

Grade Information and Grade Inflation: The Cornell Experiment

Talia Bar, Vrinda Kadiyali, and Asaf Zussman

Grade inflation and high grade levels have been subjects of concern and public debate in recent decades as student grades have increased considerably in many institutions of higher education across the United States, notably in the Ivy League (Rosovsky and Hartly, 2002). For example, the *Boston Globe* reported that in 1950, about 15 percent of Harvard students received a grade of B+ or better, whereas in 2007, more than half of all Harvard grades were in the A range (Primack, 2008). Since grades are bounded from above (in most institutions by an A or an A+), grade inflation is accompanied by a compression of grades at the top. Grades are intended to communicate valuable information to, among others, students (identifying their strengths), financial aid officers (allocating funding), graduate schools (making admission decisions), and potential employers (screening job applicants). Grade inflation and grade compression depreciate the information content of grades and undermine their usefulness.

In the mid-1990s, Cornell University's Faculty Senate had a number of discussions about grade inflation and what might be done about it. In April 1996, the Faculty Senate voted to adopt a new grade reporting policy which had two parts: 1) the publication of course median grades on the Internet; and 2) the reporting of course median grades in students' transcripts. The policy change followed the determination of a university committee that "it is desirable for Cornell University to provide more information to the reader of a transcript and produce more meaningful letter grades." Curbing grade inflation was not explicitly stated as a goal

■ *Talia Bar is Assistant Professor of Economics and Vrinda Kadiyali is Professor of Economics and Marketing in the Johnson School, both at Cornell University, Ithaca, New York. Asaf Zussman is Lecturer, Department of Economics, Hebrew University, Jerusalem, Israel. Their e-mail addresses are <tb97@cornell.edu>, <kadiyali@cornell.edu>, and <azussman@mssc.huji.ac.il>, respectively.*

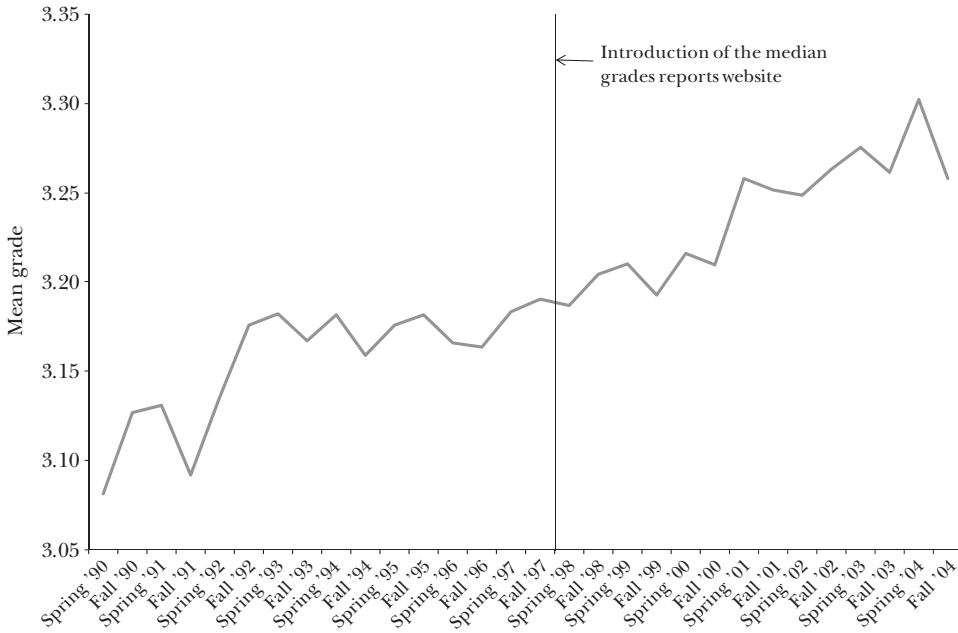
of this policy. Instead, the stated rationale was that “students will get a more accurate idea of their performance, and they will be assured that users of the transcript will also have this knowledge. A grade of B– in a course of substantial enrollment in which the median was C+ will often indicate a stronger performance than, e.g., a B+ in a large course in which the median is A. *More accurate recognition of performance may encourage students to take courses in which the median grade is relatively low* [emphasis ours].” The first “median grade report” was available online by the spring semester of 1998, and the reports can be found under “Useful Links” in the website of Cornell’s Registrar at (<http://registrar.sas.cornell.edu>). However, the reporting of course median grades in the transcripts has not been implemented yet; grades are expected to enter students’ transcripts only for the class of 2012 and beyond. Prior to the implementation of the policy, students could learn about grading practices mostly through the grapevine. Since 1998, whatever grade information was available through this informal channel was augmented by easily accessible official information.

We evaluate the effect of this policy change on patterns of course choice and grade inflation. Specifically, we test two related hypotheses (see Bar, Kadiyali, and Zussman, 2008, for the formal model that underlies these hypotheses). First, all else being equal, the availability of online grade information will lead to increased enrollment into leniently graded courses. Second, high-ability students will be less attracted to the leniently graded courses than their peers. Building on these results we perform an exercise which identifies the extent to which the change in student behavior resulted in an increase in the university-wide mean grade. We call this “compositional grade inflation,” to differentiate it from a situation in which instructors raise the grades they assign, a phenomenon that we call “classic grade inflation.”¹

Figure 1 shows the mean grade in undergraduate-level courses in Cornell’s College of Arts and Sciences from 1990 to 2004. During this period the mean grade climbed from around 3.10 (equivalent to a little higher than a B) to around 3.30 (equivalent to a B+), an increase of more than 6 percent (this rate is similar to that reported for private universities in the website (<http://gradeinflation.com>)). The early part of the period, between the spring semester of 1990 and the fall semester of 1992, witnessed a relatively significant increase in the mean grade. The mean grade then remained stable for several years. From 1998 until 2004, the period in which the policy of online reporting of median grades was in effect, the mean grade has steadily increased.

¹ The empirical literature on course choice and grade inflation is rather thin. Sabot and Wakeman-Linn (1991) and Fournier and Sass (2000) demonstrate that the likelihood that a student will enroll in advanced courses in a given department or major in it depends positively on the student’s grade in an introductory course in that department. Johnson (2003) analyzes data from an experiment conducted at Duke University during the academic year 1998–99 in which several hundred students were provided with historical course information and were then asked to choose among courses. Johnson shows that students tended to elect courses offered by leniently grading instructors.

Figure 1

Mean Grade at Cornell's College of Arts and Sciences, 1990–2004

Notes: The figure displays mean grade in undergraduate-level courses with at least one pre-policy change observation and at least one post-policy change observation.

A Preliminary Look at the Data

To study the effect of grade information on course selection and grade inflation, we utilize a large dataset of grades assigned in undergraduate level courses at the College of Arts and Sciences between the spring semester of 1990 and the fall semester of 2004. We restrict our analysis to courses that were taught at least once pre-policy change (1990–1997) and at least once post-policy change (1998–2004). Each observation in the dataset has information on an individual student taking a specific course and the final grade of that individual in the course.

Median grades have been reported online since the spring semester of 1998. There is a separate report for each semester. For each course, the report lists its title, enrollment, and the median grade. Graduate courses and courses with enrollment of less than 10 students were typically excluded from the reports. Grade posting is by course and not by instructor, and is aggregated up across various sections of the course, if they exist. This procedure prevents us from including controls for instructors and sections. By inspecting electronic file property characteristics, we were able to estimate that the lag in the online publication of grades was typically one or two semesters. However, from the fall semester of 1999 to the spring semester of 2002, the publication lag was at least three semesters.

Table 1
Descriptive Statistics

Panel A: Total Course Information					
<i>Term</i>	<i>Student–course observations</i>	<i># of courses</i>	<i>% of student grades</i>		
			<i>A</i>	<i>B</i>	<i>F–C</i>
Spring '90	23,028	567	35.5	44.7	19.7
Fall '90	25,485	564	37.8	43.7	18.5
Spring '91	22,357	589	37.7	44.3	18.0
Fall '91	25,005	575	35.8	44.7	19.5
Spring '92	21,255	623	38.2	43.9	17.9
Fall '92	23,957	610	41.4	41.9	16.7
Spring '93	21,614	667	40.8	42.8	16.5
Fall '93	24,572	648	40.5	42.5	17.1
Spring '94	21,540	671	40.2	43.9	15.9
Fall '94	24,470	679	40.4	42.4	17.2
Spring '95	21,642	698	40.9	42.4	16.8
Fall '95	25,286	701	41.7	41.7	16.6
Spring '96	21,335	700	41.2	41.5	17.3
Fall '96	24,870	708	41.2	41.2	17.6
Spring '97	22,148	739	42.3	41.2	16.5
Fall '97	25,140	776	42.0	41.7	16.3
Spring '98	22,258	773	42.4	41.2	16.3
Fall '98	25,219	739	43.2	40.7	16.0
Spring '99	22,406	744	43.2	41.1	15.8
Fall '99	24,923	710	42.3	41.1	16.5
Spring 2000	21,162	717	44.0	40.3	15.7
Fall 2000	23,544	699	43.7	40.1	16.2
Spring '01	21,113	695	46.5	38.9	14.6
Fall '01	23,206	652	46.2	38.9	14.9
Spring '02	20,530	641	45.8	39.2	15.0
Fall '02	23,502	663	46.0	40.2	13.8
Spring '03	20,719	650	48.0	38.0	14.1
Fall '03	22,673	626	46.4	39.6	13.9
Spring '04	20,052	635	49.2	38.2	12.7
Fall '04	21,688	604	46.5	39.0	14.5
1990–1997	373,704	10,515	39.8	42.8	17.4
1998–2004	312,995	9,548	45.2	39.8	15.0
Total	686,699	20,063	42.3	41.4	16.3

Table 1 provides an overview of the data. Panel A reports the total number of student–course observations, the total number of courses, and the share of student–course observations that were assigned a letter grade in a given range. Panel B reports the share of courses with a median grade in a given range (first three columns) and the share of student–course observations in courses with a median grade in a given range (last three columns).

Several facts are worth highlighting. The total number of student–course observations and the number of courses do not exhibit a clear upward or downward trend

Table 1—continued

Panel B: Median Grade Information						
Term	% of courses with median grades in the range of			% of student enrollment by course median grades		
	A	B	F–C	A	B	F–C
Spring '90	47.4	52.0	0.5	19.6	78.8	1.7
Fall '90	49.5	49.6	0.9	23.2	76.7	0.2
Spring '91	50.8	48.4	0.8	20.5	78.1	1.3
Fall '91	46.4	53.0	0.5	19.5	79.4	1.1
Spring '92	51.7	47.2	1.1	21.8	78.1	0.1
Fall '92	53.1	46.1	0.8	22.3	77.5	0.3
Spring '93	51.9	47.4	0.7	23.1	76.8	0.1
Fall '93	49.8	49.7	0.5	24.2	74.4	1.4
Spring '94	55.3	44.1	0.6	23.1	76.7	0.2
Fall '94	51.7	48.0	0.3	24.8	75.2	0.1
Spring '95	52.0	46.1	1.9	25.0	71.9	3.1
Fall '95	51.5	47.6	0.9	25.5	74.4	0.1
Spring '96	52.9	46.3	0.9	26.3	73.6	0.1
Fall '96	52.3	47.2	0.6	24.6	74.6	0.8
Spring '97	54.5	43.7	1.8	28.3	69.8	1.9
Fall '97	53.7	45.1	1.2	24.9	74.9	0.2
Spring '98	55.5	43.7	0.8	28.8	71.1	0.1
Fall '98	54.3	44.5	1.2	26.7	73.1	0.2
Spring '99	56.6	42.1	1.3	30.2	69.3	0.5
Fall '99	54.5	44.2	1.3	28.3	69.7	2.0
Spring 2000	58.3	40.4	1.3	30.4	69.3	0.3
Fall 2000	60.2	38.5	1.3	30.9	68.9	0.2
Spring '01	63.2	36.3	0.6	39.3	60.5	0.2
Fall '01	61.8	37.7	0.5	31.9	68.1	0.0
Spring '02	62.4	36.7	0.9	36.7	63.1	0.2
Fall '02	62.6	36.7	0.8	33.5	66.4	0.1
Spring '03	65.7	33.5	0.8	39.3	60.7	0.0
Fall '03	60.2	39.0	0.8	35.6	64.4	0.1
Spring '04	66.5	32.8	0.8	42.6	57.3	0.1
Fall '04	63.6	35.9	0.5	37.2	62.8	0.0
1990–1997	51.7	47.4	0.9	23.5	75.7	0.8
1998–2004	60.2	38.9	0.9	33.4	66.3	0.3
Total	55.7	43.4	0.9	28.0	71.4	0.6

Notes: The table displays descriptive statistics for undergraduate-level courses at Cornell's College of Arts and Sciences in the period 1990–2004. As in the analyses below, the sample is restricted to grades assigned in courses that have at least one pre-policy change observation and one post-policy change observation. **Panel A:** column 1 reports the total number of student–course observations; column 2 reports the total number of courses; columns 3–5 report the share of student–course observations that were assigned a letter grade in a given range. **Panel B:** columns 1–3 report the share of courses that in the relevant semester had a median grade in a given range; columns 4–6 report the share of student–course observations in courses that in the relevant semester had a median grade in a given range. The line dividing the “Fall 1997” and “Spring 1998” rows separates the pre- and post-implementation periods for the policy of posting median course grades.

over the years. In contrast, the share of students receiving grades in the A range has strongly increased. There was also an increase in the share of courses with a median grade in the A range as well as an increase in the share of students enrolled in such courses. These observations are consistent with the conventional view on grade inflation—instructors are applying more lenient grading practices. However, while the rise in the share of courses with a median in the A range seems to be relatively stable over the 1990–2004 period, the rise in the share of students enrolled in courses with a median grade in the A range seems to have significantly accelerated in the 1998–2004 period. Comparing the pre- and post-policy change periods (the second and third rows from the bottom in Panel B), the share of courses with a median in the A range increased by about 16 percent, while the share of students enrolled in such courses increased by more than 42 percent. This suggests the possibility of compositional grade inflation. We return to this issue below.

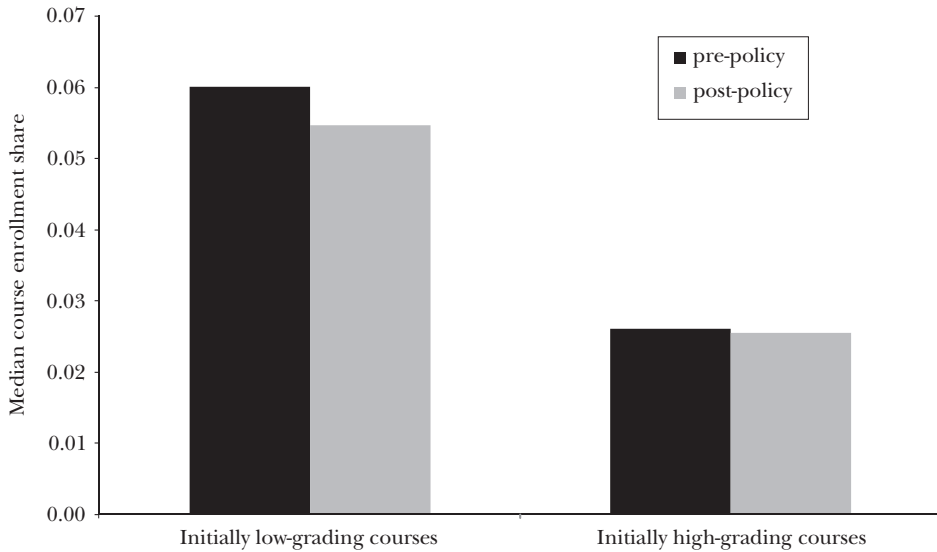
Grade Information and Course Selection

Our prediction that in the post-policy change period students will be drawn to leniently graded courses relies on the assumption that students had access to the online grade information and used it to select courses. We found support for this assumption in a survey of around 500 students from Cornell's Department of Economics conducted during the spring semester of 2006. Close to 60 percent of the students indicated that they heard about the website and visited it, and roughly 80 percent indicated that they sometimes use the information provided in the website to select courses or have heard about other students doing so. Additionally, we obtained two and half years worth of daily data on the number of visits (hits) to the website and examined whether it was visited more frequently during periods of course registration, when students are more likely to need grade information for course selection. There are two periods of course registration in every semester at Cornell: in the first ("add/drop" period) students enroll in courses offered in the current semester; in the second ("pre-enrollment" period) students pre-enroll in courses offered in the next semester. We found large and highly statistically significant differences in the number of hits between the periods: the daily number of visits was roughly twice as large in registration periods than in nonregistration periods. The rise in hits during registration periods indicates that students are aware of these posted median grades and are presumably using them to guide their enrollment choices.

In the empirical test of our first hypothesis, we attempt to exploit persistent differences in course medians that are in place prior to the policy change. Thus the "dosage" variable is the pre-existing median grade and the treatment is the interaction between this dosage and the new grade disclosure policy. Figure 2 illustrates our approach. We define two groups of courses by the value of the mean of median grades in each course in the pre-policy change period (this variable times a dummy variable that takes the value of 1 when the policy is in effect is the main explanatory variable in the regression reported in Table 2). The cutoff between "initially high-" and "initially

Figure 2

Median Course Enrollment Share Pre- and Post-Policy Change in Initially Low-versus Initially High-Grading Courses



Notes: The figure displays the median course enrollment share, pre- and post-policy change in two groups of courses that are defined by the value of the mean of median grades in each course in the pre-policy change period. The cutoff between “initially high-” and “initially low-” grading courses is 3.5, and this splits the set of courses into two equal-sized groups.

low-” grading courses is 3.5, and this splits the set of courses into two equal-sized groups. Figure 2 supports our basic conjecture in that the median enrollment share dropped in initially low-grading courses while barely changing in the initially high-grading courses. (The fact that the overall result presented in Figure 2 is not as striking as we would have hoped may be related to classic grade inflation. The data shows that grades were rising in the initially low-grading courses but were stagnant in the initially high-grading courses, thus some of the initially low-grading courses became high-grading courses in the second period, attracting more students, while initially high-grading courses would not become even more attractive over time.)²

In Table 2, we provide the results of a formal test of our first hypothesis (that the availability of online grade information will lead to increased enrollment into leniently graded courses). In column 1, the sample includes all courses. In column 2, we restrict the sample to advanced (400-level) courses. We do so because such courses are in principle electives and because they are typically not attended by freshman students. Freshman students are less knowledgeable about the median

² Note: If median enrollment goes down for one set of courses, it does not need to go up for another. For example, an increase in the number of courses offered would push down proportional enrollment in all courses.

Table 2

Effect of Grade Information on Course Enrollment Shares

	<i>Courses</i>			
	<i>All</i>	<i>Advanced</i>	<i>Annual</i>	<i>Advanced, Annual</i>
	(1)	(2)	(3)	(4)
Pre-Policy-Period Average				
Median Grade * Policy Dummy	0.014* (0.010)	0.014*** (0.003)	0.029* (0.019)	0.020*** (0.007)
Semester fixed effects	Yes	Yes	Yes	Yes
Course fixed effects	Yes	Yes	Yes	Yes
R^2	0.90	0.73	0.95	0.72
Observations	20,063	5,519	8,637	2,255

Notes: The dependent variable is the share (in percent) of students enrolled in a given course in a given semester out of total enrollment (student-course observations) in undergraduate level courses at Cornell's College of Arts and Sciences. The main explanatory variable is the interaction between the average median grade during the pre-policy change period (1990–1997) and a dummy variable for the post-policy change period (1998–2004). The sample includes all courses that have at least one pre-policy change observation and one post-policy change observation. Advanced (400-level) courses are upper-division courses. Annual courses are those in which two consecutive course observations are two semesters apart. All models include fixed effects for each course and for each semester. The regressions were estimated by ordinary least squares. Standard errors, clustered by course, are reported in brackets.

*, **, *** represent statistical significance at the 10, 5, and 1 percent levels in a one-tailed t-test.

grades website (as the results of our survey strongly indicate), and their course choice might be less a function of grades and more about learning one's tastes and abilities and fulfilling requirements for future course selection. In column 3, we restrict the sample to annual courses (those in which two consecutive course observations are two semesters apart). We make this restriction because, at least based on evidence from the Department of Economics, the likelihood that the same instructor will persist in such courses is higher than in other courses. In column 4, we restrict the sample to advanced annual courses.

The results in column 1 imply that a one unit difference (say, the difference between a B and an A) in the pre-policy change average median grade is associated with a 0.014 percentage point increase in the share of students enrolled in the course in the post-policy change period. Since the median course share was about 0.052 percent in the pre-policy change period, this represents a significant rise in the enrollment share. The results for advanced courses are much stronger, both in the statistical sense and in the economic sense: Since the median enrollment share in such courses was only about 0.026 percent in the pre-policy change period, the coefficient of 0.014 in column 2 implies a much stronger reaction. The sensitivity of enrollment shares to grades is also relatively strong for annual courses: compared to column 1, the regression coefficient in column 3 is twice as large (0.029 versus 0.014), while the pre-policy change median enrollment share is roughly the same (0.054 versus 0.052). Unsurprisingly, results are strongest in column 4: a coefficient

of 0.020 relative to a pre-policy change median enrollment share of 0.029 percent. Thus for such courses, a one unit difference in the “pre-existing” grade is associated with almost a 70 percent increase in the course’s enrollment share.

In sum, the evidence suggests that the provision of grade information led to increased enrollment into leniently graded courses and that this effect was strongest in advanced annual courses.

Student Ability and Responsiveness to Grades

Our second hypothesis is that, in the presence of grade information, the tendency to select leniently graded courses is weaker for high-ability students than for their peers. This hypothesis depends on the shape of the grading policies employed in different courses. In Bar, Kadiyali, and Zussman (2008), we provide theory and empirical evidence on this issue. We demonstrate that the best students can expect to obtain the highest grade (A+) regardless of the leniency of the grading policy (proxied by the median grade) while the weaker students can expect a much higher grade in leniently graded courses than in strictly graded ones. Thus high-ability students have less to gain than their peers by taking the leniently graded courses.

To test the second hypothesis, we rely on SAT scores as a proxy for student ability. Using SAT scores makes sense for a number of reasons: the SAT scores are based on a standardized test that almost all students take before admission to Cornell, and there is a strong positive correlation between SAT scores and grades. Of course, we cannot use Cornell grades as a measure of ability because they may be affected by the extent to which students choose leniently graded courses. In 1995, SAT scores were “re-centered” at a score of 500 by the Educational Testing Service.³ This recalibration implies that SAT scores of students who took the test before 1995 are not directly comparable to those of students who took the test later. Unfortunately, our data do not indicate whether a student’s SAT score is based on a test taken before or after the re-centering. Additionally, we are missing most data from the earlier years. Thus, in this section we focus our analysis on the period from the fall semester of 1999 to the fall semester of 2004, a period in which the vast majority of the SAT observations are likely to be re-centered.

We first identify high-ability students as those in the top percentiles in the SAT distribution of each semester (using the universe of all Cornell students in that semester for whom we have data), where the percentile variable p alternatively takes the values 40, 30, 20, and 10. For each course-semester observation, we compute the share of high-ability students in the total number of students in the course. This forms our dependent variable. The explanatory variables include the course’s

³ Re-centered grades are not a simple monotone transformation of the old scale. The re-centered math scores were raised relative to the original ones by up to 50 points, but could also drop by up to 10 points; the re-centered verbal scores were raised relative to the original ones by up to 80 points (for details see the College Board website (<http://www.collegeboard.com>)).

Table 3
Are High-Ability Students Less Sensitive to Grades?

	<i>Top 40%</i>	<i>Top 30%</i>	<i>Top 20%</i>	<i>Top 10%</i>
Lagged Median Grade	-2.531** (1.454)	-2.396** (1.314)	-3.192*** (1.253)	-1.382* (1.021)
Semester fixed effects	Yes	Yes	Yes	Yes
Course fixed effects	Yes	Yes	Yes	Yes
R^2	0.60	0.59	0.57	0.53
Observations	6,328	6,328	6,328	6,328

Notes: The dependent variable is the share (in percent) of high-ability students in the total number of students in a given course and semester. High-ability students are identified as those in the top 40, 30, 20, and 10 percentiles in the SAT distribution of Cornell Arts and Sciences students in the corresponding semester. The main explanatory variable is the last median grade reported online for the course. The sample includes courses that have at least one pre- and one post-policy change observation. Due to SAT data limitations, the period covered is from the fall semester of 1999 to the fall semester of 2004. All models include fixed effects for each course and for each semester. The regressions were estimated by ordinary least squares. Standard errors, clustered by course, are reported in brackets. The results presented in Table 3 are robust to the addition of a lead or a lag of the main explanatory variable. The additional variables were for the most part insignificant. This suggests that our results are not driven by mean reversion.

*, **, *** represent statistical significance at the 10, 5, and 1 percent levels in a one-tailed t-test.

lagged median grade, semester fixed-effects, and course fixed-effects. Our prediction is that the share of high-ability students in a course would depend negatively on the lagged median grade.

Table 3 reports the results of the analysis. We find that the relationship between the share of high-ability students and the lagged median grade is negative and statistically significant for all four SAT threshold levels. The interpretation of the coefficients is as follows (taking the “Top 20%” as an example): a one-unit increase in the lagged median grade leads to a 3.19 percentage points decrease in the share of high-ability students in the course. Since the median share of high-ability students is 16.67 percent, this represents a nonnegligible effect. The strength of the effect on student ability is similar in the column for the “Top 10%” share and somewhat weaker for the other two shares. The evidence is thus consistent with our prediction that high-ability students will be less attracted to leniently graded courses than their peers.

How Much Does Compositional Grade Inflation Matter?

To estimate the relative contributions of the compositional and classic components of grade inflation in the post-policy change period, we use a simulation based on the share of courses with a median grade in the A range and the share of student enrollment in courses with a median in this grade range (these data are provided in columns 1 and 4, respectively, of Table 1, Panel B). The behavior of the share of courses with a median grade in the A range approximates the effect of classic grade inflation. Comparing the share of student enrollment in courses

having a median in the A range (student enrollment in the A range) to the share of courses having a median in the A range (share of courses in the A range) reveals the role of compositional grade inflation—that is, students shifting toward more leniently graded courses. In the absence of grade-driven course selection, the two series are expected to have a constant relationship; conversely, grade-driven course selection will manifest itself in a breakdown of this relationship with the first series (enrollment share) rising faster than the second (course share).⁴

Figure 3 suggests that grade-driven course selection is likely to have had a large influence. The figure plots the two series mentioned above: the share of courses with a median grade in the A range and the share of student enrollment in courses with a median grade in this range. Throughout the period, the course share in the A range is much higher than the enrollment share in the A-range courses because lower-level courses typically have larger student enrollment and stricter grading policies than upper-level courses. The figure also plots three additional series. One is a counterfactual share of courses with medians in the A range for the period 1998–2004; it is an extrapolation based on the growth rate of the A-range course share in the pre-policy change period. The deviation between the actual and the counterfactual A-range course share series reflects acceleration in the classic component of grade inflation. The two remaining series display the predicted shares of student enrollment in the A-range courses based on the relationship between the enrollment-share and the course-share series in the 1990–1997 period: one uses the actual course share, and the other, the extrapolated counterfactual course share.⁵ The fact that the actual enrollment share is higher than the counterfactual enrollment share in the post-policy change period is consistent with compositional grade inflation. It suggests that students' enrollment in leniently graded courses exceeded what we would have expected based on just the share of courses in the A range.

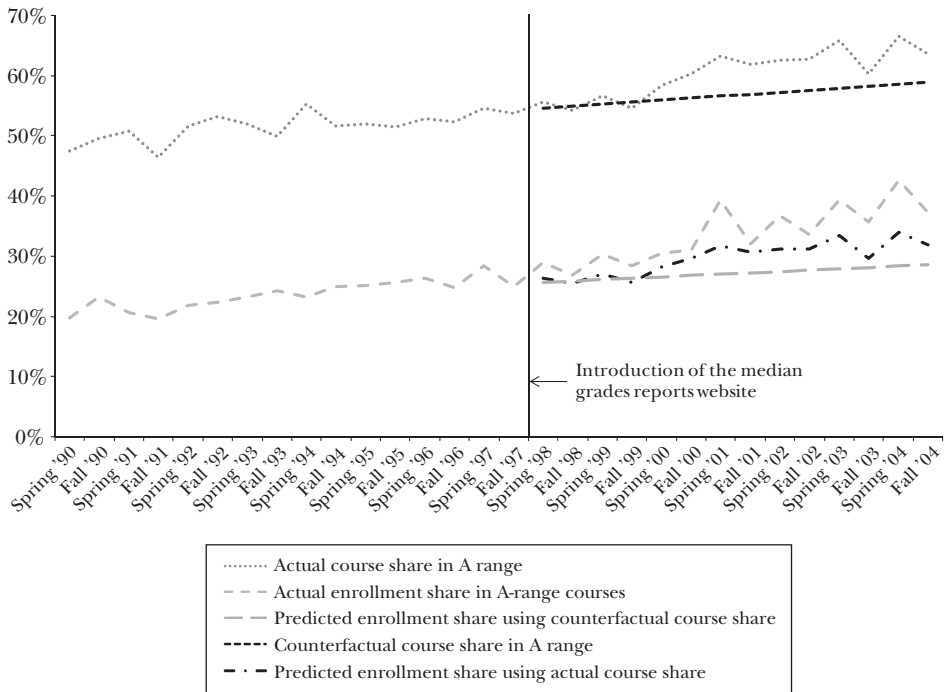
To estimate how much of the increase in grade inflation is due to acceleration of classic grade inflation and how much should be attributed to compositional effects, we do the following. First, we approximate the college-wide mean grade based on two factors: 1) the shares of student enrollment in courses with a median in the A, B, and C or below ranges; and 2) the mean grades in courses in these three

⁴ If larger classes are the ones that raise their grades faster in the post-policy change period, this could seem like compositional grade inflation. To examine this issue, we regressed the change in median grade between two consecutive periods on lagged enrollment (in logs) pre- and post-policy change using the same sample as in our baseline regressions. We found no relationship between the change in the median grade and lagged enrollment in either subperiod. Additionally, in both subperiods, we found no significant differences between the average sizes of courses whose median grade increased from the B range to the A range.

⁵ Let c_t and e_t denote, respectively, the actual share of courses with medians in the A range and enrollment share in such courses at time t . We estimated the equation $c_t = \alpha_1 + \beta_1 t + \varepsilon_t$ for 1990–97 for t equals 1 to 16. The counterfactual course share \hat{c} for 1998–2004 for t equals 17 to 30 is given by $\hat{c}_t = \hat{\alpha}_1 + \hat{\beta}_1 t$. Next, we estimated the equation $e_t = \alpha_2 + \beta_2 c_t + \varepsilon_t$ for 1990–97. The 1998–2004 predicted enrollment share series using the actual and counterfactual course shares are given by, respectively, $\hat{\alpha}_2 + \hat{\beta}_2 c_t$ and $\hat{\alpha}_2 + \hat{\beta}_2 \hat{c}_t$.

Figure 3

Estimating the Effect of Compositional Grade Inflation: Course Shares with Medians in the A Range and Enrollment Shares in Courses with Medians in the A Range



Notes: The figure displays shares of courses with medians in the A range, and enrollment shares in courses with medians in the A range for our baseline set of courses. The counterfactual course share is an extrapolation based on the growth rate of the course share in the pre-policy change period. The figure displays two sets of predicted enrollment shares for the post-policy change period. One is based on the linear relationship between the course share and the enrollment share in the pre-policy change period and on the actual course share in the post-policy change period. The other is based instead on the counterfactual course share. The text explains how to use the three enrollment share series to estimate the contribution of “compositional grade inflation” to overall grade inflation in the post-policy change period.

categories.⁶ The approximated series mimics well the behavior of the actual college-wide mean grade series (shown in Figure 1): the average rate of grade inflation per semester in the post-policy change period was 0.15 percent according to the actual series and 0.16 percent according to the approximated series. We then construct in a similar manner two counterfactual college-wide mean grade series based on the

⁶ To construct the approximated college-wide mean grade series, we rely on the fact that the share of students enrolled in courses with a median of C+ and below is miniscule (typically less than 1 percent—see column 6 in Table 1, Panel B). In our approximation, we assume that this share is constant, which implies that the share of students enrolled in B-range courses varies one-to-one with the share of students enrolled in the A-range courses. The mean grade for each of the three median grade categories is computed from the individual-level dataset using our baseline set of courses.

two predicted enrollment share series (shown in Figure 3). The average rate of grade inflation per semester in the post-policy change period is 0.05 percent according to the series that relies on the counterfactual course share and 0.09 percent according to the series that relies on the actual course share. The difference between the last two reported figures again reflects acceleration in the rate of classic grade inflation in the post-policy change period.

What then were the relative contributions of these components of grade inflation in the post-policy change period? Our results imply that about 29 percent ($0.05/0.16$) of the rise in grades was due to a continuation of the classic component of grade inflation at the same rate as in the pre-policy change period. About 26 percent ($0.04/0.16$) was due to the post-policy change acceleration in the rate of classic grade inflation. The remaining 45 percent was thus due to compositional grade inflation, which in turn was probably related to the availability of online grade information that allowed students to shop for courses. We note that there is some anecdotal evidence that during the period under investigation there was a more general trend in U.S. colleges and universities of increasing “grade grubbing” by students. Thus it is possible that some of the compositional grade inflation was due to an interaction between this general trend and the policy change at Cornell.

What Drives Classic Grade Inflation?

Our analysis demonstrates that classic grade inflation existed before the policy change and accelerated after the new policy was implemented. In background work, we empirically examined potential determinants of these phenomena and here report our main findings.

One potential explanation for classic grade inflation that we examined draws on the observation that, in many colleges and universities, the share of non-tenure-track and untenured faculty increases over time and on the belief that instructors from this group tend to grade more leniently than others (Rosovsky and Hartley, 2002). The classic grade inflation at Cornell, however, does not seem to result from this cause. The share of non-tenure-track and untenured faculty in Cornell’s College of Arts and Sciences did not display much variation over the period under investigation. Moreover, analysis of data on the Department of Economics (for which we have instructor information) demonstrated that the tenure status and academic rank of faculty were not statistically significant determinants of course grades (though the lack of significance may simply reflect small sample size).

A second explanation for classic grade inflation is that it resulted from an improvement in the quality of students. This explanation implicitly rests on the assumption that instructors compare students to previous cohorts and raise students’ grades to match their higher abilities. Student abilities could improve for various reasons. For example, an increase in the demand for college education in elite universities which is not matched by an increase in admissions would result in such an improvement. Our analysis found some support for this explanation.

University records (at the aggregate level) show that SAT scores of freshmen students entering the College of Arts and Sciences tended to increase throughout the period under investigation, and especially post-policy change.⁷ We also found that SAT scores have a positive and statistically significant effect on students' course grades. Hence, the rise in the pace of classic grade inflation in the post-policy change period may at least in part be due to a faster rate of improvement in student quality.

A third possibility is that the acceleration in the pace of classic grade inflation reflected faculty response to the policy change. However, from our conversations with many colleagues at Cornell, most professors seemed unaware of the median grade reports website during the period of our study (until we started working on this project). Moreover, given the heterogeneous preferences of professors for grading leniency as well as for course size, even if professors were informed, it is not clear how they would respond to information on median grades. Some professors might be embarrassed to be revealed to be lenient graders, but others might be proud. Professors might revert to the mean grade in their department or variously defined peer groups, or they might want to be different from their peer groups (for example, to be seen as among the toughest graders in their department).

One possible way in which faculty may respond to the policy change is to try to affect student enrollment by changing their course's median grade. In fact, even if professors are uninformed of the policy of publicizing grades, they might respond to observed changes in enrollment caused by the new policy. We examined the relationship between current median grade, this time as a dependent variable, and enrollment. We alternatively used either lagged enrollment or current enrollment (in logs).⁸ In both the pre- and the post-policy change periods, we found no association between lagged or current enrollment on the one hand and the current median grade on the one hand.

In sum, based on our findings, we cannot attribute the acceleration in the rate of classic grade inflation in 1998–2004 to either changes in the composition of the faculty or to faculty response to the policy change. Rather, our analysis suggests that this acceleration may have been driven to some extent by improvement in student quality, as well as by other trends operating within the university. At the same time we note that we have not exhausted the potential reasons for grade inflation. It is possible that other factors, Cornell-specific or more general, have contributed to this phenomenon.

⁷ While we interpret the increase in SAT scores as an indication of improvement in students' academic abilities, it is possible that students are not more able but simply better prepared for the test itself. Of course, a background that includes an ability to prepare better for an important test could also be construed as an improvement in academic ability.

⁸ Note that since the median is set at the end of the course, we can test how it is related to current enrollment without raising the question of endogeneity.

Concluding Remarks

In the spring semester of 1998, Cornell University started publishing median course grades on the Internet. Our analysis finds that the provision of grade information online induced students to select leniently graded courses—or in other words, to opt out of courses they would have selected absent considerations of grades. We also find that the tendency to select leniently graded courses was weaker for high-ability students. Finally, our analysis demonstrates that a significant share of the acceleration in grade inflation since the policy was adopted can be attributed to this change in students' course choice behavior. While the focus of our study is on a particular university, these results may have broader applicability in today's world where students have access to online peer-generated reviews of instructors and courses (like at RateMyProfessor.com). It is possible, even likely, that the effects we find of student selection into leniently graded courses exist beyond Cornell.

What are the welfare implications of our study? The introduction of information on grades increases the utility of some students but decreases the utility of others. For example, a student who chooses the strictly graded course under both information regimes is made worse off by the change because the change drives lower-ability students out of the strictly graded course and makes this student look worse compared to the remaining students. Conversely, a student who chooses the leniently graded course regardless of whether grade information is available is made better off, because the leniently graded class will attract more low-ability students, and thus the student's grade relative to the median will look better. To evaluate the overall welfare implications of the policy change rigorously we would need more information—for example, on the ways in which student learning is affected by more actively choosing leniently graded classes and on the goals of readers of students' transcripts such as future employers. For now we only note that though grade inflation may have the undesirable effect of reducing the information content of grades, it is possible that an increase in grade information could have a negative effect on social welfare.

The median grade policy has remained to date only partially implemented: median grades have been reported online since 1998 but do not yet appear in transcripts. The difference between having median grade information available only online and having it available on transcripts is potentially quite important. The current system allows students to select into leniently graded classes but does not make the information easily accessible to readers of transcripts, which may explain why the policy did not achieve one of its explicit objectives; enrollment into leniently graded courses increased, rather than decreased. It is possible that the inclusion of median grades in students' transcripts in the future will mitigate or even reverse this effect. However, in Bar, Kadiyali, and Zussman (2009), we show that theoretically, even if information is provided to students and readers of the transcript alike, enrollment into leniently graded courses as well as average grades could increase.

In May 2009, Cornell's Educational Policy Committee brought for consideration in a faculty senate meeting a resolution proposing to remove the median grades reports from Cornell's website. The rationale provided for this policy change was based on the committee's reading of the results of the working paper version of this study. The senate voted to table the resolution, so as to allow faculty to obtain more information, and is expected to vote on this in a future meeting. We hope that our analysis will stimulate research on related policy questions: What type of grade information should be provided to students, instructors, and others? Should uniform grading guidelines be imposed on instructors? Can the reporting of median grades on transcripts substitute for uniform grading policies? These questions deserve the attention of the academic community.

■ *We are grateful to numerous individuals in Cornell's administration for providing data and for offering useful advice. We are indebted to David Autor, Haim Bar, Hugo Benitez, John Cawley, Ron Ehrenberg, Jeremy Fox, David Genesove, Sachin Gupta, Jim Hines, Dean Lillard, Francesca Molinari, Daniel Simon, Timothy Taylor, Ed Vytlačil, Mike Waldman, colleagues at Cornell University, and seminar participants at Bar-Ilan University, Brigham Young University, Hebrew University, Iowa State University, Stanford University Graduate School of Business, Tel-Aviv University, and the University of Montreal for helpful conversations and insightful comments.*

References

- Bar, Talia, Vrinda Kadiyali, and Asaf Zussman.** 2008. "Quest for Knowledge and Pursuit of Grades: Information, Course Selection, and Grade Inflation." Available at SSRN: <http://ssrn.com/abstract=1019580>.
- Bar, Talia, Vrinda Kadiyali, and Asaf Zussman.** 2009. "Putting Grades in Context." Available at SSRN: <http://ssrn.com/abstract=1406681>.
- CollegeBoard.com.** 2009. Website. <http://www.collegeboard.com>.
- **Fournier, Gary M., and Tim R. Sass.** 2000. "Take My Course, Please: The Effects of the Principles Experience on Student Curriculum Choice." *Journal of Economic Education*, 31(4): 323–39.
- Gradeinflation.com.** 2009. Website. <http://gradeinflation.com>.
- Johnson, Valen E.** 2003. *Grade Inflation: A Crisis in College Education*. New York: Springer-Verlag.
- Primack, Phil.** 2008. "Doesn't Anybody Get a C Anymore?" *Boston Globe*, October 5. (http://www.boston.com/bostonglobe/magazine/articles/2008/10/05/doesnt_anybody_get_a_c_anymore/).
- Rosovsky, Henry, and Matthew Hartley.** 2002. *Evaluation and the Academy: Are We Doing the Right Thing?* Cambridge, MA: American Academy of Arts and Sciences.
- Sabot, Richard, and John Wakeman-Linn.** 1991. "Grade Inflation and Course Choice." *Journal of Economic Perspectives*, 5(1): 159–70.