Course reviews reveal gender differences and other scientific insight about the students
who submit them

David Lang, Youjie Chen, Andreas Paepcke, Mitchell L. Stevens

Stanford University

Stanford, CA

{dnlang86, minachen, paepcke, stevens4}@stanford.edu

Abstract

Student course reviews are rarely considered as research instruments, yet their ubiquity makes them potentially powerful tools for education data science. To illustrate this potential we utilize a corpus of 11,255 reviews of computer science courses submitted by students at a private research university to observe how students appraise their own learning and give advice to future students. We recover evidence of gendered self-perceptions and care-giving strategies in Computer Science courses. Among submitters: females understate their achievement of learning goals relative to males earning the same grades; females offer lengthier written advice to future students than males; advice written by females exhibits more positive tone. Findings affirm the potential of leveraging course reviews for archival, survey, and quasi-experimental research investigations going forward.

Course reviews reveal gender differences and other scientific insight about the students who submit them

Introduction

Student course reviews are a controversial subject in academia. It is not uncommon for instructors to approach reviews of their own courses with some mixture of ambivalence, bemusement and curiosity. Considerable work has addressed problems of validity and bias in the use of these instruments for assessing instructors and instruction (Stark, Ottoboni, & Boring, 2016; Uttl, White, & Gonzalez, 2017). Meta-analyses have suggested that the relationship between instructional quality and student learning as measured by teaching evaluations may be tenuous (Uttl et al., 2017). Even more critical work exists suggesting that course reviews may be negatively correlated with students' learning in subsequent courses (Carrell & West, 2010). In light of these critiques, research on course evaluations spans questions of instrumentation (Rivera & Tilcsik, 2019), student participation (Chapman & Joines, 2017), and even the medium (paper versus digital) (Avery, Bryant, Mathios, Kang, & Bell, 2006) of course reviews. Yet few have recognized reviews as potentially useful tools for education data science.

Several features of course reviews make them potentially attractive for researchers. First, reviews are ubiquitous features of teaching and learning in US higher education. Because they are so commonly solicited and so frequently submitted, data yielded from reviews represent a very wide swath of student populations. Second, reviews are routinely submitted through online platforms and conducted by administrative units supported on hard budget lines, bringing the marginal cost of acquiring research data through reviews close to zero. Third, it is technically feasible to link information obtained through reviews with institutional data describing those who submit them. Thus while course reviews may be problematic means of assessing the quality of instructors and instruction – a matter on which we take no position here – we believe these instruments hold substantial promise for education data science. Our guiding premise is that course reviews offer potentially powerful instruments for learning about

students: how they make sense of their own learning, experience particular instructional environments, and pass advice on to others.

In this paper we offer illustrative case studies of the potential for course reviews to yield scientific insight by considering a corpus of 11,255 reviews submitted by students enrolled in Computer Science (CS) classes at a private research university during the 2015-16 and 2016-17 academic years. We chose CS as our domain because it is receiving a rapidly growing proportion of undergraduate enrollments at our case school and nationally (National Center for Education Statistics, 2016). We chose gender as a demographic focus partly in light of data availability, but also because gender asymmetries in CS and STEM persistence are important scientific and policy problems (Weeden, Gelbgiser, & Morgan, 2020; Beede et al., 2011). Yet by no means is the applicability of our analytic strategy limited to any particular field of study or student characteristic.

Even while each case study is grounded in different methods and domain expertise, all three are unified by a common theme of gender difference. In the first study we don the hat of a statistician and examine response rates of course review submission by grade and gender. In the second study we consider evaluations from the standpoint of a psychometrician, assessing whether particular components of an evaluation instrument exhibit a differential response bias by gender. In the third study we consider how a computational linguist might use natural-language methods to observe gender differences in how students describe their experiences in courses. We finish by offering a set of recommendations for how our insights might inform the collaborative design of evaluations in ways that simultaneously serve the interests of administrators, education data scientists, students, and instructors.

Background and Motivation

Institutional and academic research have long been separate domains in higher education, with the former associated with compliance and reporting functions and carried out by administrators, and the latter as scientific endeavors under the purview

of research faculty (Piety, Hickey, & Bishop, 2014). This distinction is rapidly dissolving, for several reasons. First, demands for outcomes-based accountability encourage administrative staff to measure student progress more scrupulously and more often (Taylor, Hanlon, & Yorke, 2013; Delaney, 2009). Second, research faculty are becoming increasingly adept at the computational methods necessary to extract knowledge from administrative data (Varian, 2014). Third, the growing interoperability of data streams and the rapidly declining costs of computer capacity are making previously unimaginable investigations of large data corpora easily tractable for administrators and faculty alike (Einav & Levin, 2014).

Recent work at the intersection of institutional and academic research in higher education has investigated gendered patterns of academic performance in large courses (Matz et al., 2017), variation in the complexity of requirements for undergraduate majors (Baker & Huntington-Klein, 2018), sequences of course-taking (Gruver et al., 2019), enhancing course search and discovery (Pardos & Jiang, 2020), and the forecasting of majors on the basis of early course-taking (Wang, Lang, Paepcke, & Stevens, 2020), among other topics. Such inquiries leverage academic transcript data that colleges and universities have legal and fiduciary obligations to retain (Hutt, 2016).

While they typically do not face legal requirements to conduct them, institutions nevertheless routinely – indeed almost universally – maintain protocols for soliciting feedback from students about their experiences in particular courses. While soliciting such reviews from students dates as early as the 1920s, the practice became fully diffused and normalized throughout the US postsecondary sector in the 1970s (Wachtel, 1998; Murray, 2005; J. England, Hutchings, & McKeachie, 1996). This secondary literature on course reviews makes clear that reviews have only occasionally been designed as scientific instruments, by which we mean driven by learning, psychometric or other theory and subject to academic peer review. Instead, course evaluations come from a tradition of enabling student voice to academic authority. Nevertheless reviews have become important components of instructional evaluation – some have argued excessively so (Spooren, Brockx, & Mortelmans, 2013). In the absence of other

measures of postsecondary instructional effectiveness and learning, student reviews are among very few streams of routinely collected information about how students experience particular college courses.

Prior has addressed how students and instructors respond to reviews. Students are more likely to submit reviews for courses in their majors (Adams & Umbach, 2012), and that female students are relatively more likely to respond to course evaluations than males (Sax, Gilmartin, & Bryant, 2003). Research manipulating reviews' informational content found positive effects on enrollment in courses when reviews have positive valence, particularly when there were large numbers of reviews to consider (Li & Wang, 2013). Parallel work found that exposure to positive or negative reviews had modest to large effects on students' self-reported expected performance in those same courses, and in students' likelihood of recommending those courses to others in the future (Kowai-Bell, Guadagno, Little, Preiss, & Hensley, 2011). Instructors, too, absorb the valence of review content, especially regarding information about instructor mood and affect (Kowai-Bell et al., 2011).

An important concern about reviews from the psychometric literature is the potential for differential item function, a phenomenon in which respondents of equal ability exhibit different responses to a given survey item or question. Studies of differential item function in course reviews have focused on the quantitative difficulty of a class and on characteristics of the instructor (Narayanan, Sawaya, & Johnson, 2014) (DeFrain & DeFrain, 2016). Even the range of a Likert scale can have significant affects on how students assess the relative quality of male and female instructors, for example (Rivera & Tilcsik, 2019) found that reducing the number of likert points reduced the gender gap in course evaluations.

Virtually the entire corpus of scholarship on course reviews addresses either their value as instruments for assessing courses, instructors and instruction, or as mechanisms for influencing student or instructor behavior – not as mechanisms for learning about the students who submit them. Yet is good *prima facie* reason to suspect that reviews might be instrumented for scientific inquiry about their submitters. The ubiquity and

frequency of reviews make them readily available vehicles for serial data collection, their anonymity (to the instructor) creates general conditions for candor, and it is technically possible to link submissions with institutional metadata carrying potentially salient characteristics of submitters.

Reviews can only be useful for scientific inquiry, however, to the extent that at least three conditions can be met. The first is representativeness: the statistical relationship between submitter and overall student populations must be knowable. The second is instrument integrity: reviews and/or their sub-components must be amenable to psychometric scrutiny. The third is substantive and conceptual value: reviews must yield information that can be leveraged for empirical insight that contributes to educational improvement and scientific discourse.

Gender and Care Work in Course Reviews

We take the substantive problem of women's relative under-representation in Computer Science education as the substantive context for considering the scientific potential of course reviews. According to a 2019 report from the National Science Foundation, women's percentage of degree recipients in CS continues to be the lowest among the broad field of science and engineering, and has been decreasing in recent decades (NSF, 2019). This fact has far reaching consequences. As a field that often leads to lucrative careers (Gomis, Kapsos, & Kuhn, 2020), women's under-representation in the major may impede their potential incomes after graduation. It may also perpetuate male bias in beliefs about CS "talent," and in the selection and definition of practical problems to which computer scientists, and the myriad firms which employ them, attend (Beyer, Rynes, Perrault, Hay, & Haller, 2003; Kumar, 2012).

Prior research on CS education has found supporting evidence of gender variance in student experiences. For example, women are less comfortable asking questions and have less confidence in CS courses than their male peers (Redmond, Evans, & Sahami, 2013). This "gender confidence gap" grows as students take more advanced courses (Alvarado, Cao, & Minnes, 2017). Additionally, analyses of communal academic

resources in CS programs find substantial differences in how contributions by male and female users are acknowledged in *Github* and *StackOverflow* (May, Wachs, & Hannák, 2019) (Terrell et al., 2017). Course reviews may provide additional insight into gender variance in how students make sense of their academic exposure to computer science.

Further, gendered differences in skill perception may influence how students report learning gains in the same courses. While women tend to approach STEM fields with relatively less confidence than men, men tend to over-estimate their abilities.

Experimental work by Correll (2004) found that men expressed inflated perceptions of their own skill at completing quantitative tasks compared to women performing at the same level of measured accomplishment. This insight suggests that course reviews may bear traces of gendered patterns of academic self-perception.

Finally, to the extent that students submit course reviews as a service to instructors and future students, we might expect review corpora to exhibit gendered patterns common in care work. Social scientists define *care work* as that which attends to the well-being of others. It comprises activities intended to help other people develop their capabilities and pursue their goals (England, 2005). Care work is consistently associated with femininity and female role expectations, is often unpaid or poorly compensated, and disproportionately shouldered by women (Folbre, 1995). To the extent that submitting course reviews is an act of assistance – to improve classes and to inform future students – it may be theorized as a form of care work. Thus we might expect that female and male students will approach the task of course reviews with different dispositions, such that the number, extensiveness, and content of course evaluations may vary systematically by gender.

Data

To investigate the scientific potential of course reviews for informing substantive educational phenomena, we conduct a series of analyses on a data corpus comprising course reviews and institutional metadata of all students enrolled in courses offered by the Computer Science department of a private research university during the 2015-16

and 2016-17 academic years. Reviews were administered near the end of each academic term but before the beginning of the official final exam period. As an incentive for submitting reviews, submitters were given the ability to view their final course grades about a week earlier than non-submitters.

In total these data include 11,255 student responses from 251 courses. Courses range in character from very large introductory lecture-and-lab formats to small advanced seminars. Institutional metadata made available to us for analysis include each student's grade, gender (M/F), course grades, and declared major (if known). Approximately one-third of reviews submitted to CS courses during the study period were from female students, and approximately half were from undergraduates. We cannot track or identify students enrolling in multiple CS courses during the study period, however we can compute and generate response rates by course grade and gender.

We limit our analysis to responses in which submitters offered a response to the review protocol's only open-ended question, which reads:

"What would you like to say about this course to a student who is considering taking it in the future?"

This prompt is highly appropriate for investigating review submissions as a form of care work because it specifically asks submitters to give advice to a hypothetical other. Individual responses to this prompt vary substantially in length: from a single character to over 5,964 characters (the latter equivalent to 1004 words). The mean response length is 132 characters – approximately the length of a tweet. The entire corpus of responses in our study sample is 266,746 words.

We also analyze responses to two review prompts designed with five-point Likert responses:¹

- How much did you learn from this course?
- How well did you achieve the learning goals of this course?

¹ "One of the two questions was not consistently administered across courses. We limit our analysis to cases where both questions were answered. This step reduced our total to 9,002 cases"

Aggregated responses to these two prompts appear in Figure 1.

Analyses

We conducted three case studies, each of which investigates at least one of the three conditions course reviews must meet to be scientifically valuable: representativeness, instrument integrity, and substantive and conceptual value.

Study One: Representativeness

A paramount question of any data corpus is its representativeness of populations of interest. Researchers considering corpora of course reviews may worry that students who submit reviews differ from all students who took the courses in some material way. Even if this is not the case, there is still concern that certain voices may be amplified and others diminished through variation in submission. Knowing the relationship between submitters and all students is essential if researchers are to derive insights that generalize to relevant student populations. We here investigate the extent to which submitters in our study sample differ from all students on two dimensions: course grade and gender.

Response Rates by Course Grade. There are competing intuitions about how grades might influence rates of review submission. Those anticipating poor grades may be more inclined to approach the review task as an opportunity to retaliate against the grader. Alternately, students who receive low grades may opt to avoid occasions to reflect on negative experiences. Finally, students may perceive of reviews as a kind of gift to instructors; those receiving low marks may decide that instructors are not worthy recipients. We observe the relationship between submission and course grade empirically utilizing a simple chi-squared test of response rates by course grade.

Figure 2 indicates a strong positive correlation between course grade and likelihood of review submission. A student with a grade of A or higher had an 80% likelihood of submitting a review, while students who do not pass the course or receive credit without a grade responded approximately 50% of the time. Effectively this means that students who failed the relevant course are represented in our corpus about half as

often as students who excelled in the relevant course. Differences are statistically significant, with a chi-squared statistic of 395.37 and a p-value of less than 0.001.

Response Rates by Gender. Theories of care work suggest that female students will be more likely to respond to requests for reviews. To the extent that reviews represent feedback to instructors that might aid in course improvement, and advice to future students who might be considering enrollment, the task of submitting reviews may have a feminine cast that encourages female participation more strongly. We observe the relationship between submission and gender empirically using an exact binomial-two-sample test. We examine results by gender and grade simultaneously.

Rates of review submission by student gender and earned grade are reported in Figure 2. Two features are notable. First, females are more likely to submit reviews overall. On average, females submit reviews at a rate of 78.0%, vs. males, whose rate is 74.5% ($p < .001$).

Second, those receiving higher grades in a course are more likely to submit reviews. Females receiving a grade of "A" are 3.7% more likely to respond to submit than their male counterparts earning the same grades. The gender submission gap is greatest among students receiving a grade of "B," with female "B" recipients 6.5% ($p < .001$) more likely to submit than males. We do not observe statistically significant gender differences in submission rates for those receiving grades below "B," however such grades represent fewer than 5% of grades in the research sample.

Findings of this first study illustrate that course reviews meet our first condition for scientific utility. It is possible, indeed quite simple with institutional metadata, to assess the representativeness of review submitters in relation to relevant student populations. These findings also meet our third utility condition, substantive/conceptual value, in that they yield insight about course reviews as a form of care work.

In light of these findings, we might imagine collaborative inquiry involving institutional researchers and gender scholars about how the act of submitting reviews might be explicitly framed: whether as a "contribution," "obligation," or "responsibility"

of course participation, for example. Experimental manipulation of definitions of the review situation would enable observation for potential gender variation in response rates. Such work would be of both institutional and scholarly relevance.

Study Two: Item Response

A second condition for the scientific utility of course reviews is instrument integrity: the extent to which components of reviews can be assessed for validity and reliability in relation to the underlying phenomena they purport to tap. To investigate the integrity of the review instrument from which our study corpus was derived, we rely on the concept of differential item function from the field of psychometrics. The basic idea behind differential item function is that prompts may be written in ways that induce participating subjects to respond variably in relation to focal characteristic (e.g. gender, race, native language), even when subjects have the same underlying ability with the construct of interest (e.g., competencies in math or language). A famous example of differential item function is an SAT analogy task in which subjects who grew up in more affluent backgrounds were better able to correctly answer analogies that involved activities less common in lower-income milieux (Geiser, 2017). Differential item function has been a serious concern in addressing issues of equity and gender in rigorous standardized testing (Martinková et al., 2017).

As we discussed above, course reviews typically are carried out by staff administrators for institutionally specified purposes, rather than by research faculty for scientific purposes, so we would not necessarily expect review items to be tested for differential item function. Further, reviews typically have a one-size fits all design. Items might serve well for some courses and academic programs, but not others. Regardless, if reviews are to be used for scientific purposes, their items must be amenable to assessments of instrument integrity.

In this study we examine responses to two review items which prompt students to respond via Likert scales to the following:

- Learning- How much did you learn from the course?

- Achievement- How well did you achieve the learning goals of this course?

Both of these items pertain to learning, but male and female students may respond differently to them. In light of social-psychological research on academic self-concept, we suspect that males may tend to over-estimate their learning net of grades, while women may tend to under-estimate their learning net of grades (Correll, 2004). To investigate, we deploy a fixed-effect regression model with the following specification:

$$Y_{ij} = \beta_1 Male_i + Grades_{ij} + \Gamma_j + \epsilon_{ij} \quad (1)$$

The subscripts i and j correspond to indices for student and course. The Y variable corresponds to our focal outcome variable, in this case, responses to a Likert item. $Male$ corresponds to institutional metadata reporting student gender, and β_1 corresponds to the associated coefficient with this variable. We represent course effects with Γ_i to control for factors like the difficulty of the course or instructional quality. We also control for course grades with an additional fixed-effect for each possible grade a student might receive.² The error is represented by ϵ . Errors are clustered at the course level.

We observe the proportion of students reporting having **achieved the learning goals** specified by the course instructor *Extremely Well* by grade in Figure 3. Unsurprisingly, we find a strong direct correlation with course grade, such that reported goal achievement declines with grade. What is striking is that at every grade level, there is a gender gap in reported goal achievement, with males more likely to report achievement than females earning the same grade.

We extend this analysis to see if this same pattern occurs with the question of **how much students learn**, using the same specification as described in equation 1 in Table 1. We see that after controlling for grades and course, males and females exhibit no differences in self-reported measures of ‘how much they learned’ in a class.

However, when we look at the similar question about ‘achievement of learning goals,’ we see a stark difference. Male students are 8% points more likely than females

² At our case school there are over twenty grade types, including + and - variants as well as credit and noncredit courses. We aggregate to A,B,C,D, and not-passing grades for simplicity

to state they mastered the learning goals of a course. This finding suggest two concerns. First, given the similarity of these questions, it appears that subtle differences in phrasing yield substantial differences in student response by gender. Second, female students report relatively lower levels of mastery even net of grades.

Findings of this second study illustrate that course reviews meet our second condition for assessment of instrument integrity. When they are linked with metadata describing salient student characteristics (course grades and gender, in our case), reviews are readily amenable to psychometric tests for such phenomena as differential item function. This study also meets our third utility condition, providing evidence that gender is implicated in how students appraise their own learning.

In light of these findings, we might imagine a collaboration between institutional researchers, psychometricians and gender scholars about how instrument design influences self-reports of student learning. Experimental manipulation of items would enable observation of gender variation in responses, yielding insight that could inform selection of prompts that elicit the most gender-equitable replies.

Study 3: Computational Reading

Here we illustrate the potential for natural language processing techniques to yield tractable insight into student subjectivity from qualitative components of course reviews. We present three ways of quantifying and describing responses to the question, "What would you like to say about this course to a student who is considering taking it in the future?" We first observe variation in the length of responses and in the frequency with which responses mention instructors by name. Second, we use a prebuilt dictionary to quantify the tone, authenticity, and authority review texts exhibit. Third, we summarize and group responses into topical categories using unsupervised topic models. We continue to key our observations to the intuition that review submissions may betray systematic differences by gender.

Who invests more in writing course reviews? Theories of care work suggest that women may invest more time and effort in giving advice to others.

Specifically, women may write lengthier responses with more specific details. We operationalize these ideas by measuring (a) the length of submissions by number of words and (b) the frequency with which submissions mention an instructor by name. The first measure (a) is a fairly straightforward operationalization of investment. The second measure (b) provides a proxy for submitters' attention to the specifics of a particular instance of a course. We use the same regression model presented in Study 2 to observe variation between male and female submitters, once again controlling for course grade. The only difference in the model is the outcome variable (word count, instructor mention).

We report the standardized results of our analysis in Table 2. Submissions from women are lengthier – about .15 of a standard deviation or 5 words longer. We observe significant differences in the frequency with which submissions from males and females mention instructor name, with women mentioning instructor by name by approximately two percentage points. To put this number in context, mentioning of instructor name is an infrequent event where only about 1 in 10 evaluations explicitly mentions the instructor. While modest in magnitude, these statistically significant findings comport with our care work hypotheses that women may approach the task of submitting reviews with more investment and attention to specificity.

What is the tone of textual submissions? In addition to quantifying review texts based on word count and handcrafted-rules, researchers can also use dictionary methods to describe these data. The Linguistic Inquiry and Word Count (LIWC) (Pennebaker, Boyd, Jordan, & Blackburn, 2015) is a popular dictionary with a component that is particularly appropriate for use in our case. In addition to grouping words into 75 distinct categories and themes (e.g. family, power, death, e.g.), LIWC generates four psycho-social variables that were validated on college application essays through a rating process. Each of these variables is scored on a 1 to 99 interval where 1 is a complete lack of the construct and 99 is a highly pronounced form of the construct. The constructs are:

1. **Tone**- a summary variable describing the emotional quality of the text. A score of

99 reflects a positive tone and a score of 1 reflects a negative tone. A score of 50 represents neutral valence.

2. **Analytic**- a measure of how much formal logic is used in the text. A score of 1 indicates little use of formal logic and a score of 99 indicates statement with a great deal of formal logic.
3. **Authenticity**- a measure of the sincerity/honesty of a text. A score of 1 indicates insincerity and a score of 99 indicates high sincerity.
4. **Clout**- a measure of the text's authority, status, and confidence. A score of 1 suggests relatively little authority and a score of 99 suggests high authority.

We observe for potential gender differences on these four psycho-social variables after controlling for variation that can be attributed either to the course overall or to each submitter's earned course grade. To aid interpretability, we report outcomes in standardized effect sizes. In light of social-psychological theories of gender and academic self-concept, as well as Study 2's suggestive findings, we suspect that men's relatively stronger academic self-concepts may empower them to write more critical reviews, manifesting in relatively lower scores on *Tone*. At the same time, women's relatively weaker academic self-concepts may result in greater numbers of personally qualified "I" statements and tentative language, and manifest in higher scores on the LIWC measure *Authenticity*.

Results are reported in Table 3. We find modest variation by gender in how submitters describe their experience in the same course, conditional on grades. On average, submissions from males evince slightly more negative and slightly less authentic language. While these gender differences are highly significant, their magnitude is modest: on the order of a tenth of a standard deviation. Nevertheless, they are consistent with prior social-psychological insight and with the results of Study 2.

What do course reviews say? Here we utilize structural topic models to observe whether qualitative components of the corpora vary systematically by gender. The basic premise behind this method is that texts include multiple topics, and topics

comprise words. Using unsupervised methods such as Latent Dirichelet Allocation (LDA), we can group texts into topical categories. Additionally, *structural topic modeling* enables us to observe variance in topics by metadata describing other characteristics of authors (Roberts et al., 2014).

We ran structural topic models while allowing submitter gender to vary with topic prevalence. We tuned the optimal number of topics from 2 to 50 using an exclusivity measure called FREX (Airoldi & Bischof, 2012) (Reich, Tingley, Leder-Luis, & Roberts, 2014) (see Equation 2). FREX is a harmonic weighting of the frequency (F) with which a word occurs in a topic; exclusivity (E), how frequently a word occurs in a given topic relative to others. The parameter ω corresponds to a tuning parameter of the relative importance of these features. We used the default parameter of $\omega = .7$ to favor topics that had more exclusivity.

$$FREX = \left(\frac{\omega}{E} + \frac{1 - \omega}{F} \right)^{-1} \quad (2)$$

Using this criterion, we found the locally optimal parameter to be 32 distinct topics. We then hand-labelled each topic, observing the ten (10) statements that had the highest probability labels of that topic. See Figure 4 (Top). The most common topics were about a course being a tutorial, a suggestion to take the course, or a positive review. The least common topics were highly specific suggestions and issues pertaining to course prerequisites.

While we did not see substantial gender variation in topics overall, there are exceptions of note. First, submissions from males are more likely to invoke math prerequisites and claims that instruction was poorly organized or of poor quality. Submissions from men also are more likely to invoke course organization and instruction. Submissions from women are more likely to invoke workload, study practices, and attendance. These observations correspond with a consistent pattern across all three studies, which provide evidence that women invest more effort and empathy in course reviews and men evince more academic self-confidence – the latter expressed here as a greater tendency to critique courses and instructors.

There is an important caveat to these findings, however. In contrast with all of

the other above studies, the topic model results presented in this section do not control for course or submitter grades or course selections, because our corpus did not have the scale necessary to disaggregate on the basis of additional variables simultaneously.

Overall, results of Study 3 provide additional illustration that student course reviews can provide insight into phenomena that are substantively important in education and conceptually valuable for social science. In light of these findings we can imagine collaborative work between institutional researchers, linguists and other social scientists in which review texts are examined across a wider array of demographic characteristics of applicants – first-generation status or ethno-racial identity, for example – and also across disciplinary domains, to yield insight about how different kinds of students make sense of different kinds of courses.

Discussion

Even while they are controversial for evaluating instructors and instruction, course reviews are ubiquitous features of the US higher education landscape and potentially powerful tools for education data science. In the work presented here we have sought to demonstrate the promise of reviews as windows into students' perceptions of their academic experiences and their orientation to assisting others. Taking advantage of archival data that included 11,255 evaluations submitted to 251 computer science courses at a single university between 2015-2017 that was linked to administrative information describing submitters' gender (M/F) and course grades, we observed patterned variation in who submits course reviews, and how.

In three observational studies we found that (1) women and those earning high grades were disproportionately likely to submit reviews (2) the phrasing of close-ended review prompts influenced patterns of response by gender (3) responses to qualitative review prompts differed subtly but significantly by gender, with women writing somewhat more positive, individualized, and lengthier reviews. These empirical findings comport with theoretical insights from educational social psychology and feminist social science, which suggest gender variation in how men and women perceive of their own

academic accomplishments and their attendance to the well-being of others.

Our work suggests the promise of leveraging course reviews for cumulative science in at least three ways. First, the inquiries presented here are based entirely on the premise that course reviews and submitter demographic information are "found" data. While our topical focus was gender differences in Computer Science, we would be remiss if we did not discuss the potential for more capacious utilization of reviews and institutional metadata. The work we presented here is an proof-of-concept, focusing on students' experiences of courses in a single department. Research corpora could be readily expanded for studies of how gender inflects review submissions across the entire curriculum. Corpora also could be expanded temporally to investigate whether student perceptions and reported experiences change over time. It would be technically quite simple, for example, to structure review data in a panel format and observe how the same students respond to reviews in sequences of courses across their college careers. Such longitudinal inquiry might yield insight into how students' academic self-concepts evolve with their accumulating academic experiences.

To the extent that virtually every US college and university possesses data such as these, we can only imagine the number and variety of insights that might be gained from parallel investigations at other schools. To a nascent field whose promise lies substantially in observing phenomena at scale, course reviews provide exceptionally promising mechanisms for education data science. Such work may serve as an instructional companion to the work of bibliometric researchers who study scientific discourse to observe how academic careers and fields evolve (Munoz-Najar Galvez, Heiberger, & McFarland, 2020).

Second, as we suggested serially above, there is good reason to imagine that education data scientists might collaborate with school administrators to more explicitly and conscientiously instrument reviews for systematic experimental and quasi-experimental research. The basic conditions for such inquiries already are in place and sustained by established administrative rhythms: schools have offices conducting reviews, students anticipate receiving them, and they are administered multiple times a

year. This infrastructure could support substantial scientific insight through the linkage of review submissions with administrative data describing characteristics of submitters. Moreover, the fact that course evaluations are an empirical context in which the interests of institutional researchers, psychometricians, and a wide range of social scientists intersect, it is easy to imagine productive joint action. The initial efforts presented here provide an inkling of this promise.

Third, there is the strong potential for such joint action to be reciprocally productive for educational improvement and scientific advance. To the extent that students read reviews when considering courses, instructors consult or game them, and administrators incorporate them into personnel evaluations, instrumenting reviews for responsible science might be of great service to educational improvement. The findings presented here offer potentially useful takeaways for institutional researchers and academic stakeholders. Results of Study 1 suggest that re-weighting of student responses' by gender and grade may reduce some of the missingness bias that exists within this type of "found" data. Results of Study 2 on suggest that the wording of questions can substantially influence how responses are related to underlying student characteristics. Results of study 3 indicate the promise of computational text techniques to parse insight from very large corpora of review texts.

While building collaborative research programs of the kind suggested here would be straightforward to scaffold technically, pursuing them would require considerable organizational investment and will. To the extent that institutional and academic researchers occupy different organizational niches, speak somewhat disparate analytic languages, and are oriented to separate professional communities, collaboration between them will take patience and humility on both sides. Utilizing reviews to pursue open science also will raise new questions for data governance as well as instructor and student privacy that champions of education data science are only now beginning to fully specify (Brown & Klein, 2020; Kurzweil & Stevens, 2018).

The payoff for tackling these matters will be well worth it. With a little work, course reviews could be thoughtfully instrumented and combined with enrollment data

to concisely depict how students accumulate and make sense of their academic experiences over time. That might help move the needle on some of the sector's largest challenges. It is no secret that nearly half of US students seeking four-year diplomas leave college before earning their degrees (NCES, 2018). The most advanced machine learning algorithms may reveal which students are unlikely to persist, but cannot identify the classroom experiences and perceptions that precede early college exit. Carefully designed course reviews could capture those experiences and perceptions for each and every course students take. They might be the first step in transforming a problematic means of evaluating instructors into an indispensable tool of postsecondary improvement.

References

- Adams, M. J. D., & Umbach, P. D. (2012, 8). Nonresponse and Online Student Evaluations of Teaching: Understanding the Influence of Salience, Fatigue, and Academic Environments. *Research in Higher Education*, 53(5), 576–591. Retrieved from <http://link.springer.com/10.1007/s11162-011-9240-5> doi: 10.1007/s11162-011-9240-5
- Airoidi, E. M., & Bischof, J. M. (2012). A Poisson convolution model for characterizing topical content with word frequency and exclusivity. *arxiv.org*. Retrieved from <https://arxiv.org/abs/1206.4631> <http://arxiv.org/abs/1206.4631>
- Alvarado, C., Cao, Y., & Minnes, M. (2017, 3). Gender Differences in Students' Behaviors in CS Classes throughout the CS Major. In *Proceedings of the 2017 acm sigcse technical symposium on computer science education* (pp. 27–32). New York, NY, USA: ACM. Retrieved from <https://dl.acm.org/doi/10.1145/3017680.3017771> doi: 10.1145/3017680.3017771
- Avery, R. J., Bryant, W. K., Mathios, A., Kang, H., & Bell, D. (2006). Electronic course evaluations: does an online delivery system influence student evaluations? *The Journal of Economic Education*, 37(1), 21–37.
- Baker, R., & Huntington-Klein, N. (2018). Student preference for guidance and complexity in college major requirements. cepa working paper no. 18-06. *Stanford Center for Education Policy Analysis*.
- Beede, D. N., Julian, T. A., Langdon, D., McKittrick, G., Khan, B., & Doms, M. E. (2011). Women in stem: A gender gap to innovation. *Economics and Statistics Administration Issue Brief*(04-11).
- Beyer, S., Rynes, K., Perrault, J., Hay, K., & Haller, S. (2003). Gender differences in computer science students. In *Proceedings of the 34th sigcse technical symposium on computer science education* (pp. 49–53).
- Brown, M., & Klein, C. (2020). Whose data? which rights? whose power? a policy discourse analysis of student privacy policy documents. *The Journal of Higher*

Education, 1–30.

- Carrell, S. E., & West, J. E. (2010). Does professor quality matter? evidence from random assignment of students to professors. *Journal of Political Economy*, 118(3), 409–432.
- Chapman, D. D., & Joines, J. A. (2017). Strategies for increasing response rates for online end-of-course evaluations. *International Journal of Teaching and Learning in Higher Education*, 29(1), 47–60.
- Correll, S. J. (2004, 2). Constraints into Preferences: Gender, Status, and Emerging Career Aspirations. *American Sociological Review*, 69(1), 93–113. Retrieved from <http://journals.sagepub.com/doi/10.1177/000312240406900106> doi: 10.1177/000312240406900106
- DeFrain, E., & DeFrain, E. (2016). An Analysis of Differences in Non-Instructional Factors Affecting Teacher-Course Evaluations over Time and Across Disciplines. Retrieved from <https://repository.arizona.edu/handle/10150/621018>
- Delaney, A. M. (2009, 6). Institutional researchers' expanding roles: Policy, planning, program evaluation, assessment, and new research methodologies. *New Directions for Institutional Research*, 2009(143), 29–41. Retrieved from <http://doi.wiley.com/10.1002/ir.303> doi: 10.1002/ir.303
- Einav, L., & Levin, J. (2014). The data revolution and economic analysis. *Innovation Policy and the Economy*, 14(1), 1–24.
- England. (2005, aug). Emerging Theories of Care Work. *Annual Review of Sociology*, 31(1), 381–399. Retrieved from <http://www.annualreviews.org/doi/10.1146/annurev.soc.31.041304.122317> doi: 10.1146/annurev.soc.31.041304.122317
- England, J., Hutchings, P., & McKeachie, W. J. (1996). *The professional evaluation of teaching* (No. 33). American Council of Learned Societies.
- Folbre, N. (1995, 3). “Holding hands at midnight”: The paradox of caring labor. *Feminist Economics*, 1(1), 73–92. Retrieved from <http://www.tandfonline.com/doi/abs/10.1080/714042215> doi:

10.1080/714042215

- Geiser, S. (2017). Norm-referenced tests and race-blind admissions: The case for eliminating the sat and act at the university of california. research & occasional paper series: Cshe. 15.17. *Center for Studies in Higher Education*.
- Gomis, R., Kapsos, S., & Kuhn, S. (2020). World employment and social outlook: trends 2020.
- Gruver, N., Malik, A., Capoor, B., Piech, C., Stevens, M. L., & Paepcke, A. (2019). Using latent variable models to observe academic pathways. *arXiv preprint arXiv:1905.13383*.
- Hutt, E. (2016). *A brief history of the student record*. New York: Ithaka S+ R. Retrieved from <https://doi.org/10.18665/sr>.
- Kowai-Bell, N., Guadagno, R. E., Little, T., Preiss, N., & Hensley, R. (2011, 9). Rate My Expectations: How online evaluations of professors impact students' perceived control. *Computers in Human Behavior*, 27(5), 1862–1867. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0747563211000756> doi: 10.1016/J.CHB.2011.04.009
- Kumar, A. N. (2012). A study of stereotype threat in computer science. In *Proceedings of the 17th acm annual conference on innovation and technology in computer science education* (pp. 273–278).
- Kurzweil, M., & Stevens, M. (2018). Setting the table: Responsible use of student data in higher education. *Educause Review*, 53(3), 16–24.
- Li, C., & Wang, X. (2013, 7). The power of eWOM: A re-examination of online student evaluations of their professors. *Computers in Human Behavior*, 29(4), 1350–1357. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0747563213000101> doi: 10.1016/J.CHB.2013.01.007
- Martinková, P., Drabinová, A., Liaw, Y.-L., Sanders, E. A., McFarland, J. L., & Price, R. M. (2017). Checking Equity: Why Differential Item Functioning Analysis Should Be a Routine Part of Developing Conceptual Assessments. *CBE life*

- sciences education*, 16(2). Retrieved from
<http://www.ncbi.nlm.nih.gov/pubmed/28572182>
<http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5459266>
 doi: 10.1187/cbe.16-10-0307
- Matz, R. L., Koester, B. P., Fiorini, S., Grom, G., Shepard, L., Stangor, C. G., . . . McKay, T. A. (2017). Patterns of gendered performance differences in large introductory courses at five research universities. *AERA Open*, 3(4), 2332858417743754.
- May, A., Wachs, J., & Hannák, A. (2019, 8). Gender differences in participation and reward on Stack Overflow. *Empirical Software Engineering*, 24(4), 1997–2019. Retrieved from <http://link.springer.com/10.1007/s10664-019-09685-x> doi: 10.1007/s10664-019-09685-x
- Munoz-Najar Galvez, S., Heiberger, R., & McFarland, D. (2020). Paradigm wars revisited: A cartography of graduate research in the field of education (1980–2010). *American Educational Research Journal*, 57(2), 612–652.
- Murray, H. G. (2005). Student evaluation of teaching: Has it made a difference. In *Annual meeting of the society for teaching and learning in higher education* (pp. 1–15).
- Narayanan, A., Sawaya, W. J., & Johnson, M. D. (2014, 7). Analysis of Differences in Nonteaching Factors Influencing Student Evaluation of Teaching between Engineering and Business Classrooms. *Decision Sciences Journal of Innovative Education*, 12(3), 233–265. Retrieved from
<http://doi.wiley.com/10.1111/dsji.12035> doi: 10.1111/dsji.12035
- National Center for Education Statistics. (2016). *Bachelor's degrees conferred by postsecondary institutions, by field of study: Selected years, 1970-71 through 2014-15*. Retrieved from
https://nces.ed.gov/programs/digest/d19/tables/dt19_322.10.aspx
https://nces.ed.gov/programs/digest/d16/tables/dt16_322.10.asp
- NCES. (2018). *Fast Facts: Graduation rates (40)*. Retrieved 2020-08-21, from

<https://nces.ed.gov/fastfacts/display.asp?id=40>

- NSF. (2019). *Women, minorities, and persons with disabilities in science and engineering. special report nsf 19-304*. National Science Foundation Alexandria, VA.
- Pardos, Z. A., & Jiang, W. (2020). Designing for serendipity in a university course recommendation system. In *Proceedings of the tenth international conference on learning analytics & knowledge* (pp. 350–359).
- Pennebaker, J., Boyd, R., Jordan, K., & Blackburn, K. (2015). The development and psychometric properties of LIWC2015. Retrieved from <https://repositories.lib.utexas.edu/handle/2152/31333>
- Piety, P. J., Hickey, D. T., & Bishop, M. (2014). Educational data sciences: Framing emergent practices for analytics of learning, organizations, and systems. In *Proceedings of the fourth international conference on learning analytics and knowledge* (pp. 193–202).
- Redmond, K., Evans, S., & Sahami, M. (2013). A large-scale quantitative study of women in computer science at Stanford University. In *Proceeding of the 44th acm technical symposium on computer science education - sigcse '13* (p. 439). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=2445196.2445326> doi: 10.1145/2445196.2445326
- Reich, J., Tingley, D., Leder-Luis, J., & Roberts, M. (2014). Computer-assisted reading and discovery for student generated text in massive open online courses. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2499725
- Rivera, L. A., & Tilcsik, A. (2019). Scaling down inequality: Rating scales, gender bias, and the architecture of evaluation. *American Sociological Review*, 84(2), 248–274.
- Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., ... Rand, D. G. (2014, 10). Structural Topic Models for Open-Ended Survey Responses. *American Journal of Political Science*, 58(4), 1064–1082. Retrieved from <http://doi.wiley.com/10.1111/ajps.12103> doi: 10.1111/ajps.12103

- Sax, L. J., Gilmartin, S. K., & Bryant, A. N. (2003). *Assessing response rates and nonresponse bias in web and paper surveys* (Vol. 44) (No. 4). Retrieved from <https://link.springer.com/article/10.1023/A:1024232915870> doi: 10.1023/A:1024232915870
- Spooren, P., Brockx, B., & Mortelmans, D. (2013). On the validity of student evaluation of teaching: The state of the art. *Review of Educational Research*, 83(4), 598–642.
- Stark, P., Ottoboni, K., & Boring, A. (2016). Student Evaluations of Teaching (Mostly) Do Not Measure Teaching Effectiveness. *ScienceOpen Research*. Retrieved from https://www.scienceopen.com/document_file/25ff22be-8a1b-4c97-9d88-084c8d98187a/ScienceOpen/3507xE6680747344554310733.pdf doi: 10.14293/s2199-1006.1.sor-edu.aetbzc.v1
- Taylor, J., Hanlon, M., & Yorke, M. (2013). The evolution and practice of institutional research. *New Directions for Institutional Research*, 2013(157), 59–75.
- Terrell, J., Kofink, A., Middleton, J., Rainear, C., Murphy-Hill, E., Parnin, C., & Stallings, J. (2017, 5). Gender differences and bias in open source: pull request acceptance of women versus men. *PeerJ Computer Science*, 3, e111. Retrieved from <https://peerj.com/articles/cs-111> doi: 10.7717/peerj-cs.111
- Uttl, B., White, C. A., & Gonzalez, D. W. (2017, 9). Meta-analysis of faculty's teaching effectiveness: Student evaluation of teaching ratings and student learning are not related. *Studies in Educational Evaluation*, 54, 22–42. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0191491X16300323> doi: 10.1016/J.STUEDUC.2016.08.007
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3–28.
- Wachtel, H. K. (1998). Student evaluation of college teaching effectiveness: A brief review. *Assessment & Evaluation in Higher Education*, 23(2), 191–212.
- Wang, A., Lang, D., Paepcke, A., & Stevens, M. (2020). A natural-language approach to early forecasting of undergraduate majors: An investigation of efficacy and bias

for serendipity in a university course recommendation system. In *Workshop proceedings of educational data mining fairness and accountability workshop*.

Weeden, K. A., Gelbgiser, D., & Morgan, S. L. (2020, 6). Pipeline Dreams:

Occupational Plans and Gender Differences in STEM Major Persistence and Completion. *Sociology of Education*, 003804072092848. Retrieved from

<http://journals.sagepub.com/doi/10.1177/0038040720928484> doi:

10.1177/0038040720928484

	Learning	Achievement
Male	0.01	0.08***
	(0.01)	(0.01)
Num. obs.	9,002	9,002
R ² (full model)	0.15	0.18
R ² (proj model)	0.00	0.01
Adj. R ² (full model)	0.13	0.16
Adj. R ² (proj model)	−0.03	−0.02
Num. groups: Grades	20	20
Num. groups: Courses	203	203

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 1

Gender Differences in Likert Items

	ProfessorName	Word Count
Male	-0.02^*	-0.15^{***}
	(0.01)	(0.03)
Num. obs.	11, 255	11, 255
R ² (full model)	0.08	0.08
R ² (proj model)	0.00	0.00
Adj. R ² (full model)	0.06	0.06
Adj. R ² (proj model)	-0.02	-0.02
Num. groups: Grades	20	20
Num. groups: Courses	251	251

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 2

Handcrafted Features Regressions

	Tone	Analytic	Clout	Authentic
Male	-0.08^{**}	-0.01	0.01	-0.10^{***}
	(0.03)	(0.02)	(0.02)	(0.02)
Num. obs.	11, 255	11, 255	11, 255	11, 255
R ² (full model)	0.07	0.05	0.05	0.04
R ² (proj model)	0.00	0.00	0.00	0.00
Adj. R ² (full model)	0.05	0.02	0.02	0.02
Adj. R ² (proj model)	-0.02	-0.02	-0.02	-0.02
Num. groups: Grades	20	20	20	20
Num. groups: Courses	251	251	251	251

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 3

Linguistic Inquiry and Word Count (LIWC) Regressions

Figure 1. Responses to Likert item questions on 11,255 CS course reviews

Figure 2. Survey response rates by course grade

Figure 3. Claims of high achievement by grade and gender

Figure 4. (Top): Topic Proportions

(Bottom): Differences in Topic Prevalence by Gender