Student Demand and the Supply of College Courses*

Jacob Light†

February 19, 2025

Abstract

In an era of rapid technological and social change, how do universities adapt? To examine this, I extracted the information contained in the course catalogs of over 700 US colleges and universities, observing courses dating back to 1998. When there are changes in student demand, universities adjust course quantity substantially less than one-for-one, particularly in fields experiencing declining demand. Research universities (R1) are somewhat more responsive in expanding high-demand fields and less research-intensive universities in reducing declining fields. Using Natural Language Processing, I further show that while the content of existing courses remains largely unchanged, newly created courses incorporate topics related to social justice and job skills. R1 and Liberal Arts universities exhibit the most pronounced content changes, with a notable increase in emphasis on social justice.

JEL Classification: A20, I23, I24, I26, J24

*Thanks to Nick Bloom, Eric Hanushek, Caroline Hoxby, and Petra Persson for substantial guidance and support. I also thank seminar participants at CUNY, the Hoover Institution, Stanford University, the University of Zurich, the University of California Santa Barbara, Vanderbilt University, the All California Labor Economics Conference, and the Southern Economic Association Conference for helpful feedback. This research was supported by grants from George P. Shultz Dissertation Support Fund at SIEPR, the Institute for Research in the Social Sciences at Stanford University, and the Leonard W. Ely and Shirley R. Ely Graduate Student Fellowship through a grant to the Stanford Institute for Economic Policy Research. All errors are my own.

†Hoover Institution. jdlight@stanford.edu

1 Introduction

Universities have long been integral to the fabric of American society, shaping its workforce, fostering innovation, and driving the creation of knowledge. Their enduring relevance, even amid economic and societal change, reflects a capacity for adaptation. The evolution of Harvard College exemplifies this institutional flexibility. Established in 1636 to train Puritan clergy, Harvard has since transformed into a leading global research institution. This journey from narrow specialization to intellectual breadth mirrors the evolution of many of today's elite universities (MacLeod and Urquiola 2021).[1]

In recent decades, technological disruption and economic restructuring have transformed the skills that employers seek (e.g., Autor et al. 2003). At the same time, recent global conflicts and intensified public discourse on social inequities have heightened the prominence of societal issues on college campuses. Against this backdrop, higher education leaders face a fundamental challenge: how quickly and effectively can universities realign their curricula when student preferences and broader social imperatives change?

Despite widespread debate over whether universities sufficiently adapt to evolving labor market and societal demands, systematic empirical evidence on this question remains limited. To measure responsiveness, I combine novel data on course offerings and enrollment at a large sample of US universities with an instrumental variables approach that isolates changes in student demand from confounding factors like institutional policy shifts. Specifically, I use region-by-field variations in employment growth to pinpoint where demand for new skills should be highest. Then, through text analysis of course descriptions, I assess whether universities update not just the number of courses offered but the substance of what is taught. The results reveal that institutions expand and contract course offerings far less than one-for-one in response to changes in student demand — particularly when demand decreases — and that new course creation, rather than a major overhaul of existing classes, is the primary channel for curricular change. Differences by institution type further highlight how organizational priorities and resource constraints shape universities' capacity to adapt.

To measure course supply and student demand, I constructed a novel dataset containing granular course-level information for a large and nationally representative sample of US universities. The dataset includes the complete set of courses offered by 782 colleges and universities, which collectively enroll over 52% of all US baccalaureate-level undergraduate students, amounting to over 29 million course sections offered since 1998. I collected the

[1] Harvard's transformation, particularly its elevation to a premier research institution, owes much to Charles Eliot's leadership in the late 19th Century. Drawing on European models, Eliot introduced sweeping reforms, including expanding research initiatives, diversifying graduate programs, and creating a flexible course selection system that reshaped curricula in sciences, history, languages, and social sciences.

data by scraping information from online course catalogs and schedules. I observe details such as the instructor(s), section-level enrollment, instructional format (e.g., in-person or online), and a brief text description of the course content.

Using this new dataset, I analyze how universities adjust their supply of courses in response to changes in student demand. I consider two margins of course supply adjustment. Along the extensive margin, universities meet changing demand for courses in a field of study (hereafter, "field") by adjusting the quantity of courses or sections.[2] Along the intensive margin, fields within universities might meet changing demand by modifying the content of the courses they offer.

I describe supply responses on the extensive margin by estimating the elasticity of the quantity of courses a field offers in response to changes in demand for that field. Such estimation is complicated because course enrollment, which I use as a proxy for demand, is influenced by both demand and supply factors. Students cannot enroll in courses that do not exist or that are rationed. These constraints can distort the relationship between enrollment and demand, making it difficult to separate genuine shifts in student interest from supply-driven limitations. To address this, I construct a shift-share instrument based on field-specific employment growth, allowing me to isolate the component of enrollment changes driven by evolving labor market conditions. The instrument allows me to focus on student demand for fields or skills rather than universities' supply of courses that reflect those fields or skills.

While students are responsive to changing conditions in the labor market, universities adjust the quantity of courses less than one-for-one to this changing demand. On average, a 10% change in demand for a field leads to a 2.2% change in the quantity of courses and 5.5% change in the quantity of course sections in that field. The elasticity varies across fields. Course quantity is relatively more elastic when demand for a field is growing and less elastic when demand is decreasing. The results suggest that the constraints that an institution faces when it seeks to grow a field may differ from the constraints it faces when it seeks to shrink a declining field.

Universities can respond to changing student demand not just by offering more or fewer courses in a particular field, but also by modifying the content of pre-existing courses. For example, an Economics course might incorporate programming concepts to cater to surging demand for Computer Science skills. This adaptation, occurring within the content of courses rather than through the quantity of courses offered, represents a response on the university's

[2] A large lecture course in the Principles of Economics might be offered at multiple times, in multiple semesters, with multiple instructors. In this example, Principles of Economics is a single course and each instance of Principles of Economics during a term counts as a single section.

intensive margin. If fields also respond to students' changing demand on the intensive margin, the elasticities described above may understate the university's responsiveness to changing demand.

I apply Natural Language Processing (NLP) techniques to course descriptions to measure course content changes related to student demand. The central challenge in this analysis is the lack of a direct measure of students' preferences for specific topics or skills. To overcome this, I develop a method that assesses course content in relation to broader themes that students might consider when selecting courses.[3]

The text analysis proceeds in two steps. In the first step, I develop representations of course descriptions using Term Frequency-Inverse Document Frequency (TF-IDF), a widely-used method for representing text documents as vectors. In the second step, I assess each course description's "alignment" to the selected themes by scoring words in the description based on their frequency in representative documents of each theme. For example, I measure a course's alignment with current events based on how commonly key words and phrases in the course description appear in New York Times articles, or its alignment to job market demands based on the frequency with which key words and phrases appear in job descriptions.[4] To my knowledge, the approach I describe is a novel extension of NLP methods and is particularly appealing for its transparency and interpretability.

In general, course offerings are highly stable: 65% of upper-level courses offered in 2022-23 have existed for a decade or more. After their introduction, the central topics and skills emphasized in a course change infrequently. Thus, changes in the topics and skills offered by courses come primarily through the introduction of new courses. Through this channel of new courses, universities gradually incorporate content related to current events, social justice, and job relevance. Between 2013 and 2023, the average course offered at universities in my sample became 0.044 sd more aligned to social justice and 0.031 sd more aligned to job relevance. Such change occurs primarily through the introduction of new classes, rather than the modification of existing courses.[5] For comparison, the 10-year change in social justice

[3] I relate each course to five themes: current events, job relevance, scholarship, social justice, and technology. I chose themes based on three main criteria: first, their general applicability across multiple fields; second, the availability of text data sources that could help identify relevant words/phrases; and third, their relevance to intellectual interests students might possess or cultivate during college.

[4] This second step addresses an important limitation of TF-IDF: its sensitivity to specific jargon. By projecting the document vector onto thematic weights, the method accounts for evolving language use, ensuring that similar meanings—expressed through different terms—are treated consistently. For example, "global warming" and "climate change" would appear orthogonal in a pure TF-IDF representation, but the projection ensures that courses using either term score similarly if both phrases occur with comparable frequency in a thematic corpus, like New York Times articles, and enabling comparisons across course descriptions specifically in relation to each theme.

[5] A potential limitation of this analysis is that course descriptions may be updated less frequently than the courses they describe. I address this concern by measuring changes over relatively long periods and

alignment over the period of a decade is approximately 6.5% of the difference in social justice alignment between the average Sociology course and the average Business course.

My findings suggest that the quantity of courses offered adjusts far less than one-for-one to student demand and the content of these courses remains quite stable once introduced. While these results shed light on how universities adjust their course offerings and highlight certain trade-offs inherent in these decisions, determining what constitutes an "optimal" adjustment is less straightforward. The optimal elasticity of course quantity is likely greater than zero so that students can acquire the human capital needed to thrive in an evolving economic, social, and technological landscape. However, it is almost certainly less than one, given constraints such as the inelasticity of instructor supply. Similarly, while some degree of course content evolution is undoubtedly beneficial for students, rebuilding each course from scratch every semester would clearly be excessive.

In the heterogeneity analysis, I examine how different types of institutions adjust to changing student demand. Understanding these differences offers valuable insights for policymakers and labor markets by identifying which universities are most responsive to external shifts. Research-intensive universities (R1) are more elastic in adjusting their course offerings, particularly by expanding supply in high-growth fields. Less research-intensive universities respond to growing demand primarily by increasing section offerings within existing courses, while non-research-intensive universities are the most responsive in reducing course offerings in fields with declining demand.

Heterogeneity analysis of intensive margin changes reveals R1 universities made the largest shifts in course content, where new courses increasingly emphasize current events and social justice. In contrast, courses at less research-intensive institutions closely resemble those offered a decade ago. Courses at these institutions generally have a stronger focus on job-related topics compared to similar courses at R1 schools. These patterns reflect both institutional priorities and shifting student preferences, underscoring the interaction between supply and demand in shaping university curricula.

This paper makes three central contributions. First, it introduces a novel dataset that provides detailed insight into higher education instruction through course-level enrollment, supply, and course description data. Beyond this paper, the dataset will facilitate new research of how colleges shape the provision of knowledge through their course offerings. Second, the project is unique for analyzing supply-side responses to changing student demand. The results complement a larger existing literature on how students adjust to changing re-

through an index that captures the central themes of a course rather than the introduction/elimination of specific topics. I discuss the advantages and limitations of the course descriptions, as well as possible bias from infrequent course description updating, in Section 5.

turns to college degrees by highlighting the potential influence of imperfect course supply adjustments and the adaptability of course content in shaping students' course preferences. Third, I document heterogeneity in both the margins along which universities of different types respond to student demand and the kinds of courses these universities offer.

The rest of the paper proceeds as follows. Section 2 summarizes research related to this project. Section 3 describes the unique course catalog dataset used for this project. Sections 4 and 5 document adjustments of course quantity and content, respectively, to student demand. Section describes how fields adjust course content with changing enrollment. Section 6 concludes.

2 Related Literature

This paper contributes to a small but growing literature on higher education supply and a robust literature on factors influencing student demand in college. Previous work in higher education supply demonstrates how costs of instruction and instruction technology influence the supply of instructors across fields (Courant and Turner 2017, Hemelt et al. 2021). Existing research related to margins of course supply response includes work on rationing of courses in high-demand fields (Bleemer and Mehta 2022, Bleemer and Mehta 2021, Mumford et al. 2024) and grade inflation (Ahn et al. 2019, Denning et al. 2022).[6] Closest to this project is Thomas (2024), who models university preferences using instructor allocation and enrollment in introductory-level courses for a sample university. His work demonstrates the influence of course supply on students' enrollment decisions and considers the welfare tradeoffs of expanding sections in high-demand fields with the higher cost of instruction in these fields. My project extends this literature in two main ways. First, I use new data that provide insight into both course supply and content. Second, I leverage the diversity of schools in my dataset, which allows me to explore the heterogeneity in institutional responsiveness.

This paper builds on a much larger literature on factors influencing student decision-making, such as major choice, in college.[7] The work in this area most relevant to my paper studies students' responsiveness to changing conditions in the labor market, where evidence is mixed. Some have found that students respond inelastically to changing wages (Beffy et al. 2012, Wiswall and Zafar 2015, Long et al. 2015), while other work documents larger responses in terms of completed majors to occupation-specific shocks (Freeman 1976, Acton 2021, Weinstein 2020) or changing macroeconomic conditions (Blom et al. 2021). Contemporaneous work by Conzelmann et al. (2023) estimates relatively elastic responses of students'

[6] A related literature examines the supply of for-profit colleges, though it primarily focuses on institutional supply and enrollment rather than field-level adjustments (Deming et al. 2012, Gilpin et al. 2015).

[7] See Altonji et al. (2016) and Patnaik et al. (2021) for recent reviews of this literature.

enrollment across fields to changing labor market conditions, measured using job vacancies, and employs an estimation strategy similar to the one used here. My results align with their findings, but by focusing on course-level supply rather than degree completions, this paper uncovers additional margins of institutional adjustment that complement their analysis. Specifically, the asymmetric responses in course quantity to growing versus shrinking demand reveal supply-side constraints that are not apparent in aggregate degree-level studies. This finer granularity highlights how institutions dynamically adapt their curricula to shifts in student interest and demand, which may ultimately shape broader educational outcomes.

Finally, this paper contributes to a growing literature in Economics using text data (Gentzkow et al. 2019), including research that applies text methods in the economics of education (Eggenberger et al. 2018, Biasi and Ma 2022, Chau et al. 2023). Using a dataset containing a large cross-sectional sample of course syllabi, Biasi and Ma (2022) and Chau et al. (2023) use novel text analysis techniques to develop measures of course content in relation to the research frontier and skill demand, respectively. Biasi and Ma (2022) document disparities in access to frontier knowledge across institutions of varying selectivity, highlighting how institutional characteristics shape students' exposure to cutting-edge content. My results similarly reveal differences in curricular emphasis by selectivity, but with the added benefit of observing the entire set of courses within specific institutions and fields. This rich dataset enables me to disentangle broader trends in course supply and institutional responses to changing student demand, while effectively controlling for institution-specific and field-specific effects. Moreover, by focusing on course-level adjustments rather than solely the alignment with frontier knowledge, this paper considers curriculum responses to a wider range of social, economic, and technological forces that may influence course demand.

3 Data

3.1 Course catalog dataset

To analyze how higher education institutions adjust course offerings in response to changing student demand, I constructed a unique "course catalog" dataset with detailed course-level information from a sample of U.S. colleges and universities. The dataset includes 29 million course sections offered since 1998 from 782 institutions, covering 52% of baccalaureate enrollment. I collected the data by scraping universities' online course catalogs and schedules, recording details of each course offered during a specific term.[8]

For a given course, I may observe the name(s) of the instructor(s), the number of sections

[8] Appendix A summarizes the inclusion criteria for institutions in the course catalog sample and exercises to validate the data.

offered in a year, enrollment in each section, the format of instruction (whether it is in-person or online), and a brief text description of the course content. Figure 1 demonstrates the information contained in a typical observation in my dataset.

The dataset broadly reflects characteristics of the population of US universities. While the sample is not truly random, as it only includes universities with online course catalogs, the sample aligns with the broader population in several important respects. Table 1 benchmarks the characteristics of schools in the catalog sample against the characteristics of the US higher education system. While the sample aligns closely with the average US four-year institution in aspects like selectivity, cost, and resources, it does skew towards larger, public institutions. Extremely small private (often religiously affiliated) institutions are under-represented in this sample.

The course catalog dataset offers unique advantages. By capturing course-level enrollment and content details, it can measure shifts in student demand years before they manifest in completed majors. Because a major represents only a portion of a student's coursework, course-level data may also provide a more comprehensive snapshot of skill acquisition, reflecting the breadth of knowledge students engage with beyond their primary field of study. Moreover, this dataset provides insight into different margins of course supply — for example, whether a university expands enrollment in a program by expanding existing courses or by creating new courses or eliminating existing courses. These data provide a uniquely powerful lens for examining short-run supply adjustments that would not be detectable in completed major counts.

I impose a series of restrictions to refine the raw course catalog dataset for analysis.[9] I exclude non-classroom-based courses (e.g., independent study, internships), restrict the sample to undergraduate courses, and differentiate between lower- and upper-level courses based on institutional numbering conventions. I further limit the data to complete academic years, excluding summer terms. To standardize classification, I categorize over 28,000 department names into 54 standardized fields (e.g., History, Education, Economics, Engineering).[10] Throughout the analysis, I weight enrollment and course offerings by credit hours.

3.2 Supplemental data sources

I supplement the course data with institution characteristics from the National Center for Education Statistics' Integrated Postsecondary Education Data System (IPEDS). For the IV analysis, I use employment data from IPUMS using the 2009-2018 ACS 1% samples (Ruggles et al. 2023).

[9] See Appendix A for details on data processing.
[10] See Appendix B for details on field standardization.

In Section 5, I document changing course content in relation to student demand using a weighting system to gauge the significance of specific words or phrases based on their frequency in theme-specific corpora relative to a neutral corpus. I summarize these text data briefly below; further details are available in Appendix C.

New York Times articles: using the New York Times Developer API, I downloaded 938 thousand articles published between 2000-2022, capturing headlines and abstracts or text snippets.

Academic journal abstracts: following Biasi and Ma (2022), I compiled 155 thousand abstracts from 180 top-ranked academic journals (by H-index) from 2000-2022, sourced from Elsevier's SCOPUS.

Patents: the patent corpus includes the text of 2.5 million patents from the U.S. Patent and Trademark Office, covering 2000-2020.

Job descriptions: sourced from Lightcast (formerly Burning Glass Technologies), this corpus includes 2 million job descriptions from a sample of months 2010-2018, filtered for jobs requiring a college degree.

Social justice writings: this corpus includes texts from the "Issues" and "Policy Positions" pages of organizations spanning a range of social justice topics: the ACLU, the American Association of Disabled People, Amnesty International, the Brennan Center, the Democratic Socialists of America, GLSEN, the NAACP, the National Organization of Women, Oxfam, Planned Parenthood, the Southern Poverty Law Center, the Sunrise Movement, and UNICEF. In addition, the corpus includes the full text of six books listed in the top 25 books on activism and social justice, ranked by Goodreads.

Wikipedia articles: this corpus includes the full text of all English-language pages published on Wikipedia as of July 1, 2023, using the "Wikimedia dump service," totaling 3.8 million documents.

4 Extensive Margin: How universities adjust course quantity

In this section, I document how universities adjust the number of courses offered across fields of study to meet students' changing demand for those fields. I first document this result descriptively, then estimate elasticities using an instrumental variables strategy that accounts for potential endogeneity of enrollment as a measure of demand.

The focus of this section is estimating long-run elasticities of upper-level courses. Since universities typically plan over multi-year cycles, it may be impractical to expect short-term adjustments to changing enrollment. Moreover, enrollment can be noisy, and small fluctuations might not necessarily represent genuine changes in demand. In my preferred specification, I estimate course quantity elasticities over 8-year periods. The estimates presented in this section are course quantity elasticities for upper-level courses.[11] I provide estimates of course quantity elasticity for all courses in the Appendix. These estimates are substantively similar but less precise.

4.1 Trends in course quantity and enrollment

When demand for a field of study grows, an institution can respond through four strategies: increasing the number of courses offered, adding sections to existing courses,[12] expanding the capacity of existing sections, or choosing not to react and restricting enrollment.[13] The strategies vary in cost and the extent to which they benefit students. Creating new courses involves significant fixed costs but can accommodate the broadest student base. Adding sections requires marginal costs and helps accommodate students who would otherwise be rationed out. Expanding section capacity is typically the least costly option but may strain faculty workloads and affect instructional quality. The optimal response depends on the nature of the demand shock and the institution's priorities, balancing immediate student needs with broader objectives.

To empirically assess how universities navigate these complex trade-offs in response to changing student demand, Figure 2 plots enrollment and course quantity trends across various fields of study. The figure plots the growth trends in course enrollment, course quantity, and section quantity, aggregated into six field categories: Business/Economics, Education, Humanities, Social Science, STEM (excluding Computer Science), and Computer Science.[14]

[11] Conventionally numbered in the 300-400 range, typically elective courses. I impose this restriction for two reasons. First, these are the courses over which students have the most autonomy in their selection. As a result, fluctuations in enrollment for these courses should more accurately reflect students' changing demand rather than responses to, for example, a university's changing core requirements. Second, by the time students are enrolling in upper-level courses, they have acquired information about their aptitude for a given field. Any limitations on their ability to enroll in the student's preferred courses, therefore, may divert a student from the courses for which they are most suited.

[12] A course refers to an individual class, typically identified by a unique course ID (e.g., Econ 101 or Econ 102), while a section is a specific offering of that course. For instance, if an institution offers two sections each of Econ 101 and Econ 102 in both the Fall and Spring semesters, the total would be 8 sections for 2 courses in Economics.

[13] Similarly, in response to declining demand, institutions may reduce the number of courses, scale back sections, decrease section capacity, or make no adjustments.

[14] Skilled trades, professional degree-granting fields, and interdisciplinary departments are excluded. Additional detail on field classification and selection is available in Appendix B.

The figure highlights a divergence between course quantity and enrollment growth across fields in instances where enrollment is decreasing or rapidly increasing. The figure illustrates a shift in enrollment from Humanities and Education towards fields like Business/Economics, STEM, and Computer Science. In high-growth fields like Computer Science, course quantity increased modestly but lagged behind explosive enrollment growth. For fields with declining enrollment, including Education and the Humanities, course quantity remained relatively stable despite sharp enrollment declines. In contrast, fields with modest enrollment growth, such as non-Computer Science STEM and Business/Economics, saw comparable growth in enrollment and course quantity. [15]

This asymmetry points to potential rigidities in two directions. First, downward rigidities make it difficult for institutions to reduce course offerings in response to declining demand. Structural factors, such as commitments to offering foundational skills and the constraints imposed by tenure, contribute to the stability of offerings even as enrollment declines. Commitments to instructors on long-term contracts, such as tenured instructors, regardless of a field's popularity make the cost of offering courses in less popular fields relatively low. Second, the limited responsiveness of course quantity to the rapid growth in Computer Science enrollment suggests rigidities triggered by surges in demand that outpace institutional capacity for adjustment.

Although enrollment and course quantity trends align more closely in non-Computer Science STEM and Business/Economics, this correspondence does not necessarily indicate highly elastic adjustments. Enrollment alone is an imperfect proxy for demand, as it does not capture unmet demand from students unable to enroll in preferred courses due to rationing or insufficient offerings. I address this limitation in the following section.

4.2 Empirical Strategy

4.2.1 OLS Specification

In this section, I estimate how the quantity of courses responds to changes in students' demand across different fields of study. Equation 1 shows the OLS specification I use to estimate course quantity elasticity:

[15] Computer Science is a unique field for its boom-and-bust cycles. The growth in Computer Science enrollment during the period of my analysis follows a nadir in Computer Science enrollment following the Dot-Com bubble. It is possible that some institutions had surplus capacity in Computer Science to absorb the enrollment surge, attenuating some of the immediate need to grow course quantity for the 2010s enrollment wave.

$$\Delta y_{i,s,t'} = \alpha \Delta x_{i,t}^{(\text{avg})} + \beta \Delta x_{i,s,t}^{(\text{field})} + \epsilon_{i,s,t} \tag{1}$$

$$\Delta x_{i,s,t}^{(\text{field})} = \Delta x_{i,s,t} - \Delta x_{i,t}^{(\text{avg})} \tag{2}$$

The dependent variable, $\Delta y_{i,s,t'}$, denotes the percentage change in the quantity of courses offered by institution i in field s over period t'. I calculate this change as a long log difference in the credit-weighted number of courses offered across these years.[16] The log difference specification differences out any fixed institutional characteristics. Thus, any controls I introduce should pertain to time-varying attributes of universities. To this end, I control for the university's average enrollment growth rate $\Delta x_{i,t}^{(\text{avg})}$, ensuring that the analysis accounts for shifts in course quantity tied to broader university-level changes.

After controlling for the influence of overall enrollment growth on course quantity, the parameter of interest, β, represents the elasticity of course quantity to relative shifts in enrollment across fields. For clarity, the field-specific enrollment growth rate is adjusted by subtracting the institution's average enrollment growth rate, resulting in $\Delta x_{i,s,t}^{(\text{field})}$.[17]

In the data, course quantity and course enrollment are directly linked. To disentangle changes in enrollment that are demand-driven from changes that are entirely due to changes in the courses offered, I impose a one-year offset o between the period over which I measure changes in enrollment (t in Equation 1) and the period over which I measure changes in course quantity ($t' = t + o$ in Equation 1). To be concrete, if analyzing enrollment and course quantity changes over an 8-year window, I would regress the change in course quantity from 2010-11 to 2018-19 on enrollment changes from 2009-10 to 2017-18. Appendix Table A-7 estimates course quantity elasticities under alternative offset and lag windows - neither

[16] I credit-weight both changes in course quantity and changes in enrollment.

[17] De-meaning is important when testing whether course quantity elasticity differs for fields growing or shrinking relative to the university. Measuring course supply responses to relative, rather than absolute, enrollment changes isolates how universities adjust to field-specific demand shifts. For example, consider two universities where a field grows by 5%: if overall enrollment at the first university increases by 25%, that field is growing more slowly than the institution as a whole and may not warrant expansion beyond the expansion necessary to support the growing university overall. In contrast, if overall enrollment shrinks by 25%, the field is a relative outlier, making it a more natural target for investment.

However, differencing out the institution's overall growth rate may introduce some bias to the extent that overall institutional growth could itself be responsive to field-specific demand shifts. If, for example, demand for seats at technical universities increases in tandem with demand for STEM fields at those institutions, de-meaning could obscure the effect of broader trends on course supply. Given that most institutions in my sample are regional public universities, where enrollment is less likely to be driven by specialized field-specific demand shifts, the potential for this bias is limited. Nevertheless, I test an alternative specification, which I summarize in Appendix Tables A-8 and A-9, estimating course supply elasticity using absolute enrollment changes instead of de-meaned changes. These estimates are substantively similar and, if anything, suggest that universities are less elastic than the estimates presented in this section.

choice affects the results substantively.

4.2.2 IV Specification

Particularly in cases where universities choose not to accommodate students' changing demand, we might be concerned that enrollment changes are driven or constrained by course supply rather than student demand. To illustrate this concern, consider a scenario where a university's Economics Department experiences a sudden surge in demand for its courses. In response, the university does increase its course quantity but only enough to accommodate a fraction of the new demand. For example, the university might experience a demand increase equivalent to 200 new students but expand course quantity to accommodate only 100 of them. In this case, quantity is highly responsive to changing enrollment but the university only addresses half of the new demand for Economics courses.[18] These forms of non-response will bias my OLS estimates of course quantity elasticity towards making the university appear more responsive to changing demand than it actually is.

To estimate a causal relationship between changes in student demand and changes in course quantity, I use a shift-share instrumental variables (IV) strategy that identifies a portion of enrollment changes solely attributable to shifting student preferences, independent of actions taken by the university. The instrument uses two sources of variation: variation in employment growth prospects across fields (s) and differential exposure to changing employment growth prospects in different parts of the country (based on the Census Division r in which school i is located, which I hereafter refer to as a "region").[19]

Because labor market growth is driven by macroeconomic forces rather than university actions, employment shifts serve as a plausibly exogenous shock to student preferences. The instrument captures differential employment growth across fields net of region-wide labor market shocks. I estimate course quantity elasticity during the period 2009-10 to 2018-19, the period between the Great Recession and the Covid-19 pandemic. I estimate this elasticity with enrollment and course quantity changes measured over a single 8-year window with a

[18] Further, to the extent that these students are diverted to courses in other fields, the university may appear responsive to inflated demand for other fields when, in reality, students are taking classes they would ideally prefer not to take.

[19] Contemporaneous work by Conzelmann et al. (2023) uses a similar instrument to study how students and universities respond to changing demand for college graduates in the labor market. Our analyses differ in the sense that they study the direct effect of changing job demand on completed majors and course quantity. In contrast, my analysis focuses primarily on how changing labor market conditions impact course quantity through their effects on students' demand. Conzelmann et al. use job postings data to measure changing demand in local labor markets for students from different majors, then measure the exposure of each institution in their sample to these changes using shares of graduates from the institution in each labor market (using data from LinkedIn). Using their data, I confirm that on average, more than 80% of the graduates from the schools in my sample work in the same Census division where their respective institutions are located (Conzelmann et al. 2022).

one-year offset. I construct the instrument using data from the 2009 and 2017 American Community Surveys (ACS),[20] following Equation 3 below:

$$\Delta E_{s,r,t} = \sum_{j=1}^{J} \phi_{s,j,r,t_0} \left(lnE_{j,r,t_1} - lnE_{j,r,t_0} \right) \tag{3}$$

$$z_{s,r,t} = \Delta E_{s,r,t} - \overline{\Delta E_{r,t}} \tag{4}$$

The instrument fixes ACS respondents' college major (s) to occupation (j) shares (ϕ_{s,j,r,t_0}) in 2009 (t_0), then projects the change in log employment ($\Delta E_{s,r,t}$) as the average employment growth rate of college graduates in each occupation (4-digit OCC) between 2009 and 2017 (t_1). I weight the log employment growth by the fixed major-to-occupation shares.

To get the instrument $z_{s,r,t}$, I subtract from $\Delta E_{s,r,t}$ the regional average employment growth rate for college graduates, $\overline{\Delta E_{r,t}}$. The resulting instrument isolates differential employment growth across fields relative to the regional average. Because each university is small relative to its Census division — comprising 3–8 states and tens of millions of people — its direct influence on regional employment trends should be minimal. To further ensure that changing course quantity does not affect labor market measurements, I restrict the instrument to workers aged 30–65, who generally completed their education before the baseline year. The instrument values range from -0.208 to 0.188 across field-regions, with larger values indicating relatively stronger job prospects in the field.[21]

To illustrate how the instrument captures field-by-region differences in changing employment growth prospects, consider the construction of the instrument for Computer Science and Education at a single university located in the South Atlantic Division. In the 2010 ACS, approximately half of workers in the South Atlantic region with Computer Science degrees worked as programmers/engineers, nearly 20% worked in technical administrative roles, 10% worked in sales, and the remaining 20% worked in other occupations.From 2010 to 2018, employment in these occupations grew at a weighted rate of 30.9%, which was 8.1 percentage points faster than the regional average. Similarly, 70% of workers in the South

[20] With the one-year offset, I am measuring enrollment changes from 2009-10 to 2017-18. Because the ACS is collected on a calendar year cycle, the collection period bisects the academic year. I align the instrument to the end of the academic year; results are unchanged if I align the instrument to the start of the academic year.

[21] The period of my analysis contains three important trends in labor market conditions that drive much of the variation in projected employment growth across fields. First, innovation in mobile technology and growing use of data fueled growth in technology jobs. Second, stagnant earnings and declining job satisfaction contributed to declining interest in the teaching profession (e.g., Kraft et al. 2020). Third, the passage of the Affordable Care Act in 2010 created new demand in healthcare. Fields and regions differ in their exposure to these changes, which creates variation for my estimation.

Atlantic region with Education degrees worked in education or education administration and the remaining 30% worked in other occupations. Employment in these jobs grew 7.9 percentage points less than the regional average during the same period. The instrument takes values 0.081 and -0.079 for Computer Science and Education in the South Atlantic Division, respectively.

To illustrate how the instrument uses variation across regions, consider a single field, Computer Science, offered at two different universities: one located at the same institution as in the preceding example and one located in the Pacific Division. In 2010, Computer Science graduates flowed into similar jobs in similar proportions in the two regions. However, relative employment growth in these jobs was much faster in the Pacific Division (13.7 percentage points faster than the regional average) compared to the South Atlantic Division (8.1 percentage points). The instrument takes values of 0.137 and 0.081, respectively, reflecting the stronger labor market-driven push into Computer Science courses in the Pacific Division compared to the South Atlantic Division.

I estimate the IV model using two-stage least squares. In the first stage, I estimate the relationship between the de-meaned percent change in enrollment ($\Delta x_{i,s,r}^{(\text{field})}$) in field s at college i from 2009-10 to 2017-18 and the relative employment growth ($z_{s,r}$) of occupations typical for graduates of major s in region r:

$$\Delta x_{i,s,r,t}^{(\text{field})} = \phi \Delta x_{i,r,t}^{(\text{avg})} + \kappa z_{s,r,t} + \eta_{i,s,r,t} \tag{5}$$

In the second stage, I use the first stage's predicted values, denoted as $\widehat{\Delta x_{i,s,r,t}^{(\text{field})}}$, to instrument for students' changing demand. I then estimate a regression of the percentage change in the number of courses in field s at college i between 2010-11 and 2018-19 ($\Delta y_{i,s,r,t}$) on this instrumented enrollment change:

$$\Delta y_{i,s,r,t'} = \alpha \Delta x_{i,r,t}^{(\text{avg})} + \beta \widehat{\Delta x_{i,s,r,t}^{(\text{field})}} + \epsilon_{i,s,r,t} \tag{6}$$

The second stage regression produces an estimate of the causal effect of changes in demand on changes in course quantity.

The IV strategy allows me to estimate demand when enrollment is rationed (due to capacity constraints) or inflated (due to institutional steering) by first estimating the relationship between employment and enrollment growth in cases where enrollment accurately reflects demand, then using this relationship to predict demand where enrollment is constrained or

inflated. If rationing or inflation were widespread, this would bias the first-stage estimates and limit the IV's usefulness. However, as long as many fields experience demand-driven enrollment changes without severe supply constraints, the estimated relationship between employment growth and enrollment remains informative, allowing me to reasonably project demand shifts in constrained cases.

While some universities may be unable to fully accommodate rising demand, supply constraints should attenuate rather than bias the first-stage relationship unless they are systematically correlated with labor market trends. To the extent that supply constraints prevent the full realization of demand shifts, my IV estimates may still overstate university responsiveness to student demand changes, though the overall bias should be limited if rationed or inflated cases are a modest share of the sample.

For identification, the instrument must satisfy assumptions of monotonicity, independence, relevance, and the exclusion restriction. Monotonicity requires that when labor market conditions improve for a field, no school's enrollment demand in that field decreases as a direct result. In other words, every institution's demand for the field is weakly pushed up when the instrument increases. Monotonicity is inherently untestable in this setting because enrollment is an equilibrium outcome affected by both demand and supply constraints.[22] However, a violation would require a counterintuitive pattern: that improved job prospects in a field systematically reduce demand for that field. Such a violation would run contrary to theory on students' motivation for investing in human capital or require that universities actively reduce course offerings in high-demand fields.

Independence requires that employment growth be uncorrelated with unobserved factors influencing course quantity. This assumption holds if variation in employment growth across fields is driven by broader labor market forces rather than endogenous university decisions. Labor market conditions, particularly at the scale of a Census division, are plausibly external to individual universities' decision-making. To further mitigate concerns about endogeneity, I select a period of analysis — spanning from the Great Recession to the start of the Covid-19 pandemic — that represents a distinct labor market phase, reducing the likelihood that universities had advance foresight into employment trends during this window. Additionally, I construct the instrument using employment shifts and occupational shares for workers above age 30, ensuring that the observed labor market changes reflect pre-existing trends in workforce composition rather than responses to recent course supply decisions.

I demonstrate first stage relevance through a strong first stage, with results summarized

[22] Appendix Figure A-6 provides suggestive evidence that the relationship between employment growth and enrollment is overwhelmingly positive, consistent with the monotonicity assumption. While some variation exists, there is little evidence of systematic defiers.

in Appendix Table A-2. The cluster-robust first-stage F-statistic of 129, reported in Table 2, confirms the strength of the instrument.[23]

The exclusion restriction requires that changes in labor market opportunities affect course quantity solely through their impact on student demand. Admittedly, this is a big assumption, particularly for fields experiencing growing demand. In Appendix D, I describe in detail three potential violations of exclusion restriction and provide evidence of the importance of these potential violations. The first concern is that universities are better able to forecast labor market trends and adjust supply in anticipation of future demand. To test for this, I show that there is no evidence that universities are expanding course capacity in two of the highest-growth fields — Computer Science and Engineering — before enrollment growth is realized (Appendix Figure A-8). A second potential concern is that labor market changes could also affect faculty hiring, indirectly influencing course quantity. I show that excluding outlier field-regions with extreme employment shifts does not materially change the estimates, suggesting that this channel is not driving the results. A third concern would arise if labor market conditions influence local firms or donors to invest in expanding fields linked to future employment growth. I argue that such influence is far less common at the institutions and fields of focus in this paper than at two-year institutions and in trade/professional fields, which I exclude from the analysis.[24]

For the IV, I cluster standard errors at both the institution and field-by-Census division level to address the potential serial correlation within a field-region. Because my focus is on estimating how schools adjust, on average, to changing student enrollment, I assign equal weight to each school in the regressions. Within each school, I assign weight to the field-level observations in proportion to the field-level enrollment in the base year. This means I give

[23] As a validation exercise, I also estimate the first stage regression using completed majors, reported in IPEDS data, as the measure of changing enrollment. Completed majors are essentially as responsive to changing occupation growth as enrollment in upper-level courses. This may suggest that changing conditions in the labor market push marginal students to complete a major with improving employment growth prospects, but these students still take elective classes in fields with poorer employment growth prospects.

[24] Recent work in the shift-share literature formalizes the identification assumptions underlying shift-share instruments (e.g., Goldsmith-Pinkham et al. (2020), Borusyak et al. (2022)). Goldsmith-Pinkham et al. (2020) demonstrates that the Bartik instrument is analogous to using shares as instruments, with the exogenous growth rates primarily determining the instrument's relevance. In this paper, the "shares" are major-to-occupation shares rather than employment shares.

Identification could be compromised if these shares correlate with external factors affecting both student demand and course quantity. Three design features mitigate concerns of endogeneity: first, the shares are anchored to a base period, ensuring independence from contemporaneous labor market shifts; second, the instrument uses major-to-occupation shares and employment growth rates of workers aged 30-65, which exclude recent graduates who might be affected by course offerings during this period; and third, regions are defined at the Census division level, where any individual university's graduates represent a small share of the region's labor force. In Appendix D, I summarize a robustness test that uses major-to-occupation shares from all Census divisions *except* the one where the university is located. The results remain largely consistent with the main findings.

more weight within the institution to fields with greater enrollment to improve precision.

4.3 Results

Tables 2 and 3 summarize OLS and IV estimates for the course quantity elasticity regressions for the number of courses and number of sections, respectively. Columns 1-7 summarize OLS estimates from the elasticity estimation described in Equation 1. I estimate course quantity elasticities over periods ranging from 2 to 8 years. In each pair of columns, I estimate course quantity elasticity over rolling periods (e.g., 2004-05 to 2006-07; 2005-07 to 2007-08) and staggered periods (e.g., 2004-05 to 2006-07; 2007-08 to 2009-10). The first row of the table summarizes course quantity changes in response to the overall growth/decline in enrollment at the university and the second row, which contains the estimates of interest, summarizes course quantity changes in response to a field's growth/decline in enrollment relative to the university overall.

The OLS estimates demonstrate that course quantity responds less than one-for-one to changing enrollment in both the short and long run. The two-year course quantity elasticity is approximately 0.21; course quantity becomes slightly more elastic (0.40) when measured using eight-year lags. The gradual increase in course quantity elasticity suggests that universities are more responsive to sustained enrollment trends than to episodic changes in enrollment. The quantity of sections is more elastic than the quantity of courses, suggesting that universities accommodate student demand more by modifying the frequency of offerings in existing courses rather than the outright creation/elimination of courses.[25]

Having argued in Section 4.2.2 that enrollment changes may insufficiently reflect changes in student demand, I summarize my IV estimates in Column 8 of Tables 2 and 3.[26] I estimate these elasticities using a single difference on the period from 2009-10 to 2018-19, following the Great Recession and ending in the last full academic year prior to the Covid-19 pandemic. For comparison, Column 7 in each table provides the OLS estimates corresponding to quantity and enrollment changes during this same period.

The IV estimates suggest that fields expand course quantity 2.2% for a 10% increase in demand. To illustrate, a department the size of Stanford's Economics Department would add a new course if the enrollment in upper-level courses rises by 131 seats. This enrollment

[25]For example, a department may begin offering a popular course in both the Fall and Spring semesters, instead of just in the Fall, or increase the number of sections available in a single semester. This type of adjustment is more common at larger universities, which may offer many elective courses varying in size and add sections in the most popular courses, and at less-selective universities, which often offer fewer elective courses each with lower enrollment caps.

[26]I present the reduced form estimates in Appendix Table A-4. In the reduced form, a 10% increase in relative employment growth for a given field within a university's region is associated with a 9.9% increase in course quantity.

increase corresponds to an underlying demand increase for upper-level Economics classes of 249 seats.[27] Columns 7 and 8 in Table 3 summarize estimates of the elasticity of the number of sections offered on changing enrollment. Although more elastic than courses, section quantity also adjusts less than one-for-one. Fields expand section quantity 5.5% with a 10% increase in demand.

Comparing the values in Columns 7 and 8, the OLS estimates are biased higher than the IV estimates. Because enrollment is an equilibrium outcome, changes in enrollment may reflect both students' changing demand for courses and course supply decisions unrelated to student demand. I cannot observe, for example, demand from students who are rationed out of courses they would prefer to take. Without accounting for this unmet demand, course quantity responses will appear to align better with students' changing demand. Similarly, the university may introduce policies like distribution requirements that boost enrollment in courses that students otherwise might prefer not to take. Such policies would attenuate enrollment shifts from declining fields to growing fields. Considering these issues, we might expect the bias in the OLS estimates in the direction of greater course quantity elasticity relative to the IV estimates.

4.4 Asymmetry in course quantity elasticity

The model in Equation 1 imposes that course quantity response to increasing enrollment in a field is exactly the opposite of its response to a comparable decrease. However, the practical costs of growing versus shrinking a field can differ. Specifically, considering that many faculty are employed on long-term contracts, the university might incur little to no marginal cost in allowing faculty in a field experiencing declining enrollment to teach their courses. Furthermore, descriptive evidence from Figure 2 suggests potential asymmetry in course quantity responses to enrollment changes.

Thus, I consider a more flexible model that allows course quantity elasticity to differ based on whether a field is growing slower or faster than the institution average ($\Delta x_i^{(\text{avg})}$).[28] I estimate the new model:

$$\Delta y_{i,s} = \alpha \Delta x_i^{(\text{avg})} + \beta_1 \Delta x_{i,s}^{(\text{field})} \mathbb{I}(\Delta x_{i,s}^{(\text{field})} < 0) + \beta_2 \Delta x_{i,s}^{(\text{field})} \mathbb{I}(\Delta x_{i,s}^{(\text{field})} > 0) + \epsilon_{i,s} \qquad (7)$$

[27] In 2018-19, Stanford's Economics department offered 125 credits of upper-level courses and student enrollment totaled 6747 credit hours. A one-course increase in courses supplied would be equivalent to a 4% increase, which, according to the estimates in Table 2, is the result of an 10% increase in enrollment (675 student-credit hours) or a 18.4% increase in demand (1244 student-credit hours). Dividing by 5 credit hours per course gives the values cited above. Some of the increased enrollment derives mechanically from enrollment in the new course.

[28] For simplicity, I omit time subscripts, which I use identically to the base model in Equation 1.

where the parameters of interest, β_1 and β_2, represent the course quantity elasticities when enrollment is growing slower or faster than the institution average, respectively.[29]

I augment the IV model in Equation 6 to allow for asymmetry in course quantity responses to enrollment changes above and below the institution average. In the extended first stage, I specify a model that allows the coefficients linking changes in employment to changes in enrollment to vary for fields experiencing relatively growing versus shrinking employment. In the second stage, I incorporate this asymmetry by using the predicted values from the first stage as in the following equation:

$$\widehat{\Delta x_{i,s}^{(\text{field})}} = \Delta x_{i,r}^{(\text{avg})} + \kappa_1 z_{s,r} \mathbb{I}(z_{s,r} < 0) + \kappa_2 z_{s,r} \mathbb{I}(z_{s,r} > 0) + \xi_{i,s,r} \tag{8}$$

$$\Delta y_{i,s} = \alpha \Delta x_i^{(\text{avg})} + \beta_1 \widehat{\Delta x_{i,s}^{(\text{field})}} \mathbb{I}(\widehat{\Delta x_{i,s}^{(\text{field})}} < 0) + \beta_2 \widehat{\Delta x_{i,s}^{(\text{field})}} \mathbb{I}(\widehat{\Delta x_{i,s}^{(\text{field})}} > 0) + \epsilon_{i,s} \tag{9}$$

Tables 4 and 5 present OLS and IV estimates of course and section quantity elasticities from specifications that distinguish growing and shrinking fields. The first row reports the elasticity of course quantity with respect to overall enrollment changes at the institution; subsequent rows report elasticities separately for fields growing faster or slower than the institution's average enrollment growth. Columns 1–7 contain OLS estimates, and Column 8 reports IV estimates cluster-robust first-stage F-statistic. The table also includes p-values testing the equality of elasticities for growing and shrinking fields.

Overall, course quantity expands more in fields experiencing rising demand than in those facing declining demand. Consistent with results in the previous section, course quantity elasticity increases when measured over longer time horizons. Over time, the gap between course quantity elasticities for growing and shrinking fields widens, and the estimates become statistically distinguishable. This pattern suggests that short-term constraints on course supply — such as pre-commitments to instructor headcount and advance scheduling — bind asymmetrically over the long run, affecting fields with shrinking enrollment more rigidly than growing ones. Interestingly, while course quantity elasticity diverges over time, differences in section quantity elasticity between growing and shrinking fields are less pronounced and frequently not statistically distinguishable.

Column 8 summarizes IV estimates of asymmetric course quantity elasticity.[30] The IV estimates suggest an even more dramatic gap in course quantity responses to fields with

[29] I consider an alternative specification in Appendix Tables **??** and **??** that estimates elasticities based on whether the field is growing or declining in absolute terms, rather than relative to the institution average. The results are substantively similar.

[30] Due to the nonlinear transformation of the first stage estimates, I calculate bootstrapped standard errors for the IV using 1,000 repetitions of the estimation, resampling region-by-field clusters in each iteration.

growing vs shrinking demand. The IV estimates in Table 4 indicate that course quantity increases by 3.7% for a field growing 10% faster than the institution's overall rate, while it decreases by 0.8% when the growth is 10% slower. Similarly, the IV results in Table 5 suggest that section quantity rises by 6.0% when a field's growth surpasses the institution's rate by 10%, but drops by 5.1% when it lags behind by the same measure.[31] As with the linear model, the OLS estimates appear larger than the IV estimates. However, the extent of this bias is much larger in the case of declining demand relative to growing demand. This result suggests that the shifting of enrollment into less-preferred fields may prop up fields that otherwise would not have sufficient organic demand to support their full course offerings.[32]

4.5 Heterogeneity

Universities of different types face different incentives and constraints when responding to changes in student demand. To explore these differences, I partition my sample of four-year universities by their 2010-11 Carnegie classification into R1 universities (very high research activity), R2 universities (high research activity), liberal arts colleges, and teaching-focused institutions (all other four-year institutions).[33] I modify the baseline regression equations to include interaction terms between the relative change in course quantity ($\Delta x_{i,s,t}^{(\text{field})}$) and dummy variables for the university's Carnegie classification.[34]

Figure 3 shows that R1 universities are more elastic in adjusting course quantity in response to changing student demand compared to most other university types. This higher elasticity is primarily driven by their significantly greater responsiveness to growing enrollment in high-growth fields. Conversely, when enrollment declines, R1 universities are about as responsive as R2 and liberal arts colleges, while teaching-focused institutions tend to be more elastic, discontinuing courses in fields with shrinking enrollment.

The bottom panel of Figure 3 compares section-level elasticities. Here, liberal arts institutions are less responsive to changes in enrollment than R1 universities, both when enrollment

[31] In the reduced form (Appendix Table A-4), a 10% increase in employment for a given field is associated with a 20.5% increase in course quantity, while a 10% decrease in employment growth is associated with a 3.4% reduction in course quantity.

[32] This reallocation may occur through distribution requirements that require students to take courses in low-demand fields or may arise naturally when students are unable to enroll in courses in their preferred fields.

[33] The Carnegie Classification categorizes institutions based on research resources and activity levels. In 2010, 108 institutions were classified as R1 (very high research activity), including elite private universities, state flagships, and other doctoral-granting institutions. R2 institutions (high research activity) conduct less intensive research but still offer doctoral programs. Liberal arts colleges are selective, teaching-focused institutions that prioritize small class sizes. The remaining category includes a mix of teaching-oriented and less selective regional public and private universities.

[34] The analysis in this section uses 8-year course quantity elasticities with a one-year offset. I estimate the OLS version of the model to maximize the number of institutions and periods.

is growing and shrinking. In contrast, teaching-focused universities are more likely to reduce the number of sections offered in response to declining enrollment. While R1 universities exhibit greater elasticity in adding new courses for growing fields, they are similar to R2 and teaching-focused institutions in their responsiveness to section-level adjustments.

This heterogeneity in course quantity elasticity suggests that institutional objectives and resource constraints may drive how universities manage growing and shrinking demand. R1 universities typically respond to growing demand by creating new courses, which may reflect an emphasis on curriculum innovation and an ability to cater to well-prepared students who can benefit from emerging fields. These institutions are also likely motivated by the need to maintain competitive research and academic standings by offering a diverse range of cutting-edge courses.

On the other hand, R2 and teaching-focused institutions primarily add sections to existing courses in response to growing demand. This approach is less costly and serves students who would otherwise be shut out of oversubscribed courses. For these institutions, resource constraints likely limit the capacity to create entirely new courses, leading to a more efficient but less flexible response to changes in student preferences.

When facing declining demand, liberal arts colleges tend to maintain sections in shrinking fields, resulting in smaller class sizes but potentially less efficient resource allocation. This strategy may reflect a commitment to providing a broad-based education, even at the expense of larger courses in growing fields. In contrast, teaching-focused institutions are more likely to reduce course and section offerings in response to declining enrollment, suggesting a greater sensitivity to financial and operational constraints.

These findings highlight the trade-offs between innovation and efficiency in higher education. R1 universities' flexibility allows them to adapt to changing student demand by expanding course offerings, potentially providing students with greater exposure to emerging fields. However, this adaptability comes with higher costs, which less research-intensive universities mitigate by focusing on adding sections rather than new courses.

4.6 Discussion

Course quantity adjusts far less than one-for-one to shifts in student demand, particularly in fields experiencing declining enrollment. While the preceding analysis is descriptive, it underscores the trade-offs universities face in balancing responsiveness to demand with institutional constraints.

Limited flexibility in course supply affects both educational quality and student decision-making. When course inelasticity leads to non-supply, students may be deterred from pursuing their preferred fields. Estimates in Appendix E suggest that seat rationing, as implied

by the IV estimates, reduced the number of completed Computer Science and Engineering majors by 3.9% and 3.2%, respectively, between 2009–10 and 2018–19. Even when courses remain available, inelasticity can distort class sizes, leading to crowding in high-demand fields and sparse enrollment in others. While universities may expand capacity within existing courses, these adjustments have limits, particularly for upper-level courses designed for smaller cohorts.[35]

These rigidities can also contribute to curricular misalignment, particularly in fast-evolving fields where demand for new skills outpaces institutional adaptation. Students may graduate with training less aligned with current labor market needs,[36] or with limited exposure to emerging research areas. Additionally, if course expansion occurs without corresponding enrollment growth, per-student instructional costs rise. Given universities' reliance on tuition and public funding, these costs may ultimately be passed on to students through higher tuition.

5 Intensive Margin: How fields adjust course content

A field can respond to changing student demand by updating course content, such as replacing a course teaching outdated content with one that emphasizes high-demand content. In the preceding section, an institution's only response to students' changing demand was the creation or elimination of courses within a given field of study. Fields that modify their courses in response to changing demand could attenuate enrollment shifts and relieve pressure to reallocate resources from shrinking to growing fields. This section explores how college course content adapts to align with student demand.

5.1 Measuring course content through course descriptions

To measure the content of courses, I use the course description included with many course catalog entries (for example, see Figure 1). Course descriptions are short — typically fewer than 50 words — text summaries of course content that highlight topics covered in a class, skills students may develop, and the work students will produce. This rich insight into what students learn in their college classes is a unique feature of the course catalog dataset and enables me to identify what distinguishes or connects fields, compare fields across institutions,

[35] Large class sizes are associated with lower student evaluations (e.g., Bedard and Kuhn 2008, Monks and Schmidt 2011), though evidence on their effects on student performance is mixed (e.g., Kokkelenberg et al. 2008, Bandiera et al. 2010). Appendix Figure A-5 documents changes in average course sizes by field.

[36] Though there is a trade-off between preparing for immediate job market demands and developing durable skills that remain valuable long term (Deming and Noray 2020).

and track their evolution over time. Importantly, the longitudinal structure of my data facilitates comparisons within an institution and field over time.

Although course descriptions provide insight into an institution's educational offerings, they possess a few limitations. For example, instructors may not update these descriptions frequently. In such cases, the description might not reflect recent changes in course content.[37] To the extent that changes in course content are not contemporaneous with changes in course descriptions, the timing of any individual course description change may be unreliable. In my data, most of the changes to a field come through the introduction of new courses, before a course has an opportunity to diverge from the course description, and through the discontinuation of existing courses. I also study changes over a relatively long period of time to avoid reliance on changes in any individual year.

I analyze course descriptions by adapting techniques from Natural Language Processing (NLP). These methods represent the course description for each course c offered in field s at institution i in year t as a vector of words.[38] I then represent each document as a $W \times 1$ vector $v_{c,i,s,t}$ with length (W) equal to the size of the dictionary of unique tokens $(w \in W)$. Typically, tokens are single words. However, I treat common phrases as distinct single tokens. For example, I treat "climate change" as a single token distinct from "climate" or "change"; likewise, "social media" is distinct from "social" or "media."

The values in $v_{c,i,s,t}$ are assigned according to their Term Frequency-Inverse Document Frequency (TF-IDF) weight, which is a measure of the distinctiveness of a given token to a given document. TF-IDF is the product of the Term Frequency (TF), a given token's share of all tokens in a document, and Inverse Document Frequency (IDF), measuring a token's distinctiveness across all documents. Intuitively, TF captures the intensity of a given skill/topic in a course or field. For example, courses in Economics more typically include the tokens "economics" and "regression analysis" than "Shakespeare" or "cybersecurity." Variation in the occurrence of different words/phrases is captured by the TF weight applied to each token for a given course. IDF assigns more weight to uncommon tokens and lower weight to common words. This ensures that comparisons between courses emphasize meaningful differences in content rather than variations in generic language. For example, the IDF weight emphasizes the contribution of less common tokens, like "economics" and "regression analysis," over words that appear commonly in course descriptions, like "student"

[37] In my data, 62% of courses are modified or discontinued over a ten-year period (see Appendix Figure A-11).

[38] I apply standard pre-processing to each course description. For example, I remove punctuation, standardize capitalization, remove overly-common "stopwords" (e.g., "the" and "is"), and lemmatize all words (e.g., transform "learns" or "learning" to "learn"). The complete processing procedure is described in Appendix C.

or "exam."[39]

5.2 Validating course description data

To validate the effectiveness of course descriptions in assessing course content, it is essential to demonstrate that they provide meaningful insights about courses. Specifically, variation in topics or skills across fields or over time should reflect genuine changes, rather than differences in terminology describing similar concepts. This section aims to show descriptively that the text data and methods reveal differences that are both meaningful and intuitive.

Distinctive tokens in course descriptions align closely with the skills and concepts emphasized in each field, demonstrating that these descriptions capture meaningful differences in content. Figure 4 applies the NLP methods outlined in the previous section to illustrate these differences, displaying the 25 most distinctive tokens for a sample of fields based on course descriptions from the 2022–23 academic year.[40] The results align intuitively with disciplinary focus: for example, English courses emphasize literature, reading, and writing, while Computer Science courses highlight programming and data analysis. These distinctive tokens capture both skills (e.g., reading, programming) and concepts (e.g., markets, theorems), reinforcing the validity of using course descriptions to analyze curricular content.

The effectiveness of the text analysis methods depends on their ability to detect substantive changes in course content over time, rather than merely shifts in terminology. For example, adding "climate change" to a course description where no equivalent concept previously existed signifies a meaningful change. In contrast, replacing "global warming" with "climate change" would represent a terminological update rather than a substantial alteration to the course.

In Figure 5, I demonstrate that changes in course description text represent meaningful differences in course content. For each field, I list 15 tokens distinctive of courses that have been discontinued over the last decade and 15 tokens distinctive of courses that have been introduced over the last decade. The figure highlights that the text data and methods pick up substantive changes to course content rather than changes in jargon. For example, recently created Economics courses emphasize data analysis, inequality, and topics in applied economics more than discontinued courses, which emphasize topics related to international economics and monetary policy. Similarly, Computer Science has shifted from hardware-oriented courses towards data science, cybersecurity, and machine learning.

[39]For more detail for a more detail on the construction of the TF-IDF weights and stylized example of how $v_{c,i,s,t}$ is constructed, see Appendix C.

[40]I collapse course descriptions into a single document for each institution-field, then represent each document as a TF-IDF vector. I average across institutions to get field-level values, then select the 25 tokens with highest value for each field.

5.3 Characterizing how curriculum changes

Having described the course descriptions, I next develop a measure of "alignment" between course content and student demand and analyze how this measure evolves over time. Lacking a direct measure of students' preferences for specific topics or skills, I measure the changing relationship between course content and a broad set of themes relevant to students' objectives and the mission of the university. Specifically, I focus on five themes: job relevance (as a response to economic changes); current events relevance and social justice[41] relevance (as a response to societal changes); and technology and scholarship relevance (as a response to technological change and innovation).[42]

I relate course descriptions to these themes by developing a weighting scheme that captures a token's importance to texts highly connected to that theme. For example, "oppression" is a word that might appear frequently in social justice texts and "machine" is a word that might appear frequently in technology texts, but it is unlikely for the word "machine" to frequently appear in social justice texts, and vice versa. To capture the career relevance of a given token, I measure its frequency in job descriptions. To capture a token's importance to current events, I measure its frequency in the text of front-page articles published by the New York Times. To capture a token's relation to research scholarship, I measure its frequency in abstracts for top academic journals.[43] To capture a token's connection to social justice, I measure its frequency in a corpus of books and press releases from organizations oriented towards social justice causes. Finally, to capture a token's relation to technological

[41]The period of this analysis coincides with a well-documented increase in Diversity, Equity, and Inclusion initiatives at U.S. colleges and universities. The inclusion of social justice as a theme reflects evolving institutional and student priorities, which may drive or align with this broader trend. This theme encompasses topics related to justice and activism concerning race, gender, sexuality, ability, immigration, and civil liberties, broadly defined. Rather than taking a normative stance, this analysis seeks to empirically document how institutions adjust course content in response to changing societal and student demand.

[42]I chose themes based on three main criteria: 1) their general applicability across multiple fields, 2) the availability of text data sources that could help identify relevant words/phrases, and 3) their relevance to intellectual interests students might possess or cultivate during college.

[43]Important work by Biasi and Ma (2022) explores this dimension of course content alignment in much greater detail, albeit slightly differently from the analysis described in this paper. Their analysis uses repeated cross sections of course syllabi to characterize differences in exposure to content on the cutting edge of research across institutions. Their analysis documents differences across universities in the provision of courses on the research frontier, and shows how instructors contribute to the innovative content of courses when they observe a change in instructor. My analysis builds on this important work, yet there are key differences in both the data sources used and our objectives. My dataset contains longitudinal data encompassing the full set of courses offered by a field. This enables me to observe within-institution and field-specific shifts over time and in response to changing enrollment. Biasi and Ma's analysis emphasizes the novelty of course content, whereas my research measure is related to a course's general connection with research-themed topics. My approach considers the relevance of both transitory and enduring research topics, recognizing terms like "research" or "analysis" that perennially indicate research relevance. While not differentiating between cutting-edge and older research content, it provides insight into a course's alignment with research-related topics over time.

progress, I measure its frequency in patent text. The text data sources used for quantifying each of these shifts, along with the procedure used to process these data sources, are detailed in Appendix C.

I construct "relevance weights" for each token w with respect to each theme q.[44] The weights are designed to assess each token's significance to the reference text relative to a neutral text source - in this case, the corpus of Wikipedia articles. Each weight is calculated as the ratio of the token w's share in documents of type q, to the sum of the token's shares both in documents of type q and in Wikipedia articles.[45] To demonstrate, consider the construction of the current events relevance weight of a highly topical token like "climate change." Climate change represents 0.002% of tokens in the Wikipedia data and 0.0103% of tokens in New York Times featured articles. Thus, the current events relevance weight on "climate change" is:

$$weight^{\text{current events}}_{\text{climate change}} = \frac{0.000103}{0.000103 + 0.00002} = 0.837.$$

Table 6 presents relevance weights for a selection of tokens, highlighting two important features of the method. The top panel shows relevance weights for five tokens, each aligned with one of the five themes. Each token has a high relevance weight within its corresponding theme, demonstrating the method's ability to identify important terms from the thematic documents. The bottom panel shows that word pairs with similar meanings or contexts consistently have comparable relevance weights. This consistency underscores the robustness of the method, ensuring that subtle shifts in jargon or terminology — common in academic and professional texts — do not distort alignment scores. By capturing meaningful thematic alignment while remaining insensitive to superficial linguistic variations, the method provides a reliable framework for assessing course-to-theme alignments.

To measure the extent to which a course aligns with a given theme, I calculate a "curriculum alignment score" for each course, year, and theme tuple. The curriculum alignment score is the sum of the relevance weights for tokens in a field's descriptions, weighted by the TF-IDF weights. In essence, the score averages the theme-specific importance of words/phrases in the course descriptions, with greater weight given to words/phrases distinctive to each

[44] Alternative strategies for measuring the curriculum alignment of course descriptions include using a multinomial classifier or a more sophisticated embeddings model. These alternative methods are more flexible than the method described above. The primary advantage of my expression weighting approach is its transparency; it is easy to validate the weights assigned to each token and interpret how these weights contribute to the alignment scores.

[45] This weight is analogous to the conditional probability from an experiment where a thematic corpus (q or the corpus of Wikipedia articles) is randomly selected and a token w is subsequently randomly picked from that category. The relevance weight therefore represents the conditional probability that if a particular token w was chosen, it originated from the theme q.

document. Appendix C provides a detailed example of how a curriculum alignment score is calculated.

To validate the method, I plot the average alignment scores for courses offered in 2022-23, categorized by field and averaged across institutions, in Figure 6.[46] The figure confirms that different fields align with the five themes in intuitive ways. For example, courses in Economics and Business show stronger alignment with themes related to current events and job-related skills. In contrast, courses in the Humanities are less vocationally focused but exhibit a modest alignment with current events. Computer Science courses, meanwhile, reflect a blend of academic research, vocational skills identified in job descriptions, and technological advancements found in patents.[47]

Next, I describe the changes in curriculum alignment over the past 20 years. Figure 7 plots the trend in average curriculum alignment of college courses offered since 2002-03. I estimate course-level regressions of curriculum alignment scores on a vector of time dummies, controlling for institution-by-field fixed effects. The estimates are normalized as the change in curriculum alignment (in standard deviations) relative to the curriculum alignment of the average course in 2012-13.

Figure 7 demonstrates that college course descriptions have gradually incorporated topics that are related to the themes relevant to students' interests. For example, the average college course became 0.044 sd more social justice-aligned between 2012-13 and 2022-23. For comparison, the 10-year change in social justice alignment over the period of a decade is approximately 6.5% of the difference in social justice alignment between the average Sociology course and the average Business course. Given that I am controlling for institution-by-field fixed effects, this trend is not driven by shifts in the composition of course offerings across fields, but represents within-field changes in the topics covered. The greatest growth during this period is in emphasis on topics related to current events, social justice, and job relevance.

Two factors may be driving the general upward trend in curriculum alignment: courses may be increasingly incorporating topics related to the identified themes (or removing topics unrelated to them), as I argue in this section. Alternatively, the trend could be an artifact of the method itself, if the selection of reference texts systematically favors more recent courses. To assess this possibility, I conduct a robustness check that calculates alignment scores based on weights constructed using only the older or more recent halves of the thematic corpora. Trends are broadly consistent across these constructions, suggesting that the observed shifts

[46] The results are qualitatively similar when analyzing course offerings from other years.

[47] Appendix Figure A-12 summarizes an additional validation exercise comparing alignment scores to ChatGPT's rankings of course alignment with the themes. ChatGPT and the curriculum alignment scores agree in approximately 90% of pairwise comparisons when contrasting a course in the lowest quartile of a theme with one in the top quartile.

reflect substantive curricular changes rather than methodological bias. A more detailed discussion of this analysis, including differences in trends across themes and the role of changes in the composition of news coverage, is provided in Appendix C.

The process by which curricula adapt to align with these themes has significant implications for how universities disseminate knowledge. If existing courses are continually updated to reflect new developments, the persistence of course offerings may not constrain students' access to content that meets their evolving demands. However, if curricular adjustments rely primarily on introducing new courses and phasing out outdated ones, inelasticity in course offerings could restrict students' exposure to the most relevant and timely content.

I next assess the sources of growing curricular alignment for each theme. Following Foster et al. (2001), I decompose the average change in curriculum alignment into four components. The "within" component measures changes attributable to changing course content for the same course offered in both 2012-13 and 2022-23.[48] The "between" component measures changes attributable to enrollment shifts between the continuously offered courses. The "exit" component measures changes due to the discontinuation of courses offered in 2012-13 but not in 2022-23. And the "entry" component measures changes due to the creation of courses that are offered in 2022-23 but were not offered in 2012-13. I measure changes within each institution and field, aggregate these changes at the institution level weighted by each field's share of total start-of-period enrollment, and then compute an unweighted average across institutions. I describe the decomposition procedure in greater detail in Appendix F.

Figure 8 plots the decomposition. The figure demonstrates that the increasing curriculum alignment of courses in my sample arises primarily due to the entry and exit of courses, rather than changes within existing courses. For example, the average social justice alignment of courses in my sample grew by 0.044 sd, of which 65% of the change came from the entry of new courses that are highly social justice-aligned.

The creation of new courses is an important way for universities to adapt their offerings, particularly for themes like current events and social justice, where new topics emerge frequently. By introducing new courses, universities can respond to social or technological developments in real time, in ways that existing courses might not address. Additionally, many universities restrict how much instructors can modify course descriptions without re-

[48]The risk with the "within" component is that courses might undergo changes that are not reflected in their descriptions. To address this, I evaluate the decomposition over a lengthy period. In cases where I find updates to descriptions for courses that are still being offered, the new and old descriptions typically share similar alignment scores. These scores are calibrated to represent the core themes of a course, which are usually consistent even if the description changes. While minor course adjustments might cause slight variations in these scores, substantial changes often prompt the creation of a completely new course. Many institutions have guidelines that limit the scope of course description modifications; beyond a certain point, a new course is typically introduced.

view from a curriculum oversight board, making course creation a more practical mechanism for adjustment. Finally, the curriculum alignment score is designed to capture the overarching "gist" of a course, which often remains stable even when some topics or tools change. For example, a programming course that shifts from Java to Python would not register as a major thematic change because it remains, fundamentally, a programming course.

Meanwhile, in some of the themes with more modest growth in alignment, such as technology, the growth is driven primarily by course exit. This represents less of an intentional effort on the part of the university to incorporate topics than a reflection of the kinds of courses that are likely to be discontinued during this time. For example, referring to Figure 5, courses in English departments most likely to be discontinued are courses related to poetry and British literature - courses that are inherently not strongly aligned to technology, even relative to other English courses. The discontinuation of these courses may be downstream of students' preferences for these themes, but does not reflect intentional alignment on the part of the university in the same way that alignment from course entry does.

The gradual shift in course content toward increased alignment with the five themes occurs through varied pathways. Course entry and exit play a significant role in driving these changes, which helps contextualize the earlier finding that the quantity of courses adjusts far less than one-for-one to shifts in student demand. The evidence presented in this section suggests that fields with more dynamic course offerings tend to align more closely with themes that resonate with students.

5.4 Heterogeneity

Figure 9 illustrates differences in the topics emphasized by universities and how these emphases have shifted over time. The figure compares average curriculum alignment in courses offered in 2012-13 and 2022-23, broken down by topic and Carnegie classification. Estimates are presented in standard deviation differences relative to the average course at R1 institutions in 2012-13, based on course-level regressions of alignment scores on Carnegie classification-by-academic year dummies. Each institution-term is weighted equally, with courses within each term weighted by enrollment. To isolate differences beyond field composition, I control for field fixed effects, ensuring the estimates reflect alignment changes driven by course content, offerings, and shifts in enrollment patterns. These results capture the alignment of a typical course a student takes at a given institution during a given term, highlighting both temporal changes and institutional variation.

The figure suggests three insights. First, it confirms the broader curriculum shifts toward themes like social justice, job relevance, and current events between 2012-13 and 2022-23, consistent with the patterns observed in Figure 7. Second, universities differ in the topics to

which students are exposed: at R2 and less research-intensive universities, students enroll in courses that are more job-related, while students at R1 and liberal arts universities enroll in more current events-oriented courses, particularly recently. Third, changes in curriculum alignment differ notably by institution type over time. The changes have been most striking at R1 and liberal arts institutions, where a combination of changing course offerings and shifting enrollment has led to markedly greater exposure to current events and social justice. Changes have been comparatively smaller at teaching-focused and R2 institutions.

Course supply and demand interact in important ways to generate the results in this figure. On the supply side, universities have expanded course offerings in social justice and job-relevant skills. On the demand side, students are increasingly sorting into these courses. I leave for future work whether supply or demand leads these changes, or whether they are jointly pushed by outside forces. However, differences in both the level and rate of change in curriculum alignment provide insight into the skills and topics a university prioritizes in its course offerings, and its capacity to make changes to its offerings over time.

6 Conclusion

This paper examines course supply changes with respect to changing student demand within American universities over more than two decades. I use a unique dataset that I constructed by scraping online course catalogs to measure how course supply adjusts to changing demand along both extensive and intensive margins.

During a period when students' demand for different fields of study changed dramatically, universities responded less than one-for-one in adjusting the quantity of courses to meet changing demand. I estimate that a 10% change in demand for a field results in a 2.2% change in courses offered and a 5.5% change in course sections. Notably, course quantity is more elastic when enrollment in a field is growing relative to when enrollment is shrinking. Course supply is also highly persistent: 65% of courses offered in 2022-23 have been offered for at least a decade, and changes to course descriptions are modest after a course is introduced. Thus, the primary channel for universities to innovate and align with student demand is through introducing new courses. Course content gradually adopts topics relevant to students' interests, including courses related to social justice and job-relevant skills, through the introduction of new courses. In heterogeneity analysis, I show that R1 universities are more responsive to changing enrollment - particularly for fields with growing enrollment - and that their students enroll in courses that emphasize current events and social justice, whereas students at universities with less of a research emphasis enroll in a more vocational curriculum.

The analysis in this paper is purely positive, but it is possible to point to trade-offs involved in offering a more or less adaptive curriculum. Inelasticity harms students when they cannot secure seats in their preferred courses or when the course content hasn't been updated to reflect relevant topics and skills. Other students benefit from inelasticity, particularly those who get to enroll in smaller courses because of the university's continued support of less popular fields. Similarly, faculty bear a cost when creating new courses, so some inelasticity in course supply offers assurance that the upfront investment required to create a course can be recouped over time.

The extent to which the course quantity elasticities I estimate deviate from the socially optimal course quantity elasticity depends on the appropriate balance between the welfare gains and costs from course quantity modifications. The optimal course quantity elasticity is likely greater than 0, such that students can develop human capital necessary to succeed in an evolving economic, social, and technological landscape. However, it is almost certainly less than 1 due to inelasticities in instructor supply. Moreover, the university must balance the welfare of current and future students in a way that may dampen responses to short-run student demand shocks. The socially optimal elasticity depends crucially on social weights placed on the various objectives a university satisfies, and may differ across universities. In a companion project to this paper, I outline a stylized model of course supply and demand within the university and use this model to develop testable hypotheses for the sources of inelasticity in course supply, and particularly to better understand whether inelasticity is reflective of trade-offs in the production of knowledge or indicative of frictions in the university that create a wedge between the university's preferred and actual quantity of courses (Light 2024).

The findings in this paper highlight a potential misalignment between how students and policymakers perceive universities' responsiveness to labor market conditions and how universities actually adjust. From a policy perspective, the asymmetries in course supply elasticity have implications for addressing skill gaps in the labor market. Policymakers aiming to enhance workforce alignment may benefit from supporting institutional flexibility in curriculum adjustments, particularly in fields experiencing rapid growth. Additionally, these findings highlight the need to consider institutional constraints when designing interventions, such as funding incentives, that aim to influence educational supply. For university administrators, understanding the elasticity of course supply to enrollment changes can guide strategic decisions on resource allocation, faculty hiring, and curriculum design. By closely monitoring shifts in demand, institutions can better serve their students — whether by accommodating demand or steering students away from perceived fads — while maintaining relevance and competitiveness in a dynamic educational environment.

A logical extension of this research would be to link inelasticity in higher education with students' labor market outcomes. By linking inelasticity in course supply to students' outcomes, we can gauge if inelasticity adversely impacts students. Additionally, as the landscape of higher education shifts with the emergence of disruptors in higher education — like private for-profit universities and bootcamp programs — that compete with traditional four-year institutions by offering a more focused and adaptable curriculum, this paper's insights can guide universities' adjustments to an evolving higher education landscape.

Works cited

Acton, Riley K (2021). "Community college program choices in the wake of local job losses". In: *Journal of Labor Economics* 39.4, pp. 1129–1154.

Ahn, Thomas et al. (2019). *Equilibrium grade inflation with implications for female interest in stem majors*. Tech. rep. National Bureau of Economic Research.

Altonji, Joseph G, Peter Arcidiacono, and Arnaud Maurel (2016). "The analysis of field choice in college and graduate school: Determinants and wage effects". In: *Handbook of the Economics of Education*. Vol. 5. Elsevier, pp. 305–396.

Autor, David H, Frank Levy, and Richard J Murnane (2003). "The skill content of recent technological change: An empirical exploration". In: *The Quarterly journal of economics* 118.4, pp. 1279–1333.

Bandiera, Oriana, Valentino Larcinese, and Imran Rasul (2010). "Heterogeneous class size effects: New evidence from a panel of university students". In: *The Economic Journal* 120.549, pp. 1365–1398.

Bedard, Kelly and Peter Kuhn (2008). "Where class size really matters: Class size and student ratings of instructor effectiveness". In: *Economics of Education Review* 27.3, pp. 253–265.

Beffy, Magali, Denis Fougere, and Arnaud Maurel (2012). "Choosing the field of study in postsecondary education: Do expected earnings matter?" In: *Review of Economics and Statistics* 94.1, pp. 334–347.

Biasi, Barbara and Song Ma (2022). *The education-innovation gap*. Tech. rep. National Bureau of Economic Research.

Bleemer, Zachary and Aashish Mehta (2021). "College Major Restrictions and Student Stratification. Research & Occasional Paper Series: CSHE. 14.2021." In: *Center for Studies in Higher Education*.

— (2022). "Will studying economics make you rich? A regression discontinuity analysis of the returns to college major". In: *American Economic Journal: Applied Economics* 14.2, pp. 1–22.

Blom, Erica, Brian C Cadena, and Benjamin J Keys (2021). "Investment over the business cycle: Insights from college major choice". In: *Journal of Labor Economics* 39.4, pp. 1043–1082.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022). "Quasi-experimental shift-share research designs". In: *The Review of Economic Studies* 89.1, pp. 181–213.

Chau, Hung et al. (2023). "Connecting higher education to workplace activities and earnings". In: *Plos one* 18.3, e0282323.

Conzelmann, Johnathan G et al. (2022). *Grads on the go: Measuring college-specific labor markets for graduates*. Tech. rep. National Bureau of Economic Research.

— (2023). *Skills, Majors, and Jobs: Does Higher Education Respond?* Tech. rep. National Bureau of Economic Research.

Courant, Paul N and Sarah Turner (2017). *Faculty deployment in research universities*. Tech. rep. National Bureau of Economic Research.

Deming, David J, Claudia Goldin, and Lawrence F Katz (2012). "The for-profit postsecondary school sector: Nimble critters or agile predators?" In: *Journal of Economic Perspectives* 26.1, pp. 139–164.

Deming, David J and Kadeem Noray (2020). "Earnings dynamics, changing job skills, and STEM careers". In: *The Quarterly Journal of Economics* 135.4, pp. 1965–2005.

Denning, Jeffrey T et al. (2022). "Why have college completion rates increased?" In: *American Economic Journal: Applied Economics* 14.3, pp. 1–29.

Eggenberger, Christian, Miriam Rinawi, and Uschi Backes-Gellner (2018). "Occupational specificity: A new measurement based on training curricula and its effect on labor market outcomes". In: *Labour Economics* 51, pp. 97–107.

Foster, Lucia, John C Haltiwanger, and Cornell John Krizan (2001). "Aggregate productivity growth: lessons from microeconomic evidence". In: *New developments in productivity analysis*. University of Chicago Press, pp. 303–372.

Freeman, Richard B (1976). "A cobweb model of the supply and starting salary of new engineers". In: *ILR Review* 29.2, pp. 236–248.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy (2019). "Text as data". In: *Journal of Economic Literature* 57.3, pp. 535–574.

Gilpin, Gregory A, Joseph Saunders, and Christiana Stoddard (2015). "Why has for-profit colleges' share of higher education expanded so rapidly? Estimating the responsiveness to labor market changes". In: *Economics of Education Review* 45, pp. 53–63.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020). "Bartik instruments: What, when, why, and how". In: *American Economic Review* 110.8, pp. 2586–2624.

Hemelt, Steven W et al. (2021). "Why is math cheaper than English? Understanding cost differences in higher education". In: *Journal of Labor Economics* 39.2, pp. 397–435.

Kokkelenberg, Edward C, Michael Dillon, and Sean M Christy (2008). "The effects of class size on student grades at a public university". In: *Economics of Education review* 27.2, pp. 221–233.

Kraft, Matthew A et al. (2020). "Teacher accountability reforms and the supply and quality of new teachers". In: *Journal of Public Economics* 188, p. 104212.

Light, Jacob Douglas (2024). "Essays in the economics of education". PhD thesis. Chap. 2.

Long, Mark C, Dan Goldhaber, and Nick Huntington-Klein (2015). "Do completed college majors respond to changes in wages?" In: *Economics of Education Review* 49, pp. 1–14.

MacLeod, W Bentley and Miguel Urquiola (2021). "Why does the United States have the best research universities? Incentives, resources, and virtuous circles". In: *Journal of Economic Perspectives* 35.1, pp. 185–206.

Monks, James and Robert M Schmidt (2011). "The impact of class size on outcomes in higher education". In: *The BE Journal of Economic Analysis & Policy* 11.1.

Mumford, Kevin J, Richard Patterson, and Anthony Yim (2024). *College Course Shutouts*. Tech. rep. IZA Discussion Papers.

Patnaik, Arpita, Matthew Wiswall, and Basit Zafar (2021). "College majors 1". In: *The Routledge handbook of the economics of education*, pp. 415–457.

Ruggles, Steven et al. (2023). "IPUMS USA: Version 13.0 [dataset]." In: URL: https://doi.org/10.18128/D010.V13.0.

Thomas, James (2024). "What do course offerings imply about university preferences?" In: *Journal of Labor Economics* 42.1, pp. 53–83.

Weinstein, Russell (2020). "Local labor markets and human capital investments". In: *Journal of Human Resources*, 1119–10566R2.

Wiswall, Matthew and Basit Zafar (2015). "Determinants of college major choice: Identification using an information experiment". In: *The Review of Economic Studies* 82.2, pp. 791–824.

Figure 1. Sample entry in the course catalog dataset



ECON 144 **Winter 2021**
Lecture: 01
Units: 5
Class#: 31846

Family and Society

Department of Economics

Lecture, Discussion
1/11/21 - 3/19/21

To be Scheduled
12:00 AM - 12:00 AM
Remote

Instructor: Persson, P.

Enrollment Status

Open Seats: 0
Enrolled: 81 Waitlist: No waitlist
Capacity: 80 Waitlist Max: No waitlist

Course Description

The family into which a child is born plays a powerful role in determining lifetime opportunities. This course will apply tools from economics and related social sciences to study how the functioning of families is shaped by laws, social insurance, social norms, and technology. Topics will include intergenerational transmission of wealth and health, the importance of the early family environment, partnership formation, cohabitation and marriage, teen pregnancy and contraception, assisted reproduction, Tiger Moms and Helicopter Parenting, and the employment effects of parenthood. In the context of these topics, the course will cover social science empirical methods, including regression analysis, causal inference, and quasi-experimental methods. Throughout the course, we will think critically about the role of the government and how the design of public policy targeting families affect our ability to solve some of the most important social and economic problems of our time. Prerequisites: Econ 50

Grading basis
Letter or Credit/NoCredit Exce

Class level
Undergraduate

Instructional mode
Remote: Asynchronous

Final exam

Meets Requirement(s)
WAYS - Applied Quantitative Reasoning (AQR)
WAYS - Social Inquiry (SI)
SYMBO-BS Subplan: Computational Social Science
ECON-BS Core Program Requirements
ECON101 ECON101 Prerequisite

Source: Stanford University.

Table 1. Characteristics of course catalog sample

4 year institutions

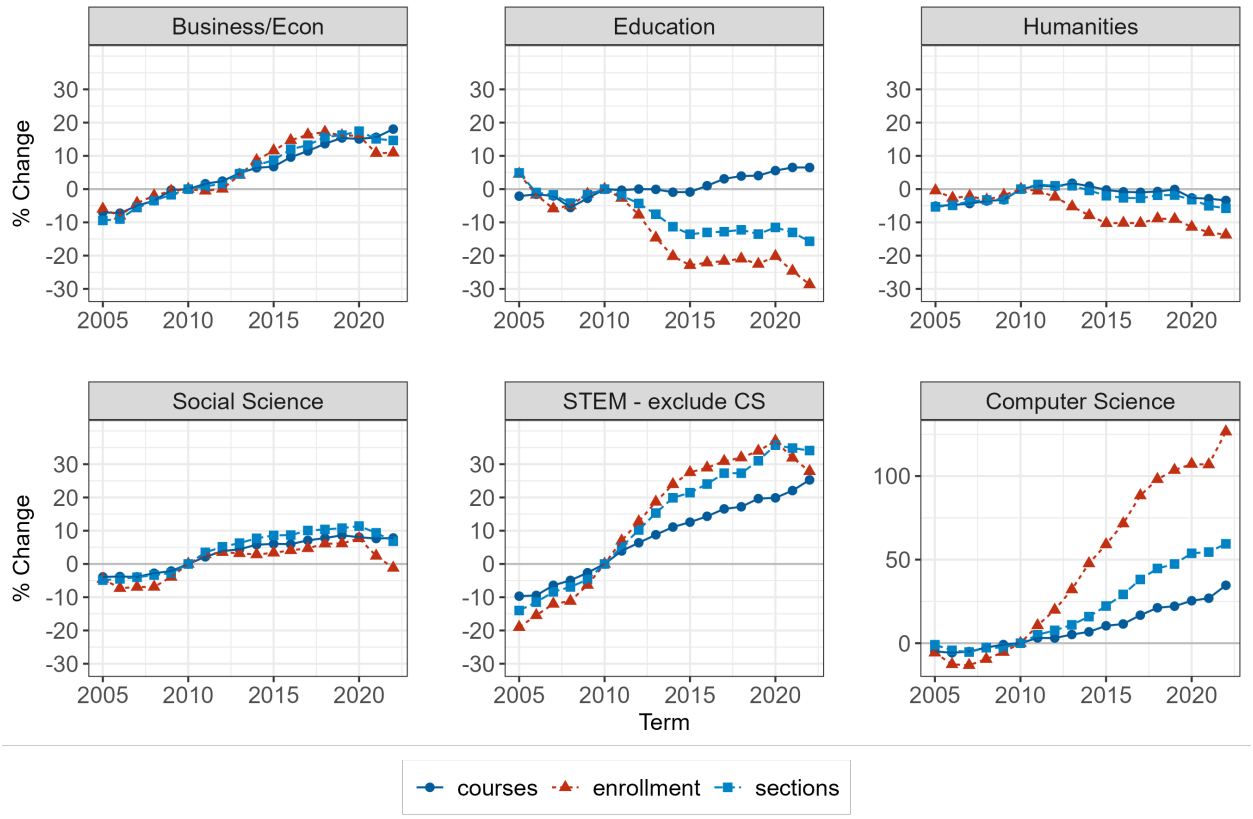
	Population		Catalog Sample		Enrollment Sample	
	mean	sd	mean	sd	mean	sd
Enrollment	4,912	20,330	9,505	20,547	9,389	22,066
Public share	72.48	44.66	80.46	39.65	80.51	39.61
Average tuition	16,774	15,370	16,013	14,713	16,013	14,933
Average price	16,786	8,455	17,045	7,650	17,176	7,520
Admit rate	71.99	22.54	70.41	23.82	71.78	23.68
Tenure share	51.45	19.29	54.48	13.90	54.43	15.15
Student-faculty ratio	17.43	5.40	17.53	4.48	17.30	4.21
6-year graduation rate	59.54	19.62	63.38	18.55	61.99	19
Endowment per student	58,785	215,681	75,081	261,983	68,207	259,094
Tuition % of revenue	34.12	19.83	31.92	17.15	32.26	17.21
Research % of spending	8.77	11.96	11.56	13.50	11.59	14.09
N	1,972		528		374	

2 year institutions

	Population		Catalog Sample		Enrollment Sample	
	mean	sd	mean	sd	mean	sd
Enrollment	5,194	16,543	6,376	14,178	5,551	8,858
Public share	99.34	8.11	100	0.00	100	0.00
Average tuition	3,495	1,978	3,468	1,468	3,350	1,462
Average price	7,973	3,079	7,669	2,479	7,674	2,721
Student-faculty ratio	19.29	5.38	19.09	4.57	18.70	4.36
N	933		254		206	

Notes: Institution characteristics from IPEDS for the 2021-22 academic year. Only non-profit, Title IV-eligible, degree-granting institutions are included. Values except for undergraduate enrollment are weighted by enrollment. Averages exclude missing values. The 'Catalog Sample' includes all institutions in the sample. The 'Enrollment Sample' includes those with course-level enrollment data.

Figure 2. Trends in course enrollment and quantity: comparison to 2010-11



Notes: This figure plots the relative growth trends in course enrollment, course quantity, and section quantity across six aggregated field categories. Enrollment and course quantity for each institution and field category are indexed to their respective levels in the academic year 2010-11. The plotted points represent the average of these indexed values, averaged across all institutions in the sample. The figure restricts to upper-level courses offered at institutions with that first appear in the catalog data in 2010-11 or earlier.

Table 2. Course quantity elasticity estimates

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change - overall	0.203	0.183	0.264	0.288	0.352	0.353	0.316	0.316
	(0.028)	(0.046)	(0.029)	(0.036)	(0.033)	(0.034)	(0.064)	(0.050)
% enrollment change - field	0.210	0.206	0.326	0.337	0.411	0.422	0.400	0.217
	(0.010)	(0.011)	(0.010)	(0.015)	(0.012)	(0.017)	(0.026)	(0.052)
First Stage F-stat								113.7
Observations	81,845	39,146	67,578	17,566	42,458	11,218	3,756	3,756
R^2	0.067	0.058	0.176	0.175	0.322	0.320	0.322	0.266

Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, standard errors are clustered at the institution and field-by-Census division level, which is the level of variation for the instrument. Significance stars are suppressed, as the natural comparison is not obviously 0.

Table 3. Section quantity elasticity estimates

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change - overall	0.379	0.350	0.523	0.573	0.662	0.678	0.651	0.651
	(0.039)	(0.063)	(0.038)	(0.043)	(0.033)	(0.032)	(0.052)	(0.032)
% enrollment change - field	0.316	0.315	0.501	0.524	0.613	0.629	0.618	0.552
	(0.011)	(0.013)	(0.011)	(0.015)	(0.010)	(0.015)	(0.026)	(0.038)
First Stage F-stat								113.7
Observations	81,845	39,146	67,578	17,566	42,458	11,218	3,756	3,756
R^2	0.137	0.119	0.342	0.352	0.552	0.554	0.585	0.581

Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, standard errors are clustered at the institution and field-by-Census division level, which is the level of variation for the instrument. Significance stars are suppressed, as the natural comparison is not obviously 0.

Table 4. Asymmetric course quantity elasticity estimates

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change - overall	0.206	0.187	0.266	0.296	0.359	0.368	0.325	0.320
	(0.028)	(0.047)	(0.030)	(0.038)	(0.035)	(0.036)	(0.068)	(0.050)
% enrollment change - growing	0.226	0.226	0.338	0.366	0.443	0.470	0.460	0.368
	(0.011)	(0.017)	(0.012)	(0.019)	(0.014)	(0.019)	(0.030)	(0.062)
% enrollment change - shrinking	0.198	0.191	0.317	0.313	0.386	0.385	0.349	0.082
	(0.016)	(0.016)	(0.017)	(0.020)	(0.019)	(0.026)	(0.037)	(0.079)
First stage F-stat								64
p-value grow = shrink	0.192	0.138	0.326	0.045	0.022	0.007	0.013	0.004
Observations	81,845	39,146	67,578	17,566	42,458	11,218	3,756	3,756
R^2	0.067	0.058	0.176	0.176	0.323	0.322	0.327	0.070

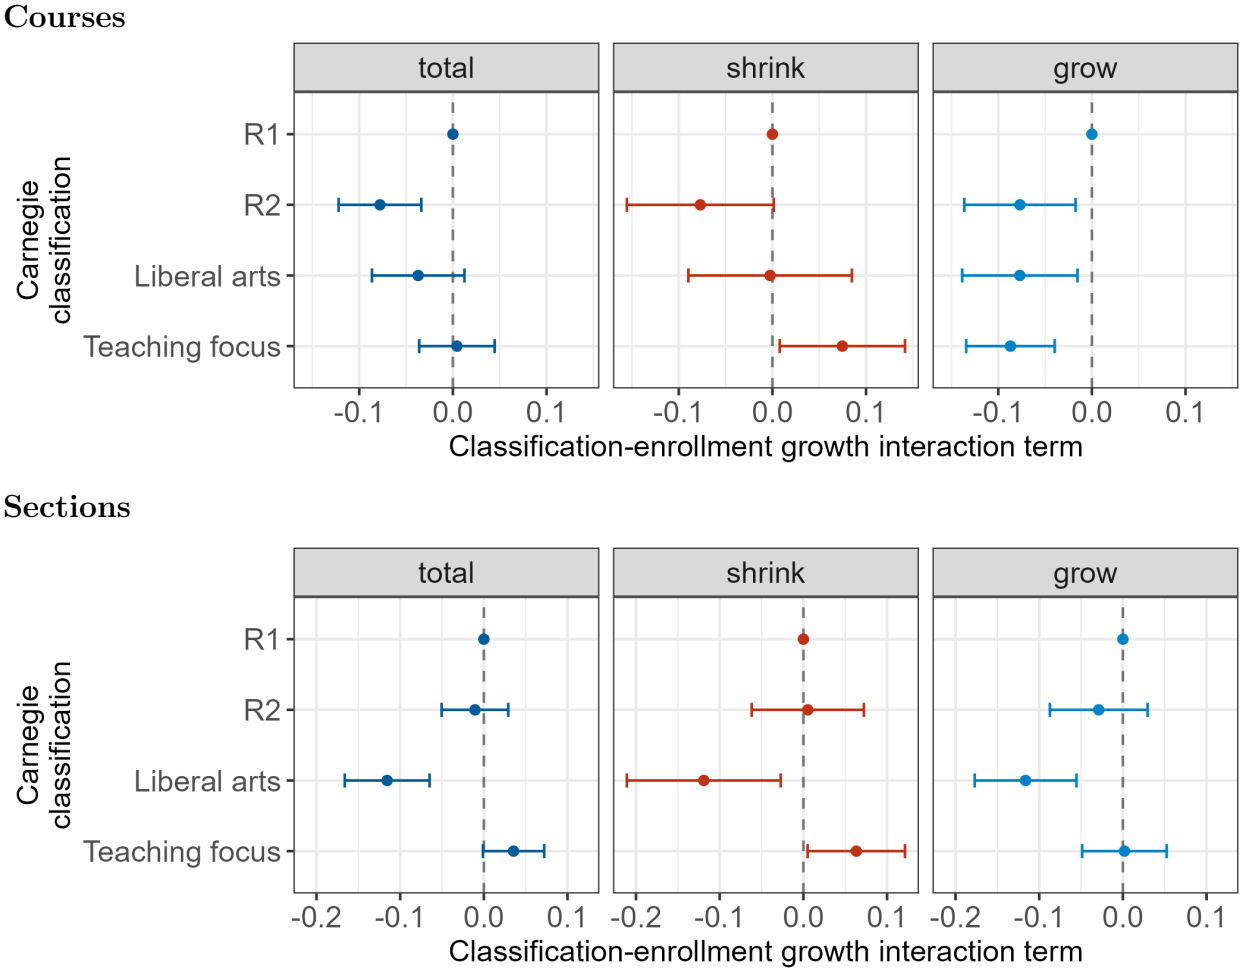
Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, bootstrapped standard errors are calculated using $1,000$ repetitions of the estimation, resampling region-by-field clusters in each iteration. Significance stars are suppressed, as the natural comparison is not obviously 0.

Table 5. Asymmetric section quantity elasticity estimates

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change - overall	0.382	0.354	0.527	0.579	0.669	0.690	0.657	0.652
	(0.040)	(0.064)	(0.039)	(0.046)	(0.035)	(0.033)	(0.055)	(0.032)
% enrollment change - growing	0.336	0.341	0.519	0.547	0.644	0.669	0.655	0.603
	(0.013)	(0.020)	(0.014)	(0.021)	(0.013)	(0.018)	(0.028)	(0.062)
% enrollment change - shrinking	0.302	0.295	0.487	0.506	0.588	0.598	0.585	0.507
	(0.017)	(0.018)	(0.016)	(0.021)	(0.016)	(0.024)	(0.037)	(0.062)
First stage F-stat								64
p-value grow = shrink	0.094	0.105	0.113	0.187	0.013	0.02	0.11	0.273
Observations	81,845	39,146	67,578	17,566	42,458	11,218	3,756	3,756
R^2	0.137	0.119	0.342	0.352	0.553	0.555	0.587	0.199

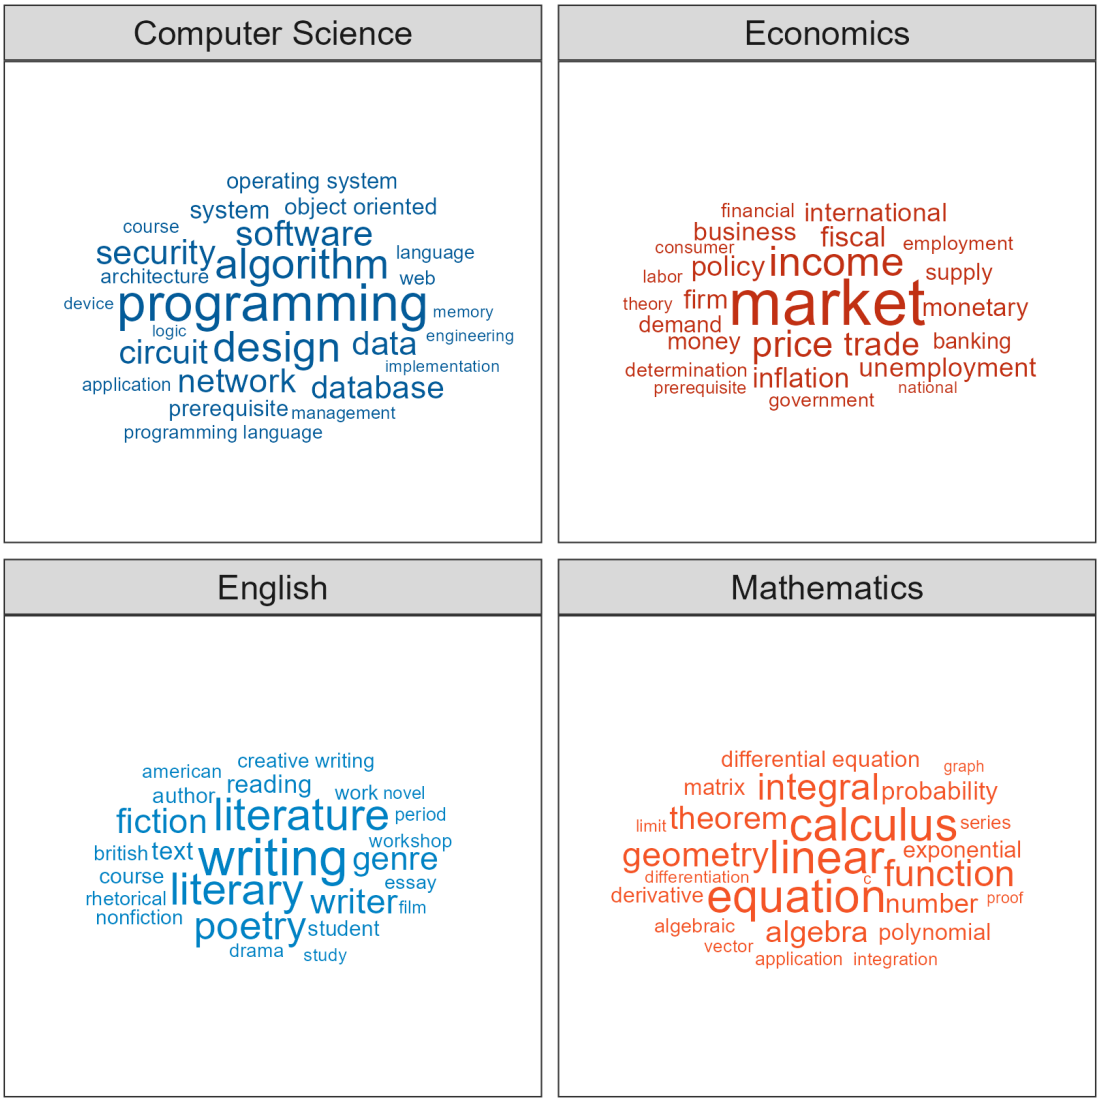
Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, bootstrapped standard errors are calculated using $1,000$ repetitions of the estimation, resampling region-by-field clusters in each iteration. Significance stars are suppressed, as the natural comparison is not obviously 0.

Figure 3. Heterogeneity in course and section quantity elasticity by Carnegie group



Notes: This figure compares the course quantity elasticities at different types of institutions. The figure plots estimates of the interaction term(s) between relative changes in enrollment and the institution's Carnegie classification, where R1 universities are the omitted institution category. The regression estimates eight-year course quantity elasticities with a one-year offset. Observations each the regression are at the institution-field-term level. Standard errors are clustered at the institution-term level.

Figure 4. Prominent words/phrases in selected fields (2022-23)



Notes: This figure showcases the top 25 tokens (words or phrases) for selected fields based on their average TF-IDF weight. For each field, courses from 2022-23 are aggregated into an institution-field document. The TF-IDF weight for each token is computed per document and then averaged within its field. Tokens containing the field's full name or common abbreviation (e.g., "Econ" for Economics) are excluded.

Figure 5. Evolution of token significance in discontinued vs introduced courses



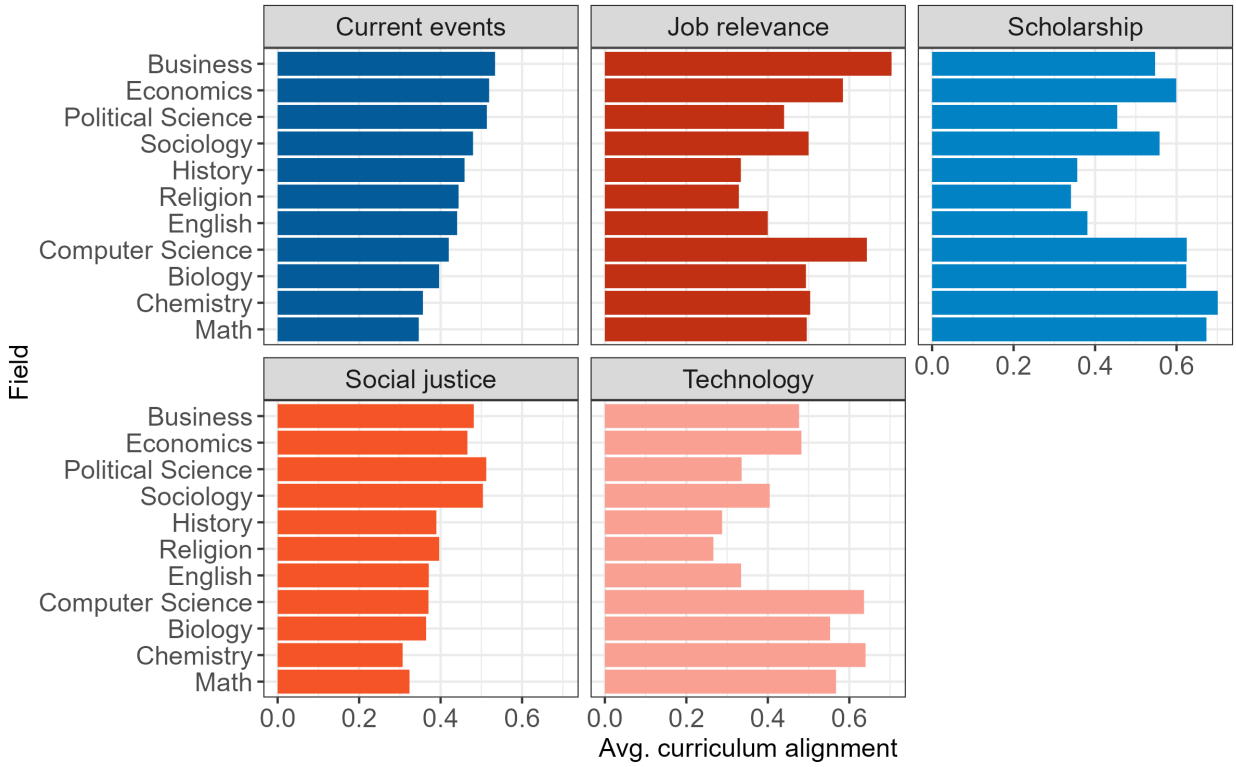
Notes: This figure contrasts the distinctive words of courses from 2012-13 to those of 2022-23. "Discontinued" courses are those offered in 2012-13 but no longer offered by 2022-23. "Introduced" courses are those not offered before 2012-13 but offered in 2022-23. Descriptions are grouped by field and course category (discontinued vs introduced). The visualization presents the top 15 tokens with the highest TF-IDF values from both course groups. Tokens with the field's full name or common abbreviation (e.g., "Econ" for Economics) are excluded.

Table 6. Relevance weights for sample tokens

	Token	Current events	Job relevance	Scholarship	Social justice	Technology
Distinctive tokens	financial crisis	0.88	0.03	0.45	0.41	0.01
	customer service	0.69	0.99	0.26	0.51	0.68
	regression	0.10	0.92	0.99	0.70	0.84
	injustice	0.67	0.06	0.29	0.93	0.00
	invention	0.38	0.17	0.03	0.48	0.99
Pairs of similar words	king	0.30	0.03	0.00	0.09	0.01
	queen	0.63	0.05	0.00	0.06	0.01
	dog	0.61	0.08	0.03	0.13	0.19
	cat	0.49	0.14	0.07	0.07	0.24
	blackberry	0.83	0.68	0.00	0.40	0.53
	iphone	0.90	0.69	0.00	0.77	0.50
	global warming	0.93	0.07	0.86	0.38	0.71
	climate change	0.84	0.17	0.83	0.36	0.14

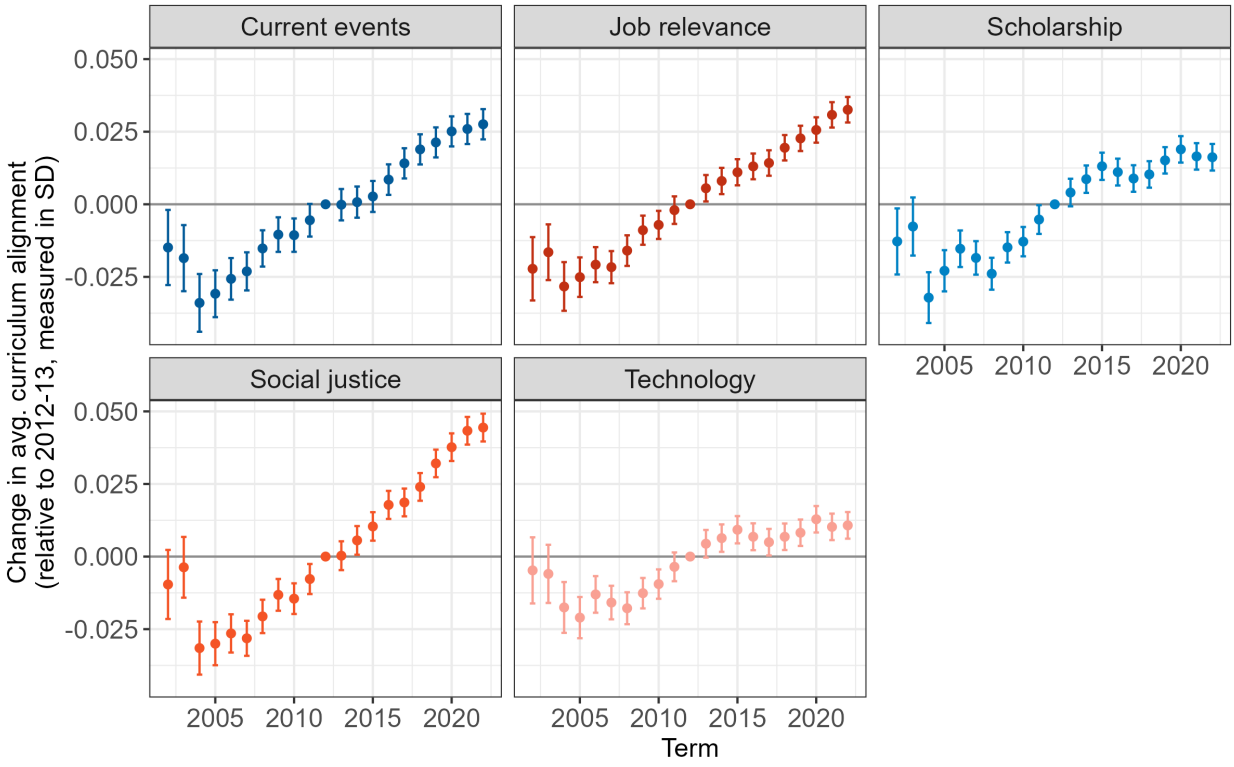
Notes: This table presents relevance weights of selected tokens. These weights measure a token's significance in a document related to a specific theme (e.g., a job description for job relevance) relative to its significance in a neutral text reference, such as the entirety of Wikipedia. Weights range from 0 to 1, with higher values indicating greater frequency in the thematic document compared to Wikipedia.

Figure 6. Average curriculum alignment across sampled fields (2022-23)



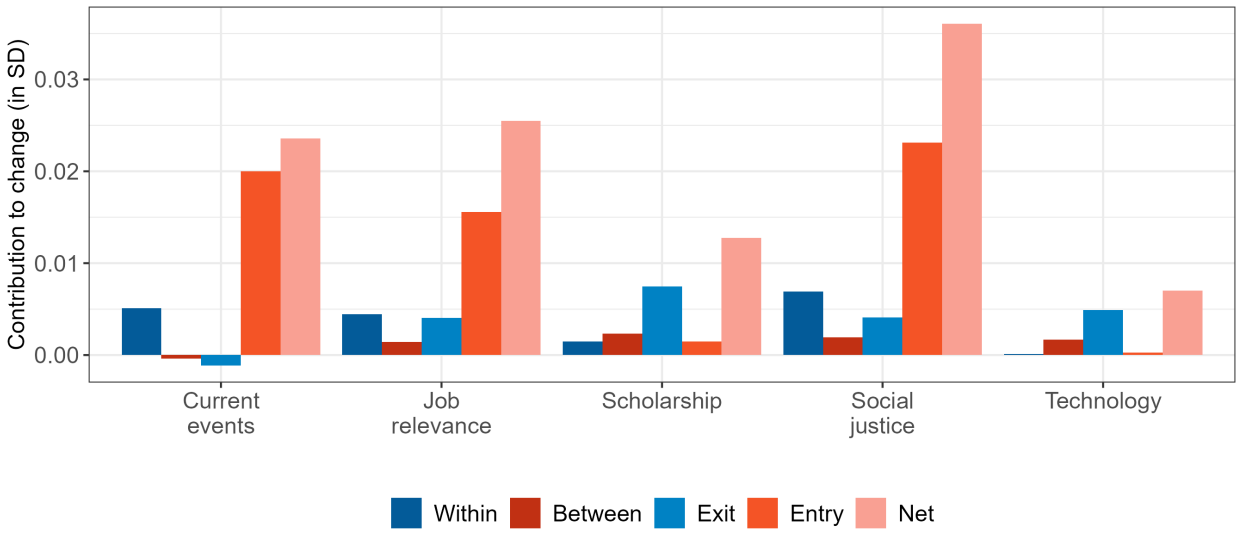
Notes: This figure plots the average curriculum alignment scores in a sample of popular fields for courses offered in 2022-23. Each course description is scored for its alignment to each of five distinct themes: current events, job relevance, scholarship, social justice, and technology. The bars plot the mean alignment within each field and theme, aggregated across institutions. Analysis restricts to upper-level courses. Fields are sorted in descending order according to their alignment with current events.

Figure 7. Change in curriculum alignment: 2002-03 to 2022-03



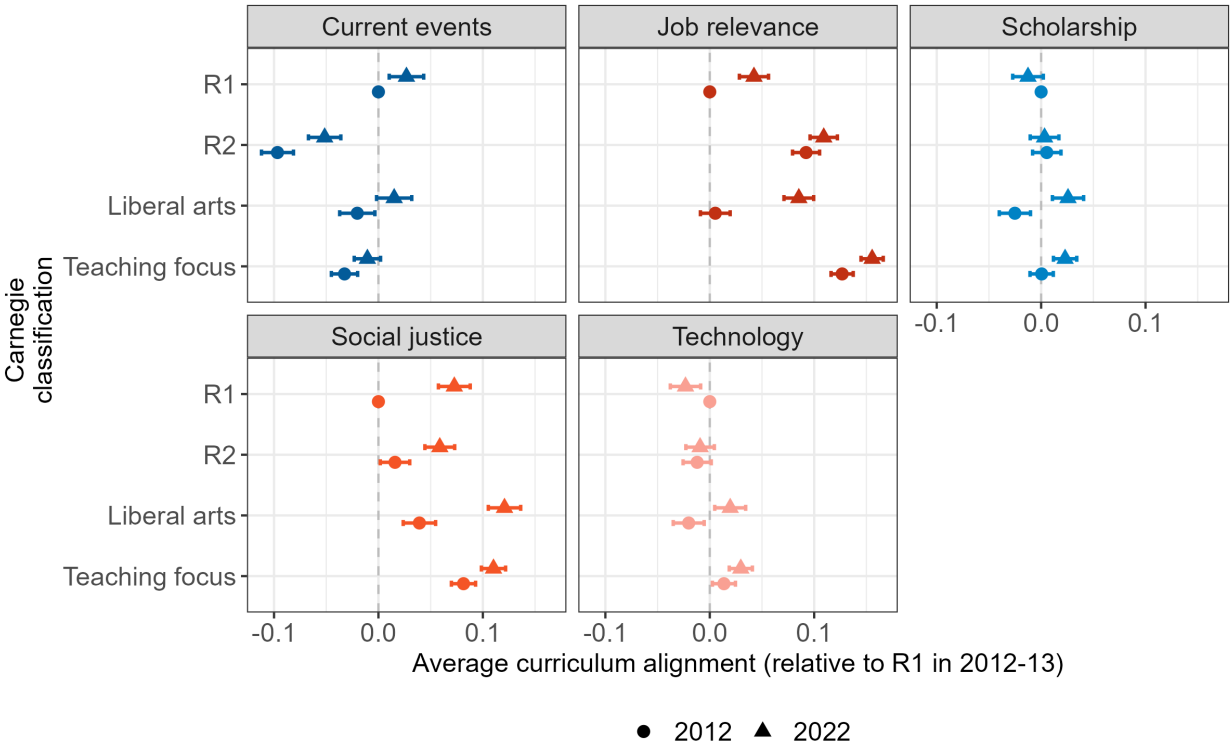
Notes: This figure plots the trend in curriculum alignment scores for courses offered from 2002-03 to 2022-23. The trend is estimated in course-level regressions of a course's curriculum alignment score for a given theme on a vector of year dummies, controlling for institution-by-field fixed effects. Changes are measured relative to the average curriculum alignment score in 2012-13 and reported in standard deviations. Analysis restricts to upper-level courses. Standard errors are clustered at the institution-field level.

Figure 8. Decomposition of curriculum alignment changes: 2012-13 to 2022-23



Notes: This figure decomposes down the shift in curriculum alignment between 2012-13 and 2022-23. Utilizing an approach based on Foster et al. (2001), the evolution in curriculum alignment at the institution-by-field level is decomposed into four components: updates within continuously-offered courses, changes due to shifts in enrollment across courses continuously offered by a field, discontinuations of courses, and introductions of new courses. Each institution receives equal weight; within each institution, fields are weighted by enrollment in 2012-13.

Figure 9. Heterogeneity in course content by Carnegie group: 2012-13 to 2022-23



Notes: This figure compares the topics emphasized in college courses at different types of institutions across different points in time. The estimates come from separate course-level regression of curriculum alignment score on Carnegie classification-by-academic year dummies. Estimates are transformed into standard deviation changes relative to R1 universities in 2012-13. For the regression, all institution-terms receive equal weight; within the institution-term, courses are weighted by enrollment. Thus, the estimates speak to variation in the topical exposure of the courses a typical student takes at different kinds of institutions and at different points in time. The regression restricts to a panel of institutions with course description and enrollment data in 2012-13 and 2022-23.

Appendix

A Dataset construction

I assembled a sample of schools for inclusion in the course catalog by using two strategies. Initially, I selected schools from the IPEDS directory to scrape their course catalogs. A manual search was conducted on over 1,000 institutions. For institutions with online course catalogs that were available in a format that could be scraped and had at least a few years of archived data, I scraped the course descriptions for all courses offered in all available years. Subsequently, I searched for institutions that used the most common course schedule templates to scrape course enrollment data, prioritizing those with at least five years of schedule data available.

The current sample comprises data from 782 institutions, including 528 4-year schools and 254 2-year schools. The 4-year schools make up 26% of schools and enroll 52% of the students at all 4-year non-profit, bachelor's degree-granting Title IV-eligible institutions. The 2-year schools make up 27% of schools and enroll 33% of the students at 2-year non-profit, degree-granting Title IV-eligible institutions. Figure A-1 plots a map of the institutions included in the sample. The focus of this paper is on course supply in the sample of 4-year institutions.

The data date back to 1998, with the most dense coverage in the last decade. Figure A-2 plots the number of institutions for which course descriptions or course enrollment data are observed annually. Data availability grows over time. 311 of institutions in my sample have data first available in 2010 or earlier, and 470 have data first available in 2015 or earlier.

Data collection began in February 2020 and is ongoing. For real-time data collection, I scrape course offerings and enrollment data mid-semester, after each school's add-drop deadline, when course offerings and section enrollment have stabilized. Scraping mid-semester ensures that I capture data for schools that remove their course schedules later in the term to replace them with the next semester's schedule. Additionally, for many universities in my sample, I am able to access archived course data from periods before the school entered my sample. In these cases, enrollment and course offerings reflect end-of-semester values. Comparisons between mid-semester and end-of-semester data reveal that course quantities and total enrollment are nearly always identical.

To validate the course catalog data, I compare it with publicly available aggregated course enrollment details from IPEDS. Institutions report to IPEDS the total number of undergraduate credit hours completed. My course-level enrollment data can be aggregated to calculate a corresponding measure for schools in my sample. Figure A-3 shows a comparison

between the total credit hours in the course catalog dataset and those in IPEDS. I plot the share of undergraduate credits from IPEDS observed in the course catalog dataset, with each observation being an institution-term and observations spanning the period 2006-2018. The histogram suggests a strong alignment between the course catalog dataset and the data reported in IPEDS: I record between 90-105% of credits for 70% of terms in the course catalog data.

There are several reasons for potential discrepancies between the total credit hours reported in the course catalog data and IPEDS data. First, some courses offer a range of credit hours. In such cases, I assign the minimum number of credits (e.g. for a course offering 3-5 credits, I assign 3 credit hours). This approach may underestimate credit hours if students select higher values within the range. Second, in the benchmark figure, no exclusions are made based on the field name; only graduate courses are excluded based on credit hours. As a result, the total undergraduate credits might be overstated if some graduate courses are given course numbers typical of undergraduate courses, especially in professional degrees like medicine or law. Detailed restrictions on the field of study, which are discussed later, are implemented in my analysis. Third, the reporting of credit hours to IPEDS may differ from how a course's credits/units are represented in the course schedule. Finally, errors either in the construction of the course catalog sample or in the data provided to IPEDS may lead to discrepancies. When such discrepancies occur, I perform detailed quality checks at the institution-by-term level, which result in the complete exclusion from the course catalog dataset of a small number of schools with clearly anomalous enrollment data.

To alleviate concerns about the potential impact of the aforementioned sources of error on the reliability of the data, I illustrate in Figure A-4 that enrollment growth trends in the course catalog data consistently align with those reported in IPEDS. For each institution-term, I index total credits in the course catalog data as a percentage of total credits in 2018-19 and plot these indexed values against the corresponding values from IPEDS. When the points lie along the 45° line, the credits reported in the IPEDS data and the course catalog data are growing at nearly the same rate. This suggests that errors from the miscounting of credits for variable-credit courses or other errors do not affect within-institution comparisons of enrollment over time. The series are highly correlated (correlation coefficient 0.91), implying that variations in the data stem from genuine enrollment trends rather than any error in data collection or processing.

Substantial processing was required to convert the scraped course catalog and schedule data into a dataset suitable for analysis. The processing of course enrollment data is outlined in this section, while the processing of course description data is detailed in Appendix Section C.

In the analysis estimating course quantity elasticity, I limit the data to the main terms offered by each institution, which typically include a Fall and Spring semester or Fall, Winter, and Spring quarters. Independent study, internship, supervised research, thesis, study abroad, student teaching, private lessons, teaching assistantship courses are excluded due to their asynchronous nature. Often, "honors" sections of a course are assigned different course numbers (e.g., Econ 101 vs Econ 101H). I treat these instances as multiple sections of the same course. Additionally, I exclude sections with fewer than 5 students enrolled due to uncertainty about whether the course actually ran.[1]

Course levels (pre-undergraduate, lower, upper, graduate) are assigned according to the institution's numbering convention. Occasionally, the course schedule distinguishes between lower/upper/graduate courses, and in these cases, I defer to the course-specific designation.

Cross-listing occurs when a single class is listed under multiple fields or levels, but such instances are not always explicitly identified in the course schedule. To address this, I infer cross-listings using details from the course catalog data. Courses are classified as cross-listed if they share the same instructor, meeting days, times, location, course title, and section number. Each associated field and level (e.g., upper-level Economics) is credited with a portion of the cross-listed course. For example, if a course is listed as both Econ 101 and Business 101 with identical cross-listing identifiers, I split quantity "credit" for this course between Economics and Business based on enrollment. When enrollment totals are reported separately for Econ 101 and Business 101, I distributed credit in proportion to the number of students enrolled in each section. When only a single enrollment total is reported for the joint Econ/Business 101, I apportion both enrollment and course credit based on the relative enrollment in other courses within the same field-level cell.[2]

B Fields of study

I manually classify the names of 28,000 departments into 54 unique fields for the analysis. A given field may be described in a number of ways depending on the institution. For example, Math may be called "Math," "Mathematics," "College Math," etc. I manually classify each department name into one of 170 sub-fields (largely at the level of a 4-digit CIP code), which I then assign to one of 54 fields. The unit of analysis in this paper is typically a field,

[1] The overwhelming majority of the sections dropped are for courses in the Humanities and Arts; to the extent that I am erroneously dropping some small courses that actually ran, I am if anything understating course quantity inelasticity by removing these small sections.

[2] For example, if Econ 101 and Business 101 are lower-level courses and 100 students are enrolled in other lower-level Economics courses while 50 students are enrolled in other lower-level Business courses, I allocate 2/3 of the enrollment and course credit for Econ/Business 101 to the Economics department and 1/3 to the Business department.

although some analyses summarized at a more aggregate field category level. Table A-1 lists the sub-field to field mapping in my analysis.

For most of my analysis, I exclude fields that do not represent departments in the conventional sense and fields associated with professional degrees or skilled trades. A number of courses are offered by administrative units (e.g. "College of Humanities" or "Office of Academic Affairs") that do not correspond to a single field of study, are often difficult to classify, and likely are not offered through the same decision-making process as courses offered within a conventional department. I exclude such courses from all parts of the analysis.

I exclude courses associated with professional degrees, including those in Medicine, Law, Nursing, Pharmacy, and Architecture. While Medicine and Law courses are rarely offered at the undergraduate level, their course numbering often does not explicitly indicate graduate-level status, so I exclude all courses in departments classified as Medicine or Law. The exclusion of Nursing, Pharmacy, and Architecture courses is motivated by their limited responsiveness to labor market changes. These programs are often siloed within universities, making it structurally challenging for students to enter or leave these fields in response to shifting demand. Additionally, the regulated nature of careers in these fields means that taking a few courses is unlikely to open job opportunities, unlike fields such as Computer Science or Business. As such, the assumptions underlying my instrumental variables strategy do not hold for these professional programs, and I exclude them from my analysis. Finally, I exclude skilled trade programs, such as Beautician or Mechanic programs. Enrollment in these fields is minimal at the baccalaureate level, and there are often too few observations in the ACS to construct a reliable instrument for employment growth in occupations tied to these majors.

C Text data processing

Supplemental text data description

In addition to text data from course descriptions scraped from college and university course catalogs, I use text data from five different types of sources to study how course content aligns with various applications of students' learning. I measure alignment with current events using data from front page articles published in the New York Times, academic advancement using data from abstracts for academic journals, technological progress using text from patent applications, skill demand using text from job descriptions, and social justice from a combination of books related to activism and online materials published by organizations oriented towards social justice and civil liberties. I also use text data from the complete set of Wikipedia articles as a neutral corpus as a benchmark for the distribution

of words against which I can identify words that are highly distinctive of each application of student learning. I describe each of these data sources in greater detail in the sections below.

C.0.1 New York Times articles

I download the complete set of articles published by the New York Times (either in print or digitally) between 2010-2019 using the New York Times Developer API. For each article, I observe the headline and either an abstract for the article or a text snippet that contains the first few paragraphs of the article. I define a document by concatenating an article's headline and the snippet or abstract (depending on which is provided). The New York Times data contain 938 thousand articles, and articles on average contain 29 words. I make no restrictions on the section of the New York Times in which an article is posted, nor do I make restrictions on whether the article was published in print or online.

C.0.2 Academic journals abstracts

I construct a corpus of abstracts from academic articles downloaded from Elsevier's SCOPUS. Following Biasi and Ma (2022), I search for abstracts from academic journals that rank in the top 10 by H-index for each field during the period 2010-2019. When available, I download the abstracts of all articles published during this period for each journal. The resulting sample includes 155 thousand abstracts from 180 journals. The average document in this corpus contains 163 words.

The distribution of tokens in academic journals will in part reflect differences across fields in the use of academic journals for publishing research. Specifically, journals in the sciences publish more editions and more articles per edition than journals in the humanities and arts. Thus, when I construct word weights using these documents, the weights will be biased towards science-oriented words and phrases simply due to the composition of this corpus. For my analysis, I typically make comparisons within an institution-field pair over time or control for field fixed effects, which will absorb some bias inherent in the construction of the corpus.

C.0.3 Patents

I download patent text from the US Patent and Trademark Office covering the period 2010-2018. The resulting corpus includes the text of 2.5 million patents, which contain on average 250 tokens per document.

C.0.4 Job descriptions

Job description data come from a dataset collected by Lightcast (previously Burning Glass Technologies) that contains the near-universe of online job posts. The full set of job descriptions is quite massive, so I build the corpus of job descriptions using job descriptions from a sample of months during my period of analysis. In particular, I include all job descriptions from March and August 2010, 2012, 2014, 2016, and 2018. I restrict to job descriptions with a requirement that applicants have at least a college degree. The resulting corpus contains 2 million documents, which contain on average 162 words per document.

C.0.5 Writings related to social justice

This corpus features the text from the 'Issues" and "Policy Positions" pages from the websites of multiple organizations spanning topics in social justice: the ACLU, the American Association of Disabled People, Amnesty International, the Brennan Center, the Democratic Socialists of America, GLSEN, the NAACP, the National Organization of Woemn, Oxfam, Planned Parenthood, the Southern Poverty Law Center, the Sunrise Movement, and UNICEF. The corpus also includes the full texts of six prominent books that are listed among the top 25 activist-related books on Goodreads. Collectively, these sources provide insights into a spectrum of topics, from racial justice, prison abolition, and women's rights to climate change and a more general exploration of civil liberties.

C.0.6 Wikipedia articles

I downloaded the text of all English-language pages published on Wikipedia as of July 1, 2023 using the "Wikimedia dump service." The dataset contains the full text of all Wikipedia pages. I restrict to articles (e.g. filter out redirect pages and media). I process the raw article entries to exclude lists of references, links, metadata not included in the article, and section headers. The resulting corpus contains 3.8 million documents, which contain on average 183 words per document.

Text processing

I apply consistent pre-processing procedures to all the text corpora, including the course descriptions. These procedures involve removing all punctuation and numbers, converting all strings to lowercase, eliminating URLs, removing stopwords, and lemmatizing the text (i.e., transforming "regressions" to "regression").

However, my approach incorporates two non-standard pre-processing steps. First, I exclude "boilerplate" language from the text data. I am concerned about capturing phrases

that are overly common in a specific text but lack relevance to the essence of the content. For instance, many job descriptions include nearly identical non-discrimination clauses at the end. Including these texts in my analysis could mistakenly suggest that phrases like "gender," "sexual orientation," and "discrimination" are highly important tokens for job skill demand, even though their usage in job descriptions is unrelated to the skills demanded of the jobs. To handle boilerplate language, I exclude sentences that are identically repeated across numerous documents within a given corpus from my analysis. Specifically, if a particular sentence appears identically more than 10 times across all documents in a specific tranche of documents, it is removed during pre-processing.

Second, I create a dictionary with tokens of varying word length based on the co-occurrence of words in the Wikipedia corpus. The objective here is to distinguish common n-grams (e.g., "machine learning" or "regression analysis") from their component words. This procedure essentially allows for all possible n-grams but removes sparse tokens and n-grams that frequently co-occur due to being composed of common words, rather than representing a distinct concept. Specifically, I combine any two-word pair into a single token if the two words appear consecutively at least 500 times and if the co-occurrence of the two-word pair occurs for at least 4% of all instances of the less frequent word in the pair. For example, in the Wikipedia corpus, the word "machine" appears 59,799 times, and the word "learning" appears 37,991 times. The words "machine" and "learning" appear consecutively 1,583 times (4.1% of the time "learning" appears in the Wikipedia corpus). Consequently, I consider "machine learning" a token distinct from "machine" and "learning."

This approach allows for tokens of varying word lengths. For example, if the words "university" and "michigan" co-occur frequently enough ("of" is removed as a stopword), and the words "michigan" and "wolverine" co-occur with sufficient frequency, the phrase "university [of] michigan wolverine" would be included in the dictionary.[3]

Finally, to reduce the size of the dictionary and minimize the impact of words that are distinctive due to misspellings or unique to specific types of documents, I project all corpora onto a dictionary of tokens that appear at least 500 times in the complete Wikipedia text. As a result, the focus is on commonly recognized words rather than theme-specific jargon, which aids in drawing meaningful comparisons between different text data sources.[4]

Enrollment and course description data often come from different sources. In some in-

[3]Incidentally, the longest phrases counted as a single token are "church [of] jesus christ latter day" and "united nation[s] security council resolution."

[4]To illustrate, consider the frequent appearances of specific terms like a website URL or the name of a job board in job descriptions. Including these "jargony" terms in the analysis might yield the misleading impression that they are distinctive features of job-related language, when, in reality, they are simply artifacts of the source or format of the content.

stances, overlap between the enrollment data and the course description data is imperfect. For example, it is somewhat common for a new course to not have a course description in the course catalog during the first year it is offered. In instances where a course is continuously offered (enrollment is nonzero) but the course description appears inconsistently in the course catalog, I backfill from next term a course description is available. For continuously-offered courses, course descriptions change somewhat infrequently and rarely change substantively (see, for example, Figure A-11).

Details on TF-IDF weights

The TF-IDF of a word w in document $d_{i,s,t}$ is the product of Term Frequency (TF) and Inverse Document Frequency (IDF). The TF for a given token in a given document is equal to the number of times w occurs in $d_{i,s,t}$ ($c_{w,d_{i,s,t}}$), normalized by the token count of $d_{i,s,t}$:

$$TF(w, d_{i,s,t}) = \frac{c_{w,d_{i,s,t}}}{\sum_{w' \in W} c_{w',d_{i,s,t}}}$$

The IDF for a given token w measures the distinctiveness of w across all documents. In other words, $IDF(w)$ reflects how rare w is in the complete corpus (D) of field descriptions. The IDF for a given token w is calculated:

$$IDF(w) = log\left(\frac{\sum_{d \in D} \mathbb{I}(w \in d)}{||D||}\right)$$

The TF-IDF value applied to a token w in document $d_{i,s,t}$ is the product of the two values:

$$v_{i,s,t}(w) = TF\text{-}IDF(w, d_{i,s,t}) = TF(w, d_{i,s,t}) \times IDF(w)$$

I use the TF-IDF representations of field curricula to construct a series of measures of a field's changing curriculum. My preferred measure calculates the syntactic distance between a field's curriculum in 2018 relative to 2010. Let $v_{i,s,t}$ be the vector representation of the TF-IDF weights from course descriptions at institution i in field s in year t. I normalize each $v_{i,s,t}$ to have a magnitude of 1. Then, I calculate the cosine distance between the vector representation of the field's content:

$$dist_{i,s} = \frac{v_{i,s,2018} \cdot v_{i,s,2010}}{||v_{i,s,2018}||\ ||v_{i,s,2010}||}$$

Details on curriculum alignment scores

I construct curriculum alignment scores as a means to quantify the level of overlap between course content and specific applications of student learning. These scores are derived from a combination of TF-IDF weights and a "relevance weight" assigned to each token in a course description based on its importance to a particular application of student learning. The relevance weight aims to highlight tokens that are distinctive to a given application of student learning. While Term Frequency helps in identifying commonly appearing tokens in a corpus, it does not address the need to downweight tokens that are commonly used in general language. To overcome this, I create weights that compare a token's importance within a particular corpus (linked to a specific application of student learning) to its importance in a "neutral corpus," which consists of the complete text of Wikipedia articles. Tokens that are part of common language (e.g. "the," "a") should appear with similar frequency in any corpus. When a token appears significantly more often in a corpus related to an application of student learning than in the Wikipedia text, it is likely to be of greater importance to that specific application.

To calculate the relevance weights b_w^q for each token w with respect to each application of student learning q, I divide token w's share of all tokens in the q corpus (W^q) by token w's share of all tokens in the Wikipedia corpus (W^{Wiki}):

$$b_w^q = \frac{\frac{\sum_{w' \in W^q} \mathbb{I}(w' = w)}{||W^q||}}{\frac{\sum_{w' \in W^q} \mathbb{I}(w' = w)}{||W^q||} + \frac{\sum_{w' \in W^{Wiki}} \mathbb{I}(w' = w)}{||W^{Wiki}||}}$$

Tokens with relevance weights closer to 0.5 have similar frequencies in both the Wikipedia corpus and the application corpus. Tokens with higher relevance weights hold more significance in the application corpus compared to the Wikipedia corpus. Table 6 provides the relevance weights of some example tokens for reference.

I calculate the curriculum alignment score for each field's curriculum with each application of student learning by taking the TF-IDF-weighted sum of relevance weights specific to that application. To ensure consistent and interpretable scores, I normalize the weights in the TF-IDF vector representation of each field's curriculum, making sure they add up to 1. This normalization guarantees that each curriculum alignment score falls within the range of 0 to 1, providing a meaningful measure of alignment. Higher scores indicate a stronger connection between the curriculum and the intended student learning application, while lower scores imply less relevance between the two.

C.1 Sensitivity to temporal overlap of thematic documents

Two factors may be driving the general upward trend in curriculum alignment: courses may be increasingly incorporating topics related to the identified themes (or removing topics unrelated to them), as I argue in this section. Alternatively, the observed trend could be an artifact of the method itself, if the selection of reference texts biases the results toward more recent courses. In Appendix Figure A-13, I provide evidence against this second explanation. I replicate the analysis from Figure 7, splitting each thematic corpus into an "older period" (typically 2000-2010) and a "recent period" (typically 2011-2019).[5]

If the increasing alignment were merely a mechanical artifact of the method, then alignment trends would be steeper for course years overlapping with the thematic corpus period and flatter outside of it. This pattern does not emerge. Instead, curriculum alignment trends remain remarkably stable across different constructions of the relevance weights. This consistency reflects the relatively high correlation ($0.8 - 0.95$) between the relevance weights derived from older versus recent corpora, indicating that the observed changes are primarily driven by durable thematic associations (e.g., "elections" always being linked to current events) rather than transient trends.[6]

A notable exception is current events alignment, where alignment trends are uniformly steeper when based on recent news articles compared to older ones. This difference is largely attributable to changes in the composition of the New York Times corpus, which initially featured more sports and human interest stories among its top articles but gradually shifted toward a heavier emphasis on political and economic news. Encouragingly, the systematic difference in slope between older and more recent news-based alignment trends suggests that these compositional shifts do not alter the broader finding: college curricula have increasingly reflected contemporary themes over time.

D Discussion of the IV assumptions and robustness

A potential concern with my identification strategy is that local labor market conditions might induce direct institutional responses, rather than working through students' enrollment decisions. If, for example, a major employer or state government preemptively funds new faculty lines or facilities in a burgeoning field, the expansion in course offerings would be driven by university-level initiatives rather than by increased student demand. Under these

[5] I exclude social justice alignment from this exercise due to limited time variation in the corpus.

[6] That said, the method does capture some transitory shifts. For example, the token "Iraq War" had a relevance weight of 0.90 for current events between 2000-2010 but dropped to 0.35 between 2011-2019, while "solar power" became increasingly associated with technology, with its weight rising from 0.11 (2000-2010) to 0.69 (2011-2020).

circumstances, my shift-share strategy would violate the exclusion restriction, as it would affect course supply directly rather than only shifting student enrollment patterns.

However, there are several reasons to believe this type of violation is unlikely to be widespread in my setting. First, universities, particularly in four-year institutions, are constrained in their ability to proactively expand or contract course offerings due to structural features such as budget cycles, tenure constraints, and facility constraints. The descriptive evidence summarized in Figure 2 supports this hypothesis: course offerings often lag behind surges in enrollment — indicating that institutions respond to students' immediate course-taking decisions, rather than lead them —- while course offerings remain relatively stable in fields experiencing declining enrollment.

To further strengthen the case for the exclusion restriction, I discuss three potential channels by which it could be violated and provide evidence that none of these factors drive the main findings. First, I consider more seriously the claim that universities may better anticipate changing labor market conditions and respond to these conditions directly, rather than meeting growing student demand in these areas. If this were happening, we would expect to see instances where course quantity expands before demand for these courses manifests in higher enrollment. I test for this by calculating the share of courses in the two highest-growth fields — Computer Science and Engineering — that run at capacity between 2005-06 and 2018-19. In contrast with this story, I find that the share of courses offered at capacity is on a secular increase that pre-dates the surge in CS/Engineering enrollment and that the share of courses offered at capacity levels out late into the period.

A second violation of the exclusion restriction would be that the growing labor market opportunities associated with certain fields increase the cost of recruiting instructors in those fields, and therefore course quantity responses are constrained by labor market conditions. While plausible, I demonstrate that this does not drive my results by Winsorizing 5% on the value of my instrument, thereby excluding from the estimation some of the region-field pairs that would be most exposed to this possible exclusion restriction violation, and demonstrate that the estimates are not sensitive to these highly-exposed values. I compare the estimates from these regressions to the main results in Tables A-10 and A-11. Comparing Columns 1 and 2, and 4 and 5, the course and section quantity elasticity estimates are nearly identical when I exclude the kinds of fields for which this potential exclusion restriction concern may be most acute.

A third case where the exclusion restriction would be violated would occur when a local employer or donor, recognizing a skill gap in the local labor force, contributes to the creation or expansion of a field to build an employee base that fills this gap. Philanthropic or governmental funding specifically targeted at expanding academic departments is not uniformly

distributed across fields and institutions, and tends to be more prevalent in two-year colleges or professional/vocational programs, which are excluded from my analysis. My focus on academic departments in four-year institutions avoids consideration of fields and institutions that may be particularly susceptible to this exclusion restriction violation. Moreover, while universities are exposed to the influence of donors and (in the case of public universities) political leaders, it is not obvious that these influences are systematically correlated with labor market conditions. While this potential exclusion restriction is genuinely of concern and may occur at a small number of schools, there is not strong evidence of systemic violations of this kind among academic disciplines in four-year universities.

Tables A-10 and A-11 present a robustness test addressing potential endogeneity in the instrument. In Columns 3 and 6, I construct an alternative shift-share instrument that assigns each university the major-to-occupation shares from all Census divisions *except* its own. While endogeneity in the main specification is unlikely—given that any single university contributes minimally to a multi-state region—this alternative approach further mitigates concerns by ensuring that a university's course supply cannot meaningfully influence the composition of the national labor force outside its own Census division.

The IV estimates from this exercise closely align with the main results. As in the primary analysis, the IV estimates are generally lower than the OLS estimates, reinforcing the conclusion that universities adjust course offerings at rates well below one-for-one with changes in student demand. Notably, for fields experiencing growing demand, the IV estimates are slightly higher than in the main specification. This likely reflects the fact that the alternative instrument does not fully capture the strength of regional pipelines into high-growth fields, leading to a slightly weaker first stage that predicts smaller demand-driven growth. As a result, the estimated course supply response appears somewhat more proportional to changing demand. However, these differences are minor and do not undermine the broader robustness demonstrated by this exercise.

E Counterfactual major completions without course rationing

I conduct a counterfactual exercise using the OLS and IV estimates from Tables 2 and 3 to project enrollments and completed majors under a scenario where students are neither rationed out of preferred courses nor diverted into less-preferred alternatives. In this scenario, universities respond to student demand with the elasticity implied by the OLS estimates, including unobserved demand stemming from rationing and diversion. To generate these projections, I adjust course supply upward based on the bias suggested by the OLS estimates. I then inflate enrollment at the institution-field level by applying the ratio of added

(or removed) courses to average upper-level course enrollment in the base year (2009–10). To convert these enrollment changes into completed majors, I scale the institution-field enrollment projections by the average number of majors per upper-level course enrollment, calculated using IPEDS data for major completions.

I present estimates based on both course and section offering adjustments, with my preferred specification relying on changes in section offerings. From a policy perspective, modifying section availability represents a more flexible and cost-effective approach to alleviating course rationing compared to adding entirely new courses. Figure A-10 displays the median and interquartile range of projected changes in major completions for a selection of relevant fields.

The simulation results suggest that, in a world where universities respond to student demand as elastically as implied by the OLS estimates — but without rationing or diversion — U.S. institutions would have produced notably more majors in Computer Science and Engineering, while producing fewer in the Humanities and Education. Under this counterfactual, Computer Science major completions would have been 3.9% higher, and Engineering major completions 3.2% higher. Conversely, fields like History would have experienced declines, with major completions decreasing by 6.0%. It is important to note that these estimates likely represent upper-bound reductions in low-demand fields; tenure constraints and institutional inertia can limit the capacity to reduce faculty headcount, making it relatively costless to continue offering such courses.

The inelasticity of course supply creates both winners and losers within the university. Students in high-demand fields face rationing or overcrowded classes, restricting their access to preferred courses. In contrast, students pursuing low-demand fields benefit from a broader array of course options that might shrink under a more elastic supply regime, along with smaller class sizes that may enhance their learning environment.

F Curriculum alignment decomposition

Following Foster et al. (2001), I decompose the total change in average curriculum alignment over the ten-year period 2012-13 and 2022-23 into changes resulting from entry, exit, within, and between. Within course changes measure the contribution from changing course descriptions for courses offered continuously over this period. Between course changes measure the contribution from changing student enrollment across continuously offered courses but within the same field of study. Exit measures the contribution from courses that were offered in 2012-13 but were not offered in 2022-23. Entry measures the contribution of courses that were offered in 2022-23 but were not offered in 2012-13.

The decomposition proceeds as follows: For each institution i and field s, let $S_{i,s}$ be the set of courses offered continuously between 2012-13 and 2022-23, $E_{i,s}$ be the set of courses offered in 2022-23 but not offered in 2012-13 or earlier, and $X_{i,s}$ be the set of courses offered in 2012-13 but discontinued before 2022-23. I denote a course belonging to any of these groups as x (for simplicity, I will omit the i and s subscripts when referring to a course).

Let $S_{i,s,t}$ be the set of courses offered at institution i in field s during year $t \in \{1, 2\}$, and denote courses by $x \in S_{i,s,t}$ (for simplicity, I will omit the i and s subscripts when I refer to a course). Let $s_{x,t}$ be course x's share of enrollment at institution i in field s, and $\varphi^q_{x,t}$ be course x's curriculum alignment to theme q. Finally, let $\Phi^q_{i,s,t}$ be the average curriculum alignment to theme q across all courses at institution i in field s during term t. The decomposition proceeds as follows:

$$\Delta\Phi^q_{i,s} = \underbrace{\sum_{x\in S} s_{x1}\left(\varphi^q_{x2} - \varphi^q_{x1}\right)}_{\text{within}} + \underbrace{\sum_{x\in S}\left(s_{x2} - s_{x1}\right)\left(\varphi^q_{x1} - \Phi^q_1\right) + \sum_{x\in S}\left(s_{x2} - s_{x1}\right)\left(\varphi^q_{x2} - \varphi^q_{x1}\right)}_{\text{between}}$$
$$+ \underbrace{\sum_{x\in E} s_{x2}\left(\varphi^q_{x2} - \Phi^q_1\right)}_{\text{entry}} - \underbrace{\sum_{x\in X} s_{x1}\left(\varphi^q_{x1} - \Phi^q_1\right)}_{\text{exit}}$$

Having computed the components $\Phi^q_{i,s}$ for each institution-field pair, I aggregate first up to the institution level, then average across institutions for the final decomposition. I aggregate up to the institution level as the average of each component of $\Phi^q_{i,s}$ weighted by each field s's share of enrollment in the base period. This gives me the components of Φ^q_i for each school. I average each of the components across schools to produce the values plotted in Figure 8. Note that the analysis restricts the sample of institutions to the subset for which I observe course offerings continuously from 2012-13 to 2022-23. As a consequence, the 10-year change in curriculum alignments plotted in Figure 7 differs slightly from the changes plotted in Figure 8.

G Supplemental figures and tables

Figure A-1. Geographic coverage of the course catalog dataset

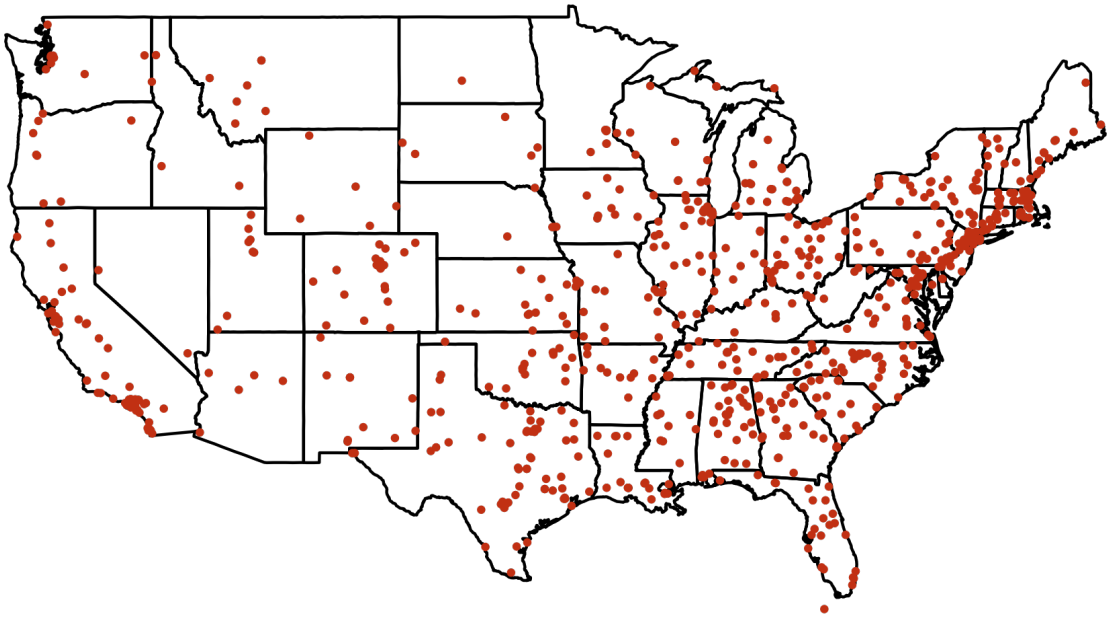
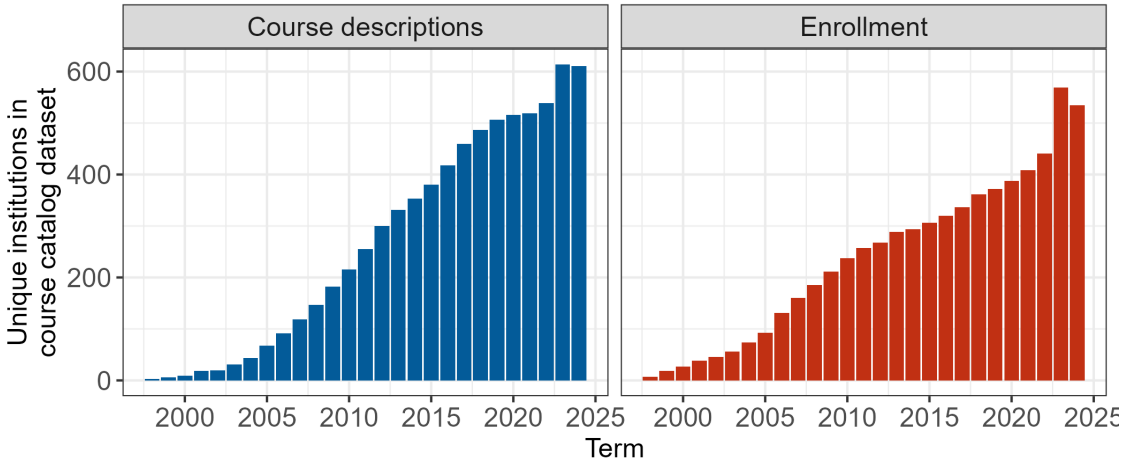
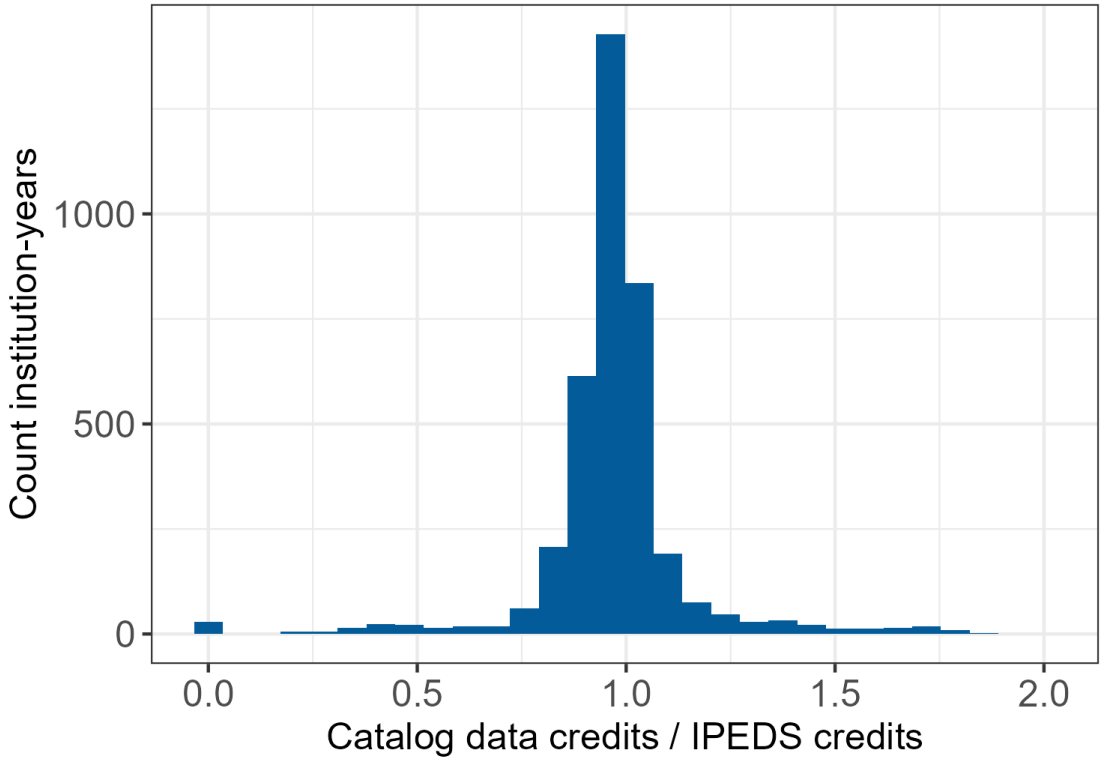


Figure A-2. Annual coverage of course catalog data



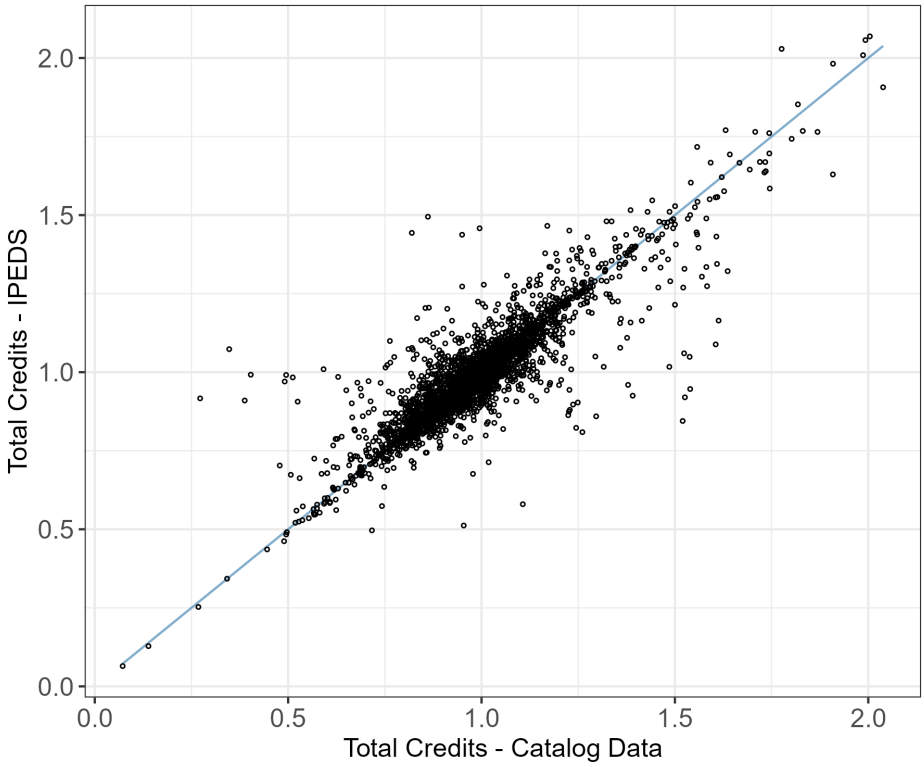
Notes: The figure counts the number of institutions in the course catalog dataset by year. The left panel counts the number of institutions with course description data; the right panel counts the number of institutions with enrollment data. For many institutions, the data record both enrollment and course descriptions.

Figure A-3. Compare total credits in catalog data to IPEDS



Notes: The figure compares total credits for enrollment in undergraduate courses in the course catalog data to total undergraduate credits reported in IPEDS. Observations divide total credits in the course catalog data by credits reported in IPEDS at the institution-year level.

Figure A-4. Compare indexed credit growth rates in catalog data to IPEDS

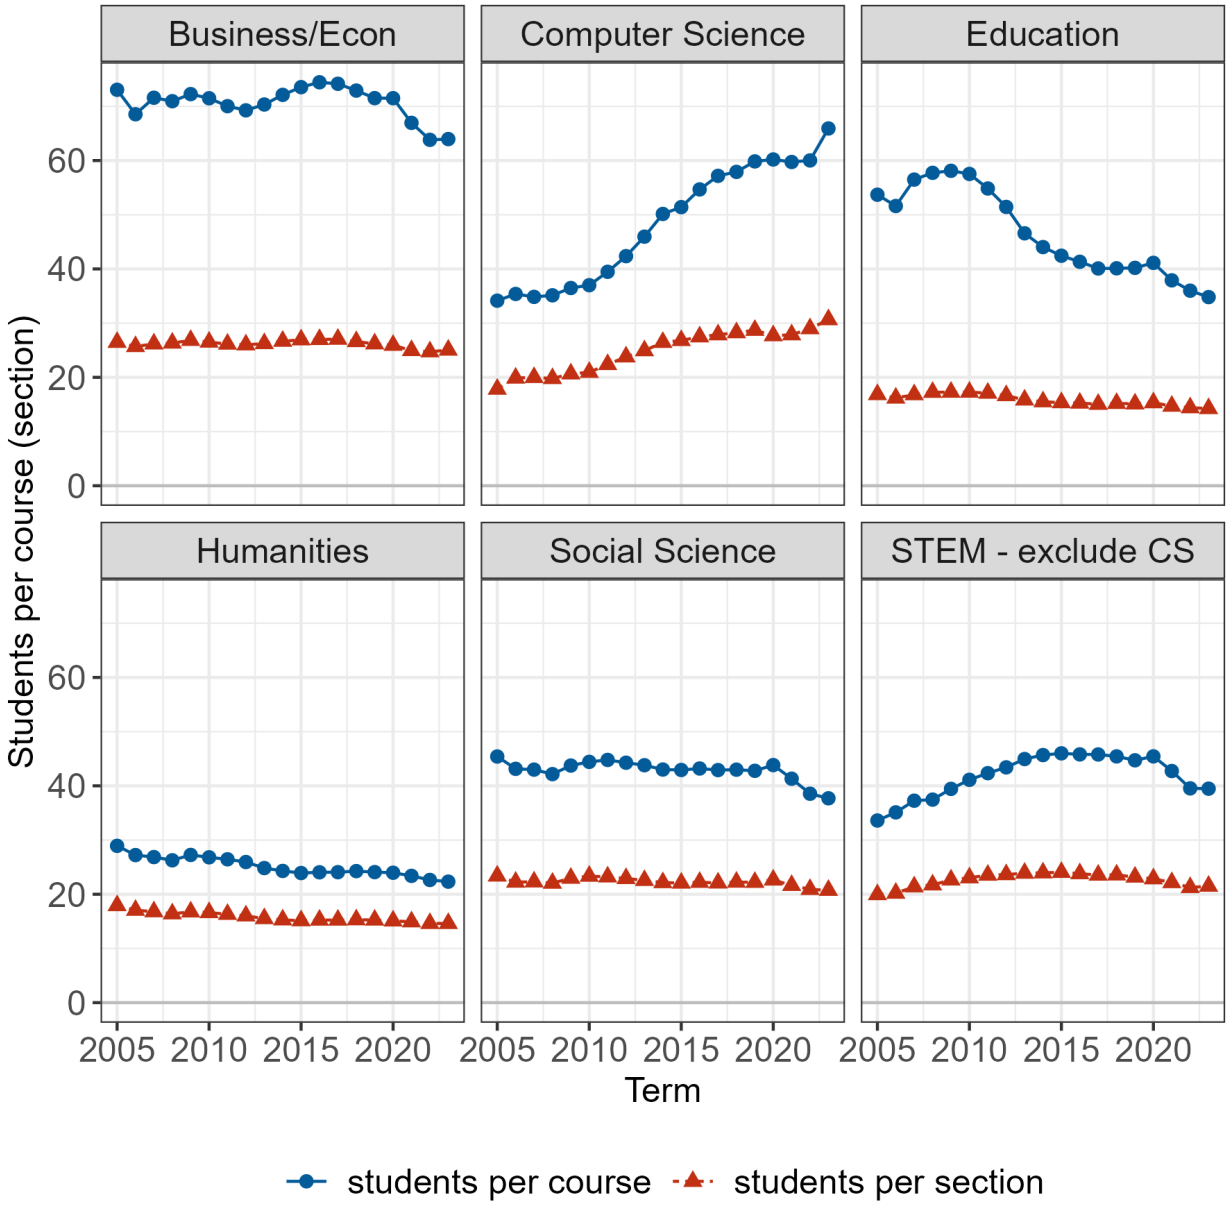


Notes: The figure compares the growth in total credits for enrollment in undergraduate courses in the course catalog data to the growth in total undergraduate credits reported in IPEDS. Observations are at the institution-year level. Catalog credits are indexed as the percent change relative to undergraduate credits completed in 2018-19; IPEDS credits are indexed as the percent change relative to undergraduate credits completed in 2018-19 in IPEDS. Because it is used as the index, enrollment in 2018-19 is omitted from the plot. The plotted line is the 45° line. The correlation coefficient between the two series is 0.94.

Table A-1. Field classification

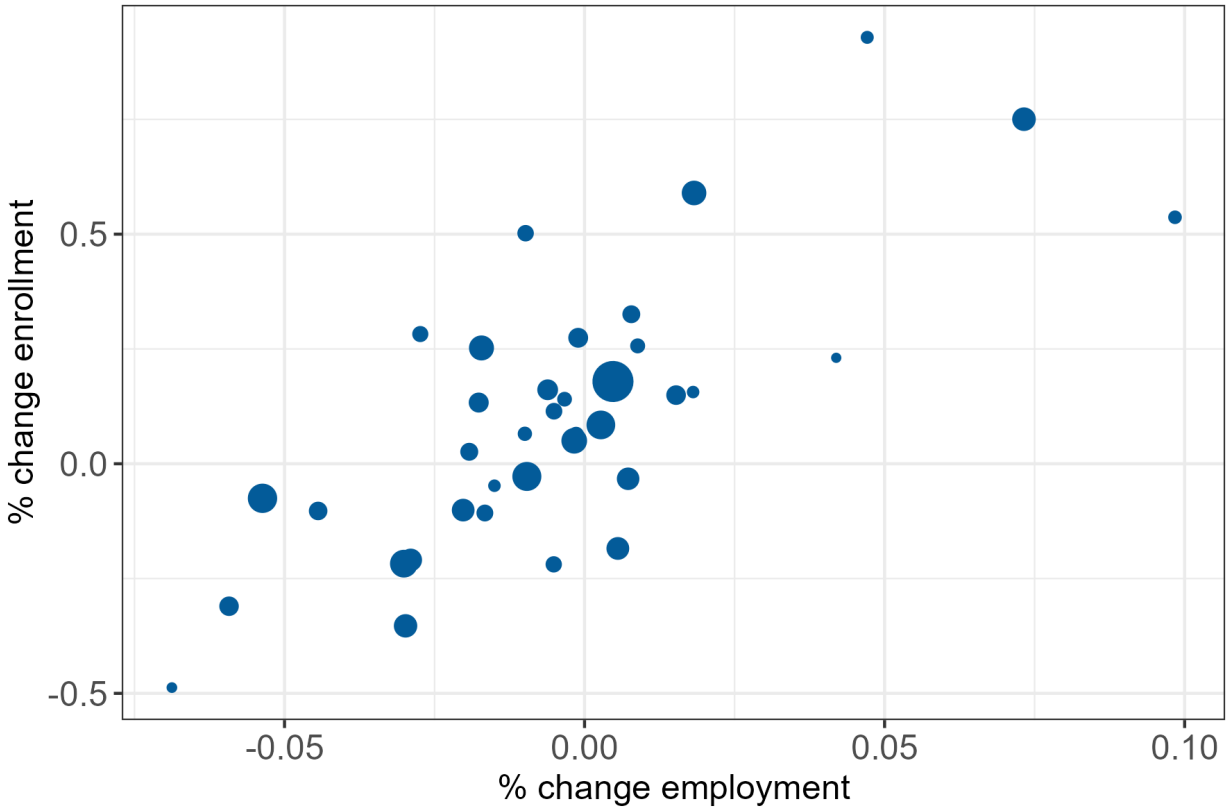
Field category	Field	Sub-field	Field category	Field	Sub-field	Field category	Field	Sub-field	Field category	Field	Sub-field
Business	Business	Accounting	**Social Science**	Communication	Advertising	**STEM - exclude CS**	Agriculture	Agriculture	**Skilled Trade**	Architecture	Architecture
		Business Administration			Communication			Agriculture Economics		Arts	Animation/Game Design
		Business Math			Journalism			Animal Science			Graphic Design
		Finance			Media Studies			Botany		Hospitality	Hospitality
		Leadership		Community Studies	Community Studies			Food Science		Medicine	EMT
		Management		Consumer Science	Consumer Studies			Horticulture		Other	CAD
		Marketing		Criminal Justice	Criminal Justice			Plant Science		Physical Education	Kinesiology
		Operations		Ethnic/Cultural Studies	American Studies		Biology	Biochemistry		Rehabilitation	Rehabilitation
		Organization Studies			Ethnic/Cultural Studies			Biology		Security Studies	Security Studies
		Real Estate			Gender Studies			Cognitive Science		Skilled Trade	Automotives
		Statistics - Business		International Studies	International Studies			Neuroscience			Aviation
	Consumer Science	Decision Science		Law	Law		Chemistry	Chemistry			Construction
	Economics	Economics		Political Science	Political Science		Engineering	Aerospace			HVAC
	Human Resources	Human Resources		Psychology	Counseling			Bioengineering			Manufacturing
	Math	Risk Management			Psychology			Chemical Engineering			Skilled Trade
	Other	Admin		Public Policy	Public Administration			Civil Engineering			Tax
Humanities	Anthropology	Anthropology			Public Policy			Engineering			Vocational
		Archeology		Security Studies	Peace Studies			Industrial Engineering	**Other**	Audiology	Audiology
	Arts	Art		Social Science	Social Science			Mechanical Engineering		ESL	ESL
		Art History			Social Science - Other			Nuclear Engineering		Other	Adult Learning
		Dance			Social Studies			Systems Engineering			Apprenticeship
		Film		Social Work	Social Work			Technology - Other			Cannabis
		Music		Sociology	Sociology		Environmental Studies	Energy Science			General Studies
		Theater			Urban Studies			Environmental Engineering			Graduate
	Consumer Science	Fashion		Urban Planning	Urban Planning			Environmental Studies			Military
		Human Development	**Education**	Education	Early Childhood Education			Forestry			Other
	English	English			Education			Natural Resources			Professional Development
		Literature			Elementary Education			Naval Studies			Remedial
		Writing			Higher Education		Health	Health			Student Affairs
	History	History			Secondary Education			Nutrition			Study Abroad
	Humanities	Classics			Special Education		Math	Math			Univeristy
		Humanities			Teaching		Medicine	Allied Health			University
	Language	Asian Languages	**STEM - CS**	Computer Science	Computer Science			Dentistry			University - Other
		Asian Studies			Computer Science - Other			Medicine			Wine
		Germanic Languages			Electrical Engineering			Optometry		Physical Education	Physical Education
		Language - Other			Informatics			Physiology			Recreation
		Mideast Languages			IT		Nursing	Nursing			Sports
		Romance Languages			Statistics - CS		Other	Mining		Rehabilitation	Occupational Therapy
		Slavic Languages						Oil and Gas			
	Library	Library					Pharmacy	Pharmacy			
		Library Science					Physics	Physics			
	Linguistics	Linguistics					Public Health	Public Health			
	Other	Museum Studies					Science - Other	Astronomy			
	Philosophy	Philosophy						Earth Science			
	Religion	Religion						Science - Other			
	Social Science	Geography					Stats/Data Science	Data Science			
								Statistics			

Figure A-5. Change in average course size



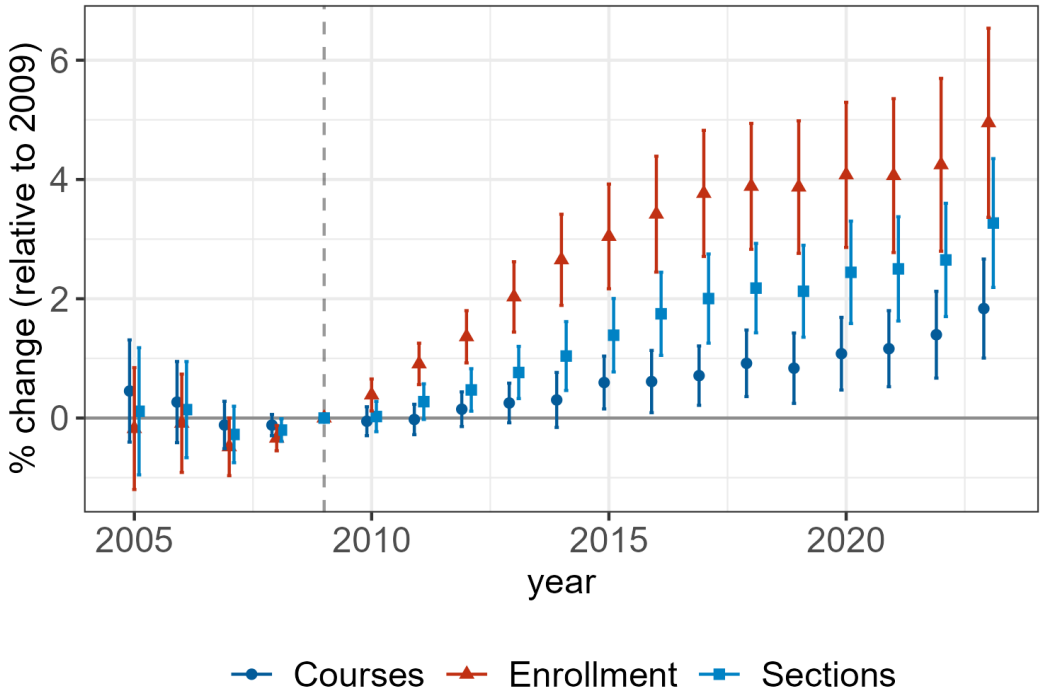
Notes: This figure plots the trend in average number of students per course and section between for six aggregated field categories. For each institution, I calculate the average number of students per course and section in each of the field categories by year. The figure plots the average of these values across institutions.

Figure A-6. First-stage monotonicity



Notes: This figure compares relative enrollment growth by field against the shift-share instrument capturing region-by-field employment growth. Points are at the field level. For each institution and field, I calculate the relative enrollment growth from 2009-10 to 2017-18, relative to the overall enrollment growth at the institution. For each field, I then calculate the average relative enrollment growth and plot it against the average value of the shift-share instrument. Each point in the figure is proportional in size to the sum of the weight each field receives in the regression analysis, where weights are allocated by field within each institution in proportion to 2010-11 enrollment and each institution receives total weight equal to 1.

Figure A-7. Test for quantity and enrollment anticipation of labor market changes



Notes: The figure plots point estimates from separate regressions of log enrollment, course quantity, and section quantity on the value of the shift-share instrument. Observations are at the institution-field-year level. Regression controls for institution-by-year fixed effects. Standard errors are clustered at the institution level.

Table A-2. First-stage estimates

	Catalog data		IPEDS
	All undergraduate	Upper-level	Completed degrees
	(1)	(2)	(3)
Constant	0.025**	0.034**	0.011
	(0.010)	(0.013)	(0.016)
% enrollment change - overall	-0.017	-0.018	-0.036
	(0.027)	(0.037)	(0.057)
Employment change	2.701***	4.771***	4.287***
	(0.294)	(0.419)	(0.524)
F-stat	84	129	67
Observations	4,059	3,647	3,656
R^2	0.063	0.118	0.072

Notes: Observations are at the institution-by-field level. I regress the log change in enrollment on the shift-share instrument reflecting major-typical employment growth in the Census division where the institution is located. Columns 1-2 measure changing enrollment using the course catalog data; Column 3 measures changing enrollment using completed degrees data from IPEDS. In the regression, each institution is uniformly weighted. Within an institution, subjects receive weights proportional to the start-of-period enrollment. In all columns, standard errors are clustered at the institution and Census division-by-field level.

Table A-3. F-stats for alternative instruments

	Lag length							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All periods	3.32	9.27	14.43	19.15	35.23	47.50	63.14	78.62
Only period ending 2018-19	0.28	10.41	14.31	17.00	31.74	37.01	53.23	129.41

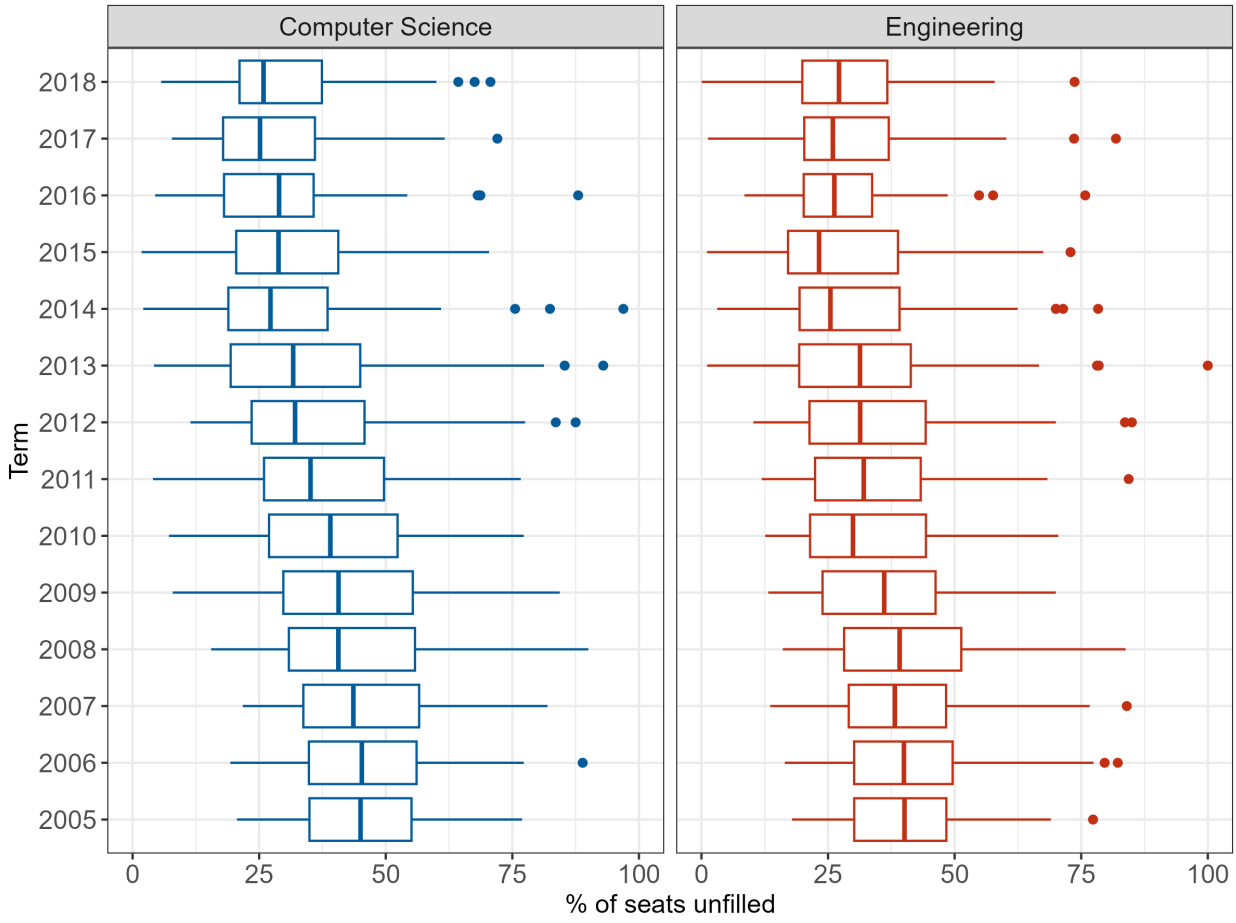
Notes: This table presents the first-stage F-statistics of log enrollment changes on the employment growth instrument calculated over various periods. "Lag length" refers to the duration (in years) over which these changes are calculated. The table includes F-statistics for three specific time windows: a single period ending in the 2018-19 school year; overlapping periods from 2009 to 2022 (e.g., a 4-year lag would cover intervals like 2009-2013, 2010-2014, etc.); and non-overlapping periods with specified lag lengths from 2009 to 2022 (e.g., a 4-year lag would cover intervals like 2009-2013, 2014-2018, etc.).

Table A-4. Reduced form

	# of Courses		# of Sections	
	(1)	(2)	(3)	(4)
% enrollment change - overall	0.3098	0.3143	0.6351	0.6387
	(0.0513)	(0.0522)	(0.0403)	(0.0410)
Employment change	0.9929		2.532	
	(0.2639)		(0.2979)	
Employment change - growing		2.047		3.380
		(0.3605)		(0.4341)
Employment change - shrinking		0.3430		2.009
		(0.3109)		(0.4092)
Observations	3,756	3,756	3,756	3,756
R^2	0.05	0.06	0.19	0.19

Notes: Observations are at the institution-by-field level. The analysis regresses change in upper-level course quantity on the institution average change in enrollment, represented as log differences from 2010-11 to 2018-19, and the shift-share instrument capturing region-by-field variation in changing occupation growth from 2010 to 2018. Quantity and enrollment are credit hour-weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment.

Figure A-8. Testing for leads in Computer Science course quantity change



Notes: The figure plots the distributions of the share of Computer Science seats unfilled by year. For each school, I calculate the share of seats unfilled based on total enrollment in Computer Science courses and the listed capacity for these courses.

Table A-5. Course quantity elasticity regression, all undergraduate courses

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change - overall	0.247	0.286	0.281	0.312	0.340	0.325	0.353	0.352
	(0.029)	(0.070)	(0.024)	(0.038)	(0.025)	(0.025)	(0.031)	(0.024)
% enrollment change - field	0.208	0.202	0.318	0.325	0.397	0.401	0.391	0.244
	(0.011)	(0.014)	(0.013)	(0.014)	(0.014)	(0.014)	(0.020)	(0.073)
First Stage F-stat								80
Observations	92,632	44,334	76,466	19,920	48,345	12,783	4,234	4,234
R^2	0.067	0.067	0.160	0.169	0.265	0.263	0.284	0.255

Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, standard errors are clustered at the institution and field-by-Census division level, which is the level of variation for the instrument.

Table A-6. Section quantity elasticity regression, all undergraduate courses

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change - overall	0.472	0.521	0.605	0.687	0.714	0.718	0.721	0.720
	(0.040)	(0.046)	(0.030)	(0.019)	(0.031)	(0.022)	(0.030)	(0.020)
% enrollment change - field	0.358	0.329	0.554	0.561	0.667	0.670	0.676	0.535
	(0.012)	(0.016)	(0.009)	(0.014)	(0.010)	(0.012)	(0.022)	(0.050)
First Stage F-stat								80
Observations	92,632	44,334	76,466	19,920	48,345	12,783	4,234	4,234
R^2	0.174	0.161	0.404	0.438	0.595	0.597	0.630	0.613

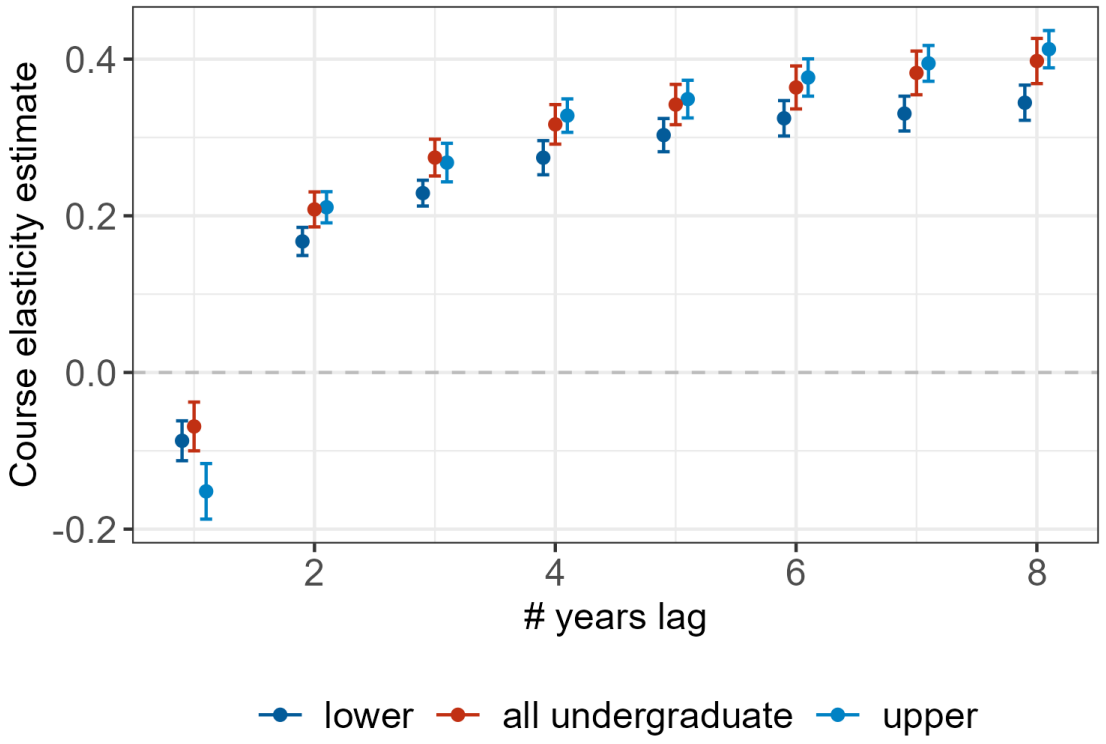
Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, standard errors are clustered at the institution and field-by-Census division level, which is the level of variation for the instrument.

Table A-7. Robustness on course quantity elasticities using different offsets and lags

	Lag length							
offset	1	2	3	4	5	6	7	8
0	0.56	0.52	0.51	0.50	0.50	0.50	0.51	0.51
	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
1	-0.15	0.21	0.27	0.33	0.35	0.37	0.39	0.41
	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
2	0.05	-0.05	0.17	0.23	0.29	0.32	0.35	0.36
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
3	-0.02	0.03	-0.04	0.14	0.20	0.25	0.28	0.31
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
4	0.02	0.01	0.04	-0.01	0.13	0.18	0.23	0.26
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)

Notes: This table summarizes OLS estimates of course quantity elasticity using different lag lengths and offsets. Lag length refers to the length of time in years over which percent changes in enrollment and course quantity are calculated. Offset refers to the lag imposed when regressing changes in course quantity on earlier changes in enrollment. Estimates come from regressions of change in course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Standard errors are clustered at the institution-by-year level.

Figure A-9. Growth in course quantity elasticity with longer lags



Notes: The figure plots OLS estimates of course quantity elasticity estimated using different "lag lengths" of change in course quantity and enrollment. Lag length refers to the number of years over which changes in enrollment and course quantity are calculated. Enrollment and course quantity changes are offset by 1 year. Estimates are plotted separately for lower-level courses, upper-level courses, and all undergraduate courses. Standard errors are clustered at the institution-year level.

Table A-8. Alternative course elasticity specification - no controls

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change	0.210	0.206	0.326	0.336	0.411	0.422	0.401	0.244
	(0.010)	(0.011)	(0.010)	(0.014)	(0.012)	(0.017)	(0.026)	(0.050)
First Stage F-stat								112.3
Observations	81,845	39,146	67,578	17,566	42,458	11,218	3,756	3,756
R^2	0.050	0.046	0.140	0.141	0.243	0.252	0.268	0.227

Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, standard errors are clustered at the institution and field-by-Census division level, which is the level of variation for the instrument.

Table A-9. Alternative section elasticity specification - no controls

	2-year diffs (1998-2022)		4-year diffs (1998-2022)		8-year diffs (1998-2022)		Single 8-year diff (2009-2018)	
	Rolling	Staggered	Rolling	Staggered	Rolling	Staggered	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% enrollment change	0.316	0.315	0.500	0.524	0.612	0.628	0.618	0.609
	(0.012)	(0.013)	(0.011)	(0.015)	(0.010)	(0.015)	(0.026)	(0.052)
First Stage F-stat								112.3
Observations	81,845	39,146	67,578	17,566	42,458	11,218	3,756	3,756
R^2	0.091	0.086	0.239	0.253	0.364	0.383	0.429	0.429

Notes: Observations are at the institution-field-period level, where a period is a pair of comparison years differenced to measure percent changes. The analysis regresses change in upper-level course quantity on change in enrollment, each represented as long log differences. Quantity and enrollment are credit-hour weighted. Each institution receives equal weight; within each institution, fields are weighted by start-of-period enrollment. Columns 1-2 estimate elasticities using two-year differences; Columns 3-4 estimate elasticities using four-year differences; Columns 5-8 estimate elasticities using eight-year differences. Columns 1, 3, and 5 use overlapping periods (e.g. 2010-2014, 2011-2015); all other columns use adjacent periods or only a single period. In Columns 1-7, standard errors are clustered at the institution-by-period level; in Column 8, standard errors are clustered at the institution and field-by-Census division level, which is the level of variation for the instrument.

Table A-10. Course quantity IV robustness to alternative specifications

	Base (1)	Winsorized (2)	Alternative IV (3)	Base (4)	Winsorized (5)	Alternative IV (6)
% enrollment change - overall	0.316	0.307	0.312	0.320	0.311	0.316
	(0.050)	(0.048)	(0.048)	(0.050)	(0.049)	(0.049)
% enrollment change - field	0.217	0.221	0.333			
	(0.052)	(0.065)	(0.057)			
% enrollment change - growing				0.368	0.379	0.451
				(0.062)	(0.103)	(0.078)
% enrollment change - shrinking				0.082	0.113	0.165
				(0.079)	(0.094)	(0.076)
First stage F-stat	113.7	76.1	104.5	64	41.1	60.3
p-value grow = shrink				0.004	0.049	0.008
Observations	3,756	3,585	3,756	3,756	3,585	3,756
R^2	0.267	0.270	0.316	0.069	0.065	0.077

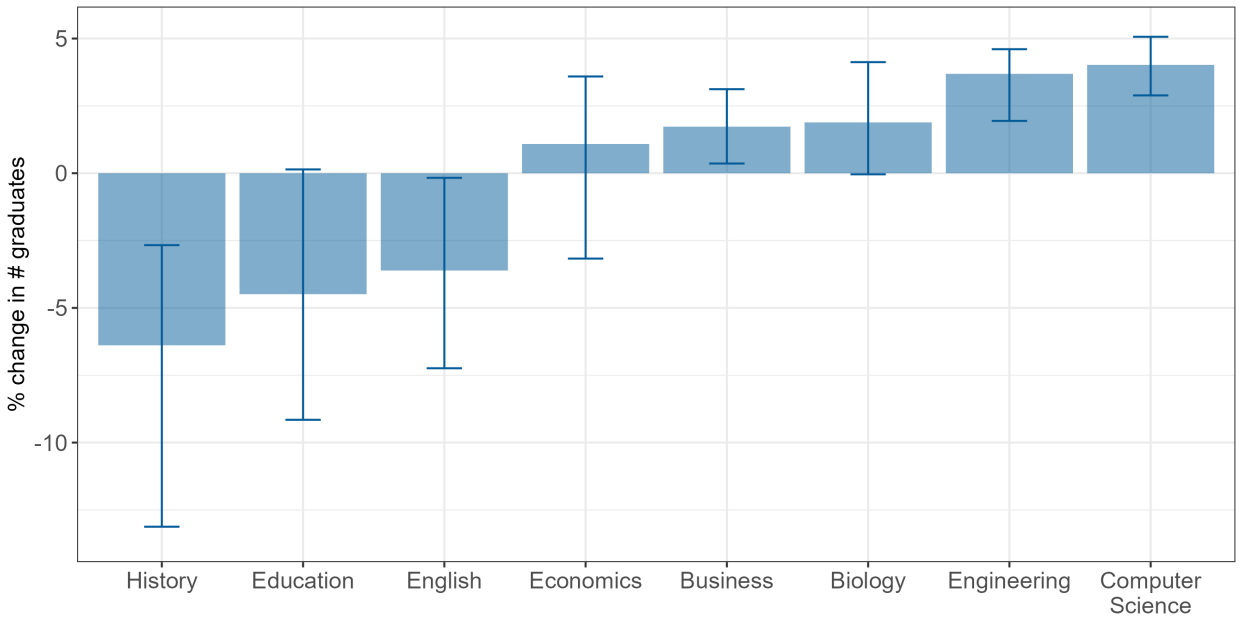
Notes: Columns 1 and 4 contain estimates from the main IV specification presented in the paper. Columns 2 and 5 drop the 2.5% of region-fields with the highest and lowest relative employment growth. Columns 3 and 6 estimate the IV with an alternative instrument based on field-level employment occupation growth in all Census divisions excluding the Census division in which the university is located. For all regressions, standard errors are clustered at the region-by-field level. For the estimates in Columns 4-6, bootstrapped standard errors are calculated using $1,000$ repetitions of the estimation, resampling region-by-field clusters in each iteration. Significance stars are suppressed, as the natural comparison is not obviously 0.

Table A-11. Section quantity IV robustness to alternative specifications

	Base (1)	Winsorized (2)	Alternative IV (3)	Base (4)	Winsorized (5)	Alternative IV (6)
% enrollment change - overall	0.651	0.649	0.651	0.652	0.651	0.652
	(0.032)	(0.032)	(0.031)	(0.032)	(0.032)	(0.032)
% enrollment change - field	0.552	0.550	0.616			
	(0.038)	(0.048)	(0.048)			
% enrollment change - growing				0.603	0.601	0.656
				(0.062)	(0.102)	(0.076)
% enrollment change - shrinking				0.507	0.514	0.560
				(0.062)	(0.075)	(0.057)
First stage F-stat	113.7	76.4	104.8	64	41.2	61.2
p-value grow = shrink				0.273	0.485	0.315
Observations	3,756	3,585	3,756	3,756	3,585	3,756
R^2	0.581	0.581	0.585	0.199	0.192	0.209

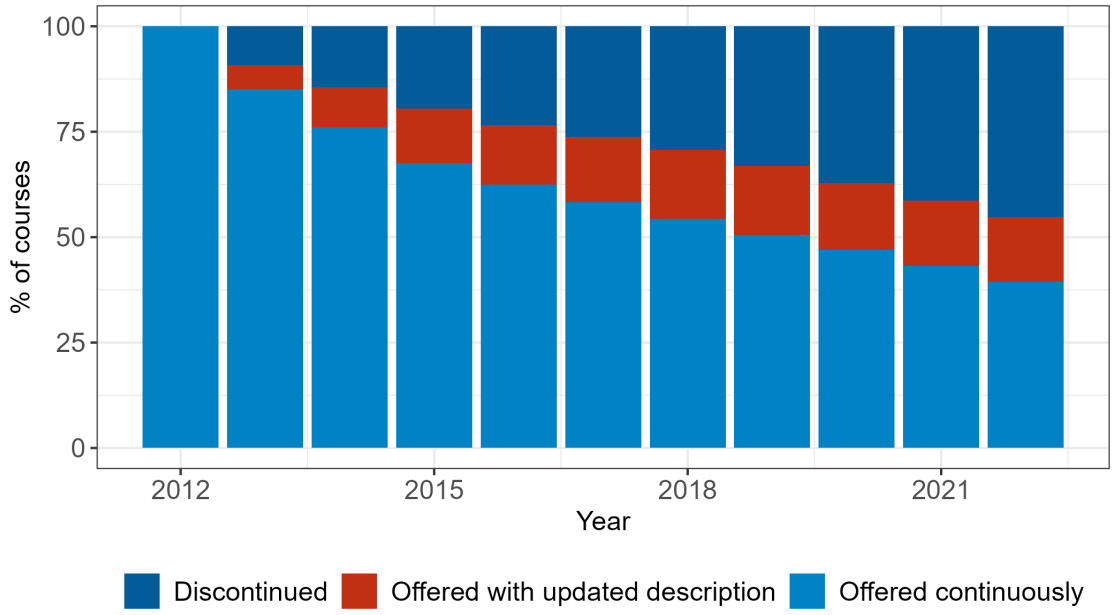
Notes: Columns 1 and 4 contain estimates from the main IV specification presented in the paper. Columns 2 and 5 drop the 2.5% of region-fields with the highest and lowest relative employment growth. Columns 3 and 6 estimate the IV with an alternative instrument based on field-level employment occupation growth in all Census divisions excluding the Census division in which the university is located. For all regressions, standard errors are clustered at the region-by-field level. For the estimates in Columns 4-6, bootstrapped standard errors are calculated using $1,000$ repetitions of the estimation, resampling region-by-field clusters in each iteration. Significance stars are suppressed, as the natural comparison is not obviously 0.

Figure A-10. Counterfactual major completions without course rationing



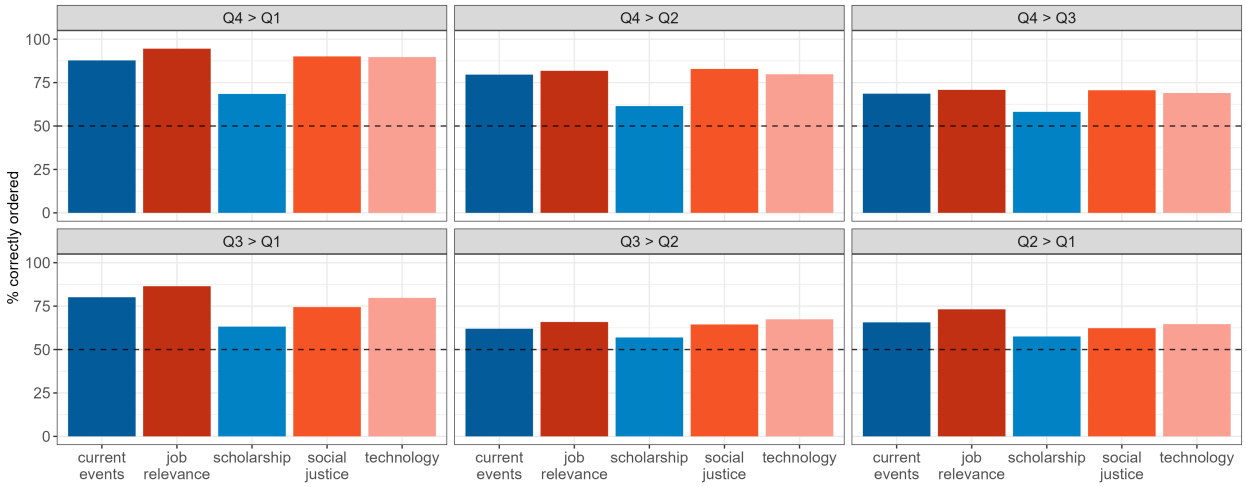
Notes: Figure plots estimates of counterfactual change in completed majors in the absence of rationing of seats in high demand sections. Values are estimated by estimating the unobserved percent change in demand for each institution-field pair between 2010-11 and 2018-19, using the difference in the section quantity point estimates between the IV and OLS specifications in Table 4. I translate the percentage change in demand into credit hours, then estimate the corresponding change in majors by multiplying by the 2010-11 ratio of number of completed majors (from IPEDS) per credit hour. The figure plots the average and inter-quartile range of estimated change in completed majors, calculated within field and across institutions, for a sample of fields.

Figure A-11. Survival of courses offered in 2010-11



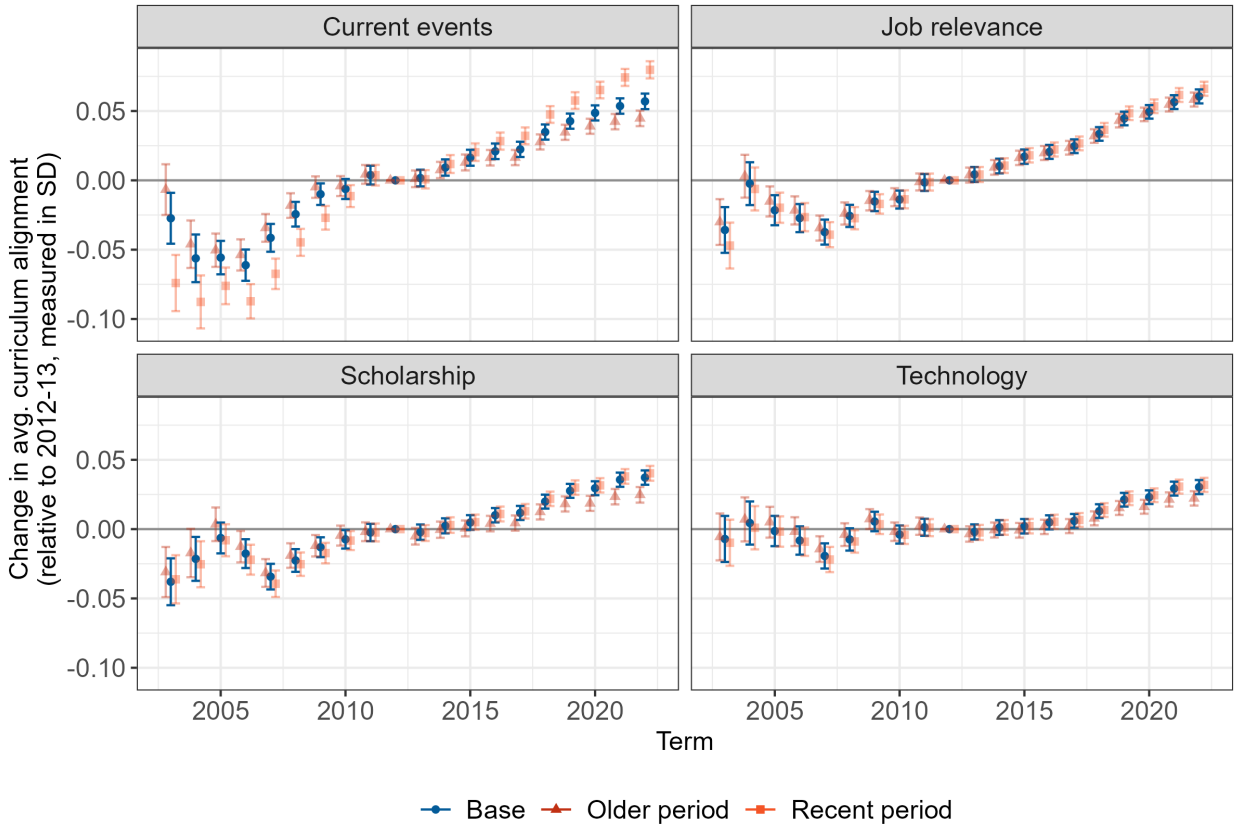
Notes: Figure plots the survival path of courses offered in 2010-2011. In each year, the course can occupy one of three states; a course is "Discontinued" if it is offered in a given year but never offered subsequently, a course is "Offered with updated description" if the course is offered in a given year but with a course description that does not match its description in 2010-11. A course is "Offered continuously" if it is not discontinued or offered with updated description. The figure cuts off in 2019-2020 to ensure that courses are not erroneously counted as discontinued when they are in fact offered infrequently. Each course receives equal weight in this analysis.

Figure A-12. Validation of curriculum alignment scores using ChatGPT



Notes: This figure plots results from a validation exercise testing how the curriculum alignment measures compare to "manual" review. For each theme, 500 four-course "menus" of courses were generated by randomly selecting one course description from each quartile of the distribution of alignment scores for that theme. ChatGPT was asked to order the courses in order from least to most aligned to the given theme. The bars plot the share of pairwise comparisons for which ChatGPT's ordering matches the ordering from the curriculum alignment scores.

Figure A-13. Sensitivity of curriculum alignment trend to alternative alignment weights



Notes: This figure compares the trend in curriculum alignment scores for courses offered from 2002-03 to 2022-23 under alternative constructions of the curriculum alignment weights. For each series, the trend is estimated in course-level regressions of a course's curriculum alignment score for a given theme on a vector of year dummies, controlling for institution-by-field fixed effects. Changes are measured relative to the average curriculum alignment score in 2012-13 and reported in standard deviations. The "base" estimates for each theme use alignment weights constructed from the full thematic corpus; the "older period" estimates use alignment weights constructed only using documents from the first half of the thematic corpus; and the "recent period" estimates use alignment weights constructed only using documents from the second half of the thematic corpus. The figure omits alignment to social justice, which is not constructed from serialized documents. Analysis restricts to upper-level courses. Standard errors are clustered at the institution-field level.