this relation is not additive. That is, PO does not equal POJ+POW+POS. Instead, PO is the percent of text that overlaps with at least one of these sources. Thus, if a student's paper overlaps 20% with an earlier student's paper (POS), and 15% with journal articles (POJ), the PO could be as high as 35% if the two sets of text were different or as low as 20% if they are the same.

For each of the next several academic years, partly as a result of our use of Turnitin, between 2 and 5 students were charged with violations of academic integrity for submitting papers that were substantially written by someone else (See Table 1). The degree of text overlap in these cases ranged from 54% to 100%, as identified by Turnitin. Notably, the 15 cases brought to the attention of the Office of Academic Integrity were not evenly spread in time. Nine (9) of them were detected in either Spring 2006 or Fall 2006. In addition to these, there were a few other cases that we would have pursued, except that the student either dropped the course, or had earned an “F” anyway. Based on retrospective evaluation of the papers added to SafeAssign in the Summer of 2008, at least 13 and perhaps 14 of the 15 students charged with academic integrity violations using Turnitin would also have been charged using SafeAssign. Likewise, none of the students not charged with Turnitin would have been charged based on SafeAssign. Of course, neither of these are unbiased tests since Turnitin was used earlier, and hence it would have been at a disadvantage given ongoing database growth. Still, despite the test limitations, for purposes of supporting academic integrity charges, that is, addressing the most extreme cases, the two software packages yield surprisingly similar results.

As noted above, over the next several semesters, the verbal and written warnings were made progressively more explicit. Of the 15 students charged with infractions, none disputed the charges. The usual sanction was a grade of “D” for the course. One had already been caught before, and two asked for special consideration (mercy?) on the basis of extenuating circumstances. The finding that only 15 cases warranting charges of academic integrity over 11 semesters and 1069 papers does not indicate that the remaining students were above reproach. In the early years, we were not sure how much support we could expect from the University, or how much controversy would be produced in addressing these cases, and so chose to focus only on the most blatant, easily proven ones. The rest of the data are discussed more fully below.

It is also important to note that overlap is not necessarily evidence of improper citation. Neither Turnitin nor SafeAssign were configured to automatically exclude properly cited sources from computation of overlap. For these papers, that can best be done by examining the output directly and perhaps rerunning the scan with those sources excluded from the software’s databases. At the time of these studies, 150 scans cost around \$150 via individual subscription. Instead, we usually examined the output directly for questionable papers. For practical reasons, we did not pursue cases where small parts of student submissions were not properly cited. Had we done so, we probably could have charged the vast majority of our students. To our thinking, it serves no purpose to define a crime so broadly that nearly everyone is a criminal. While the average overlap was near 23%, many of the overlapping sources were actually properly cited. Based on spot checks, perhaps half of the 23% were not. With such ubiquitously faulty citations, it is important to address the issue of where to draw the line.

Failure to cite journal articles that overlap the student’s text is an obvious lapse. However, given the incomplete coverage in the two databases Cengage and Proquest, such that a given journal might not be represented in the database, no mechanical cutoff is warranted. Often, we would see overlap with secondary sources, simply because the primary sources are not in the databases. Further, since these lapses are typically small - perhaps up to 5% of a typical students paper - it seems more reasonable to suspect (and grade) for carelessness more than intentional plagiarism. Such lapses might also be regarded as an opening for teaching better citation practices.

Small failures to cite web sites are even less likely to suggest deliberate plagiarism.

Many authors will have earlier versions of published papers on personal or institutional web sites. If such papers are published in journals missed by the Turnitin or SafeAssign databases, it may appear that students are plagiarising author web sites when they are properly citing the source journals. On the other hand, wholesale overlaps with Wikipedia or newsmagazines suggest an intent to deceive.

Small failures to cite student papers may also be innocuous, for similar reasons. Again, wholesale overlaps are a different story.

In fact, once we eliminate the easy cases, the most appropriate procedure is less clear. Initially, when we began screening papers against Turnitin in 2004, we did not want to “show our hand”. The grossly plagiarized papers could be easily identified and sanctioned. On the other hand, the papers with some faulty citations surely included some that were guilty of no more than sloppy or careless editing. In such cases, we did nothing beyond slightly lowering the grade for “sloppy editing”. This may have been too lenient. Galles, et al. argue that typical policy approaches to plagiarism spend too much effort on monitoring, and do not provide sufficiently tough or consistent sanctions.¹⁸

On the other hand, several scholars have suggested that part of the problem stems from student uncertainty about what constitutes plagiarism, faculty uncertainty about where to draw the line, and inconsistent responses to perceived plagiarism¹⁹.

Further, as the University Writing Program became fully operational in 2006, we could tell that student writing was improving. Going forward, we anticipate that either Turnitin, or SafeAssign could have more of a role in teaching - showing where citations are needed, rather than just flagging where they are missing after the fact.

There are other reasons to be restrained with charges of academic misconduct. Over time, the databases used by Turnitin and SafeAssign are becoming more extensive. However, neither is complete. Some of the most important finance journals are not yet included in these databases. One result is that properly cited passages from missing journals often appear to be improperly cited passages from related working papers or from subsequent citations in later works.

In response to a clear problem, we have used Turnitin and SafeAssign to detect wholesale plagiarism. Conversely, some scholars suggest a more proactive and preventative approach²⁰. Devlin notes: “As well as ensuring appropriate consequences for plagiarists, several (Australian) universities are beginning to formalize the inclusion of learning and teaching strategies in anti-plagiarism related policy and practice, as well as paying closer attention to the communication of unambiguous definitions of plagiarism.” Likewise, Howard and Davies (2009) argue that with the tremendous variety of good and bad sources available online, the need is to teach better writing skills rather than to simply “villify the internet (and students who use it improperly)” In hindsight, the combination of a punitive approach for extreme abuses, and a focus on improving

¹⁸ See Galles, et al., 2003, “Monitoring Costs and Tolerance Levels for Classroom Cheating.

¹⁹ See Ma, Wan and Lu 2008, p. 201, Selwyn, 2008, p. 473-475, Ellery, 2008, and particularly Tennant and Dugan, table 4..

²⁰See Devlin (2006), “Policy Preparation and Prevention: Proactive Minimization of Student Plagiarism”.

writing and clearly defining appropriate citation practices otherwise may be the best approach.

Lastly, the odd timing of academic integrity cases merits discussion. Initially, we expected the rumours to fly with the first few cases, ending the problem. If that happened, the effect was not very long-lasting. In fact, Tim Terpstra of the GW Office of Academic Integrity conjectured that the rumours die down with graduation of each class, so that every few years or so, some of our students seek to test that limit. One further pattern is apparent. Early on, it would have been easy to convince ourselves that devices such as Turnitin, and warnings to use them were ineffective in stopping plagiarism. However, it is notable that the drop in academic integrity cases after each surge came after use of stronger warnings in the syllabus, or more open use of such technology. One might conclude that we might make substantial progress in the effort against plagiarism, if we take heart, and take it more seriously in the future.²¹

Conclusions

We have begun to explore data on 1192 papers submitted to Turnitin and SafeAssign. While much more remains to be done, some conclusions are warranted. First, despite the different journal databases and student paper databases, the most blatant cases of plagiarism are likely to be caught by either software package. Second, small instances of apparent plagiarism ought to be interpreted cautiously. The databases behind these software packages are still quite incomplete, so that proper citation of missing journals may look like plagiarism. Third, in part because of the student databases, Turnitin appears to detect a somewhat higher percentage overlap with most submitted papers. Finally, given the ease with which blatant plagiarism can be detected, and its apparent responsiveness to blunt warnings, we might take heart that with reasonable efforts, plagiarism can be controlled, and need not prevent careful instructors from requiring papers where they are warranted.

²¹ A recent article in the Times Higher Education Supplement reported that in the United Kingdom, the apparent incidence of “serious plagiarism” based on Turnitin similarity reports, seems to be declining, from a 2005 peak through 2012.

Works Cited

- Attwood, Rebecca, "Institutions limit access to anti-cheat software", Times Higher Education Supplement, no. 1851, June 26, 2008, p. 8
- Joanna Bull, Carol Collins, Elizabeth Coughlin and Dale Sharp, "Technical Review of Plagiarism Detection Software Report, Computer Assisted Assessment Centre, University of Luton, on behalf of the Joint Information Systems Committee (JISC), 2008(?).
- Callahan, David. "The Cheating Culture: Why More Americans are Doing Wrong to get Ahead", Harcourt Books, 2004.
- Cross, Michael. "Policing Plagiarism." *BMJ (British Medical Journal?)* 335 (Nov. 2007): 963-964.
- Devlin, Marcia. "Policy, Preparation, and Prevention: Proactive minimization of Student Plagiarism." *Journal of Higher Education Policy and Management* 28.1 (2006): 45-58.
- Ellery, Karen, "Undergraduate Plagiarism: A Pedagogical Perspective", *Assessment and Evaluation in Higher Education*, v. 33no. 5, October 2008, p. 507-516.
- Enders, Walter, and Gary Hoover. "Plagiarism in the Economics Profession." *Challenge* 49.5 (2006): 92-107.
- Galles, Gary, et al. "Monitoring Costs and Tolerance Levels for Classroom Cheating." *American Journal of Economics and Sociology* 62.4 (2003): 713-719.
- Godlee, Fiona. "Plagiarism and Punishment (Editorial)." *BMJ (British Medical Journal?)* 335 (Nov. 2007).
- Goffe, William L., and Kim Sosin. "Teaching with Technology: May You live in Interesting Times." *Journal of Economic Education* (Summer 2005): 278-291.
- Hoover, Gary A. "A Game-Theoretic Model of Plagiarism." *Atlantic Economic Journal* 34 (2006): 449-454.
- Howard, Rebecca Moore, and Laura J. Davies, "Plagiarism in the Internet Age" *Educational Leadership*, March 2009, p. 64-67.
- Jump, Paul, "Turnitin is turning up fewer cases of Plagiarism", Times Higher Education Supplement, April 4, 2013, p. 11.
- Kaner, Cem, and Rebecca L. Fiedler. "A Cautionary Note on Checking Software Engineering Papers for Plagiarism." *IEEE Transactions on Education* 51.2 (2008): 184-188.
- Lipka, Sara, "Colleges Sharpen Tactics for Resolving Academic Integrity Cases", *Chronicle of Higher Education*, 4/10/09, v. 55, n. 31 p. a20.
- Mangan, Katherine, "Survey Finds Widespread Cheating in MBA Programs", *Chronicle of Higher Education*, September 29, 2006, v. 53, n. 6, a44.
- Mangan, Katherine, "Lessons Business Graduates Apply to the Real World May Include Cheating", February 10, 2009, v?, n.?, Student Affairs section.
- Ma, Hongyan Jane, Guofang Wan and Eric Yong Lu, "Digital Cheating and Plagiarism in Schools", *Theory into Practice*, 2008, v. 47, p. 197-203.
- MacMillan, Douglas. "Looking Over Turnitin's Shoulder." *Business Week Online* 13 Mar. 2007. Academic Search Premier. EBSCO. 21 July 2009 <<http://search.ebscohost.com/>>.
- Purdy, James P. "Anxiety and the Archive: Understanding plagiarism Detection Services and Digital Archives." *Computers and Composition* 26 (2009): 65-77.
- Read, Brock, and Paula Wasley. "Anti-Cheating Crusader Vexes Some Professors." *Chronicle of Higher Education* 28 Feb. 2008: A1-A12.

Scaife, Bryan. IT Consultancy Plagiarism Detection Software Report. N.p.: JISC Plagiarism Advisory Service, September, 2007.

Selwyn, Neil, “‘Not Necessarily a Bad Thing...’: A Study of Online Plagiarism Amongst Undergraduate Students”, *Assessment and Evaluation in Higher Education*, v. 33, n. 5, October 2008, p. 465-479.

Sheridan, Janie, Raid Alany and Dulcie-Jane Brake, “Pharmacy Students’ Views and Experiences of Turnitin - an Online Tool for Detecting Academic Dishonesty”, *Pharmacy Education*, September/December 2005, 5 (3/4): 241-250.

Stowers, Robert H., and Hummel, Julie Y., “The Use of Technology to Combat Plagiarism in Business Communication Classes”, *Business Communications Quarterly*, V. 74, n. 2, June 2011, p. 164-169.

Tennant, P, and F Duggan. Academic Misconduct Benchmarking Research Project: Part II The Recorded Incidence of Student Plagiarism and the Penalties Applied. N.p.: Higher Education Academy - JISC, May 2008.

Whittle, Sue R. and Deborah G. Murdoch-Eaton, “Learning about Plagiarism using Turintin Detection Software”, *Medical Education*, 2008, v. 42, p.528.

Young, Jeffrey R., “Judge Rules Plagiarism Detection Tool falls under Fair Use”, *Chronicle of Higher Education*, 4/4/2008, v. 54 Issue 30.

Tables

Table 1.
Academic Integrity Violations 2004-2009:
(For my courses, and school-wide)²²

	My courses (violations/papers submitted)			School-wide plagiarism charges		
	Spring	Summer	Fall	Spring	Summer	Fall
2003	-	-	0/65	-	-	10
2004	2/58 ²³	na	2/46	5	2	3
2005	0/92	na	0/59	3	0	0
2006	4/93	na	5/99	8	2	7
2007	0/102	0/32	1/150	4	5	3
2008	3/132	na	0/132	3	0	3
2009	0/132	na	na	3	2	na
Total	9/609	0/32	8/551	26	11	30

From 2003 to the present, student paper submissions were submitted to Turnitin, where they were scanned for originality. Turnitin generates reports that indicate what portion of a paper overlaps with known outside sources. While there are several reasons why there might be high overlap; for example, when a student writes a follow-up to an earlier paper, in some cases, it appeared that the student was submitting a paper that had already been submitted earlier by a previous student. The first 123 papers were scanned experimentally, and after a grade had been rendered, hence, no grade was changed even though two pairs of these 123 papers appeared to have been virtually duplicates, with the names changed. Out of 1069 papers subject to scrutiny for which grades had not already been given, on further investigation, 15 were charged with violations of academic integrity.

These 17 cases are among a total 67 cases brought to the attention of the Office of Academic Integrity concerning student plagiarism in the school of business during this period.

²²School-wide data were provided by Tim Terpstra of the Office for Academic Integrity.

²³Spring 2004 was the second semester added to Turnitin, during the Summer of 2004. Since grades had already gone out and no warning notice was in the Spring syllabus, no charges were made. However, had we done so using the criteria subsequently applied to the other classes, two cases out of 58 submissions would have been pursued.

Table 2.

Summary of Originality Reports. Papers Submitted to Turnitin (for my classes) 2003-2009.

Classes	# of Reports	Percentage Overlap (percent of semester sample)				
		75+	50-74	25-49	0-24	0
Fin. 122, Fall 2003	65	-	3 (4)	13 (20)	48 (74)	1 (2)
Fin. 122, Spr. 2004	58	3 (5)	5 (9)	2 (3)	47 (81)	1 (2)
Fin. 122, Fall 2004	46	2 (4)	4 (9)	7 (15)	32 (70)	1 (2)
Fin. 122, Spr. 2005	92	4 (4)	4 (4)	15 (16)	68 (74)	1 (1)
Fin. 122, Fall 2005	59	1 (2)	3 (5)	17 (29)	38 (64)	-
Fin. 122, Spr. 2006	82	5 (6)	10 (12)	28 (34)	39 (48)	-
Fin. 135, Spr. 2006	11	-	2 (18)	4 (36)	5 (45)	-
Fin.122 Fall 200699	5 (5)	3 (3)	27 (27)	64 (65)	-	-
Fin. 122, 135 Spr. 2007	102	5 (5)	3 (3)	21 (21)	71 (70)	2 (2)
MBAAd 250, Sum. 2007	32	2 (6)	5 (16)	10 (31)	15 (47)	-
Fin. 122, Fall 2007	150	3 (2)	6 (4)	46 (31)	93 (62)	2 (1)
Fin. 127, Spr. 2008	132	10 (8)	9 (7)	23 (17)	89 (67)	1 (1)
Fin. 127, Fall 2008	132	-	1 (1)	36 (27)	95 (72)	-
Fin. 127, 190, Spr. 2009	132	2 (2)	3 (2)	50 (38)	77 (58)	-
Total	1192	42 (4)	61 (5)	299(25)	781(66)	9 (1)

Turnitin generates “originality reports”, which indicate how much a paper overlaps with known outside sources. Papers range from 100% overlap to none. For most semesters after the first one, the average overlap was between 18 and 24%. Percent overlaps are strongly skewed. The average spiked at 29% for Spring 2006 (see table 3).

Table 3
Summary Statistics by Semester, Turnitin Data

	Mean Overlap	Standard Deviation	Number of Papers
Fall 2003	16.6%	0.160	65
Spring 2004	17.2%	0.235	58
Fall 2004	23.3%	0.238	46
Spring 2005	20.6%	0.219	92
Fall 2005	22.3%	0.149	59
Spring 2006	29.3%	0.213	93
Fall 2006	23.3%	0.210	99
Spring 2007	22.3%	0.203	102
Summer 2007	33.7%	0.194	32
Fall 2007	23.1%	0.171	150
Spring 2008	25.7%	0.240	132
Fall 2008	19.4%	0.095	132
Spring 2009	24.4%	0.136	132
Total (ex. Sum 07)	22.7%	0.191	1160
Total (w. Sum 07)	23.0%	0.192	1192

The first column, mean overlap indicates the percent of a student submission that overlaps known external documents. The second column is its standard deviation, the third column is the sample size. The percent overlap rises gradually, punctuated by minor spikes at Fall 2004 and Spring 2008, and a major spike at Spring 2006. These are also the three semesters that produced the preponderance of the academic integrity charges. The Summer of 2007 class is an outlier. However, it is worth noting that this is also the only MBA class in my sample, and so perhaps should not be compared to undergraduates.

Table 4
Disaggregated Percent Overlap, Turnitin Data
Last 3 semesters

	Overall Overlap	Publication Overlap	Web site Overlap	Student Paper Overlap
Spring 2008 (n=132)				
Mean	25.7	7.8	14.0	16.9
Median	17.5	5.0	9.0	10.0
Std. Dev.	24.0	9.2	14.3	22.7
Fall 2008 (n=132)				
Mean	19.4	9.7	11.7	11.9
Median	19.5	8.5	10.0	10.0
Std. Dev.	9.5	6.5	8.3	8.1
Spring 2009 (n=132)				
Mean	24.4	17.0	12.8	14.2
Median	21.0	14.0	10.5	12.0
Std. Dev.	13.6	11.7	9.6	9.5
Combined (n=396)				
Mean	23.2	14.2	10.1	14.3
Median	19.0	11.0	8.0	11.0
Std. Dev.	17.0	11.9	8.8	15.1

Turnitin reports a consolidated overlap as well as overlap with specific sources. The former is not simply the sum of the latter. For example, if a given 5% of the text were detected in both a journal article as well as a web site, both would be recorded, but the resulting overall overlap would still be only 5%. Percent overlap is generally skewed to the right. That is, the median is consistently lower than the mean because a small number of high overlap papers pulls the mean up compared to the median. Note further that Student Papers is the largest and most variable of the three categories, particularly for Spring 2008. Since SafeAssign does not include student papers beyond the host institution, this difference could account for the greater overall percentage for Turnitin compared to SafeAssign noted in the text above.