WHAT JOBS COME TO MIND? STEREOTYPES ABOUT FIELDS OF STUDY*

John J. Conlon and Dev Patel

We test for stereotyping—the exaggeration of representative traits—in a high-stakes economic environment. Using surveys administered among undergraduates at the Ohio State University as well as large-scale nationally representative data, we measure how U.S. first-year students perceive the relationship between college majors and occupations. We show that students greatly overestimate the likelihood that majors lead to their representative jobs (e.g., counselor for psychology, journalist for journalism). Using an implicit association test, we show that students associate majors with their representative careers and that these associations strongly predict belief biases, in line with a stereotyping mechanism. A simple equilibrium model of the labor market predicts that stereotyping reduces welfare by increasing misallocation, which we corroborate with correlational evidence on job/major mismatch. In a field experiment, we test a light-touch policy to reduce stereotyping and find significant effects on students' intentions about what to study as well as the classes and majors they enroll in. *JEL codes:* D83, D91, I23, J24, C93.

* We are grateful to the editor, Stefanie Stantcheva, and four anonymous referees for comments that greatly improved the article. We thank the Exploration Program at Ohio State for their support in implementing our surveys, and Jenna Anders, Zach Bleemer, Pedro Bordalo, Emily Breza, Katie Coffman, Lucas Coffman, Meg Edison, Ben Enke, Nicola Gennaioli, Ed Glaeser, Claudia Goldin, Larry Katz, Gabriel Kreindler, Ross Mattheis, Muriel Niederle, Matthew Rabin, Gautam Rao, Josh Schwartzstein, Andrei Shleifer, Evan Soltas, and Basit Zafar for helpful feedback. This study underwent human subjects review at Harvard University under IRB20-1581. We are grateful for financial support from the Lab for Economic Applications and Policy, the Mind Brain Behavior Initiative, the National Science Foundation Graduate Research Fellowship under grant DGE1745303, and the Alfred P. Sloan Foundation Predoctoral Fellowship in Behavioral Macroeconomics, awarded through the National Bureau of Economic Research. The experiments described in this article were preregistered on the AEA RCT Registry (IDs AEARCTR-0008339 and AEARCTR-0015937).

The Quarterly Journal of Economics (2026), 1–49. https://doi.org/10.1093/qje/qjag031.
Advance Access publication on June 22, 2026.

I. Introduction

College major choice is among the most consequential economic choices that many people ever make.[1] Undergraduates choose their major in part with an eye toward its effect on their future occupation, but several facts suggest frictions in this process. Almost half of employed U.S. college graduates report that their job is not closely related to their degree, and mismatch between major and occupation strongly predicts job dissatisfaction, lower salaries, and the rate at which graduates say they regret their choice of what to study.[2] What explains these patterns? One possibility is that they primarily reflect uncertainty—for instance, about undergraduates' future job preferences—resolving into ex post regret. But a second possibility is that they in part reflect systematic ex ante mistakes, for example, due to student misperceptions about the consequences of major choice. Identifying the nature and magnitude of any such mistakes would point toward policy interventions to improve student outcomes.

In this article, we study whether stereotypes about the link between college majors and careers distort undergraduates' beliefs and thereby their choices of what to study. We draw on a recent literature in economics that conceptualizes stereotypes as exaggerations based on a "kernel of truth": a trait that is more common in one group than another, even if it is uncommon in absolute terms in both, may come to mind readily when thinking about that group (Gennaioli and Shleifer 2010; Bordalo et al. 2016).[3] This kind of stereotyping arising from relative likelihoods or representative traits has been widely applied to study misperceptions sur-

1. College major choice plays a large and increasing role in shaping the economic prospects of college graduates (Altonji, Kahn, and Speer 2014). Differences in, for example, earnings across majors often rival or exceed the wage premium from attending college at all, and they appear to primarily reflect causal effects rather than selection (Hastings, Neilson, and Zimmerman 2013; Kirkeboen, Leuven, and Mogstad 2016; Bleemer and Mehta 2022). These effects appear driven at least in part by the actual classes that students take, rather than (just) the official major they graduate with (Arteaga 2018).

2. Our own calculations from the National Survey of College Graduates and the Survey of Household Economics and Decision Making. See Section III.

3. Note that this cognitive approach to stereotyping is distinct from other approaches that consider stereotypes as reflecting taste or rational statistical inferences (Becker 1957; Phelps 1972), as essentially negative caricatures of groups driven by prejudice (Adorno et al. 1950), or as tools used to perpetuate negative views of historically discriminated-against groups (Glaeser 2005).

rounding social groups, such as gender, race, and political parties (e.g., Bordalo et al. 2019; Arnold, Dobbie, and Yang 2018; Coffman, Exley, and Niederle 2021; Arnold, Dobbie, and Hull 2022; Alesina, Miano, and Stantcheva 2023; Bohren et al. 2025).

The same logic that grounds stereotypes in likelihood ratios—how common a trait is in one group divided by its prevalence in a reference group—also makes stark predictions in our setting. We can think of graduates with different majors as comprising different "groups" and their careers as "traits" that may be more common in one group than another. We can then define for each college major its representative career: the one that is most common among that major's graduates compared to other majors. We compute representativeness using career/major shares in U.S. census data, which yields intuitive classifications for each major's most representative career: for example, lawyers for political science majors, journalists for communication majors, teachers for education majors. We show that for such careers, absolute frequencies of representative careers conditional on major are often modest—between 2% and 60% depending on field—but relative frequencies are nonetheless very extreme: graduates are 155% to 1,751% more likely to have their major's representative career than graduates with other majors. Motivated by these statistics, we describe a simple model of belief formation (drawn from Bordalo et al. 2016) whereby representative careers come to mind easily and thereby get overweighted: if so, we should expect exaggerated beliefs about representative careers.

We test this prediction first through surveys run in partnership with an academic program at the Ohio State University (OSU). Our respondents are first-year undergraduates who are undecided about their major. The main survey questions simply ask students their beliefs about the share of U.S. graduates with different majors who are working in various careers. We find large and systematic stereotyping: students overestimate the prevalence of almost every major's representative career. For example, the average student believes that 53% of art majors work as artists (17% do), that 47% of journalism majors work as journalists (4% do), that 38% of political science majors work as lawyers (16% do), and that 43% of psychology majors work as counselors (21% do). Very similar patterns appear in students' beliefs about majors they are versus are not likely to pursue and in their beliefs about the careers they themselves are likely to have, facts that rule out other mechanisms like overconfidence or motivated reasoning as driving our main result. We find similar large and systematic stereotyping in a sample of U.S. adults, both among college-educated and non-college-educated respondents and for younger and older respondents.

To provide evidence on the breadth and persistence of such stereotyping, we analyze a large-scale, nationally representative sample of U.S. college first-year students including more than 9 million respondents over 40 years. We find that dramatically more first-years expect to attain their major's representative career than actually end up working in that job: 63% of biology majors expect to be doctors (versus 23% in reality), 62% of psychology majors expect to be counselors (versus 21%), 65% of art majors expect to be artists (versus 17%), 42% of communications/journalism majors expect to be writers or journalists (versus 4%), and so on. Pooling across majors, these facts combine to produce large gaps between the expected and actual careers that college graduates pursue. For professions that are primarily the rare-but-representative outcome of particular majors—doctor, counselor, journalist—two to four times more college first-years expect to work in these jobs than actually do. In contrast, many fewer expect to be teachers, work in business, or be nonemployed than ultimately are, because these are the common alternatives to the representative career of many majors. These overall gaps—amounting to 40,000 to 200,000 students a year—appear largely unchanged since at least the 1970s.

Our interpretation of these results is that representative careers come to mind easily when thinking about majors. This kind of recall process—what is immediately retrieved quickly from memory—is downstream of the associations between experiences that connect them in memory (Kahana 2012; Bordalo et al. 2021, 2023). To provide evidence on this potential mechanism, we adapt to our setting an extremely prevalent method of uncovering such associations from the psychology literature—the implicit association test (IAT) (Greenwald, McGhee, and Schwartz 1998), which is increasingly used in economics to document stereotyping (Bertrand, Chugh, and Mullainathan 2005; Lowes et al. 2015; Carlana 2019; Avitzour et al. 2020; Chetty et al. 2020; Kline, Rose, and Walters 2022; Alesina et al. 2024; Owen and Rury 2025). In the IAT, participants must sort stimuli into categories using two response keys. Their speed is taken as evidence of underlying associations: if participants sort more quickly when two categories share a key than when they do not, this is evidence of an implicit link or association in memory.

In our adaptation of the IAT, the categories are major and career groups, and the stimuli are individual occupation and degree fields. For each round, participants sort words as either belonging versus not to a focal major group (e.g., "humanities") or as belonging versus not to a focal career group (e.g., "writers and journalists"). The key comparison is between matched blocks, where the focal major and career share a response key, and unmatched blocks, where each is instead paired with the other's alternative category. The difference in completion times across these two blocks provides a measure of an implicit association between the focal major and career. We compute a difference-in-difference—comparing the matched versus unmatched blocks across different rounds of the IAT—to measure the relative strength of implicit associations across majors and careers.

Using this design, we find that participants strongly and systematically associate majors with their representative careers: implicit associations are 0.30–0.36 standard deviations higher for representative major–career pairs than for nonrepresentative pairs ($p < .01$). This effect remains (0.24–0.28 standard deviations, $p < .01$) even controlling for the true conditional frequency of careers. Furthermore, a one standard deviation increase in the associations between a major and its representative career predicts 4.1 percentage points greater stereotyping as measured by beliefs ($p < .01$). This relationship remains large and statistically significant even controlling for career-by-major fixed effects (2.8 percentage points, $p < .01$); that is, for a given major, those who reveal a greater association between that field of study and its representative career also have more biased beliefs. These results strongly suggest that the same associations that drive stereotyping in other settings also underpin the biases about college majors that we uncover. In addition, psychologists (and increasingly behavioral economists) argue that such associations and the belief biases they lead to are intimately linked to past experiences stored in memory (e.g., Greenwald and Banaji 1995; Gawronski and Bodenhausen 2011; Bordalo et al. 2016, 2023). Consistent with this idea, we find that heterogeneity across students in their experiences—the careers and majors of people they know personally—strongly predicts biases in their beliefs at the individual level.

Having provided evidence on the severity and origins of stereotyping in our context, we turn to its implications for economic outcomes and policy. We embed our model of representativeness-based belief distortions into a stylized setting where students decide their major and then—after graduating—choose their career. The model makes two intuitive points. First, if students hold internally consistent views of the economy, then stereotypes about the major–career mapping imply misperceptions about jobs' unobserved attributes: that is, students must infer that a representative career is better—for example, it requires less search to find a job, pays better, or provides better amenities—than it really is. The logic behind this result is that if students had correct beliefs about the distribution of job amenities and worker/student preferences, they would realize that their stereotyped beliefs about the distribution of careers must be wrong. Second, stereotyped beliefs lead to greater misallocation—given the occupations graduates ultimately end up pursuing, they would have been better off had they majored in a different field. Intuitively, this effect arises through selection: stereotyping draws in students on the margin who are less well suited to a major's representative job. Larger belief biases therefore cause more graduates to end up in jobs other than those for which their major prepared them.

We test whether stereotyping is correlated with misallocation by compiling data from multiple nationally representative surveys of college graduates. We find that graduates whose majors (in our OSU beliefs data) are the subject of greater stereotyping are significantly more likely to be dissatisfied with their jobs, feel their job does not fit their skills and experience, report that their job is unrelated to their field of study, and report that they regret the field of study they chose ($p < .05$ for all variables). These results are of course only correlational (and limited to the coarse major groups that we define), but they nonetheless provide suggestive evidence of the potential importance of stereotyped beliefs in distorting postgraduation labor market outcomes.

What policies can help students to make better informed choices? We conclude by analyzing the effect of a light-touch information intervention on students' self-reported intentions as well as their actual course enrollments and college major declarations up to three years later, measured using administrative data. We find that the effect of information depends critically on students' heterogeneous preferences over their future jobs. In response to news that the representative career of their ex ante most preferred major is 10 percentage points less likely, students reduce their intentions toward that major by 3.5 percentage points ($p < .01$), enroll in 0.22 fewer courses in that major in the next semester ($p < .05$), and lower their probability of having declared that major a year later by 6.1 percentage points ($p = .23$). Although the effects appear to partially fade out over the following years, we also find a large and significant effect on when students declare their major: treated students spend on average 0.21 more semesters undecided before declaring a major ($p < .05$).[4] In contrast to the effects we find on students' top-ranked major, reducing stereotyping if anything boosts the attractiveness of less preferred fields, consistent with heterogeneous preferences over careers: every 10 percentage point reduction in stereotyping about students' second-ranked major boosts intentions toward it by +2.1 percentage points ($p = .17$), immediate class-taking by +0.20 courses ($p < .10$), and major declaration within a year by +9.9 percentage points ($p < .01$). All these differences between first and second majors are significant at the $p < .05$ level.

This article contributes first to a large literature exploring stereotypes surrounding race, immigration, political affiliation, and gender (e.g., Arnold, Dobbie, and Yang 2018; Bordalo et al. 2019; Coffman, Exley, and Niederle 2021; Arnold, Dobbie, and Hull 2022; Alesina, Miano, and Stantcheva 2023; Bohren et al. 2025). We show that the same logic of stereotyping as exaggerations of a "kernel of truth" also applies to beliefs about human capital investments. Furthermore, our evidence connecting stereotyping, implicit associations, and role model effects adds to a growing literature on beliefs arising from associations, memory, and what comes easily to mind (e.g., Malmendier and Wachter 2022; Augenblick et al. 2023; Bordalo et al. 2023; Enke, Schwerter, and Zimmermann 2024; Graeber, Roth, and Zimmermann 2024; Bordalo et al. 2025; Conlon and Kwon 2025).

Our study also contributes to a rich literature on beliefs and human capital investment (see Giustinelli 2023 for a review). Misperceptions surrounding schooling choices are particularly important to identify because, unlike other premarket determinants of later economic success, such as accumulated skills (Neal and

4. If anything, treated students are slightly more likely to still be taking classes two or three years later, so the effects on major declaration do not appear to be driven by students dropping out or becoming discouraged with college.

Johnson 1996; Speer 2017), they may be more easily addressed through simple policy interventions such as information campaigns or other nudges. The importance of field of study for economic outcomes has led many studies to investigate whether biased beliefs distort students' college major choices, but these papers tend to focus on beliefs about average salary conditional on major (e.g., Betts 1996; Arcidiacono, Hotz, and Kang 2012; Wiswall and Zafar 2015; Baker et al. 2018; Conlon 2021).[5] We show that students are systematically miscalibrated about the occupations that different majors lead to, not just about the salaries those occupations pay. Our results also echo a small but growing literature showing that students in vocational programs in developing countries greatly overestimate their own employment prospects postgraduation (e.g., Bandiera et al. 2025) and that reducing these misperceptions can have positive effects (e.g., Alfonsi, Namubiru, and Sara 2025).

II. Do Students Stereotype Majors?

This section tests whether students stereotype majors. We begin by describing a representativeness-based model of belief formation that delivers a simple test for stereotyping, which we then take to the data. We present our main survey evidence testing for stereotyping, then examine robustness in other samples, and finally use implicit association tests to shed light on the psychology of stereotyping more directly. In Section III, we embed this framework in a larger model to study other predictions of this belief formation process, focusing on implications for major choice and other postgraduation outcomes.

II.A. A Model of Belief Formation

Let $p_{c|M}$ be the true share of graduates with major M who are working in career c. Let $\pi_{c|M}$ denote the student's belief about this share. We assume that the student may overweight careers that come to mind easily when considering M. Following Gennaioli and

5. A notable exception is Arcidiacono et al. (2020), who decompose beliefs about the salary returns to majors into their effects on salaries in occupations and on the likelihood of attaining certain occupations. Their study does not attempt to test whether students' beliefs are consistent with rational expectations. See also Wiswall and Zafar (2021) and Ersoy and Speer (2025) for beliefs about non–labor market consequences of major choice.

Shleifer (2010) and Bordalo et al. (2016), we assume that a career c comes to mind more easily when it is more common in M than in other majors (for microfoundations and more experimental evidence, see Bordalo et al. 2021, 2023). This produces a version of the representativeness heuristic (Kahneman and Tversky 1972) that gives a central role to the relative likelihood of traits in groups (in our case, careers within majors). More precisely, we define the representativeness $R(c, M)$ of career c for major M according to the likelihood ratio rule in equation (1), where $p_{c|-M}$ is the share of graduates with majors other than M that work in c.

$$R(c, M) \equiv \frac{p_{c|M}}{p_{c|-M}}. \tag{1}$$

We define the representative career (sometimes called the exemplar) of major M to be the career that maximizes this likelihood ratio: $c^*(M) \equiv \arg\max_c R(c, M)$.

We assume the student may overweight M's representative career when forming their beliefs because it comes to mind more easily than other careers. In particular, we assume the student's beliefs about M are given by:

$$\pi_{c|M} = (1 - \theta)\, p_{c|M} + \theta \cdot \mathbf{1}[c = c^*(M)]. \tag{2}$$

Equation (2) nests two important benchmarks. First, when $\theta = 0$, we recover the rational benchmark where the student's beliefs are correct. When $\theta = 1$, the equation reduces to the rank-based truncation formulation of stereotyping from Bordalo et al. (2016). We might also expect more intermediate levels of stereotyping, with $\theta \in (0, 1)$.

Equation (2) suggests a simple empirical strategy to test for stereotyping in this context: we can regress beliefs about the prevalence of c among M majors on the true prevalence $p_{c|M}$ and an indicator for c being M's representative career. We follow this strategy—along with other analyses exploring robustness—below.

II.B. Testing for Stereotyping

Building on the empirical test generated from the belief formation model, this section presents our main survey data and evidence on the magnitude of stereotyping.

1. Study Sample and Survey Design. Our primary data come from surveys of freshmen students in the Exploration program at OSU. Entering OSU students are automatically enrolled in this program if they have not yet officially declared a major. The Exploration program includes a mini-course that students take their first semester. This course is typically taught by an academic advisor and includes self-assessments (meant to point students toward majors that suit them), information about degree requirements, and help planning course schedules. At the time we ran our surveys, the program did not provide students quantitative information about employment outcomes by major. Students received extra credit in this course for completing our survey, ensuring a high response rate of around 80%.

We ran three surveys among students in the Exploration program. Our main survey occurred during fall semester of 2020. The other two surveys both occurred during fall semester of 2021 among a new cohort of Exploration students. The first 2021 survey served as a replication of our 2020 results, and the second included an information-provision experiment to evaluate the effect of reducing stereotyping on students' intentions and choices (see Section IV). See Online Appendix B.1 for more details about the surveys and their implementation. Online Appendix Table A.I, columns (1) and (2) give self-reported demographic information about the students in our OSU samples, which are broadly similar demographically to the overall student body at OSU, though with a somewhat higher share of first-generation college students.

The 2020 survey began by displaying 10 groups of college majors and asking students to rank them by how likely they thought they were to graduate from OSU with a degree in each. See Online Appendix Table A.II for a list of these major groups (henceforth, just "majors") and how we map them onto degree fields in the American Community Survey (ACS). The survey then asked detailed questions about a subset of these majors: the two they rated as being most likely to pursue and two other randomly selected majors. This structure allows us to estimate average beliefs among all students (i.e., unconditional on their ranking of majors) using inverse probability weights.[6]

The most important question asked students their "best guess about the percent of Americans aged 30–50 (note, not just from Exploration or OSU) who graduated with a major in M that are . . ." It

6. In particular, a student's two top majors receive a weight of one, while the two other majors receive a weight of four (because there were eight other majors, and thus a one in four chance that each was selected).

then listed nine careers, plus “Working in any other job” and “Not working for pay”. See Online Appendix Table A.III for the list of careers and how we map them onto occupations in the ACS. After answering this question, students were asked what they thought their own chance of working in each career would be, if they graduated with major M. We call these two sets of beliefs—about others’ outcomes and one’s own possible career—students’ “population” and “self” beliefs, respectively. Note that the population beliefs questions were worded to allow us to compare students’ beliefs to an objective benchmark: true career shares in the ACS, which we calculate using 2017–2019 data (Ruggles et al. 2022). To match the wording of these OSU survey questions, we restrict to college-graduate respondents born between 1958 and 1997 who are between 30 and 50 years old when answering the ACS.

2. Testing for Stereotyping. We begin by simply showing average beliefs about representative careers. We define the representative career of each major using the likelihood ratio according to equation (1). The careers that are most representative of each major by this definition are intuitive: doctors for biology/chemistry, lawyers for government, counselors for psychology, teachers for education, and so on (see Online Appendix Table A.IV for the complete list).

The gray bars in Figure I show students’ average beliefs about the share of each major’s graduates working in its most representative career. The dotted lines show the true fraction of college graduates with each major who are working in the representative career. We see that for 9 of the 10 majors (all except nursing) students believe that a major’s representative career is substantially (and statistically significantly) more common among graduates with that major than it actually is.[7] These differences are often quite large: students exaggerate the share of artists among art majors by 36 percentage points or 211%, of doctors among biology and chemistry majors by 11 percentage points (48%), of counselors among psychology majors by 22 percentage points (105%), of writers and journalists among communications majors by 43

7. Even the exception to this pattern is instructive. Although students underestimate the share of nursing majors working as nurses, this is largely because they dramatically overstate the share of such majors who eventually become doctors, which is nursing’s second most representative outcome. In fact, 4% of nursing majors work as doctors, but the average belief is 23%.

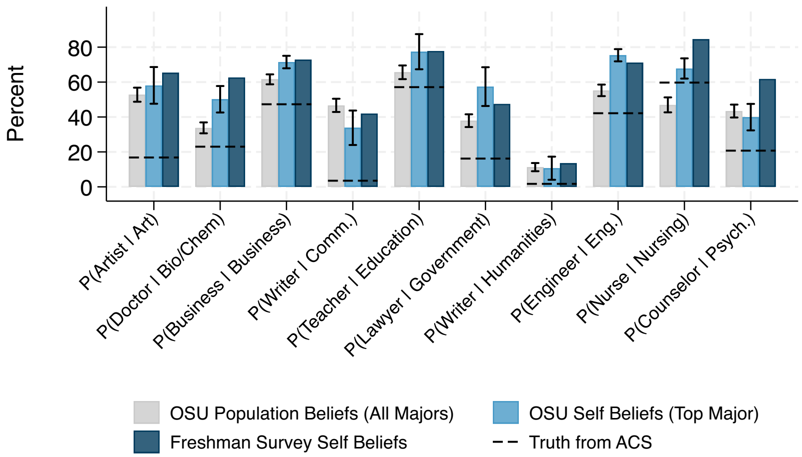

FIGURE I

Stereotyping Representative Careers

Figure I presents average statistics regarding the most representative career (as defined in Section II) for each major. The dashed horizontal lines denote the actual proportion of college graduates with each major between the ages of 30 and 50 that are working in that major's most representative career, based on data from the 2017–2019 American Community Survey. The dark blue bar shows, among students in the Freshman Survey who expect to pursue each major, what fraction list that major's representative outcome as their probable career occupation. The light blue bar plots the average belief for the 2020 OSU sample about the probability that they would be working in each career at age 30 if they graduated from OSU with their top-ranked (i.e., most likely) major. The pale gray bars show the average belief among our 2020 OSU sample about the fraction of Americans between the ages of 30 and 50 who graduated college with each major (not only their top-ranked major) that are working in each occupation. Error bars show 95% confidence intervals for the mean of the OSU beliefs.

percentage points (1,075%), and so on. All of these differences are statistically significant at the $p < .01$ level.

The light blue bars (color version online) in Figure I show average self beliefs, restricting to the major that students say they are most likely to pursue. We see that self beliefs tend to be quite similar to population beliefs, and for every major students believe they have a statistically significantly better chance of attaining their major's representative career than the true share of graduates working in it ($p < .05$ for all 10 comparisons). Note that by themselves, these patterns in students' self beliefs about their top-ranked major could have been interpreted as only reflecting, say, overconfidence or motivated reasoning. The fact that students also overestimate the share of others in representative careers, including for majors other than their own (pale gray bars), indicates that stereotyping is a large part of the explanation.[8]

The model described in Section II.A suggests a simple regression strategy for quantifying the extent of stereotyping. Table I shows OLS regressions of equation (3), clustering standard errors at the individual and career-by-major level.

$$\pi^{i}_{c|M} = \gamma p_{c|M} + \theta \mathbb{1}\Big(c = c^{*}(M)\Big) + \epsilon^{i}_{c,M}. \tag{3}$$

If participants react only to the true (not the relative) frequency of careers by major, we should expect the estimate of θ to be zero. Note that this is true not only under the rational benchmark but also under any explanation whereby errors in beliefs take a form unrelated to the likelihood ratio. For example, if students' beliefs are only noisy or underreactive to true frequencies, then we might expect an estimate of $\gamma \neq 1$ but would still expect θ to be zero.

Table I, column (1) instead shows a large, positive, and statistically significant estimate of θ, indicating substantial stereotyping in students' population beliefs. Given that these regressions control for the true frequency of careers, the coefficient of 0.32 ($p < .01$) indicates that on average students think that majors' most representative careers are 32 percentage points more common than equally prevalent but nonrepresentative careers. Column (2) adds career fixed effects, where we see a similarly sized estimate of stereotyping (28 percentage points, $p < .01$). Furthermore, Online Appendix Table A.VI shows very similar levels of stereotyping by gender, ethnicity, and first-generation status. Table I, columns (3) and (4) show that similar patterns appear for self beliefs, restricting to students' top-ranked major: students believe that they are between 36 and 43 percentage points ($p < .01$, with and without career fixed effects) more likely to have their major's representative career than equally prevalent nonrepresentative careers.

8. We can also decompose the relative importance of stereotyping compared with these other potential mechanisms. Online Appendix Table A.V presents the results of a Shapley-Sharrocks decomposition. This analysis conducts a horse race, comparing the extent to which stereotyping versus other potential mechanisms explain the variance in students' beliefs (see the notes of Online Appendix Table A.V for details). We find that stereotyping accounts for a greater share of the variation in beliefs than any of the other mechanisms we explore.

TABLE I
TESTING FOR STEREOTYPING

	2020 Ohio State sample				Online sample of U.S. adults			
	Population beliefs		Self beliefs		Population beliefs			
	(1)	(2)	(3)	(4)	Noncollege (5)	College (6)	Younger (7)	Older (8)
P(Career | major)	0.39***	0.50***	0.51***	0.72***	0.46***	0.49***	0.44***	0.51***
	(0.09)	(0.17)	(0.08)	(0.06)	(0.13)	(0.11)	(0.11)	(0.13)
$\mathbb{1}$(Most representative)	0.32***	0.28***	0.43***	0.36***	0.23***	0.19***	0.20***	0.22***
	(0.05)	(0.06)	(0.04)	(0.02)	(0.05)	(0.05)	(0.05)	(0.05)
Constant	0.03***	0.02*	0.01	−0.01*	0.03***	0.03***	0.03***	0.02***
	(0.01)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
Career fixed effects	No	Yes	No	Yes	Yes	Yes	Yes	Yes
Majors Included	All	All	Top	Top	All	All	All	All
Observations	33,220	33,220	8,305	8,305	4,664	6,072	4,928	5,808
Individuals	755	755	755	755	212	276	224	264
R^2	0.45	0.46	0.65	0.66	0.48	0.47	0.44	0.50

Notes. The table presents equation (3) using the 2020 OSU sample (columns (1)–(4)) and a sample of U.S. adults recruited online through CloudResearch (columns (5)–(8)). In columns (1) and (2), the dependent variable is respondents' population beliefs about the fraction of graduates with each major working in each career. In columns (3) and (4), the dependent variable is respondents' self beliefs about the probability they would work in each career if they graduated with their top-ranked major. All regressions include all majors that respondents were asked about and, for each of these majors, all 11 careers (where nonemployment is one of the "careers"). "P(Career | major)" is the true fraction of graduates with a major that are working in that career, calculated from the 2017–2019 American Community Survey. $\mathbb{1}$(Most representative) is a dummy variable indicating whether an occupation is the most representative outcome for a major. All regressions cluster standard errors at the individual level and at the career-by-major level. Standard errors are shown in parentheses. Columns (5) and (6) split the CloudResearch sample by whether the respondent has a four-year college degree, and columns (7) and (8) split the same sample by whether the respondent is above median age (39 years). * $p < .10$, ** $p < .05$, *** $p < .01$.

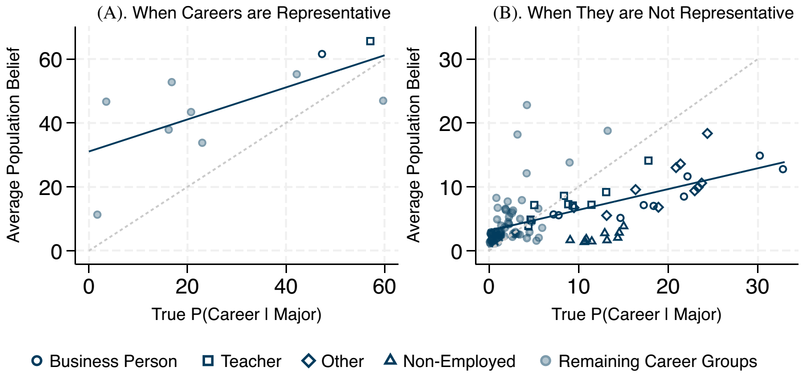

FIGURE II

Beliefs About Careers Conditional on Majors

Each dot represents a career–major pair (where nonemployment is also one of the "careers"). Panel A restricts these pairs to when the career is most representative of the major, and Panel B restricts them to when the career is not representative of the major. The x-axis of both panels is the share of graduates with that major who are working in that career in the ACS. The y-axis is the average population belief, among the 2020 OSU sample, about the fraction of graduates with that major who are working in that career. Lines show OLS regressions including all career–major pairs within each panel.

We can visualize the coefficients in the population-beliefs regressions by looking at Figure II, which plots the average population belief for each career–major pair against its true conditional likelihood. Panel A shows that students overestimate 9 of the 10 representative careers, but within these representative outcomes, the relationship between beliefs and true frequencies is attenuated (i.e., the line of best fit is flatter than 45 degrees). Panel B shows a similar attenuation between beliefs and true frequencies among nonrepresentative outcomes, such that sufficiently rare nonrepresentative outcomes are on average overestimated while sufficiently common ones are neglected. These trends explain why the estimate of γ from equation (3) is less than one. The effect of stereotyping is also obvious from Figure II; comparing objectively similarly frequent careers that are representative (Panel A) versus not (Panel B), we see that representative careers are perceived to be dramatically more common.

Note that students substantially overestimate business and teaching jobs for business and education majors (respectively) but underestimate these outcomes for almost every other major. This result shows that these biases are not just about careers per se (e.g., students simply having not heard of certain professions) but rather about the relationship between careers and majors. It also partly explains why our estimate of θ is so large and significant despite the inclusion of career fixed effects.

Figure II also shows that students greatly underestimate the share of college graduates who are not working for pay. Across majors, students believe 1%–4% of graduates are not employed. Note that these numbers are lower even than the nonemployment rate for male college graduates (6.3%). Furthermore, in our 2021 replication (more details below), we updated the wording of these questions to explicitly ask about "not working for pay (e.g., unemployed or a full-time parent)." We find very similar misperceptions, indicating these results are not driven only by a failure to consider leaving the workforce for reasons related to child care.

These results so far pertain to the career categories we have defined, but one might reasonably wonder whether biases specific to more fine-grained occupations drive our results. For example, if students are unaware of many nonrepresentative jobs in some categories and fail to correct for this when forming their beliefs about majors, this could plausibly lead to biases like those we find. In Online Appendix B.12, we describe a preregistered survey in which we explicitly test for—and fail to find evidence of—this mechanism. In particular, we ask current college students and recent graduates whether they are aware of the 100 most common nonrepresentative occupations. We find that participants are aware of the overwhelming majority of such jobs and that the small amount of unawareness we find is actually negatively correlated with stereotyping, the opposite of what we would expect if it were driving our main results. This speaks against unawareness of specific occupations as a driving force behind stereotyping.

In addition, the first 2021 OSU survey served as a replication exercise for the results from the 2020 sample. For brevity, this survey asked students only about their two top-ranked majors, but otherwise asked the same self and population beliefs about career likelihoods as the 2020 survey did. Online Appendix Table A.VII shows very similar average population beliefs in this sample compared with the 2020 sample. In particular, students again exaggerate the share of graduates working in representative jobs for 9 of the 10 majors (all but nursing).

II.C. Stereotyping in Other Samples

Are these effects limited to our OSU sample? We address this question by collecting data from online samples of U.S. adults and then analyzing nationally representative survey data of U.S. first-year college students.

1. Online Sample of Adults. First, we conducted an online survey of U.S. adults recruited through CloudResearch, where respondents were incentivized to provide accurate responses to the same population-beliefs questions as our OSU survey (see Online Appendix B for details on this survey). Table I, columns (5)–(8) show somewhat smaller but still large and statistically significant stereotyping even among these respondents, both when restricting to college- versus non-college-educated and younger versus older adults. Thus, even people with substantial labor market experience and who have attended college beyond their freshman year exhibit similar patterns of stereotyping as in our main OSU sample.[9]

We test whether biases appear/persist even for college graduates' own major, where we might have expected learning to be greatest. The CloudResearch sample did not include data on participants' own career and major, but in Online Appendix B.12 we describe a survey of recent college graduates (aged 30 or younger) recruited from Prolific in which we do collect such data. There we replicate our main stereotyping result (19.0 percentage points, $p < .01$, Online Appendix Table B.XI, column (1)) and show that stereotyping appears even when focusing on graduates' own major (21.4 percentage points, $p < .01$, column (2)). We also find some imprecise evidence that those who do not have their major's representative job stereotype it somewhat less (-5.2 percentage points, $p = .24$), despite stereotyping other majors more ($+4.4$ percentage points, $p < .05$ for difference between own and other majors). We also find no systematic age gradient: graduates aged 26 and older (the median age among this sample) do not stereotype more or less than younger respondents (Online Appendix Table B.XII). These results provide further evidence that stereotyping remains substantial even after college and even with labor market experience.

9. Note that these results also rule out inattention stemming from lack of incentives being the driving factor behind stereotyping in our main surveys, as the same patterns appear when respondents are paid for accuracy.

2. Nationally Representative Survey of Freshmen. Second, we analyze data from the CIRP Freshman Survey administered by the Higher Education Research Institute (henceforth the Freshman Survey), which surveys incoming first-year students typically during the first weeks of the school year. We pool survey data between 1976 and 2015.[10] We restrict the data to students younger than 24 years with nonmissing location (home ZIP code), race, gender, expected career, and expected major, which leaves 9,068,064 students from 1,587 schools (95.9% of students are at four-year institutions). Online Appendix Table A.I, column (3) shows self-reported demographic information about students in the Freshman Survey. Throughout the analysis, we use census data to weight the Freshman Survey data to match U.S. residents of the same birth cohorts with at least some college education on race, gender, and census division of birth.[11] In that sense, we call this sample nationally representative of incoming college first-years.

We focus on two questions from the Freshman Survey. First, students are asked to mark their "probable field of study" from a list of around 80 options, which vary from year to year, including "Other" and "Undecided." We group these fields into the same 10 groups of majors described above, plus "other" and "undecided," as shown in Online Appendix Table A.VIII. Similarly, students are asked to report their "probable career occupation" from a list of approximately 45 options, which we group into the same nine career groups described above, plus "other," "nonemployment," and "undecided," as shown in Online Appendix Table A.IX.

The dark blue bars in Figure I show the fraction of students in the Freshman Survey who list their expected major's representative career as their probable career occupation. We see a clear pattern, which closely matches our OSU results: students in every major are significantly more likely to expect to work in that major's most representative career than actually do. For example, 65% of prospective art majors expect to be artists (only 17% are), 60% of biology majors expect to be doctors (23% are), 42%

10. We use data from all years in this range except 1977 and 1978. We choose these years because they include information on students' home ZIP code, which we use for weighting. See https://heri.ucla.edu/instruments/ for a list of participating schools and survey instruments by year.

11. For people born outside the United States, we use current location as a proxy for birthplace. We include students in the Freshman Survey data that are noncitizens so long as they self-report a U.S. ZIP code.

of communications/journalism majors expect to be writers or journalists (4% are), 62% of psychology majors expect to be counselors (21% are). All of these differences are statistically significant at the $p < .001$ level.[12] In Online Appendix B.7, we bring in data from the CIRP Senior Survey, which allows us to link a subset of our Freshman Survey data to the same students' major and career expectations when they are seniors. This yields 258,134 senior students spanning 1994 to 2015. There we find that although there is substantial heterogeneity across majors, overall rates of expecting representative jobs are largely unchanged from freshman to senior year, falling by only 1.1 percentage points.

Next we examine how the gaps between expected and actual careers have changed over time. Because the U.S. Census data that include respondents' college major only began in 2012, we instead compare students' expected careers (unconditional on their major) to the actual distribution of occupations from the Annual Social and Economic Supplement to the Current Population Survey (CPS) (Flood et al. 2021). We restrict the data to those aged 33–37 because by this time the vast majority of people are no longer students and have started their career, though our results are not sensitive to this specific age range. This also matches well with the age of 35 that we ask about in the 2021 OSU survey, described further below. We match occupation codes from the CPS to the same nine occupation groups (see Online Appendix Table A.III).

12. None of these results are driven by ambiguity in the mapping between occupations and career groups. See Online Appendix B.5, where we explore reclassifying occupations across career groups. Note that these survey answers are qualitative: respondents simply answer "yes" or "no" to whether they are likely to have a given career or major. A reasonable worry is that these answers do not reflect mistaken beliefs but a mismatch between qualitatively elicited survey answers and probabilistic "ground-truth" data. For example, if all students knew they had a 30% chance of working in business, then none might report it as their "probable" career despite 30% being found with that career after graduation. However, the close correspondence between these qualitatively elicited beliefs and our quantitative OSU beliefs suggests this is not the case. More specifically, we can ask whether the differences between actual and expected careers in the qualitative Freshman Survey (shown in Figure III) are consistent with being driven by biases in population beliefs of the same magnitude as in our OSU sample (see Online Appendix B.4 for more details on this exercise). We find that the implied error based on stereotyping alone from the OSU survey strongly correlates with the actual error in the Freshman Survey (0.81 with a p-value $< .01$). We conclude from this exercise that the pattern of overestimation of careers in the Freshman Survey is quite close to what we would expect if it were driven by the same errors as in the OSU students' beliefs.

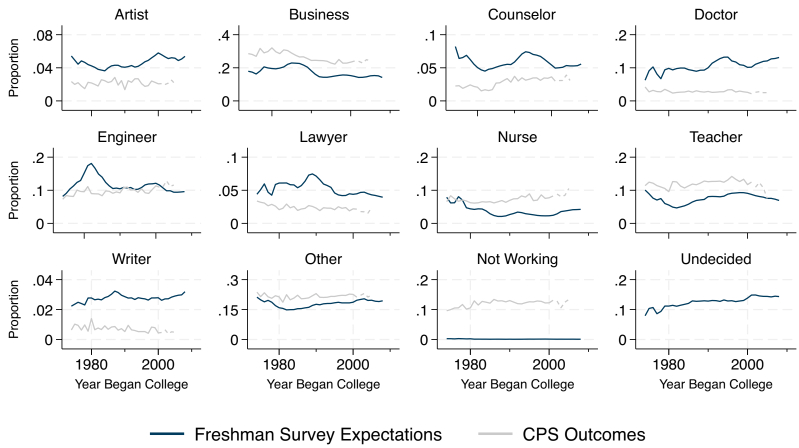

FIGURE III

Career Expectations Versus Outcomes over Time

The figure compares the share of first-year undergraduates by the year they began college (birth year plus 18) who expect to have a career in each occupation (dark blue line) according to the Freshman Survey data along with the share of college graduates (light gray line) aged 33–37 who work in that occupation in the same cohort, according to the Current Population Survey. The light gray line becomes dashed when CPS outcomes begin to only include graduates younger than 37.

Figure III shows large, systematic, and persistent differences between the careers that freshmen expect to attain and the actual occupations they go on to have. The dark blue lines show the share of first-year students each year who expect to have each career. The light gray lines show the share of college graduates in the same cohort that are working in that occupation in the CPS. Online Appendix Table A.X, Panel A shows the corresponding share expecting and actually working in each career, pooling across cohorts.[13] We see that twice as many students expect to become artists, counselors, and lawyers (about 5% each) than actually do (2%–3% each). Four times as many students expect to become writers and doctors (2.7% and 11.1%)

13. If students' beliefs were on average correct, we might expect the outcomes data to lag the expectations data. For simplicity, we do not attempt to correct for this, and the two series are plotted contemporaneously in Figure III. Note that there is no lag in outcomes for which the patterns would match.

than do (0.7% and 2.8%).[14] Note that these gaps between expected and actual careers are exactly what we would expect from students stereotyping majors based on their representative careers. The occupations with the largest positive gaps—for example, doctors, writers, lawyers, artists—are those that are primarily the (rare) representative careers of particular majors. In contrast, those with negative gaps (teaching, business, nonemployment) tend to be the most common alternatives to other majors' representative jobs. Further, these patterns appear quite stable over time, suggesting that the apparent stereotyping in Figure I reflects a long-standing pattern in students' expectations.

II.D. Evidence on Stereotyping from Implicit Associations

Our preferred explanation for the patterns of belief biases documented here is that representative careers come to mind easily because they are strongly associated with their corresponding majors (Gennaioli and Shleifer 2010; Bordalo et al. 2016). This hypothesis suggests that alternative measures of associations should also be influenced by representative careers and should be tightly linked to the belief biases we document. To provide evidence in this direction, we adapt the IAT, an extremely prevalent measure from the psychology literature studying stereotyping (Greenwald, McGhee, and Schwartz 1998), to our setting.

In two preregistered experiments, we design and conduct IATs to measure whether people implicitly associate majors with their representative careers and whether these associations predict belief biases. Experiment 1 asks whether participants exhibit especially strong implicit associations between majors and their representative careers. Experiment 2 then asks whether implicit associations between majors and representative careers correlate with belief biases at the individual level. See Online

14. Online Appendix Table A.X, Panel B shows that there do not appear to be similarly large differences between the fraction of students who expect to pursue each major and the fraction of students who actually attain such majors, calculated from the 2017–2019 American Community Survey (ACS) (Online Appendix Table A.II shows how we categorize majors from the ACS into our 10 major groups). Thus, differences between expected and actual careers are unlikely to be driven by systematic biases in the majors with which students expect to graduate.

Appendix B.11 for details on this data collection, including incentives, full experimental design, and preregistration.

1. IAT Implementation. Experiments 1 and 2 included IATs following a similar design modeled closely off the design from Project Implicit[15]. A typical IAT requires participants to sort words into two pairs of mutually exclusive categories. For example, an IAT might ask participants to sort distinctively Black versus distinctively white names (e.g., “Latonya” versus “Heather”) as well as to sort pleasant versus unpleasant words (“lucky” versus “poison”). Different rounds of the test differ according to which categories share a response key. For example, in one round, both distinctively Black names and pleasant words need to be sorted to the left (by pressing the Q key) while distinctively white names and unpleasant words need to be sorted to the right (by pressing the P key). In a second round, these pairings are reversed. A participant reveals an implicit association between Black names and unpleasant words if she sorts words faster when these two categories are paired together than when Black names are paired with pleasant words.

Figure IV shows a stylized illustration of the design of our IATs measuring associations between careers and majors. Participants underwent multiple rounds, each of which had a focal major group (e.g., biology and chemistry) and a focal career group (e.g., doctors). They were then presented one by one with individual occupations or majors and had to decide whether they belonged to these focal groups or instead belonged to “other majors” or “other careers” (i.e., majors/careers that did not fall into any of the 10 major groups/9 career groups).[16] Participants always sorted words to the left or right using the Q or P keys. Whenever participants successfully sorted a word in the correct direction, a green checkmark appeared and the screen progressed to the next word. Whenever they incorrectly sorted a word, a red X appeared, and participants had to wait one second before trying again. Thus, though all participants had to correctly sort all words, incorrect responses slowed them down. Partici-

15. https://implicit.harvard.edu/.

16. The individual majors included all those where at least 0.5% of graduates with a degree in that group had that major. Similarly, the individual occupations were all those for which at least 0.5% of college graduates within that career group had that occupation.

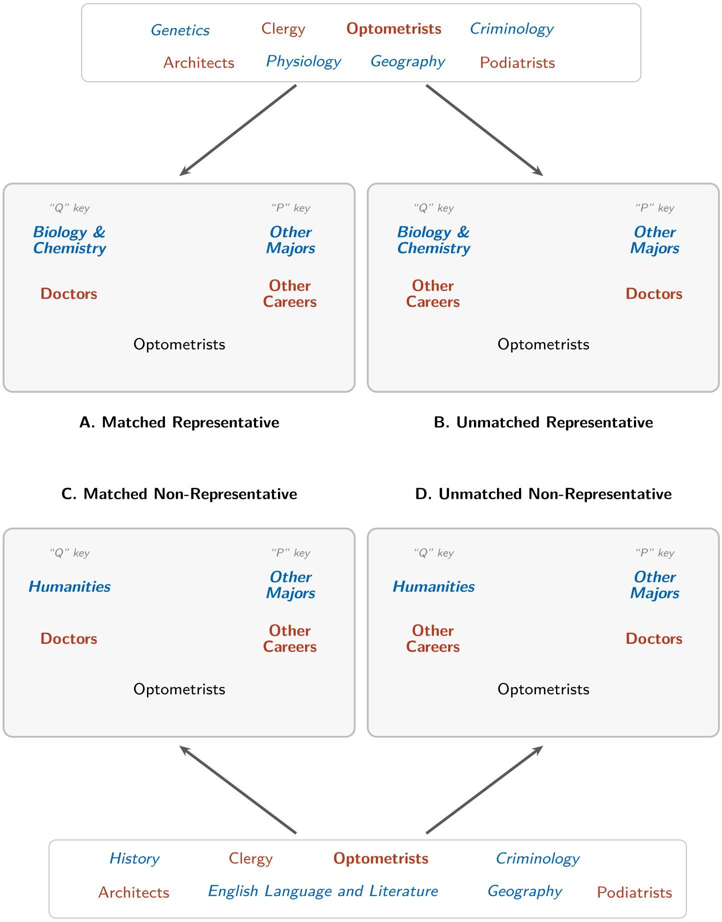

FIGURE IV

Implicit Association Test Design

This figure illustrates the design of our implicit association test. In each round, participants sort individual majors and careers into category groups using the Q (left) and P (right) keys. Panels A and B show a round with a representative career-major pairing (biology and chemistry with doctors). In the matched condition (Panel A), the focal major and career share a response key; in the unmatched condition (Panel B), they do not. Panels C and D show a round with a nonrepresentative pairing (humanities with doctors). The difference in response times between matched and unmatched conditions measures the strength of implicit associations. Our main test asks whether this difference is larger for representative pairings (A versus B) than nonrepresentative pairings (C versus D). For illustration, majors are shown in blue italics and careers in red regular font; in the actual experiment, all category labels were the same color. Task images are simplified in this figure; See Online Appendix B.11 for links to screenshots of and the actual survey used for these experiments.

pants were incentivized to complete all rounds as quickly as possible, as the fastest 10% of participants earned an additional $5 bonus.

Each round included two crucial sorting tasks, which occurred in a random order.[17] These tasks differed in whether the focal career and major groups either shared a response key (the matched task) or not (the unmatched task). The left and right panels of Figure IV show stylized examples of matched and unmatched tasks, respectively. The main outcome variable we construct is $\text{Difference}_{i,c,M}$, the difference between the time taken in the unmatched compared with the matched task.[18] Participant i reveals an implicit association between c and M to the extent that $\text{Difference}_{i,c,M}$ is positive.

For example, suppose participant i has a strong association between the major group biology/chemistry and doctors (its representative career). She might therefore find it very easy to quickly categorize the word "optometrist" as belonging to "biology/chemistry or doctor," but she may find it more difficult to categorize "optometrist" when choosing between "biology/chemistry or other careers" and "other majors or doctors." Figure IV, Panels A and B show such a case. In contrast, Panels C and D show a round where doctor is paired with humanities, a major for which it is not the representative career. Here, intuitively, the student may find it no more difficult to categorize "optometrists" in the unmatched task compared to the matched task. This difference-in-difference, comparing $\text{Difference}_{i,c,M}$ across rounds with representative versus nonrepresentative focus careers, allows us to test whether implicit associations are linked to representativeness.

2. Implicit Associations Between Majors and Representative Careers. Experiment 1 tests whether participants implicitly associate majors with their representative career. To this end, participants completed four rounds of the IAT. Each round included one of our 10 major groups as the focal major, and the focal career was chosen from that major's representative career group plus "teach-

17. Each round also began with three practice sorting tasks to familiarize participants with the careers and majors in each group (see Online Appendix B.11 for full details).

18. To avoid outliers driving our results, we winsorize this variable at the 95th percentile, and then we standardize it to have mean zero and standard deviation one.

ers" and "businesspeople" (unless this group is already included as the representative career for that major). We chose these focal career–major pairs to ensure that we have measures of associations between each major and its most common careers, allowing us to measure whether participants have stronger associations between majors and their representative careers than other common careers. This process results in 28 pairs of focal careers–majors, and each participant underwent rounds for a randomly chosen four pairs. We recruited 100 participants on Prolific to complete Experiment 1.

Table II, Panel A shows OLS regressions where the dependent variable is $\text{Difference}_{i,c,M}$, our measure of i's implicit association between the focal career c and major M. Column (1) regresses this variable on an indicator for whether the round's focal career is the representative career of its focal major. We see that for rounds with representative careers, $\text{Difference}_{i,c,M}$ is 0.36 standard deviations (6.5 seconds) larger than for nonrepresentative careers ($p < .01$). Column (2) adds individual fixed effects, which do not qualitatively change the main result (0.30 standard deviations, or 5.4 seconds, $p < .01$). Columns (3) and (4) add to this regression the true population frequency of the focal career among graduates with the focal major. These frequencies also predict implicit associations between the focal career and major, but the magnitude of this difference is much smaller than for representativeness: the coefficient on the representativeness dummy is 45–55 times as large as that on true frequency, indicating that the boost in implicit associations from representativeness is equivalent to a 45–55 percentage point boost in true frequency. Finally, columns (5) and (6) add as controls the (winsorized) time participants took in the earlier practice sorting tasks in each round, which have no effect on our main conclusions.[19]

We conclude from these results that participants appear to implicitly associate majors with their representative career, and

19. Note that one might have worried that the names of some majors partially overlap with their representative career (e.g., "psychology" and "psychologist"). This worry prompted us to include a sorting task (part 3, see Online Appendix B.11 for details) in which participants must sort just between the focal major and focal career. If it is simply difficult to tell apart members of the focal major and career, and this were driving any of our results, we should expect inclusion of time to complete part 3 to significantly predict $\text{Difference}_{i,c,M}$ and reduce the coefficient on representativeness. In fact, neither of these turn out to be true, suggesting that this is not an important confound.

TABLE II
IMPLICIT ASSOCIATIONS

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Dependent variable = Difference$_{i,c,M}$						
$\mathbb{1}$(Most representative)	0.36***	0.30***	0.28***	0.24***	0.30***	0.23***
	(0.09)	(0.07)	(0.08)	(0.06)	(0.09)	(0.08)
P(Career | Major)			0.01**	0.00*	0.01**	0.00
			(0.00)	(0.00)	(0.00)	(0.00)
Part 1 time					0.01	0.06
					(0.05)	(0.12)
Part 2 time					0.08*	0.12
					(0.04)	(0.08)
Part 3 time					−0.05	0.06
					(0.07)	(0.11)
Constant	−0.12*	−0.10***	−0.20**	−0.16***	−0.22**	−0.13**
	(0.06)	(0.02)	(0.08)	(0.04)	(0.10)	(0.06)
Observations	400	400	400	400	400	400
Individuals	100	100	100	100	100	100
Individual fixed effects	No	Yes	No	Yes	No	Yes

TABLE II
CONTINUED

	(1)	(2)	(3)	(4)	(5)	(6)
Panel B: Dependent variable = population belief of P(career | major)						
$\mathbb{1}$(Most representative) × $\text{Difference}_{i,M}$	4.34***		3.04***	2.99***		0.92**
	(0.78)		(0.74)	(0.75)		(0.40)
$\mathbb{1}$(Most representative) × LoM $\text{Difference}_{i,M}$		3.27***	1.71		3.25***	2.65**
		(0.94)	(1.11)		(0.95)	(1.02)
$\mathbb{1}$(Most representative)	16.44***	15.74***	15.93***			
	(3.31)	(4.13)	(3.28)			
P(Career | major)	0.35***	0.38***	0.35***			
	(0.07)	(0.09)	(0.07)			
Observations	21,780	43,560	21,780	21,780	43,560	21,780
Individuals	396	396	396	396	396	396
Individual-by-career fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual-by-major fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Major-by-career fixed effects	No	No	No	Yes	Yes	Yes

Notes. Panel A presents OLS estimates where the dependent variable is the difference between the time participants took to complete parts 5 and 4 (unmatched pairing minus matched pairing) of the implicit association test in Experiment 1. We winsorize this variable at the 95th percentile and then normalize it to have mean zero and a standard deviation of one. "P(Career | major)" is the true percent of graduates with the IAT round's focal major that are working in its focal career, calculated from the 2017–2019 American Community Survey. $\mathbb{1}$(Most representative) is a dummy variable indicating whether the focal career is the most representative outcome for the focal major. The time control variables are the times that participants took to complete parts 1–3 of the same IAT round (winsorized at the 5th and 95th percentiles and then normalized). Panel B presents OLS estimates where the dependent variable is Experiment 2 participants' population beliefs about the fraction of graduates with a certain major with a certain career. $\text{Difference}_{i,M}$ is (winsorized and normalized) difference between the time participants took to complete parts 5 and 4 (unmatched pairing minus matched pairing) of the IAT for that major and its representative career. The leave-out-mean (LoM) difference is calculated by taking the average of $\text{Difference}_{i,M'}$ for all other majors M' that participant took the IAT for. All regressions cluster standard errors at the individual level and at the career-by-major level. Standard errors are shown in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

that the size of these associations is much greater than the true conditional frequency of those careers alone would predict. These results appear to mirror our main results concerning students'/participants' beliefs about the frequency of careers. We turn to Experiment 2 to explore this connection more directly.

3. Implicit Associations Predict Belief Biases. Experiment 2 tests whether heterogeneity across participants in associations between majors and their representative career predicts stereotyping in their beliefs (i.e., overestimation of the representative career's conditional frequency). The IAT portion of this experiment was identical to Experiment 1 except that participants did five rounds (instead of four) and the focal career group was always the representative career of the randomly chosen focal major. In addition, participants in Experiment 2 answered population beliefs questions about all 10 major groups. The IAT and belief modules of the survey occurred in a random order. We recruited 396 participants on Prolific to complete Experiment 2, which we conducted separately from Experiment 1 to keep the length of each survey manageable.

Table II, Panel B shows OLS regressions asking whether participants' population beliefs are predicted by their implicit associations between majors and their representative career, as measured by our IAT. In column (1), we regress population beliefs about $P(c|M)$ on the true conditional frequency, an indicator for whether c is M's most representative career, and an interaction between this indicator and participants' (winsorized and normalized) implicit association between M and the representative career. The coefficient on this interaction is positive and statistically significant ($p < .01$), indicating that participants who reveal a stronger implicit association between a major and its representative career have a more exaggerated belief about that career's conditional frequency. The coefficient of 4.34 can be interpreted as saying that every standard deviation increase in implicit associations predicts an increase of 4.34 percentage points in participants' stereotyped beliefs.

A natural question is whether these results are driven by heterogeneity across majors—some majors being stereotyped to a greater degree and participants holding stronger implicit associations between those majors/careers—or heterogeneity across people. To investigate this, column (4) adds career-by-major fixed effects to these regressions, meaning that all variation in implicit associations and beliefs stems from differences across participants within career–major pairs. We see that the correlation between implicit associations and stereotyping is still large and statistically significant even controlling for these fixed effects: a one standard deviation increase in associations predicts an increase of 2.99 percentage points in beliefs ($p < .01$).

A second natural question is whether the results presented so far reflect differences across participants in an overall tendency to stereotype majors—and to implicitly associate them with their representative career—or instead reflect within-participant differences across majors. To answer this question, we look at whether stereotyping of a major M is predicted by participants' average implicit association between other majors and their representative career. Column (2) shows that these leave-out-mean (LoM) implicit associations are strongly correlated with participants' beliefs ($p < .01$). This remains true controlling for career-by-major fixed effects (column (5), $p < .01$). Finally, columns (3) and (6) regress beliefs about $\mathrm{P}(c|M)$ on both participants' implicit association between M and its representative career as well as the LoM constructed from each participant's associations for other majors. We see that when including major-by-career fixed effects, each variable continues to predict stereotyping conditional on the other ($p < .05$ for both). We conclude from this analysis that the correlation across participants between associations and stereotypes is partly driven by there being types who stereotype more and have stronger associations between majors and their representative careers.

4. Explicit Associations and Stereotypical Descriptions. The above results show that participants appear to have implicit associations between majors and their representative careers and that these associations in part appear to drive stereotyped beliefs. Note that such implicit measures need not reflect unconscious attitudes or views (Gawronski, Hofmann, and Wilbur 2006), in the sense of people believing they do not hold such associations. Indeed, we show in Online Appendix B.12 that directly asking participants to select which careers they associate with each major, and to describe in their own words their stereotypical image of graduates with each major, yields very similar findings as the IAT. In particular, participants tend to explicitly associate majors with their representative careers, and their stereotypical description of someone with a major tends to include that career. Furthermore, both of these measures are correlated at the individual level with beliefs even controlling for major-by-career fixed effects. Thus the implicit associations we uncover go hand in hand with more direct measures of participants' stereotypes.

5. Heterogeneity by Students' Experiences. Implicit associations like those described above are thought to be tightly linked to memory (e.g., Greenwald and Banaji 1995; Phelps et al. 2000), and recent work in economics has begun to formalize the connections between associations, memory, and belief biases (e.g., Bordalo et al. 2023; Enke, Schwerter, and Zimmermann 2024; Conlon and Kwon 2025). A basic implication of these ideas is that heterogeneity across students in their experiences should have outsized effects on their beliefs. In Online Appendix B.10, we show that beliefs in the Freshman Survey and in our OSU data about the likelihood of careers—self beliefs but also (more tellingly) population beliefs—are strongly correlated with the careers and majors of students' parents and other role models. Furthermore, these results appear inconsistent with some natural rational-learning stories: for example, students with a role model in a major's representative career—who we might think should have better information about it—stereotype it more (6.13 percentage points, $p < .01$). These results instead look consistent with a simple memory framework in which beliefs depend on who comes readily to mind. See Online Appendix B.10 for full details.

III. Implications of Stereotyping

What are the implications of stereotyping for labor market outcomes and for welfare? We ask how students who form beliefs according to the model of belief formation described in Section II.A choose their majors. We test the model's predictions to explore how stereotypes about majors may distort education choices and labor market outcomes.

III.A. Setup

For simplicity, assume there are only two majors $M \in \{A, B\}$ and two careers $c \in \{a, b\}$. We assume that students can only enter job a if they major in A, but students of either major can choose to enter job b. Intuitively, job a requires special training—for example, being a graphic designer might require having majored in art (A)—while the other job b (e.g., working in business) is open to people from multiple educational backgrounds but may benefit from having majored in an aligned field.

The model has two periods. In the first period, student j irreversibly chooses which field to major in. In the second period—corresponding to after the student has graduated with their chosen major—they choose which occupation to enter. The student's utility function includes three additive terms, only some of which are observed in the first period. First, the student cares about the wages $w_{c,M}$ that they would earn with career c and major M. We assume that wages are observable in both periods, and that $w_{b,B} > w_{b,A}$: that is, majoring in B boosts wages in its associated job b.

Next, the student has two career-specific utility terms, ψ_j and ϵ_j, that they derive if they choose to work in career a. The first, ψ_j, is observable in the first period, when they choose their major. The second, ϵ_j, is only observable after graduation. The student must therefore choose a major in the first period given their beliefs about the distribution of ϵ_j. We primarily refer to ψ_j and ϵ_j as relating to the person-specific "amenities" that career a delivers, though other interpretations deliver identical results. For example, either/both ψ_j and ϵ_j could also be interpreted as capturing the students' idiosyncratic productivity (and therefore individual-specific wages) in job a or as capturing the difficulty of finding a job in a (conditional on its wage). The crucial distinction between these two objects is only that ψ_j captures what the student knows in advance about their fit with job a, whereas ϵ_j captures aspects of the job that they learn only after graduating.

To summarize, student j's utility ex post (after graduating) from choosing career c is:

$$u_{j,c,M} = \begin{cases} w_{a,M} + \psi_j + \epsilon_j & \text{if } c = a \\ w_{b,M} & \text{if } c = b. \end{cases}$$

For simplicity, we assume that $\psi_j \sim U(-\sigma_\psi + \mu_\psi, \sigma_\psi + \mu_\psi)$, $\epsilon_j \sim U(-\sigma_\epsilon + \mu_\epsilon, \sigma_\epsilon + \mu_\epsilon)$ and that ψ_j and ϵ_j are independent of each other. In the main text, we assume that wages are fixed, but in Online Appendix C, where we prove the results presented in this section, we show that our main conclusions are robust to allowing for endogenous wages (i.e., diminishing marginal productivity—and therefore downward-sloping demand curves—within industries).

1. Discussion of Assumptions. Besides introducing necessary simplifications, the model incorporates two crucial assumptions. The first is that college majors matter for productivity and therefore wages, which is consistent with many studies finding that college majors have large effects on earnings (e.g., Hastings, Neilson, and Zimmerman 2013; Kirkeboen, Leuven, and Mogstad 2016; Bleemer and Mehta 2022).

The second crucial assumption is that students care about their future occupation—over and above the average income it provides—when choosing their major. Besides being intuitive, several pieces of suggestive evidence corroborate this assumption. First, in the second 2021 survey, we asked the control group of the experiment (described below) to rate on a scale from 0 to 100 how important various factors were to them as they decided what to major in. The "effect on what job I will get after college" received the highest average ranking (83 out of 100), above enjoyability of classes (70), difficulty of classes (58), and the opinions/choices of family, friends, and peers (all below 30).

Second, in Online Appendix B.8, we use our 2020 survey data to structurally estimate students' preferences when choosing their major. There, we find that students care both about the expected salary their major will earn them but also have strong (indeed, much stronger) nonpecuniary preferences over the occupations they might pursue depending on their choice of major.

Finally, Online Appendix Table A.XI shows that in two data sets—the National Survey of College Graduates (NSCG) and Survey of Consumer Expectations (SCE)—having a job that does not match one's field of study or skills/experience is strongly positively correlated with college graduates' job dissatisfaction ($p < .01$ in both samples). This kind of "mismatch" is also negatively correlated with income, but the correlation between it and dissatisfaction appears even controlling for income ($p < .01$ in both samples). Note that all these regressions control for major fixed effects, so these correlations reflect differences in major across respondents who do versus do not experience mismatch. These correlational results are consistent with both pecuniary and nonpecuniary career concerns being important to graduates' welfare.[20]

20. In addition, the NSCG data ask respondents who report their job is unrelated to their field of study why they do not: 72.5% report a factor related to pay, amenities, location, or the availability of such a job as being the most important reason, while only 18.5% report a "change in career or professional interest." This

III.B. Stereotyping

Following equation (2), we assume students stereotype majors: they exaggerate the share of A majors that end up working in career a: $\pi_{a|A} = p_{a|A} + \theta(1 - p_{a|A})$.[21] How should this bias in students' beliefs about others' occupations affect which major they themselves choose? Answering this question requires us to consider how students' occupation-major beliefs relate to their expectations about the rest of the economy. To pin down these additional beliefs, we assume that students believe the economy is in equilibrium: they understand that graduates are choosing the career and major that maximize their subjective expected utility, given their beliefs. Students' beliefs about the distribution of ψ and ϵ, the career-specific utility terms, must therefore rationalize their stereotyped beliefs. In other words, we assume that students hold an internally consistent view of the world.

To formalize this, let $p_{a|A}(\mu_\psi, \mu_\epsilon)$ denote the share of A graduates who would choose career a if the means of ψ and ϵ were μ_ψ and μ_ϵ respectively. Let $\widehat{\mu}_\psi$ and $\widehat{\mu}_\epsilon$ be the student's beliefs about these means. The equilibrium condition describing the relationship between stereotyped beliefs about occupations and biases about amenities is given by:

$$p_{a|A}(\widehat{\mu}_\psi, \widehat{\mu}_\epsilon) = \pi_{a|A}. \tag{4}$$

Equation (4) formalizes our assumption that students hold a consistent internal model of the economy. Stereotyped beliefs about majors must be reflected in biased beliefs about the factors that lead people to choose majors and careers (i.e., about μ_ψ and μ_ϵ). Proposition 1 describes implications of this consistency requirement:

PROPOSITION 1. Suppose students stereotype majors: $\theta > 0$. Then, under the internal consistency condition of equation (4),

(i) Students cannot have accurate beliefs about the distributions of both amenities: $(\widehat{\mu}_\psi, \widehat{\mu}_\epsilon) \neq (\mu_\psi, \mu_\epsilon)$.
(ii) If students' self beliefs do not differ from their population beliefs, then their perceived distribution of ψ is accurate,

suggests that the major driver of such mismatch is related jobs being worse than anticipated, rather than changing preferences (where welfare analyses would be less straightforward).

21. By assumption, $p_{b|B} = 1$, so equation (2) implies no bias in beliefs about major B.

while their beliefs about ϵ are overly optimistic. That is, let $\pi^{\text{self}}_{a|A}$ be the average self belief among A graduates: that is, their ex ante perceived probability of choosing career a after graduating with major A. Then $\pi^{\text{self}}_{a|A} = \pi_{a|A}$ implies $\widehat{\mu}_{\psi} = \mu_{\psi}$ and $\widehat{\mu}_{\epsilon} - \mu_{\epsilon} = 2\sigma_{\epsilon}\theta(1 - p_{a|A})$.

Part (i) of Proposition 1 says that stereotyping implies distorted beliefs about at least one underlying distribution of the career-specific amenities. Part (ii) says that comparing students' self and population beliefs allows us to pin down which belief is biased: if self beliefs are the same as population beliefs—which we saw in Figure I is approximately the case—then students' stereotyped beliefs about occupations must lead them to be overly optimistic about ϵ, the unobserved amenities associated with a (e.g., it being easier to find a job in that sector than it truly is).

To see the intuition for this result, suppose students have upwardly biased beliefs about the distribution of ψ: i.e., they think other students are more interested ex ante in career a than they are. This could rationalize their stereotyped belief about the share of A majors who choose career a. But students observe their own value of ψ, so this overestimation about others' choices would not have any implications for their own career or major choice. Therefore, biases in beliefs about ψ imply a divergence in students' population and self beliefs: they should think that the share of A majors entering career a is larger than their own probability of choosing that job. In Section II, we saw instead that, if anything, the reverse is true. When interpreted through the lens of this model, this suggests that students are instead biased about μ_{ϵ}: they think that career a turns out to have better ex post amenities (e.g., it being easier to find a job) than is in fact the case. Note that this interpretation is also consistent with our finding (described in Section IV) that changing participants' population beliefs also has significant effects on their self beliefs. This implies that students are indeed using their beliefs about population career outcomes to infer some unobserved common source of utility from different careers.

For the remainder of this section, we assume that $\widehat{\mu}_{\epsilon} = \mu_{\epsilon} + 2\sigma_{\epsilon}\theta(1 - p_{a|A})$. That is, stereotyping of major A is accompanied by proportional overoptimism about some unobserved attribute of career a.

III.C. Misallocation

Proposition 2 describes how stereotyping leads to misallocation:

PROPOSITION 2. Assume, following Proposition 1.ii, that $\widehat{\mu}_\epsilon = \mu_\epsilon + 2\sigma_\epsilon\theta(1 - p_{a|A})$. Then stereotyping increases misallocation:

$$\frac{dp_{b|A}}{d\theta} = \frac{1 - p_{a|A}}{2 - \theta} > 0. \tag{5}$$

The misallocation result in Proposition 2 follows from the fact that stereotyping leads students to have an overly optimistic view of the unobserved amenities (e.g., ease of finding a job) of career a. This draws more students into major A, but the marginal student induced to major in A is negatively selected on ψ and therefore is less likely to actually end up choosing career a than inframarginal A majors are. Misallocation, which we define as the share of A majors who end up working in b (and who therefore would have benefited from instead majoring in B), therefore increases with stereotyping.

To test Proposition 2, we use multiple data sources to construct six different proxies for misallocation at the major level. First, we use the share of graduates not working in their major's representative career in the ACS. Second, we use average responses in the 2013 NSCG to the question "To what extent was your work on your principal job related to your highest degree?," where answers are on a scale from 1 (meaning closely related) to 3 (meaning not related). Third, we use average responses in the New York Fed's SCE to the question "On a scale from 1 to 7, how well do you think this job fits your experience and skills?" (we reverse the scale so higher numbers imply a poorer fit).

Our next three proxies reflect the idea that misallocation implies a form of regret or dissatisfaction with one's career path. The fourth proxy is average answers in the 2013 NSCG to the question "How would you rate your overall satisfaction with [your] principal job" on a scale from 1 (very satisfied) to 4 (very dissatisfied). Fifth, we use average answers in the SCE to the question "Taking everything into consideration, how satisfied would you say you are, overall, in your [current/main] job?" on a scale from 1, very dissatisfied, to 5, very satisfied (we reverse the scale so that higher numbers indicate more dissatisfaction). Finally, we use the share of respondents in the 2021 Survey of Household Economics and Decision Making who reported that they would have chosen a dif-

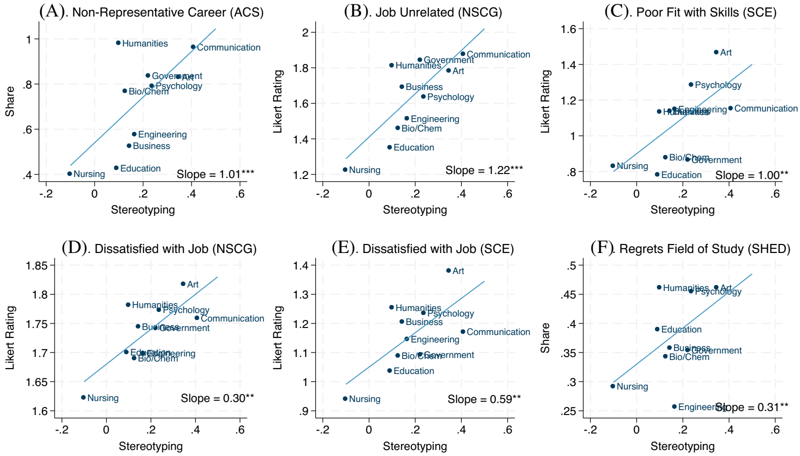

FIGURE V

Stereotyping Predicts Misallocation

The x-axis in each panel shows average stereotyping of each major (population beliefs minus truth from the 2020 OSU data). The y-axes in each panel are the share of ACS respondents who are not working in that major's representative career (Panel A); the average rating (on a scale from 1 to 3) of how unrelated employed NSCG respondents' job is to their highest degree (Panel B); the average rating (on a scale from 1 to 7) of how well employed SCE respondents say their job fits their skills and experience, reverse coded so high ratings indicate a poor fit (Panel C); the average rating (on a scale from 1 to 4) employed NSCG respondents give to how satisfied they are with their job, coded so that high ratings indicate dissatisfaction (Panel D); the average rating (on a scale from 1 to 5) employed SCE respondents give to how satisfied they are overall with their job, coded so that high ratings indicate dissatisfaction (Panel E); and the share of SHED respondents who report that they would have chosen a different field of study if they could go back and make their education decisions again. The slope reported is the coefficient on stereotyping in a univariate OLS regression. * $p < .10$, ** $p < .05$, *** $p < .01$.

ferent field of study if they "could go back and make your education decisions again."

Figure V shows that all six proxies for misallocation are positively and significantly ($p < .05$) correlated at the major level with average stereotyping in our 2020 OSU data (i.e., average population beliefs about representative careers minus the truth).[22] In majors where there is greater stereotyping, people are more likely to work in nonrepresentative careers (Panel A), report their jobs

22. Note that each data set has somewhat different categories of majors, and sometimes the mapping between our major groups in the ACS to those in the other data sets is not always perfect or one for one.

are more unrelated to their degree (Panel B), report their jobs are a worse fit with their skills and experience (Panel C), report being more dissatisfied with their jobs (Panels D and E), and report wishing they could go back and choose a different field of study (Panel F). All of these correlations are significant at the $p < .05$ level. We conclude that these results, though of course only correlational, appear to corroborate the intuitive prediction from the model that stereotyping produces misallocation (and also thereby dissatisfaction and regret about education choices) by pushing people to choose majors whose associated jobs they end up not pursuing.

III.D. Welfare

Finally, Proposition 3 describes how overall ex ante welfare changes as stereotyping becomes more severe.

PROPOSITION 3. Stereotyping reduces average welfare:

$$\frac{dW}{d\theta} = -\frac{dp_A}{d\theta}\,\theta(1 - p_{a|A})\Big[2\sqrt{\sigma_\epsilon(w_{b,B} - w_{b,A})} - \sigma_\epsilon\theta(1 - p_{a|A})\Big] < 0,$$

$$\text{where } \frac{dp_A}{d\theta} > 0. \tag{6}$$

The marginal welfare loss is increasing in $w_{b,B} - w_{b,A}$: errors in major choice are most costly when having the aligned major matters most for career-b earnings.[23] The marginal welfare loss is also increasing in the magnitude of stereotyping θ.

Proposition 3 suggests that the welfare losses from stereotyping—and therefore the majors for which combating its negative effects may be more urgent—are those where belief biases (θ) are larger and ones where the wage loss from misallocation ($w_{b,B} - w_{b,A}$) is larger. Although we can compute average stereotyping across majors from our survey data, note that the wage loss term is counterfactual (how much less is student j earning in their job than she would have if she majored in something else) and therefore unobservable. We can nonetheless construct a plausible proxy for this term by looking at wages

23. This statement holds for not-too-large stereotyping (θ) and wage differences ($w_{b,B} - w_{b,A}$). With endogenous wages ($\gamma > 0$), the welfare derivative acquires an additional wage-externality term that can partially offset the misallocation loss; see Online Appendix C.

among graduates not working in their major's representative career.

Online Appendix Figure A.I shows average stereotyping in our OSU sample along the x-axis for our 10 major groups. The y-axis shows average earnings among those without their major's most representative career (net of career fixed effects to account for the fact that alternative careers differ across majors).[24] Proposition 3 suggests that majors to the bottom right of the figure, with more severe belief biases and less lucrative alternative careers, are ones where stereotyping may be especially harmful on the margin. As is perhaps intuitive, the majors with the largest marginal welfare costs are fields such as communication, art, and psychology, each of which (as we've seen) have very large average biases and also tend to result in lower-paying nonrepresentative jobs. In contrast, majors like nursing and engineering have lower implied costs from stereotyping, either because the average belief bias is small or because graduates with nonrepresentative careers nonetheless earn high salaries.

Note that all of these welfare analyses concern the effect on the marginal student: that is, one who perceives themselves to be indifferent between the major they are stereotyping and a less risky alternative major. This analysis therefore suggests that there may be particularly large benefits from nudging such students toward less risky majors.

1. Discussion. Of course, the model necessarily makes several simplifying assumptions. For example, we assume that the alternative major students might pursue automatically leads to its associated job. If instead a student would choose an equally risky other major—one for which she would also end up in a mismatched job—the misallocation and welfare effects of stereotyping could well be different. As shown in Figure V, the fact that stereotyping correlates with multiple proxies for misallocation and welfare provides suggestive evidence that the main forces our model captures may be at play. But caution—and more research—is warranted.

24. More precisely, we estimate by OLS a regression of earnings among those without their major's most representative career on fixed effects for major group and for occupation (note this is more fine-grained than our career groups). Online Appendix Figure A.I plots the major-group fixed effects on the y-axis.

IV. Testing a Light-Touch Policy Intervention

The previous section suggested that stereotyping may have substantial long-run labor market consequences. We turn to evidence on a light-touch policy intervention: a field experiment in which we tested an information intervention embedded in the second 2021 OSU survey. The survey began by asking students the percent chance that they would graduate with the two majors they selected as being most likely to pursue (their “top-ranked” and “second-ranked” majors). It then asked their self and population beliefs about the likelihood of each career group conditional on these two majors. Students were randomly sorted into a control group and a treatment group. Those in the control arm answered questions about their classes so far that semester and how they had (or had not) contributed to their major and career plans. These questions were designed to be similar in overall length and broadly about the same topic as the information module in the treatment arm but without providing students any new objective information.

In the treatment arm, information modules provided students with the actual distribution of careers conditional on each of their top two majors according to data from the ACS. For each major, it told them several headline numbers about the frequency of the careers they had listed as their most likely jobs if they graduated with that major.[25] We provided interactive infographics depicting the share of graduates with each major who were working in each career group (plus “other” and nonemployed). A further graphic broke down these groups into more detailed occupation titles. After showing this information for each major, we reasked students how likely they thought they would be to have each job if they graduated with that major. The information module (filled in with fictitious previous answers) can be accessed at this link: https://cmu.ca1.qualtrics.com/jfe/form/SV_b8zyALwAH0Nsy2O.

Because the treatment group saw information about the two majors they thought they were most likely to pursue, in practice each student was assigned to either zero (in the control group) or two (in the treatment group) of 10 possible information modules. A natural first question is the extent to which each of these treat-

25. These headline numbers were always about the two careers they said they would be most likely to be working in if they graduated with that major, plus (if not already included) that major’s most representative career.

ments succeeded in changing students' beliefs about their own chances of attaining each major's representative career. Online Appendix Figure A.II shows the average revision, among those in the treatment group who saw information about each major, in students' beliefs about their own chances of attaining that major's representative career. It plots this statistic against the average error in students' population beliefs: that is, the average belief about the share of Americans with that major working in its representative career minus the true share. Two facts stand out. First, while students on average revise in a sensible direction (reducing self beliefs when they overestimate population outcomes), this updating is far from one for one: a regression of average revisions on average errors yields a coefficient of only -0.28. Second, average revisions are heterogeneous across majors: for some majors, average revisions are well above or well below the trend predicted by the -0.28 coefficient, and we can reject the null hypothesis that each major's average revision is on this trend line ($p < .01$).[26] This result is perhaps not surprising: each module in the treatment group necessarily provided many distinct pieces of information (about each career group and more fine-grained occupations) that differed across majors, and we intentionally did not provide students guidance on how this information should affect their self beliefs.

Our question of interest is whether reducing students' beliefs about their likelihood of attaining representative careers influences their intentions and behavior. Following the literature on information interventions, we account for heterogeneity in treatment-induced belief changes (see Haaland, Roth, and Wohlfart 2023; Stantcheva 2023). This is particularly important in our setting because students received differing (and high-dimensional) information. We adopt a data-driven approach to test the causal effect of interest by constructing a simple measure of treatment intensity: the leave-out mean (LoM; i.e., the average excluding the student herself) reduction by major in treated students' self beliefs about their chances of having that major's rep-

26. This p-value comes from regressing revisions in self beliefs about a major's representative career on the average population error for that career as well as major fixed effects. If average errors simply reflected the average error across majors, we would expect all these fixed effects to be zero, and the p-value is for the null hypothesis that the major fixed effects are jointly zero.

resentative career.[27] Note that using the LoM reduction ensures that we are not conditioning on any posttreatment data for student i. We estimate equation (7) by OLS:

$$Y_{i,M} = \beta \cdot T_i \cdot AvgReduction_{i,M} + \alpha_1 Intentions_{i,M,pre} + \alpha_2 Classes_{i,M,pre} + \epsilon_{i,M}. \tag{7}$$

In the equation, $Y_{i,M}$ is an outcome variable of interest for a given major M for student i. T_i is a dummy variable for treatment status. $AvgReduction_{i,M}$ is the measure of treatment intensity described above: the LoM reduction in self beliefs about the representative career for major M. Finally, we include two controls: $Intentions_{i,M,pre}$ is i's pretreatment belief about their likelihood of graduating with major M, and $Classes_{i,M,pre}$ is the number of courses in M in their (pretreatment) fall 2021 class schedule.

Figure VI shows the coefficient on the interaction term from equation (7) for different outcome variables (Online Appendix Tables A.XII and A.XIII show the full estimates). The dark and light blue dots correspond to regressions where the dependent variable involves students' top- and second-ranked majors respectively. The leftmost dark blue dot shows that reducing stereotyping decreases intentions toward students' top-ranked majors: the coefficient of -1.1 implies that reducing students' self beliefs about their chance of having their top-ranked major's representative job by 10 percentage points decreases intentions toward that major by 0.11 standard deviations (about 3.5 percentage points, $p < .05$). The next six dark blue dots in Figure VI show analogous results but where the dependent variable is the number of classes students took in their top major every semester between spring 2022 (immediately posttreatment) and fall 2024. We find very similar initial effects as with intentions: reducing stereotyping regarding students' top-ranked major by 10 percentage points decreases class-taking in that major by 0.22 standard deviations (about 0.21 classes, $p < .05$) in the semester immediately following treatment.[28]

27. More precisely, we compute, for every student i, the average reduction in self beliefs about $P(c|M)$ among all students other than i. For students in the control group, this is simply the average reduction, since control group students were not asked to update these self beliefs. For treated students, we exclude i herself when computing these LoMs for them.

28. We are underpowered to investigate which if any majors are driving this result, but we do not see large differences in this coefficient if we drop any indi-

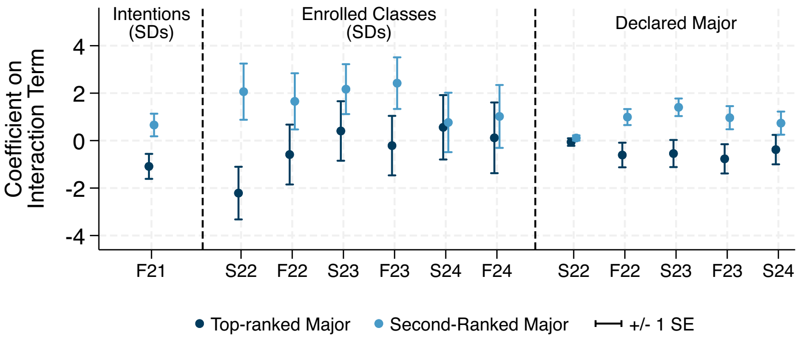

FIGURE VI

The Effect of Information

Each dot shows the estimate for the interaction term in equation (7) in a separate OLS regression. Dark blue dots indicate regressions where M is participants' top-ranked major, while for light blue dots it is their second-ranked major. The leftmost dots indicate regressions where the dependent variable is students' post-intervention updated belief about their likelihood of graduating with M. The middle dots indicate regressions where the dependent variable is the number of classes in M that students took during spring semester 2022 (S22), fall 2022 (F22), and so on. The dots on the right indicate similar regressions where the dependent variable is an indicator for whether students' had declared a major in M during those semesters.

In contrast to these effects on top-ranked majors, the light blue dots in Figure VI show that, if anything, reducing beliefs about attaining the representative job of students' second-ranked major directionally increases intentions toward it (0.07 standard deviations from a 10 percentage point belief change, $p = .17$) and immediate class-taking in it (0.21 standard deviations, or about 0.20 classes, $p < .10$). Both coefficients are statistically significantly different from the analogous results regarding top-ranked majors ($p < .01$ for intentions and $p < .05$ for classes). Although these results may appear surprising, in Online Appendix B.8 we show that they are consistent with a simple counterfactual simulation given estimates of students' heterogeneous preferences over careers: in particular, this is the pattern we would expect if students have much stronger preferences toward their top-ranked

vidual major. We can also investigate major switching in the administrative data. We find that the magnitudes of this phenomenon are far too small to meaningfully affect our results, however. By the eighth semester of college, only 2.46 percent of students have ever switched their declared major across our major groups.

major's representative career than that of their second-ranked major.[29]

The rightmost dots in Figure VI show that both treatment effects eventually fade, becoming smaller and statistically insignificant by fall 2024, three years after treatment. This fading is much more rapid for students' top-ranked major, disappearing entirely by fall 2022. While caution is warranted in interpreting these coefficients ex post, one possibility is that students manage to selectively forget the "bad news" about their top-ranked major more readily than the on-average good news about their second-ranked majors, consistent with "asymmetric updating" or motivated memory (e.g., Möbius et al. 2022; Amelio and Zimmermann 2023; Conlon 2025).[30]

These results investigate the effect of changing students' beliefs about the likelihood of attaining representative jobs by leveraging variation, across information modules, in how much stereotypical beliefs changed in response to the treatment. Online Appendix Tables A.XV and A.XVI instead regress various outcomes on controls and an indicator for treatment assignment, without interacting treatment status with this measure of treatment intensity. This specification tests not the extent to which changes in students' beliefs about their chances of representative careers affect their choices but the average effect of being assigned to the treatment group. These average effects are of course generally smaller (and usually not statistically significant), reflecting the fact that our intervention often had small (and for some majors null) effects on what jobs students expected themselves to have. However, Online Appendix Table A.XVII, column (1) shows that the treatment did have a significant average effect on the timing of students' major decisions: among students who have declared a major by spring 2024, those in the treatment group spent an additional 0.21 semesters ($p = .01$) undecided before doing so. This appears to be because they are somewhat less likely to have declared a major during their sophomore year but have done so at equal or even somewhat higher rates by the end of their junior year (a 6 percentage point increase, $p < .10$). This latter effect ap-

29. Of course, although these patterns are consistent with this story, there may be other explanations.

30. Note that all these results exclude from each semester's data students who are no longer taking classes at OSU. Online Appendix Table A.XIV shows very similar results if we instead impute zeros for such students.

pears to be driven by the fact that those in the treatment group are somewhat less likely to have dropped out by that time (5 percentage point effect, $p < .10$), which we measure by whether they are still enrolling in classes (see Online Appendix Table A.XVII, columns (7)–(12)).

V. Discussion

Across multiple survey samples, time periods, and elicitation methods, we find that U.S. undergraduate students greatly oversimplify the college-to-career process. Students appear to stereotype majors (“Art majors become artists,” “Political science majors become lawyers”), exaggerating the share of college graduates who are working in their major’s most representative job. These stereotypes in beliefs are strongly mirrored in implicit associations between majors and careers and seem partly driven by past experiences. A stylized model predicts—and empirically we confirm—that greater stereotyping produces misallocation, with more graduates ending up in jobs unrelated to their field of study. In a field experiment testing a light-touch intervention, we find that information changing these beliefs can have significant effects on students’ intentions and later choices.

One natural follow-up question to our results is why stereotyping persists despite the apparently large incentives that students have to make informed decisions about their education. For example, one could imagine students seeking out information (online, from better-informed friends, etc.) to correct their biased initial perceptions or learning from better-informed peers. We find, however (see Online Appendix B for more details), that the students in our data who are most confident that their beliefs about careers are correct are also the ones who stereotype most. These results, though only suggestive, point toward the possibility that biased students may fail to correct their beliefs because they are confident in their misperceptions (Enke, Graeber, and Oprea 2023). Another possibility, of course, is that the effect of any information may fade over time, for example if students forget it and revert to relying on what comes to mind intuitively. The fade-out of treatment effects is consistent with this interpretation, as is recent evidence on the short-term effects of quantitative information due to forgetting (Graeber, Roth, and Zimmermann 2024).

Even more speculatively, our results may help partly explain several striking and perhaps puzzling facts about students’ human capital decisions. For example, more U.S. undergraduates are currently pursuing a bachelor's degree in journalism than there are journalists in the entire country. Psychology majors outnumber accounting majors in the United States, and yet there are eight times as many accountants as psychologists. Students take on considerable debt to fund master's programs with appealing but unlikely associated careers (e.g., film studies).[31] Ex ante, of course, rational mechanisms could have fully explained these patterns: for example students with correct beliefs might rationally pursue certain career paths which, though very unlikely to pan out, they feel are worth the risk (e.g., journalism or film), or students may realize that certain majors (e.g., psychology) provide a general education not intended for use in any particular sector. Our findings suggest that mistaken beliefs may also contribute to these patterns: certain fields of study may appear especially appealing because students overestimate the likelihood that they will lead to attractive representative jobs. These human capital investments carry substantial monetary and opportunity costs, and therefore it may be beneficial to find ways to help students make better informed decisions or to nudge them toward less risky academic paths.

Supplementary Material

An Online Appendix for this article can be found at the *The Quarterly Journal of Economics* online.

Data Availability

The data and code underlying this article is available in Conlon and Patel (2026), in the Harvard Dataverse, https://doi.org/10.7910/DVN/LJMB5M.

The Ohio State University, United States
Brown University, United States

31. Shares of majors come from the ACS (authors' calculation), and the number of college graduates comes from the National Center for Education Statistics. Counts of occupations come from the Bureau of Labor Statistics' Occupational Outlook Handbook. See Korn and Fuller (2021) for the article on film studies master's programs.

References

Adorno, T. W., Else Frenkel-Brunswik, Daniel J. Levinson, and R. Nevitt Sanford, *The Authoritarian Personality*, (New York: Harper & Row, 1950).

Alesina, Alberto, Michela Carlana, Eliana La Ferrara, and Paolo Pinotti, "Revealing Stereotypes: Evidence from Immigrants in Schools," *American Economic Review*, 114 (2024), 1916–1948. https://doi.org/10.1257/aer.20191184.

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva, "Immigration and Redistribution," *Review of Economic Studies*, 90 (2023), 1–39. https://doi.org/10.1093/restud/rdac011.

Alfonsi, Livia, Mary Namubiru, and Sara Spaziani, "Meet Your Future: Experimental Evidence on the Labor Market Effects of Mentors," Working Paper, Harvard University, Cambridge, Massachusetts, 2025.

Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer, "Trends in Earnings Differentials across College Majors and the Changing Task Composition of Jobs," *American Economic Review Papers and Proceedings*, 104 (2014), 387–393. https://doi.org/10.1257/aer.104.5.387.

Amelio, Andrea, and Florian Zimmermann, "Motivated Memory in Economics—A Review," *Games*, 14 (2023), 15. https://doi.org/10.3390/g14010015.

Arcidiacono, Peter, V. Joseph Hotz, and Songman Kang, "Modeling College Major Choices Using Elicited Measures of Expectations and Counterfactuals," *Journal of Econometrics*, 166 (2012), 3–16. https://doi.org/10.1016/j.jeconom.2011.06.002.

Arcidiacono, Peter, V. Joseph Hotz, Arnaud Maurel, and Teresa Romano, "Ex Ante Returns and Occupational Choice," *Journal of Political Economy*, 128 (2020), 4475–4522. https://doi.org/10.1086/710559.

Arnold, David, Will Dobbie, and Peter Hull, "Measuring Racial Discrimination in Bail Decisions," *American Economic Review*, 112 (2022), 2992–3038. https://doi.org/10.1257/aer.20201653.

Arnold, David, Will Dobbie, and Crystal S. Yang, "Racial Bias in Bail Decisions," *Quarterly Journal of Economics*, 133 (2018), 1885–1932. https://doi.org/10.1093/qje/qjy012.

Arteaga, Carolina, "The Effect of Human Capital on Earnings: Evidence from a Reform at Colombia's Top University," *Journal of Public Economics*, 157 (2018), 212–225. https://doi.org/10.1016/j.jpubeco.2017.10.007.

Augenblick, Ned, Kelsey Jack, Supreet Kaur, Felix Masiy, and Nicholas Swanson, "Retrieval Failures and Consumption Smoothing: A Field Experiment on Seasonal Poverty," Working Paper, University of California, Berkeley, California, 2023.

Avitzour, Eliana, Adi Choen, Daphna Joel, and Victor Lavy, "On the Origins of Gender-Biased Behavior: The Role of Explicit and Implicit Stereotypes," Working Paper no. 27818, National Bureau of Economic Research, Cambridge, MA, 2020. https://doi.org/10.3386/w27818.

Baker, Rachel, Eric Bettinger, Brian Jacob, and Ioana Marinescu, "The Effect of Labor Market Information on Community College Students' Major Choice," *Economics of Education Review*, 65 (2018), 18–30. https://doi.org/10.1016/j.econedurev.2018.05.005.

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali, "The Search for Good Jobs: Evidence from a Six-Year Field Experiment in Uganda," *Journal of Labor Economics*, 43 (2025), 885–935. https://doi.org/10.1086/728429.

Becker, Gary S., *The Economics of Discrimination*, (Chicago: University of Chicago Press, 1957).

Bertrand, Marianne, Dolly Chugh, and Sendhil Mullainathan, "Implicit Discrimination," *American Economic Review*, 95 (2005), 94–98. https://doi.org/10.1257/000282805774670365.

Betts, Julian R., "What Do Students Know about Wages? Evidence from a Survey of Undergraduates," *Journal of Human Resources*, 31 (1996), 27–56. https://doi.org/10.2307/146042.

Bleemer, Zachary, and Aashish Mehta, "Will Studying Economics Make You Rich? A Regression Discontinuity Analysis of the Returns to College Major," *American Economic Journal: Applied Economics*, 14 (2022), 1–22. https://doi.org/10.1257/app.20200447.

Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope, "Inaccurate Statistical Discrimination: An Identification Problem," *Review of Economics and Statistics*, 107 (2025), 605–620. https://doi.org/10.1162/rest_a_01367.

Bordalo, Pedro, Giovanni Burro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, "Imagining the Future: Memory, Simulation, and Beliefs," *Review of Economic Studies*, 92 (2025), 1532–1563. https://doi.org/10.1093/restud/rdae070.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, Frederik Schwerter, and Andrei Shleifer, "Memory and Representativeness," *Psychological Review*, 128 (2021), 71–85. https://psycnet.apa.org/doi/10.1037/rev0000251.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. "Stereotypes," *Quarterly Journal of Economics*, 131 (2016), 1753–1794. https://doi.org/10.1093/qje/qjw029.

———, "Beliefs about Gender," *American Economic Review*, 109 (2019), 739–773. https://doi.org/10.1257/aer.20170007.

Bordalo, Pedro, John J. Conlon, Nicola Gennaioli, Spencer Y. Kwon, and Andrei Shleifer, "Memory and Probability," *Quarterly Journal of Economics*, 138 (2023), 265–311. https://doi.org/10.1093/qje/qjac031.

Carlana, Michela, "Implicit Stereotypes: Evidence from Teachers' Gender Bias," *Quarterly Journal of Economics*, 134 (2019), 1163–1224. https://doi.org/10.1093/qje/qjz008.

Chetty, Raj, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter, "Race and Economic Opportunity in the United States: An Intergenerational Perspective," *Quarterly Journal of Economics*, 135 (2020), 711–783. https://doi.org/10.1093/qje/qjz042.

Coffman, Katherine B., Christine L. Exley, and Muriel Niederle, "The Role of Beliefs in Driving Gender Discrimination," *Management Science*, 67 (2021), 3551–3569. https://doi.org/10.1287/mnsc.2020.3660.

Conlon, John J., "Major Malfunction: A Field Experiment Correcting Undergraduates' Beliefs about Salaries," *Journal of Human Resources*, 56 (2021), 922–939. https://doi.org/10.3368/jhr.56.3.0317-8599R2.

———, "Memory Rehearsal and Belief Biases," Working Paper, Carnegie Mellon University, Pittsburgh, PA, 2025.

Conlon, John J., and Spencer Y. Kwon, "Beliefs from Cues," Working Paper, Carnegie Mellon University, Pittsburgh, PA, 2025.

Conlon, John J., and Dev Patel, "Replication Data for: 'What Jobs Come to Mind? Stereotypes about Fields of Study'," (2026), Harvard Dataverse. https://doi.org/10.7910/DVN/LJMB5M.

Enke, Benjamin, Thomas Graeber, and Ryan Oprea, "Confidence, Self-Selection, and Bias in the Aggregate," *American Economic Review*, 113 (2023), 1933–1966. https://doi.org/10.1257/aer.20220915.

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann, "Associative Memory, Beliefs and Market Interactions," *Journal of Financial Economics*, 157 (2024), 103853. https://doi.org/10.1016/j.jfineco.2024.103853.

Ersoy, Fulya, and Jamin D. Speer, "Opening the Black Box of College Major Choice: Evidence from an Information Intervention," *Journal of Economic Behavior & Organization*, 231 (2025), 106800. https://doi.org/10.1016/j.jebo.2024.106800.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry, "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0," [data set], IPUMS, Minneapolis, MN, 2021. https://doi.org/10.18128/D030.V9.0.

Gawronski, Bertram, and Galen V. Bodenhausen, “The Associative–Propositional Evaluation Model: Theory, Evidence, and Open Questions,” *Advances in Experimental Social Psychology*, 44 (2011), 59–127. https://doi.org/10.1016/B978-0-12-385522-0.00002-0.

Gawronski, Bertram, Wilhelm Hofmann, and Christopher J Wilbur, “Are ‘Implicit’ Attitudes Unconscious?” *Consciousness and Cognition*, 15 (2006), 485–499. https://doi.org/10.1016/j.concog.2005.11.007.

Gennaioli, Nicola, and Andrei Shleifer, “What Comes to Mind,” *Quarterly Journal of Economics*, 125 (2010), 1399–1433. https://doi.org/10.1162/qjec.2010.125.4.1399.

Giustinelli, Pamela, “Expectations in Education,” in *Handbook of Economic Expectations*, Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw, eds. (Cambridge, MA: Academic Press, 2023), 193–224.

Glaeser, Edward L., “The Political Economy of Hatred,” *Quarterly Journal of Economics*, 120 (2005), 45–86. https://doi.org/10.1162/0033553053327434.

Graeber, Thomas, Christopher Roth, and Florian Zimmermann, “Stories, Statistics, and Memory,” *Quarterly Journal of Economics*, 139 (2024), 2181–2225. https://doi.org/10.1093/qje/qjae020.

Greenwald, Anthony G., and Mahzarin R. Banaji, “Implicit Social Cognition: Attitudes, Self-Esteem, and Stereotypes,” *Psychological Review*, 102 (1995), 4–27. https://psycnet.apa.org/doi/10.1037/0033-295X.102.1.4.

Greenwald, Anthony G., Debbie E. McGhee, and Jordan L. K. Schwartz, “Measuring Individual Differences in Implicit Cognition: The Implicit Association Test,” *Journal of Personality and Social Psychology*, 74 (1998), 1464–1480. https://psycnet.apa.org/doi/10.1037/0022-3514.74.6.1464.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart, “Designing Information Provision Experiments,” *Journal of Economic Literature*, 61 (2023), 3–40. https://doi.org/10.1257/jel.20211658.

Hastings, Justine S., Christopher A. Neilson, and Seth D. Zimmerman, “Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile,” Working Paper no. 19241, National Bureau of Economic Research, Cambridge, MA, 2013. https://doi.org/10.3386/w19241.

Kahana, Michael, *Foundations of Human Memory*, (New York: Oxford University Press, 2012).

Kahneman, Daniel, and Amos Tversky, “Subjective Probability: A Judgment of Representativeness,” *Cognitive Psychology*, 3 (1972), 430–454. https://doi.org/10.1016/0010-0285(72)90016-3.

Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad, “Field of Study, Earnings, and Self-Selection,” *Quarterly Journal of Economics*, 131 (2016), 1057–1111. https://doi.org/10.1093/qje/qjw019.

Kline, Patrick, Evan K. Rose, and Christopher R. Walters, “Systemic Discrimination Among Large U.S. Employers,” *Quarterly Journal of Economics*, 137 (2022), 1963–2036. https://doi.org/10.1093/qje/qjac024.

Korn, Melissa, and Andrea Fuller, “‘Financially Hobbled for Life’: The Elite Master’s Degrees That Don’t Pay Off,” *Wall Street Journal*, July 8, 2021.

Lowes, Sara, Nathan Nunn, James A. Robinson, and Jonathan Weigel, “Understanding Ethnic Identity in Africa: Evidence from the Implicit Association Test (IAT),” *American Economic Review*, 105 (2015), 340–345. https://doi.org/10.1257/aer.p20151075.

Malmendier, Ulrike, and Jessica A Wachter, “Memory of Past Experiences and Economic Decisions,” in *Oxford Handbook of Human Memory*, Michael J. Kahana and Anthony D. Wagner, eds. (Oxford: Oxford University Press, 2022), 2228–2266.

Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya Rosenblat, “Managing Self-Confidence: Theory and Experimental Evidence,” *Management Science*, 68 (2022), 7793–7817. https://doi.org/10.1287/mnsc.2021.4294.

Neal, Derek A., and William R. Johnson, "The Role of Premarket Factors in Black-White Wage Differences," *Journal of Political Economy*, 104 (1996), 869–895. https://doi.org/10.1086/262045.

Owen, Stephanie, and Derek Rury, "Implicit Gender-STEM Stereotypes and College Major Choice," Working Paper, Colby College, Waterville, Maine, 2025.

Phelps, Edmund S., "The Statistical Theory of Racism and Sexism," *American Economic Review*, 62 (1972), 659–661.

Phelps, Elizabeth A., Kevin J. O'Connor, William A. Cunningham, E. Sumie Funayama, J. Christopher Gatenby, John C. Gore, and Mahzarin R. Banaji, "Performance on Indirect Measures of Race Evaluation Predicts Amygdala Activation," *Journal of Cognitive Neuroscience*, 12 (2000), 729–738. https://doi.org/10.1162/089892900562552.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Megan Schouweiler, and Matthew Sobek, "Integrated Public Use Microdata Series, USA: Version 12.0," [data set], IPUMS, Minneapolis, MN, 2022. https://doi.org/10.18128/D010.V11.0.

Speer, Jamin D., "Pre-Market Skills, Occupational Choice, and Career Progression," *Journal of Human Resources*, 52 (2017), 187–246. https://doi.org/10.3368/jhr.52.1.0215-6940R.

Stantcheva, Stefanie, "How to Run Surveys: A Guide to Creating Your Own Identifying Variation and Revealing the Invisible," *Annual Review of Economics*, 15 (2023), 205–234. https://doi.org/10.1146/annurev-economics-091622-010157.

Wiswall, Matthew, and Basit Zafar, "How Do College Students Respond to Public Information about Earnings?" *Journal of Human Capital*, 9 (2015), 117–169. https://doi.org/10.1086/681542.

———, "Human Capital Investments and Expectations about Career and Family," *Journal of Political Economy*, 129 (2021), 1361–1424. https://doi.org/10.1086/713100.

CALL FOR NOMINATIONS

$300,000 Nemmers Prize in Economics

Northwestern University invites nominations for the Erwin Plein Nemmers Prize in Economics, to be awarded during the 2026–27 academic year. The prize pays the recipient $300,000. Recipients of the Nemmers Prize present lectures, participate in department seminars, and engage with Northwestern faculty and students in other scholarly activities.

Details about the prize and the nomination process can be found at **nemmers.northwestern.edu**. Candidacy for the Nemmers Prize is open to those with careers of outstanding achievement in their disciplines as demonstrated by major contributions to new knowledge or the development of significant new modes of analysis.Individuals of all nationalities and institutional affiliations are eligible except current or recent members of the Northwestern University faculty and past recipients of the Nemmers or Nobel Prize. **Nominations will be accepted until January 14, 2026.**

The Nemmers prizes are made possible by a generous gift to Northwestern University by the late Erwin Esser Nemmers and the late Frederic Esser Nemmers.

Northwestern

Nemmers Prizes • Office of the Provost • Northwestern University • Evanston, Illinois 60208
nemmers.northwestern.edu