Quantifying professors' effect on student grades.

Sebastian Gregorius Winther-Larsen and John M. Aiken

Department of Physics and Center for Computing in Science Education, University of Oslo, N-0316 Oslo, Norway

Marcos D. Caballero

*Department of Physics and Astronomy, Department of Computational Mathematics,
Science, and Engineering, and CREATE for STEM Institute,*

Michigan State University, East Lansing, Michigan 48824 and

Department of Physics and Center for Computing in Science Education, University of Oslo, N-0316 Oslo, Norway

In this work, we gauge the effect different lecturers have on student grades. To do this we extract student grades in introductory physics courses as well as ethnicity, gender, high school GPA and math placement scores from the registrar data of a large Midwestern university. We match the course data to instructors over a period from 2002 to 2017, yielding a data set with 8641 observations. We build three different linear models, taking into account different effects and controls in order to quantify the impact individual lecturers have on student grades. We find statistically significant difference of more than half a grade that is directly attributable to the lecturers student's had in a course. Moreover, significant effects associated with genders and ethnic groups are generally lower than effects associated with lecturers. We see an increased gender grade difference when we sub-categorize students taught by individual lecturers into genders.

I. INTRODUCTION

The professors you have when you are working towards a degree could have a profound impact on student engagement and performance. Tinto’s theory of drop-out [1] suggests that a student’s engagement in their academic studies is one of the primary factors that can mediate their desire to persist in their chosen studies at a chosen university. Because a student’s academic engagement is so important to a student persisting through a degree program, it is imperative that the grades student’s receive in said program be impartial across instructors. If there is an instructor based effect this can have a direct impact on student persistence in a physics degree program.

There has been a few studies that try to quantify and isolate the effect of instructor quality on student performance more directly. For instance, a study from the US Air Force Academy, where students are assigned to lecturers randomly in introductory STEM courses, the assigned professors was shown to affect student’s performance in both current courses and following courses [2]. In a study from a large Canadian university [3] instructors are matched to metrics of teaching quality, in order to predict perceived student performance and subject interest. They found that high-quality instructors would make students less likely to drop out, but effects on student grades were limited. The findings from a similar study from an Italian university, found that instructor quality has a significant effect on student grades [4]. In the past few years, there has been an increase in adjunct teaching staff working part time. Part-time lecturers are found to affect the likelihood of success and subject interest for students when compared to full-time faculty members in a study of colleges in Ohio [5]. Such studies of adjunct faculty are of isolated cases and worthy to examine more broadly, however.

In this study we do not seek to uncover what the distinction is between a high-quality instructor from a low-quality instructor, but we try to show that there is a significant difference in student grades that is attributable to instructors. The questions we seek to answer are the following. Can the grade a student achieves in a course be attributed to the lecturer the student had in that course? What is the magnitude of such an effect? How large is the expected grade difference between a student that had a lecturer with a strong positive grade effect versus a student that had a lecturer with a strong negative grade effect? In order to answer these questions, we employ a data set built from registrar data from a large Midwestern university. The data is fitted to three comparable and easily interpretable models, using ordinary least squares.

II. DATA AND MODELS

For this study, we augment the “pathways database” [6], with data about lecturers. The pathways database is constructed using data about the physics bachelor of science program from the registrars office and contains three tables describing student demographics, time-stamped course data,

TABLE I. Descriptions of variables included in the models. In all the models, grades was the outcome variable, high school GPA and math placement scores were control variables. The base model includes the instructor binary variables as well. In the extended model all variables were added. In the gender product model, gender binary variables were multiplied with the instructor variables to sub-categorize students taught by a specific professor.

Variable name	Explanation
grade	Grade per student per course. Difference from average.
hsgpa	High school GPA per student. Difference from average.
math_placement	Math placement score per student. Difference from average.
semester	The number of semesters after students start before they take the course.
[gender]	binary variable for gender of students, divided into male and female.
[ethnicity]	binary variable for 4 ethnicity categories.
instructor_n	binary variable for instructor n, there are 19 instructors.

and time-stamped major registration data. The database catalogs every course a student takes, and for this study was augmented with a fourth table describing the lecturers teaching the courses in the courses table. The database already contains grades for each student for each course, and for most students also include high school GPA and performance on a mathematics placement test.

Even though data is available since 1993 and until 2017, only data from students taking courses after 2002 have been included in this study. The grades were found to be stationary in time after 2002, while showing a slight upwards trend before this year. Only data for Introductory Physics 1 and Introductory Physics 2, are included in the analysis. These courses have a high number of enrolled students, providing a good foundation for statistical analysis.

Data from the augmented pathways database has been winnowed down to a data set consisting of the handful of variables listed in Table I. Numerical grades for each student in each course are included as the difference from the average grade in the variable `grade`. In a similar manner, high school GPAs and math placement scores are included as a difference from the average in the variables `hsgpa` and `math_placement`, respectively. The `semester` variable is the number of semesters since the student has begun university until the student takes the course. Several binary variables has been constructed from the data about students, like the variables `male` and `female`, which sorts students into genders that were determined by their selection in their college application. Similarly, we have constructed four binary variables representing ethnic groups that students again selected in their application; `white`, `black`, `asian` and `combined`. The last of these, `combined`, includes differ-

ent underrepresented groups like Hispanic, multiracial, Native American, Alaska Native and Hawaiian / Pacific Islander, as well as those that did not report their ethnicity. We acknowledge that there are deep complexities in gender and ethnic identity, but our data were limited by the student's application, which uses the gender binary and the IPEDS definitions for ethnicity. Further, we recognize that constructing a binary variable containing several different ethnic groups is problematic. This group represents less than 5% of the overall data and thus the statistical power of our linear models would be significantly reduced with finer grain variables.

Each instructor for each course for a given year and matched to the student data, are included in the data set as binary variables and numbered, `instructor_n`. Thus, each observation in the data set gives information about a student's performance in a certain course, which lecturer the student had in that course, the student's performance in high school and on the math placement exam, the student's gender and ethnic group, as well as the semester the student took the course.

While we are mainly interested in the effect that lecturers have on student grades, we know from previous quantitative studies that a student's gender and ethnic identity can be significant factors in predicting the likelihood that a student finishes a physics major [6], as well as being variables of high "sensitivity" when included in models to predict grades [7]. Qualitative research studies find a variety of reasons for this significance related to the culture of science and a student's developing identity as a scientist [8]. We also know that GPA and placement test scores to some extent predict student grades [9]. Including these variables would therefore add to the robustness of any results concerning lecturers.

We have constructed three models, all with `grade` as the outcome variable, using ordinary least squares with the `Statsmodels` library in Python [10]. The high school GPA (`hsgpa`) and math placement scores (`math_placement`) are used in all models as control variables, as we assume that performance in high school and on entry exams would correlate the grades student achieve in college. The first model,

$$\text{grade} = \sum_n \delta_n \text{Instructor}_n + \alpha_0 \text{hsgpa} + \alpha_1 \text{math_placement} + u, \quad (1)$$

is the base model with only the control variables, `hsgpa` and `math_placement`, in addition to the instructor variables as input variables. We use a notation where δ are binary variable coefficients, α for other coefficients and u for residuals. We call this model "the base model" from now on.

In the second model, we extend the base model with a few more binary variables, namely the student ethnicity variables and the gender variables,

$$\begin{aligned} \text{grade} = & \sum_n \delta_n \text{Instructor}_n \\ & + \sum_m \delta_m \text{ethnicity}_m \\ & + \delta_0 \text{male} + \delta_1 \text{female} \\ & + \alpha_0 \text{hsgpa} + \alpha_1 \text{math_placement} + u. \end{aligned} \quad (2)$$

TABLE II. Summary statistics of the models. Explanatory power increases slightly, while quality decreases through the models. No sign of heteroscedasticity in the residuals from White's test.

N = 8641		Base	Extended	Gender Product
R^2		0.164	0.184	0.187
R^2_{adj}		0.162	0.182	0.183
P(F-statistic)		<0.001	<0.001	<0.001
White's test	f-statistic	2.582	1.956	1.580
	p-value	<0.001	<0.001	<0.001

While gender and ethnicity effects may change over the period we study, we keep the binary variables associated with genders and ethnicities the same for the entire data set. We will refer to this model as "the extended model".

In the third model we take the product of the student gender binary variables and the instructor binary variables,

$$\begin{aligned} \text{grade} = & \sum_n \delta_{n0} (\text{Instructor}_n \times \text{female}) \\ & + \sum_n \delta_{n1} (\text{Instructor}_n \times \text{male}) \\ & + \sum_m \delta_m \text{ethnicity}_m \\ & + \alpha_0 \text{hsgpa} + \alpha_1 \text{math_placement} + u. \end{aligned} \quad (3)$$

This has the effect of splitting students into genders for each lecturer. We call this "the gender product model".

Because the sum of the our category variables; e.g. `male` and `female`, or `Instructor_n`; has perfect multicollinearity with a presumptive intercept variable, none of the models include an intercept variable.

Dividing students into subcategories of ethnicity, i.e. a multiplication of lecturer binary variables and ethnicity binary variables, is not possible, because of the homogeneous nature of the student corpus. Since White students make up 83 % of the total, this would leave few observations per lecturer and in turn make such a proposed fourth model unable to provide any interesting results.

The introductory physics courses at the university in our study has undergone several transformations since the fall of 2013. This includes a more active learning oriented class, a group-based course for engineering students [11], and a specialised course for Life Science students [12]. Regrettably, very few observations were left for each teaching method after cleaning the data. A proposed fifth model including teaching methods was not possible to build for this reason.

To clean the data, we strip away all observations that are incomplete, i.e. we perform a listwise deletion of the data. Subsequently, all observations involving a single lecturer with less than 50 data points is deleted as well. This yields a data set with a total of 8641 observations.

III. RESULTS

The summary statistics of the three fitted models can be found in Table II, while coefficients for the variables in the

fitted models, is shown in Table III with different indicators for significance level.

From Table II, we see that the input variables explain relatively little of the total variance in the grades of students, with an adjusted R^2 of 0.162, 0.184 and 0.184 for the base model, the extended model and the gender product model, respectively. When testing for joint significance of the explanatory variables, we get a very low p-value for all the models. This means that each model explain more than an intercept-only model would. In other words, even though the explanatory power of all models is moderate, the explanatory power is statistically significant when looking at each model as a whole.

We see that the explanatory power of the models increase slightly from the base model as variables are added in order to construct the extended model and the gender product model. Reasonably, we see that the quality of the models fall slightly as variables are added. However, these effects are very limited and the integrity of the three models are similar.

To make sure that we have trustworthy models we check that the residuals follows the same distribution, i.e. that there is no sign of heteroscedasticity. White's test [13] suggest no heteroscedasticity in the residuals. Even so, we got somewhat low p-values when conducting Lilliefors' test [14] which could make it necessary to reject the null hypothesis that the residuals are normally distributed. Indeed, all of Kim and White's skewness measures [15] are slightly negative, i.e. $(SK_1, SK_2, SK_3, SK_4) = -(1.045, 0.058, 0.150, 0.114)$. This means that all models explain lower grades relatively poorly, and can in turn make direct statistical inference problematic.

IV. DISCUSSION

Generally, the three models do not explain much of the grades. In order to be able to predict grades better, we would have to include other in-class variables that has been shown to be better predictors of grades, like performance on homework [9]. That said, in this study we mainly seek to ascertain if lecturer-grade effects exist on a higher level, not to explain or predict student grades entirely. Therefore, we are nevertheless confident in the models, bearing their purpose in mind.

The non-binary control variables, high school GPA and math placement scores (i.e. `hsgpa` and `math_placement` respectively), are highly significant in all three models. The coefficients for these variables are also very close in all three models. Looking at the fitted coefficients in Table III, we can see that starting a physics undergraduate degree with a full grade higher than average GPA out of high school should result in half a grade better grade on average in introductory physics courses. Likewise, scoring a point better than average on the math placement grade would on average yield a student a full grade better than average.

In the extended model and gender product mode, we have another highly significant non-binary variable, `semester`, which is the number of semesters since starting college the

TABLE III. Fitted coefficients of the three models; base model, extended binary variable model and gender product model. a: $p < 0.001$, b: $p < 0.01$, c: $p < 0.05$, d: $p < 0.1$

	Base	Extended	Gender Prod.	
<code>hsgpa</code>	0.474 ^a	0.492 ^a	0.492 ^a	
<code>math_placement</code>	1.045 ^a	0.983 ^a	0.983 ^a	
<code>semester</code>		-0.138 ^a	-0.138 ^a	
<code>male</code>		0.104 ^c		
<code>female</code>		-0.075 ^d		
<code>white</code>		0.064 ^b	0.109 ^d	
<code>black</code>		-0.122 ^a	-0.075	
<code>asian</code>		0.001	0.048	
<code>combined</code>		0.086 ^b	0.132 ^d	
			m	f
<code>instructor_18</code>	-0.361 ^a	-0.347 ^a	-0.254 ^a	-0.624 ^a
<code>instructor_02</code>	-0.202 ^b	-0.156 ^d	-0.028	-0.823 ^a
<code>instructor_07</code>	-0.257 ^b	-0.115	-0.063	-0.183
<code>instructor_08</code>	-0.202 ^a	-0.114 ^d	-0.063	-0.207 ^b
<code>instructor_17</code>	-0.119 ^a	-0.076	0.020	-0.318 ^a
<code>instructor_01</code>	-0.176 ^a	-0.033	0.048	-1.223 ^c
<code>instructor_06</code>	-0.056	-0.021	0.033	-0.130
<code>instructor_09</code>	-0.101 ^c	0.052	0.085	-0.001
<code>instructor_05</code>	-0.087 ^c	0.057	0.139 ^d	-0.137
<code>instructor_16</code>	-0.071 ^c	0.072	0.128 ^d	-0.036
<code>instructor_13</code>	-0.034	0.076	0.108	0.099
<code>instructor_10</code>	0.072	0.100	0.156 ^d	-0.016
<code>instructor_14</code>	0.088 ^d	0.114	0.143	0.088
<code>instructor_11</code>	0.037 ^d	0.126 ^c	0.176 ^b	0.030
<code>instructor_15</code>	-0.004	0.128	0.176 ^d	0.034
<code>instructor_04</code>	0.006	0.147 ^c	0.222 ^b	-0.036
<code>instructor_03</code>	0.091	0.237 ^b	0.282 ^b	0.166
<code>instructor_12</code>	0.211 ^a	0.274 ^a	0.329 ^a	0.165 ^b
<code>instructor_19</code>	0.105 ^c	0.279 ^a	0.309 ^a	0.241 ^c

student takes the course pertaining to the data point. We see that for the average student, the expected grade decreases by 0.138 for each semester of waiting to take the courses. That said, there may be reasons that a student choose not to take introductory physics straight away (when the value of the variable would be zero), for instance if the student is switching from another major.

Looking at the statistically significant lecturer binary variables in the base model, we see that they have a relatively large magnitude, i.e. they are far from zero. This is an artifact of the hypothesis testing involved, as the model is testing whether the coefficient is different from zero or not. If a professor would have any one-sided effect on the grades of students it would either be a large negative effect or a large positive effect. The maximum difference between the significant lecturer coefficients, instructor 18 and instructor 12 with -0.361 and 0.211 respectively, is 0.572. This is more than half a grade difference, and an astounding result on its own. What is more, we see that we have more negative lecturer variable coefficients than positive in the base

model. Since the outcome variable, `grade`, is defined as difference from the mean grade, this may be a resulting counterbalancing of the positive coefficients associated with `hs_gpa` and `math_placement`.

When we add the binary variables for gender and ethnicity to the base model, we get the extended model. In this model we see a significant drop in the number of lecturer binary variables that are significant, but we still see a very significant and large difference in grades given under the “best” and “worst” lecturer of 0.626. We also see that the newly added binary variables are all significant to a certain degree except for the binary variable representing Asian students. In the data we had 6486 male students and 2155 female students; 7178 students were White, 459 were Black, 511 were Asian and 493 students were in the combined category. From the coefficients in the extended model, we see that Black students and female students are predicted to perform slightly worse than their peers. However, we find that the statistically significant effect of certain lecturers is much larger than effects attributed to ethnicity and gender in the extended model.

The difference of the gender product model when compared to the extended model is that the gender binary variables have been multiplied with the binary variables representing lecturers. This groups the students taught under a specific lecturer into genders. The result is that we see more pronounced gender differences between students taught under single lecturers. The greatest difference is between the male and female students under instructor 1. For all lecturers we see a higher expected grade for male students, compared with female students. In the gender product model we see that the significance of ethnicity variables drops off, when compared to the extended model.

There are several studies on the role of gender as well as ethnicity in physics education research. For instance, McCullough [16] show possible effects of a male context on student responses to conceptual exams. Employing the Physics Anxiety Rating Scale, Agra *et al.* [17] show that female and male students show statistically different levels of anxiety in different situations throughout a physics course. Henderson and Stewart [18] found that White students outperformed Black students on the Force Concept Inventory test.

We would argue that seeing lecturer variable coefficients that are not statistically significant is also an interesting result, and perhaps the desired result. A non-significant coefficient implies that the null hypothesis that a coefficient is equal to zero is not rejected, meaning that the lecturer does not have significant effect on the grades of students. This may be a preferable situation, compared with having a lecturer that has a profound negative impact on female student’s grades.

Apart from achieving a more certain effect per lecturer if the lecturers has had many students, we might expect to see some effect related to experience. That is, a positive relationship between the number of students a lecturer has had and the value of coefficients. Such a relationship does not seem to exist. Arguably the most experienced lecturer in the data set, instructor 12 with 2670 students in total, has a high pos-

itive effect on student grades. On the other hand we have instructor 8, who has taught 1068 students, but have one of the largest negative effect on grades. Another lecturer worth mentioning is instructor 19, who has had the highest effect on student grade in two of the models, but has only taught 196 students. This suggests that there might be some underlying cause that leads to the difference in impact on student grades. One is tempted to suggest that there are lecturers who might be better at teaching, who spends more time to prepare for their teaching obligations or who simply likes teaching better. That said, we also cannot overlook the possibility of different grading practices in different courses, which may influence our results.

For future research, we would like to extend this study to include lecturer effect on persistence of students and their likelihood of finishing a physics degree. There is evidence that certain course choices students make affect the student’s probability of finishing a physics degree [6]. While we opted for a linear model in this study, because of the model’s easily interpretable coefficients, the use of machine learning on educational data has grown considerably in recent years [9, 19]. An alternative method, like the very prolific gradient boosting model [20], may increase fit and yield better predictions.

V. CONCLUSION

In this study we show that which instructor a student has in an introductory physics course has an effect on the grade of the student in that course. We built a large data set of student grades, high school GPAs, math placement scores, instructors, genders and ethnicities from registrar data from a large Midwestern university. We fit this data to three linear models, ensuring ease of interpretability and statistical robustness. We find that there is a statistically significant difference of more than half a grade when comparing “best” and “worst” professors. When effects from gender are kept separate from lecturers, the statistically significant effects of lecturers are much greater than those attributable to genders. However, sub-grouping students such that each professor has a coefficient associated with male and female students gives a very large significant difference between grades achieved by students of different genders under certain lecturers. Effects attributable to ethnicity is always lower than statistically significant effects attributable to lecturers.

ACKNOWLEDGMENTS

This project was supported by the Michigan State University College of Natural Sciences including the STEM Gateway Fellowship and the Lappan-Phillips Foundation, the Association of American Universities, and the Norwegian Agency for Quality Assurance in Education (NOKUT), which supports the Center for Computing in Science Education. This project was also supported by the INTPART project

of the Research Council of Norway (Grant No. 288125).

-
- [1] V. Tinto, Dropout from higher education: A theoretical synthesis of recent research, *Review of educational research* **45**, 89 (1975).
- [2] S. E. Carrell and J. E. West, Does professor quality matter? evidence from random assignment of students to professors, *Journal of Political Economy* **118**, 409 (2010).
- [3] F. Hoffmann and P. Oreopoulos, Professor qualities and student achievement, *The Review of Economics and Statistics* **91**, 83 (2009).
- [4] M. De Paola, Does teacher quality affect student performance? evidence from an italian university, *Bulletin of Economic Research* **61**, 353 (2009).
- [5] E. Bettinger and B. T. Long, *Do college instructors matter? The effects of adjuncts and graduate assistants on students' interests and success*, Tech. Rep. (National Bureau of Economic Research, 2004).
- [6] J. M. Aiken, R. Henderson, and M. D. Caballero, Modeling student pathways in a physics bachelor's degree program, *Physical Review Physics Education Research* **15**, 010128 (2019).
- [7] S. DeVore, J. Yang, and J. Stewart, Extending machine learning to predict unbalanced physics course outcomes, *arXiv preprint arXiv:2002.01964* (2020).
- [8] E. Seymour and N. M. Hewitt, *Talking about leaving* (Westview Press, Boulder, CO, 1997).
- [9] C. Zabriskie, J. Yang, S. DeVore, and J. Stewart, Using machine learning to predict physics course outcomes, *Physical Review Physics Education Research* **15**, 020120 (2019).
- [10] S. Seabold and J. Perktold, statsmodels: Econometric and statistical modeling with python, in *9th Python in Science Conference* (2010).
- [11] P. W. Irving, M. J. Obsniuk, and M. D. Caballero, P3: a practice focused learning environment, *European Journal of Physics* **38**, 055701 (2017).
- [12] V. Sawtelle and S. Gleason, [How a physics class is changing student attitudes for the better](#) (2017), [Online; posted 5-December-2017].
- [13] H. White, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica: journal of the Econometric Society*, 817 (1980).
- [14] H. W. Lilliefors, On the kolmogorov-smirnov test for normality with mean and variance unknown, *Journal of the American statistical Association* **62**, 399 (1967).
- [15] T.-H. Kim and H. White, On more robust estimation of skewness and kurtosis, *Finance Research Letters* **1**, 56 (2004).
- [16] L. McCullough, Gender differences in student responses to physics conceptual questions based on question context, *ASQ Advancing the STEM Agenda in Education* (2011).
- [17] E. Agra, S. M. Fischer, and S. L. Beilock, The role of students' gender and anxiety in physics performance, *PERC Proceedings* (2017).
- [18] R. Henderson and J. Stewart, Racial and ethnic bias in the force concept inventory, *PERC Proceedings* (2017).
- [19] S. B. Kotsiantis, Use of machine learning techniques for educational proposes: a decision support system for forecasting students' grades, *Artificial Intelligence Review* **37**, 331 (2012).
- [20] T. Chen and C. Guestrin, XGBoost: A scalable tree boosting system, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16 (ACM, New York, NY, USA, 2016) pp. 785–794.