

When Research Meets Teaching: Evidence from Student-Level Data*

Annalisa Cristini¹, Federica Origo¹, Piera Bello¹, and Costanza Marconi¹

¹University of Bergamo

Abstract

This paper studies whether the alignment between instructors’ research agendas and the content of the courses they teach affects student outcomes in higher education. Using administrative data from a large Italian university linked to publication records, we construct a novel text-based measure of research–teaching alignment that captures topic similarity between course syllabi and abstracts of instructors’ recent publications. Our empirical strategy exploits within-student and within-instructor variation to address selection concerns. We find that greater research–teaching similarity is associated with higher student performance, with stronger effects in advanced courses and among less-prepared students. Moreover, alignment positively affects graduation rates and final scores, as well as employment probability and wages one year after graduation. We provide suggestive evidence that these effects operate through increased exposure to frontier knowledge and methodologies and through higher perceived teaching quality, as students report greater satisfaction when courses are more closely aligned with instructors’ research.

Keywords: Education; Syllabi; Research; Text Analysis; Teaching Quality

JEL Codes: I23, I26, J24

1 Introduction

Universities are central institutions for both the production and the diffusion of new knowledge. Faculty members are simultaneously researchers, engaged in advancing the scientific frontier, and teachers, responsible for transmitting knowledge to students. While these two roles are formally complementary, the actual relationship between research and teaching remains an open empirical question. A growing body of work examines whether research activity enhances or detracts from teaching quality (Rodríguez and Rubio, 2016; García-Gallego et al., 2015). A key argument in favor of complementarity is that research-active instructors may introduce more recent ideas into their courses, thereby facilitating the diffusion of innovation within higher education. In this regard, Biasi and Ma (2025) documents the existence of an “education innovation gap” driven by frictions in the diffusion of new research into university curricula: even when scientific progress occurs, it does not automatically translate into updated teaching materials. This raises the possibility that the alignment between an instructor’s research agenda and course content could be a crucial channel through which new knowledge reaches students. At the same time, research productivity may also crowd out teaching effort (Braga et al., 2014) or interact non-linearly with pedagogical effectiveness (Artés et al., 2017). Thus, understanding when and how

*The research is funded by the European Union -Next Generation EU, in the framework of the GRINS - Growing Resilient, Inclusive and Sustainable Project (GRINS PE00000018 – CUPE63C22002140007). The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union, and the European Union cannot be held responsible for them.

research affects teaching remains essential for evaluating the role of the university as a vehicle of knowledge transmission.

In this paper, we contribute to this debate by shifting the focus from the quantity of research performed to the alignment between research and teaching. We hypothesize that instructors who teach courses closely related to their research interests deliver higher-quality teaching because alignment enables deeper subject-matter expertise, more recent knowledge, and greater intrinsic motivation. Importantly, we distinguish the content of research from its quantity: we control for publication output, grants, and citations, while directly measuring the similarity between instructors’ research and the syllabi of the courses they teach

To test this mechanism, we construct a text-based measure of similarity between each instructor’s recent publications and the syllabi of the courses they teach. We do so by representing syllabi and abstracts of research articles as term-frequency vectors, reweighting terms using term-frequency–backward inverse document frequency (TF-BIDF) to emphasize informative vocabulary, and computing cosine similarity to obtain a continuous measure of research–teaching alignment.¹

To do so, we use detailed administrative data from a large Italian university between academic years 2015/16 and 2023/24, which include students’ course choices, grades, and academic progression, as well as information on teaching staff and course syllabi. We link these data to instructors’ publication records obtained from the Scopus database.

Our empirical strategy exploits within-student and within-teacher variation. By comparing the same student across courses taught by instructors with different research–teaching similarity, and the same instructor across courses that differ in how closely they align with their research agenda, our estimates are not biased by unobserved student or teacher characteristics that are constant across courses. This approach therefore mitigates the main sources of selection bias, including student sorting, instructor teaching style, and time-invariant differences in research productivity.

Our results show that greater alignment between an instructor’s research agenda and the content of the courses they teach is associated with higher student performance. Students earn higher grades and are more likely to pass exams when taught by instructors whose research is closer to the course material, even after controlling for course characteristics, and student background. Moreover, the effect is stronger in more advanced courses and for students with lower prior preparation.

The positive impact of alignment extends beyond short-term academic outcomes. Students exposed to higher research–teaching similarity over their academic career are more likely to graduate on time, obtain higher final grades, and receive more points on their dissertation. In addition, alignment is positively associated with early labor market outcomes, increasing both the probability of being employed and the likelihood of earning higher wages one year after graduation.

We also provide evidence on the mechanisms underlying these effects. We show that courses with higher research–teaching similarity are characterized by more up-to-date bibliographic references, indicating greater exposure to recent knowledge and frontier developments. In addition, students exposed to higher alignment report greater satisfaction with the teaching staff, consistent with increased engagement and motivation. Together, these findings suggest that alignment improves learning outcomes by both updating course content and enhancing students’ academic motivation.

Our evidence contributes to a broad literature documenting the importance of instructors in shaping student performance in higher education. Prior work has linked student

¹Similar measures of similarity have been employed to analyze research content to various ends, e.g., to estimate the impact of syllabus contents on student mobility and credits validation (Orellana et al., 2018); to evaluate the effects of similarity between senior and junior research agendas on promotion in academia and gender diversity (Bello et al., 2025); and to construct a measure of distance between higher education content from frontier knowledge (Biasi and Ma, 2025).

outcomes to various instructor characteristics, including subject-matter knowledge and formal qualifications (González-Regalado, 2025; Sancassani, 2023), cognitive skills and academic records (Bhai and Horoi, 2019; Hanushek et al., 2016; Kukla-Acevedo, 2009), employment conditions such as tenure status (Figlio et al., 2015; Jones, 2015; Sav, 2012), and demographic match between students and instructors (de Gendre et al., 2024; Gershenson et al., 2022). However, many of these characteristics reflect institutional structures and labor market sorting rather than the content of what is taught (Burtless, 2011). Another strand of research explicitly examines the dual role of university faculty as researchers and teachers by focusing on the effect of research productivity (Rodríguez and Rubio, 2016; Palali et al., 2018; Nur-tegin et al., 2020). Within this literature, findings remain mixed: research productivity may enhance teaching by bringing more current knowledge to the classroom, but it may also reduce effectiveness when high-impact research crowds out time and effort devoted to students (Braga et al., 2014; Coccoresse et al., 2024).

We advance this literature by providing the first empirical evidence that the match between what instructors research and what they teach plays a central role in shaping student achievement in higher education. By shifting attention from how much faculty publish to what they research in relation to what they teach, we identify a previously overlooked channel through which new knowledge is transmitted to students.

The remainder of the paper is organized as follows. Section 2 presents the data and describes the variables on student outcomes. Section 3 details the methodology for computing similarity in research and teaching content, and tests the validity of our metrics. Section 4 describes the main specification of the model for estimating the effects on student outcomes. Section 5 presents first descriptive evidence and discusses the results of the model. Finally, Section 7 offers concluding remarks and summarizes the key findings.

2 Data

The empirical analysis relies on a novel matched student-teacher dataset for a mid-sized Italian university covering academic years 2015/16 to 2023/24, which combines administrative, textual, and survey data. Specifically, student-level administrative records are linked to the full textual content of all course syllabi taught at the university and to instructors’ publication data from Scopus. We further complement these data with student-level survey information on post-graduation outcomes from AlmaLaurea. Below, we provide a detailed description of each data source.

Student Background and Academic Careers

Student-level administrative data include information on academic careers and socio-demographic characteristics for 43,862 students.

Regarding the academic career, we observe the type of degree, field of study (major), the title, date, credits, and score associated with every exam passed. We also have information on the title and abstract of each dissertation for both undergraduate and graduate students. Each dissertation is also associated with the advisor’s name and unique university identifier, as well as with the date of graduation, final score and points received for the dissertation. With respect to socio-demographic characteristics, we observe student gender, year of birth, citizenship, high school education and family socioeconomic status (proxied by the ISEE index, an official indicator of household income that students may choose to share and is used for allocating financial aid). Summary statistics for this sample are shown in Table 1, Panel A. According to the table, 64% of students in our dataset are women, 78% are enrolled in undergraduate programs, and 85% are observed for up to five years.² The panel is unbalanced, since every year we include individuals who enroll for

²Table A1, Panel A, reports analogous summary statistics by broad fields of study, namely Economics and

the first time and we drop individuals who graduate, drop out or transfer to another institution. We complement this dataset with student-level survey data on post-graduation outcomes. These data come from AlmaLaurea, an Italian graduate survey system that is administered to students regularly after graduation and covers employment status, wage, sector of occupation, job stability, further education, as well as satisfaction with their academic experience. We are able to successfully match 7,300 graduate students with their AlmaLaurea questionnaires, collected respectively one, three and five years after graduation.

University Course Syllabi

We link this information to a novel dataset of university syllabi, which we construct by collecting and structuring raw textual information, covering all courses taught between a.y. 2015/16 and a.y. 2023/24. The syllabus of each course is compiled by the instructor in each academic year. The initial sample includes 1,393 syllabi.

Most syllabi share a standard structure, i.e., which includes details on the content and topics of the course, course’s learning goals, and a list of recommended readings and reference texts. For the last three years of our panel, they also report the 2030 Agenda Goals for Sustainable Development corresponding to the course’s subjects. Each syllabus is also associated with basic course details, such as the name and code of the exam, as well as the instructor’s name and academic rank. Syllabi also contain information on partition criteria, when applicable: this usually concerns courses with a large number of students, who are split in different classes, associated to different instructors, based on either their surname’s initials (e.g., A-D, E-N and M-Z) or the final number of their student ID (either even or odd).

Given this general structure, we parse each syllabus and extract four sets of information, which we describe in detail below: (i) basic course details, (ii) the course content, (iii) the list of required and recommended readings, and (iv) partition criteria, when present. We are able to match each exam a student has passed during his/her career with the corresponding syllabus and instructor. We focus on the content section of the course to compute topic similarity between the course and the instructor’s research interests (see Section 3). Table 1, Panel B, shows summary statistics for the syllabi, including the distribution of syllabi across three broad fields (Economics and Management, Engineering, and Humanities and Law), and the number of instructors per course over the nine academic years included in our sample.

Instructors’ Research Output

From the course syllabi, we identify 584 distinct instructors teaching across the eight departments of the university. Table 1, Panel C, summarizes their main characteristics: 28% are assistant professors (approximately 55% of whom are men), 22% are associate professors (about 60% of whom are men), and 14% are full professors (around two-thirds of whom are men).³

For these instructors, we collect information on academic publications using data from Scopus, an Elsevier-owned repository providing titles, abstracts, and citation information for academic articles. The resulting dataset includes 19,866 articles published between 1967 and 2024. For each article, we extract information on the title, abstract, keywords, authors, authors’ affiliations, and research funding acknowledged in the publication.

The distribution of publications per instructor is highly heterogeneous, reflecting the wide range of academic disciplines in the sample and their differing publication practices.

Management, Engineering, and Humanities and Law.

³Table A1, Panel B, reports analogous summary statistics by broad fields of study, namely Economics and Management, Engineering, and Humanities and Law.

It is also strongly right-skewed, with nearly 50% of instructors having authored five or fewer publications. In terms of research funding, more than half of the instructors in our sample have not received any external funding for research projects, and the average number of grants acknowledged across publications is below five (Table 1, Panel C).

Table 1: Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Panel A: Students' Characteristics					
Female	43,862	0.64	0.48	0	1
Italian	43,862	0.95	0.22	0	1
High School Final Score (weighted)	42,683	0.78	0.11	0.60	1
Family Income (ISEE)	30,089	22,734.01	12,691.34	0	57,994.48
Undergraduate	43,862	0.78	0.42	0	1
Average Exam Score	40,029	25.1	3.1	18	30
Graduation Score	23,836	100	8.3	71	110
Points for Dissertation	23,836	4.7	2.0	0	22
Satisfied with Teachers	21,135	0.93	0.26	0	1
Enrolled in Ph.D.	2,713	0.01	0.10	0	1
Intend to Remain in Education	21,036	0.65	0.48	0	1
In Employment	13,380	0.57	0.50	0	1
Wage > €750	6,859	0.73	0.44	0	1
Panel B: Syllabi Characteristics					
Field: Economics and Management	1,393	0.12	0.33	0	1
Field: Engineering	1,393	0.19	0.39	0	1
Field: Humanities and Law	1,393	0.68	0.47	0	1
Number of Teachers per Course	4,181	1.42	0.71	1	6
Panel C: Teachers' Characteristics					
Female	584	0.38	0.49	0	1
Post-doc	584	0.02	0.15	0	1
Assistant Professor	584	0.28	0.45	0	1
Associate Professor	584	0.22	0.41	0	1
Full Professor	584	0.14	0.35	0	1
Publications	584	16.2	39.8	1	637
Citations	584	402.3	1,910.9	0	33,122
Grants	584	4.7	33.2	0	617
Workload: CFU per Year	4,181	2.3	1.03	1	6
Number of Thesis Students per Year	875	1.2	0.43	1	3

Notes: This table reports summary statistics for students, syllabi, and teachers included in the analysis. Panel A refers to student-level characteristics. Family income is proxied by the ISEE index, an official indicator of household income used for allocating financial aid and voluntarily reported by students. Panel B reports characteristics of course syllabi. Panel C refers to teacher-level characteristics. The number of observations varies across variables due to missing information and differences in the unit of observation (student, course, or teacher).

3 The Research-Teaching Similarity Measure

This section introduces the Research–Teaching Similarity Measure and explain how we quantify the alignment between course content and instructors’ research activity. We exploit Natural Language Processing (NLP) techniques applied to textual information from course syllabi and publication abstracts to construct a novel metric capturing the

semantic proximity between what instructors teach and what they study. The section first outlines the methodological steps used to build the similarity index, then presents a set of validation exercises aimed at assessing its empirical relevance. Finally, we provide descriptive evidence on the distribution of the similarity measure across courses, fields of study, and instructor characteristics.

3.1 Methodology

Text Pre-Processing

To analyze the textual data, we follow the approach developed by [Kelly et al. \(2021\)](#) and [Biasi and Ma \(2025\)](#), who use text-based methods to measure technological innovation and frontier knowledge, respectively.

As a first step, we convert each text document—namely, course syllabi and publication abstracts—into a term-frequency vector TF_{dw} , where each element represents the frequency with which term w appears in document d . We address the concern that very common words, such as prepositions, tend to appear in nearly all documents and may therefore receive excessive weight based solely on their frequency, despite having limited explanatory power for document content. To account for this, we downweight common terms by assigning each word a weight that is inversely proportional to its frequency across the entire corpus. Specifically, we construct an inverse document frequency measure IDF_w , defined as the inverse of the share of documents in which term w appears.

Combining these two components, we transform the original term-frequency vectors into term-frequency–inverse-document-frequency (TF–IDF) representations, where the weight assigned to each term is given by:

$$TFIDF_{dw} = TF_{dw} \times IDF_w. \quad (1)$$

This measure captures the relative importance of a term within a given document, accounting for its prevalence across the full corpus.

The Research-Teaching Similarity Index

After representing each document as a TF–IDF vector, we compute cosine similarity, which measures the similarity between a course syllabus and a publication abstract by comparing the direction of their respective vectors rather than their magnitude.⁴ Computation is defined as:

$$\cos(\theta) = \frac{\mathbf{V}_\alpha \cdot \mathbf{V}_\beta}{\|\mathbf{V}_\alpha\| \|\mathbf{V}_\beta\|} \quad (2)$$

where V_α indicates the $TFIDF$ for document α .

Formally, the cosine similarity between any V_α and V_β is the dot product of the pair’s vectors normalized by the vectors’ lengths. The value normally ranges from -1 to 1, where 1 indicates identical direction, 0 indicates orthogonality (no similarity) and -1 indicates completely opposite direction. However, since all our $TFIDF_d$ are non-negative, this measure of similarity lies in the interval [0,1]: a cosine index $\theta_{\alpha,\beta} = 0$ indicate that two vectors have no term in common, whereas $\theta_{\alpha,\beta} = 1$ signal that the two vectors use the exact same set of words.

We compute this index for each course syllabus and for each abstract of the instructor’s publications. We then aggregate the measure at the instructor level by calculating both the average similarity and the maximum similarity between the instructor’s publications

⁴Cosine similarity has been previously used to measure semantic similarity in publication and syllabus data; see, for instance, [Bello et al. \(2025\)](#) and [Biasi and Ma \(2025\)](#).

and the syllabi of the courses they teach. Therefore, for each course taught by a given instructor in a given academic year, we measure the degree of alignment between the content of the course syllabus and the instructor’s research interests.

In the main analysis, we use the maximum similarity index as our primary variable of interest. We rely on the maximum similarity measure rather than the mean because instructors often pursue multiple, distinct lines of research. Using the maximum similarity index allows us to capture whether a course syllabus is closely aligned with at least one of the instructor’s research lines.

Figure A1 shows the distributions of the maximum and mean similarity indices. The two distributions are very similar, although the maximum similarity distribution is shifted to the right, as expected.


3.2 Validating the Teaching-Research Similarity Measure

To validate our similarity measure and assess whether it captures meaningful semantic proximity between course syllabi and publication abstracts, we conduct two simple validation exercises.

First, we test whether the similarity index is higher for syllabus–publication pairs within the same department than for pairs across different departments. To this end, we associate each instructor in our sample with a randomly selected author, subject to the condition that the author’s department of affiliation differs from that of the instructor. For each instructor, we then compute the cosine similarity between the course syllabi they teach and 50 publications authored by the matched (random) author, and construct the average similarity across these pairs.

Finally, we compare the distribution of this placebo similarity measure with the distribution of the similarity index computed using the true instructor–syllabus pairs.

Figure 1: Similarity Validation



Notes: The figure reports kernel density distributions for values of a measure of teaching-research similarity computed between syllabi and abstracts from the same teacher-author, and values of teaching-research similarity computed between syllabi and abstracts from a different teacher-author.

The two measures display markedly different distributions, as shown in Figure 1. The placebo index ranges from 0 to 0.33, with a median value of approximately 0.10. In contrast, our similarity index shows a much wider dispersion, with a median value of 0.26 and values reaching as high as 0.74. This comparison confirms that similarity scores are substantially lower when a syllabus is paired with publications from unrelated fields. Table

A2 also presents examples of syllabus–abstract pairs with low and high textual similarity from this validation exercise.

As a second validation exercise, we examine whether the similarity measure varies systematically with the level of course specialization. Figure A2 reports the full distribution of the similarity index by degree type. The distribution for Master’s courses is slightly shifted to the right relative to Bachelor’s courses, consistent with the more specialized content typically covered at the graduate level.

3.3 Descriptive Evidence on the Similarity Measure

Table 2 reports summary statistics for the similarity measure between course syllabi and instructors’ research output. Overall, the similarity index exhibits substantial variation, with an average value around 0.44, indicating meaningful heterogeneity in the alignment between teaching content and research activity. Noticeable differences emerge across degree types, confirming the previous visual evidence. We distinguish between standard Bachelor’s courses, first-year Bachelor’s courses (which typically cover more introductory and general content), Master’s courses, and five-year single-cycle Master’s programs. The highest levels of similarity are observed in courses belonging to Master’s and five-year Master’s degrees, whereas Bachelor’s courses—particularly those taught in the first year—exhibit lower similarity values.

Differences in similarity are pronounced across broad fields of study. Courses in Humanities and Law display, on average, higher similarity values relative to Economics and Management and Engineering, while also exhibiting greater dispersion. These patterns are consistent with the distributional evidence shown in Figure A3, which reports the full distributions of the similarity index by field of study.

Greater heterogeneity emerges when the similarity measure is disaggregated by instructors’ gender and academic rank. On average, syllabi taught by female instructors exhibit a higher degree of alignment with their research output. Differential patterns are also observed across most academic ranks, from postdoctoral researchers to assistant professors, with lower-ranking positions having the highest mean value of similarity, probably due to the small number of publication outputs and less diversification in research interests at the beginning of their career. Male associate and full professors tend to teach courses that are less closely aligned with their research areas in both Bachelor’s and Master’s programs. This pattern reverses in five-year degree programs,⁵ where male professors exhibit higher research–teaching similarity than female professors.

We then examine how the degree of similarity between teaching and research relates to research productivity. Figure A4 shows how the similarity index varies across quintiles of the publication distribution. It suggests a nonlinear relationship: similarity initially increases with the number of publications, but beyond a certain point, greater research output is associated with lower similarity. One possible explanation is that more senior faculty—who typically have a larger publication record—often maintain a stable teaching load and may invest less in revising or updating their syllabi, resulting in weaker alignment with their evolving research interests. At the same time, this pattern may partly reflect heterogeneity in similarity levels across fields, which differ substantially in their average number of publications.

Lastly, we explore the correlation between research-teaching alignment and student performance. Figure A5 plots average exam scores across quintiles of the similarity distribution. We observe a positive and approximately linear relationship: higher similarity between a course syllabus and the instructor’s research agenda is associated with higher exam scores.

In the next section, we define our identification strategy to provide a causal estimation of the effect of similarity between teaching and research on student performance.

⁵At the University of Bergamo, five-year degree programs include Law and Primary Teacher Education.

Table 2: Similarity Teaching-Research Between Course Syllabi and Instructors' Publications

Variable	Obs	Mean	Std. Dev.	Min	Max
Similarity (All Courses)	7,437	0.44	0.16	0	0.82
Similarity – Bachelor	4,268	0.43	0.16	0	0.82
Similarity – Bachelor (First-year courses)	328	0.42	0.17	0	0.76
Similarity – Master	2,611	0.45	0.15	0	0.81
Similarity – Five-year Degree	558	0.50	0.16	0	0.82
Similarity – Econ. & Management	1,061	0.39	0.15	0	0.74
Similarity – Engineering	2,084	0.37	0.14	0	0.76
Similarity – Humanities & Law	4,292	0.49	0.15	0	0.82
Similarity – Female Teacher	1,908	0.46	0.14	0	0.82
Similarity – Male Teacher	2,923	0.43	0.16	0	0.82
Similarity – Post Doc	87	0.48	0.14	0.22	0.72
Similarity – Researcher	1,709	0.44	0.16	0	0.78
Similarity – Associate Prof.	2,877	0.44	0.15	0	0.82
Similarity – Full Prof.	1,804	0.44	0.14	0	0.79

Notes: This table reports summary statistics of the similarity measure between course syllabi and instructors' research output. Similarity is computed using cosine similarity between TF-IDF representations of syllabus content and publication abstracts. Statistics are reported for the full sample, by degree type (Bachelor's, Master's, and five-year degree programs), by broad field of study (Economics and Management, Engineering, and Humanities and Law), and by gender and academic ranking of the teacher.

4 Empirical Strategy

To estimate the effect of research-teaching similarity on student outcomes, we proceed in two steps. First, we estimate the effect of similarity on short-term academic performance using student-exam level data, focusing on exam outcomes such as exam scores and the probability of passing an exam. Second, we analyze the impact of teaching-research similarity on longer-term academic outcomes and labour market outcomes at the student level.

We begin by estimate the following specification relating student performance to our similarity index:

$$Y_{ict} = \alpha + \beta \text{Similarity}_{jct} + X_i' \phi + X_j' \delta + \eta_t + \zeta_d + \epsilon_{ict}, \quad (3)$$

where Y_{ict} denotes the performance, measured by the probability of passing the exam or the exam score, for student i in course c taught by instructor j in year t . Similarity_{jct} is our similarity index measuring the alignment between course content and the instructor's recent research outputs. X_i and X_j are vectors of observable student (e.g., gender, age, citizenship, high school graduation score, and family income) and instructor characteristics (e.g., publications and grants), respectively. We include academic year fixed effects η_t to account for any time-specific change, as well as type of degree fixed effects ζ_d to level out any degree-specific dynamics affecting both students' outcomes and similarity. Standard errors are clustered at the student level.

A causal interpretation of β in Equation (3) requires that the degree of research-teaching similarity be as good as randomly assigned to students, conditional on the included controls. In practice, however, this assumption is unlikely to hold. First, the allocation of students to instructors is typically non-random: higher-ability or more motivated students may be more likely to enroll in courses taught by highly research-active faculty, or to select courses whose content is more closely aligned with their interests. Second, the extent

to which instructors teach courses that match their research agenda is itself endogenous and may correlate with teaching effectiveness. For example, instructors who are more motivated, more confident in their subject expertise, or who receive greater departmental support may be both more likely to design courses aligned with their research and more effective in the classroom.

If such student- and instructor-level unobservables are correlated with both research–teaching alignment and student achievement, estimates of β based on Equation (3) will be biased. The direction of this bias is generally ambiguous: the effect may be upward biased if high-performing students self-select into aligned courses, but downward biased if research-aligned courses are more advanced or demanding. Likewise, if research-intensive faculty devote less time to teaching-related activities, the estimated relationship between similarity and performance may be negatively confounded.

To mitigate these concerns, we progressively introduce fixed effects. We first include student fixed effects to control for all time-invariant determinants of student achievement:

$$Y_{ict} = \alpha + \beta \text{Similarity}_{jct} + X'_i \delta + \gamma_i + \eta_t + \epsilon_{ict}, \quad (4)$$

where γ_i represents student fixed effects. This specification compares the performance of the *same student* across courses taught by different instructors, thereby removing bias from unobserved differences in ability, motivation, or socioeconomic background. However, unobserved instructor characteristics that are constant across courses and potentially correlated with similarity may still confound the estimates.

We therefore estimate a model that includes instructor fixed effects, while controlling for all student observables:

$$Y_{ict} = \alpha + \beta \text{Similarity}_{jct} + X'_i \phi + \theta_j + \eta_t + \epsilon_{ict}, \quad (5)$$

with θ_j denoting teacher fixed effects. Instructor fixed effects absorb all time-invariant instructor characteristics—such as general teaching skill, pedagogical style, enthusiasm, or professional reputation—that could otherwise jointly influence both research activity and student outcomes.

In the university setting, instructors are more likely to self-select into courses that they particularly enjoy or that closely align with their research interests, and this likelihood tends to increase with seniority and departmental experience. In contrast, student-level sorting is less pronounced, especially in Bachelor’s programs, where many courses are mandatory and class assignments are often random. Some degree of student selection may occur later in their academic trajectory—for example, in Master’s programs—if students choose to take courses with professors they previously encountered and found particularly lenient or effective in their teaching. In this context, student outcomes could be biased if unobserved characteristics that influence grades—such as course preference or responsiveness to teaching style—are correlated with unobserved instructor traits, such as teaching ability. For instance, if high-performing students tend to be matched with instructors who are both highly effective and teach courses closely aligned with their research, estimates of the effect of research–teaching similarity on student outcomes could be upwardly biased.

To rule this option out, we estimate a final model as follows:

$$Y_{ict} = \alpha + \beta \text{Similarity}_{jct} + \mu_{ij} + \eta_t + \epsilon_{ict}, \quad (6)$$

where μ_{ij} denote student by instructor fixed effects. The interaction between student and teacher level fixed effects allows us to compare the same student paired with a teacher for more than one course during the course of her academic career⁶ and to exploit the variation in similarity among courses taught by one teacher. Identification

⁶28% of students observed in our sample meet the same teacher more than once and around 7% of all exam records correspond to a student–teacher pair that occurred more than once

in Equation (6) therefore relies on comparing (i) the same student across courses taught by different instructors and (ii) the same instructor across courses that differ in the degree of research–teaching similarity. This within-student and within-instructor variation substantially reduces the scope for bias arising from non-random matching and selection. Equation (6) is thus our preferred specification.

Then, we then analyze the effect of teaching–research similarity on longer-term academic and labor market outcomes by estimating an analogous specification at the student level. In this case, we aggregate the similarity measure across all courses taken by a given student and relate it to student-level outcomes.

The specification is as follows:

$$Y_i = \alpha + \beta \text{Similarity}_i + \mathbf{X}_i' + \delta_f + \phi_c + \eta_t + \epsilon_i, \quad (7)$$

where Y_i denotes delay at graduation, the final degree score, the dissertation grade, the probability of being employed, and the wage level after graduation. δ_f are type of degree fixed effects, ϕ_c are cohort fixed effects and the other variables are defined as before, with the only difference that the teaching–research similarity measure is aggregated at the student level by calculating the simple average across all courses taken by the student.

5 Results

5.1 Short-Term Academic Outcomes

We begin by analyzing the effect of teaching–research similarity on short-term academic outcomes, focusing on exam scores. The results are reported in Table 3.

Column (1) presents baseline estimates including the similarity index and basic course controls. Column (2) augments the specification with a rich set of student-level characteristics, such as gender, prior academic performance, and other observable background variables. Columns (3) and (4) further introduces student and teacher fixed effects, thereby exploiting within student–teacher variation and absorbing time-invariant heterogeneity in student ability, motivation, teaching quality, and grading standards.

Across all specifications, the coefficient on similarity is positive and statistically significant. However, its magnitude decreases as progressively richer fixed effects are introduced. This pattern indicates that part of the raw correlation between similarity and exam scores reflects sorting across students and instructors. Once we account for these sources of selection, the estimated effect becomes smaller but remains statistically significant even in the most demanding specification with student–teacher fixed effects. This suggests that higher alignment between teaching content and research activity improves exam performance beyond compositional differences.

In terms of magnitude, a one-standard-deviation increase in the similarity index (0.16) is associated with an increase in exam scores of approximately 3–5 pp. While modest in relative terms—around 0.1–0.2 percent of the mean grade—this effect is economically meaningful given the limited variation in grades and the stringent fixed effects structure.

Table A3 report coefficients of all the variables included in the model. In addition to teaching–research similarity, several student characteristics are strongly associated with academic performance. Students with higher secondary school grades and Italian citizenship exhibit a higher probability of higher exam scores. Socio-economic background, proxied by family income (ISEE), displays a strong and statistically significant association with student performance, indicating that students from more advantaged backgrounds achieve systematically better academic outcomes.

Table 3: Exam Score and Similarity Teaching-Research

	Exam Score			
	(1)	(2)	(3)	(4)
Similarity Teaching-Research	2.428*** (0.052)	1.153*** (0.050)	0.192** (0.084)	0.370** (0.167)
Student FE	NO	YES	NO	YES
Teacher FE	NO	NO	YES	YES
StudentxTeacher FE	NO	NO	NO	YES
R^2	0.15	0.43	0.29	0.76
Obs	238409	233119	173027	38594
Mean of dep. var.	26.14	25.66	25.63	26.33

Notes: The table reports OLS regressions of exam scores. The key independent variable is Similarity Teaching-Research, measuring the alignment between the course syllabus and the teacher’s research. Columns (1)–(4) progressively include fixed effects: student, teacher, and student $\times$ teacher. Standard errors are clustered at the student level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. The mean of the dependent variable and the number of observations are reported for reference.

Heterogeneity Analysis

We then conduct a heterogeneity analysis by degree level, academic fields, instructor gender, and student characteristics to assess whether the effect of teaching–research similarity varies across groups.


Bachelor vs Master Students. We start by analyzing differences in the effect of teaching–research similarity across years of study in each degree program. In particular, we want to assess to what extent the effect of similarity on student performance changes in the last year of their study, when courses become more specialized, compared to the first years. To this end, we estimate Equation (5) for each degree type, augmented with interactions between the similarity index and indicators for enrollment in the final year of each degree. The omitted category is the first years. Figure 2, panel (a), reports the coefficients on the interaction terms.

The results show substantial heterogeneity across academic stages. For Bachelor students in their last year, the effect of similarity on grades is significantly stronger than for first-year students. The increase is sizable and corresponds to an additional 0.09-point increase in exam scores due to one-standard deviation increase in similarity ($p < 0.001$). In contrast, for final-year Master students the effect is positive but not statistically distinguishable from zero, indicating that there is no evidence that similarity matters differently for last-year compared first-year Master students. These results do not surprise if we think about the difference in the content of courses across years of a Bachelor’s degree, with the first years being dedicated to several introductory and compulsory disciplines. This distinction between courses in the first and second year is less stark in a Master’s degree, where course are homogeneously more advanced and specialized than at Bachelor’s level. Lastly, for five-year Master students in their last year, the effect of similarity is significantly weaker than for first years, and even detrimental, although the effect is relatively noisy ($p = 0.041$).

These findings suggest that the alignment between teaching and research becomes increasingly important at more advanced stages of study, when courses are typically more specialized and closer to the research frontier. At the Bachelor level, many courses in the first and second year serve as introductory overviews of a discipline, leaving limited room for instructors to tailor content to their own expertise; as a result, syllabi tend to be standardized and undergo few substantive revisions over time. In contrast, last-year Bachelor’s and Master’s courses typically focus on more specialized topics and allow greater flexibil-

ity for instructors to design content—such as readings, materials, and assignments—that reflects their specific research interests. Overall, these findings suggest that the alignment between teaching and research is particularly beneficial in advanced stages of study, where students engage with more specialized content and can better appreciate the depth of expertise that research-active instructors bring to the classroom.

Figure 2: Grading Outcomes and Similarity Teaching-Research, By Course Characteristics



Notes: OLS estimates. The dependent variable is the exam score. Each coefficient is an estimate of β interacted respectively with indicators for type of degree (panel (a)); and with indicators for field of study (panel (b)). The regression model controls for exam credits, student characteristics (i.e., gender, citizenship, family income), and includes academic year and teacher fixed effects. Standard errors are clustered at student level.

Academic Fields. Figure 2, panel (b), reports the heterogeneous effects of teaching–research similarity across academic fields, revealing substantial differences across disciplines. The omitted category is Humanities and Law.

Relative to this group, the effect of similarity is significantly stronger in Engineering. In particular, a 0.1-unit increase in teaching–research similarity is associated with an additional increase of approximately 0.25 points in exam scores for Engineering students

($p < 0.001$). By contrast, we do not detect statistically significant differences between Economics and Humanities, suggesting that the baseline effect is comparable across these fields.

These findings indicate that alignment between teaching and research is particularly important in Engineering. One possible explanation is that engineering courses tend to be more technically demanding and cumulative, often requiring a deep understanding of advanced methodologies and problem-solving techniques. In such contexts, instructors who are actively conducting research in closely related areas may be better equipped to clarify complex concepts, provide cutting-edge applications, and guide students through challenging material. When course content is highly specialized and quantitatively rigorous, the instructor’s frontier expertise may therefore play a more central role in shaping student performance.


Student’s Ability. The impact of teaching–research similarity may be heterogeneous across students’ observable characteristics, such as academic ability and family income. On the one hand, lower-ability students may benefit from greater alignment between teaching and research if instructors’ deeper subject knowledge translates into clearer explanations and more effective pedagogy. On the other hand, these students may be disadvantaged if highly research-oriented syllabi introduce concepts that are excessively advanced for their preparation, such that exposure to frontier methods and topics damages effective learning. For higher-ability students, greater teaching–research similarity may instead be beneficial if delivery by motivated and well-prepared instructors enhances engagement and intellectual stimulation. Alternatively, the marginal effect may be limited for students who are already highly prepared and intrinsically motivated, for whom additional alignment between teaching and research may yield little incremental benefit.

To test for these hypotheses, Figure 3, panel (a) shows estimates from a version of equation 5 where β is allowed to vary by quintiles of each student’s high school graduation score and the middle quintile (Q3) is used as the reference category⁷. The estimates reveal a non-linear pattern. For students in the lowest quintile (Q1), the effect of similarity is statistically indistinguishable from that of the baseline group ($\beta = 0.143$, $p = 0.39$), indicating that the benefit of similarity for low-ability students is extremely similar to the benefit for those with intermediate ability. Slightly better-performing students (Q2) exhibit a somewhat larger effect ($\beta = 0.332$, $p = 0.056$), although the difference is only marginally significant. In contrast, the effect declines for higher-ability students: the coefficient for Q4 is negative and marginally significant ($\beta = 0.292$, $p = 0.093$), while the highest-ability group (Q5) experiences a significantly smaller effect than the baseline ($\beta = 0.553$, $p = 0.002$). Taken together, these results suggest that the benefits of teaching–research similarity are concentrated among students with poor to intermediate levels of prior ability, whereas the returns are little for top-performing students. This is in line with our hypothesis that high-ability students are less dependent on course-specific alignment to perform well thanks to higher skills and better preparation, while students at the bottom and middle of the distribution may benefit more from instructors with a deep and up-to-date knowledge of their field of teaching.


⁷For simplicity, the analysis is restricted to observations with graduation scores between 60 and 100L.

Figure 3: Grading Outcomes and Similarity Teaching-Research, By Student Characteristics

(a) Exam Score by High School Quintiles



(b) Exam Score by Income Quintiles



Notes: OLS estimates. The dependent variable is the exam score. Each coefficient is an estimate of β interacted respectively with indicators for quintiles of high school graduation scores weighted against the maximum possible score (panel (a)); and with indicators for quintiles of the ISEE index (panel (b)). The third quintile (Q3) is used as the reference group. The regression model controls for exam credits, student characteristics (i.e., gender, citizenship, family income), and includes academic year, type of degree and teacher fixed effects. Standard errors are clustered at student level.

Student’s Socioeconomic Status. We also explore the differential effect of exposure to high research-teaching similarity for students with diverse socio-economic background. The relationship is not a priori straightforward. The effect may be weaker for students from lower socio-economic backgrounds if they face greater obstacles in fully benefiting from their academic programs, due to more limited family resources, higher opportunity costs of time, or part-time work commitments. At the same time, it may be stronger if research-active instructors provide access to knowledge, and skills that these students are less likely to acquire through their family environment. Figure 3 panel (c) reports estimates from a version of equation 5 where β is interacted with quintiles of family income categories. The results indicate that the effect of similarity is significantly larger for students in the lowest income quintile ($\beta = 0.347$, $p = 0.046$), implying that teaching–research alignment benefits economically disadvantaged students more strongly. In contrast, the estimates for the remaining quintiles are small and statistically indistinguishable from zero, suggesting no meaningful differences relative to the baseline group. Overall, the evidence points to heterogeneity concentrated at the bottom of the income distribution, with similarity primarily contributing to performance among lower-income students, while effects are broadly similar across middle- and high-income groups.

5.2 Long-Term Outcomes

Graduation

We then investigate whether exposure to greater teaching–research similarity has persistent effects on long-term academic outcomes, in particular on students’ graduation results.

Table 4 reports the coefficients from Equation (7), where we use a student-level measure of similarity aggregated across all courses attended during the academic career. We examine its effect on several outcomes: the probability of graduating on time (as opposed to graduating with delay), the final graduation grade, and the number of points awarded for the dissertation.

The results indicate that greater teaching–research similarity significantly increases the likelihood of graduating on time, as well as both the final graduation score and the dissertation evaluation.

In terms of magnitude, a one-standard-deviation increase in similarity (0.16) raises the final graduation grade by approximately 0.97 points. Relative to a mean graduation score of 99.8, this corresponds to roughly a 1 percent increase. The same increase in similarity is associated with an additional 0.11 points in the dissertation evaluation, which represents about 2 percent of the average dissertation score (4.7 points).

Overall, these findings provide consistent evidence that higher alignment between teaching and research not only improves short-term academic performance, but also translates into meaningful gains in long-term outcomes, including graduation performance.

These long-term effects likely operate through a research-training channel. Sustained exposure to research-aligned teaching may gradually shape students’ ability to think critically, apply advanced methodologies, and engage with frontier topics. This accumulated research-oriented training becomes particularly valuable at the dissertation stage, where independent work and methodological rigor are essential. The stronger effect observed on dissertation points suggests that teaching–research alignment primarily enhances students’ capacity to produce higher-quality thesis work, which ultimately translates into better graduation outcomes.

Labour Market Outcomes

Finally, we examine whether exposure to a high degree of teaching–research similarity has significant effects on early labor market outcomes. To do so, we rely on information from the AlmaLaurea survey described in Section 2, which is administered to graduates one

Table 4: Long-Term Academic Outcomes and Similarity Teaching-Research

	Delayed Grad.	Final Score	Points to Dissertation
Similarity Teaching-Research (Agg.)	-0.138*** (0.040)	6.113*** (0.476)	0.732*** (0.132)
R^2	0.71	0.38	0.26
Obs	12481	15043	15043
Mean of dep. var.	0.44	99.78	4.70
Student Controls	YES	YES	YES
Year FE	YES	YES	YES
Cohort FE	YES	YES	YES
Field FE	YES	YES	YES

Notes: The table reports the relationship between Similarity Teaching-Research (a measure of the alignment between the course syllabus and the instructor's research) and three long-term academic outcomes: delayed graduation (binary), final score (continuous), and points toward the dissertation (continuous). All regressions include student-level controls, year, cohort, and field fixed effects. Standard errors are clustered at the student level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels. The mean of the dependent variable and the number of observations are reported for reference.

Table 5: Labour Market and Similarity Teaching-Research

	Employed	Wage over €750
Similarity Teaching-Research (Agg.)	0.374*** (0.050)	0.369*** (0.059)
R^2	0.04	0.07
Obs	7289	3998
Mean of dep. var.	0.59	0.74
Student Controls	YES	YES
Year FE	YES	YES
Cohort FE	YES	YES
Field FE	YES	YES

Notes: The table reports the relationship between Similarity Teaching-Research (a measure of the alignment between the course syllabus and the instructor's research) and two labour market outcomes: probability of being employed (binary) and probability of having a wage higher than €750 (binary) one year after graduation. All regressions include student-level controls, year, cohort, and field fixed effects. Standard errors are clustered at the student level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels. The mean of the dependent variable and the number of observations are reported for reference.

year after completion of their degree and collects detailed information on their current labor market status.

Table 5 reports the estimated impact of the aggregated similarity index on two post-graduation outcomes: the probability of being employed one year after graduation and the probability that the first post-graduation wage exceeds €750⁸.

Following the approach adopted in the previous section, we regress each outcome on the student-level aggregated similarity index, controlling for gender, citizenship, high school graduation score, and family income, as well as academic year, cohort and degree-type fixed effects.

The estimates indicate that a one-standard-deviation increase in teaching–research similarity (0.16) raises the probability of being employed one year after graduation by approximately 0.59 percentage points. Relative to a mean employment probability of 59 percent, this corresponds to about a 10 percent increase. The same increase in similarity raises the probability that the first post-graduation wage exceeds €750 by 5.9 percentage points, which represents an increase of roughly 8 percent relative to the mean probability of 75 percent.

Overall, these findings suggest that greater alignment between teaching and research not only improves academic performance and graduation outcomes, but also translates into economically meaningful gains in early labor market success.

This positive effect could stem from the fact that greater alignment between teaching and research enhances students’ enthusiasm and commitment to their academic path, while fostering skills that increase their productivity and make them more competitive candidates upon entering the labor market.

6 Discussion

In this section, we discuss the mechanisms that may explain our findings. Based on the empirical literature, teaching–research similarity may affect student outcomes through several channels.

First, instructors who actively conduct research in the same areas they teach may be more effective educators. Their deeper expertise and intrinsic interest in the subject can translate into clearer explanations, richer examples, and greater enthusiasm in the classroom, thereby fostering student engagement and participation. Second, research-active instructors are more likely to be familiar with the latest methodologies and developments in their field. This may allow them to deliver more up-to-date and rigorous course content, equipping students with advanced analytical tools and skills that enhance both academic performance and labor market competitiveness.

Although we cannot fully disentangle these channels—and they are likely not mutually exclusive—we provide suggestive evidence on the underlying mechanisms. First, we examine whether students whose academic careers were characterized by greater exposure to teaching–research similarity report higher overall satisfaction with their university experience. Second, we investigate whether courses with higher teaching–research similarity are more likely to incorporate innovative content and frontier knowledge (Biasi and Ma, 2025).

We begin by testing whether teaching–research similarity operates through higher perceived teaching quality. To do so, we exploit a question included in the AlmaLaurea survey in which students are asked to report their level of satisfaction with the teaching staff encountered during their academic career.⁹ Table 6 reports the analysis for this variable.

⁸The survey asks respondents about their net monthly wage and record responses on an ordinal scale: €750 corresponds to the first quartile of the wage’s distribution.

⁹The survey question asks: “Are you satisfied about the relationship you had with the teaching staff in general?” Responses are recorded on an ordinal scale and recoded into a binary indicator equal to one for high levels of satisfaction.

Across all specifications—progressively including year, cohort, and field fixed effects—the coefficient on teaching–research similarity is positive and statistically significant. The magnitude of the estimates is stable, indicating that a one-standard-deviation increase in similarity raises the probability of reporting high satisfaction by approximately 1 percentage point. These findings provide suggestive evidence that alignment between teaching and research improves students’ perceptions of teaching quality, which may represent one channel through which similarity affects academic and labor market outcomes.

Table 6: Satisfaction with Teaching Staff and Similarity Teaching-Research

	Satisfaction with Teachers			
	(1)	(2)	(3)	(4)
Similarity Teaching-Research (Agg.)	0.076*** (0.019)	0.078*** (0.019)	0.098*** (0.025)	0.081*** (0.026)
R^2	0.003	0.01	0.01	0.01
Obs	42,705	42,705	42,705	42,705
Mean dep. var.	0.93	0.93	0.93	0.93
Year FE	NO	YES	YES	YES
Cohort FE	NO	NO	YES	YES
Field FE	NO	NO	NO	YES

Notes: The table reports OLS estimates for exam scores. The key explanatory variable is *Similarity (Syllabus–Research)*, which measures the degree of alignment between the content of a course syllabus and the instructor’s research activity, aggregated at student $\times$ year level. Columns (1)–(4) progressively include year fixed effects, cohort fixed effects, and academic field fixed effects. All specifications control for a rich set of student characteristics, unless absorbed by fixed effects. Standard errors are clustered at the student level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

In order to verify our other channel, namely, that similarity between teaching and research entails a degree of innovation and frontier knowledge included in course curricula, we rely on the bibliography portion of the syllabus. From this section, we extract all years of publication of articles and volumes that are indicated as reading materials for the exam. We compute what we called the *bibliographic distance*, i.e., the distance between the year of the course and the most recent year of publication mentioned among the course’s reading materials.

We explore this relation by regressing similarity on our indicator of bibliography distance. We include course fixed effects and cluster standard errors at course level. Estimates reported in Table 7 show that similarity is negatively related to the gap between the course year and the most recent publication year in the bibliography: a larger lag is associated with lower similarity. Quantitatively, a one-year increase in this distance reduces similarity by approximately 0.005 percentage points, corresponding to about 0.45% relative to the mean level of similarity, and the effect is statistically significant at the 5% level.

Figure A6 further explores the relationship between teaching–research similarity and the time gap between the publication year of the bibliography and the year in which the course is offered. The figure confirms that greater alignment between teaching and research is associated with a smaller time gap, indicating a more up-to-date bibliography.

Together, these analyses allow us to assess whether alignment between teaching and research operates through improved student perceptions and greater exposure to cutting-edge material.

While preliminary and exploratory in nature, these findings offer consistent and suggestive evidence that teaching–research alignment may operate through higher perceived teaching quality and greater exposure to frontier knowledge.

Table 7: Bibliographic Distance and Teaching–Research Similarity

	Similarity Teaching–Research	
	(1)	(2)
Bibliographic Distance	-0.005*** (0.001)	-0.002** (0.019)
R^2	0.02	0.89
Observations	124,629	124,626
Mean dep. var.	0.44	0.44
Exam FE	NO	YES

Notes: The table reports OLS estimates for bibliographic distance, which is the difference between the year of the course and the most recent year of publication of any bibliographic material in the course’s syllabus. The key explanatory variable is *Similarity (Teaching–Research)*, which measures the degree of alignment between the content of a course syllabus and the instructor’s research activity, aggregated at student $\times$ year level. Column (2) include exam fixed effects. Standard errors are clustered at the exam level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

7 Conclusion

This paper investigates the relationship between teaching–research alignment and student performance using administrative data from a university in Northern Italy covering the academic years 2015/16 to 2023/24. We introduce a novel measure of teaching quality based on the textual similarity between course syllabi and instructors’ research output. The underlying idea is that greater similarity reflects a closer match between what instructors teach and their area of expertise, potentially capturing depth of knowledge, engagement with the subject, and exposure to recent developments in the field.

To identify the causal effect of similarity, we estimate models with progressively richer fixed effects. Our preferred specification includes interacted student and teacher fixed effects, absorbing all time-invariant characteristics at both levels and mitigating concerns about sorting and unobserved heterogeneity.

We consistently find that higher teaching–research similarity improves student outcomes. In the short run, it increases exam performance, with stronger effects in more advanced programs and in technically demanding fields such as Engineering. At the aggregate level, sustained exposure to aligned teaching raises the likelihood of timely graduation, improves final grades and dissertation evaluations, and is positively associated with early labor market success.

Beyond documenting these effects, our analysis sheds light on the mechanisms at play. Teaching–research alignment appears to operate through a frontier-exposure channel: when instructors teach topics close to their active research, students are more likely to be exposed to frontier knowledge, updated methodologies, and research-oriented ways of thinking. This exposure enhance both perceived teaching quality and the acquisition of advanced skills, particularly those relevant for independent work such as the dissertation. In turn, stronger analytical and methodological training can translate into higher productivity and better labor market performance after graduation.

Overall, our findings provide new evidence on the research–teaching nexus, suggesting that alignment between teaching and research is not merely a feature of academic organization, but a meaningful determinant of student achievement. These results have implications for course allocation policies and for broader debates on how universities can balance research activity and teaching effectiveness.

Appendix

Table A1: Summary Statistics by Field of Study

Variable	Economics and Management			Engineering			Humanities and Law		
	Obs	Mean	SD	Obs	Mean	SD	Obs	Mean	SD
Panel A: Students' Characteristics									
Female	8,058	0.50	0.50	6,864	0.25	0.44	28,940	0.78	0.41
Italian	8,058	0.92	0.27	6,864	0.94	0.23	28,940	0.95	0.21
High School Final Score (weighted)	7,762	0.77	0.11	6,647	0.79	0.11	28,274	0.78	0.11
Family Income (ISEE)	5,069	23,452.9	13,535.5	4,349	24,225.5	13,245.7	20,671	22,243.9	12,320.6
Undergraduate	8,058	0.77	0.42	6,864	0.84	0.37	28,940	0.77	0.42
Average Exam Score	7,459	24.4	3.1	6,854	23.7	3.0	25,716	25.7	2.9
Graduation Score	5,736	96.8	9.2	3,758	97.2	9.7	14,342	101.8	6.8
Points for Dissertation	5,736	3.9	1.8	3,758	4.8	2.4	14,342	5.0	1.8
Satisfied with Teachers	4,778	0.92	0.27	2,892	0.92	0.27	13,465	0.93	0.25
Enrolled in PhD	736	0.01	0.10	443	0.03	0.16	1,534	0.01	0.08
Intend to Remain in Education	4,743	0.63	0.48	2,885	0.68	0.47	13,408	0.65	0.48
In Employment	3,171	0.57	0.50	1,932	0.50	0.50	8,277	0.58	0.49
Wage > €750	1,597	0.80	0.40	874	0.79	0.41	4,388	0.70	0.46
Panel B: Teachers' Characteristics									
Female	81	0.42	0.50	149	0.28	0.45	354	0.42	0.49
Post-doc	81	0.00	0.00	149	0.02	0.14	354	0.03	0.17
Assistant Professor	81	0.28	0.45	149	0.36	0.48	354	0.25	0.43
Associate Professor	81	0.37	0.49	149	0.23	0.42	354	0.18	0.38
Full Professor	81	0.22	0.42	149	0.15	0.36	354	0.12	0.32
Publications	81	10.3	9.0	149	39.1	71.5	354	7.9	12.6
Citations	81	188.4	269.5	149	1,155.1	3,635.7	354	134.4	388.5
Grants	81	2.3	3.7	149	14.5	64.6	354	1.2	4.3
Number of Courses per Year	570	2.2	0.98	1,062	1.9	0.83	2,549	2.47	1.06
Number of Thesis Students per Year	174	1.25	0.47	305	1.17	0.39	396	1.19	0.44


Notes: This table reports summary statistics by field of study. Panel A refers to student-level characteristics, while Panel B refers to teacher-level characteristics. Family income is proxied by the ISEE index, an official indicator of household income used for allocating financial aid and voluntarily reported by students. The number of observations varies across variables due to missing information and differences in the unit of observation. SD denotes standard deviation.

Table A2: Examples of Syllabus–Abstract Similarity (Low vs High)

Course	Similarity Index	Excerpt (Syllabus / Abstract Match)
24133 – LIN-GUA E LET-TERATURA LATINA B	0.00000	Syllabus: Augustus, <i>Res gestae</i>. Creation of skills to deal with translation and commentary of Latin texts and their historical and cultural contextualization. Abstract: The ratio of production cross-sections of $\psi(2S)$ over J/ψ mesons as a function of charged-particle multiplicity in proton-proton collisions at $\sqrt{s} = 13$ TeV is measured with a data sample collected by the LHCb detector, corresponding to an integrated luminosity of 658 pb^{-1}.
87092 – DIRITTO TRIBUTARIO	0.66240	Syllabus: Tax law and the concept of tax. The sources of Tax Law. The constitutional principles in tax law: the legality principle and the principle of the ability to pay tax. The effectiveness of tax rules over time and space. The interpretation of the rules and tax avoidance. Power to tax and taxable subjects. The tax return. The tax assessment. Abstract: This chapter discusses comparative tax law. The basic purpose of comparative law (including that of comparative tax law) is knowledge of foreign and national law: comparative law looks at different rules and institutions to establish to what extent they are similar or different. A curious and well-known result of comparative law is that it may lead to a better understanding of the law of one’s own country and its evolution by comparison with other legal systems. The boundaries of comparative tax law, however, are not easily defined. Comparative tax law may be applied to income tax, VAT, inheritance tax, and other taxes or compulsory contributions with regard to the specific rules, institutions, and structures of a certain tax.


Notes: Examples of low and high similarity values computed via TF–IDF cosine similarity between course syllabi and research abstracts.

Figure A1: Similarity: Maximum vs. Average




Notes: The figure reports kernel density distributions for values of two different measures of research-teaching similarity, respectively: the average similarity value between syllabi and abstracts of one author, and the highest similarity values between syllabi and publications of the same author.

Figure A2: Similarity: By Type of Degree




Notes: The figure reports kernel density distributions for values of teaching-research similarity for Bachelor's and Master's courses.

Figure A3: Similarity: By Field




Notes: The figure reports kernel density distributions for values of teaching-research similarity by field of study, respectively: Economics and Management, Engineering, Humanities and Law.

Figure A4: Similarity Teaching-Research and Nr. Publications




Notes: Points show the mean teaching-research similarity for courses taught by instructors in each quintile of total publications. Vertical lines represent 95% confidence intervals. Quintiles are computed based on professors' total number of publications.

Figure A5: Similarity Teaching-Research and Exam Score



Notes: Points show the mean exam score for courses in each quintile of teaching-research similarity. Vertical lines represent 95% confidence intervals. Quintiles are computed based on the similarity between teaching content and research activity of the instructor.

Figure A6: Similarity and Bibliographic Distance



Notes: The figure shows the relationship between teaching-research similarity and bibliographic distance. Each point represents the average within a bin. Higher similarity is associated with smaller bibliographic distances.

Table A3: Exam Score and Similarity Teaching-Research

	Exam Score			
	(1)	(2)	(3)	(4)
Similarity Teaching-Research	1.760*** (0.119)	1.153*** (0.050)	0.192** (0.084)	0.370** (0.167)
CFU	0.013* (0.007)	0.007** (0.003)	-0.003 (0.004)	-0.001 (0.007)
Gender	0.086 (0.057)		0.068** (0.031)	
Italian Citizenship	0.851*** (0.137)		0.895*** (0.067)	
Final Grade Secondary School	5.917*** (0.202)		6.864*** (0.115)	
ISEE	0.000*** (0.000)		0.000*** (0.000)	
Female Instructor	-0.349*** (0.035)	-0.112*** (0.015)		
Post Doc	0.000 (.)	0.000 (.)		
Researcher	0.541*** (0.197)	0.225*** (0.080)		
Associate Professor	0.700*** (0.195)	0.251*** (0.079)		
Full Professor	0.582*** (0.195)	0.098 (0.080)		
Nr. Publications	0.001 (0.001)	-0.001*** (0.000)		
Student FE	NO	YES	NO	YES
Teacher FE	NO	NO	YES	YES
StudentxTeacher FE	NO	NO	NO	YES
R^2	0.15	0.43	0.29	0.76
Obs	38594	38594	38594	38594
Mean of dep. var.	26.14	25.66	25.63	26.33

Notes: The table reports OLS estimates for exam scores. The key explanatory variable is *Similarity (Teaching-Research)*, which measures the degree of alignment between the content of a course syllabus and the instructor's research activity. Columns (1)–(4) progressively include student fixed effects, teacher fixed effects, and student $\times$ teacher fixed effects. All specifications control for a rich set of student and instructor characteristics, unless absorbed by fixed effects. Standard errors are clustered at the student level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

References

- Artés, J., F. Pedraja-Chaparro, and M. del Mar Salinas-Jiménez (2017). Research performance and teaching quality in the spanish higher education system: Evidence from a medium-sized university. *Research policy* 46(1), 19–29.
- Bello, P., A. Casarico, and D. Nozzay (2025). Research similarity and women in academia. *The Economic Journal*, ueaf113.
- Bhai, M. and I. Horoi (2019). Teacher characteristics and academic achievement. *Applied Economics* 51(44), 4781–4799.
- Biasi, B. and S. Ma (2025). Frontier knowledge in college and student success.
- Braga, M., M. Paccagnella, and M. Pellizzari (2014). Evaluating students’ evaluations of professors. *Economics of Education Review* 41, 71–88.
- Burtless, G. (2011). *Does money matter?: The effect of school resources on student achievement and adult success*. Brookings Institution Press.
- Coccorese, P., R. Dell’Anno, and M. Restaino (2024, None). Are outstanding researchers also top teachers? exploring the link between research quality and teaching quality. *Socio-Economic Planning Sciences* 96(C), None.
- de Gendre, A., J. Feld, N. Salamanca, and U. Zölitz (2024). Same-sex teacher effects. *Working paper series/Department of Economics* (438).
- Figlio, D. N., M. O. Schapiro, and K. B. Soter (2015). Are tenure track professors better teachers? *Review of Economics and Statistics* 97(4), 715–724.
- García-Gallego, A., N. Georgantzís, J. Martín-Montaner, and T. Pérez-Amaral (2015). (how) do research and administrative duties affect university professors’ teaching? *Applied Economics* 47(45), 4868–4883.
- Gershenson, S., C. M. Hart, J. Hyman, C. A. Lindsay, and N. W. Papageorge (2022). The long-run impacts of same-race teachers. *American Economic Journal: Economic Policy* 14(4), 300–342.
- González-Regalado, C. (2025, Feb). Teacher quality, family background and student performance in ecuador. Working Papers hal-05130196, HAL.
- Hanushek, E. A., S. G. Rivkin, and J. C. Schiman (2016). Dynamic effects of teacher turnover on the quality of instruction. *Economics of Education Review* 55, 132–148.
- Jones, M. D. (2015). How do teachers respond to tenure? *IZA Journal of Labor Economics* 4, 1–19.
- Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights* 3(3), 303–320.
- Kukla-Acevedo, S. (2009). Do teacher characteristics matter? new results on the effects of teacher preparation on student achievement. *Economics of Education Review* 28(1), 49–57.
- Nur-tegin, K., S. Venugopalan, and J. Young (2020, October). Teaching load and other determinants of research output among university faculty. *The American Economist* 65(2), 300–311.

- Orellana, G., M. Orellana, V. Saquicela, F. Baculima, and N. Piedra (2018). A text mining methodology to discover syllabi similarities among higher education institutions. In *2018 International Conference on Information Systems and Computer Science (INCISCOS)*, pp. 261–268. IEEE.
- Palali, A., R. van Elk, J. Bolhaar, and I. Rud (2018, None). Are good researchers also good teachers? the relationship between research quality and teaching quality. *Economics of Education Review* 64(C), 40–49.
- Rodríguez, R. and G. Rubio (2016). Teaching quality and academic research. *International review of economics education* 23, 10–27.
- Sancassani, P. (2023, None). The effect of teacher subject-specific qualifications on student science achievement. *Labour Economics* 80(C), None.
- Sav, G. T. (2012). Female faculty, tenure, and student graduation success: Efficiency implications for university funding. *International Journal of Business Management & Economic Research* 3(6).