# The Hidden Curriculum

Michael G. Cuna[*] Elaine Shen
University of Chicago UC Berkeley

September 4, 2026

*Please click here for the latest version.*

**Abstract**

Despite expanded access to selective colleges, first-generation students often trail their peers in educational and early-career outcomes. We study one mechanism: the hidden curriculum – implicit, unwritten practices critical to educational and career success such as networking. First, we use data from over 70,000 students across 20 public universities to identify a first-generation gap in investment in hidden-curriculum actions. Second, in a field experiment at UC Berkeley, we show an information intervention increases take-up of hidden-curriculum actions by 13pp, closing roughly 77% of the baseline gap. Finally, an online experiment with AI college advisors shows that first- and continuing-generation students exert similar information-acquisition effort, but respond differently to hidden-curriculum nudges. Proactive information about returns eliminates this gap. The results highlight that asymmetric information about potential strategies and returns may prevent equal access from delivering equal opportunity and equal outcomes.

**JEL Codes: I23, I24, J24, D83**

*Keywords*: hidden curriculum; information frictions; educational inequality; social mobility; human capital investments.

[*]Corresponding author: Michael G. Cuna (michaelcuna@uchicago.edu).

# 1 Introduction

First-generation and low-income students continue to experience large gaps in education and labor market outcomes: these students tend to do worse academically, are less likely to finish college, graduate with lower average wages, and are less likely to hold C-suite or CEO positions (Chen, 2005; Angrist et al., 2022; Zimmerman, 2019). One way organizations have tried to close these gaps is by offering these traditionally disadvantaged students membership in these organizations, yet gaps in outcomes persist, suggesting formal membership alone is insufficient. Existing research has shown social ties and social networks contribute to the persistence of these gaps; we propose a new mechanism: the hidden curriculum.

We argue that a key driver of persistent gaps in educational and labor market outcomes is the "hidden curriculum": unwritten, context-specific norms, actions, and expectations that are important for success, but that are neither formally required nor explicitly communicated. These norms span a wide range of educational and labor market settings – from knowing how to cold-email an alumnus about a job or cultivate a relationship with a professor who may later write a reference letter, to recognizing which tasks are rewarded in competitive workplaces or understanding conference etiquette and paper-submission norms in academia. We study the hidden curriculum in an important setting: the career and academic choices of first-generation college students. We show that, beyond formal qualifications, skills, and constraints, differences in access to tacit knowledge generate persistent differences in students' actions across educational and labor market domains.

This paper starts with a simple conceptual framework. The framework helps clarify the definition of the hidden curriculum, the mechanisms that govern it, and the consequences for those affected by it. We model the phenomenon as a traditional skill-augmenting investment problem with a hidden-curriculum input and a formal-curriculum input. We then discuss two distinct informational frictions: (i) awareness of actions, and (ii) beliefs about the returns to such actions. The simple framework predicts that for those impacted by the hidden curriculum we would see: (i) lower investment in the hidden curriculum, and (ii) generally higher investment in the formal curriculum.

We then combine evidence from three empirical strategies. First, we use observational and survey data for more than 70,000 students across 20 large U.S. public colleges. We first document that first-generation students invest less in hidden-curriculum actions and more in formal-curriculum actions. In particular, they are .19 points lower on an index of hidden-curriculum actions, which is 30% of a standard deviation of the raw index[1]. We also document that these differences persist over time. We then show that these differences cannot be fully ruled out by accounting for gender, ethnicity, U.S. nationality, or home state. We also control for constraints such as socioeconomic status, high-school GPA, and many other controls. We then show that first-generation students are not less likely to be involved in formal but social activities, providing suggestive evidence that the pattern is not simply a preference issue.

From the observational data, we also show that first-generation students achieve worse outcomes. In particular, first-generation students are 18.5% less likely to get a for-credit internship. When controlling for the usual observables, the gap is reduced but remains statistically significant ($p \leq 0.01$). However, when controlling for hidden-curriculum investments, we find a gap of $+0.02 p.p.$ that is not statistically significant ($p = .947$). This evidence suggests that hidden-curriculum actions can explain part of the post-enrollment disparity in outcomes.

Since the observational data alone cannot determine the causal impact of information on actions, we control for the assignment mechanism by running a pre-registered field experiment with 645 students at a large U.S. public university (UC Berkeley). Students first report baseline information and are randomly assigned to either a control group or one of two information interventions. After receiving their assigned information, they report whether they are willing to pay for a service that provides help with a hidden-curriculum action: cold networking. We design the first information intervention to provide students with the information that cold networking is an available strategy. We design the second information intervention to update their beliefs about the returns to such actions.

The experimental results show that first-generation students are 16.7 percentage points less likely to take up the hidden-curriculum action, in this case a networking service. The randomized information treatment shows that information frictions contribute to these differences in actions: first-generation students who receive the information treatment are 12.9 percentage points more likely to take up the hidden-curriculum action. We also cannot reject that the two treatments reduce the gap by the same amount ($p = .39$).

We show that information also plays a role for other groups that may not have the same insider knowledge about the institution and does not play a role for traditionally disadvantaged groups that face constraints other than information (e.g., gender). Regarding the former, we show that the more detailed treatment changes demand for the service

[1] The SD of the index is 0.632

among community-college transfer students, while the lighter treatment changes demand by less. We provide an explanation of this heterogeneity in treatment response in Section 4.2. The existence of heterogeneity in mechanisms across groups implies the need for a more scalable solution compared with traditional interventions or one-off nudges to increase access to on-campus resources.

Finally, we explore one of the mechanisms that keep the hidden curriculum hidden: search frictions. In a pre-registered online experiment, we provide students with an AI college advisor tailored to help them with questions regarding behaviors and strategies that they can adopt both within and outside of college to succeed in the labor market. We randomize students into one of two conditions: (i) a Passive AI that mainly discusses topics directly related to the student's question, and (ii) an Active AI that purposefully informs students of hidden-curriculum strategies in relation to the question they asked. We then use the transcripts to measure students' own search effort, the content of their questions, and their next-turn continuation after the AI introduces a hidden-curriculum topic.

We first find through qualitative coding that permission or legitimacy concerns and social-risk language are several times more common among first-generation students (20.0 versus 4.5 percent of coded excerpts in the passive condition). We then find that first-generation students are much less likely to follow up on hidden-curriculum topics within the Passive AI than they are within the Active AI nudges. When the Passive AI raises a hidden-curriculum topic, first-generation students have a lower point estimate in continuation of hidden-curriculum topics. Active guidance reverses this pattern: under the Active AI, first-generation students have a statistically significantly larger increase in relative continuation of hidden-curriculum topics, represented by a relative change of 29.4 percentage points.

These patterns help explain why information gaps persist in equilibrium: the hidden curriculum persists not because first-generation students fail to invest effort to search for information, but because hidden strategies must be made visible and at the same time be perceived as legitimate and easy to act on before students engage with them.

Our results also point to a distinction between equal access and equal opportunity. Students may have equal access to an institution while differing in their knowledge of how to take advantage of this access. Because educational and career outcomes are a product of a series of actions taken over time, disparities in these intermediate actions can compound into disparities in later outcomes. We do not observe these long-run consequences directly, but our findings show that an information intervention can substantially reduce one such disparity in the short run, providing evidence for a mechanism through which differences in information may contribute to persistent gaps in educational and labor market outcomes.

This paper makes two main contributions. First, we build on the literature that seeks to understand barriers to upward mobility by showing how gaps in hidden-curriculum actions may contribute to gaps in educational achievement and labor market outcomes (Zimmerman, 2019; Michelman et al., 2022; Angrist et al., 2022). By documenting a gap in hidden-curriculum actions between first-generation and continuing-generation students, our results show that hidden information and costly learning may prevent disadvantaged students from succeeding, even after they gain formal membership to elite organizations.

Second, we provide evidence that differences in post-enrollment investments are shaped by information frictions. Existing work shows that beliefs about returns, norms, and available resources shape education and labor-market choices (Zafar, 2011; Jensen, 2010; Wiswall and Zafar, 2015; Baker et al., 2018; Dizon-Ross, 2019; Bursztyn et al., 2020; Conlon, 2021). We focus on a related but different type of friction: an individual must know that the action exists before investing, and must also understand why it matters, view it as legitimate, and know how to carry it out. Our field experiment shows that making an informal strategy visible and relevant disproportionately changes first-generation students' consequential demand. The AI experiment takes the next step by distinguishing passive availability of the information from engagement with the information. First-generation students differ in the content of their questions and in their relative continuation of hidden-curriculum topics once the advisor introduces them; active guidance changes this continuation margin. We therefore extend research on information provision, attention, and incomplete take-up from formal programs and explicitly stated choices to strategies that are typically transmitted informally (Altmann et al., 2018; Belot et al., 2019; Bhargava and Manoli, 2015; Finkelstein and Notowidigdo, 2019; Manzini and Mariotti, 2014).

Beyond these two main contributions, the paper speaks to several broader literatures. It contributes to research on opportunity and social mobility by identifying a post-entry channel through which equal access may fail to generate equal opportunity and, subsequently, equal outcomes. Students can enter the same institutions yet differ in their knowledge of and investment in the unwritten practices that convert enrollment into faculty relationships (Canaan et al., 2022), professional networks (Marmaros and Sacerdote, 2002; Ost et al., 2026), internships, and career opportunities (Arellano-Bover et al., 2026). This adds a tacit-knowledge channel to research on college access and intergenerational mobility (Hoekstra, 2009; Chetty et al., 2017, 2020; Dynarski and Scott-Clayton, 2013), complements classic accounts emphasizing the intergenerational transmission of financial and human capital (Becker and Tomes, 1979; Solon, 1999), and could explain first-generation frictions in graduate-school labor-market outcomes (Stansbury and Rodriguez, 2026). In doing so, we bridge sociology research on the hidden curriculum and cultural capital (Jackson, 1968; Bourdieu, 1986; Calarco, 2011; Jack, 2019) with economic work on attention, consideration, and decision-making under imperfect information (Masatlioglu et al., 2012; DellaVigna, 2009; Gabaix, 2019) by translating tacit know-how into measurable investments and experimentally varying information about it. The paper also extends the literature on incomplete take-up (Bhargava and Manoli, 2015; Finkelstein and Notowidigdo, 2019) from formal programs to informal strategies, for which nominally available information need not translate into engagement.

The paper proceeds as follows. In Section 2 we introduce a simple conceptual framework for the hidden curriculum. In Section 3 we use the observational data to document differences in hidden-curriculum actions and outcomes. In Section 4 we show causal evidence through a field experiment of information frictions leading to under-utilization of hidden-curriculum strategies. In Section 5 we use a randomized AI advisor to separate search effort from the uptake of hidden-curriculum information. Finally, we conclude in Section 6 and discuss the policy relevance and potential areas for further research.

# 2 The Hidden Curriculum: Definition and Framework

In this section we formally define the hidden curriculum and introduce a simple conceptual framework. The framework provides a foundation for understanding: i) patterns in the data, ii) mechanisms that define this information friction, and iii) policy implications.

## 2.1 Definition and Examples

We define the hidden curriculum as actions or strategies that are consequential for success and are not formally required or communicated.

In U.S. colleges first-generation students are the main group of interest: parents who did not attend a four-year college could not, by construction, transmit institution-specific knowledge about the opportunity structures and informal practices of four-year colleges.

When the literature in sociology started in the 1960s the term referred to the purely academic side of college. It was used to explain why first-generation students had lower grades and dropped out at a higher rate. In most of this paper we instead focus on actions students use during college to get hired into the career and position of their choosing[2].

With this definition in mind[3] we can now categorize some actions into being either part of the hidden curriculum or part of the formal curriculum.

Actions in the formal curriculum appear in a class syllabus. They are formally required, and students whose parents stopped their schooling after high school, or after an equivalent degree, have already experienced them. GPA is one such requirement: the university states it and most jobs ask for it.[4] Clearly stated actions that raise grades are therefore not part of the hidden curriculum. Examples include doing the class readings, spending an extra hour on homework, and revising an assignment when the instructor allows it.

On the other hand, many high-skilled jobs value signals of passion and preparedness that they do not formally state in the application process. For example, cold networking[5] or joining a student organization are both valued by hiring managers. However, they are usually not formally stated requirements from either the university or the job application.

## 2.2 A Simple Framework

The conceptual framework in this section gives intuition on the problem the student faces and its implications. It shows the two frictions, i) an awareness friction and ii) a belief friction, that we treat as the hidden curriculum, and it sets up the policy discussion in Section 2.3.

A student enters college with a given baseline ability $a$. The student then has two inputs available. One is the hidden-curriculum action $h \geq 0$ and the other is the formal-curriculum action $v \geq 0$. The above three components enter the following quality index

$$Q = f(h, v, a). \tag{1}$$

[2]This is a sound approach for two reasons: i) if modeled in its entirety this step would come after a student solves a Roy model to determine career choice and is now optimizing how to reach that career, ii) the literature has shown that first-generation students are less likely to switch careers and majors (Aucejo et al., 2025); therefore, studying what they do after having picked a path is not only relevant but also theoretically sound.

[3]Unstated and never experienced by the student or their family.

[4]Most job applications for high-skilled labor require submitting transcripts and GPA.

[5]Cold networking is the action of reaching out and networking with someone you do not have a direct connection to. One very common case of cold networking is reaching out to alumni from your institution and leveraging the indirect connection as a way to increase chances of positive engagement and interactions.

The student will receive monetary benefits from future labor market earnings. These are captured by $B(Q)$ and should be interpreted as integrated over labor market uncertainty. The student will also face costs $C(h, v)$, which include both material and psychological costs of investments. In a perfect information environment the student will then choose a pair $(h, v)$ that maximizes the student's net payoff

$$U(h, v; a) = B\big(f(h, v, a)\big) - C(h, v). \tag{2}$$

These scalar inputs can be interpreted as a representation of all opportunities.

We assume continuing-generation students are informed and first-generation students face information frictions. The intuition behind this conceptual framework does not need every first-generation student to be uninformed or every continuing-generation student to be fully informed.

**Two information frictions.** The hidden curriculum can affect students' choices in two ways.

First, a student may not be aware of an action. An alternative interpretation is that the student is not aware that the action is available in a specific context. If that is the case, the student will not consider it when optimizing.

Second, the student may know of the action, but have a distorted view of the benefits. This is represented as

$$\widetilde{U}(h, v; a) = U(h, v; a) - \psi(h), \qquad \psi(0) = 0, \quad \psi \text{ strictly increasing}, \tag{3}$$

The wedge $\psi$ represents the friction in beliefs. This wedge is the difference between the real and perceived value of hidden-curriculum actions. We represent students' various costs (e.g., effort and time) with $C$, and they enter the true payoff $U$.

**Prediction 1: hidden-curriculum investment.** Both frictions reduce the investments in hidden-curriculum actions. The result is mechanical when the awareness friction is present: a student cannot choose an action they do not know about. The result for the valuation friction simply derives from lower expected marginal gains. Holding ability, opportunities, and true costs fixed, either friction therefore implies

$$h^{FG} \leq h^{CG}, \tag{4}$$

where $h^{FG}$ denotes the choice under either friction. (See the Online Appendix for a formal proof.)

**Prediction 2: the allocation of visible effort.** How the hidden curriculum changes investments in $v$ depends on whether the two inputs are complements, substitutes, or independent.

If the two inputs are net substitutes, a student who is unaware of $h$ or who undervalues $h$ reallocates toward the visible margin, thus lower $h$ is accompanied by weakly higher $v$. If the two inputs are net complements, lower $h$ lowers the marginal value of formal investment, and $v$ weakly falls.

The main takeaway from the conceptual framework is that hidden-curriculum-type frictions lower investment in those actions. Whether formal investment is higher depends on whether the two inputs are complements or substitutes, which is an open empirical question that is not the focus of this paper. In our setting one interpretation is that a student who is uncertain about networking or a professional organization can still see that readings and assignments raise grades, which the market rewards. While we show that the empirical patterns in the next section are consistent with this form of portfolio reallocation, we stress that the implications of this article do not require it.

## 2.3 Why the Mechanism Matters for Intervention

The distinction between not knowing that an action exists and incorrectly valuing the action's returns has direct implications for intervention design, even though the framework is static. In Section 5, we show that students struggle to invest time and effort in learning about actions they do not know exist. This would imply that if most of the hidden curriculum is coming through the consideration-set mechanism, we would see slower learning across a student's tenure in college. If instead the issue were mainly subjective beliefs, the student could, if incentives were aligned, invest effort into learning more about a certain action.

To close the information gap, addressing a consideration-set issue would require more than simply offering students resources. If the student does not know they lack knowledge of inputs, then they are less likely to invest time and effort into figuring out how and when to interact with on-campus services such as first-generation offices and career services.

The two frictions likely coexist. Furthermore, information interventions likely move both objects. Making students aware that a strategy exists and is legitimate may also raise their beliefs about its returns. An information intervention on high returns mechanically forces the student to be aware of the action. For this reason, we design our treatments to load differently on the two mechanisms rather than to isolate them perfectly. The light-touch information intervention discussed in Section 4 is designed to mainly load on the first friction, while the more heavy-handed information intervention is designed to update beliefs on returns. Likewise, active AI guidance can both introduce an option that the student did not raise and present it in a way that makes its value easier to assess and its execution easier to undertake.

While not presented formally in the framework, in settings where early hidden-curriculum investments open access to later opportunities, discovering an action only after a deadline or recruiting stage may have lasting consequences. Such cumulative effects are a natural extension of the framework and should be a topic of future research.

# 3 Evidence from Observational Data

In this section, we use a large survey linked with administrative data to show motivating facts about first-generation students and how their behaviors cannot be fully characterized by preferences and constraints. We first show differences in effort across our two student groups across hidden- and formal-curriculum actions. We then show differences in outcomes and how effort investments in hidden-curriculum actions may explain this difference in outcomes. Finally, we report patterns in hidden curriculum across college years.

## 3.1 SERU Data

This section uses data from the "Student Experience in the Research University" (SERU) Consortium, specifically focusing on the 2022 cross-sectional dataset. The SERU Consortium conducts large-scale surveys targeting undergraduate students enrolled in major research-intensive universities.

The full sample includes more than 100,000 students, while the most conservative sample includes 73,989 students from twenty institutions, primarily characterized by their high research activity, substantial undergraduate enrollments, and distinguished academic programs. These institutions range geographically across various regions of the United States, enabling us to examine patterns broadly representative of large research universities. The universities in our sample are also on average highly ranked. According to nationally recognized rankings we have 11 schools that are ranked in the top 50 undergraduate programs in the nation, 17 in the top 100. We therefore view first-generation students in our sample as possibly highly positively selected.

The dataset includes extensive student-level survey responses, complemented by administrative data on student academic and demographic characteristics, such as cumulative GPA, high school GPA, tuition payments, credits taken, residency status, family social class, race, and gender.

Of particular interest for our analysis are variables directly linked to the hidden and formal curriculum. The dataset includes many variables related to actions within purely academic settings, such as various measures of faculty interactions (e.g., how many faculty students feel comfortable asking for a reference letter), participation in study groups, and seeking help from faculty or peers. Additionally, the dataset includes actions that are extra-academic, such as peer help with networking, participation in internships, joining professional clubs, and taking on leadership roles. Given the extensive nature of the dataset in terms of inputs, in the next sections we focus on a subset of actions, namely those survey items that are included at most schools.

Additionally, the dataset contains optional survey modules that further enrich our understanding of student engagement, experiences, and behavior. While our main empirical analysis focuses specifically on measures directly related to the hidden curriculum and academic investments, additional variables capturing broader dimensions of student well-being, overall satisfaction, and college experience are available and discussed in supplementary analyses provided in the online Appendix.

Appendix Table A.1 provides summary statistics comparing first-generation (FG) and continuing-generation (CG) students, highlighting key differences across observable characteristics. Notably, first-generation students tend to be older and are substantially more likely to come from lower-income and working-class backgrounds compared to their continuing-generation peers. Additionally, first-generation students are more likely to be Hispanic or Latino (47% for first-generation students versus 10% for continuing-generation students when combining Chicano/Latino and Latino) whereas continuing-generation students are disproportionately White (16% versus 6% of first-generation students). First-generation students are also more likely to be in-state students (87% versus 79%).

The demographic and academic background variables let us condition on the observables that the leading alternative explanations turn on. We nonetheless treat these results as descriptive, and we turn to a randomized experiment for causal evidence.

## 3.2 First-Generation Hidden Curriculum Gap

Our framework predicts that first-generation students invest less in hidden-curriculum actions and invest weakly more in formal-curriculum actions. We test this prediction directly before turning to outcomes. If limited knowledge of what actions to take during college drives the observed patterns, first-generation students should report lower effort in hidden-curriculum activities and higher effort in formal-curriculum activities.

We pick a mix of academic and non-academic actions in both sets to show that this pattern holds across domains. We summarize these actions as portfolio indices and report individual components in the Appendix.

We construct each index as the weighted average of standardized hidden-curriculum and formal-curriculum components. For the hidden-curriculum index we use the following actions: faculty relationships outside of class activities[6], whether a professor knows the student's name, assisting faculty research, and student organization participation. We chose these actions because they fit the definition of not being formal requirements and because prior literature in sociology and education has shown differences across socioeconomic classes and other disadvantaged groups (Jack, 2016; Stuber, 2009; Michelman et al., 2022; Kim and Sax, 2009)[7].

For the formal-curriculum index we use the following actions: readings, presentations, methods coursework, independent study, and increased academic effort.

[6]This should be interpreted as the student creating a relationship with the faculty for reference and career purposes.

[7]As a robustness check, we fed ChatGPT and Claude the list of variables reported by all colleges in the dataset and asked the LLMs to identify actions that they believed to be part of the hidden curriculum and visible curriculum. The LLM-recommended variables tended to match the variables discussed in the previous literature.

Figure 1: First-Generation Gaps in HC and FC Portfolio Indices and Components

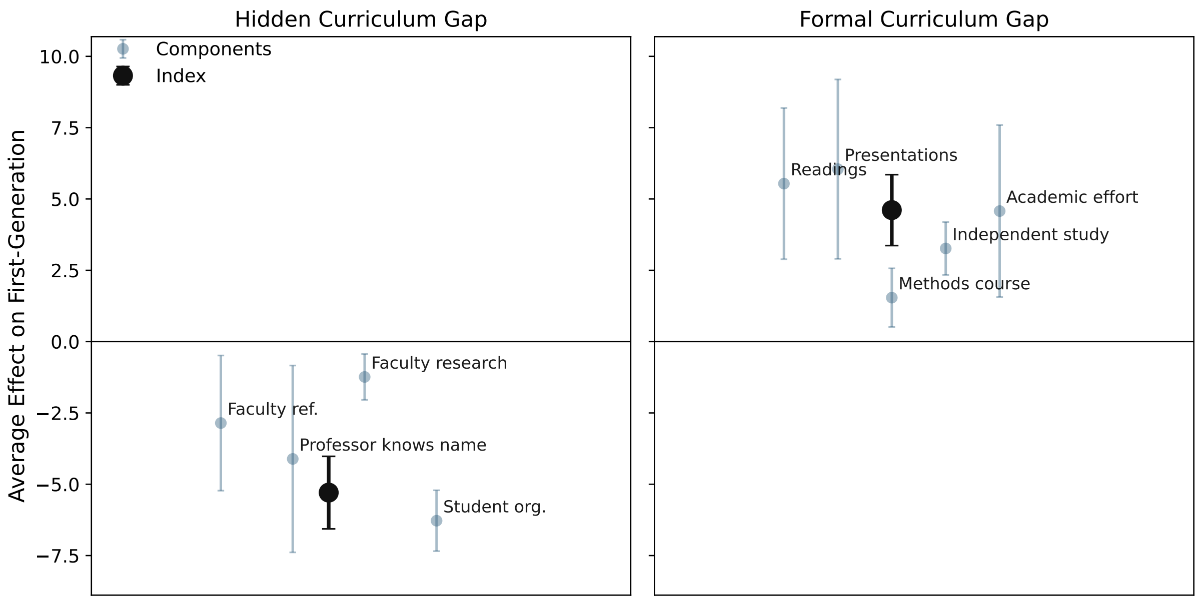


*Notes:* The figure reports OLS coefficients from the controls specification for the HC and FC portfolio indices and components. Large dark points are portfolio-index estimates; smaller transparent points are component estimates. Controls include gender, race/ethnicity, social class, in-state status, U.S. national status, university fixed effects, major fixed effects, and year in college. Vertical bars show 95 percent confidence intervals.

In Figure 1 we report the results of an OLS regression of actions on first-generation status, including controls, university fixed effects, major fixed effects, and approximate college-year controls. The results are more formally presented in Appendix Table A.2.

In raw differences, the hidden-curriculum index is 0.191 lower for first-generation students, and it is 0.053 lower in the preferred specification. The formal-curriculum index is 0.048 standard deviations higher in raw differences and 0.046 standard deviations higher in the preferred specification. These patterns hold true for other specifications as well, including coarsened exact matching and propensity-score balancing.

The individual action table in Appendix Table A.3 shows the same pattern in components: first-generation students report lower faculty reference relationships, lower name recognition by professors, lower rates of assisting faculty research, and lower student organization participation, while formal actions are generally higher.

We tested the same relationship for placebo actions such as class discussions, active participation, participation in study groups, or asking questions outside of class. These are actions that are social in nature, that measure engagement with overall school life, and are not geared towards professional progression. Here we see none of the clear negative pattern seen in the hidden-curriculum actions.

## 3.3 First-Generation Outcome Gaps

We now move on to characterizing differences in outcomes between first-generation and continuing-generation students. First, we show that first-generation students perform worse across a range of outcomes. We then provide motivating evidence that constraints are not the sole cause of these differences.

We start by looking at college GPA and internship outcomes. In raw differences, first-generation students are 3.7 percentage points less likely to report a for-credit internship, 9.2 percentage points less likely to report a non-credit internship, and have college GPA that is lower by 0.247 grade points. To test the robustness of these differences we regress each outcome on first-generation status. In our preferred specification we include controls for gender, race and ethnicity, social class, in-state status, U.S. national status, university fixed effects, major fixed effects, and approximate college year. In Appendix Table A.5 we report the results for internships and GPA. The preferred controls reduce the for-credit internship gap to 1.1 percentage points, the non-credit internship gap to 3.5 percentage points, and the GPA gap to 0.079 grade points.

We now focus on for-credit internships and leave the other outcomes for the Online Appendix. We focus on for-credit internships for many reasons. First, for elite jobs there has been an increase in companies considering internships as a necessary step in order to secure their full-time role. Second, the recruiting process is similar, if not the same, as for the full-time role, making it a great proxy for full-time labor employment. Finally, compared to non-credit internships, these are largely paid, so they do not raise the concern that first-generation students cannot afford an unpaid internship.

Figure 2: For-Credit Internship Gap With and Without Portfolio Controls

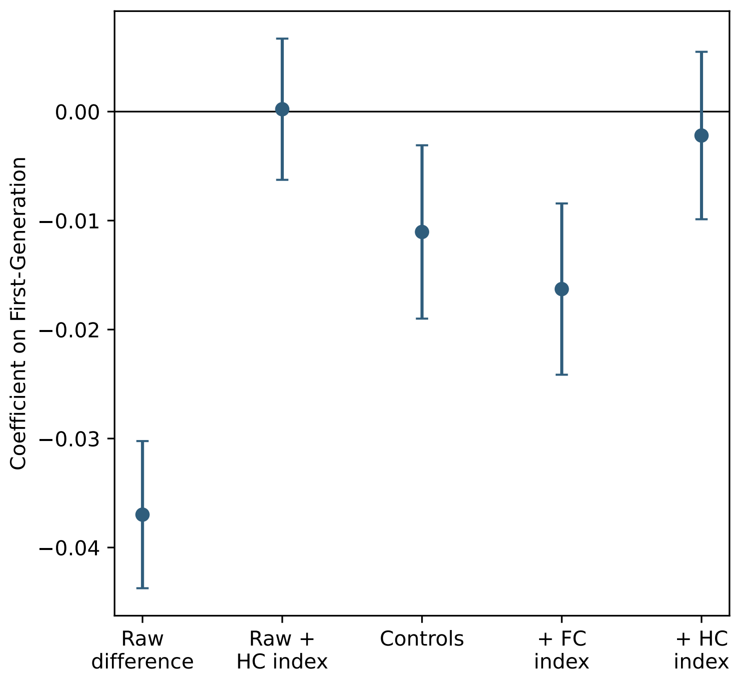


*Notes:* The figure reports the descriptive first-generation gap for for-credit internships. The raw + HC index column estimates the gap after adding only the generated HC index, without demographic controls or fixed effects. Controls columns control for gender, race/ethnicity, social class, in-state status, U.S. national status, university fixed effects, major fixed effects, and year in college; portfolio-control columns add the FC or HC index. Vertical bars show 95 percent confidence intervals.

We add in turn to our specification first actions from the formal curriculum, and then actions from the hidden curriculum. From Figure 2 we see that controlling for formal-curriculum actions leaves the gap essentially unchanged. However, controlling for hidden-curriculum actions reduces the gap to the point that it is no longer statistically significant. The same is true if we take out all other controls and include only the hidden-curriculum index. In the appendix we see a similar, yet not as stark, pattern for GPA and non-credit internships.

Our controls address the constraints most common in the literature, such as financial constraints or a language barrier driving the differences in outcomes. They do not address baseline ability. We therefore add high school GPA as a proxy for college preparedness.[8]

In Appendix Table A.4 we first see that the actions gap remains in both formal- and hidden-curriculum actions. In Appendix Table A.6 we also see that internship outcomes remain stable and significant.

Finally, we report the results from an Oaxaca-Blinder decomposition for receiving a for-credit internship in Appendix Figure A.1. We estimate a pooled linear model for for-credit internships, using the hidden-curriculum and formal-curriculum action components, controls, university fixed effects, broad major group, and year in college, and then decompose the raw continuing-generation minus first-generation gap into the part associated with differences in those covariates and a remaining unexplained component. If we normalize the gap to 100%, the model accounts for about 86% of the gap, leaving about 14% unexplained. Among individual actions, assisting faculty research accounts for about 27% of the total gap, student organization participation accounts for about 26%, having a faculty reference relationship accounts for about 12%, and professor knows name accounts for about 2%. Taken together, the hidden-curriculum action components account for about 68% of the total for-credit internship gap. From Appendix Figure A.1, we see that the largest action-level shares of the explained gap are concentrated in hidden-curriculum actions.

## 3.4 Gaps Across Cohorts

We now focus our attention on how students differ across years. While our data do not allow us to have a panel-like structure, we intend to give a snapshot of how students in different years behave. The reasoning is that even if students differ because of cohort characteristics they should still show some convergence toward similar means. Four years of college should be enough time to see behaviors of first-generation students converge toward the means of their continuing-generation peers, even if they never fully reach them because of unobservable characteristics.

Figure 3 provides descriptive evidence that first-generation students do not catch up in hidden-curriculum investment during their college careers. The estimated gap is negative in every year and becomes more negative in years three and four. Because the data are

[8]For high school GPA we have to drop observations from students who do not have traditional GPAs, for example those reported out of 100 or out of 60, since we cannot weight them properly.

Figure 3: HC Index Gap by Approximate College Year

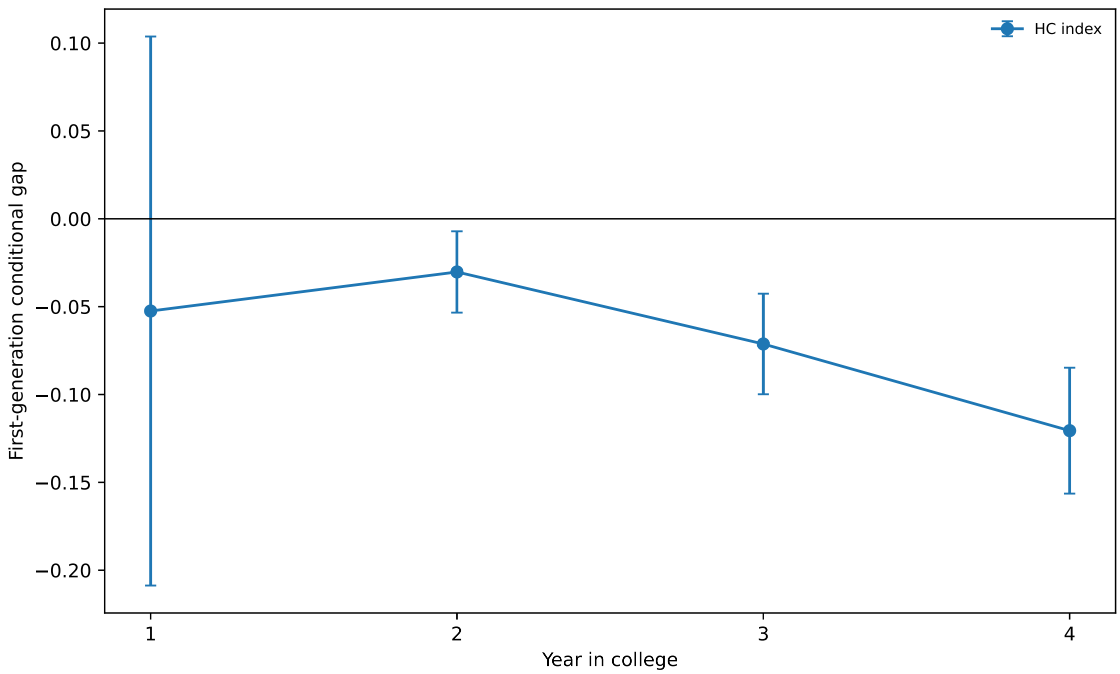


*Notes:* The figure reports the descriptive first-generation gap in the generated HC index among non-transfer students in years 1 through 4. Estimates control for gender, race/ethnicity, social class, in-state status, U.S. national status, university fixed effects, and major fixed effects.

repeated cross-sections, we cannot interpret this as direct evidence of a lack of learning, but it is suggestive of a lack of convergence in behavior. This is consistent with multiple types of mechanisms: for one, students may never learn these strategies, or if they learn them too late they may not be allowed to participate in these activities, or they may face cumulative disadvantages because certain opportunities can only be taken conditional on having achieved certain requirements early on in their careers.

# 4 The Role of Information

The administrative data analysis in Section 3 allows us to document substantial differences in hidden curriculum (HC) investment between first-generation and continuing-generation students and rule out several potential explanations; however, the data can only bring us so far. To pinpoint the causal link between information on actions and first-generation students' choices in investments we move to a field experiment. This allows us to precisely control the assignment mechanism and better understand the mechanisms driving differences in hidden-curriculum knowledge.

## 4.1 Experimental Design

To conduct the field experiment, we recruited students through two methods to reach a diverse, representative group of participants. The first method was by distributing the survey across campus with the help of a large group of major offices and student organizations. We also publicized the survey through mailing lists and on-campus flyers. The second was by distributing the survey through the Xlab Berkeley (Experimental Social Science Laboratory). We report pooled results since student demographics do not look significantly different by distribution method[9].

We screen students to make sure they are primarily interested in careers in business, finance, or technology.[10] The final sample consisted of 645 students. In the control group, 66.2% of students are female, 23.0% are first-generation, and 17.8% are transfer students. The two largest ethnic groups were Asian and White. The former group is over-represented in our experiment compared to the overall underlying population. However, we do not

[9] Specifications that include distribution method fixed effects do report the same results.

[10] We actively excluded students interested in an academic career because our outcome variable of interest for this specific design would not have been natural for such students. However, future research should also include students interested in such careers since similar achievement gaps between first-generation and continuing-generation students have been documented in academia (Stansbury and Rodriguez, 2026).

know the distribution of ethnicity for those interested in the careers we select on.

Once recruited, students completed a survey on demographic and baseline characteristics to use as controls and to identify first-generation status. The experiment is pre-registered in the AEA RCT Registry.

Participants were randomly assigned, using complete randomization, to one of three groups: a control group, a light-touch information treatment (T1), and a detailed, higher-intensity information treatment (T2). Assignment balance is reported in Table 1.

Table 1: Experimental Balance by Assignment Arm

| Variable | Control | SD | T1 | SD | T2 | SD | T1 - C | T2 - C |
|---|---|---|---|---|---|---|---|---|
| Observations | 213 | | 212 | | 220 | | | |
| Female | 0.662 | 0.474 | 0.637 | 0.482 | 0.668 | 0.472 | -0.025 | 0.006 |
| First-Gen | 0.230 | 0.422 | 0.292 | 0.456 | 0.236 | 0.426 | 0.062 | 0.006 |
| Transfer | 0.178 | 0.384 | 0.184 | 0.388 | 0.232 | 0.423 | 0.006 | 0.053 |
| Asian | 0.695 | 0.462 | 0.670 | 0.471 | 0.668 | 0.472 | -0.025 | -0.027 |
| Black | 0.014 | 0.118 | 0.014 | 0.118 | 0.014 | 0.116 | 0.000 | -0.000 |
| White | 0.254 | 0.436 | 0.278 | 0.449 | 0.282 | 0.451 | 0.025 | 0.028 |

*Notes:* This table shows aggregate means and standard deviations by assignment arm. Differences are treatment-arm mean minus control mean. Stars indicate significance levels: ${}^{*}p < 0.10$, ${}^{**}p < 0.05$, ${}^{***}p < 0.01$.

Students in the control group did not receive any informational intervention and were simply asked to report their beliefs about effective strategies to secure a job. The first treatment group (T1) received information emphasizing networking as a strategy to secure employment, designed to increase the salience of networking in their consideration set. The second treatment group (T2) received more explicit and detailed information about the returns and practical importance of networking, including statements such as "Engaging with alumni can give you a competitive edge and help you get noticed in a crowded job market," and "The best candidates often come from referrals." We also note evidence that some elite firms require cold networking.[11]

After baseline measurement, random assignment, and delivery of the assigned information, our primary outcome is take-up of a networking service offered at a posted price. We asked students whether they were willing to pay $5, deducted from their potential gift card compensation, for a service providing five alumni contacts in their desired career field. This measure serves as a behavioral proxy for students' valuation of cold networking

[11] This information comes from qualitative interviews with hiring managers and from the literature (e.g., Rivera, 2016).

as a career success strategy. We use cold networking for this experiment for two reasons. First, for the careers of the students in the sample it is a necessary condition for finding a job, therefore it may be undervalued by students who are not familiar with the inner workings of such high-skilled and high-paying jobs. This highlights the subjective beliefs channel[12]. Second, disadvantaged students may only think of networking as an activity performed by those with existing connections, for example a family member or a friend who introduces them to someone. As a result, first-generation students may not know that cold networking is an acceptable strategy, which highlights our second mechanism.

We designed the two treatment groups with the idea of teasing out the two mechanisms we introduced in the model. The light-touch information intervention (T1) is designed to mainly load on the consideration-set mechanism, while the higher-intensity information intervention (T2) loads on both mechanisms. However, we cannot rule out that either mechanism is at play in both T1 and T2 because by implying that a strategy is available we may also be updating beliefs on returns, and by heavily updating beliefs on returns through T2 we are increasing the likelihood of the action being in the consideration set.

Finally, we highlight some considerations of the current design. First, even though we do not explicitly use the word networking in the control group, offering access to five contacts could induce students to think of networking, therefore including it in their consideration set. Therefore, the treatment effect we find is a lower bound on the total effect of not having an action in the consideration set, and it mainly identifies the difference between knowing and not knowing that cold networking is an acceptable strategy.

Second, UC Berkeley students are likely to be positively selected versus the general population. It is therefore reasonable to assume that first-generation students at UC Berkeley are more informed than the average or marginal first-generation student in U.S. colleges.

## 4.2 Experimental Results

We first document substantial baseline differences in networking valuation across student types.[13] We document differences in take-up of the networking service by first-generation background in the control group. First-generation students take up the service at a rate of

[12]The fact that this is a requirement in these high-skill, high-paying jobs also eliminates ambiguity around the optimal approach. The action itself is optimal and the only consideration the student has to make is whether their outside option is better or not.

[13]Throughout this section we present OLS linear-probability results for ease of discussion. However, the results are robust to different specifications, including Probit and Logit.

approximately 28 percent, compared with 40 percent for continuing-generation students.

These differences alone could reflect credit or preference constraints among first-generation students. We now look at whether treatment had any effect in order to identify the causal impact of information.

Figure 4: Raw Take-Up Rates and Adjusted Differential Information Effects

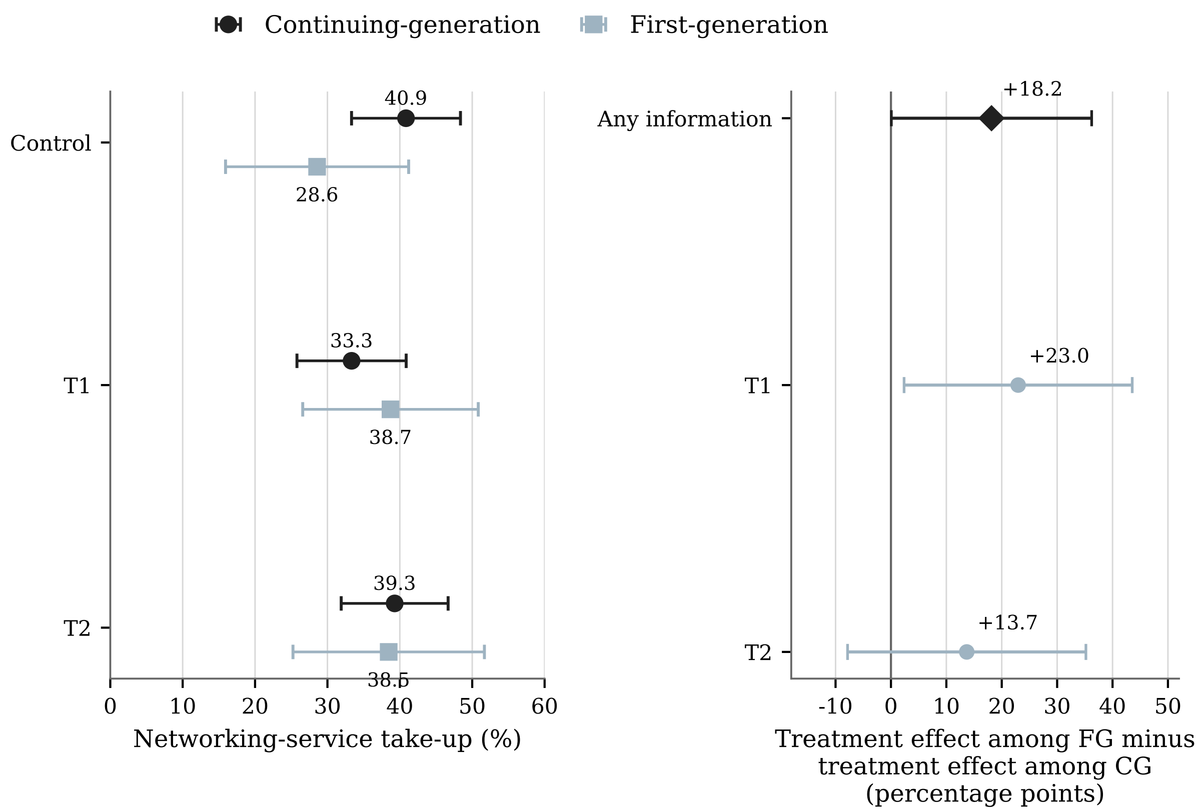


*Notes:* Panel A reports unadjusted take-up rates by randomized assignment arm and first-generation status. Panel B reports covariate-adjusted differences in treatment effects; positive values indicate a more positive response among first-generation students. The pooled-information comparison uses pooled-treatment model; T1 and T2 are secondary estimates from one joint arm-specific model. Horizontal bars show 95 percent confidence intervals. Controls include race indicators, GPA-category dummies, internship indicator, career-category dummies, female status, and transfer-student status. Cell sizes are: Control, 164 continuing-generation and 49 first-generation students; T1, 150 continuing-generation and 62 first-generation students; and T2, 168 continuing-generation and 52 first-generation students.

Figure 4 presents our first causal tests. Panel A reports unadjusted take-up rates by randomized assignment arm and first-generation status. Panel B reports covariate-adjusted differences in treatment effects. Among continuing-generation students, the estimated effect of any treatment is −5.3 percentage points and is statistically indistinguishable from zero. In contrast, the interaction between any treatment and first-generation status is 18.2 percentage points in our preferred specification and is statistically significant at the 5 percent level. The point estimates therefore imply that information raises take-up among first-generation students by 12.9 percentage points. Relative to the adjusted 16.7 percentage point control-group gap, this corresponds to closing approximately 77 percent of the gap through the first-generation treatment response alone. Thus, information largely reduces the baseline difference in demand for networking assistance.

We now move on to each treatment individually as can be seen again in Figure 4. We start with the effects of the light-touch information treatment (T1). As a reminder, the light-touch treatment emphasized networking as an available strategy, hence mainly loading on the consideration-set mechanism.

For first-generation students, T1 increased take-up by 14.1 percentage points in the controls specification, closing 92.5% of the control gap. The detailed pairwise T1 regression appears in Appendix Table A.7, while the main comparison across arms uses the direct-contrast test in Table 2.

Table 2: Direct Tests Comparing Treatment 1 and Treatment 2

| Contrast | No controls | Controls | N |
|---|---|---|---|
| T1 minus T2, continuing-generation | -0.060<br>(0.054) | -0.071<br>(0.054) | 645 |
| T1 minus T2, first-generation | 0.002<br>(0.092) | 0.022<br>(0.095) | 645 |
| T1 interaction minus T2 interaction | 0.062<br>(0.107) | 0.093<br>(0.109) | 645 |
| Controls | No | Yes | |

*Notes:* Contrasts are direct linear combinations from the joint T1/T2 model: T1 minus T2 for continuing-generation students, T1 minus T2 for first-generation students, and the T1 interaction minus the T2 interaction. Controls include race indicators, GPA-category dummies, internship indicator, career-category dummies, female status, and transfer-student status. Standard errors are robust and reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

As expected from our discussion of baseline means, the coefficient on first-generation students is negative, economically large, and statistically significant. The interaction term between T1 and first-generation status is positive, statistically significant, and large enough to decrease baseline differences, such that the two groups are now not statistically different from each other.

This result is consistent with a consideration-set channel playing an important role in the baseline first-generation gap in demand for the networking service.

The higher-intensity treatment (T2), which provided explicit information about networking returns and their necessity for competitive positions, generated a different pattern of responses. Appendix Table A.8 presents the results from the same specification used for T1. We still find a positive interaction for first-generation students, 0.136 in the controls specification. However, this estimate is imprecise and not statistically significant.

One possible explanation is that the more detailed message may lead students who were unaware of the importance of the action to conclude that it is too late to act.

To test the robustness of our findings we compare our results from our first-generation students to those of other groups. In particular we look at two types of heterogeneity: i) groups that should not have this specific form of informational friction but that are usually disadvantaged in higher education, and ii) a group that instead could potentially suffer from the hidden curriculum. Figure 5 shows results for our population of interest and the two comparison populations.

Figure 5: Exploratory Heterogeneity in Differential Information Effects

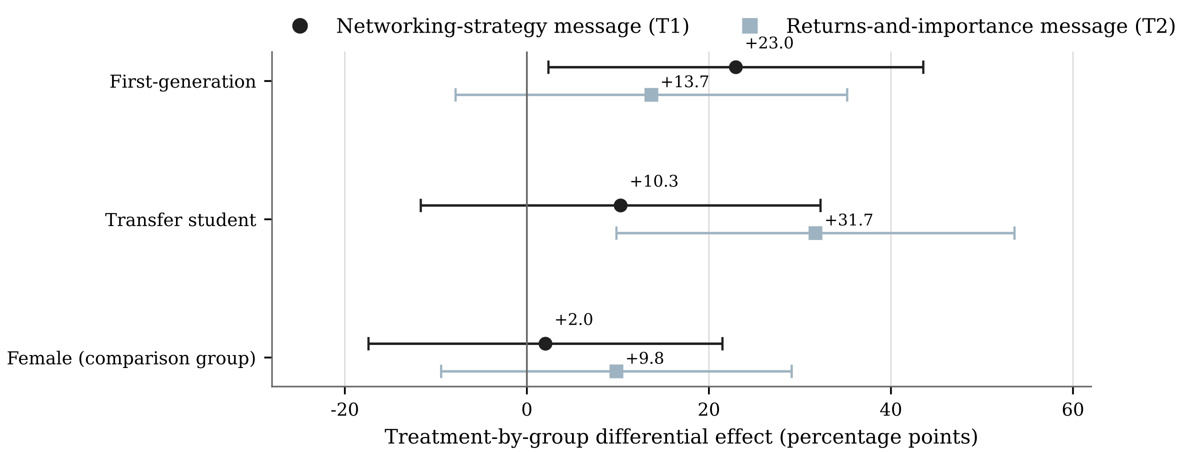


*Notes:* Estimates come from parallel all-arm joint models using the same full sample, control group, covariate set, and interaction structure. Horizontal bars show 95 percent confidence intervals based on robust standard errors. The arm-specific patterns are exploratory and do not establish distinct mechanisms or sharply rank T1 and T2 across groups.

For the "traditional" groups we look at gender.[14] We do not find any significant effect

[14]For ethnic minorities we only have enough power to look at Asian-Americans. We do not see any difference there either.

in gender for either of the two treatments.

For the "other" group that may be affected by the hidden curriculum we look at transfer students from two-year community colleges. These individuals may have acquired the insider knowledge of the hidden curriculum specific to their previous institutions. However, moving to a highly ranked four-year college implies working with a completely new set of rules in the hidden curriculum.

For transfer students we find small treatment effects from T1. As reported in Appendix Table A.9 these differences are not statistically significant. On the other hand we find that T2 has large and statistically significant effects on take-up for transfer students: the transfer by T2 interaction is 0.307 in the controls specification. This is suggestive evidence that for transfer students the information friction stems largely from lower subjective beliefs.

The experimental results provide several insights into the mechanisms underlying differences in willingness to invest in hidden-curriculum actions. Baseline differences in networking valuation and positive treatment interactions are consistent with information frictions contributing to hidden-curriculum gaps, while the larger T1 response for first-generation students points toward availability or consideration-set information as a central margin in this setting.

Comparing the first-generation results with the transfer-student results shows that each group faces a different friction, with transfer students showing signs of lower subjective beliefs about the returns to these actions. The direct T1-versus-T2 contrast for first-generation students is imprecise, so this heterogeneity should be read as suggestive mechanism evidence rather than a sharp ranking of the two information designs.

These results suggest that targeted interventions should account for students' existing knowledge and institutional experience. When the friction is lack of awareness, information that makes an action visible may matter more than simply expanding access to on-campus and on-demand resources. When students have some prior exposure, more detailed information about returns may be necessary.

# 5 Endogenous Information and Learning Under Hidden Curriculum

In the previous sections, the observational evidence and the information experiment documented hidden-curriculum gaps and showed that information provision can shift demand for a concrete networking service. We now ask why this lack of information and perception of low returns persists, despite the seeming abundance of information and resources available to students at universities and online. To answer this question we design a novel experiment that uses AI tools to observe how students search for and learn from information.

Studying information acquisition using interactive AI-powered chatbots allows us to create a stigma-free, low-stakes environment where students can ask questions and search for information free from human judgment. This reduces potential barriers to asking for information or advice on sensitive topics and reduces concerns about signaling lower ability. Additionally, the use of a chatbot allows us to track the path-dependent, sequential nature of information acquisition. Similar to real-life information acquisition, the dialogue with the LLM requires students to actively formulate their questions, process multiple complex and sometimes ambiguous responses, and make a conscious decision on the next question to ask and the next search step. Thus, the AI tool closely mirrors real-world information acquisition.

Given the rise in AI-powered career information and education tools, we believe this is a realistic proxy for how students increasingly search for information.

Finally, because the mechanisms driving the hidden curriculum differ across groups, a scalable policy tool must adapt the information it provides to the group[15]. AI with its ability to adapt to the needs of the student and answer in a natural manner a variety of questions has the potential of being the scalable tool that universities and policymakers can use.

[15] The results on first-generation students alone would point to an awareness campaign as the best intervention. However, if then launched at scale for all disadvantaged groups we would run into the problems discussed in Al-Ubaydli et al. (2017). In particular we showed how each group may have unique needs, and this would create a drop in returns at scale.

## 5.1 Experimental Design

We recruit U.S. college students using Prolific, an online survey platform. Students first answer some basic demographic questions and a short survey. Students are then asked to "beta test" a new AI tool designed to help college students like them succeed in college and their careers. The primary nudge-level analysis uses 993 AI nudges from 478 conversations.[16] The experiment is pre-registered in the AEA RCT Registry.

We randomize students into either a treatment or a control AI college advisor. The Treatment AI tool, or Active AI, is designed to answer questions that the students ask, while *actively* nudging students to consider hidden-curriculum actions (e.g., professional networking, leveraging office hours, obtaining mentorship). The Control AI, or Passive AI, is designed to mainly answer with topics directly related to the question the student has asked. Thus, the Passive AI can nudge the student, but does so to a much lesser degree than the Active AI. Apart from the likelihood of receiving a hidden-curriculum nudge the two AI tools also differ in how the nudge is presented. In the treatment tool the nudges are much more central to the response, while in the control tool nudges can sometimes be embedded within a long response.

The Active AI is designed to add hidden-curriculum actions to students' consideration sets. The design mimics the effects in our information experiment with the added benefit that this is a targeted and flexible intervention: the information shock should be adaptive to the needs of the students. Whether it is depends on how students respond. Students may not trust the AI, they may not know how to interpret the information provided, or they may be confident enough in their own priors to ignore any nudge.

The Passive AI is designed to understand what students' search efforts would look like if they did not receive active guidance and had to decide their own path of discovery. By the nature of AI and by design we allow some variance in responses, which lets us see how students respond to nudges when their consideration set has not been changed as it has for students in the treatment arm.

We create the two AI tools using an AI education platform. Both tools are powered by Anthropic's LLM. We design the AI tools with low variability and largely restrict the information provided to students to the information we explicitly feed into the AI tools. This gives us a large amount of control over the information provided, allowing us to document differences in search and topic continuation while holding the information

[16] The difference between the conversation-level and nudge-level samples comes from the requirement that the nudge-level analysis observe an AI nudge and a next substantive user message.

itself largely fixed. We then use the transcripts from the chat history to quantify students' consideration sets and continuation efforts as they search for academic and career information.

The key outcomes of interest include:

*Initial Consideration Sets and AI Treatment Effect:* We analyze the share of student-initiated messages focused on hidden-curriculum topics (or the User HC Rate). This measures the student's starting consideration set and their willingness to seek information about hidden-curriculum actions. We compare the incidence of HC topics between the Active AI Tool and the Passive AI Tool to measure the treatment effect of expanding the consideration set.

*Differences in Search Behavior:* We use LLM text analysis to identify patterns in how first-generation and continuing-generation students ask questions and search for information. Specifically, we are interested in the specificity of questions, the ability of the students to guide the conversation and extract relevant, actionable information from the LLM. We also document differences in language related to confidence and uncertainty.

*Nudges and Learning:* To better understand how learning happens, we track how often students "follow-up" on AI-initiated "nudges" on HC topics, which we use as a proxy for a student's willingness to explore HC topics. We call this the HC Follow-Up Rate. A student has "followed-up" on a nudge if the student has asked a follow-up question or otherwise continued to discuss the HC topic following the AI Tool's nudge. Combined with the baseline rate of HC topic discussion, this metric gives us a sense of how much the student is willing to engage with and learn about an HC topic.

## 5.2 Results

We first ask whether first-generation students search less for hidden-curriculum information when the choice of topics is theirs. They do not. In the Passive AI condition, first-generation students devote 18.0 percent of their own messages to hidden-curriculum topics, compared with 18.5 percent for continuing-generation students; in conversation-level regressions with standard errors clustered by participant, the first-generation coefficient is small and imprecise (Appendix Table A.10). The Active AI raises user hidden-curriculum message shares by roughly 25 percentage points for both groups, to 43.9 percent for first-generation students and 46.4 percent for continuing-generation students, with no differential response by first-generation status. Search effort, measured by the share of messages initiated by the student on hidden-curriculum topics, is therefore not the margin on which the two groups differ.

The share of student-initiated messages on hidden-curriculum topics does not by itself identify what a student knows. A student who already knows the hidden-curriculum actions has no reason to ask about them, and a student who does not know they exist cannot ask about them. Both students initiate no hidden-curriculum conversation, so equal shares across the two groups do not imply equal information. Two other margins do separate the groups: the content of the questions students ask, and whether they continue a hidden-curriculum topic after the AI raises it.

We next code the content of the questions themselves. Figure 6 summarizes qualitative question-content outcomes: hidden-curriculum strategy questions[17], social-capital orientation, and permission or legitimacy language. In the Passive AI condition, first-generation students ask fewer hidden-curriculum strategy questions than continuing-generation students (11.5% versus 14.5%) and show less social-capital orientation (4.5% versus 9.0%), while permission or legitimacy language is somewhat more common among first-generation students (8.0% versus 5.0%).

[17] A message is specifically coded as an HC strategy question if it is a question and includes an actionable HC strategy. If the student mentioned a networking tool but did not ask for information on how to network or implement an action, then it would not be coded here.

Figure 6: Student question quality by group

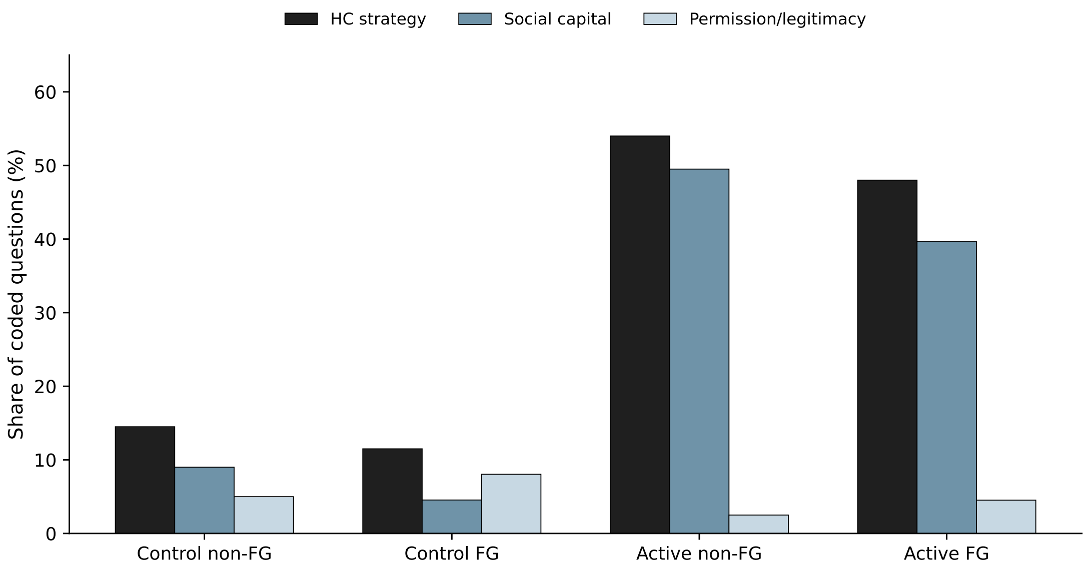


*Notes:* The figure reports aggregate shares from blinded qualitative coding from LLM of redacted student questions. The plotted outcomes are hidden-curriculum strategy questions, social-capital orientation, and permission/legitimacy seeking.

The Active AI condition moves both groups toward more hidden-curriculum strategy questions and more social-capital-oriented questions. Hidden-curriculum strategy questions go from 11.5% to 48.0% for first-generation students and from 14.5% to 54.0% for continuing-generation students. Social-capital questions go from 4.5% to 39.7% for first-generation students and from 9.0% to 49.5% for continuing-generation students. These patterns are consistent with active guidance moving the content of student search, but on their own they do not show whether the students have fully internalized the hidden-curriculum strategies.

We produce these question-content measures using LLM text analysis of the redacted transcripts.

We also find that first-generation students often frame their questions in a way that suggests social risk or uncertainty. In the passive condition, social-risk concerns appear in 20.0% of coded first-generation messages compared to 4.5% for continuing-generation students. These patterns suggest that hidden-curriculum search is not only about topic availability but also about whether a strategy feels legitimate and actionable.

The higher prevalence of permission or legitimacy language and social-risk concerns among first-generation students is suggestive of greater self-doubt when engaging with hidden-curriculum information. This may make first-generation students more reticent to update their beliefs in response to new information than the aggregate topic shares alone suggest.

Figure 7 reports next-turn topic-continuation levels by AI arm, first-generation status, and whether the AI nudge concerns a hidden-curriculum topic. Hidden-curriculum continuation is lower for first-generation students in the Passive AI condition, the first-generation difference is much smaller for non-hidden-curriculum nudges, and hidden-curriculum continuation is higher under Active AI. Active AI therefore shifts next-turn continuation toward hidden-curriculum content rather than raising continuation uniformly across all topics.

Figure 7: Adjusted Next-Turn Topic Continuation by Nudge Type

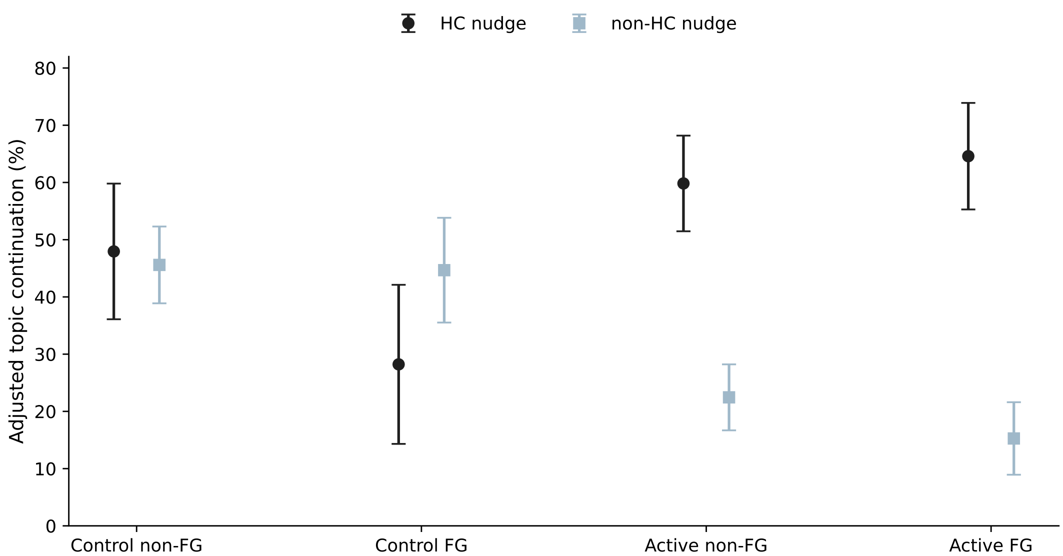


*Notes:* The figure reports adjusted next-turn topic-continuation rates by treatment arm, first-generation status, and nudge type. Estimates control for age, sex indicators, and a White/non-White indicator.

This evidence matters because the AI can introduce a hidden-curriculum topic without the student choosing to continue it. Next-turn continuation therefore captures a distinct margin from topic exposure or the overall share of hidden-curriculum messages.

While the Active AI raised hidden-curriculum content for both groups, the two groups differ sharply in whether they continue a hidden-curriculum topic once the AI introduces it.

Figure 8 subtracts adjusted non-hidden-curriculum continuation from adjusted hidden-curriculum continuation within each group and AI arm. The relative hidden-curriculum continuation differences are +2.4 percentage points for continuing-generation students in Passive AI, -16.5 percentage points for first-generation students in Passive AI, +37.4 percentage points for continuing-generation students in Active AI, and +49.3 percentage points for first-generation students in Active AI.

Figure 8: Relative next-turn topic continuation on HC versus non-HC nudges

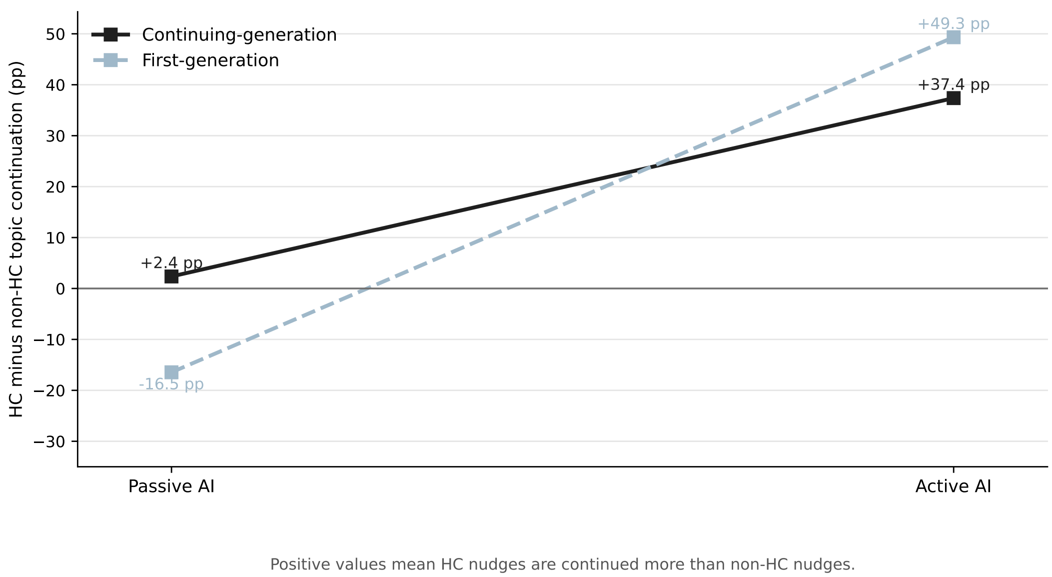


*Notes:* The figure reports adjusted HC topic-continuation rates minus adjusted non-HC topic-continuation rates, in percentage points. Positive values mean HC nudges are continued more than non-HC nudges. Estimates use the same adjusted nudge-level data as the adjusted topic-continuation figure, controlling for age, sex indicators, and a White/non-White indicator.

Table 3: Nudge topic continuation and hidden-curriculum content

| | (1) Follow-Up | (2) Follow-Up |
|---|---|---|
| First-gen | -0.001 | -0.009 |
| | (0.058) | (0.058) |
| Active AI | -0.228*** | -0.232*** |
| | (0.045) | (0.045) |
| HC nudge | 0.025 | 0.018 |
| | (0.074) | (0.073) |
| First-gen × Active AI | -0.063 | -0.063 |
| | (0.073) | (0.073) |
| First-gen × HC nudge | -0.172 | -0.176 |
| | (0.112) | (0.112) |
| Active AI × HC nudge | 0.346*** | 0.353*** |
| | (0.093) | (0.092) |
| First-gen × Active AI × HC nudge | 0.290** | 0.294** |
| | (0.137) | (0.136) |
| N nudges | 993 | 993 |
| N conversations | 478 | 478 |
| N participants | 456 | 456 |
| Controls | no | yes |

*Notes:* Unit is AI nudge. Outcome is next-turn topic continuation, equal to one if the next user message is assigned the same topic as the AI nudge. Controls are age, sex indicators, and a White/non-White indicator. Standard errors are clustered by participant. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The nudge-level regression in Table 3 formalizes the same pattern. In the controls specification, the triple interaction between first-generation status, Active AI, and hidden-curriculum nudge is 0.294, or 29.4 percentage points. We interpret this as evidence that active guidance changes next-turn continuation for hidden-curriculum content, not as evidence that the AI directly measures learning or fully closes hidden-curriculum gaps. However, when first-generation students are less likely to devote their next turn to a hidden-curriculum topic, they are also less likely to invest additional search effort in that topic. Repeated differences in responsiveness can therefore generate lower cumulative engagement and a risk of lower learning.

Finally, the heterogeneity in mechanisms driving the hidden curriculum motivates studying flexible tools that can adapt information to the student's question and level of prior exposure. AI may be one scalable tool for this task because it can answer a variety of student questions in a natural way, but the evidence here is about search behavior and next-turn continuation rather than a complete policy evaluation[18].

The differences in hidden-curriculum continuation between Passive and Active AI are consistent with several frictions operating after information is introduced: processing costs, uncertainty about how to use the information, lower starting confidence, skepticism about AI advice, or stronger priors about which strategies are worth pursuing. This interpretation connects the AI evidence to the awareness and subjective-return margins in the model, but it should be read as mechanism evidence rather than a direct elicitation of those margins.

First-generation and continuing-generation students devote similar shares of their own messages to hidden-curriculum topics, so the two groups do not differ in search effort. They differ on two other margins: first-generation students ask fewer hidden-curriculum strategy questions and use permission and legitimacy language more often, and under the Passive AI they are less likely to continue a hidden-curriculum topic in the turn after the AI raises it. Hidden-curriculum gaps can therefore persist even when the information is available to both groups.

# 6 Conclusion

This paper reveals how the hidden curriculum perpetuates inequality even when formal barriers to opportunity have been removed. Our evidence shows that unwritten rules and informal strategies create systematic disadvantages for first-generation students that compound over time, contributing to persistent gaps in economic and social outcomes.

Our findings challenge conventional explanations for inequality. The differences we document between first-generation and continuing-generation students cannot be fully explained by preferences, constraints, or ability. Instead, they stem from differential access to tacit knowledge about how institutions actually work. Advantaged students can inherit this knowledge through family networks, while disadvantaged students must discover it through costly trial and error, if at all.

Three aspects of our results are particularly striking for understanding inequality. First, the hidden curriculum creates a perverse dynamic in which those who work hardest may achieve least. First-generation students compensate for their lack of informal strategies

[18]The results on first-generation students would have pointed to an awareness campaign as being the best intervention. However, if then launched at scale for all disadvantaged groups we would run into problems of implementation and targeting. In particular we showed how each group may have unique needs, and this would create a drop in returns at scale.

by overinvesting in formal requirements, such as doing more readings and presentations, and yet still achieve lower GPAs and fewer internships. This mismatch between effort and reward violates basic notions of meritocracy and fairness.

Second, our experiments show that these inequalities are easy to perpetuate yet simple to address. A brief information intervention closed roughly 77 percent of the first-generation gap in take-up of a paid networking service. The problem is not that disadvantaged students lack ability or motivation, but that success requires knowledge of rules that are never explicitly stated.

Third, the AI experiment shows that the difficulty is not in finding information. First-generation students devote the same share of their own messages to hidden-curriculum topics as continuing-generation students, but they ask different questions, express more uncertainty, and are less likely to continue a hidden-curriculum topic once the AI raises it. This creates an inequality trap in which initial disadvantages in cultural capital translate into persistent differences in information acquisition.

Our experiments do not allow us to measure whether closing these gaps in hidden-curriculum actions translates into improved long-run educational or labor market outcomes. Such effects may take years to materialize and depend on the accumulation of many subsequent choices. Nevertheless, our findings identify an important intermediate margin: first-generation students are less likely to invest in actions that can shape educational and career opportunities, and information about these actions and their returns substantially narrows these gaps. To the extent that such actions are inputs into later educational and career decisions, closing these intermediate gaps is a necessary step toward closing the longer-run gaps that motivate our analysis.

The implications extend far beyond individual students. When capable first-generation students systematically miss opportunities due to information they were never given, society loses potential innovations and diverse perspectives in highly consequential settings.

The path forward requires recognizing that equal access may fail to generate equal opportunity and therefore equal outcomes when success depends on unwritten rules. Institutions serious about equity must make the hidden curriculum visible. However, our results on heterogeneity of mechanisms suggest that each disadvantaged group would need its own tailored intervention. For example, groups affected by consideration-set issues will not be moved by opportunities to engage with on-campus offices and resources, since they are not aware of their own lack of knowledge.

Future work should examine how hidden curricula operate in other domains: workplace advancement, entrepreneurship, political participation. Understanding these mechanisms across contexts could inform broader efforts to create genuinely equal opportunity. The hidden curriculum represents a fundamental challenge to meritocratic ideals. By documenting its existence and effects, we hope to spark conversation about the gap between formal equality and substantive opportunity. Only by making the invisible visible can we begin to address the subtle mechanisms that perpetuate inequality across generations.

## Acknowledgments

We are grateful to Ned Augenblick, Alec Brandon, Leonardo Bursztyn, Stefano DellaVigna, Brent Hickman, Justin Holz, Alex Imas, John A. List, Gautam Rao, and Andrew Simon for helpful comments. We are also grateful to Igor Chirikov for facilitating access to the SERU Consortium Data and for providing context.

IRB of record: University of Chicago Protocol #IRB24-1552; UC Berkeley Protocol #2024-09-17838 (Experiment); UC Berkeley Protocol #2025-07-18799 (UCES Data); AEA Trial Registration #AEARCTR-0014970; AEA Trial Registration #AEARCTR-0016671.

# REFERENCES


**Al-Ubaydli, Omar, John A List, and Dana L Suskind**, "What can we learn from experiments? Understanding the threats to the scalability of experimental results," *American Economic Review*, 2017, *107* (5), 282–286.

**Altmann, Steffen, Armin Falk, Simon Jäger, and Florian Zimmermann**, "Learning about job search: A field experiment with job seekers in Germany," *Journal of Public Economics*, 2018, *164*, 33–49.

**Angrist, Joshua, David Autor, and Amanda Pallais**, "Marginal Effects of Merit Aid for Low-Income Students," *The Quarterly Journal of Economics*, April 2022, *137* (2), 1039–1090.

**Arellano-Bover, Jaime, Carolina Bussotti, John M Nunley, and R Alan Seals**, "Unbundling the effects of college on first-job search: Returns to majors, minors, and extracurriculars," *Labour Economics*, 2026, p. 102846.

**Aucejo, Esteban M, Jacob French, Paola Ugalde Araya, and Basit Zafar**, "Understanding gaps in college outcomes by first-generation status," Technical Report, National Bureau of Economic Research 2025.

**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*, 2018, *65*, 18–30.

**Becker, Gary S. and Nigel Tomes**, "An equilibrium theory of the distribution of income and intergenerational mobility," *Journal of Political Economy*, 1979, *87* (6), 1153–1189.

**Belot, Michele, Philipp Kircher, and Paul Muller**, "Providing advice to jobseekers at low cost: An experimental study on online advice," *The review of economic studies*, 2019, *86* (4), 1411–1447.

**Bhargava, Saurabh and Dayanand Manoli**, "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment," *American Economic Review*, November 2015, *105* (11), 3489–3529.

**Bourdieu, Pierre**, "The Forms of Capital," in John G. Richardson, ed., *Handbook of Theory and Research for the Sociology of Education*, New York: Greenwood, 1986, pp. 241–258.

**Bursztyn, Leonardo, Alessandra L González, and David Yanagizawa-Drott**, "Misperceived social norms: Women working outside the home in Saudi Arabia," *American economic review*, 2020, *110* (10), 2997–3029.

**Calarco, Jessica McCrory**, ""I need help!"" social class and children's help-seeking in elementary school," *American Sociological Review*, 2011, *76* (6), 862–882.

**Canaan, Serena, Antoine Deeb, and Pierre Mouganie**, "Adviser value added and student outcomes: Evidence from randomly assigned college advisers," *American Economic Journal: Economic Policy*, 2022, *14* (4), 151–191.

**Chen, Xianglei**, "First Generation Students in Postsecondary Education: A Look at Their College Transcripts," Technical Report NCES 2005-171, U.S. Department of Education, National Center for Education Statistics, Washington, DC 2005.

**Chetty, Raj, John N. Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan**, "Mobility report cards: The role of colleges in intergenerational mobility," NBER Working Paper 23618, National Bureau of Economic Research 2017.

**_ , Nathaniel Hendren, Maggie R Jones, and Sonya R Porter**, "Race and economic opportunity in the United States: An intergenerational perspective," *The Quarterly Journal of Economics*, 2020, *135* (2), 711–783.

**Conlon, John J**, "Major malfunction: A field experiment correcting undergraduates' beliefs about salaries," *Journal of Human Resources*, 2021, *56* (3), 922–939.

**DellaVigna, Stefano**, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, May 2009, *47* (2), 315–372.

**Dizon-Ross, Rebecca**, "Parents' Beliefs about Their Children's Academic Ability: Implications for Educational Investments," *American Economic Review*, August 2019, *109* (8), 2728–2765.

**Dynarski, Susan M. and Judith Scott-Clayton**, "Financial aid policy: Lessons from research," *The Future of Children*, 2013, *23* (1), 67–91.

**Finkelstein, Amy and Matthew J. Notowidigdo**, "Take-Up and Targeting: Experimental Evidence from SNAP," *The Quarterly Journal of Economics*, 2019, *134* (3), 1505–1556.

**Gabaix, Xavier**, "Behavioral inattention," in "Handbook of behavioral economics: Applications and foundations 1," Vol. 2, Elsevier, 2019, pp. 261–343.

**Hoekstra, Mark**, "The effect of attending the flagship state university on earnings: A discontinuity-based approach," *The review of economics and statistics*, 2009, *91* (4), 717–724.

**Jack, Anthony Abraham**, "(No) harm in asking: Class, acquired cultural capital, and academic engagement at an elite university," *Sociology of Education*, 2016, *89* (1), 1–19.

_ , *The Privileged Poor: How Elite Colleges Are Failing Disadvantaged Students*, Cambridge, MA: Harvard University Press, 2019.

**Jackson, Philip W.**, *Life in Classrooms*, New York: Holt, Rinehart and Winston, 1968.

**Jensen, Robert**, "The (perceived) returns to education and the demand for schooling," *The Quarterly Journal of Economics*, 2010, *125* (2), 515–548.

**Kim, Young K and Linda J Sax**, "Student–faculty interaction in research universities: Differences by student gender, race, social class, and first-generation status," *Research in Higher Education*, 2009, *50* (5), 437–459.

**Manzini, Paola and Marco Mariotti**, "Stochastic Choice and Consideration Sets," *Econometrica*, 2014, *82* (3).

**Marmaros, David and Bruce Sacerdote**, "Peer and social networks in job search," *European economic review*, 2002, *46* (4-5), 870–879.

**Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Y Ozbay**, "Revealed Attention," *American Economic Review*, August 2012, *102* (5), 2183–2205.

**Michelman, Valerie, Joseph Price, and Seth D Zimmerman**, "Old boys' clubs and upward mobility among the educational elite," *The Quarterly Journal of Economics*, 2022, *137* (2), 845–909.

**Ost, Ben, Weixiang Pan, and Douglas Webber**, "College Networks: The Importance of Employer Connections," *Journal of Labor Economics*, 2026, *44* (2), 351–381.

**Rivera, Lauren A**, "Pedigree: How elite students get elite jobs," 2016.

**Solon, Gary**, "Intergenerational mobility in the labor market," in "Handbook of labor economics," Vol. 3, Elsevier, 1999, pp. 1761–1800.

**Stansbury, Anna and Kyra Rodriguez**, "The class gap in career progression: Evidence from us academia," *Econometrica*, 2026, *94* (4), 1345–1373.

**Stuber, Jenny M**, "Class, culture, and participation in the collegiate extra-curriculum," in "Sociological forum," Vol. 24 Wiley Online Library 2009, pp. 877–900.

**Wiswall, M. and B. Zafar**, "Determinants of College Major Choice: Identification using an Information Experiment," *The Review of Economic Studies*, April 2015, *82* (2), 791–824.

**Zafar, Basit**, "How do college students form expectations?," *Journal of Labor Economics*, 2011, *29* (2), 301–348.

**Zimmerman, Seth D.**, "Elite Colleges and Upward Mobility to Top Jobs and Top Incomes," *American Economic Review*, January 2019, *109* (1), 1–47.

# A Appendix

## A.1 Additional Tables and Figures

Table A.1: Descriptive Balance by First-Generation Status

| Variable | First-Gen Mean (SD) | Continuing-Gen Mean (SD) | Difference |
|---|---|---|---|
| Female | 0.658 (0.474) | 0.612 (0.487) | 0.047*** |
| In-state | 0.871 (0.335) | 0.786 (0.410) | 0.085*** |
| U.S. national | 0.931 (0.254) | 0.921 (0.269) | 0.009*** |
| Low income or poor | 0.438 (0.496) | 0.066 (0.248) | 0.372*** |
| Working class | 0.368 (0.482) | 0.158 (0.365) | 0.209*** |
| Middle class | 0.166 (0.372) | 0.454 (0.498) | -0.288*** |
| Upper-middle class | 0.026 (0.159) | 0.293 (0.455) | -0.267*** |
| Wealthy | 0.003 (0.056) | 0.029 (0.169) | -0.026*** |
| Asian | 0.243 (0.429) | 0.252 (0.434) | -0.009** |
| Black | 0.021 (0.144) | 0.026 (0.159) | -0.005*** |
| Chicano/Latino | 0.396 (0.489) | 0.071 (0.256) | 0.325*** |
| Latino | 0.067 (0.250) | 0.027 (0.163) | 0.039*** |
| Pacific Islander | 0.010 (0.100) | 0.036 (0.186) | -0.026*** |
| White | 0.059 (0.236) | 0.158 (0.365) | -0.099*** |

*Notes:* This table shows means, standard deviations in parentheses, and differences across first-generation and continuing-generation students in the common main-text sample. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.2: First-Generation Gaps in Opportunity Portfolio Indices

| | HC index raw gap | HC index controls gap | FC index raw gap | FC index controls gap | Placebo index raw gap | Placebo index controls gap |
|---|---|---|---|---|---|---|
| First-Generation | -0.191*** | -0.053*** | 0.048*** | 0.046*** | -0.052*** | -0.002 |
| | (0.006) | (0.006) | (0.005) | (0.006) | (0.006) | (0.008) |
| Observations | 73,989 | 73,989 | 73,989 | 73,989 | 73,989 | 73,989 |
| Controls | No | Yes | No | Yes | No | Yes |
| University FE | No | Yes | No | Yes | No | Yes |
| Major FE | No | Yes | No | Yes | No | Yes |
| Year-in-college control | No | Yes | No | Yes | No | Yes |

*Notes:* Shows results from OLS regressions of opportunity portfolio indices on first-generation status. Raw-gap columns report first-generation minus continuing-generation differences without controls. Controls-gap columns control for gender, race/ethnicity, social class, in-state status, U.S. national status, university fixed effects, major fixed effects, and year in college. Robust standard errors are reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

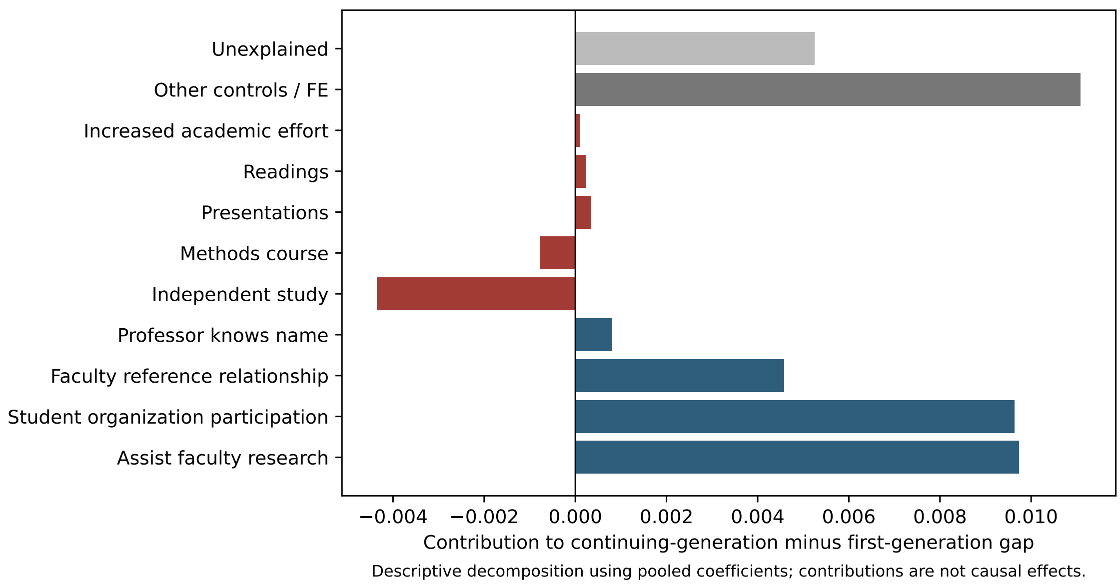


Figure A.1: Oaxaca-Blinder Decomposition for For-Credit Internships

*Notes:* The figure decomposes the continuing-generation minus first-generation outcome gap using pooled linear coefficients. Individual hidden and formal action components are shown separately, while other controls and fixed effects are grouped. The decomposition is descriptive.

Table A.3: Individual Action Gaps by Portfolio Family

| Action | Family | Raw | Preferred | N |
|---|---|---|---|---|
| Faculty reference | HC | -0.171*** | -0.029** | 73,989 |
| | | (0.010) | (0.012) | |
| Professor knows name | HC | -0.235*** | -0.041** | 73,989 |
| | | (0.015) | (0.017) | |
| Faculty research | HC | -0.047*** | -0.012*** | 73,989 |
| | | (0.003) | (0.004) | |
| Student organization | HC | -0.168*** | -0.063*** | 73,989 |
| | | (0.005) | (0.005) | |
| Readings | Formal curriculum | 0.078*** | 0.055*** | 73,989 |
| | | (0.011) | (0.014) | |
| Presentations | Formal curriculum | -0.027** | 0.060*** | 73,989 |
| | | (0.014) | (0.016) | |
| Methods course | Formal curriculum | 0.061*** | 0.015*** | 73,989 |
| | | (0.005) | (0.005) | |
| Independent study | Formal curriculum | 0.039*** | 0.033*** | 73,989 |
| | | (0.004) | (0.005) | |
| Academic effort | Formal curriculum | -0.030** | 0.046*** | 73,989 |
| | | (0.013) | (0.015) | |
| Study group | Placebo | -0.134*** | 0.040** | 73,989 |
| | | (0.015) | (0.018) | |
| Class discussion | Placebo | -0.183*** | -0.098*** | 73,989 |
| | | (0.012) | (0.014) | |
| Active participation | Placebo | 0.027*** | 0.010 | 73,989 |
| | | (0.010) | (0.012) | |
| Instructor outside class | Placebo | -0.007 | 0.052*** | 73,989 |
| | | (0.013) | (0.016) | |

*Notes:* This table shows results from OLS regressions of each action on first-generation status. Raw columns report first-generation minus continuing-generation differences. Preferred columns include controls for gender, race/ethnicity, social class, in-state status, U.S. national status, university fixed effects, major fixed effects, and approximate college year. Robust standard errors are reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.4: Opportunity Portfolio Indices, HS-GPA U.S.-National Sample

| | HC index raw gap | HC index controls gap | FC index raw gap | FC index controls gap |
|---|---|---|---|---|
| First-Generation | -0.239*** | -0.070*** | 0.012* | 0.010 |
| | (0.007) | (0.008) | (0.006) | (0.008) |
| Observations | 52,946 | 52,946 | 52,946 | 52,946 |
| Controls | No | Yes | No | Yes |
| University FE | No | Yes | No | Yes |
| Major FE | No | Yes | No | Yes |
| Year-in-college control | No | Yes | No | Yes |

*Notes:* The sample is restricted to U.S. national students with high-school GPA below 5. Controls include gender, race/ethnicity, social class, in-state status, high-school GPA, university fixed effects, major fixed effects, and year in college; U.S.-national status is omitted because it is constant in the restricted sample. All estimates are descriptive. Robust standard errors are reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.5: Outcome Gaps With and Without Portfolio Controls

| Specification | For-credit internship | Non-credit internship | GPA |
|---|---|---|---|
| Raw difference | -0.037*** | -0.092*** | -0.247*** |
| | (0.003) | (0.004) | (0.006) |
| Controls | -0.011*** | -0.035*** | -0.079*** |
| | (0.004) | (0.004) | (0.007) |
| Controls + FC index | -0.016*** | -0.038*** | -0.082*** |
| | (0.004) | (0.004) | (0.007) |
| Controls + HC index | -0.002 | -0.027*** | -0.070*** |
| | (0.004) | (0.004) | (0.007) |
| Observations | 73,989 | 73,989 | 73,989 |

*Notes:* This table shows results from OLS regressions of internships and GPA on first-generation status. The first row reports raw differences. The controls row controls for gender, race/ethnicity, social class, in-state status, U.S. national status, university fixed effects, major fixed effects, and year in college. The last two rows add the FC index and the HC index to the controls. Robust standard errors are reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.6: Outcome Gaps, HS-GPA U.S.-National Sample

| Specification | For-credit internship | Non-credit internship | GPA |
|---|---|---|---|
| Raw difference | -0.051*** | -0.107*** | -0.278*** |
| | (0.004) | (0.004) | (0.006) |
| Controls | -0.020*** | -0.040*** | -0.074*** |
| | (0.005) | (0.005) | (0.007) |
| Controls + FC index | -0.021*** | -0.041*** | -0.074*** |
| | (0.005) | (0.005) | (0.007) |
| Controls + HC index | -0.008* | -0.029*** | -0.064*** |
| | (0.005) | (0.005) | (0.007) |
| Observations | 52,946 | 52,946 | 52,946 |

*Notes:* The sample is restricted to U.S. national students with high-school GPA below 5. Controls include gender, race/ethnicity, social class, in-state status, high-school GPA, university fixed effects, major fixed effects, and year in college; U.S.-national status is omitted because it is constant in the restricted sample. All estimates are descriptive. Robust standard errors are reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.7: Treatment 1 Effects by First-Generation Status

| | No controls | Controls |
|---|---|---|
| Treatment 1 | -0.075 | -0.086 |
| | (0.055) | (0.055) |
| First-generation | -0.123 | -0.152** |
| | (0.075) | (0.076) |
| Treatment 1 x First-generation | 0.177* | 0.227** |
| | (0.105) | (0.105) |
| Controls | No | Yes |
| Observations | 425 | 425 |

*Notes:* The table reports OLS linear probability models on the Control vs T1 sample. Treatment 1 equals one for assignment to T1 and zero for Control. Controls include race indicators, GPA-category dummies, internship indicator, career-category dummies, female status, and transfer-student status. Standard errors are HC1 robust and reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.8: Treatment 2 Effects by First-Generation Status

| | No controls | Controls |
|---|---|---|
| Treatment 2 | -0.016 | -0.027 |
| | (0.054) | (0.054) |
| First-generation | -0.123 | -0.170** |
| | (0.075) | (0.078) |
| Treatment 2 x First-generation | 0.115 | 0.136 |
| | (0.108) | (0.110) |
| Controls | No | Yes |
| Observations | 433 | 433 |

*Notes:* The table reports OLS linear probability models on the Control vs T2 sample. Treatment 2 equals one for assignment to T2 and zero for Control. Controls include race indicators, GPA-category dummies, internship indicator, career-category dummies, female status, and transfer-student status. Standard errors are HC1 robust and reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.9: Appendix Heterogeneity and Placebo Checks

| Check | Comparison | No controls | Controls | N |
|---|---|---|---|---|
| Transfer x T1 | Control vs T1 | 0.125 | 0.098 | 425 |
| | | (0.111) | (0.113) | |
| Transfer x T2 | Control vs T2 | 0.310*** | 0.307*** | 433 |
| | | (0.110) | (0.113) | |
| Female x T1 | Control vs T1 | 0.011 | 0.027 | 425 |
| | | (0.098) | (0.100) | |
| Female x T2 | Control vs T2 | 0.086 | 0.098 | 433 |
| | | (0.099) | (0.098) | |
| Controls | | No | Yes | |

*Notes:* Transfer is a secondary theory heterogeneity check; gender is a placebo/contrast check. Controls include race indicators, GPA-category dummies, internship indicator, career-category dummies, female status, and transfer-student status. Standard errors are HC1 robust and reported in parentheses. Stars indicate significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.10: HC message shares and any-HC rates

| Group | User HC share | Provider HC share | All HC share | Any user HC | Any all HC |
|---|---|---|---|---|---|
| Control × Non-FG | 18.5% | 19.5% | 19.0% | 70.9% | 82.9% |
| Control × First-gen | 18.0% | 19.8% | 18.9% | 71.4% | 83.0% |
| Active AI × Non-FG | 46.4% | 56.1% | 51.6% | 92.8% | 97.6% |
| Active AI × First-gen | 43.9% | 52.6% | 48.6% | 91.4% | 95.2% |

*Notes:* Unit is conversation. The table reports hidden-curriculum (HC) message prevalence by experimental arm and first-generation status. User HC share is the fraction of student messages classified as HC; provider HC share is the fraction of AI/provider messages classified as HC; all HC share pools student and provider messages. Any user HC indicates the share of conversations with at least one student HC message. Any all HC indicates the share of conversations with at least one HC message by either the student or provider.