

NBER WORKING PAPER SERIES

SKILLS, MAJORS, AND JOBS:
DOES HIGHER EDUCATION RESPOND?

Johnathan G. Conzelmann
Steven W. Hemelt
Brad Hershbein
Shawn M. Martin
Andrew Simon
Kevin M. Stange

Working Paper 31572
<http://www.nber.org/papers/w31572>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023

This research was supported by the National Science Foundation (Collaborative Grant #1919387), Russell Sage Foundation (Grant #1811-09737), and Institute of Education Sciences (Grant R305B150012). We are grateful to Nathan Sotherland, Jonathan Hartman, and Hayden Le for stellar research assistance. We received helpful comments and suggestions from Joe Altonji, John Bound, Charlie Brown, Brian Cadena, Blake Heller, Brian Jacob, and Sarah Turner—as well as seminar participants at the Federal Reserve Bank of New York, Federal Reserve Bank of Philadelphia, University of Michigan, University of North Carolina at Chapel Hill, Upjohn Institute, Villanova University, LMU Munich/CESifo, Institute for Labor Economics (IZA), and meetings of the American Economic Association (AEA), Association for Education Finance and Policy (AEFP), Southern Economic Association, Association for Public Policy Analysis and Management (APPAM), Society of Labor Economists (SOLE), and the National Bureau of Economic Research (NBER). All errors and any opinions are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Johnathan G. Conzelmann, Steven W. Hemelt, Brad Hershbein, Shawn M. Martin, Andrew Simon, and Kevin M. Stange. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Skills, Majors, and Jobs: Does Higher Education Respond?

Johnathan G. Conzelmann, Steven W. Hemelt, Brad Hershbein, Shawn M. Martin, Andrew Simon, and Kevin M. Stange

NBER Working Paper No. 31572

August 2023

JEL No. I23,J23,J24

ABSTRACT

How do college students and postsecondary institutions react to changes in skill demand in the U.S. labor market? We quantify the magnitude and nature of response in the 4-year sector using a new measure of labor demand at the institution-major level that combines online job ads with geographic locations of alumni from a professional networking platform. Within a shift-share setup, we find that the 4-year sector responds. We estimate elasticities for undergraduate degrees and credits centered around 1.3, generally increasing with time horizon. Changes in non-tenure-track faculty allocations and the credits they teach partially mediate this overall response. We provide further evidence that the magnitude of the overall response depends on both student demand and institutional supply-side constraints. Our findings illuminate the nature of educational production in higher education and suggest that policy efforts that aim to align human capital investment with labor demand may struggle to achieve such goals if they target only one side of the market.

Johnathan G. Conzelmann
Department of Public Policy
131 S. Columbia Street
Abernethy Hall, Campus Box 3435
Chapel Hill, NC 27599
jconzelm@unc.edu

Steven W. Hemelt
Department of Public Policy
University of North Carolina at Chapel Hill
131 S. Columbia Street
Abernethy Hall, Campus Box 3435
Chapel Hill, NC 27599
hemelt@email.unc.edu

Brad Hershbein
W.E. Upjohn Institute for
Employment Research
300 S. Westnedge Ave.
Kalamazoo, MI 49007
hershbein@upjohn.org

Shawn M. Martin
University of Michigan
shawmmm@umich.edu

Andrew Simon
University of Chicago,
Department of Economics
1126 E. 59th St
Chicago, IL 60637
andrewsimon@uchicago.edu

Kevin M. Stange
Gerald R. Ford School of Public Policy
University of Michigan
5236 Weill Hall
735 South State Street
Ann Arbor, MI 48109
and NBER
kstange@umich.edu

I. Introduction

The U.S. labor market increasingly rewards skilled workers, as technological change and outsourcing have reduced the need for workers to perform routine cognitive and manual tasks (Acemoglu & Autor, 2011; Autor et al., 2003, 2008). Skill demand varies both across and within occupations, varies across labor markets for the same occupation (Deming & Kahn, 2018), and accelerates during recessions (Hershbein & Kahn, 2018).² For example, the demand for social skills—an adeptness at productively working with others in flexible, team-based settings—has become increasingly necessary for the coordination and teamwork tasks central to the modern skilled workplace, and returns to jobs intensive in both cognitive and social skills have risen sharply (Deming, 2017).

Despite greater understanding of the evolution of skill demand, little research has focused on how students and workers acquire these skills (Altonji, Blom, & Meghir, 2012). The most direct way for individuals to build specific skills, at least for the nearly two-thirds who attend college, is through their choice of curriculum and field of study.³ Whether, and how quickly, human capital investment responds to evolving skill demand shapes economic growth, welfare, and inequality (Autor et al., 2020). Most empirical work on these questions, however, has focused on particular fields or sectors, going all the way back to studies of engineers and scientists a half century ago (Freeman, 1975, 1976).

In this paper, we quantify the magnitude and nature of human capital responses to shifts in labor market demand for nearly all undergraduate programs at 4-year colleges and universities across the United States—a sector of higher education that is often criticized for its inattention to labor market needs (e.g., Hansen, 2021). We measure labor demand at the institution-by-major level by combining the near-universe of online job ads with geographic locations of alumni from a professional networking platform. We ask, for example, how human capital production in chemistry at the University of Maryland, Baltimore County (UMBC) changes relative to other programs when the demand for chemistry majors changes more in Baltimore, Washington D.C., and New York City, three popular destinations for UMBC graduates, than in other areas. A

² Hershbein and Kahn (2018) demonstrate that these factors have manifested in skill and task changes within specific occupations over the past decade, and Atalay et al. (2018) show that task change within occupations has been occurring since at least the 1960s.

³ In their review of literature that explores heterogeneity in returns to college majors, Altonji et al. (2012) argue that notable portions of the differences in returns across majors are likely due to differences in “the market value of tasks that require specific knowledge and skills particular majors develop” (p. 218).

shift-share instrumental variables strategy then isolates the arguably exogenous portion of such variation due to differences in the baseline geographic concentration of industries—and by extension majors—in the geographic areas most relevant for each institution. That is, a national boom for industries that employ chemistry majors will result in relatively greater increases in demand for chemistry majors in labor markets where such employment is concentrated, and thus boost effective demand for chemistry majors relatively more at institutions that send larger shares of their graduates to those locales.

We find that postsecondary human capital investment responds strongly to changes in major-specific labor demand. Our preferred two-stage least squares (2SLS) estimates suggest an average elasticity of 4-year degree production with respect to labor demand of about 1.3. This elasticity generally rises and then plateaus as we extend the time horizons over which we measure changes in skill demand and human capital production. These results are robust to including additional controls for residual demand informed by our model of major choice and to various validation exercises suggested by Goldsmith-Pinkham, Sorkin, and Swift (2020).

To better understand the nature of the degree response, we investigate intermediate mechanisms, including credits taken and faculty composition, for a subset of institutions with available department-level data. The response of undergraduate credits is similar to that for degrees, and colleges appear to increase production by relying on non-tenure-track faculty. However, we find little change in the average number of course sections, which—combined with the prior finding of increases in undergraduate credits—suggests larger class sizes. Although major-specific human capital increases with labor demand, these supply-side adjustments in faculty composition and class size raise potential concerns about changes to inputs that shape the student experience.

While students and institutions are quite responsive on average, the magnitude of the equilibrium response depends on both student demand and a department's ability to accommodate additional demand given financial and other constraints. We investigate treatment heterogeneity in order to better understand how student preferences and supply constraints can influence the total response in ways that align with a theoretical framework of major choice. We find that less-selective and less-research-intensive institutions are much more responsive to skill demand changes than selective and research-intensive institutions. Less-selective colleges may be more likely to have excess capacity and therefore be better able to accommodate changing

student demand (Hoxby, 2009), consistent with the importance of supply-side constraints.⁴ When faced with excess demand and capacity constraints, institutions may instead ration slots in specific majors (for example, by implementing major-specific GPA requirements), which may have consequences for stratification across race or socioeconomic status (Bleemer & Mehta, 2021). Moreover, the role of capacity constraints is also consistent with the timing dynamics, where the elasticity of degree production is larger when institutions have more time to respond.

Educational production cost differences across majors could also affect the extent of response to shifts in demand. Because marginal costs from increasing class sizes or adding new sections (holding instructional quality fixed) differ across fields of study (Hemelt et al., 2021b), supply curves may be heterogeneous. We find that the overall elasticity is driven by fields of study in the lowest and middle terciles of the average cost-per-credit-hour distribution, while fields in the most costly tercile do not respond over the time period captured by our data. When we examine broad major groupings, social sciences, health, and communications display the most elastic responses to changes in skill demand. While these major groups generally consist of lower-cost majors, there are exceptions, with even higher-cost engineering including some lower-cost majors that may individually be more responsive. How much each field responds depends on student demand and major-specific production functions, as well as university priorities for relaxing major-specific constraints (Thomas, 2022). This is further evidence that supply-side constraints play an important role in determining which types of fields expand in response to additional labor market demand.

Heterogeneity in the magnitude of responses by institutions and fields of study could partially reflect differences in student composition and associated behavior, in addition to institutional supply-side behavior. We find that women are more responsive on average than men, but the gender difference in response shrinks notably—but does not disappear—when the distribution of fields is balanced between men and women. Taken altogether, our results indicate that both sides of the market—student demand for majors and institutions’ ability to expand access—shape the human capital response to changes in skill demand. Policies intended to increase the responsiveness of human capital accumulation to demand shocks that target only one

⁴ The heterogeneity by selectivity does not reflect differences between public and private nonprofit colleges, as we find similar levels of responsiveness across ownership.

side of the market, such as differential tuition for students or increased state subsidies for institutions, may thus have limited success.

Our main contribution is to provide causal evidence that human capital production in the U.S. 4-year college sector responds to field-specific changes in labor demand, which has not been previously shown despite the important implications. Our universal coverage of both institutions and fields of study offers a comprehensive view of the extent and nature of human capital responses on multiple margins, including degrees and undergraduate credits. In addition, novel department-level data on supply-side mediators, including faculty allocations by type and the number of course sections, permit us to investigate several channels through which institutions may facilitate (or impede) educational investment responses on the part of students. The broad coverage also allows us to consider heterogeneity by institution type and field of study, as well as across students, to highlight key mechanisms of the postsecondary educational production function.

Closely related work has largely focused on a few specific fields of study, institutions, or local labor markets (Acton, 2020; Carranza, Ferreyra, & Gazmuri, 2023; Foote & Grosz, 2020, 2020; Gilpin et al., 2015; Weinstein, 2022). With the exception of Weinstein (2022), this work also has mainly considered the 2-year sector, perhaps due to the large role community colleges play in training workers for middle-skill jobs (Deming & Noray, 2020; Grubb, 1996), or for-profit institutions, due to their perceived “nimbleness” (Deming et al. 2012). Understanding the responses of 4-year institutions is imperative given the level of resources invested toward the production of bachelor’s degrees (Ma & Pender, 2022), and sizable structural shifts in skill demand documented over the past 10–15 years (Blair & Deming, 2020; Hershbein & Kahn, 2018).⁵ Our framework and data also allow us to explore the timing of how students and institutions respond, building on prior work that focused on the production of engineers and scientists (Bound et al., 2015; Freeman, 1975, 1976; Ryoo & Rosen, 2004; Siow, 1984).⁶

A larger literature has examined students’ responses to earnings differences across fields of study, either quasi-experimentally (Altonji et al., 2016; Baker et al., 2018; Long et al., 2015;

⁵ Indeed, several states in recent years have proposed redirecting public subsidies in higher education toward fields believed to be in greater demand by employers (e.g., Cohen, 2016; Kumar, 2021).

⁶ This work highlights the notion that human capital investments take time, are forward-looking, and respond to persistent changes in demand. Findings from these early studies also suggest caution in the use of either current prices (i.e., wages) or quantities (i.e., employment) to characterize educational investment response, as neither is likely exogenous and instead reflects lagged shocks. These insights motivate our use of an augmented shift-share instrument to carve out plausibly exogenous variation in major-specific skill demand changes.

Webber, 2014; Wiswall & Zafar, 2015) or through experimental informational interventions (Baker et al., 2018; Hastings et al., 2015; Wiswall & Zafar, 2015), finding relatively modest responses in major choice. In contrast, our measure of labor demand derives from employers' explicitly stated preferences for specific majors and latent preferences embedded in the features of job ads, such as occupation, industry, and highly granular skills. Job ads are potentially more salient to students than earnings (Betts, 1996). In addition, when measured over suitable intervals, job ads may better reflect the “career prospects” of different fields of study—a construct that is theoretically consistent with optimization and which Ryoo and Rosen (2004) emphasize as critical to shaping human capital responses and educational investment decisions more broadly.⁷ This new measure, in tandem with our validated shift-share instrument, enables us to identify the causal effect of more meaningful labor demand shocks on educational investments.⁸

The paper proceeds as follows. In the next section, we describe the data sources we use to construct our measure of skill demand and to capture outcomes that reflect postsecondary educational investment. Section III presents our conceptual framework, describes our stacked long-differences empirical setup, details the construction of our instrument, and assesses the identifying assumptions on which a causal interpretation of our estimates stands. Section IV presents the main findings and discusses heterogeneity in educational investment responses to changes in skill demand by institutional characteristics, across fields of study, and by student gender. Section V concludes.

II. Data Sources, Core Measures, and Analytic Sample

We combine data from several sources, described below, to construct our measure of skill demand. We use nationwide institution-level data from the Integrated Postsecondary Education Data System (IPEDS) to capture the high-level outcome of interest—bachelor's degree completions, and we use department-level data from the The Cost Study at the University of

⁷ In testament to the saliency of non-wage measures of labor demand, recent work from Norway finds that high school students shift their postsecondary curricular choices away from vocational areas and toward academic areas in response to declines in the routine-task concentration of relevant occupations (Bennett et al., 2023). Of course, differences across students in ability and preferences also shape choices about field of study (Altonji et al., 2012).

⁸ As the instrument is constructed from publicly released data, other researchers can easily implement this design.

Delaware (Cost Study) for a subset of institutions to examine more granular outcomes, such as the number of credits, course sections, and faculty positions by type.⁹

A. Measuring Labor Demand

The construction of our demand measure is motivated by insights from early work that sought to model the supply of workers to skilled professions. Namely, Freeman (1976) concluded that while salaries did a decent job of explaining the supply of engineers, more “direct” measures of “market-determining factors” would better identify causal responses to demand changes. This sentiment was echoed in follow-on work that emphasized the importance of labor market entrants’ forward-looking behavior in terms of their career prospects (e.g., Ryoo & Rosen, 2004; Zarkin, 1985).

We attempt to measure “career prospects” through a manner directly observable to job seekers: the near universe of online job ads in the United States between 2010 and 2017, obtained from Burning Glass Technologies (BGT or Burning Glass).¹⁰ Job ads precede both employment decisions and salary offers and are designed to be highly visible; they thus constitute a much more direct and salient signal of demand conditions. Indeed, recent experimental work finds that college students’ choice of major responds much more to information about employment prospects than earnings conditional on employment (Ersoy & Speer, 2022).

BGT scours about 40,000 online job boards and company websites to aggregate job postings, parse and deduplicate them into a systematic, machine-readable form, and create labor market analytics products. The data contain detailed information on over 70 standardized fields including occupation, geography, skill requirements, education and experience demands, and firm identifiers. There are over 15,000 individual skills standardized from the open text across job postings. The data cover the entire United States and contain roughly 153 million postings.

Since the database covers only vacancies posted online, the jobs represent only a subset of the employment demand in the entire economy. Coverage of the BGT data has been examined in prior work. Hershbein and Kahn (2018) find that although BGT postings are

⁹ More information about The Cost Study can be found at <https://ire.udel.edu/cost/>. Because coverage of private for-profits is limited in this source (and some other sources we use, described below), we focus our analyses on public and private nonprofit 4-year institutions in the United States.

¹⁰ In 2021 BGT merged with competitor Emsi, and the joint company is now known as Lightcast. Our data predate this merger.

disproportionately concentrated in occupations and industries that typically require greater skill, the distributions are relatively stable across time, and both the aggregate and industry-specific trends in the number of vacancies track other sources reasonably closely.¹¹

We restrict our sample to job postings that list at least one skill, require exactly 16 years of education (e.g., a bachelor's degree), and are posted in a Metro or Micro Statistical Area (MSA). Our focus on job ads requiring a bachelor's degree will cause the skill skew to be of less concern. Hemelt, Hershbein, Martin, and Stange (2021a) find that, unsurprisingly, the job ads included in this sample are disproportionately in professional occupations and less likely to be in Sales, Office Administration and Support, and Food Preparation.¹²

We aim to construct the demand for postsecondary education at a program (institution-by-major) level, so a key variable in the BGT data is college major (which is provided under the Classification of Instructional Programs (CIP) taxonomy). College major is listed in 54 percent of all job ads requesting a bachelor's degree as the education requirement, with about 55 percent of such postings listing a single major, 30 percent listing two, and 15 percent listing three or more. Hemelt et al. (2021a) investigate the differences between ads with and without a major explicitly listed. They find that the distribution of observables—occupation, industry, MSA, skills—between job postings with a major differs from those without a major. However, even with a very detailed set of observable controls, almost three-quarters of the variation in whether a posting lists a major remains unexplained.

For the purposes of analyzing skill demand by major, we aggregate college majors into 71 categories, though we use only 66 in our analysis.¹³ Our aggregation procedure, detailed in Hemelt et al. (2021a), attempts to produce categories that reflect fields that students confront when making decisions about paths of study in college and that have meaningful quantities of both job ads and degrees granted according to IPEDS.

¹¹ See online Appendix A of Hershbein and Kahn (2018).

¹² Appendix Table A1 describes the occupational distribution of the job ad sample as various restrictions are imposed. Although we have not imposed a maximum experience restriction (to focus on recent college graduates), relatively few ads call for more than five years of experience, so gains in sample size may warrant slight deviations from representativeness.

¹³ Omitted majors include Construction Management; Mental and Social Health Services; Allied Health Diagnostic, Intervention, and Treatment; and Urban Planning—as these do not have readily identifiable matches in the American Community Survey, which we use to map majors and industries as described below. We also omit college majors that are traditionally sub-baccalaureate or remedial programs (e.g., Basic Skills and Developmental/Remedial Education), that are predominantly post-baccalaureate or graduate programs (e.g., Residency Programs), or trade-specific (e.g., Mechanic and Repair Technologies/Technicians).

B. Imputing Labor Demand

There are several issues with characterizing demand for different majors from job ads. One is that many ads list multiple majors and while it is straightforward to treat each listed major as a separate observation, it is not clear whether employers have a preference ordering across these majors. Another issue is the fact that nearly half of all bachelor’s degree-seeking ads do not explicitly state a specific major. Using a sample restricted to ads that *do* explicitly state a major could mischaracterize the total demand for that major (or for other majors). For example, it is possible that the absence of a listed major indicates indifference by the employer such that *any* major would be suitable. Moreover, employers are staffed by humans, and ads may neglect to list majors that are indeed demanded, either at the extensive margin (no major was listed, but one should have been) or the intensive margin (at least one major was listed, but not all demanded majors were listed). These forms of measurement error could mischaracterize true demand for majors when using a sample based only on ads with a major listed explicitly.¹⁴ To address these issues, we impute demand for each major, for each ad, using the rich information contained in the ad and a machine learning classifier.¹⁵

This multi-class classification problem aims to assign the probability that each ad would be appropriate for each individual major. Standard metrics used to determine the best algorithm in binary classification problems, such as precision, recall, and F1, are potentially misleading in our setting for three reasons.¹⁶ First, as a multi-class categorization problem (i.e., each job ad can have up to 71 labels), we require a metric that can accommodate assignment to multiple classes. Second, the “truth” data against which we train our algorithm may be incomplete. Many ads list only one or two majors, but may actually represent demand for additional majors—for example, a job ad that calls for Communications majors might also be appropriate for Journalism majors. Consequently, we want to avoid metrics that sharply penalize “false positive” predictions—that is, predictions that may be valid but are not classified as such in our training data. Finally, we are interested in predicted probabilities for each major-ad combination (rather than binary assignment), since majors with a reasonable likelihood of being appropriate for a given ad (but

¹⁴ Similarly, ads open to majors with what employers believe are closely related skills (e.g., “business, management, or a related field”) would not yield information about the implicit related fields.

¹⁵ Additional details of our classification approach can be found in Hartman, Hemelt, Hershbein, Sotherland, and Stange (2022).

¹⁶ Suppose c is the true class and c' is the predicted class. Precision is then $P(c' = c | c')$, the share of predictions that are “correct”, recall is $P(c' = c | c)$, the share of true instances that are accurately predicted, and F1 is the harmonic mean of precision and recall.

perhaps not above an arbitrary 50 percent threshold) should not be treated the same as majors that are completely unrelated to a given ad. The metric referred to as Label Ranking Average Precision (LRAP) is appropriate for such settings. LRAP accounts for multiple classes, incorporates predicted probabilities, and does not harshly penalize “false positives.”¹⁷

Using the LRAP metric, we evaluated three different algorithms (penalized logistic regression, decision tree, and random forest), for three different training sample sizes (1, 3, and 5 percent random samples), and for four different sets of features. Our preferred approach uses a random forest trained on a 5 percent random sample with the following features: indicators for occupation (6-digit SOC codes), industry (4-digit NAICS codes), MSA, year, and month, and the 1,250 most predictive unigrams from tokenized text data on job title, employer name, and skill requirements.¹⁸

We use our estimated model to predict the probability of each major being appropriate for each ad in our analytic sample of job ads. We then aggregate these probabilities to construct our main measure of major-specific demand for each higher education institution in each year: the number of total ads demanding each major, including those imputed through the above process. We also separately aggregate the number of ads explicitly listing each major as an alternative demand measure. As shown in Appendix Figure A1, there is a reasonably close correspondence between the *aggregate change* in field-specific demand measured with or without the imputations.¹⁹ However, we believe the imputation-based measure better captures

¹⁷ We calculate LRAP as follows. For each ad, we rank the majors by predicted probability. For each observed major for each ad, we determine (a) the rank of that major among the predictions, as well as (b) the number of predicted majors that are among the observed set of majors for that ad *and* are of at least the rank of (a). We divide (b) by (a), and then repeat for each major of a given ad. We then average these ratios at the major-by-ad level. An LRAP value closer to one means that the model has predicted more of its true labels with higher probability, and has avoided false negative predictions. False positives will not be penalized by this metric unless the model predicts them as more likely than the true labels.

¹⁸ The 1,250 unigrams included in our preferred feature set are selected from 5,000 each of the most common job title, employer name, and skill tokenized text unigrams. We use a “chi-2 feature selection” method that is common in the natural language processing field. We operationalized this method by conducting chi-2 tests between each pair of the 15,000 features and 71 possible majors. The “most predictive” 1,250 unigrams are those with the highest sum of chi-2 statistics across the 71 majors. Appendix Table A2 provides performance metrics for other models, feature sets, and sample sizes. We train each model, feature set, and sample size combination using 80 percent of the data and then calculate test metrics using the withheld 20 percent. The models are built in Python using sklearn’s OneVsRestClassifier. Each model is scored on the held-out test set.

¹⁹ In a companion paper (Hartman, Hemelt, Hershbein, Sotherland, & Stange, 2022), we explore the degree of disconnect between a demand measure based solely on majors explicitly stated in job ads and another that incorporates latent demand using the classification methods summarized above. We find that latent demand is greatest among majors with broad sets of “soft skills,” and that accounting for latent demand reduces the measured disconnect between supply and demand.

demand, and this may matter more at our preferred level of analysis of institution by major. Henceforth, we include the sum of imputed and explicitly listed majors whenever we refer to the number of ads for each major-institution-year observation.

C. Defining Markets

Institutions in our sample vary by geographic location and by unique ties to different labor markets across the country. Rather than assume employment demand shocks are felt equally by all colleges (i.e., national changes) or that the area closest to an institution defines its “market,” we measure institution-specific labor markets more granularly based on where recent alumni live and work.

We use college-specific labor market catchment areas from Conzelmann et al. (2022), who aggregate data from the social networking platform LinkedIn (LI). Specifically, these data capture institution-specific alumni counts, among the classes of 2010 through 2018 from the 15 most popular metropolitan destinations in the U.S. for each institution, all in-state locations, and a subset of other geographies identified through nearest-neighbor matching to other institutions with similar characteristics.²⁰ The processed data consist of a set of shares for each institution that capture the distribution of that institution’s total U.S. alumni residing across 278 LI geographies. These geographies roughly correspond to individual or aggregations of Core-Based Statistical Areas (CBSAs) from the Census Bureau.²¹ Although one may be concerned about representativeness of the data, Conzelmann et al. (2022) subject the data to several validity checks, including against more representative (if geographically limited) sources, such as the Census Bureau’s Post-Secondary Employment Outcomes, and find that the LI data perform quite well.

For our purposes, the geographic differences in college-specific labor markets provide additional variation that helps us identify how changes in employer demand for certain majors produce a supply response. To visualize some of this geographic variation across different types of colleges, in Figure 1 we separately map the geographic distributions of graduates from North Carolina’s public (Panel A) and private nonprofit 4-year institutions (Panel B). A large concentration of graduates from both types of institutions remain in North Carolina locales (e.g.,

²⁰ For more details on the data collection process, representativeness, and validation of these data, please see Conzelmann et al. (2022). The data are publicly available at <https://www.openicpsr.org/openicpsr/project/170381/>.

²¹ The data from Conzelmann et al. (2022) include a crosswalk between the LI geographies and CBSAs.

Charlotte and the Raleigh-Durham area); however, institutions also send notable proportions of their graduates outside the state and to other metropolitan areas across the country. This variation is more pronounced for private institutions, as depicted by a larger number of areas outside of North Carolina where graduates reside. For example, larger shares of alumni from private North Carolina institutions end up in Washington-Baltimore, the Northeast, and major cities outside the state—such as Chicago, Denver, and Los Angeles—relative to public alumni, who tend to stay closer to their institutions.

These contrasts support defining each college’s labor market according to where their graduates end up. The patterns suggest, for example, that a computer science demand boom in Texas is unlikely to have a meaningful effect on students from North Carolina public institutions, but a boom in the New York area might, since a large share of North Carolina’s graduates tend to locate there. These geography shares for graduates of each college thus provide weights that let us map demand shocks at the level of geographic labor markets to the level of institution-specific labor markets.

More specifically, our main explanatory variable is the aggregation of BGT job postings to the major-institution-year level: $\log JobAds_{tms}$. Using information in each job ad on the advertised majors (including the imputation process described above), geography, and posting date, we compute the number of job ads for graduates of major m in area g in time t , A_{tmg} , as our measure of demand. We then aggregate to the program (i.e., institution-major) level by summing across areas, weighting by the institution-specific LI market shares described above, ω_{gs} , and taking the natural log:

$$\log JobAds_{tms} = \ln \sum_g \omega_{gs} A_{tmg} \quad (1)$$

Our measure of demand therefore computes the effective number of job ads for a graduate from institution s who majored in m in year t based on the number of ads targeted to her major in a given area and the likelihood a graduate from institution s moves to that area.²²

²² Note that our LI shares are not major-specific because geographic locations of graduates are not available in LI separately by institution and field. Aside from the Census Bureau’s Post-Secondary Employment Outcomes, which have limited scope in coverage of institutions and geographic granularity, we are unaware of any large-scale source that provides geographic location of alumni for institutions separately by major.

D. Outcomes

The extensive and intensive margins of any postsecondary response to labor market changes likely differ in magnitude and timing, and we focus on the latter.²³ Our main outcome of interest is the number of bachelor's degrees awarded each year by major for each institution in our sample. We obtain this information from IPEDS completion files, crosswalking counts at the 6-digit 2010 CIP code level to a condensed list of 66 major categories.

We construct these major aggregates in a manner that preserves the CIP code hierarchy, ensures a sufficient number of degrees granted and job ads in each aggregate, and combines majors that display a similar skill profile in the BGT job ads. Hemelt et al. (2021a) describe this process in more detail. We use these yearly counts to generate long differences in degrees granted at the program (i.e., institution-major) level.

In addition, we obtain program-level data on undergraduate credits, instructional costs, course sections, and faculty allocations from The Cost Study, which is organized and managed by the University of Delaware. The Cost Study has collected program-level data from 4-year institutions on costs, faculty, credits produced, and other measures of productivity since the late 1990s. Participation in The Cost Study is voluntary, and institutions sometimes move in and out of the sample, but there is no reason to think participation is tied to either major-specific or institution-specific labor demand shocks for graduates.²⁴

From this data source we obtain the total number of undergraduate credits produced by program, which provides a more malleable measure of supply than degrees completed from IPEDS. For example, students may respond to labor demand shocks by taking more credits in a given field, even if they do not switch majors, which can be costly, especially late in one's college career. The Cost Study separately reports lower-division and upper-division undergraduate credits, allowing us to distinguish whether response comes more from introductory level classes or from advanced courses. We additionally observe potential supply-side adjustment measures by the program, including the number of tenure-track and non-tenure-track faculty—that is, faculty allocations—and the number of sections taught. From a

²³ For instance, offering a new degree program (or eliminating one) is a time-consuming and burdensome process involving multiple layers of institutional, and often system-level, deliberation and approval. We would thus expect response at this extensive margin to be rarer, occurring under sustained demand shifts and over long time horizons. Given our relatively short panel, we focus on changes to programs that operate in both a beginning and an ending period.

²⁴ See Hemelt et al. (2021b) for additional information on institutional participation in The Cost Study, as well as differences in instructional costs by field of study.

production perspective, such data permit us to investigate how institutions produce more (or fewer) credits, including adjustments in the numbers and types of faculty as well as (implicit) class size, with each of these margins implying different marginal cost structures.²⁵ As with our IPEDS sample, all Cost Study data are reported by CIP code, which we crosswalk to our 66 major categories and then generate long differences at the program level.

E. Analytic Sample and Descriptives

We begin with the universe of public and private nonprofit 4-year colleges and universities in IPEDS, 1,754 institutions, each of which have up to 66 different fields, for a total of 115,764 possible programs. However, few institutions grant degrees in all fields in all years: only about 30 percent of potential programs have a positive number of degrees. Because our analysis focuses on long differences, we exclude institution-field combinations with no degrees granted in either the base year or final year, resulting in an analytic sample of 32,554 individual programs (institution-major) at 1,681 institutions.

Table 1 summarizes the main features of our analytic sample. The IPEDS sample is representative of public and nonprofit 4-year colleges and universities in the United States. The sample covers about three-quarters of all U.S. public and nonprofit institutions and an even greater share of all 4-year enrollments and degrees, since the sample includes only programs with a positive number of degrees granted. The data capture institutions that differ in their control, level of research activity, and selectivity, which allows us to consider the extent to which responses to changes in skill demand vary across different types of institutions. The Cost Study covers fewer institutions, only 114, and more heavily represents public institutions, as well as research universities. This limits the scope of our analysis to the average effects of changes in skill demand on credits, faculty, and other inputs. The programs (and institutions) in The Cost Study also are much larger, on average, granting more than twice the number of degrees as the IPEDS sample.

Appendix Table A3 shows the distribution of degrees granted by field for the IPEDS analytic sample for the graduating cohorts of 2017 through 2019. The ten most common fields account for 44 percent of degrees granted and more than one-third of all institution-by-major observations.

²⁵ Although class size is not explicitly reported in the Cost Study data, we infer implicit changes in class size by comparing changes in total (undergraduate) credits taken relative to changes in the number of course sections.

III. Empirical Framework and Method

A. Motivating Conceptual Framework

We are interested in how major-specific demand shocks at different postsecondary institutions affect subsequent human capital production, measured by degrees or credits. We briefly sketch a model of program choice to guide the factors we include in our empirical model. Consider the decision of student l , from cohort c , to enroll in institution s and major in m .²⁶ Her decision is based on the average net value of each major—that is, the major’s (institution-specific) present discounted value of future earnings and additional non-pecuniary benefits less net costs, V_{csm} —as well as her own preferences, ε_{clsm} , which we assume follows a Type I extreme value distribution. More formally, she chooses m^* to maximize her utility:

$$m_{cls}^* = \arg \max_m U_{clsm} = \arg \max_m V_{csm} + \varepsilon_{clsm}, \quad (2)$$

where we normalize the value of not enrolling in college to be 1, without loss of generality.

Following Blom et al. (2021), we decompose V into three components. First, there is a fixed component of completing major m at institution s , η_{ms} , which captures time-invariant benefits and costs to the specific major and college, such as access to professional networks or the difficulty of coursework. Second, we include a structural component, μ_{cm} , that represents time-varying changes in the value of each major that are common across institutions, such as skill-related technical change or demand shifts brought by evolving demographics (e.g., health training for an aging population). Finally, and the focus of our paper, there is a program-specific time-varying component that captures relative labor demand for graduates from cohort c of institution s who majored in m , γ_{cms} . How much the student responds to changes in labor demand depends on β , which may reflect, in part, the salience of γ_{cms} .²⁷ Therefore we can rewrite

Equation (2) as:

²⁶ In this framework we focus on students’ decisions, though our empirical application will uncover the combined responses of institutions and students. One can think of students as solving the maximization problem subject to institutional decisions to alter course offerings, resources, and limitations on major choice in response to changes in demand and various constraints. We return to this issue of separating student from institutional response when we present results.

²⁷ For simplicity, we assume here that students’ responses to changes in labor demand do not vary across individuals. However, in reality, different types of students may respond differently based on their preferences or based on supply-side constraints that vary across institution types. We later relax the assumption of constant responsiveness and allow for heterogeneity in our empirical results in section IV.D.

$$m_{cls}^* = \arg \max_m \eta_{ms} + \mu_{cm} + \beta\gamma_{cms} + \varepsilon_{clsm} \quad (3)$$

From this expression, the likelihood that an individual student chooses major m is:

$$\Pr(m_{cls}) = \frac{\exp(\eta_{ms} + \mu_{cm} + \beta\gamma_{cms})}{1 + \sum_k \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})} \quad (4)$$

Aggregating across all students in a cohort who are considering s , N_{cs} , the number who major in m is $Y_{cms} = N_{cs} \Pr(m_{cls})$. And so:

$$\log(Y_{cms}) = \log(N_{cs}) + \eta_{ms} + \mu_{cm} + \beta\gamma_{cms} - \log(\pi_{cs}), \quad (5)$$

where $\pi_{cs} = 1 + \sum_k \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})$ is institution-cohort specific. Therefore, we can express the number of degrees produced in a major m by institution s for cohort c as a function of labor demand for that major and other cohort-, program-, and institution-specific factors.

With this framework, we now consider how students respond to changes in labor demand. We take the long difference across cohorts from Equation (5):

$$\Delta \log(Y_{ms}) = \beta\Delta\gamma_{ms} + \Delta\mu_m + \Delta \log(N_s) - \Delta \log(\pi_s), \quad (6)$$

Note that the fixed component η_{ms} differences out, such that the change in educational investment across cohorts can thus be decomposed into changes in demand ($\Delta\gamma_{ms}$), changes in major-specific attributes common across institutions ($\Delta\mu_m$), and institution-level factors ($\Delta \log(N_s) - \Delta \log(\pi_s)$).

B. Empirical Implementation

We map this model-based equation to our data by parameterizing $\Delta\gamma_{ms}$ as the log change in major-relevant job ads for students at s between time t_0 and time t_j

$\log JobAds_{t_j ms} - \log JobAds_{t_0 ms} + u_{ms}$, where u_{ms} is sampling error. We increase precision by stacking several long differences of program-year (i.e., institution-by-major-by-year) data, pooling all institutions and majors:

$$\log Y_{t_k ms} - \log Y_{t_0 ms} = \beta(\log JobAds_{t_j ms} - \log JobAds_{t_0 ms}) + \mu_{t_0 m} + \Theta_{t_0 s} + u_{t_0 ms}. \quad (7)$$

Treating each program-year long difference as a separate observation, this model estimates the association between the growth in demand from year t_0 to t_j and the growth in the output (e.g., undergraduate degrees granted) from t_0 to t_k . The parameter β captures the elasticity of educational investment with respect to changes in employer skill demand.

We weight by the number of degrees granted in t_0 so that the elasticity reflects the change in aggregate supply by individuals rather than by programs. Identification comes from variation in changes in demand across institutions and their effective labor markets within a given major. For instance, how much more does computer science investment increase at Arizona State relative to other institutions when the demand for computer science majors increases more in Phoenix and Los Angeles (two popular locales for Arizona State graduates) than in other areas?

The long-difference nets out any fixed (time-invariant) differences across programs that may also correlate with demand and degree/course production (e.g., average reputation). Equation (6) indicates that to estimate β , we need to additionally control for $\Delta\mu_m$ and for $\Delta \log(N_s) - \Delta \log(\pi_s)$. To address the former, we include major-by-base-year fixed effects, $\mu_{t_0 m}$. Within each long difference, these effects account for any national trends in field-specific demand or supply that may correlate with the production of degrees and courses.²⁸ For example, aggregate trends in student preferences for majors—say, toward those with easier grading—could affect supply and also be spuriously correlated with market demand; the inclusion of $\mu_{t_0 m}$ helps control for such possible confounders.

In (6), $\Delta \log(N_s)$ captures factors influencing the change in overall degree production at an institution—such as aggregate population growth (or decline) in the areas served by an institution. For example, population growth might induce increases in both job ads and college completions, especially among institutions that primarily serve local communities, incidental to major-specific demand. This could induce a spurious positive correlation between educational investment and labor market demand. To the extent this term is relatively fixed across institutions, it can be approximated with a constant in our long-differences setup (which is functionally absorbed by the vector of major-by-base-year fixed effects). $\Delta \log(\pi_s)$ captures the

²⁸We stack three long-differences and so, following the logic of the equation, each fixed effect is a long-difference-specific major fixed effect.

change in the value of all majors from a demand shock, including m . Thus in some specifications we include base-year-specific institution fixed effects, Θ_{t_0s} , which capture both $\Delta \log(N_s)$ and $\Delta \log(\pi_s)$, or alternatively, the change in degrees granted in other fields at the same institution.²⁹ Since $\Delta \log(\pi_s)$ is a function of $\beta \Delta \gamma_{ms}$, however, Θ_{t_0s} likely captures part of the effect of interest. Moreover, by capturing $\Delta \log(N_s)$, the inclusion of institution fixed effects also importantly changes the identifying variation to instead reflect within-institution compositional changes rather than total increases in degrees produced. We thus present estimates with and without controls for these institution-wide factors.

In our preferred specification, we set $(t_j - t_0)$ to 5, and $(t_k - t_0)$ to 7. To capture sustained demand shifts to which students and institutions can respond, we measure changes in demand over a 5-year horizon. Based on the typical timing of degree choices and prior work, we allow the supply outcome (e.g., degrees granted or credits taken) to lag demand changes by 2 years.³⁰ In our job postings data, we observe $\log JobAds_{tms}$ for each year from 2010 through 2017, and thus our preferred specification includes three sets of stacked long-differences with base years (t_0) of 2010, 2011, and 2012.³¹ Consequently, we use the periods 2010–2015, 2011–2016, and 2012–2017 to measure changes in demand, and the periods 2010–2017, 2011–2018, and 2012–2019 to measure changes in outcomes. To better understand dynamics, however, we also consider specifications that vary the length over which we measure both changes in demand and changes in outcome measures; although theory and prior empirical work both suggest there should be a lag between a change in demand and the response to it, the length of this lag may vary across time, major, and other institutional characteristics.

The single elasticity estimate from the pooled specification described above likely masks heterogeneity in response across fields, institutional type, and context. Thus, we explicitly examine such heterogeneity across the dimensions of institution control (private versus public), research intensity, and selectivity. We also explore differences in responsiveness across fields of

²⁹ This leave-out control should account for institution-level factors influencing the desirability of all majors other than the focal one, as we show in an extension to our conceptual framework in Appendix B.

³⁰ For example, Stange (2015) finds that changes in degrees granted by major follow changes in major-specific tuition prices by 2 years.

³¹ In practice, the base years for demand and the outcome do not need to be the same. We use a single base year for tractability.

study characterized by their instructional costs per credit hour (described in more detail below), as well as ten broad groups of majors that reflect institutional organization (reflecting colleges within a university, for example). Finally, given large differences in field of study between men and women, we investigate heterogeneity by student gender.

C. Limitations of OLS Approach

OLS estimates of β from equation (7) will be biased if there are unobserved factors influencing educational investment that are correlated with demand, conditional on the included fixed effects. One source of endogeneity is simultaneity or reverse causation, wherein institutions ramp up (or ratchet down) course offerings and degree production in certain areas due to knowledge of possible future business openings or closures. Virginia's promise of educational investments to lure Amazon's second headquarters serves as a prominent example (Svrluga, 2018). Similarly, exogenous changes in the supply of majors in a given area (e.g., prospective teachers deterred by an inhospitable work climate) might induce changes in employers' ad-posting behavior—possibly leading to a negative relationship in the OLS setup.

Finally, declining costs of online job-ad posting, the emergence of additional online job search platforms (e.g., Indeed), and recruiting norms, all of which could vary by major, may introduce measurement error into the variable we use to capture the underlying construct of skill demand. For example, IT job postings were likely exclusively online earlier in the sample period than were postings for nursing jobs, so growth in postings for the latter may capture both real changes in demand as well as changing coverage in the data. Such measurement error would likely attenuate estimates from OLS estimation.

D. Exogenous Demand Shocks

To address these concerns, we instrument for institution-major-specific demand. We exploit national shocks in the demand for the skills embedded in certain majors interacted with pre-existing differences in industry structure across labor markets to create a shift-share instrument (Bartik, 1991). Like the standard application of the Bartik instrument, our instrument relies heavily on industry-area shocks. However, different from classical adaptations of this instrument, we convert these industry-area shocks into major-area shocks and then into institution-major shocks using several mappings described below.

Our instrument takes the following form:

$$\sum_g \omega_{gs} \sum_i \frac{Z_{im} \times \frac{E_{igt_0}}{E_{gt_0}}}{\sum_h Z_{hm} \times \frac{E_{hgt_0}}{E_{gt_0}}} \times (\ln E_{it_j} - \ln E_{it_0}) \quad (8)$$

Here, i refers to industry, g is the LI geography (roughly CBSA), m is major, s is institution, and t is time. We describe the components of our instrument, moving from right to left.

The right-most component measures the log change in national employment from year t_0 to t_j —where E_{it_j} denotes the national employment of workers in industry i during year t , which we measure with data from the County Business Patterns (CBP).³² As before, we use 5-year horizons to operationalize t_j and t_0 .

The next term is a fraction that contains two elements. First, we map industries—the guts of the shift-share instrument—to majors. The term Z_{im} represents the national share of workers employed in industry i who majored in field m , based on a pooled cross-section of American Community Survey (ACS) data from 2010–2018.^{33,34} Second, we need a baseline measure of the local industry mix for each geography. Accordingly, E_{igt_0}/E_{gt_0} is the share of employment in area g at baseline (i.e., t_0) working in industry i . We then sum this measure over all industries (within each geographic area). The result—combined with the first term—is a measure of major-specific employment shocks to an area due to varying exposure to common, national employment changes. However, because the numerator is the product of two shares that do not sum to one within major field and geography, we rescale that product so that the values sum to one within a major-geography cell—which is accomplished by the denominator in this second term.

³² We map consistent 2012-vintage NAICS industries in the CBP into 239 detailed industries available in the American Community Survey. We use industries rather than occupations (which would make sense from a task or skill viewpoint) because consistent, occupation-based employment at the metro level is not available at the granularity we need.

³³ Although it would be ideal to have geography-specific mappings between industries and majors, the ACS dataset, sizable as it is, regrettably does not permit such detailed mappings.

³⁴ As an alternative conceptualization, we could define the relevance as the total number of workers who majored in m and work in industry i divided by the total number of workers who majored in m . However, this is not our preferred approach, as we discuss in Appendix C, since it breaks the direct link between changes in industry job ads and labor demand through employment counts. Under this alternative conceptualization, the weights are instead tied to the size of the majors. Our main 2SLS results are qualitatively similar using this alternative weighting approach, though less precise and with much less power in the first stage.

Finally, the left-most term in our instrument translates the major-area shocks to institution-major-specific demand shocks based on the locations of each institution's recent graduates. That is, for each institution, ω_{gs} is a vector of shares of graduates residing in each geography. We sum over the major-area shocks, weighting by ω_{gs} , to arrive at the final value of our instrument, the institution-major demand shock.

At a high level, identification for the instrument is similar to that for the direct measure of job postings in that it stems from different geographic exposure of institutional labor markets. In this case, however, it comes from changes in demand that vary across place solely due to the underlying composition of industries—and by extension majors—which in turn matter to different degrees for institutions based on their markets. For instance, a national boom for the IT industry will disproportionately increase the demand for majors that feed into IT in labor markets where employment in IT is concentrated. This will be particularly true for postsecondary institutions that send a sizable portion of graduates to IT-heavy locations. Thus, the exclusion restriction is that local industry-major shares—which apportion national employment changes across place—influence educational investment only through changes in skills and majors demanded on job ads.³⁵

E. Instrument Diagnostics

Recent econometric work provides refined guidance on the identification that undergirds Bartik-like IV approaches (Goldsmith-Pinkham et al., 2020). The validity of all instruments turns on their relevance and their ability to satisfy the context-specific exclusion restriction (i.e., that the instrument is correlated with the outcome only through its influence on the focal independent variable). In subsequent results, we report F-statistics from first-stage regressions of our measure of skill demand (the change in job postings) on our instrument in support of the relevance criterion. Here, we explore several analytic checks to assuage concerns that our baseline industry-major shares might correlate with other forces that influence both skill demand and measures of educational investment, such as degree production.

A key insight from Goldsmith-Pinkham et al. (2020) is that identification in a setup using a Bartik-like instrument stems from the baseline industry shares. Hence, the base period must be

³⁵ Put differently, the first stage of regressing the change in job ads on the instrument isolates the component of the change in job ads that is due to shifts in predicted employment, thus purging some (if not all) of the measurement error and endogeneity issues described in the previous subsection.

specified sufficiently early such that the initial industry shares are conditionally exogenous.³⁶ While the diagnostics presented in their paper are helpful, they do not map clearly to our setting for a few reasons. First, we are not focused on a particular field of study nor a specific sectoral employment shock, as is the case in related prior work. Thus, the recommended checks for assessing the face validity of the baseline industries with the largest weights (i.e., those that drive the overall 2SLS estimate) hold less relevance for our setting. Second, the piece of our instrument that apportions national employment changes across areas in an arguably exogenous manner is a function of industry shares mapped to college majors and weighted by each institution's relevant market, complicating the application of other recommended diagnostics (as well as diluting the role of any given individual industry share).

Given these caveats, we attempt to follow the spirit of Goldsmith-Pinkham, Sorkin, and Swift (2020)'s validity exercises by first calculating the Rotemberg weights adjusted to be analogous to our setting. These weights permit the researcher to determine the industries with the greatest contributions to the overall, average 2SLS estimate.³⁷ We then flag the top 10 industries based on those weights, drop each industry from the construction of our instrument one at a time, and re-estimate the overall 2SLS regression. We also use data on degree production from periods that entirely precede those captured by our instrumented demand measure to assess pre-trends via a falsification analysis. We report and discuss the results from these robustness and validity checks in Section IV.C. below.

IV. Findings

A. Aggregate Patterns

We first explore aggregate relationships, by major, between the change in the number of degrees granted and the change in the number of job postings. Figure 2 plots, for 66 majors, the average, demeaned 7-year change in the log of degrees granted against the average, demeaned 5-year change in the log of job postings, with the size of the markers proportional to the average baseline number of degrees granted in the major.³⁸ At the aggregate field level, there is a clear

³⁶ An alternative framework provides support for settings where the causal plausibility of shift-share designs rests on the exogeneity of the shifts (Borusyak, Hull, & Jaravel, 2022).

³⁷ That is, the largest weights reflect industries whose 2SLS estimates, if produced by using solely that industry's baseline shares as the instrument, account for the most weight in terms of influencing the overall 2SLS estimate.

³⁸ For each major and long-difference period, we compute the change in each log measure for each institution over the specified time period. We then subtract the average of these changes, where this average is weighted by the baseline number of degrees granted. We then average across each long-difference period to yield a single ordered pair for each major.

negative relationship between the change in demand (job ads) and supply (degrees). The correlation is -0.40 when weighting fields by the number of degrees granted at baseline, and -0.31 when unweighted. Teacher Education, for instance, experiences a very large increase in the number of job ads but a large (0.30 log point) reduction in the number of degrees granted relative to other majors. English—the field offered by the most institutions—also experiences a large drop in the number of degrees despite having average demand growth. Computer Science and Other Engineering experience very large increases in degrees granted despite lower-than-average demand. Some large fields move counter to this pattern, consistent with a positive relationship between supply and demand. Nursing has both larger than average demand growth and an increase in the number of degrees granted, while Business has the opposite pattern, with below-average changes in degree production and number of job ads.

While potentially informative of aggregate trends in supply-demand imbalances, these field-level patterns will confound any secular trends in the desirability of majors and their labor market prospects. Furthermore, as discussed above, if the composition of the ads contained in the BGT database is changing differentially by field due to improvements in their data collection technology unrelated to the underlying labor market, then these patterns will again mischaracterize the responsiveness of postsecondary programs to labor market demand.³⁹

To address these concerns, we exploit variation in demand changes experienced by different institutions for the same field. Figure 3 depicts such cross-institution variation within fields, plotting different quantiles of the raw distribution of changes in log job ads across institutions. Large aggregate differences across fields (shown in Figure 2) are apparent in the medians; however, there is also substantial variation in demand *within* field. For instance, education programs experienced substantial growth in demand across the board, but the interquartile range equals more than 0.25 log points. Similar cross-institution ranges appear across many other fields. We use this cross-program, within-field variation in the context of our instrumental variables setup to estimate the responsiveness of postsecondary educational investment to changes in skill demand.⁴⁰

³⁹ Specifically, it is possible that BGT may have disproportionately collected ads in technology fields (e.g., computer science, engineering, etc.) at the beginning of the analysis sample, but process improvements now do a better job of picking up ads in other fields. This may reflect changing BGT coverage or changes in true posting behavior by field. Either way, this will manifest as slower-than-average demand growth in technology sectors and higher-than-average demand growth in other (e.g., liberal arts) sectors.

⁴⁰ Appendix Figure A2 presents a graph similar to Figure 3, showing the variation in the instrument used to predict program-specific demand changes, after netting out field-by-base-year fixed effects. Although this nets out much of

B. Main Results and Proximate Mechanisms

Table 2 reports our preferred 2SLS estimates of the responsiveness of postsecondary investment to labor market demand, pooling across fields, institutions, and long differences. We measure postsecondary investment in two ways: degrees awarded (Panel A), which is available for all institutions from IPEDS; and undergraduate credits taken (Panel B), which is available for 114 institutions participating in multiple years of The Cost Study.

We report estimates from our base model, which includes major-specific fixed effects, in column 1. This specification lets the average 5-year change in the outcome differ across fields, but restricts the field-specific trend to be the same across each of the three long differences stacked in the regression. The first stage is highly significant, with an F-statistic exceeding 100.⁴¹ We find a large and positive relationship between demand and degrees granted: a one percent increase in the number of job ads for a given program over five years results in a 0.98 percent increase in degrees granted in that program over seven years (95% CI: 0.54 to 1.41). Our preferred estimate, in column 2, lets each stacked difference have its own major-specific time trend. Although the strength of the first stage is slightly reduced, the F-statistic still exceeds 100; moreover, the second stage point estimate is now about one-quarter larger, with an elasticity closer to 1.3 (95% CI: 0.75 to 1.76).

Columns (3) and (4) control for the relative desirability of the focal institution in two different ways, as described earlier. In (3), we control for the change in the number of other degrees granted at the institution and in (4) we control for institution-by-base-year fixed effects. Point estimates are qualitatively similar to our preferred estimates in column (2). The finding of elastic response to changes in program-specific demand is robust to various alternative specifications, including not weighting by program size, dropping very small programs, or including state or state-by-base-year fixed effects (Appendix Table A5). These latter specifications address the concern that location-specific shocks (e.g., population growth) may affect both degree supply and skill demand. Results are also qualitatively similar when the

the cross-field variation at the median, by construction, there remains substantial variation in predicted demand *across* programs in the same field. In addition, Appendix Figure A3 and Appendix Table A4 show that meaningful levels of variation in our instrument remain even after we condition on increasingly stringent vectors of fixed effects, culminating with our preferred specification that includes major-by-base-year fixed effects.

⁴¹ We find that, at the program level, a 1-log-point increase in the instrument is associated with a 2.84-log-point increase in the number of ads. Note that the instrument is a stock (total employment) whereas the outcome and endogenous variables are “flows” (new graduates, job openings) so the scale of first-stage coefficients is not comparable to that of the second-stage coefficients.

instrument is constructed using the total number of workers who majored in m and work in industry i divided by the total number of workers who majored in m to define industry-major mappings.

As a point of comparison, column (5) presents OLS estimates. The implied elasticity of degrees granted to field-specific demand is precisely estimated and close to zero, with confidence intervals implying that elasticities as small as 0.1 in magnitude are still rejected as being too large. As described above, these patterns likely mischaracterize the causal effect of labor market demand on postsecondary investment for a number of reasons; we therefore prefer the 2SLS estimates.

In Panel B we examine undergraduate credit-taking for the institutions participating in The Cost Study. The responsiveness to credit-taking is similar to that estimated for bachelor's degrees awarded. A 1 percent increase in the number of job ads for a given program leads to between a 0.9 and 1.3 percent increase in undergraduate credits taken, depending on the fixed effects included. The much smaller sample size reduces power (and precision), but the first stage still has a reasonable F-statistic, between 16 and 21 depending on specification.⁴² Again, Appendix Table A5 shows that point estimates are similar for alternative specifications, albeit less precise (particularly when including state fixed effects, which is very demanding on the data).

A key modeling choice is how to specify the dynamics of treatment effects, or how sensitive estimates are to plausible alternative lag structures. We desire to measure demand shifts that are perceived as sufficiently persistent so as to solicit an investment response. We also recognize that such a response may take variable lengths of time because of when students choose their majors as well as supply-side constraints on the part of institutions.

Our base model quantifies the relationship between demand shifts over five years and changes in outcomes measured over seven. In Figure 4 we present 2SLS point estimates, analogous to column (2) of Table 2, from specifications that alter both the length of the demand change (t_0 to t_j) and the horizon over which the outcome is measured (t_0 to t_k).⁴³ Panel A shows elasticity estimates for degrees (IPEDS), and Panel B shows elasticity estimates for undergraduate credits taken (Cost Study). We apply color shading to indicate both the magnitude

⁴² Similar to the IPEDS sample, we find in the Cost Study sample that a 1-log-point increase in the instrument is associated with a 2.97-log-point increase in the number of ads at a program level.

⁴³ Note that the number of stacked long differences we are able to include changes across these specifications.

(blue scale) and statistical significance (pink scale) of the estimated elasticities. Moving horizontally reveals that effects generally become larger in magnitude when the outcome difference is measured over a greater time horizon, as students and schools have more time to respond. Importantly, our base specification of a five-year treatment difference and seven-year outcome difference does not appear unusual relative to other similar time frames, as point estimates in adjacent cells are quite close to our preferred specification. Estimates for undergraduate credits (panel B) show similar patterns but are less precise.

To investigate proximate mechanisms for the supply response, we turn to the department-level information on course-taking, faculty, and sections available in The Cost Study. Given the similarity between degrees and credit-taking estimates reported in Table 2, we view this analysis as broadly representative of the IPEDS universe.⁴⁴

Table 3 presents 2SLS estimates of the effects of changes in skill demand on intermediate outcomes. We find shifts in both upper- and lower-division coursework (Panel A), with similar elasticities (about 1.0 to 1.2). These increases in credits appear to be accommodated by growth in non-tenure-track faculty and increases in the number of credits they teach (Panels C and D), although these estimates are somewhat noisy. Moreover, we fail to detect meaningful movements in the number of upper-level course sections (Panel B)—which, combined with the prior results, suggests larger upper-level undergraduate courses. We do see a modest increase in the number of lower-division course sections, but estimates are imprecise.

C. Validity and Robustness of the IV Estimator

As discussed in the previous section, we check that 2SLS estimates derived from a shift-share instrument are not unduly driven by just a few shares, as this calls into question their exogeneity. In Appendix Figure A4, we plot the series of 2SLS estimates that emerge from an exercise in which we leave out select industries with the largest Rotemberg weights (Goldsmith-Pinkham et al., 2020), both positive and negative, when constructing our instrument and subsequent elasticity estimates. We see two encouraging patterns. First, no single industry is critical to our ability to predict job ad changes (demand) with our instrument, as the first-stage F-statistic exceeds 60 in all cases and is fairly close to our benchmark F-statistic ($F = 111$) in

⁴⁴ We have also estimated an IPEDS degree specification on a sample restricted to programs and institutions available in The Cost Study; we did not find meaningfully different estimates from our baseline specification in Table 2, Panel A.

most cases. Second, the elasticity produced when dropping any one of these industries is close to our preferred estimate; the one, modest exception occurs when we drop the vehicle manufacturing industry, in which case the elasticity estimate rises to above 2. We thus conclude that the variation isolated by our instrument reflects industry-related changes in employment broadly and is not driven by any single sector.

We can also assess the validity of our instrument by exploring whether it is systematically related to outcome changes in periods that entirely precede the changes in skill demand captured by the instrument. That is, we would not expect—for example—changes in effective demand for Biology majors at a given institution between 2010 and 2015 to predict changes in Biology degrees produced at that institution between 2003 and 2010. Such an analysis falls under the consideration of “pre-trends” discussed by Goldsmith-Pinkham et al. (2020). Appendix Table A6 presents results from such falsification regressions, estimated using an analogous stacked long-differences approach. Estimates of the association between our instrumented demand measure and outcomes that capture pre-period degrees (columns 1-2) or undergraduate credits (columns 3-4) are small and insignificant. We view these falsification results as broadly supportive of the assumption that our instrument is indeed carving out plausibly exogenous variation in skill demand.

D. Why Are Some Programs More Responsive Than Others?

Combining all programs, we have found that postsecondary investment is quite responsive, on average, to changes in labor market demand.⁴⁵ However, the conceptual framework suggests that human capital investment in some programs will be more responsive to changes in demand than in others due to differences in supply-side constraints or differences in student preferences, which influence the valuation of labor market demand. We thus examine heterogeneity in such responsiveness by institutional, program, and student characteristics in effort to illuminate the ways that supply constraints versus student preferences mediate the overall response.

Supply-side features may impose different constraints across the diverse public and private institutions in our analysis. Gilpin, Saunders, and Stoddard (2015) hypothesize that

⁴⁵ Recent research has shown that 2SLS estimates do not necessarily produce local average treatment effects under heterogeneity when specifications include covariates unless these are included in saturated form (Blandhol et al., 2022). In our case, our fixed-effects specification is essentially saturated, and thus the pooled estimate is likely close to a weighted average across programs, although these implicit weights may not be the policy-relevant ones.

structural features differentiating for-profit from public institutions—faculty composition, governance structure, resources, campus size—influence the greater responsiveness among for-profit schools. Similar differences could apply between public and nonprofit private schools, and across institutions that vary in selectivity or research focus.

In Table 4 we examine differences in responsiveness by institutional characteristics. We find quite similar estimated elasticities between public and private nonprofit institutions (Panel A), 1.15 and 0.99, respectively. Greater differences exist by institutional focus and selectivity (Panels B and C). Highly selective and research-intensive doctoral institutions are nearly unresponsive to changes in labor demand, but less-selective institutions and those offering master’s (but not many doctoral) degrees are very responsive.

These relationships merit further attention, but it seems plausible that less-selective and less-research-intensive institutions face fewer capacity constraints to expansion when demand increases. In addition, such institutions are more dependent on tuition revenue and publicly appropriated funds for their operations, while highly selective and research-intensive institutions rely more heavily on research and endowment funds. Tuition and public appropriations are subject to greater market pressure than are funding streams from research and financial assets, so less-selective and non-doctoral institutions may face stronger incentives to be responsive to labor demand.

The geographically diffuse nature of the labor markets served by more selective institutions (Conzelmann et al., 2022) may also make labor market shocks less salient than for colleges serving narrow markets. Relatedly, graduates from doctoral and more-selective institutions may face weaker earnings premia across different majors (Quadlin, Cohen, & VanHeuvelen, 2021), possibly because of the signaling roles played by college selectivity and major (Hershbein, 2013). Thus, weaker response at more-selective and research-intensive institutions may stem in part from structural market factors (sorting and reputation) as well as education production constraints. The muted responsiveness at these institutions, however, is not due to the omission of online job ads by employers that tend to recruit heavily on such campuses. Indeed, in Appendix Table A7, we confirm that job ads from the vast majority of prestigious consulting firms, banks, and top technology companies appear in our job ad data.⁴⁶

⁴⁶ We use the list of prestigious consulting firms and banks provided in Table 1 of Weinstein (2022). Several of the firms that do not appear in our data were previously acquired by other firms or are located outside the United States.

The nature of production and cost structure also differ across fields (Altonji & Zimmerman, 2017; Hemelt et al., 2021b), which may make it easier to expand postsecondary supply in response to an increase in demand in some fields more than others. For instance, many science courses require labs which are difficult to expand quickly. On the other hand, some fields may have excess capacity due to downward enrollment trends combined with employment rigidities (Johnson & Turner, 2009); these fields might more easily accommodate additional graduates if demand increases. There also could be differences in how labor market demand influences student demand. Fields that are closely tied to specific jobs or for which students' pursuits are employment-driven may be more responsive than fields students pursue out of passion.

To investigate this, we estimate 2SLS versions of equation (7) separately for each of ten aggregate fields. Figure 5 plots these major-group-specific estimates along with F-statistics from the accompanying first stages; for comparison, we also indicate our overall pooled estimate across all fields.⁴⁷ We find that the broad fields of communications, social sciences, and health are the most responsive to relevant changes in skill demand, with each of these broad fields having an estimated elasticity greater than 2. Other broad fields, such as education and humanities, are less responsive, with estimated elasticities less than 0.5 and not statistically different from zero. Indeed, besides communications, social sciences, and health, the only broad major group for which we can reject a zero elasticity is engineering (point estimate = 0.92).⁴⁸ Although precision is a concern, even at this aggregated level of heterogeneity, first-stage F-statistics are reasonably large for all groups other than agriculture and arts, both of which are excluded from Figure 5 (but appear in Table 5).

Panel A of Table 5 groups detailed majors into terciles by average program cost, operationalized as instructional expenditures per student credit hour taken from The Cost Study. We find marked responsiveness among programs in the middle- and low-cost terciles, with elasticities of 2.3 and 1.4, respectively. We see no evidence of responsiveness among the high-cost programs (an elasticity estimate near zero). This pattern accords with institutions more

⁴⁷ Panel B of Table 5 presents the analogous point estimates and additional details.

⁴⁸ We observe very similar patterns of results by broad field of study if we drop the highly selective institutions from the analytic sample, which suggests that differences in results by broad field of study are not driven by differences in the mix of fields studied across institutional selectivity.

easily expanding supply in fields with lower costs.⁴⁹ Moreover, we show in Appendix Figure A5 that this result is not simply an artifact of more expensive fields taking longer to respond: we find greater response for less-expensive fields over all reasonable horizons, while elasticities for the group of fields in the upper tercile remain near zero.

The heterogeneity in response by program cost is also informative for understanding the responses by the broad major groups in Figure 5 (as well as panel B of Table 5).⁵⁰ For example, although we find that the broad area of “health” is quite responsive to changes in labor demand, we also see in Appendix Table A8 that most of the specific majors that constitute the “health” group are high-cost fields. Taken together, this suggests that the responsiveness of “health” fields is driven by lower-cost majors such as Dietetics and Clinical Nutrition Services, rather than costly majors such as Nursing. Similarly, while most component majors of the “engineering” group are high-cost fields, two are not: Computer and Information Science and Engineering Technology. These two fields drive the positive elasticity for the broad group of “engineering.” Thus, certain broad major groups—as categorized by both the U.S. Department of Education’s Classification of Instructional Programs and by college units within universities—can contain detailed majors that vary considerably in their cost per credit hour, and this variation may in turn mask heterogeneity in response to labor demand shifts. On the other hand, the vast majority of majors in the “social sciences” and “communications” bucket are low-cost fields, and it makes sense that these broad groups are highly responsive.⁵¹

Although these patterns by cost and field of study imply an important role for supply-side constraints, the response heterogeneity across fields may also reflect, in part, differences in student preferences. For instance, women are overrepresented in the health and communications fields, which are among the most responsive. Table 6 reports response estimates separately by student gender. We find that female students are more responsive (elasticity of 1.61) than male students (elasticity of 1.02), and this difference is statistically different from zero (p-value =

⁴⁹ We are not able to consistently measure marginal cost by field and instead use average cost. However, we also find that “service” fields (those that have a much greater share of credits taught than degrees granted) also are more responsive. These fields may have greater scale economies and thus face lower marginal cost. The rank correlation between the number of credit hours taught per degree and the average cost per credit hour is -0.28, consistent with the notion that these “service” fields are lower cost.

⁵⁰ Appendix Table A8 lists the component majors of each broad major group along with the cost tercile into which each specific major falls.

⁵¹ Public Policy and Public Administration are the exceptions in social sciences; both are relatively less common programs that grant a modest number of bachelor's degrees and are often located in departments with master's programs (i.e., MPP, MPA), which drive up average instructional costs per credit hour.

0.02). This is consistent with Blom, Cadena, and Keys (2021), who find that women’s choice of major is especially responsive to shifts in the national unemployment rate.⁵²

In Figure 6, we report estimated elasticities separately for men and women within each broad field of study (Panel A) and within tiers of institutional selectivity (Panel B). Although limited power precludes us from statistically rejecting their equivalence, the elasticity estimate for women is greater than that for men in nearly all broad fields of study. However, women also disproportionately select fields that exhibit stronger overall responses to labor demand changes. For example, female students represent 85 percent of bachelor’s degrees in health, 63 percent in communications, and 54 percent in social sciences (*Digest of Education Statistics*, 2022).

To understand the importance of such selection in the gender differences in elasticity, we re-estimate equation (7) for females using weights that reflect the proportion of baseline degrees held by *males* in each broad field. The elasticity estimate for females falls from 1.6 to about 1.2, implying an important role for field-specific, supply-side features in shaping human capital responses to changes in labor demand.⁵³ However, even reweighted, the female-specific elasticity remains larger than that for men, indicating that demand-side preferences also play a role.

Thus, supply-side constraints and demand-side preferences interact to produce the realized response to changes in skill demand.

V. Conclusion

In this paper, we investigate how educational investment by postsecondary institutions and their students responds to labor demand shocks that are specific to each institution and field of study. Using millions of online job ads, we characterize changes in labor market demand for individual majors at nearly all U.S. 4-year public and private nonprofit postsecondary institutions between 2010 and 2017. Institutions vary considerably in the labor markets in which their students work after graduation, and these labor markets are differentially affected by demand shifts based on their pre-existing industry mixes. We exploit this cross-sectional variation, along with industry-major mappings, to develop an instrument for institution-major labor demand shocks that we use to isolate demand-driven variation in job postings, a salient signal of employment opportunities to college graduates and their colleges.

⁵² Specifically, Blom, Cadena, and Keys (2021) find that during recessions women are relatively more likely than men to choose “gender-atypical” majors and fields with lower average grades.

⁵³ If we instead run the reverse exercise and reweight the regression for males using weights that correspond to baseline degree shares for females, the male-specific elasticity estimate increases, but by a relatively modest amount.

Using this variation, we find that the number of bachelor's degrees granted by postsecondary programs responds robustly to changes in major-specific demand, with an average elasticity of about 1.3. Moreover, department-level data show a nearly identical elasticity in the number of credits taken, as well as suggestive evidence that the number of non-tenure-track faculty and the number of credits they teach each rise, both corroborating the overall estimate and illustrating a prominent mechanism through which response occurs. Heterogeneous responses across types of institutions and programs also highlights the importance of supply-side constraints in mediating how responsive postsecondary investment is to changes in labor market demand. Less-selective and non-doctoral institutions evince elasticities higher than average while more-selective and doctoral institutions exhibit negligible responses. We also find that lower-cost fields are relatively more responsive to shifts in labor demand—with programs in communications, social sciences, and health being particularly responsive. Finally, we show that female students are more responsive than their male peers, even when accounting for their choice of majors. This suggests that latent differences in preferences are also important in understanding responsiveness.

Our results reinforce the small body of prior evidence showing that postsecondary investment in the 2-year sector is moderately responsive to changes in labor market demand (Acton, 2020; Gilpin et al., 2015; Grosz, 2022). Importantly, we show that the large 4-year sector—which represents nearly two-thirds of degree-seeking undergraduate enrollment and over 80 percent of expenditures—is also quite responsive over the medium term.⁵⁴ This core finding counters one major critique of the 4-year sector—namely, that it does not adequately prepare students for work (Chamorro-Premuzic & Frankiewicz, 2019), with adverse consequences for productivity and U.S. economic prosperity. Weinstein (2021) also finds the production of bachelor's degrees responded to large, localized, sector-specific shocks in four fields. Our study is the first to show that this pattern generalizes to demand shifts averaged across all locales and fields (albeit with heterogeneity) and even when those shifts are not as nationally salient as the fracking boom or dot-com bust that Weinstein studies.

Whether one views the patterns of responsiveness we document among the 4-year sector of higher education in a positive or negative light turns on normative questions about the role and

⁵⁴ Digest of Education Statistics 2021, Tables 301.10, 303.25 and 334.10, modified by authors' calculations of IPEDS data to adjust for associate's degrees granted by 4-year colleges.

mission of postsecondary institutions. Such normative tensions cannot be resolved by the tools we employ. However, our findings can feed into broader discussions among policy and institutional leaders as they seek to balance the multiple missions of higher education in society. Businesses and employers will naturally encourage colleges to produce graduates who can fill specific roles or jobs. However, many stakeholders believe colleges have broader obligations to students, parents, and society—especially in the case of public institutions. Those obligations may extend beyond the development of industry-specific skills, critical thinking capacities, and the ability to interpret and weigh evidence in decision-making to the development of social and communication skills necessary to function productively in a pluralistic society—as well as the cultivation of an appreciation for art, culture, music, and the human condition more broadly.⁵⁵

Even when viewed solely through an economic lens, the *a priori* optimal response is not clear. An elasticity of 1, for instance, may be suboptimal for several reasons. First, new domestic college graduates do not fill all job openings for college graduates; rather, some job openings are taken by experienced college graduates switching jobs (and occupations), while others are taken by immigrant college graduates whose degrees were earned in other countries. An increase in total demand in a given major of 1 percent, therefore, could rationally be met with an increase in supply of greater than 1 percent among new college graduates when their share of new positions filled is relatively small.

The cobweb models of Freeman (1976) and others provide another reason that elasticities away from 1 could be optimal and could diverge across fields. Long training times—or supply constraints in producing more graduates or courses taken—can lead to lumpy responses as the lag structures between shock and response change. Indeed, Freeman (1976) finds for engineers long-run elasticities nearly twice the size of short-run elasticities, and this was during the Cold War when federal financial support for engineering programs was much higher than today.

Determining the optimal rate of postsecondary response, how this differs with field and institutional characteristics, and the conditions that moderate institutions' ability to respond are important directions for future research. Regardless, we have shown that both supply-side constraints and demand-side preferences are important in shaping human capital responses to

⁵⁵ Nussbaum (2016) offers a powerful argument for the role of the humanities in developing young adults with the capacity to empathize, think critically about issues of the day, and productively participate in healthy democracies.

skill demand. Policy efforts that aim to align educational investment with labor demand may struggle to achieve such goals if they target only one side of the market.

References

- Acemoglu, D., & Autor, D. (2011). Chapter 12—Skills, Tasks and Technologies: Implications for Employment and Earnings. In D. Card & O. Ashenfelter (Eds.), *Handbook of Labor Economics* (Vol. 4, pp. 1043–1171). [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
- Acton, R. (2020). Community College Program Choices in the Wake of Local Job Losses. *Journal of Labor Economics*, 712555. <https://doi.org/10.1086/712555>
- Altonji, J. G., Arcidiacono, P., & Maurel, A. (2016). Chapter 7—The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects. In E. A. Hanushek, S. Machin, & L. Woessmann (Eds.), *Handbook of the Economics of Education* (Vol. 5, pp. 305–396). Elsevier. <https://doi.org/10.1016/B978-0-444-63459-7.00007-5>
- Altonji, J. G., Blom, E., & Meghir, C. (2012). Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers. *Annual Review of Economics*, 4(1), 185–223. <https://doi.org/10.1146/annurev-economics-080511-110908>
- Altonji, J., & Zimmerman, S. (2017). *The Costs of and Net Returns to College Major* (w23029). National Bureau of Economic Research. <https://doi.org/10.3386/w23029>
- Atalay, E., Phongthientham, P., Sotelo, S., & Tannenbaum, D. (2018). The Evolving U.S. Occupational Structure. *Working Paper*.
- Autor, D., Goldin, C., & Katz, L. (2020). *Extending the Race between Education and Technology* (w26705). National Bureau of Economic Research. <https://doi.org/10.3386/w26705>
- Autor, D. H., Katz, L. F., & Kearney, M. S. (2008). Trends in U.S. Wage Inequality: Revising the Revisionists. *The Review of Economics and Statistics*, 90(2), 300–323.
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333.
- Baker, R., Bettinger, E., Jacob, B., & Marinescu, I. (2018). The Effect of Labor Market Information on Community College Students' Major Choice. *Economics of Education Review*, 65, 18–30. <https://doi.org/10.1016/j.econedurev.2018.05.005>
- Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute; <http://www.jstor.org/stable/j.ctvh4zh1q>
- Betts, J. R. (1996). What do students know about wages? Evidence from a study of undergraduates. *The Journal of Human Resources*, 31(1), 27–56.
- Blair, P. Q., & Deming, D. J. (2020). Structural Increases in Demand for Skill after the Great

- Recession. *AEA Papers and Proceedings*, 110, 362–365.
<https://doi.org/10.1257/pandp.20201064>
- Blandhol, C., Bonney, J., Mogstad, M., & Torgovitsky, A. (2022). *When is TSLS Actually LATE?* (w29707). National Bureau of Economic Research. <https://doi.org/10.3386/w29709>
- Bleemer, Z., & Mehta, A. S. (2021). *College Major Restrictions and Student Stratification*. EdWorkingPaper 21-502. <http://www.edworkingpapers.com/ai21-502>
- Blom, E., Cadena, B. C., & Keys, B. J. (2021). Investment over the Business Cycle: Insights from College Major Choice. *Journal of Labor Economics*, 39(4), 1043–1082.
<https://doi.org/10.1086/712611>
- Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, 89(1), 181–213. <https://doi.org/10.1093/restud/rdab030>
- Bound, J., Braga, B., Golden, J. M., & Khanna, G. (2015). Recruitment of Foreigners in the Market for Computer Scientists in the United States. *Journal of Labor Economics*, 33(S1), S187–S223. <https://doi.org/10.1086/679020>
- Carranza, Juan Esteban, Ferreyra, Maria Marta, & Gazmuri, Ana Maria. (2023). *The Dynamic Market for Short-Cycle Higher Education Programs*. EdWorkingPaper 23-712.
<https://www.edworkingpapers.com/ai23-712>
- Chamorro-Premuzic, T., & Frankiewicz, B. (2019, January 7). Does Higher Education Still Prepare People for Jobs? *Harvard Business Review*.
<https://hbr.org/2019/01/does-higher-education-still-prepare-people-for-jobs>
- Cohen, P. (2016, February 22). A Rising Call to Promote STEM Education and Cut Liberal Arts Funding. *The New York Times*.
<https://www.nytimes.com/2016/02/22/business/a-rising-call-to-promote-stem-education-and-cut-liberal-arts-funding.html>
- Conzelmann, J. G., Hemelt, S. W., Hershbein, B., Martin, S. M., Simon, A., & Stange, K. M. (2022). *Grads on the Go: Measuring College-Specific Labor Markets for Graduates* (w30088). National Bureau of Economic Research. <https://doi.org/10.3386/w30088>
- Deming, D. J. (2017). The Growing Importance of Social Skills in the Labor Market*. *The Quarterly Journal of Economics*, 132(4), 1593–1640. <https://doi.org/10.1093/qje/qjx022>
- Deming, D. J., & Noray, K. (2020). Earnings Dynamics, Changing Job Skills, and STEM Careers. *The Quarterly Journal of Economics*, 135(4), 1965–2005.


- <https://doi.org/10.1093/qje/qjaa021>
- Deming, D., & Kahn, L. B. (2018). Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals. *Journal of Labor Economics*, 36(S1), S337–S369. <https://doi.org/10.1086/694106>
- Digest of Education Statistics, 2022*. (2022). National Center for Education Statistics. https://nces.ed.gov/programs/digest/d22/tables/dt22_322.40.asp
- Ersoy, F. Y., & Speer, J. (2022). *Opening the Black Box of College Major Choice: Evidence from an Information Intervention*. EdWorkingPaper 22-666. <https://doi.org/10.26300/839G-W730>
- Foote, A., & Grosz, M. (2020). The Effect of Local Labor Market Downturns on Postsecondary Enrollment and Program Choice. *Education Finance and Policy*, 15(4), 593–622. https://doi.org/10.1162/edfp_a_00288
- Freeman, R. B. (1975). Supply and Salary Adjustments to the Changing Science Manpower Market: Physics, 1948-1973. *The American Economic Review*, 65(1), 27–39.
- Freeman, R. B. (1976). A Cobweb Model of the Supply and Starting Salary of New Engineers. *Industrial and Labor Relations Review*, 29(2), 236–248.
- Gilpin, G. A., Saunders, J., & Stoddard, C. (2015). Why has for-profit colleges' share of higher education expanded so rapidly? Estimating the responsiveness to labor market changes. *Economics of Education Review*, 45, 53–63. <https://doi.org/10.1016/j.econedurev.2014.11.004>
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8), 2586–2624. <https://doi.org/10.1257/aer.20181047>
- Grosz, M. (2022). Do Postsecondary Training Programs Respond to Changes in the Labor Market? *Journal of Human Capital*. <https://doi.org/10.1086/722264>
- Grubb, W. N. (1996). *Working in the Middle: Strengthening Education and Training for the Mid-Skilled Labor Force* (1st edition). Jossey-Bass.
- Hansen, M. (2021, May 18). The U.S. Education System Isn't Giving Students What Employers Need. *Harvard Business Review*. <https://hbr.org/2021/05/the-u-s-education-system-isnt-giving-students-what-employers-need>

- Hartman, J., Hemelt, S. W., Hershbein, B. J., Sotherland, N., & Stange, K. M. (2022). *The Latent Demand for College Majors*. Working paper.
- Hastings, J., Neilson, C., & Zimmerman, S. (2015). *The Effects of Earnings Disclosure on College Enrollment Decisions* (w21300). National Bureau of Economic Research. <https://doi.org/10.3386/w21300>
- Hemelt, S. W., Hershbein, B., Martin, S. M., & Stange, K. M. (2021a). *College Majors and Skills: Evidence from the Universe of Online Job Ads* (w29605). National Bureau of Economic Research. <https://doi.org/10.3386/w29605>
- Hemelt, S. W., Stange, K. M., Furquim, F., Simon, A., & Sawyer, J. E. (2021b). Why Is Math Cheaper than English? Understanding Cost Differences in Higher Education. *Journal of Labor Economics*, 39(2), 397–435. <https://doi.org/10.1086/709535>
- Hershbein, B. (2013). *Worker Signals among New College Graduates: The Role of Selectivity and GPA*. Upjohn Institute Working Papers. <https://doi.org/10.17848/wp13-190>
- Hershbein, B., & Kahn, L. B. (2018). Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings. *American Economic Review*, 108(7), 1737–1772. <https://doi.org/10.1257/aer.20161570>
- Hoxby, C. M. (2009). The Changing Selectivity of American Colleges. *Journal of Economic Perspectives*, 23(4), 95–118.
- Johnson, W. R., & Turner, S. (2009). Faculty without Students: Resource Allocation in Higher Education. *Journal of Economic Perspectives*, 23(2), 169–189. <https://doi.org/10.1257/jep.23.2.169>
- Kumar, D. (2021, July 5). Florida law offering ‘BOGO’ tuition for STEM majors raises faculty concerns. *Tampa Bay Times*. <https://www.tampabay.com/news/education/2021/07/05/florida-law-offering-bogo-tuition-for-stem-majors-raises-faculty-concerns/>
- Long, M. C., Goldhaber, D., & Huntington-Klein, N. (2015). Do completed college majors respond to changes in wages? *Economics of Education Review*, 49, 1–14. <https://doi.org/10.1016/j.econedurev.2015.07.007>
- Ma, J., & Pender, M. (2022). *Trends in College Pricing and Student Aid 2022*. College Board.
- Nussbaum, M. (2016). *Not for Profit: Why Democracy Needs the Humanities*. Princeton University Press.

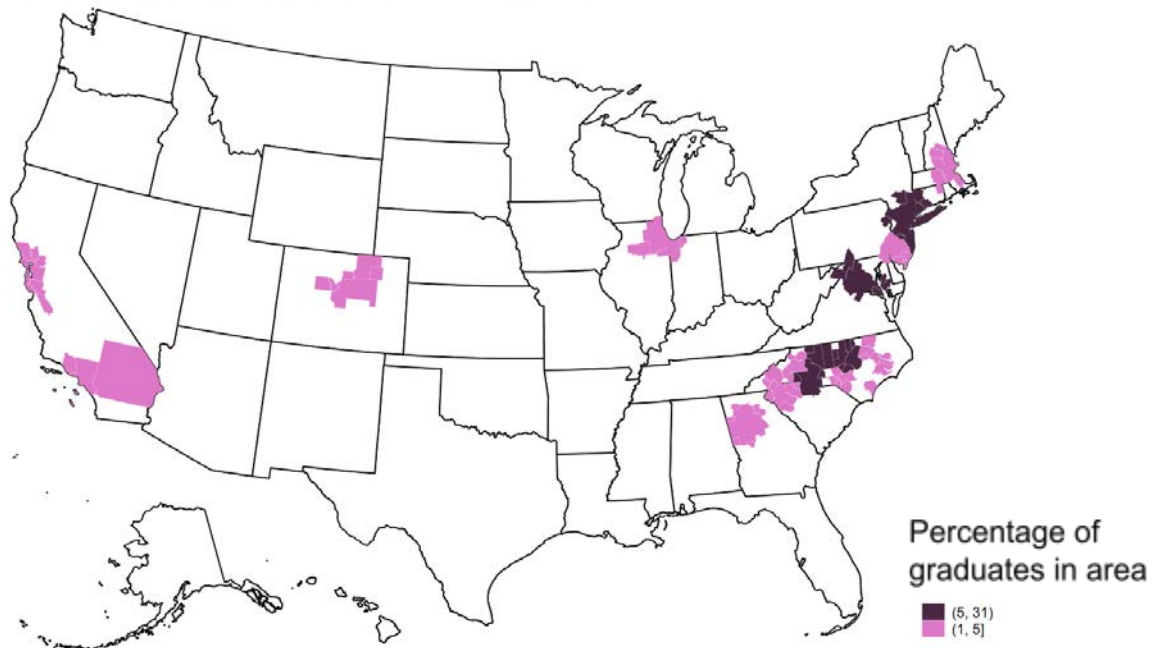
- Quadlin, N., Cohen, E. D., & VanHeuvelen, T. (2021). Same Major, Same Economic Returns? College Selectivity and Earnings Inequality in Young Adulthood. *Research in Social Stratification and Mobility*, 75(100647). <https://doi.org/10.1016/j.rssm.2021.100647>
- Ryoo, J., & Rosen, S. (2004). The Engineering Labor Market. *Journal of Political Economy*, 112(S1), S110–S140. <https://doi.org/10.1086/379946>
- Siow, A. (1984). Occupational Choice under Uncertainty. *Econometrica*, 52(3), 631. <https://doi.org/10.2307/1913468>
- Stange, K. (2015). Differential Pricing in Undergraduate Education: Effects on Degree Production by Field. *Journal of Policy Analysis and Management*, 34(1), 107–135.
- Svrluga, S. (2018, November 18). For universities in Virginia, Amazon’s HQ2 came at the perfect moment. *The Washington Post*. <https://www.washingtonpost.com/education/2018/11/18/universities-virginia-amazons-hq2-came-perfect-moment/>
- Thomas, J. (2022). What Do Course Offerings Imply about University Preferences? *Journal of Labor Economics*. <https://doi.org/10.1086/722087>
- Webber, D. A. (2014). The lifetime earnings premia of different majors: Correcting for selection based on cognitive, noncognitive, and unobserved factors. *Labour Economics*, 28, 14–23. <https://doi.org/10.1016/j.labeco.2014.03.009>
- Weinstein, R. (2022). Firm Decisions and Variation across Universities in Access to High-Wage Jobs: Evidence from Employer Recruiting. *Journal of Labor Economics*, 40(1), 1–46. <https://doi.org/10.1086/714367>
- Wiswall, M., & Zafar, B. (2015). How Do College Students Respond to Public Information about Earnings? *Journal of Human Capital*, 9(2), 117–169. <https://doi.org/10.1086/681542>
- Zarkin, G. A. (1985). Occupational Choice: An Application to the Market for Public School Teachers. *The Quarterly Journal of Economics*, 100(2), 409. <https://doi.org/10.2307/1885389>

Figure 1. Geographic Distribution of North Carolina 4-Year College Graduates, by Control

A. North Carolina 4-year Public

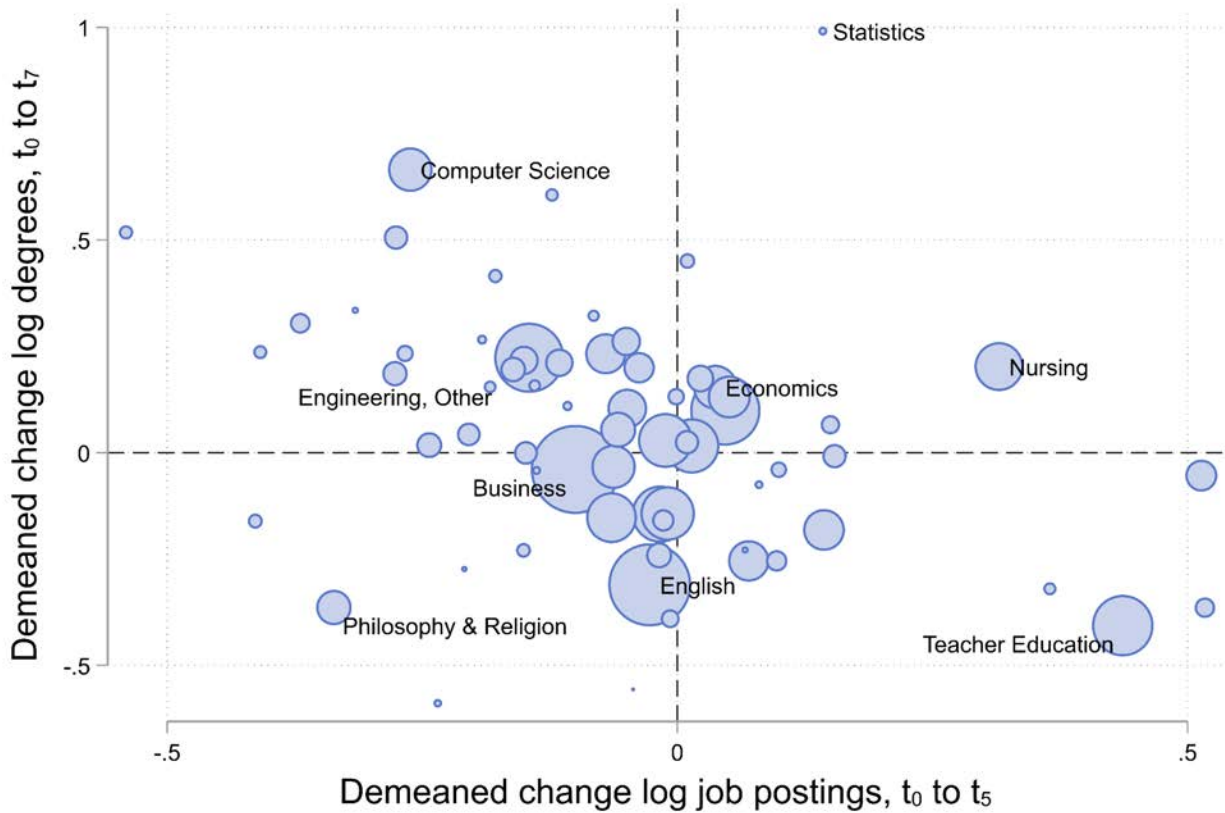


B. North Carolina 4-year Private Nonprofit




Notes: Only areas accounting for a high (greater than 5 percent) or moderate (1 to 5 percent) percentage of graduates from North Carolina 4-year colleges are highlighted on the maps, though all areas with alumni are used in the analysis. Data on the destinations of college graduates come from Conzelmann et al. (2022) and roughly reflect bachelor's degree graduates from the classes of 2010 through 2018.

Figure 2. Changes in Bachelor's Degrees Granted and Demand, by Field of Study



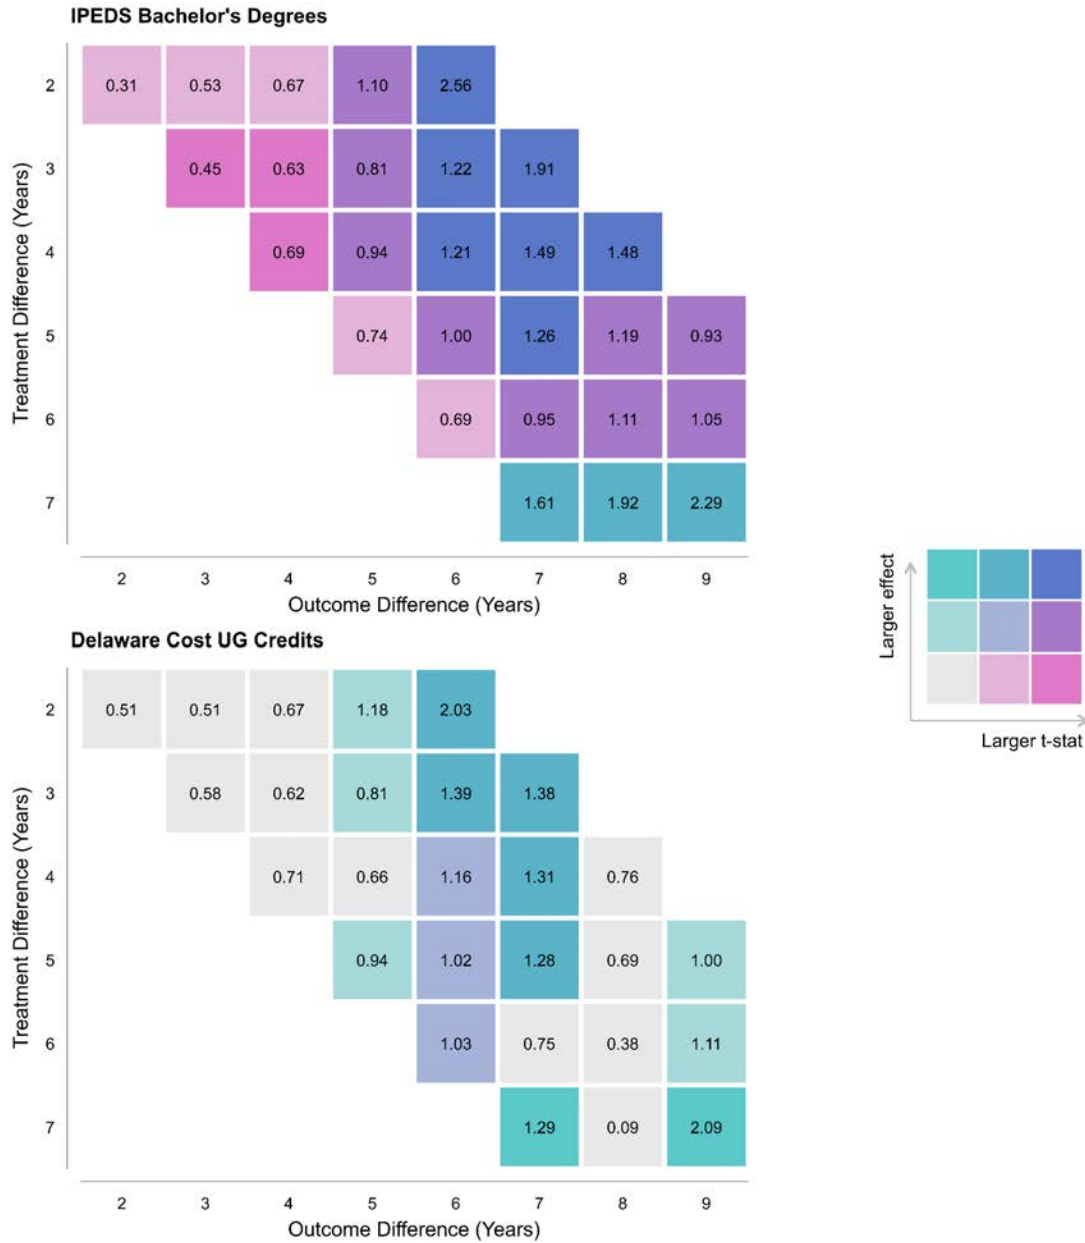
Notes: For each major and long-difference period, we compute the change in each log measure across all institutions over the specified time period. We then subtract the average of these changes, where this average is weighted by the baseline number of degrees granted. We then average across each long-difference period to yield a single ordered pair for each major. The x-axis thus plots, for each of 66 majors, the average, demeaned change in the log of job postings over three, stacked 5-year horizons, and the y-axis plots for the same majors the average, demeaned change in the log of degrees granted over three, stacked 7-year horizons. Marker size is proportional to the average number of degrees granted in the baseline years.

Figure 3. Cross-Institution Variation in Demand Shifts by Field of Study




Notes: The figure plots selected quantiles of the raw distributions of changes in log job ads, where each observation is a long difference of program-specific changes in the log of the demand measure (from t_0 to t_0+5), weighted by base-year degrees. The x-axis plots the range of program-specific changes in the log of the demand measure (from t_0 to t_0+5) within each field of study. This measure includes imputed demand. The figure includes fields offered by at least 200 institutions (which cover roughly two-thirds of the 66 fields in our main analyses).

Figure 4. Dynamics of Degree and Undergraduate Credit Response to Skill Demand Shifts



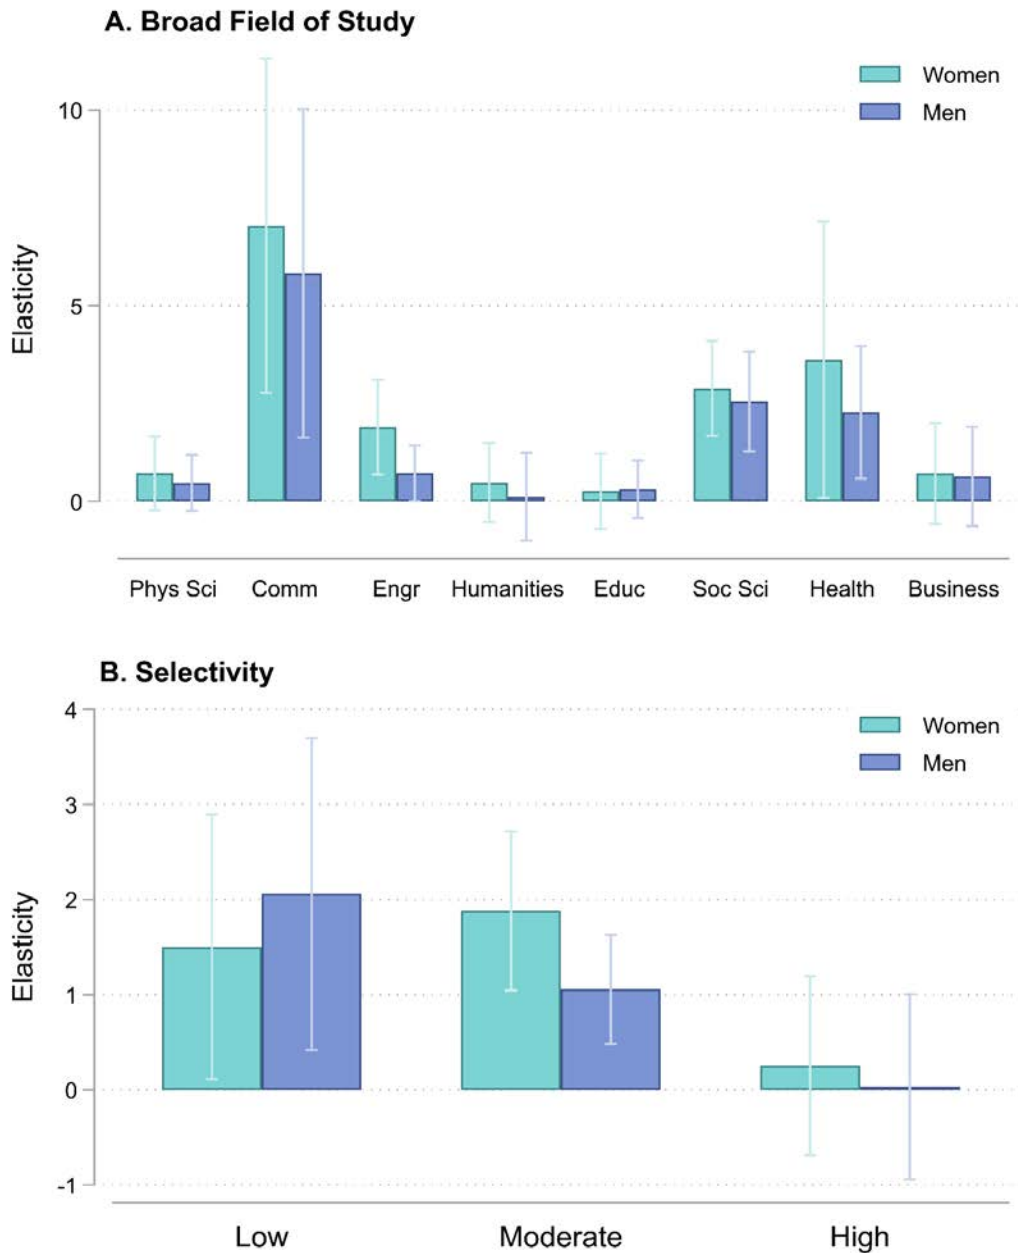
Notes: Figures report 2SLS elasticity estimates for different combinations of timing for the long difference in treatment (vertical axis) and outcome (horizontal axis), analogous to column (2) of Table 2. Our base specification in Table 2 uses a five-year treatment difference and seven year outcome difference. The cutoff values that determine the three categories for elasticities (effect sizes) are 0.8 and 1.2. The t -statistics refer to those from the second-stage coefficient on the treatment variable. The cutoff values for t -statistic categories are 1.96 and 4. UG = Undergraduate.

Figure 5. Bachelor's Degree Responsiveness by Broad Field of Study



Notes: Figure excludes “Agriculture” and “Arts” from the ten field aggregates due to small or null first-stage F-statistics. Estimates come from a 2SLS specification that includes major-by-year fixed effects. Markers are proportional to the average number of degrees awarded in the baseline years. Each field aggregate contains a subset of the 66 total majors we work with in the full sample. Please see Appendix Table A8 for a list of the majors included in each aggregate field.

Figure 6. Gender-Specific Bachelor’s Degree Responsiveness within Broad Field and Institutional Selectivity



Notes: Error bars represent the 95% confidence interval of the elasticity estimates generated from 2SLS. In Panel A, we run separate models for each broad major group and gender including detailed major-by-year fixed effects in each specification. We weight by the baseline degrees earned by each gender. These exclude “Agriculture” and “Arts” from the ten field aggregates due to small or null first-stage F-statistics. Engr=Engineering. Panel B plots gender-specific elasticity estimates within groups of institutions defined by selectivity: "High" combines the two highest competitiveness categories from the Barron's competitiveness index (Most and Highly competitive); "Moderate" groups the next two Barron's categories (Very competitive and Competitive); "Low" includes all others, including institutions without a Barron's value.

Table 1. Summary Statistics for Analytic Sample

	IPEDS Degree Sample		Delaware Cost Sample	
	Mean	Std. Dev.	Mean	Std. Dev.
<i>Institution Characteristics</i>				
Public	0.362	0.481	0.816	0.389
Most selective	0.109	0.312	0.088	0.284
Moderately selective	0.541	0.498	0.798	0.403
Research university	0.162	0.368	0.518	0.502
Locale: City	0.481	0.500	0.535	0.501
Locale: Suburb	0.240	0.427	0.254	0.437
Locale: Town or rural	0.279	0.449	0.211	0.409
Average FTE: Less than 1,000	0.236	0.425	0.009	0.094
Average FTE: 1,000-4,999	0.491	0.500	0.254	0.437
Average FTE: 5,000 or greater	0.273	0.446	0.737	0.442
Average degrees granted	1,045	1,567	2,652	1,959
Average number of programs offered	21.43	11.90	33.70	8.77
Number of institutions	1,681		114	
<i>Program Outcomes</i>				
Average degrees granted	47	87	112	130
Average UG credits			5,669	5,471
Average low division UG credits			3,431	4,144
Average upper division UG credits			2,131	2,386
Number of unique programs	32,554		3,081	

Notes: Averages for institution characteristics in IPEDS and The Cost Study are based on data from 2010-2019. "Most selective" groups the two highest competitiveness categories from Barron's competitiveness index (Most and Highly competitive). "Moderately selective" groups the next two Barrons' categories (Very competitive and competitive). Excluded selectivity categories include less and non-competitive, special institutions, and those with no Barron's competitiveness value. Research university refers to schools classified by the Carnegie Foundation as having very high or high research activity, and other doctoral institutions.

Sources: IPEDS, The Cost Study at the University of Delaware, authors' calculations.

Table 2. Responsiveness of Educational Investment to Changes in Skill Demand

	2SLS (1)	2SLS (2)	2SLS (3)	2SLS (4)	OLS (5)
Panel A. Outcome = Change in log(4-year degrees awarded), t0 to t0+7					
Change log(ads), t0 to t0+5	0.978*** (0.221)	1.258*** (0.257)	0.794*** (0.202)	1.236*** (0.379)	-0.024 (0.036)
F-stat from first stage	142.68	111.26	111.04	89.88	--
N(program-years)	92,501	92,501	92,175	92,138	92,501
N(institutions)	1,681	1,681	1,570	1,559	1,681
Panel B. Outcome = Change in log(total undergraduate credits), t0 to t0+7					
Change log(ads), t0 to t0+5	1.100** (0.452)	1.277** (0.561)	0.999* (0.516)	0.864 (0.585)	0.170 (0.136)
F-stat from first stage	20.68	16.28	15.77	26.92	--
N(program-years)	6,177	6,177	6,177	6,177	6,177
N(institutions)	114	114	114	114	114
Fixed effects or other controls	Major, Year	Major-by-base-year	Major-by-base-year, Δ Other Degrees	Major-by-base-year, School-by-base-year	Major-by-base-year

Notes: The outcome is the change in the measure designated by each panel for a given program (institution-by-field cell) over a 7-year period for one of three long-difference intervals (i.e., 2010-2017, 2011-2018, or 2012-2019). Degree data come from IPEDS and data on undergraduate credit hours come from The Cost Study. See Table 1 and the text for details on construction of the analytic samples. The change in ads, the key independent variable, is based on an aggregation of job-ad data at the institution-major-year level, weighted by shares of an institution's graduates living and working in areas from which the job ads originate. This demand change is calculated over a 5-year period that uses the same base year as the corresponding outcome horizon (i.e., 2010-2015, 2011-2016, or 2012-2017). Please consult the text for additional information on this measure. All estimates are weighted by the number of degrees in the baseline year. Standard errors, clustered by institution, appear in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 3. Skill Demand Shifts and Intermediate Outcomes: Credits by Level, Course Sections, and Faculty Staffing

	2SLS (1)	2SLS (2)
<u>Panel A. Outcome = Change in log (credit types), t0 to t0+7</u>		
	UG Lower Division	UG Upper Division
Change log(ads), t0 to t0+5	1.167** (0.580)	1.003** (0.481)
F-stat from first stage	17.81	16.29
<u>Panel B. Outcome = Change in log(section types), t0 to t0+7</u>		
	UG Lower Division	UG Upper Division
Change log(ads), t0 to t0+5	0.429 (0.843)	0.142 (0.555)
F-stat from first stage	17.81	16.34
<u>Panel C. Outcome = Change in log(faculty types), t0 to t0+7</u>		
	TT Faculty FTE	Non-TT Faculty FTE
Change log(ads), t0 to t0+5	0.225 (0.358)	0.832 (0.901)
F-stat from first stage	16.26	16.25
<u>Panel D. Outcome = Change in log(credits by faculty types), t0 to t0+7</u>		
	TT Faculty FTE	Non-TT Faculty FTE
Change log(ads), t0 to t0+5	0.438 (0.551)	1.819** (0.915)
F-stat from first stage	16.22	16.27

Notes: The outcome is the change in the measure designated by each panel for a given program (institution-by-field cell) over a 7-year period for one of three long-difference intervals (i.e., 2010-2017, 2011-2018, or 2012-2019). Data come from The Cost Study. See Table 1 and the text for details on construction of the analytic samples. The change in ads, the key independent variable, is based on an aggregation of job-ad data at the institution-major-year level, weighted by shares of an institution's graduates living and working in areas from which the job ads originate. This demand change is calculated over a 5-year period that uses the same base year as the corresponding outcome horizon (i.e., 2010-2015, 2011-2016, or 2012-2017). Please consult the text for additional information on this measure. All specifications include major-by-base-year fixed effects. All estimates are weighted by the number of degrees in the baseline year. Standard errors, clustered by institution, appear in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 4. Heterogeneity in Degree Responsiveness to Skill Demand by Institutional Characteristics

Outcome = Change in log(4-year degrees awarded), t0 to t7	2SLS	First-stage F-stat	N (Institutions)	N (Program-years)
	(1)	(2)	(3)	(4)
<u>Panel A. Change log(ads), t0 to t5 by Institution Control</u>				
Control = Public	1.153*** (0.282)	68.8	609	43,052
Control = Private nonprofit	0.990** (0.435)	76.3	1,072	49,449
<u>Panel B. Change log(ads), t0 to t5 by Institution Type</u>				
Type = Doctoral (high research activity)	-0.448 (0.650)	12.7	107	11,339
Type = Other Doctoral	0.820*** (0.298)	45.1	165	15,225
Type = Master's	1.847*** (0.519)	51.4	604	39,365
Type = Baccalaureate and other	1.203* (0.714)	24.3	805	26,567
<u>Panel C. Change log(ads), t0 to t5 by Institution Selectivity</u>				
Selectivity = High	0.119 (0.483)	18.5	183	12,761
Selectivity = Moderate	1.411*** (0.323)	58.3	909	61,933
Selectivity = Low	1.618** (0.673)	31.4	589	17,806

Notes: All models include major-by-base-year fixed effects and are weighted by the number of 4-year degrees awarded in the base year. Standard errors, clustered by institution, appear in parentheses. See notes to Table 2. Outcome data (degrees) are from IPEDS. Selectivity categories are based on Barron's data: "High" combines the two highest competitiveness categories (Most and Highly competitive); "Moderate" groups the next two categories (Very competitive and Competitive); "Low" includes all others, including institutions without a Barron's value.

*** p<0.01, ** p<0.05, * p<0.1.

Table 5. Heterogeneity in Degree Responsiveness to Skill Demand by Program Costs and Broad Field of Study

Outcome = Change in log(4-year degrees awarded), t0 to t7	2SLS	First-stage F-stat	N (Institutions)	N (Program-years)
	(1)	(2)	(3)	(4)
<u>Panel A. Change log(ads), t0 to t5 by Average Program Costs</u>				
Avg cost per credit hour = Bottom tercile	1.420*** (0.363)	105.2	1,580	52,431
Avg cost per credit hour = Middle tercile	2.300*** (0.534)	48.0	1,583	28,504
Avg cost per credit hour = Top tercile	0.001 (0.338)	53.2	1,214	11,566
<u>Panel B. Change log(ads), t0 to t5 by Broad Field of Study</u>				
Field = Agriculture	-15.00 (49.59)	0.1	671	2,166
Field = Physical Sciences	0.476 (0.348)	43.9	1,308	11,951
Field = Communications	6.469*** (2.069)	13.3	1,093	4,422
Field = Engineering	0.919** (0.388)	40.4	1,160	9,124
Field = Education	0.481 (0.506)	24.4	1,151	4,674
Field = Humanities	0.312 (0.428)	106.8	1,469	11,213
Field = Social Sciences	2.618*** (0.577)	63.5	1,426	22,526
Field = Arts	9.905 (7.110)	1.8	1,283	6,798
Field = Health	3.011** (1.414)	20.3	1,199	7,562
Field = Business	0.574 (0.637)	86.5	1,404	12,065

Notes: Panel A presents estimates from the stacked long differences using our IPEDS sample, allowing the effect of changes in demand on degrees to differ by a major's average instructional costs per credit hour split into terciles (regressions estimated separately by subsample). We compute the major-specific credit hour costs from The Cost Study as total expenditures divided by total credit hours produced using all available data from 1998 to 2010. Panel B presents estimates from the stacked long-differences approach using the IPEDS sample, and divides majors into 10 broad fields of study. All models include *detailed* major-by-base-year fixed effects and are weighted by the number of 4-year degrees awarded in the base year. Standard errors, clustered by institution, appear in parentheses. Please see Appendix Table A6 for the component majors within each broad field grouping from Panel B.

*** p<0.01, ** p<0.05, * p<0.1

Table 6. Heterogeneity in Degree Responsiveness to Skill Demand by Gender


Outcome = Change in log(4-year degrees awarded in gender), t0 to t7	2SLS	First-stage F-stat	N (Institutions)	N (Program-years)
	(1)	(2)	(3)	(4)
<u>Change log(ads), t0 to t5</u>				
Gender = Female	1.615*** (0.310)	102.49	1,662	85,351
Gender = Male	1.023*** (0.261)	95.03	1,649	82,311

Notes: Estimates are from the stacked long-differences approach using our IPEDS sample, allowing the effect of changes in demand on degrees to differ by gender. All models include detailed major-by-base-year fixed effects and are weighted by the number of 4-year degrees awarded within the category of interest in the base year. Sample sizes differ slightly between female and male samples and from Table 2 due to exclusion of programs with zero graduates in the gender. Standard errors, clustered by institution, appear in parentheses.

*** p<0.01, ** p<0.05, * p<0.1


Appendix A. Additional Figures and Tables

Appendix Figure A1. Change in Implied vs. Actual Major Demand via Job Postings, t_0 to t_5



Notes: Marker size is proportional to the number of degrees granted. The figure plots demeaned values of the 5-year change in the log of the demand measure that includes both explicitly stated and imputed majors (x-axis) against demeaned values of the demand measure that only includes explicitly stated majors (y-axis). Figure includes fields with at least 200 programs (institution-major tuples; this covers roughly two-thirds of the 66 fields in our main analyses).


Appendix Figure A2. Cross-Institution Variation in Demand Shifts by Field, Net of Major-by-Year Fixed Effects




Notes: These distributional statistics were calculated by regressing the pooled long differences of program-specific changes in the instrument (from t_0 to t_0+5) on major-by-year fixed effects, weighted by base-year degrees. Then we predict the residual and add back the grand (weighted) mean to calculate the percentiles of interest from the resulting distribution. The measure includes imputed demand. The figure includes fields offered by at least 200 institutions (which cover roughly two-thirds of the 66 fields in our main analyses).

Appendix Figure A3. Variation in Labor Demand Instrument

A. IPEDS Sample

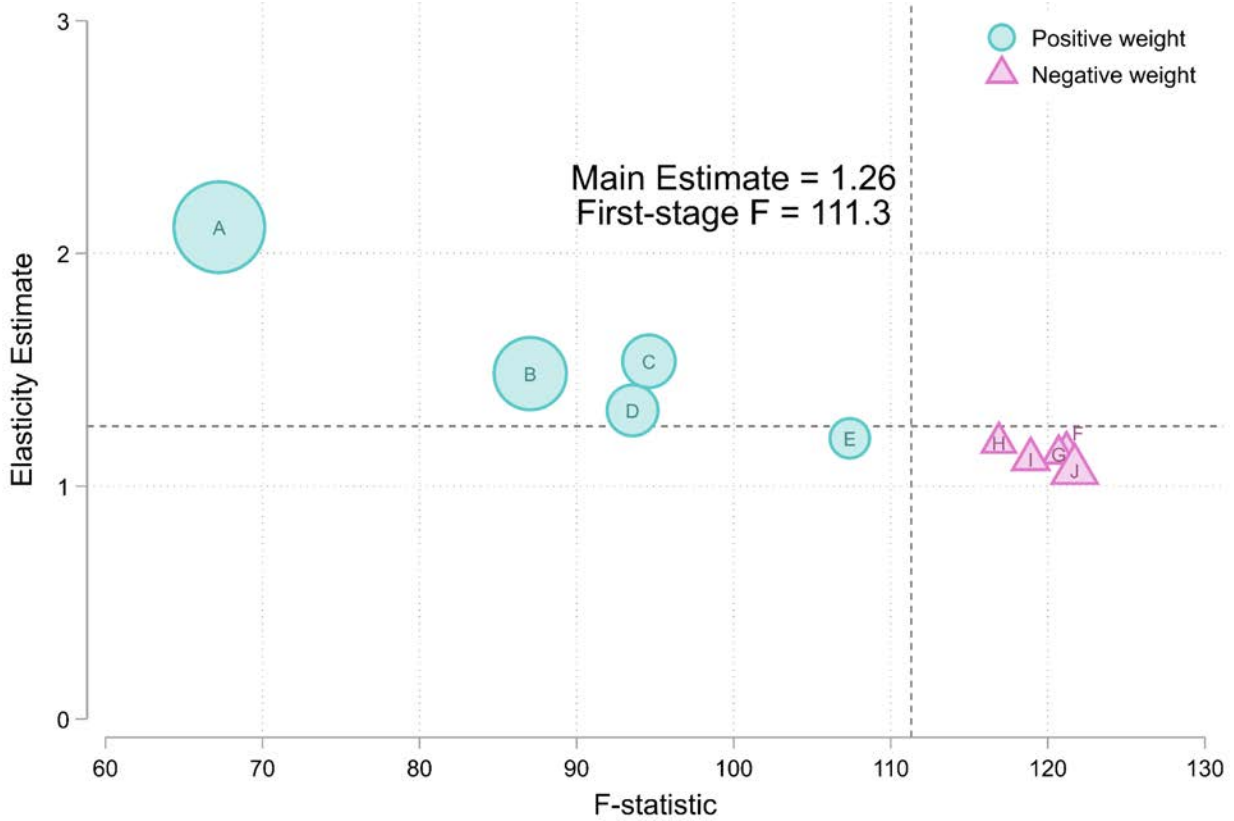


B. Delaware Cost Sample



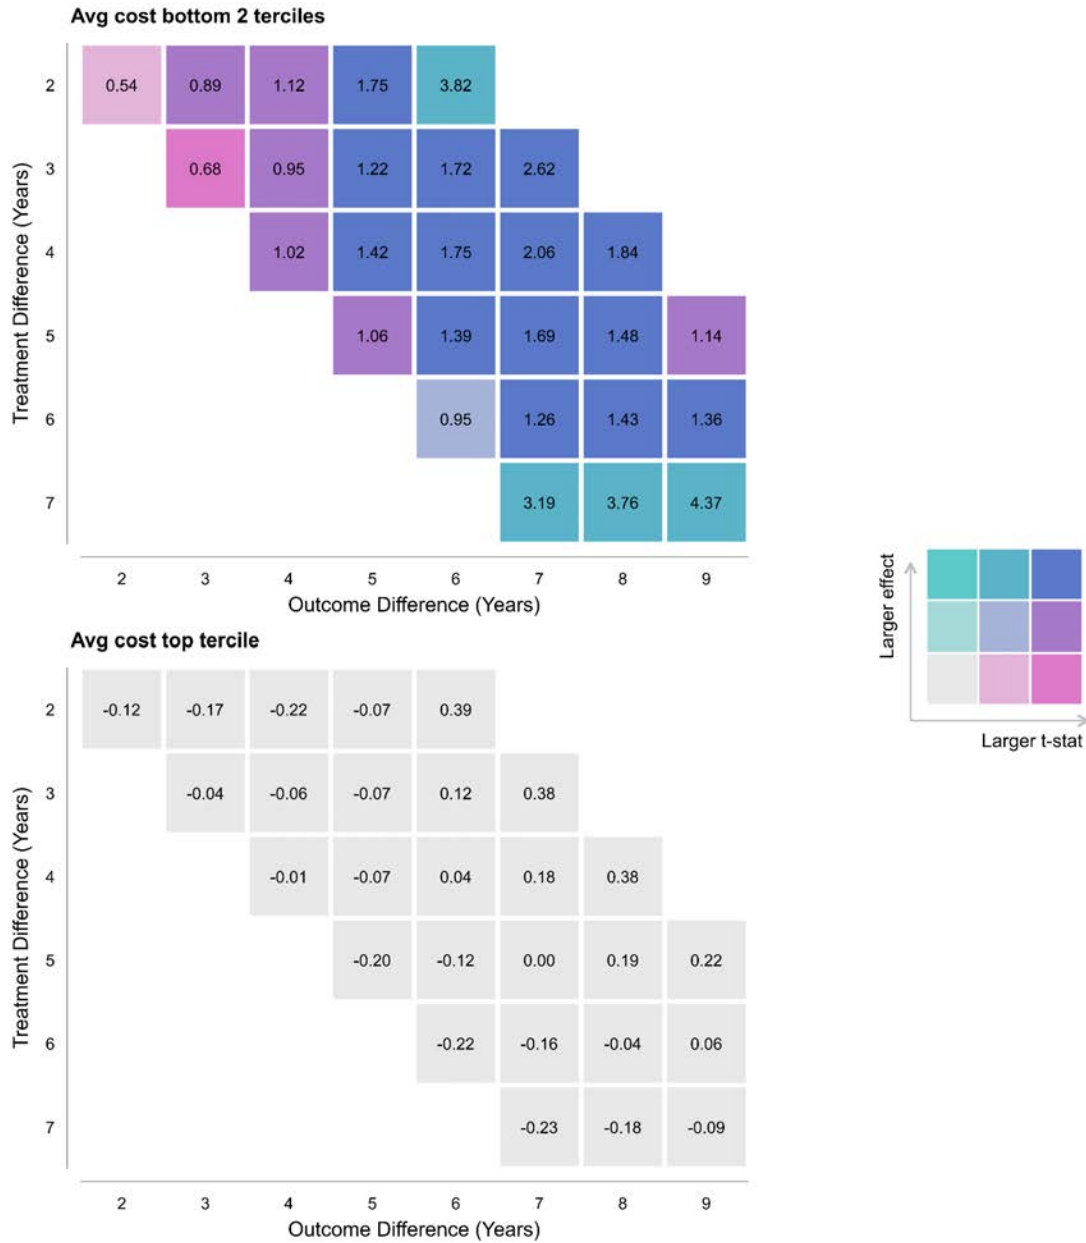
Notes: Panels A and B plot selected percentiles of the shift-share instruments for the IPEDS and DCS samples, respectively. In each panel, the top row presents the percentiles for the unadjusted 5-year instrument from Equation 3. The following 4 rows show the remaining variation after conditioning on year, institution, major, or year-by-major fixed effects. The remainder of each panel repeats this for the 4-, 3-, and 2-year difference instruments in each sample.

Appendix Figure A4. Instrument Diagnostics: Leave-One-Industry-Out Analysis



Notes: Each marker corresponds to a 2SLS elasticity estimate that leaves out the specified industry (indicated by the letter A-J) from the construction of the instrument and subsequent 2SLS estimation. The x-axis indicates the resultant first-stage F-statistic, and the y-axis shows the resultant 2SLS elasticity estimate. Each marker is weighted by the absolute value of the industry’s “alpha” or “Rotemberg” weight calculated per Goldsmith-Pinkham, Sorkin, and Swift (2020) using the full analytic sample and estimates. This analysis highlights the industries with the five largest positive and five largest negative weights. These include: A=Vehicle Manufacturing, B=Employment Services, C=Support activities for Mining, D=General Merchandise Stores, E=Food Services, F=Management and Technical Consulting Services, G=Higher Education, H=Video and disk rental stores, I=Computer system design and services, J=Construction.

Appendix Figure A5. Dynamics of Degree Response to Skill Demand Shifts, by Average Instructional Costs per Credit Hour



Notes: Average instructional costs per undergraduate credit hour were calculated using all available Delaware Cost data from 1998 to 2010. We generate an unweighted ranking of the 66 fields based on these average costs and split the fields into terciles. The top panel shows pooled elasticities of degree production using only fields in the first and second terciles, and the bottom panel shows elasticities using only the most expensive tercile of programs. The cutoff values that determine the three categories for elasticities (effect sizes) are 0.8 and 1.2. The t-statistics refer to those from the second-stage coefficient on the treatment variable. The cutoff values for t-statistic categories are 1.96 and 4.

Appendix Table A1. Occupational Distribution by Sample

	Sample				
	All Postings	At least 1 skill	1 Skill and Education = 16	Educ = 16 At least 1 skill At least 1 major	Analysis Educ = 16 At least 1 skill At least 1 major In Metro CBSAs
Count of unique ads	153,031,199	148,000,000	35,938,213	19,519,480	18,471,199
Count of unique ad-major (4-digit CIP)				32,847,216	31,153,536
% of original sample remaining		96.71%	23.48%	12.76%	12.07%
Mean experience level (years)			3.391	3.649	3.682
<i>Occupation</i>					
Management (11)	11.70%	11.92%	22.22%	21.93%	21.84%
Business/Financial (13)	6.64%	6.80%	14.30%	14.82%	15.02%
Computer/Math (15)	11.54%	11.85%	22.13%	25.23%	25.83%
Architecture/Engineering (17)	3.15%	3.22%	6.70%	9.50%	9.26%
Life/Physical/Social Science (19)	1.00%	1.03%	1.69%	2.04%	1.97%
Community/Social Service (21)	1.09%	1.09%	1.38%	1.40%	1.28%
Legal (23)	0.85%	0.87%	0.41%	0.25%	0.26%
Education/Training/Library (25)	2.49%	2.52%	2.48%	1.31%	1.25%
Arts/Design/Entertainment (27)	2.37%	2.42%	2.53%	2.29%	2.32%
Healthcare Practitioners (29)	12.27%	12.24%	7.58%	8.21%	8.01%
Healthcare Support (31)	2.03%	2.06%	0.01%	0.01%	0.01%
Protective Service (33)	1.00%	0.99%	0.33%	0.22%	0.21%
Food Prep/Serving (35)	3.38%	3.24%	0.24%	0.23%	0.23%
Building/Cleaning/Maintenance (37)	1.11%	1.11%	0.06%	0.04%	0.04%
Personal Care (39)	1.75%	1.75%	0.27%	0.21%	0.20%
Sales (41)	11.76%	12.03%	8.20%	4.37%	4.38%
Office/Admin Support (43)	9.96%	10.17%	4.28%	3.02%	3.02%
Farming/Fishing/Forestry (45)	0.06%	0.06%	0.02%	0.02%	0.02%
Construction/Extraction (47)	0.97%	0.98%	0.09%	0.11%	0.11%
Installation/Maintenance/Repair (49)	2.94%	3.00%	0.31%	0.27%	0.25%
Production (51)	2.45%	2.45%	0.64%	0.56%	0.52%
Transportation/Material Moving (53)	5.81%	4.51%	0.14%	0.09%	0.09%
Military (55)	0.07%	0.07%	0.03%	0.02%	0.02%
Missing (0)	3.61%	3.61%	3.93%	3.84%	3.85%
<i>Sample Restrictions</i>					
Year >= 2010	Y	Y	Y	Y	Y
At least one skill	N	Y	Y	Y	Y
Seeking Bachelor's Degree	N	N	Y	Y	Y
At least one major	N	N	N	Y	Y
Only Metropolitan Statistical Areas	N	N	N	N	Y

Source: Authors' analysis of Burning Glass Technologies (BGT) job postings data.

Note: Occupations are two-digit Standard Occupational Classification (SOC) codes. Unique ad-majors treat ads with multiple majors listed as multiple observations, one for each major listed. Statistics for the last two columns represent this level of observation.

Appendix Table A2. Major Classification Model Performance

A. Model Comparisons: Feature Set 4 & 1% Sample					
	<i>Avg Precision</i>	<i>Avg Recall</i>	<i>Macro F1</i>	<i>Micro F1</i>	<i>LRAP</i>
Standard Logit	0.57	0.50	0.52	0.69	0.845
Penalized Logit	0.65	0.52	0.57	0.70	0.861
SGD Logit	0.72	0.44	0.52	0.69	0.865
Decision Tree	0.45	0.45	0.45	0.60	0.523
Random Forest (preferred)	0.90	0.38	0.50	0.72	0.885
B. Feature Set Comparisons: Random Forest - 1% Sample					
	<i>Avg Precision</i>	<i>Avg Recall</i>	<i>Macro F1</i>	<i>Micro F1</i>	<i>LRAP</i>
Feature Set 1	0.73	0.19	0.27	0.54	0.749
Feature Set 2	0.89	0.34	0.44	0.71	0.877
Feature Set 3	0.85	0.38	0.48	0.72	0.868
Feature Set 4 (preferred)	0.90	0.38	0.50	0.72	0.885
C. Sample Size Comparisons: Random Forest - Feature Set 4					
	<i>Avg Precision</i>	<i>Avg Recall</i>	<i>Macro F1</i>	<i>Micro F1</i>	<i>LRAP</i>
1% Sample	0.90	0.38	0.50	0.72	0.885
3% Sample	0.92	0.48	0.60	0.76	0.904
5% Sample (preferred)	0.93	0.53	0.65	0.79	0.913

Notes: Statistics presented compare performance of algorithms to assign majors across BGT jobs as discussed in the text. Feature set 1 include indicators for six-digit SOC occupation, four-digit NAICS industry, CBSA (metro/micro area), and year-month dummies. Feature set 2 adds a cubic in the number of skills present in the ad as well as indicators for the 1000 most frequently occurring skills. Feature set 3 includes only the 1,000 most predictive unigrams from tokenized text data on job title, employer name, and skill requirements. Feature set 4 adds to feature set 1 the 1,250 most predictive unigrams from tokenized text data on job title, employer name, and skill requirements.

Appendix Table A3. Distribution of Degrees Granted, by Field, 2017–2019

Field of Study	Programs	Number of degrees granted						Share of degrees	
		Mean	p10	p25	p50	p75	p90	Across all institutions	Within institution
English Lib Arts & Humanities	1391	81	7	14	35	90	218	0.062	0.078
Business	1385	136	15	31	65	149	327	0.104	0.143
Psychology	1340	91	10	20	41	100	224	0.067	0.077
Biology	1298	82	8	18	37	83	193	0.059	0.068
Mathematics	1189	22	2	4	9	22	53	0.014	0.016
Other Visual/Performing Arts	1153	40	3	7	16	41	97	0.025	0.046
Computer & Info Science	1114	65	5	10	23	68	162	0.040	0.037
Teacher Education	1104	61	5	13	29	73	161	0.037	0.048
Applied Arts	1092	39	4	7	17	43	96	0.023	0.071
Chemistry	1067	15	2	4	8	17	35	0.009	0.031
Poli Sci/Gov & Intl Relations	1039	47	4	8	19	50	120	0.027	0.013
Communication & Media Studies	1036	66	5	11	26	75	181	0.037	0.047
Accounting	930	53	6	12	26	67	138	0.027	0.041
Sociology	929	33	3	7	15	34	73	0.016	0.023
Philosophy & Religion	918	14	1	3	8	16	28	0.007	0.020
Foreign Language & Linguistics	915	28	2	5	13	33	67	0.013	0.022
Other Social Sciences	868	52	3	7	20	56	130	0.024	0.034
Nursing	767	137	22	46	100	180	273	0.058	0.061
Fitness & Leisure Studies	767	67	9	17	32	78	177	0.028	0.171
Physics	694	13	2	4	8	16	26	0.005	0.032
Economics	693	58	3	8	19	56	157	0.022	0.009
Protective Services	665	73	9	18	37	86	175	0.027	0.072
Natural Resources	616	31	3	6	14	32	74	0.010	0.046
Social Work	615	43	7	13	29	58	96	0.015	0.036
Marketing	591	67	6	14	34	91	168	0.021	0.026
Finance	520	83	7	16	39	109	223	0.024	0.038
Biochem & Molecular Biology	506	23	3	5	10	23	51	0.006	0.049
Allied Health	494	66	6	15	37	76	149	0.018	0.015
Geological & Earth Sciences	410	16	3	6	12	19	33	0.004	0.009
Mgmt Info Systems & Science	371	41	2	7	18	55	105	0.008	0.024
Special Educ & Teaching	359	22	2	6	13	28	51	0.004	0.020
Other Engineering	304	45	3	6	22	49	120	0.007	0.026
Electrical Engineering	303	55	11	20	40	69	112	0.009	0.050
Journalism	299	45	3	7	23	54	115	0.007	0.019
Theology	296	26	2	4	12	30	57	0.004	0.194
Mechanical Engineering	291	115	27	51	94	164	224	0.019	0.042
Family & Consumer Sciences	280	83	6	14	39	112	220	0.013	0.040
Engineering Technology	261	62	5	15	37	81	144	0.009	0.007
Geography	257	19	4	7	13	23	39	0.003	0.048
PR & Advertising	244	60	2	7	22	75	159	0.008	0.023
Human Resources Mgmt & Services	239	40	3	8	20	44	90	0.005	0.029
Civil Engineering	231	59	15	25	47	80	126	0.008	0.014
Other Education	226	21	2	4	10	23	51	0.002	0.029
Computer Engineering	224	40	5	12	24	46	96	0.005	0.025
Health & Medical Admin Services	206	42	5	10	23	52	96	0.005	0.036
Agriculture	170	121	7	15	59	170	334	0.011	0.060
Hospitality Admin/Mgmt	166	69	6	15	35	78	160	0.006	0.026
Legal Studies	159	22	2	4	12	22	47	0.002	0.032
Chemical Engineering	157	71	25	39	63	90	132	0.006	0.024
Rehab & Therapeutic Professions	155	29	4	7	13	41	74	0.002	0.019

Public Health	148	73	7	17	39	81	214	0.006	0.019
Dietetics & Nutrition Services	135	40	8	13	25	61	90	0.003	0.019
Systems Engineering	128	54	9	21	39	69	121	0.004	0.031
Statistics	110	34	5	8	15	37	81	0.002	0.032
Biomedical Engineering	108	59	16	33	53	79	102	0.004	0.022
Architecture	106	42	9	19	39	59	81	0.002	0.016
Public Administration	106	30	2	4	12	36	70	0.002	0.008
Microbiology	77	35	7	13	29	49	65	0.001	0.007
Aeronautical Engineering	58	67	27	37	58	89	122	0.002	0.018
Materials Science & Eng	58	30	8	15	25	43	62	0.001	0.008
Atmospheric Sci & Meteorology	54	11	3	5	10	15	24	0.000	0.005
Other Physical Sciences	54	12	1	2	4	10	24	0.000	0.008
Public Policy	52	34	3	7	20	42	88	0.001	0.020
Pharm Sciences & Admin	26	72	5	14	45	108	146	0.001	0.086
Culinary Arts	17	20	1	4	9	18	37	0.000	0.012
Library Science	13	7	1	3	4	9	19	0.000	0.035

Notes: The table shows statistics relating to the number of institutions in our IPEDS sample providing bachelor's degrees across our categorization of 66 majors. Program counts and degrees awarded are aggregated over the three academic years 2017–2019 and include only those with positive degrees granted in any given year. For example, 1391 institutions awarded at least one bachelor's degree in English over 2017–2019; the average institution awarded 81 degrees over this period, while the median institution awarded 35. Source: IPEDS and authors' calculations.

Appendix Table A4. Variation in the Labor Demand Instrument

	Mean	SD	P10	P25	P50	P75	P90
<u>Panel A: 5-year changes (e.g., 2010-2015)</u>							
Program-years (n = 94,440)	0.0956	0.0146	0.0787	0.0877	0.0959	0.1045	0.1124
w/ Base year FE		0.0138	0.0801	0.0891	0.0957	0.1037	0.1104
w/ Institution FE		0.0134	0.0807	0.0879	0.0955	0.1038	0.1114
w/ Major FE		0.0092	0.0853	0.0898	0.0952	0.1015	0.1069
w/ Base-year-by-major FE		0.0066	0.0886	0.0920	0.0954	0.0991	0.1026
<u>Panel B: 4-year changes</u>							
Program-years (n = 126,280)	0.0779	0.0132	0.0616	0.0700	0.0786	0.0866	0.0930
w/ Base year FE		0.0119	0.0645	0.0726	0.0785	0.0846	0.0904
w/ Institution FE		0.0122	0.0633	0.0701	0.0785	0.0860	0.0922
w/ Major FE		0.0093	0.0667	0.0715	0.0784	0.0844	0.0887
w/ Base-year-by-major FE		0.0058	0.072	0.0749	0.0777	0.0809	0.084
<u>Panel C: 3-year changes</u>							
Program-years (n = 158,170)	0.0584	0.012	0.0437	0.0501	0.0598	0.0666	0.0720
w/ Base year FE		0.0101	0.0480	0.0543	0.0589	0.0637	0.0684
w/ Institution FE		0.0114	0.0444	0.0504	0.0597	0.0663	0.0713
w/ Major FE		0.0098	0.0461	0.0513	0.0597	0.0654	0.0693
w/ Base-year-by-major FE		0.005	0.0537	0.0560	0.0582	0.0608	0.0633
<u>Panel D: 2-year changes</u>							
Program-years (n = 190,059)	0.0383	0.0101	0.0258	0.0317	0.0399	0.0452	0.0496
w/ Base year FE		0.0081	0.0301	0.0353	0.0389	0.0424	0.0462
w/ Institution FE		0.0099	0.0259	0.0319	0.0401	0.0450	0.0493
w/ Major FE		0.0091	0.0274	0.0322	0.0398	0.0447	0.0480
w/ Base-year-by-major FE		0.0042	0.0347	0.0366	0.0382	0.0401	0.0421

Notes: The labor demand instrument is calculated at the program (institution-by-major) level for a given long-difference interval. Each panel presents the distributional statistics across multiple (stacked) long differences of a given length (e.g., 5 years). For example, Panel A includes all the intervals: 2010-2015, 2011-2016, and 2012-2017. The differences never go past 2017 as this is the most recent complete year for which job ad data are available. The first row presents statistics across programs. The subsequent rows display residuals plus the grand mean after controlling for the indicated fixed effects (FE). All estimates are weighted by the base year number of degrees awarded in each program.

Appendix Table A5. Responsiveness of Educational Investment to Changes in Skill Demand, Robustness

	Unweighted	Unweighted and drop small programs (i.e., < 10 degrees at baseline)	Alternate vectors of fixed effects		Instrument constructed with industry-major employment as share of aggregate major employment		
	2SLS (1)	2SLS (2)	2SLS (3)	2SLS (4)	2SLS (5)	2SLS (6)	2SLS (7)
Panel A. Outcome = Change in log(4-year degrees awarded), t0 to t0+7							
Change log(ads), t0 to t0+5	1.115*** (0.196)	1.444*** (0.224)	0.620*** (0.233)	1.076*** (0.352)	5.029*** (1.581)	3.249*** (1.128)	1.783*** (0.492)
F-stat from first stage	152.72	143.83	131.18	99.68	12.82	12.26	90.40
N(program-years)	92,501	67,257	92,501	92,501	92,501	92,175	92,138
N(institutions)	1,681	1,648	1,681	1,681	1,681	1,570	1,559
Panel B. Outcome = Change in log(total undergraduate credits), t0 to t0+7							
Change log(ads), t0 to t0+5	0.967* (0.541)	0.923* (0.544)	0.594 (0.519)	0.765 (0.731)	1.822 (2.274)	2.163 (2.303)	0.564 (0.618)
F-stat from first stage	8.33	10.54	20.74	12.44	2.65	3.5	28.55
N(program-years)	6,177	5,654	6,177	6,177	6,177	6,176	6,177
N(institutions)	114	113	114	114	114	114	114
Fixed effects or other controls	Major-by-base-year	Major-by-base-year	Major-by-base-year, State	Major-by-base-year, State-by-base-year	Major-by-base-year	Major-by-base-year, Δ Other Degrees	Major-by-base-year, School-by-base-year

Notes: The outcome is the change in the measure designated by each panel for a given program (institution-by-field cell) over a 7-year period for one of three long-difference intervals (i.e., 2010-2017, 2011-2018, or 2012-2019). Degree data come from IPEDS and data on undergraduate credit hours come from The Cost Study. See Table 1 and the text for details on construction of the analytic samples. The change in ads, the key independent variable, is based on an aggregation of job-ad data at the institution-major-year level, weighted by shares of an institution's graduates living and working in areas from which the job ads originate. This demand change is calculated over a 5-year period that uses the same base year as the corresponding outcome horizon (i.e., 2010-2015, 2011-2016, or 2012-2017). Please consult the text for additional information on this measure. Estimates in columns 3-7 are weighted by the number of degrees in the baseline year. Standard errors, clustered by institution, appear in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Appendix Table A6. Assessing Instrument Validity: Skill Demand Changes and Pre-Period Educational Investment

	Outcome = Change in log(4-year degrees awarded), t0-7 to t0		Outcome = Change in log(total undergraduate credits), t0-7 to t0	
	(1)	(2)	(3)	(4)
Change log(ads), t0 to t0+5	0.299 (0.200)	0.272 (0.195)	-0.179 (0.518)	-0.138 (0.528)
Fixed effects	Major-by-year	Major-by-year	Major-by-year	Major-by-year
Other controls	None	Contemporaneous 5-year shock	None	Contemporaneous 5-year shock
F-stat from first stage	100.26	101.55	19.65	18.43
N(program-years)	86,654	86,654	6,208	6,208
N(institutions)	1,604	1,604	131	131

Notes: This table presents stacked long-difference 2SLS results where the first stage is the same as that of our preferred equation from the main paper and the second stage outcome is a 7-year change in degrees or credits lagged by 7 years, such that it has no overlap to our preferred 7-year outcome difference. Columns 2 and 4 further add a control for a 5-year shock to demand corresponding to the placebo (lagged) 7-year outcome difference, as proxied by an instrument constructed like that in the first stage.

Appendix Table A7. Presence of Elite Firms in Job Ads

Firm Name	In sample?	Number of job ads (2017 calendar year)	Firm Name	In sample?	Number of job ads (2017 calendar year)
<i>A. Technology Firms</i>			<i>C. Banks</i>		
Apple	X	8,530	ABN AMRO	X	20
Amazon	X	28,670	Bank of America	X	21,800
Google	X	7,170	Brown Brothers Harriman	X	220
Facebook (Meta)	X	4,740	Citi	X	12,170
Microsoft	X	6,400	Cowen Group	X	<10
<i>B. Consulting Firms</i>			Deutsche Bank	X	890
A. T. Kearney			Evercore Partners	X	<10
Accenture	X	48,380	Gleacher & Company (closed in 2014)		
Advisory Board	X	1,120	Jefferies & Company	X	30
Analysis Group			Lazard	X	10
Arthur D. Little			Macquarie Group (Australia-based company)		
Bain & Company	X	260	Morgan Stanley	X	3,730
Booz & Company (acquired by PricewaterhouseCoopers (PwC) in 2014)			Perella Weinberg Partners	X	20
Booz Allen Hamilton	X	26,000	Piper Jaffray Companies	X	100
Cambridge Associates	X	50	Raymond James Financial	X	2,380
Charles River Associates	X	10	Robert W. Baird & Company	X	450
Cornerstone Research	X	10	Rothschild		
Corporate Executive Board	X	100	Thomas Weisel Partners Group		
Dean & Company			U.S. Bancorp	X	18,610
First Manhattan Consulting Group			Wells Fargo	X	18,080
FTI Consulting	X	250	William Blair & Company	X	170
Gallup	X	20			
Hewitt Associates (now Aon Hewitt)					
Huron Consulting Group	X	470			
Kurt Salmon (part of Accenture)					
Marakon (London-based company)					
McKinsey & Company	X	1,530			
Mercer	X	1,130			
Mitchell Madison Group					
Navigant	X	2,740			
NERA Economic Consulting	X	<10			
OC&C Strategy Consultants					
Oliver Wyman	X	10			
Parthenon Group					
PRTM (acquired by Pricewaterhouse Coopers (PwC) in 2011)					
Putnam Associates					
The Boston Consulting Group	X	<10			
ZS Associates	X	270			

Notes: The list of elite consulting firms and banks come from Table 1 in Weinstein (2022), with one edit: Wells Fargo acquired Wachovia in 2008, and thus we search for job ads from Wells Fargo. Job ad counts represent the number of times a given firm was listed on a job ad that appeared in the 2017 calendar year, rounded to the nearest 10.

Appendix Table A8. Component Majors of Broad-Major-Field Groups

Broad Group Name	Component Majors	Avg Cost Tercile of Major
Agriculture	Agriculture	2
	Natural Resources	2
Physical Sciences	Biology	1
	Chemistry	2
	Geological and Earth Sciences/Geosciences	2
	Physics	2
	Other Physical Sciences	2
	Biochemistry, Biophysics and Molecular Biology	3
	Microbiology	3
	Atmospheric Sciences and Meteorology	3
Communications	Materials Science and Engineering	3
	Journalism	1
Engineering	Public Relations, Advertising, and Applied Communication	1
	Communication and Media Studies	1
	Computer and Information Science	2
Engineering	Engineering Technology	2
	Aeronautical Engineering	3
	Biomedical Engineering	3
	Chemical Engineering	3
	Civil Engineering	3
	Computer Engineering	3
	Electrical, Electronics and Communications Engineering	3
	Mechanical Engineering	3
	Systems, Industrial, Manufacturing, and Operations Engineering	3
	Other Engineering	3
Education	Teacher Education	2
	Special Education and Teaching	2
	Other Education	2
Humanities	Foreign Language and Linguistics	1
	Family and Consumer Sciences	1
	English, Liberal Arts, Humanities	1
	Philosophy and Religion	1
	Theology	2
	Legal Studies	3
	Library Science	3
Social Sciences	Statistics	1
	Mathematics	1
	Psychology	1
	Protective Services	1
	Economics	1
	Geography	1
	Political Science, Government, and International Relations	1
	Sociology	1
	Other Social Sciences	1
	Social Work	2
	Public Administration	3
	Public Policy	3
	Arts	Design, Photography, Video, and Applied Arts
Other Visual/Performing Arts		2
Architecture		2
Culinary Arts		2
Health	Fitness, Recreation and Leisure Studies	1
	Dietetics and Clinical Nutrition Services	2
	Allied Health	2
	Health and Medical Administrative Services	3
	Pharmacy, Pharmaceutical Sciences, and Administration	3
	Public Health	3
	Rehabilitation and Therapeutic Professions	3
	Registered Nursing, Nursing Administration, Nursing Research and Clinical Nursing	3
Business	Accounting and Related Services	1
	Marketing	1
	Business, general	1
	Finance and Financial Management Services	2
	Hospitality Administration/Management	2
	Human Resources Management and Services	2
	Management Information Systems and Science	2

Notes: Each of the 66 majors used in the main analyses are categorized into one of the 10 broad fields listed above. The broad field groupings represent those used for field-specific estimates in Table 5 and Figure 5. Each major is grouped into a tercile of average costs based on the process described in Table 5 (1 = lowest average instructional costs per credit hour; 3 = highest average instructional costs per credit hour).

Appendix B. Extension to Conceptual Framework

From our motivating framework in Section III.A, the log number of degrees produced by cohort c in major m at institution s can be expressed as follows:

$$\log(Y_{cms}) = \log(N_{cs}) + \eta_{ms} + \mu_{cm} + \beta\gamma_{cms} - \log(\pi_{cs})$$

When estimating β , one concern is that $\beta\gamma_{cms}$ appears both directly and in π_{cs} . Estimating the above equation and its long-differenced version with institution-year fixed effects to capture π_{cs} will partially control for the effect of interest, while additionally altering the identifying variation to within institution compositional changes. In this appendix, we show that our model implies that other degrees outside of m produced at institution s serve as a potential alternative control to appropriately estimate β .

First note that the total number of degrees produced by cohort c at institution s is:

$$D_{cs} = N_{cs} \left(\frac{\sum_k \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})}{1 + \sum_k \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})} \right)$$

and similarly, the total number of degrees other than in major m produced at institution s is:

$$D_{cs, \neq m} = N_{cs} \left(\frac{\sum_{k \neq m} \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})}{1 + \sum_k \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})} \right)$$

Now, we reconsider our expression for the negative of $\log(\pi_{cs})$:

$$\begin{aligned} -\log(\pi_{cs}) &= \log \left(\frac{1}{1 + \sum_k \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})} \right) \\ &= \log \left(\frac{1}{1 + \sum_k \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})} \times \frac{\sum_{k \neq m} \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})}{\sum_{k \neq m} \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})} \right) \\ &= \log \left(\frac{D_{cs, \neq m}}{N_{cs}} \right) - \log \left(\sum_{k \neq m} \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks}) \right) \end{aligned}$$

$$= \log(D_{cs, \neq m}) - \log(N_{cs}) - \log\left(\sum_{k \neq m} \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks})\right)$$

$$\text{Let } \pi_{cs, \neq m} = \sum_{k \neq m} \exp(\eta_{ks} + \mu_{ck} + \beta\gamma_{cks}).$$

We can now plug the above expression into our main long-difference specification:

$$\Delta \log(Y_{ms}) = \beta \Delta \gamma_{ms} + \Delta \mu_m + \Delta \log(N_s) + \Delta \log(D_{s, \neq m}) - \Delta \log(N_s) - \Delta \log(\pi_{s, \neq m})$$

$$\Delta \log(Y_{ms}) = \beta \Delta \gamma_{ms} + \Delta \mu_m + \Delta \log(D_{s, \neq m}) - \Delta \log(\pi_{s, \neq m})$$

Based on this result, in some specifications, we control for the change in the log of other degrees (or credits), $\Delta \log(D_{s, \neq m})$, but we cannot directly control for $\Delta \log(\pi_{s, \neq m})$. However, note that $\pi_{s, \neq m}$ is in the numerator of $D_{s, \neq m}$, and so it is perhaps reasonable to assume that $\Delta \gamma_{ms}$ is uncorrelated with $\Delta \log(\pi_{s, \neq m})$ after already conditioning on $\Delta \log(D_{s, \neq m})$ and $\Delta \mu_m$.

Appendix C. Indirect Demand for Majors via Industry Employment

In this appendix we describe two candidate approaches to mapping industries to majors. Consider a demand shock for industry i and area j between cohorts 0 and t that increases labor demand from γ_{0ji} to γ_{tji} . The effective change in demand for major, m , then depends on the importance of γ_{cji} for γ_{cjm} . We define the relationship as $\gamma_{cjm} = \sum_i \phi_{im} \gamma_{cji}$ where ϕ_{im} is the relevance of major m in industry i .

We operationalize ϕ_{im} (called Z_{im} in the main text) as the total number of workers who majored in m and work in industry i divided by the total number of workers in industry i . Therefore, in our empirical work where γ reflects job ads and employment, labor demand for major m increases proportional to the size of the increase in the *number* of ads, a likely salient measure of labor demand for students. For example, if an industry's (college graduate) employment is 50 percent mechanical engineering majors, after an increase of 1,000 jobs for that industry we would expect the number of jobs for mechanical engineering majors to increase by 500. Our preferred conceptualization creates a direct link between the number of job ads in major-relevant industries and labor demand for a given major.

An alternative construction would operationalize ϕ_{im} as the total number of workers who majored in m and work in industry i divided by the total number of workers in major m . This conceptualization instead weights labor demand in industry i by its relevance among only workers in major m , regardless of m 's size overall or its size within the industry. Therefore, if an industry's employment increases by 1,000 jobs, then this alternative measure does not directly translate into the expected increase in employment or labor demand for m . An important implication here is that, compared to our preferred construction, the first-stage of our 2SLS specification under this alternative construction is considerably weaker. Since this alternative construction does not as cleanly link new job ads to labor demand, it is not our preferred measure. We do, however, present estimates using this alternative construction as a robustness exercise, noting caution given the low power for the first stage.