Paying for Prestige: Compensating Differentials in Faculty Job Choice

Alex Bell and Todd R. Jones*

March 19, 2025

Abstract

We investigate how academics trade off salary and job amenities, particularly the prestige of their departments, using a novel anti-instrumental variable approach (Bell 2022). While compensating differential theory predicts such trade-offs, standard regression techniques often fail to account for unobserved variation in faculty offer sets, leading to biased estimates. Our approach addresses this challenge by leveraging the relationship between a faculty member's PhD-granting university quality, current institution salary, and current university prestige to better identify the salary-prestige tradeoff. Using de-identified data on over 180,000 faculty-field pairs in 325 U.S. universities across 107 fields, we find strong evidence that faculty are willing to forgo substantial earnings for positions at more prestigious institutions. The average faculty member must give up $2,000 in pay to move to the next-most-prestigious school. These findings contribute to a broader understanding of labor market dynamics in high-skilled professions, where non-monetary factors play a critical role in job selection. Our methodological approach offers new insights into the pricing of job amenities and how institutions compete for talent beyond academia.

*Bell: Georgia State University. *ambell@gsu.edu*. Jones: Mississippi State University, CESifo, and IZA. *trj234@msstate.edu*.

1 Introduction

Do faculty "pay" for prestige? Economists dating back to Smith (1776) have long theorized that workers forego earnings for other desirable facets of a job, but empirical evidence in favor of that view scarce. This paper applies new methods in labor economics to granular data on U.S. faculty rosters and backgrounds to document and explore the trade-off prospective faculty face when choosing to work at more prestigious schools versus schools that pay better.

This paper contributes to our understanding of incentives in labor markets for high-skilled workers. As evidenced by Stern (2004), scientists and researchers often exhibit non-financial motivations for their work beyond pay, but the extent of this preference variation was not previously documented across the broader academic labor market. This paper builds on a broad literature pricing job amenities in various contexts (Ashenfelter and Greenstone, 2004; Greenstone et al., 2015; Lavetti, 2015; Cook et al., 2018). This paper also complements surveys indicating that workers have preference over job attributes by showing a salient case in which worker preferences can actually translate to labor market factor prices (Mas and Pallais, 2017; Maestas et al., 2018; Wiswall and Zafar, 2018; He et al., 2019).

While evidence is growing that non-wage attributes are important facets of the labor market, little is known about the role of job amenities in high-skilled markets. On one hand, by focusing solely on the academic labor market, we are already considering a very specialized type of job. Relative to other jobs, US faculty jobs offer autonomy, flexible schedules, and the potential for long-term job security through tenure. Faculty also benefit from intellectual freedom, access to research funding, campus amenities, and perks like tuition benefits and sabbaticals. However, faculty salaries often lag behind industry, and the workload—especially before tenure—can be intense.

As unique as the academic career path may be, variation in job types still exists within the academic job market. For instance, the investigation of faculty hiring practices by

Goolsbee and Syverson (2023) concluded that faculty view "jobs at prestigious schools as meaningfully differentiated from alternative positions." Other ways to classify or differentiate academic jobs exist; research universities tend to offer more funding and lighter teaching loads, whereas liberal arts colleges emphasize student engagement. Geographic locations of universities, such as rural versus urban, also shape how prospective faculty evaluate the utility conferred by a job offer. Although we find some evidence that these other amenities matter for faculty, our analysis suggests that the salary versus prestige trade-off is surprisingly distinct and robust to controlling for other amenity measures.

Our analysis reveals that across US colleges and universities, the average price for prestige is about $2,000. More precisely, the average professor wishing to move to a marginally one-rank-better school can expect a pay cut of $2,000 to be associated with that move.

To arrive at these estimates, we first need a measure of university prestige. Our preferred specification uses SpringRank, calculated by Wapman et al. (2022), which operationalizes the notion that the most prestigious schools are those that hire from the most prestigious PhD programs. Appendix Figure A.1 plots distributions of this measure of prestige. Faculty in the sample tend to work at universities with higher SpringRank scores (Panel A). Strikingly, faculty disproportionately received their PhDs at highly-ranked universities (Panel B).

Still, even with a state-of-the-art measure of school prestige, traditional economic analysis fails to accord with intuition. Panel C of Figure 1 shows the raw relationship between average salary at a school and the job's prestige (as measured by SpringRank). Following the theory of compensating differentials, if faculty trade off between pay and prestige, a downward-sloping relationship should emerge here. Instead, the relationship is strongly upward-sloping. One hypothesis for the wrong-signed result is that some jobs are just better than others—in terms of both pay and prestige—but we lack an adequate measure of worker ability for which to control. In Panel D, we test the "ability bias" hypothesis by

including the prestige measure of the PhD granting institution for each faculty member. The resulting relationship is still upward-sloping, but slightly less so, which is evidence consistent with ability being a confounding factor in the relationship in Panel C, but we do not observe the full set of ability-relevant controls.

The approach we apply in this paper is to treat the quality of each faculty member's PhD not as a control but instead as an anti-instrument for their job prestige (Bell, 2022). As with other recent literature, the anti-IV methodology adequately controls for the latent index, and reveals the expected downward-sloping market tradeoff between pay and prestige (Folke and Rickne, 2022; Burbano et al., 2023; Ferreira et al., 2024; De Schouwer and Kesternich, 2024; Bell et al., 2024).

The details of that empricial strategy are the subjection of the next section. Subsequently, we discuss our data on faculty and their PhD-granting institutions, interpret our results, and conclude with a discussion of open questions in this literature.

2 Empirical Strategy

2.1 The Anti-IV Approach

The key confounding factor when studying job choice among academics is that the market does not operate such that a given worker has access to all jobs. Instead, following Bell (2022), we assume the market is segmented vertically such that some workers have access to better jobs than other workers, both in terms of pay and prestige.

The ideal experiment in this case would be to observe two workers who could have chosen among the same sets of jobs, but differing preferences for pay versus prestige led to different choices. If we could observe and condition on the sets of jobs available to all workers in academia, this variation in preferences would trace out the trade-off workers face when choosing between jobs with greater pay versus prestige. While surveys of job applicants like that of Stern (2004) among biology post-docs are helpful for studying the

behavior of small groups of individuals, we seek to study the broader academic labor market in the US.

Following Bell (2022), we use observed qualifications of workers as anti-instruments for the quality of their offer sets. The intuition for the anti-instrumental approach is similar to the typical approach of controlling for qualifications in hedonic regressions, but it differs in that the assumptions are weaker: Instead of needing to observe all productivity-relevant attributes of the worker, we only need to observe one. In our preferred specification, we take the ranking of the program from which the worker received their PhD as the anti-instrument for the quality of their offer set.

Implementing this approach requires 1) regressing the anti-instrumental variable (PhD quality) on the dependent variable (average salary) and amenity (job prestige); 2) creating predicted values; and then 3) regressing the dependent variable on the amenity, while controlling for the predicted values. The coefficient on the amenity is the coefficient of interest.

This approach estimates the wage-amenity frontier commonly attributed to Rosen (1986). The intuitive interpretation of the frontier is the price of the amenity that arises from a perfectly competitive frictionless sorting equilibrium between workers and firms. It is an equilibrium price in the sense that if an additional worker were added to the market, the price dictates the schedule of wage-amenity combinations that the worker would be able to choose from. Our approach differs from the literal textbook model of compensating differentials only in that we allow for a single, partially latent index of vertical job quality (whereas the textbook model assumes no such vertical differentiation of jobs or workers).

While the anti-IV approach removes the need for the researcher to observe and condition on the full set of productivity-relevant attributes, a remaining limitation (as with the traditional hedonic approach) is that the researcher still needs to observe the set of job attributes that generate utility for the worker. At present, we primarily model the worker's

decision problem as trading off only between pay and prestige because we feel this is an interesting and under-explored dimension of job choice. Still, future work should explore robustness to incorporating additional workplace amenities into the job choice model.

2.2 Approach to Aggregation

Although the quality of one's PhD granting institution is measured at the individual-level, other key variables are observed only at more aggregated levels. The primary job attribute of interest, prestige, is measured at the level of the school-by-field. We also measure pay at the school level from IPEDS. We therefore conduct the analysis at the individual level, and cluster our standard errors at the school-by-field level. In additional analysis that estimates the price of prestige within a field, we cluster our standard errors at the school level. We also explore robustness of our estimates to estimating the price of prestige on datasets that are collapsed to the school-by-field (or school) level.

2.3 Ranking Schools

Following a growing wave of recent applied work in labor economics, we use the network of worker flows across universities as a revealed-preference indicator for university prestige. The specific algorithm used to rank universities is called SpringRank.[1] While the exact details of the implementation differ slightly from PageRank (Sorkin, 2018) or "poaching rank" (Bagger and Lentz, 2019), it is also recursive in nature in that departments are ranked higher if their graduates place at higher-ranked universities. Moreover, recent research in network sciences has found SpringRank to be more robust than alternatives such as PageRank in a variety of real-world and simulated datasets, particularly when the underlying data imply a natural hierarchy (De Bacco et al., 2018).

[1]The algorithm is named after an analogy to a physical system in which universities (i.e., the nodes of the graph) are connected by springs. The resting lengths of these springs (i.e., the edges of the graph) are determined by hiring patterns across the universities. The algorithm then finds the ordering of universities that minimizes the resting energy of the system.

3 Data

3.1 Faculty Data

Our baseline data on a faculty's current university (which we refer to as job), the university at which they received their PhD, and SpringRank come from Wapman et al. (2022) (hereafter, WZCL) as part of their publicly-available replication package.[2] WZCL obtained their original data from Academic Analytics Research Center. These original data, for each year between 2011 and 2020, included a census of the tenure or tenure-track faculty at the 392 US universities that grant PhD degrees. The data include the institution at which the faculty received their PhD (or other terminal degree). WZCL conducted an extensive data cleaning process, which we summarize in the Data Appendix (and which is outlined in more detail in the Methods section of their paper). Here we note two things. First, that our unit of observation is the individual-field. This is because WZCL allowed a person to appear in multiple fields; for instance, someone in a department of economics and finance would appear in both fields. And second, that a field-job contains all relevant departments. For instance, if a university contained both a department of economics and a department of business economics, then faculty from each would be grouped together.

3.2 Job and PhD Ranking Metrics

For our primary ranking metric, we use SpringRank, as computed by WZCL. The data contain two versions: the raw SpringRank variable, which we don't use, and an ordinal rank of SpringRank, which we do. We reverse the ranking of ordinal SpringRank such that higher numbers are better and call it 'SpringRank, 1–N,' where N is the highest number observed within a field, noting that the N varies widely across fields, as depicted in Panel A of Appendix Figure A.3. The ordinal rank skips numbers when the

[2] We downloaded data from `https://github.com/LarremoreLab/us-faculty-hiring-networks`. Accessed January 17, 2025.

skipped-schools do not appear in the data. Thus, N is almost always higher than the number of schools we observe (Panel B). We also construct 'SpringRank, 1–100' by scaling 'SpringRank, 1–N' by N.[3] This has the advantage of placing all fields on a common scale, which is especially useful for graphing. After doing this, we match the job (university) and PhD for each faculty member to SpringRank. We refer to a job's SpringRank as 'job prestige' and a PhD's SpringRank as 'PhD quality.'

We use two additional measures of prestige: production and undergraduate acceptance. Production is also computed by WZCL and is an ordinal ranking of how many PhDs a given university-field produced that are among the faculty-field pairs in the sample. This variable also skips universities. In addition, we use the undergraduate acceptance rate at the university.[4] We convert this to an ordinal rank and do so within field, only among the universities that appear in the final sample.[5]

In Appendix Tables A.1, A.2, and A.3, we present examples for three disciplines—economics, history, and computer science—of the top 30 ranked SpringRank universities. For comparison, we include the production rank, the USNWR rank (which we currently have for only these three disciplines), and the accept rate. For ease of interpratability, in this table, we present the rankings such that 1 is the top-ranked school and we present the raw undergraduate acceptance rate.

3.3 Salary Data

We use salary data from IPEDS. A main weakness of our paper is that because we do not have individual-level faculty salary data, we rely on average salary data at the university level (i.e., the data are not reported separately by field). We use the average

[3]This variable actually ranges from 0 to 100.

[4]We access the variables in IPEDS used to construct this variable, as well as various other variables such as salary data, using the educationdata Stata package (Urban Institute, 2024).

[5]Thus, the N for this variable is the same as the number of universities (jobs) or PhD universities in the field. This is in contrast to N for SpringRank and production, which is almost always larger than the number of schools observed after sample selections.

salary for assistant professors for the 2020-2021 school year.

3.4 Sample Selections and Summary Statistics

We exclude observations that are missing data on salary, school rank, and/or PhD rank. We are unable to match a small number of schools to the average salary data in IPEDS. The way the SpringRank algorithm runs necessitates that some schools (within a field) do not receive a ranking. As stated above, we rely on the same SpringRank for a given school, whether it serves as the faculty member's current job or school at which they received their PhD. This means that some individuals will be missing either their job prestige and/or PhD quality. If we consider the field-level data and expand to individual, then 1.4% of individuals are missing average salary, 15.9% are missing school prestige, and 31.4% are missing PhD quality. (These numbers are smaller at the domain and academia levels.) We begin with 317,009 observations. After sequentially dropping those missing salary data, school ranking, and PhD ranking, our sample falls to 312,426, 262,725, and 181,448, respectively.

4 Results

In Table 1, we compare naive regressions to the anti-IV approach. We begin by regressing income on job quality in Column 1 of Table 1. Contrary to what theory would predict, we find a positive coefficient, which suggests that faculty accept $221 to go to a one higher-ranked job. In Column 2, we control for PhD quality. Doing so decreases the coefficient on job rank slightly, but it remains positive. We then implement the anti-IV approach, using PhD rank as the anti-instrument.[6] In Column 3, we find a trade-off of $2,137 to move to a school ranked one rank higher.

To help visualize the result, we proceed to conduct a nonparametric analysis in Fig-

[6]To do so, we use the `aivreg` Stata package, available at https://github.com/ZhongShusheng/aivreg$_stata_package$.

ure 2. Panel A shows the non-parametric analogue of our preferred specification. Within quantile bins of salary and prestige, we calculated the mean PhD quality, the levels of which are depicted by different colors. The result is a downward-sloping relationship between salary and school prestige, revealing that the market imposes a tradeoff between these two forms of compensation. (Panel B also shows this non-parametric relationship for bins of mean undergraduate acceptance rank at one's PhD.)

We test robustness in several different ways. First, in Appendix Table A.4, we collapse the data to the field-school level. We find similar results (noting that the drop-off in the school rank coefficient between Columns 1 and 2 is larger than in Table 1).

Second, in Table 2, we add several different amenities. In column 2, we these are mean January temperature, public school ranking, urbanicity (all corresponding to the county of the university, borrowing from the geographic amenities collected by Bell et al. (2024)), and the Carnegie ranking of the university. We find a large coefficient on urbanicity (which ranges between 4 and 12, with larger numbers being more urban), indicating a strong preference to live in cities. We find a positive coefficient for January temperature (where higher is warmer), which is contrary to our expectation. In column 3, we add county PageRank, a proxy for how desirable a county is. Importantly, the coefficient on job prestige is robust to the addition of these additional control variables.

Third, in Table 3 we probe robustness of our preferred specification to different choices of anti-instruments as an intuitive test of over-identification. Reassuringly, all anti-instruments that we have investigated yield a negative sign, indicating (in contrast to hedonic approaches) that a trade-off exists between pay and prestige. The exact magnitude changes somewhat between -$1,000 to -$2,100 in the first three columns, which use different parameterizations of PhD quality as measured by SpringRank or the number of PhD's produced in the PhD-granting institution.

While the signs of all estimates align with theoretical expectations, the trade-off appears notably smaller when using the undergraduate acceptance rate of the PhD-granting

institution as the anti-instrument. Panel B of Figure 2 shows the nonparametric version, with a noticeably flatter slope. We suspect this occurs because undergraduate acceptance rates are not perfect anti-instruments for PhD program quality, but may instead be contaminated by preferences for pay vs prestige themselves. Notably, as shown in Tables A.1–A.3, schools with higher acceptance rates (relative to other prestige measures) tend to be state schools. Since state schools typically charge lower tuition, they may attract students who prioritize financial considerations over prestige. Following Bell (2022), the bias in our second-stage estimate corresponds to the coefficient in a hypothetical regression of salary on the undergraduate acceptance rate of one's PhD institution, controlling for true PhD program quality. Given that state schools tend to have both higher acceptance rates and graduates more focused on salary over prestige, this coefficient is likely positive. Thus, it is unsurprising that this specification yields a more positive estimate. For this reason, we mainly stick to our preferred specification using SpringRank as the anti-instrument, but still include undergraduate acceptance rates as a reasonable alternative one might a priori consider as relevant.

Fourth, Table 4 evaluates robustness of our estimates to how we measure the prestige amenity. We consider prestige as measured by SpringRank, ProductionRank, or undergraduate acceptance rates. There is some variation in the estimated tradeoff depending on which prestige measure we use, although all prestige measures consistently reveal trade-offs associated with them when priced using the anti-IV framework.

Fifth, we test robustness to using an alternate measure of job quality and PhD ranking: US News & World Report (USNWR). In Table 5, we consider the fields of economics, history, and computer science, the three fields for which we have USNWR data.[7] To increase comparability, we include only individuals for whom we observe both SpringRank and USNWR. We also re-rank each variable such that the number of universities (N) for a given ranking (SpringRank, 1–N and USNWR, 1–N) is consistent for job prestige

[7]Clauset et al. (2015) collected the USNWR data for history and computer science, and Jones and Sloan (2024) collected the USNWR data for economics.

and is constent for PhD quality. We then vary the measure of job prestige and the anti-instrument, finding that the results are similar for each of the four combinations.

Next, we separately consider the eight individual domains, Applied Sciences; Education; Engineering; Humanities; Mathematics and Computing; Medicine and Health; Natural Sciences; and Social Sciences. A mapping between domains and fields is found in The Data Appendix. We first subset to a given domain, and then conduct the anti-IV approach, using field fixed effects and clustering at the field-school level. In Panel A Figure 3, we present the results, using job quality on the 1–N scale. Based on the size of the confidence intervals, many of the fields are indistinguishable from each other, the exception being that medicine and health does overlap with several of the fields. We also present in Panel B the results when using job quality on the 1–100 scale. There is more variation, though the interpretation of the coefficients is harder to compare across fields based on having differing numbers of jobs in the fields in each domain.

We repeat the above exercise in Figure 4, but instead subset to each individual field. The sample are relatively small at this point, contributing to, in some cases, small partial F-statistics. As such, we present results only for those fields with a value of at least 10.[8] For the most part, in both panels, the fields are not statistically different from each other, limiting the conclusions we can make about individual fields.

5 Conclusion

This paper provides new evidence on the trade-off between salary and prestige in faculty job choices, using a novel anti-IV approach to correct for unobserved heterogeneity in offer sets. We find that faculty are willing to forgo on the order of $2,000 of earnings for the opportunity to work at a more prestigious institution. Our results challenge the

[8]We also truncate the lower bounds at -7,000. We present a version of these figures in Appendix Figure A.4 that does not impose the partial F-statistic of 10 restriction, instead showing all fields for which a finite confidence interval was estimated; for ease of interpretability, we truncate the lower bounds for these plots at -10,000.

assumption that salary is the dominant factor in academic job decisions and highlight the role of non-monetary job amenities in shaping career trajectories.

By leveraging anti-instruments, our methodology overcomes key limitations of standard hedonic regression approaches, offering a more robust estimate of compensating differentials in academia. This has broader implications for understanding labor market sorting in high-skilled professions, where prestige, mission alignment, or job security may be as valuable as pay.

Despite the strengths of our approach, some limitations remain. While our sample is very large, it covers only faculty members at U.S. PhD-granting institutions; we are further limited in that we can only consider a portion of these faculty based on needing to observe data on both a faculty member's job prestige and PhD quality. Our reliance on university-level salary data constrains our ability to measure individual wage variation (and matching offers, etc.) , and our model primarily focuses on pay and prestige, omitting other key job attributes. We prioritize prestige in our framework because it is an especially salient factor in academic job choice, often considered a central motivation for faculty decisions, and because its measurement benefits from greater consistency across institutions. Future work could extend this framework by incorporating more detailed salary data, exploring alternative measures of prestige, and applying the anti-IV approach to other professional labor markets.

Overall, our findings contribute to a broader understanding of how highly skilled labor markets function, particularly in academia, where non-monetary job attributes play a significant role in career decisions. Researchers often prioritize factors such as research autonomy and mission alignment, willingly accepting lower salaries in exchange for positions that offer greater academic freedom and prestige. By refining estimates of the pay-prestige tradeoff, this paper provides new evidence for how preferences for institutional prestige and career incentives shape labor market outcomes across different high-skilled professions. More directly, our estimates will be useful for universities wishing to tailor

their recruitment and retention incentives by understanding the prices the academic labor market imposes on prestige.

References

Ashenfelter, O. and M. Greenstone (2004). Using Mandated Speed Limits to Measure the Value of a Statistical Life. *Journal of Political Economy 112*(S1).

Bagger, J. and R. Lentz (2019, January). An Empirical Model of Wage Dispersion with Sorting. *The Review of Economic Studies 86*(1), 153–190.

Bell, A. (2022). Job amenities and earnings inequality. *Available at SSRN 4173522*.

Bell, A., S. B. Billings, S. Calder-Wang, and S. Zhong (2024, July). An Anti-IV Approach for Pricing Residential Amenities: Applications to Flood Risk *.

Burbano, V. C., O. Folke, S. Meier, and J. Rickne (2023, December). The Gender Gap in Meaningful Work. *Management Science*. Publisher: INFORMS.

Clauset, A., S. Arbesman, and D. B. Larremore (2015). Systematic inequality and hierarchy in faculty hiring networks. *Science advances 1*(1), e1400005.

Cook, C., R. Diamond, J. Hall, J. A. List, and P. Oyer (2018, June). The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers. Working Paper 24732, National Bureau of Economic Research.

De Bacco, C., D. B. Larremore, and C. Moore (2018, July). A physical model for efficient ranking in networks. *Science Advances 4*(7), eaar8260. Publisher: American Association for the Advancement of Science.

De Schouwer, T. and I. Kesternich (2024). Work Meaning and the Flexibility Puzzle.

Ferreira, S., M. Moro, and H. Welsch (2024). Using Life Satisfaction and Happiness Data for Environmental Valuation: An Experienced Preference Approach. *SSRN Electronic Journal*.

Folke, O. and J. Rickne (2022, November). Sexual Harassment and Gender Inequality in the Labor Market*. *The Quarterly Journal of Economics 137*(4), 2163–2212.

Goolsbee, A. and C. Syverson (2023, October). Monopsony Power in Higher Education: A Tale of Two Tracks. *Journal of Labor Economics 41*(S1), S257–S290. Publisher: The University of Chicago Press.

Greenstone, M., S. P. Ryan, and M. Yankovich (2015). The Value of Statistical Life: Evidence from Military Retention Incentives and Occupation-Specific Mortality.

He, H., D. Neumark, and Q. Weng (2019, January). Do Workers Value Flexible Jobs? A Field Experiment on Compensating Differentials. Working Paper 25423, National Bureau of Economic Research.

Jones, T. R. and A. A. Sloan (2024). The academic origins of economics faculty. *The Journal of Economic Education 55*(4), 434–454.

Lavetti, K. (2015, February). Estimating Preferences in Hedonic Wage Models: Lessons from the Deadliest Catch. SSRN Scholarly Paper ID 1697243, Social Science Research Network, Rochester, NY.

Maestas, N., K. J. Mullen, D. Powell, T. von Wachter, and J. B. Wenger (2018, October). The Value of Working Conditions in the United States and Implications for the Structure of Wages. Working Paper 25204, National Bureau of Economic Research.

Mas, A. and A. Pallais (2017, December). Valuing Alternative Work Arrangements. *American Economic Review 107*(12), 3722–3759.

Rosen, S. (1986). Chapter 12 The theory of equalizing differences. In *Handbook of Labor Economics*, Volume 1, pp. 641–692. Elsevier.

Smith, A. (1776). *The Wealth of Nations*.

Sorkin, I. (2018, August). Ranking Firms Using Revealed Preference. *The Quarterly Journal of Economics 133*(3), 1331–1393.

Stern, S. (2004, June). Do Scientists Pay to Be Scientists? *Management Science 50*(6), 835–853.

Urban Institute (2024). IPEDS, Education Data Portal (Version 0.22.0). Made available under the ODC Attribution License.

Wapman, K. H., S. Zhang, A. Clauset, and D. B. Larremore (2022). Quantifying hierarchy and dynamics in us faculty hiring and retention. *Nature 610*(7930), 120–127.

Wiswall, M. and B. Zafar (2018, February). Preference for the Workplace, Investment in Human Capital, and Gender. *The Quarterly Journal of Economics 133*(1), 457–507.

Table 1: Comparison of Naive and AIV Approaches

	(1)	(2)	(3)
Job Prestige	220.56***	208.81***	-2137.24***
	(5.26)	(5.28)	(125.87)
PhD Quality		36.89***	
		(2.04)	
Constant	73859.75***	70746.50***	
	(432.28)	(454.22)	
Observations	181448	181448	181448
AIV	No	No	PhD Quality

Notes: Column 1 shows a regression of average salary on job prestige (SpringRank, 1–N). Column 2 adds PhD quality (SpringRank, 1–N). Column 3 regresses average salary on job prestige (SpringRank, 1–N), with PhD quality (SpringRank, 1–N) as the anti-instrument. All regressions include field fixed effects. We cluster at the school-field level for all columns and use Anderson-Rubin standard errors for column 3. Standard errors in parentheses. * 0.10; ** 0.05; *** 0.01.

Table 2: Adding Additional Amenities

	(1)	(2)	(3)
Job Prestige	-2137.24***	-1835.11***	-2232.60***
	(125.87)	(117.69)	(139.32)
Mean January Temperature		2807.11***	
		(785.70)	
Public School Ranking		-1016.62	
		(781.89)	
Urbanicity		-2841.22***	
		(462.15)	
Carnegie		-2033.85	
		(2349.26)	
County PageRank			-726.13***
			(232.31)
Observations	181448	156239	181448

Notes: Each column includes a regression of average salary on job prestige (SpringRank, 1–N), with PhD quality (SpringRank, 1–N) as the anti-IV. Column 2 also includes mean January temperature, public school ranking, urbanicity (all at the county-level), and Carnegie classification for the job (university). Column 3 instead adds the count PageRank. All regressions include field fixed effects. We cluster at the school-field level and use Anderson-Rubin standard errors. Standard errors in parentheses. * 0.10; ** 0.05; *** 0.01.

Table 3: Comparison of Different Anti-IV Variables

	(1)	(2)	(3)	(4)
Job Prestige, 1-N	-2137.24***	-1020.40***	-1668.54***	-482.27***
	(125.87)	(61.74)	(112.09)	(25.25)
Observations	181448	181448	181448	181448
AIV	PhD Quality, 1-N	PhD Quality, 1-100	Production, 1-N	UG Acceptance at PhD
Mean of AIV	111.0	78.4	25.8	69.2
SD of AIV	61.8	22.0	29.9	44.9

Notes: This table shows a regression of average salary on job prestige (SpringRank, 1–N). We vary the anti-IV we use across columns: we use PhD quality (SpringRank, 1–N) in column 1, PhD quality (SpringRank, 1–100) in column 2, PhD production (ranking of the number of PhDs one's PhD school produced that are in the sample) in column 3, and the undergraduate acceptance rate at one's PhD school in column 4. The means and standard deviations of the anti-IVs are presented at the bottom of the the tables. All regressions include field fixed effects. We cluster at the school-field level and use Anderson-Rubin standard errors. Standard errors in parentheses. * 0.10; ** 0.05; *** 0.01.

Table 4: Comparison of Different Measures of Prestige

	(1)	(2)	(3)	(4)
Job Prestige, 1-N (Preferred)	-2137.24***			
	(125.87)			
Job Prestige, 1-100		-2244.65***		
		(126.57)		
Job PhD Production, 1-N			-353.95***	
			(20.41)	
UG Acceptance Rate at Job, 1-N				-190.08***
				(28.22)
Observations	181448	181448	181448	181448
Mean of Job Prestige Var.	83.5	58.8	120.4	55.9
SD of Job Prestige Var.	56.9	26.9	53.2	41.0

Notes: This table shows a regrssion of average salary on job prestige, with PhD quality (SpringRank, 1–N) as the anti-instrument. It varies the definition of job prestige, with our preferred measure (SpringRank, 1–N) in column A, SpringRank, 1–100 in column B, the PhD production (ranking of the number of PhDs one's PhD school produced that are in the sample) in column 3, and the undergraduate acceptance rate at the job in column 4. All regressions include field fixed effects. The mean and standard deviation of the relevant measure of job prestige are presented at the bottom of the table. We cluster at the school-field level for all columns and use Anderson-Rubin standard errors for column 3. Standard errors in parentheses. * 0.10; ** 0.05; *** 0.01.

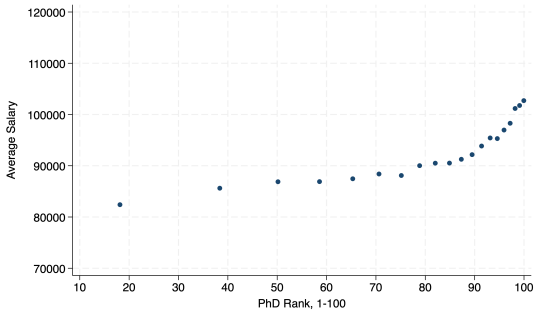
Table 5: Comparison of SpringRank and USNWR Rank for Three Fields

	(1)	(2)	(3)	(4)
Job Prestige, 1–N	-1932.30***	-1683.87***		
	(404.24)	(365.05)		
Job Prestige, USNWR, 1–N			-1690.41***	-2002.21***
			(370.08)	(452.07)
Observations	13779	13779	13779	13779
AIV	PhD Quality, 1-N	USNWR	PhD Quality, 1-N	USNWR

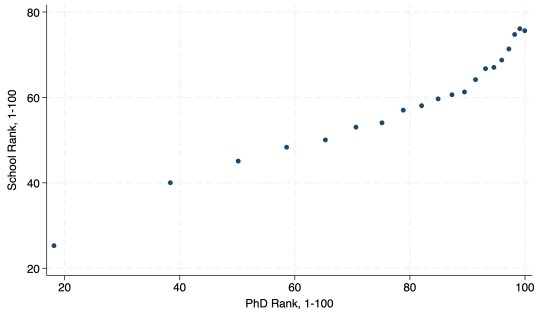
Notes: The dependent variable in each column is average salary. The measure of job prestige is our preferred measure (SpringRank, 1–N) in columns 1 and 2, and USNWR (1–N) in columns 3 and 4. We vary the anti-IVs, with PhD quality defined as SpringRank, 1–N in columns 1 and 3, and as USNWR, 1–N in columns 2 and 4. This table considers the fields of economics, history, and computer science. To increase comparability across datasets, we use only observations that have non-missing observations for each of the variables mentioned above. We also re-rank each variable. All regressions include field fixed effects. We cluster at the school-field level and use Anderson-Rubin standard errors. Standard error in parentheses. * 0.10; ** 0.05; *** 0.01.

Figure 1: Binscatters of Salary, Job Prestige, and PhD Quality

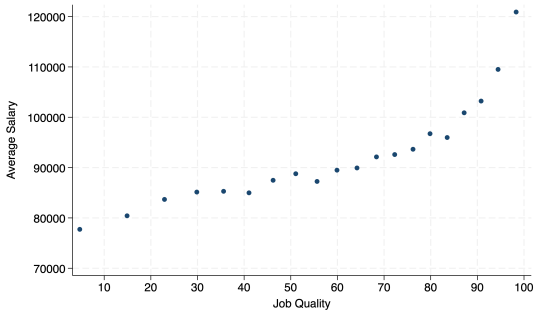
(a) Salary on PhD Quality



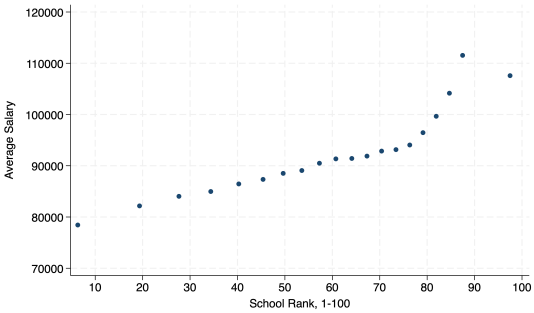
(b) Job Prestige on PhD Quality



(c) Salary on Job Prestige



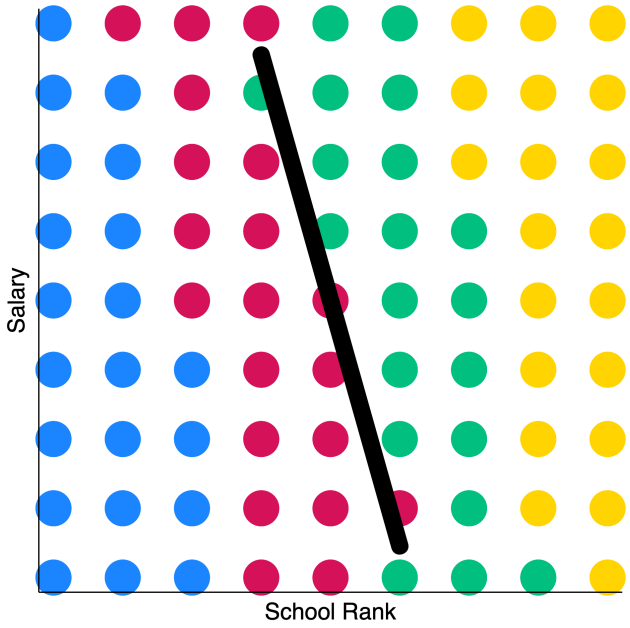
(d) Salary on Job Prestige, Residualize on PhD Quality



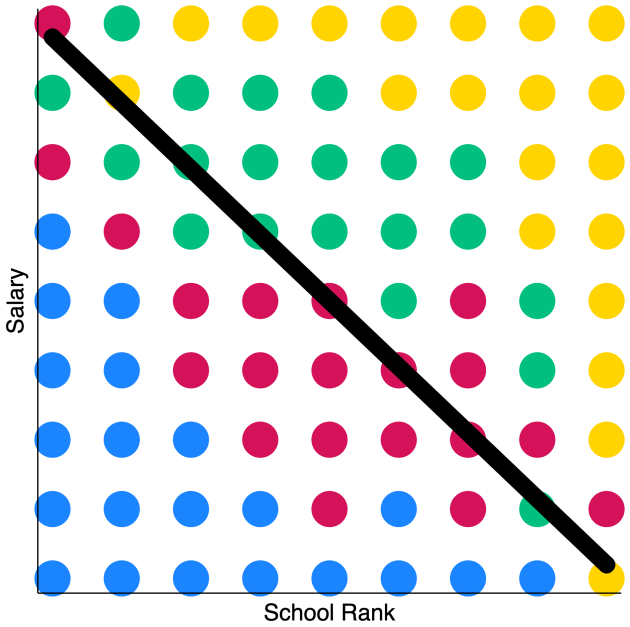
Notes: Each panel includes a binscatter: Panel A) average salary on PhD quality (SpringRank, 1–100); Panel B) job prestige (SpringRank, 1–100) on PhD quality (SpringRank, 1–100); Panel C) average salary on job prestige (SpringRank, 1–100); and Panel D) average salary on job prestige (SpringRank, 1–100), residualizing on PhD quality (SpringRank, 1–100).

Figure 2: Nonparametric

Panel A: PhD Quality as AIV

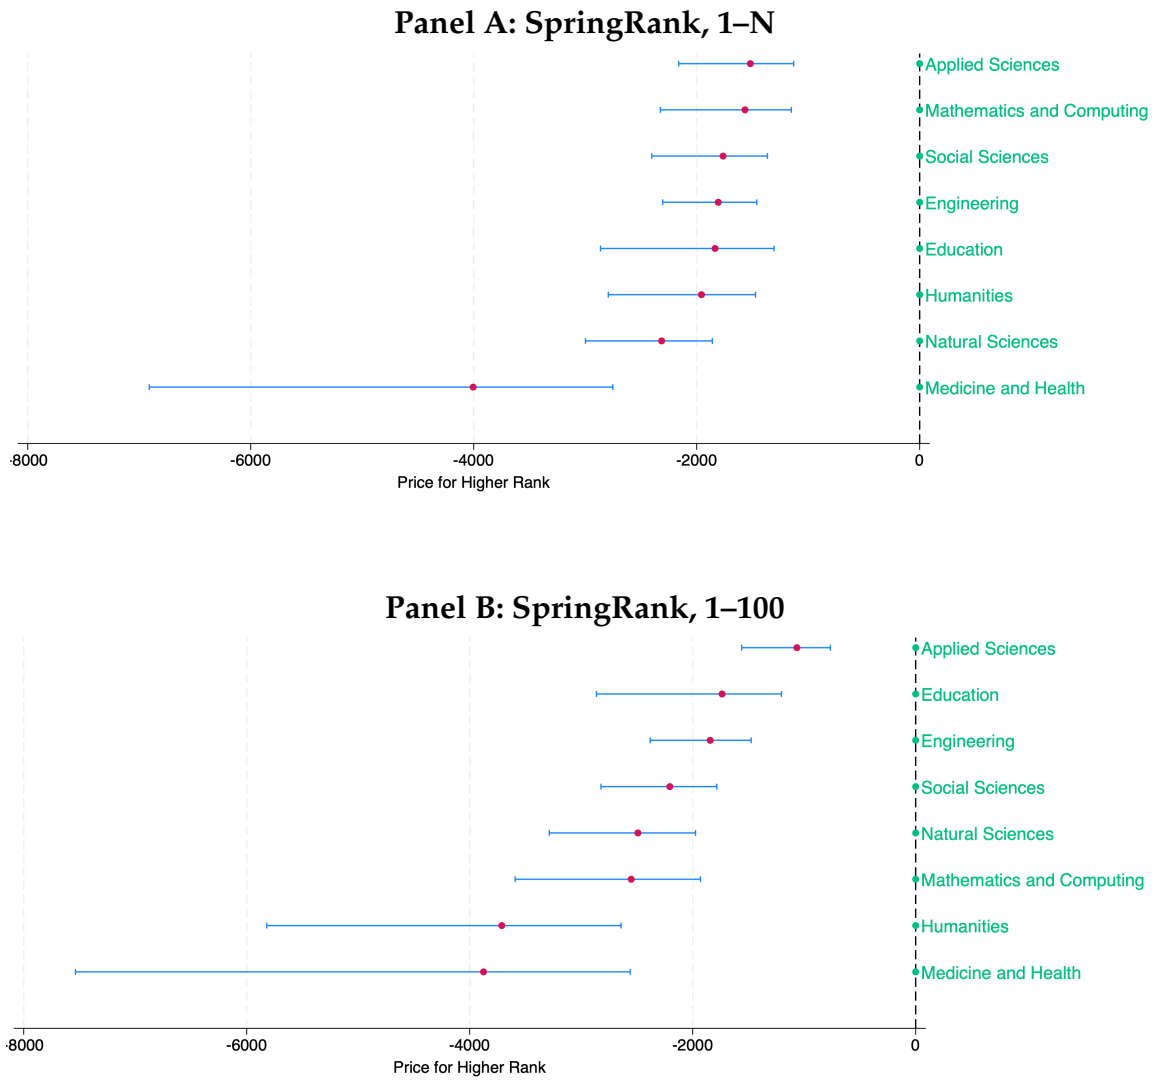


Panel B: Undergraduate Acceptance Rate as AIV



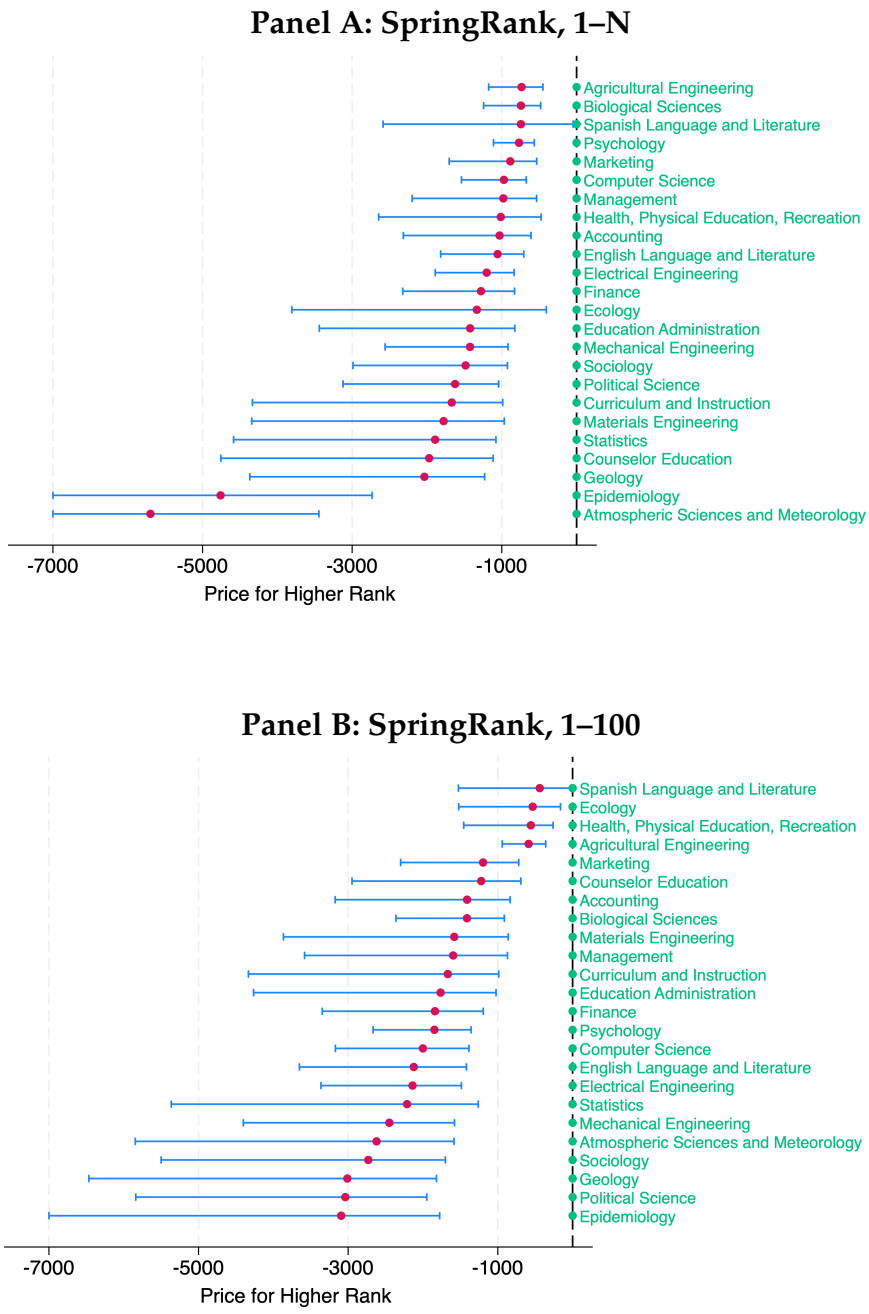
Notes: To create the figure in Panel A, we take the following steps: 1) By field, de-mean average salary, job prestige (SpringRank, 1–N), and PhD quality (SpringRank, 1–100). 2) Create quantiles (9 categories) for average salary and for job prestige. 3) Collapse the data to the average salary quantile -by- job prestige quantile level, computing means for average salary, job prestige, and PhD quality. 4) Create quantiles (4 categories) of PhD quality. 5) Graph, where we assign each average salary quantile -by- job prestige quantile level to its corresponding PhD quality quantile, shading them blue, red, green, and yellow, from lowest to highest. Finally, we show a line obtained from a applying the anti-IV method to the data (after constructing the quantiles, not the raw data). Panel A follows the same process, except that it uses as anti-IV the undergraduate acceptance rate at the PhD university.

Figure 3: By Domain



Notes: For each domain, we subset to fields within that domain and use the anti-IV approach of average salary on job quality, with PhD quality (SpringRank, 1–N) as the anti-instrument. In Panel A (B), we use SpringRank, 1–N (SpringRank, 1–100) as our measure of job quality. We cluster at the field-school level and use Anderson-Rubin standard errors. The dots are the point estimates, with the blue lines representing the confidence intervals.

Figure 4: By Field



Notes: For each field, we subset to that field and use the anti-IV approach of average salary on job quality, with PhD quality (SpringRank, 1–N) as the anti-instrument. In Panel A (B), we use SpringRank, 1–N (SpringRank, 1–100) as our measure of job quality. The dots are the point estimates, with the blue lines representing the confidence intervals. We present only fields that have a partial F statistic of at least 10. We also truncate confidence intervals at -7,000 where applicable.

A Appendix Figures and Tables

Table A.1: Comparison of Measures of Prestige, Economics

	Spring	Production	USNWR	Accept Rate
Massachusetts Institute of Technology	1	1	2	7.3
Harvard University	2	1	1	5.0
Princeton University	3	1	8	5.6
Stanford University	4	1	5	5.2
Yale University	5	1	6	6.5
University of Chicago	6	7	3	7.3
Northwestern University	7	7	7	9.3
University of California-Berkeley	8	1	4	17.5
University of Minnesota-Twin Cities	9	16	11	69.6
University of Pennsylvania	10	10	10	9.0
University of Rochester	11	21	18	35.4
Carnegie Mellon University	12	20	37	17.3
University of Michigan-Ann Arbor	13	12	12	26.1
University of Wisconsin-Madison	14	12	9	57.2
University of California-Los Angeles	15	12	14	14.3
University of California-San Diego	16	12	23	36.6
Brown University	17	19	28	7.7
Cornell University	18	16	13	10.7
New York University	19	11	21	21.1
Columbia University in the City of New York	20	9	17	6.7
University of Iowa	21	42	36	84.3
University of Illinois Urbana-Champaign	22	29	16	63.3
Johns Hopkins University	23	23	35	11.1
University of Maryland-College Park	24	21	30	51.0
Boston University	25	23	41	20.1
Washington University in St Louis	26	27	34	16.0
Duke University	27	16	15	7.7
The University of Texas at Austin	28	27	25	32.0
The New School	29	91	63	69.1
University of Virginia-Main Campus	30	29	27	22.6

Notes: This table compares measures of prestige for 30 top-ranked universities in economics. It is sorted by SpringRank, and the top 30 universities along this measure are included. We also present US News & World Report rankings, PhD production rankings (how many PhDs a university produced that are in economics in the Wapman et al. (2022) data), and the undergraduate acceptance rate (which has not been converted to a rank in this table).

Table A.2: Comparison of Measures of Prestige, History

	Spring	Production	USNWR	Accept Rate
Harvard University	1	5	1	5.0
University of California-Berkeley	2	1	3	17.5
Stanford University	3	1	10	5.2
Yale University	4	1	2	6.5
Columbia University in the City of New York	5	7	4	6.7
Princeton University	6	1	7	5.6
University of Chicago	7	5	5	7.3
University of Michigan-Ann Arbor	8	7	8	26.1
Brandeis University	9	36	30	33.4
University of Pennsylvania	10	9	11	9.0
University of Wisconsin-Madison	11	14	6	57.2
Johns Hopkins University	12	9	12	11.1
University of California-Los Angeles	13	9	9	14.3
New York University	14	17	16	21.1
Brown University	15	17	21	7.7
Massachusetts Institute of Technology	16	28	51	7.3
Northwestern University	17	14	20	9.3
Cornell University	18	12	17	10.7
University of North Carolina at Chapel Hill	19	12	15	25.0
University of Rochester	20	64	42	35.4
University of Minnesota-Twin Cities	21	24	23	69.6
University of California-San Diego	22	33	29	36.6
Rutgers University-New Brunswick	23	20	19	66.9
University of Virginia-Main Campus	24	20	14	22.6
The University of Texas at Austin	25	17	13	32.0
University of Washington-Seattle Campus	26	26	31	55.9
University of Illinois Urbana-Champaign	27	22	25	63.3
Indiana University-Bloomington	28	22	18	80.4
University of California-Davis	29	28	35	46.4
Duke University	30	14	22	7.7

Notes: This table compares measures of prestige for 30 top-ranked universities in history. It is sorted by SpringRank, and the top 30 universities along this measure are included. We also present US News & World Report rankings, PhD production rankings (how many PhDs a university produced that are in history in the Wapman et al. (2022) data), and the undergraduate acceptance rate (which has not been converted to a rank in this table).

Table A.3: Comparison of Measures of Prestige, Computer Science

	Spring	Production	USNWR	Accept Rate
Stanford University	1	1	4	5.2
Massachusetts Institute of Technology	2	1	1	7.3
University of California-Berkeley	3	1	2	17.5
Carnegie Mellon University	4	1	5	17.3
Harvard University	5	17	18	5.0
California Institute of Technology	6	11	47	6.7
Cornell University	7	5	9	10.7
Princeton University	8	8	10	5.6
University of California-Los Angeles	9	14	17	14.3
The University of Texas at Austin	10	8	7	32.0
University of Wisconsin-Madison	11	11	14	57.2
University of Illinois Urbana-Champaign	12	5	3	63.3
University of Pennsylvania	13	17	12	9.0
University of Rochester	14	47	54	35.4
Rice University	15	20	37	10.9
Yale University	16	20	36	6.5
Georgia Institute of Technology-Main Campus	17	10	6	21.3
University of Washington-Seattle Campus	18	7	15	55.9
Brown University	19	20	32	7.7
Columbia University in the City of New York	20	17	25	6.7
University of Michigan-Ann Arbor	21	13	13	26.1
University of Maryland-College Park	22	14	11	51.0
University of Cincinnati-Main Campus	23	99	79	76.4
University of Massachusetts-Amherst	24	20	16	65.3
University of California-San Diego	25	14	28	36.6
University of North Carolina at Chapel Hill	26	20	23	25.0
University of California-Santa Barbara	27	35	24	36.7
University of Minnesota-Twin Cities	28	35	20	69.6
University of California-Irvine	29	28	27	29.9
Duke University	30	27	56	7.7

Notes: This table compares measures of prestige for 30 top-ranked universities in computer science. It is sorted by SpringRank, and the top 30 universities along this measure are included. We also present US News & World Report rankings, PhD production rankings (how many PhDs a university produced that are in computer science in the Wapman et al. (2022) data), and the undergraduate acceptance rate (which has not been converted to a rank in this table).

Table A.4: Comparison of Naive and AIV Approaches, Data Are Collapsed

	(1)	(2)	(3)
Job Prestige	232.49***	111.01***	-2371.85***
	(4.36)	(7.48)	(138.35)
PhD Quality		373.65***	
		(19.81)	
Constant	74339.89***	46220.95***	
	(318.87)	(1548.45)	
Observations	7612	7612	7612
AIV	No	No	PhD Quality

Notes: The data are collapsed to the school-field level. Column 1 shows a regression of average salary on job prestige (SpringRank, 1–N). Column 2 adds PhD quality (SpringRank, 1–N). Column 3 regresses average salary on job prestige (SpringRank, 1–N), with PhD quality (SpringRank, 1–N) as the anti-instrument. All regressions include field fixed effects and are unweighted. We cluster at the school-field level for all columns and use Anderson-Rubin standard errors for column 3. Standard errors in parentheses. * 0.10; ** 0.05; *** 0.01.

Table A.5: Comparison of Different Anti-IV Variables, Using SpringRank 1–100

	(1)	(2)	(3)	(4)
Job Prestige, 1-100	-2244.65***	-3276.34***	-2043.01***	-401.00***
	(126.57)	(221.46)	(142.13)	(23.58)
Observations	181448	181448	181448	181448
AIV	PhD Quality, 1-N	PhD Quality, 1-100	Production, 1-N	UG Acceptance at PhD
Mean of AIV	111.0	78.4	25.8	69.2
SD of AIV	61.8	22.0	29.9	44.9

Notes: This table shows a regression of average salary on job prestige, where we use the SpringRank, 1–100 version. We vary the anti-IV we use across columns: we use PhD quality (SpringRank, 1–N) in column 1, PhD quality (SpringRank, 1–100) in column 2, PhD production (ranking of the number of PhDs one's PhD school produced that are in the sample) in column 3, and the undergraduate acceptance rate at one's PhD school in column 4. The means and standard deviations of the anti-IVs are presented at the bottom of the the tables. All regressions include field fixed effects. We cluster at the school-field level and use Anderson-Rubin standard errors. Standard errors in parentheses. * 0.10; ** 0.05; *** 0.01.

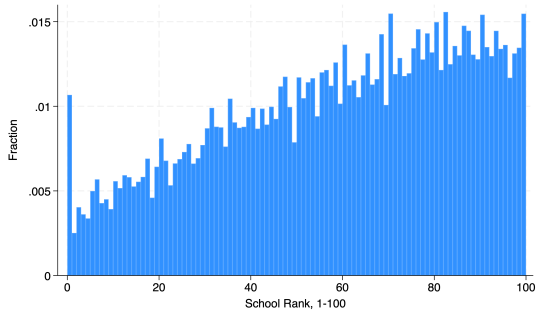
Table A.6: Comparison of Different Measures of Prestige, Naive Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Job Prestige, 1-N (Preferred)	220.56***	208.81***						
	(5.26)	(5.28)						
Job Prestige, 1-100			335.07***	316.41***				
			(7.26)	(7.17)				
Job PhD Production, 1-N					165.77***	146.27***		
					(5.02)	(4.76)		
UG Acceptance Rate, 1-N							287.69***	270.15***
							(6.22)	(6.05)
PhD Quality		36.89***		43.07***		101.45***		87.28***
		(2.04)		(2.27)		(3.14)		(3.60)
Constant	73859.75***	70746.50***	72563.51***	68881.24***	72315.30***	63402.92***	76199.46***	67491.61***
	(432.28)	(454.22)	(416.09)	(457.89)	(557.69)	(627.66)	(319.42)	(460.55)
Observations	181448	181448	181448	181448	181448	181448	181448	181448

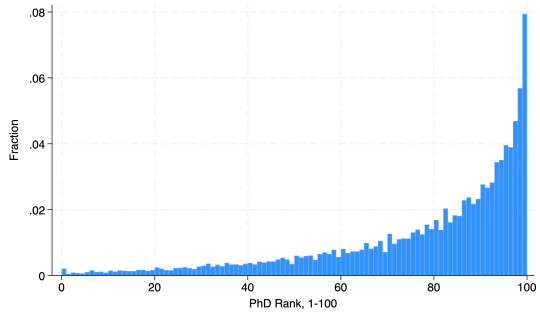
Notes: This table varies the definition of job prestige, with our preferred measure (SpringRank, 1–N) in columns 1 and 2; SpringRank, 1–100 in columns 3 and 4; PhD production (ranking of the number of PhDs one's PhD school produced that are in the sample) in columns 5 and 6; and the undergraduate acceptance rate at the job in columns 7 and 8. The odd-numbered columns show a regression of average salary on the relevant measure of job prestige. The even-numbered columns add PhD quality (SpringRank, 1–N). All regressions include field fixed effects. We cluster at the school-field level for all columns. Standard errors in parentheses. * 0.10; ** 0.05; *** 0.01.

Figure A.1: Distribution of SpringRank, 1-100 for School; SpringRank, 1-100 for PhD, and Undergraduate Acceptance Rate at PhD Institution

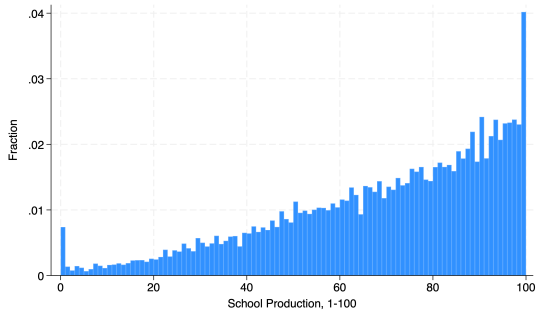
(a) SpringRank, 1-100 for School



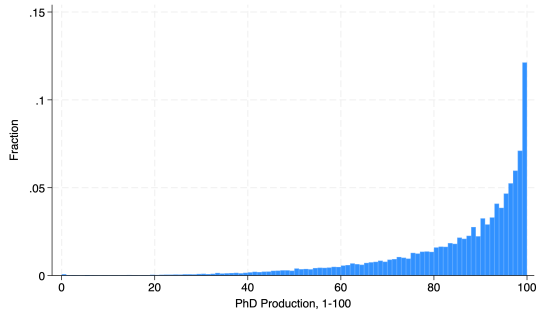
(b) SpringRank, 1-100 for PhD



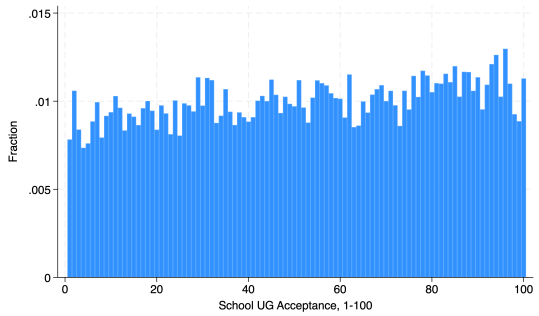
(c) Production, 1-100 for School



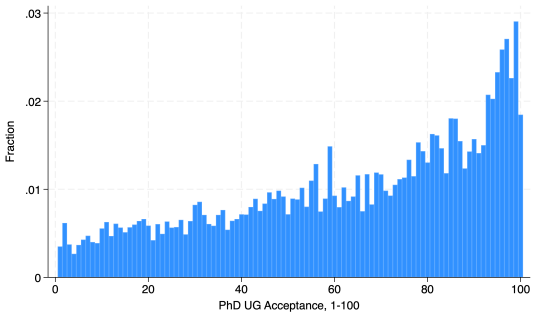
(d) Production, 1-100 for PhD



(e) Undergraduate Acceptance Rate, 1-100 for School



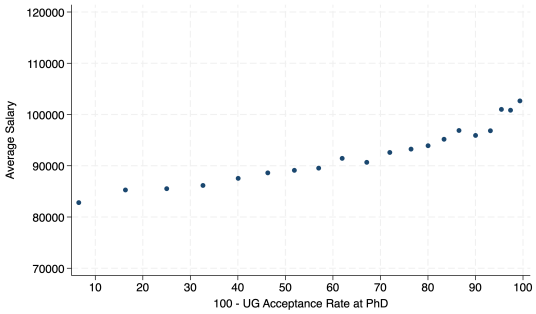
(f) Undergraduate Acceptance Rate, 1-100 for PhD



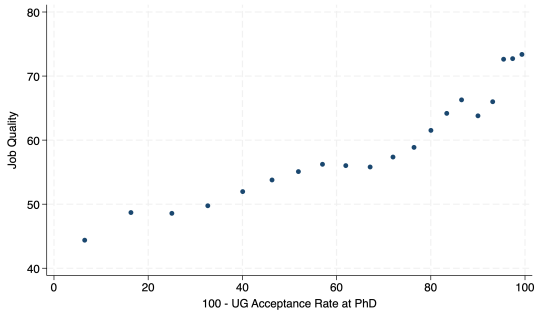
Notes: These figures are histograms of SpringRank, 1–100 at a faculty member's school (Panel A) and PhD (Panel B); PhD production (Panels C and D); and undergraduate acceptance rate (Panels E and F).

Figure A.2: Binscatters of Salary, Job Prestige, and Undergraduate Acceptance Rank at PhD

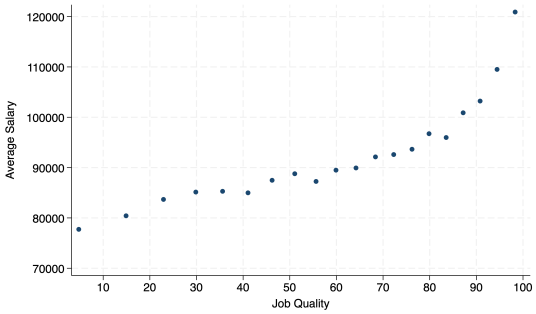
(a) Salary on PhD UG Acceptance Rate



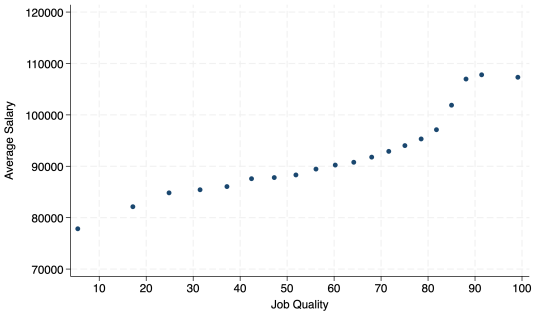
(b) Job Prestige on PhD UG Acceptance Rate



(c) Salary on Job Prestige



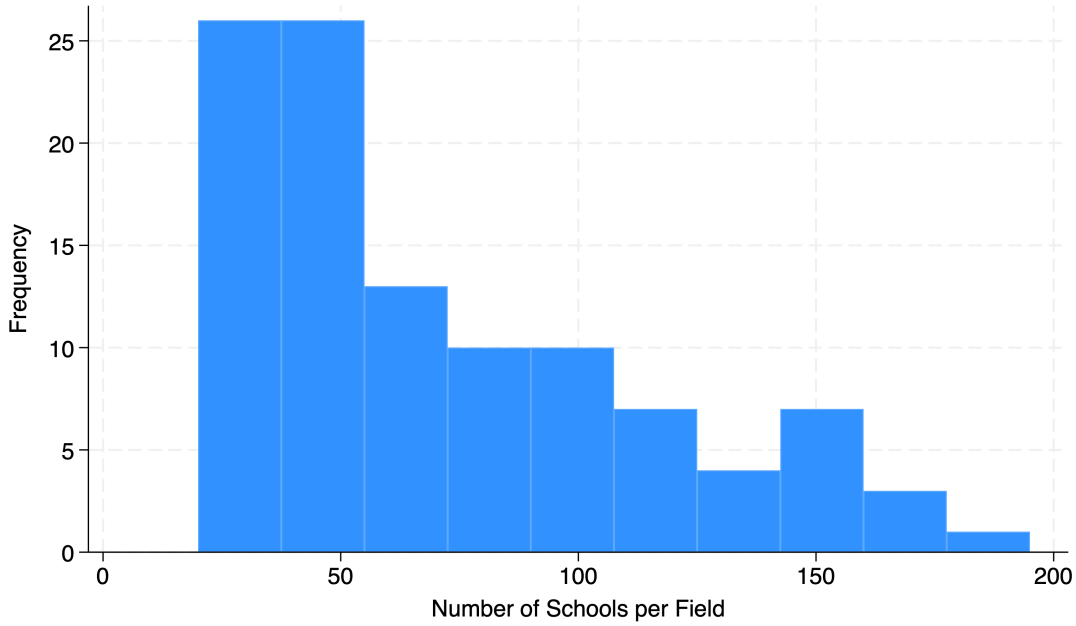
(d) Salary on Job Prestige, Residualize on PhD UG Acceptance Rate



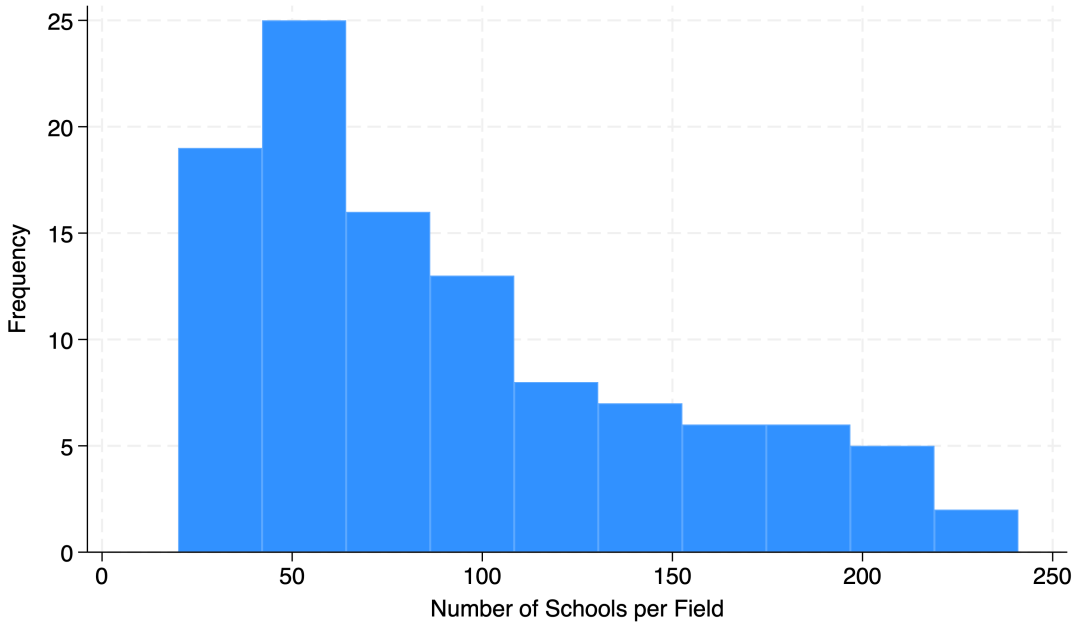
Notes: Each panel includes a binscatter: Panel A) average salary on the undergraduate acceptance rank at the PhD university; Panel B) job prestige (SpringRank, 1–100) on the undergraduate acceptance rank at the PhD university; Panel C) average salary on job prestige (SpringRank, 1–100); and Panel D) average salary on job prestige (SpringRank, 1–100), residualizing on the undergraduate acceptance rank at the PhD university.

Figure A.3: Distribution of Number of Schools per Field

(a) Based on Highest Observed Ordinal Rank



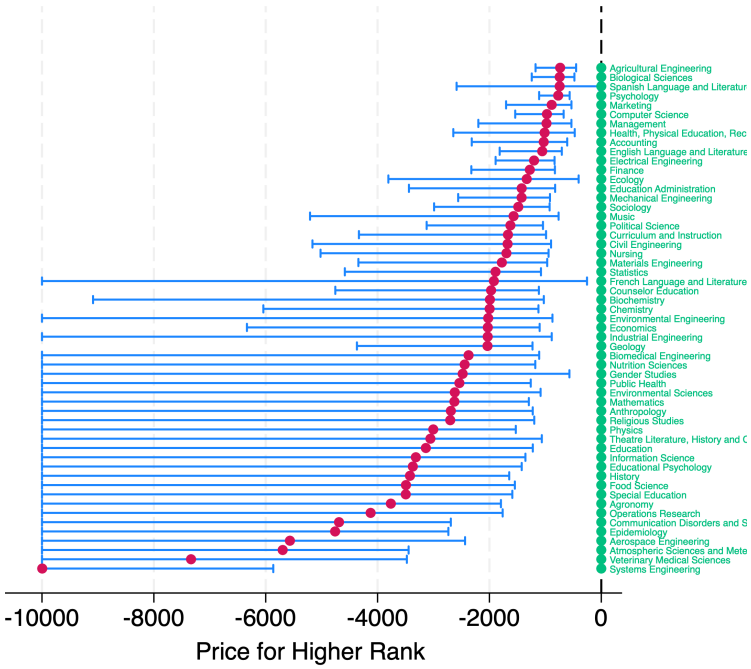
(b) Based on Number of Observed Schools



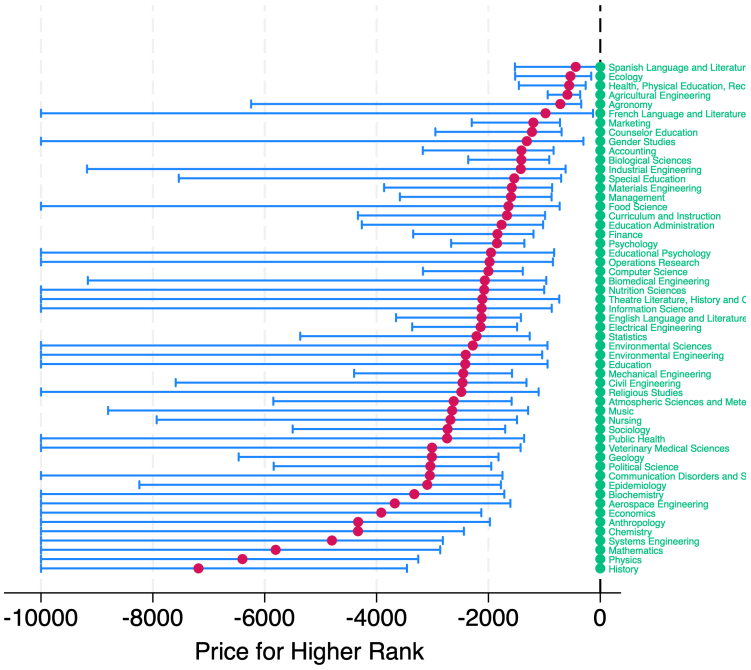
Notes: These panels show histograms for the number of schools per field. Panel A is based on the highest-observed ordinal SpringRank in the data, while Panel b is based on the number of schools observed in the data.

Figure A.4: By Field

Panel A: SpringRank, 1–N



Panel B: SpringRank, 1–100



Notes: For each field, we subset to that field and use the anti-IV approach of average salary on job quality (SpringRank, 1–N), with PhD quality (SpringRank, 1–N) as the anti-instrument. We cluster at the field level and use Anderson-Rubin standard errors. The dots are the point estimates, with the blue lines representing the confidence intervals. We present all fields with a non-infinite confidence interval. We truncate confidence intervals at -10,000 where applicable.

B Data Appendix

B.1 Additional Details on WZCL Data

First, they standardized the names of the PhD-granting universities, necessary because sometimes a single university appeared under different names. Second, they cleaned the terminal degree. For analyses involving prestige, they considered only faculty with a PhD. Third, they de-duplicated departments, combining a department that is coded as having different names in different years. Fourth, they mapped departments to the more granular field-level, such as physics and astronomy, and to the less-granular domain, such as natural sciences.[9] Each department gets mapped to one field, or, in some cases, to more than one field—for instance a department of physics and astronomy would get mapped to both physics and astronomy (both in the natural sciences). For this reason, the unit of observation is a person-field (or person-domain), entailing that a person can appear multiple times in the dataset. If a university has multiple departments for a given field, then all of these departments are considered. For instance, if there is a department of economics and a department of business economics, then both would be considered as economics. Fifth, they cleaned cases in which a faculty member appeared to be employed by two universities in a given year, due to mid-career moves. Sixth, they addressed cases where a faculty member disappears in a year(s) from the data, only to re-appear later. Seventh, they attempted to identify a faculty member's primary appointment and discarded other appointments (such as courtesy appointments). They do not consider faculty (1.2%) for which they cannot determine the primary appointment in their analysis (other than when considering academia as whole). Eighth, they assign a professor to the university at which they were most recently employed (relevant if faculty switch universities mid-sample). Ninth, they removed departments that did not appear in most of the sample years. In addition, they employed several steps to identify the gender of faculty, and were able to do so for 85% of the sample.

We make use of a file that has been collapsed (separately) to the field-job-PhD level, to the domain-job-PhD level, and to the academia-job-PhD level, where job is the university of the professor and PhD is the university at which they earned their terminal degree.[10] Each row contains the total number of faculty at that level, as well as the total number of women and the total number of men. After expanding the data at the academia level, there are 274,474 observations.[11] At the domain level, there are 293,905 observations, and at the field level, there are 317,009 observations. A key reason why there are more observations at finer levels of taxonomy is that, as mentioned above, an individual can appear in multiple fields and in multiple domains. Another complication in comparing across levels is that some fields are excluded from the domain and academia levels (due to being too small), and two domains are excluded from academia.

[9] A field maps into only one domain.

[10] We are unable to observe the departments of the schools and PhDs, only the universities.

[11] This is after dropping 348 cells with schools but no PhDs included; these cells have a missing value for total. The 274,474 number is somewhat lower than the 295,089 total that the authors report in their paper.

B.2 Mapping of Fields to Domains

Applied Sciences: Accounting, Agronomy, Animal Sciences, Architecture, Business Administration, Finance, Food Science, Horticulture, Management, Management Information Systems, Marketing, Soil Science, Urban and Regional Planning.

Education: Counselor Education, Curriculum and Instruction, Education, Education Administration, Special Education.

Engineering: Aerospace Engineering, Agricultural Engineering, Biomedical Engineering, Chemical Engineering, Civil Engineering, Electrical Engineering, Environmental Engineering, Industrial Engineering, Materials Engineering, Mechanical Engineering, Operations Research, Systems Engineering.

Humanities: Art History and Criticism, Asian Languages, Asian Studies, Classics and Classical Languages, Comparative Literature, English Language and Literature, French Language and Literature, Germanic Languages and Literatures, History, Linguistics, Music, Near and Middle Eastern Languages and Cultures, Philosophy, Religious Studies, Slavic Languages and Literatures, Spanish Language and Literature, Theatre Literature, History and Criticism, Theological Studies.

Mathematics and Computing: Computer Engineering, Computer Science, Information Science, Information Technology, Mathematics.

Medicine and Health: Anatomy, Communication Disorders and Sciences, Environmental Health Sciences, Epidemiology, Exercise Science, Kinesiology, Rehab, Health, Genetics, Health, Physical Education, Recreation, Human Development and Family Sciences, Immunology, Nursing, Pathology, Pharmaceutical Sciences, Pharmacology, Pharmacy, Physiology, Public Health, Speech and Hearing Sciences, Veterinary Medical Sciences.

Natural Sciences: Astronomy, Atmospheric Sciences and Meteorology, Biochemistry, Biological Sciences, Biophysics, Biostatistics, Cell Biology, Chemistry, Ecology, Entomology, Environmental Sciences, Evolutionary Biology, Forestry and Forest Resources, Geology, Marine Sciences, Microbiology, Molecular Biology, Natural Resources, Neuroscience, Nutrition Sciences, Physics, Plant Pathology, Plant Sciences, Statistics.

Social Sciences: Agricultural Economics, Anthropology, Criminal Justice and Criminology, Economics, Educational Psychology, Gender Studies, Geography, International Affairs, Political Science, Psychology, Social Work, Sociology.