Do Credentials Sort by Cost? A Field Experiment on Educational Signaling

Gedeão Locks*   Matías Giaccobasso†

Job Market Paper

Click here for most recent version

Abstract

We combine a field experiment that corrects misperceptions about credential returns with pre-measured individual effort costs to test whether educational signaling operates through cost-based self-selection, as predicted by Spence (1973). We find that supply-side information frictions are large and consequential for sorting: nearly 80% of students underestimate the returns to academic honors. Correcting these beliefs increases honors attainment, but only among students who are both near the credential threshold and face low effort costs. This is the pattern predicted by the separating equilibrium: once beliefs are corrected, only low-cost types acquire the credential. Our results show that even well-designed credential systems can fail to produce efficient sorting when students are uninformed about the returns, and that a simple information provision can improve sorting at negligible cost.

*We are deeply grateful to Amanda Pallais, Ricardo Perez-Truglia, Pamela Giustinelli, Pia Pinger, Paola Giuliano, Melanie Wasserman, Julie Tréguier, Renke Schmacker, David Margolis and participants in seminars at the University of Chicago, University of California Los Angeles (UCLA), DIW Berlin, and Paris 1 Panthéon-Sorbonne University for helpful comments. German Institute for Economic Research (DIW), Mohrenstraße 58, 10117, Berlin, Germany. Email: glocks@diw.de

†VATT Institute, Helsinki, Finland.

1 Introduction

Educational credentials are designed and often used to sort workers by productivity. In the canonical framework of Spence (1973), this sorting arises from a separating mechanism: high-ability individuals, facing lower costs of acquiring credentials, choose to invest in them, while low-ability individuals do not. This cost-based self-selection is what makes credentials informative and the reason why employers can infer productivity from a diploma or an honor degree (Stiglitz and Weiss, 1990). This mechanism has influenced research on education (Card, 1999), labor markets (Altonji and Pierret, 2001), and information economics (Riley, 2001), and yet it has never been tested using exogenous variation in perceived returns. Regression discontinuity designs identify the returns to credentials at the margin (Clark and Martorell, 2014), but are silent on *why* individuals end up on either side of a cutoff, or whether the sorting actually reflects the process described by Spence (1973). In this paper, we provide the first causal evidence that educational signaling operates through this mechanism, and that misperceptions about returns distort the sorting educational credentials are designed to produce.

We conduct a field experiment among master's students at a large French public university. France's academic honor system provides a particularly clean setting to test predictions of signaling theory for several reasons. First, honors are awarded at fixed, well-known GPA thresholds (12, 14, and 16 out of 20), creating sharp and transparent cutoffs. Second, recent evidence documents a 12.5% earnings premium at the High Honors threshold of 14/20 (Astruc-Le Souder et al., 2024). Finally, the system is in place since the 19th century,[1] making it an institutionally stable environment where the signaling mechanism should operate at its strongest.

At midsemester, we randomly inform students about the estimated earnings premium for High Honors (14/20 threshold). Crucially, before the treatment, we elicit each student's perceived GPA–study-hours mapping by asking: *"You said your GPA is [X]; what would it be if you studied [h] more hours per week?"*. This provides a pre-treatment measure of the individual effort cost of raising GPA, and allows us to test two pre-registered predictions from the separating equilibrium: that treatment effects on credential attainment concentrate

[1] France's honor system (*mentions*) dates to 1840, when the Minister Victor Cousin proposed that diplomas carry graded assessments to "sustain the attention of examiners and stimulate the effort of students." French graduates typically report their honors, rather than their GPA, on CVs and LinkedIn profiles, making honors a primary signal visible to employers in job applications (Le Parisien, 2021).

near the threshold, and that among near-threshold students, effects are driven by those with lower effort costs.

We find that students hold large misperceptions about the returns to academic honors. Nearly 80% of students believe that academic honors matter little or not at all to employers, and a similar share perceives honors carry zero return. This is not because students misunderstand the honor system, nearly all students correctly identify that honors depend solely on GPA with fixed cutoffs. Students understand the rules, but not the rewards. Prior beliefs about the honors premium average only 3.5%, and providing information increases perceived returns, closing roughly 60% of the gap between prior beliefs and the true return. This first stage confirms that our information treatment successfully corrected a widespread misperception.

We then examine whether these updated beliefs translate into behavioral responses. On self-reported intentions, the treatment does not raise study hours, but it increases willingness to display honors on a CV or LinkedIn profile by about 8%. This asymmetry is consistent with strategic behavior across margins of different cost: students adjust the low-cost display margin while leaving the more costly effort margin unchanged. On academic achievement, we find no significant average treatment effect on GPA or High Honors attainment. However, our model predicts that effects should be heterogeneous. Students far from the 14 threshold face high costs of crossing regardless of their beliefs, meaning that even accurate information about returns may not lead to additional effort. Consistent with this, we find that the treatment increased High Honors attainment by 32% (11.5 percentage points) among students who were near the threshold at midsemester. We rule out that this concentration reflects differential belief updating: treated students near and far from the threshold revised their beliefs by similar magnitudes, confirming that the heterogeneity comes from costs, not from differential learning.

Our central result tests the cost-based self-selection mechanism directly. Among students near the 14/20 threshold, we split the sample by pre-treatment effort costs using the elicited GPA-hours mapping. The treatment effect on High Honors attainment is concentrated entirely among high-productivity students, those for whom crossing the threshold requires the least additional effort. Our treatment increases honors attainment by 17.8 percentage points, roughly 45% relative to baseline. For low-productivity students at similar distance from the cutoff, the point estimate is not statistically significant. This means that, once information frictions are removed, the students who acquire the credential are precisely those whose signaling costs fall below the return. To our knowledge, this is the first direct causal evidence for the mechanism at the core of Spence (1973): credentials are informative because low-cost types self-select into obtaining them. Our experimental design identifies this mechanism by exogenously shifting beliefs while holding effort costs fixed, allowing us to observe the sorting process in response to changes in perceived returns.

Our results have broader implications for how credential systems work. Much of the literature on academic standards has focused on credential design (Betts, 1998). Our findings suggest that even in a transparent system with fixed, well-known thresholds, credentials fail to produce efficient sorting when students underestimate the returns. If information frictions distort signaling in this setting, they are likely more consequential in systems where both the thresholds and the returns are harder to observe. We use the individually elicited production functions and pre-treatment beliefs to decompose the population of near-threshold students into those who would cross under correct beliefs but do not (undersignalers), and those whose perceived return already exceeds their measured cost but who do not cross. We find that our experiment recovers roughly half of the potential welfare gains from eliminating information frictions. The remainder is explained by incomplete belief updating and non-monetary effort costs. Extrapolating our results to a national scale suggests that information frictions at a single credential threshold generate annual welfare losses of tens of millions of euros, which could be reduced at nearly zero fiscal cost.

Our paper makes three contributions. First, we provide direct causal evidence for the supply-side mechanism in educational signaling. A growing experimental literature has documented the importance of demand-side signals, showing that providing employers with credible information about worker productivity improves hiring outcomes (Pallais, 2014; Deming et al., 2016; Abel et al., 2020; Abebe et al., 2021; Bassi and Nansamba, 2022; Heller and Kessler, 2024). We show that the supply side is equally important: students misperceive what credentials employers reward, and correcting these misperceptions induces the cost-based self-selection that the credential system is designed to produce. Second, we contribute to the literature on student beliefs about educational returns (Jensen, 2010; Wiswall and Zafar, 2015; Nguyen, 2008; Boneva and Rauh, 2017; Arcidiacono et al., 2020). Most closely related, Ehrmantraut et al. (2020) document large perceived signaling returns to degree completion. Their evidence is descriptive, while we examine whether those beliefs affect credential attainment. By exogenously shifting beliefs about the return to a credential and tracing the effects through the signaling mechanism, we show not only if students respond or not, but whether low-cost (high-ability) students respond differently, as the model predicts. Third, our results shed light on the mixed findings in the RD literature on credential returns. Studies exploiting discontinuities at credential thresholds have found mixed results, with sizable wage premia in some settings (Feng and Graetz, 2017; Khoo and Ost, 2018; Barrera-Osorio and Bayona-Rodríguez, 2019; Astruc-Le Souder et al., 2024) but limited effects in others (Tyler et al., 2000; Bedard, 2001; Clark and Martorell, 2014). Our results suggest that the signaling value of the credential depends on whether high-productivity students are aware of the returns. When they are not, some high-productivity students who would optimally signal fail to do so, and the credential becomes less informative to employers.

The paper proceeds as follows. Section 2 reviews the institutional background and the theoretical predictions. Section 3 details the theoretical framework, while Section 4 preesntes experimental design and data. Section 5 presents results on effort, GPA, and the cost–productivity. Section 6 shows robustness checks and Section 7 discusses welfare implications. Section 8 concludes.

2 Institutional Background

2.1 France's Educational System

In France, education is compulsory from ages 3 to 15 and consists of three stages: primary school (up to age 11), middle school (ages 11–15), and high school (ages 15–18). At the end of high school, students earn the baccalaureate diploma, which grants access to higher education. Grades range from 0 to 20, with 10 as the minimum passing mark. The honors system dates back to in 1840, when the Minister of Public Instruction reformed the high school system and introduced honors (*mentions*) (Le Parisien, 2021). These are awarded based on strict GPA thresholds: 12 (*Assez bien*), 14 (*Bien*), and 16 (*Très bien*); and do not consider (relative) rankings or extracurricular activities and are widely applied across institutions. As a reference, these distinctions correspond broadly to North American Latin honors: *cum laude*, *magna cum laude*, and *summa cum laude*. According to Campus France, official translations are "Honors", "High Honors", and "Highest Honors" (CGF, 2015).

Overall, academic honors are a salient feature of France's education system as this standardized system has been used from secondary school through Master's degrees. We confirm that most students are aware of how honors are obtained in Section 5.

2.2 Master's Degrees

Bachelor programs in Europe are usually three years long. This is the result of the Bologna Process, launched in 1999 and completed by 2010, which harmonized higher education across Europe. It introduced a standardized three-cycle structure (Bachelor–Master–Doctorate) and a common credit system. The Master's degree is often considered the "new Bachelor" in Europe with enrollment shares now comparable to Bachelor completion rates in the 1960s. In 2020, there were 137,000 Master's graduates in France compared to 56,000 in 2000, a 145% increase over two decades (INA, 2024).

In France, access to graduate education follows a selective two-step structure. After completing the three-year undergraduate program (*Licence*, equivalent to a Bachelor's degree), students apply to Master's programs, which typically span two years: Master 1 (M1) and Master 2 (M2). Admission into M1 is competitive, and the progression from M1 to M2 is not automatic. Students must reapply at the end of M1, and universities often impose additional academic or field-specific requirements for M2 entry. This double step selection process is different from the American system, where students are generally admitted directly into a two-year Master's track. As a result, M2 cohorts are smaller and more selective.

Overall, Master's degrees can be considered as the main entry credential for labor markets in continental Europe. In this sense, Master's graduates in Europe are similar to Bachelor's graduates in the United States. However, because admission from M1 to M2 is selective, European Master's graduates are closer in profile to the upper tail of U.S. Bachelor's graduates.

3 Theoretical Framework

In this section, we develop a simple model of educational signaling with information frictions to fix ideas and guide the interpretation of our empirical findings. The model generates testable predictions about which students respond to information about credential returns.

3.1 Setup: Baseline Signaling Model

Student can obtain an academic credential (e.g., High Honors or *mention Bien*) by achieving a GPA at or above a threshold $\bar{g} = 14/20$. Let g_i denote student i's mid semester GPA and $d_i = \max\{\bar{g} - g_i, 0\}$ the distance to the threshold.

Students differ in their marginal productivity θ_i, which determines how efficiently additional effort translates into GPA gains. The cost of obtaining honors is given by:

$$c_i = \frac{d_i}{\theta_i} \tag{1}$$

Students with high productivity (high θ_i) or who are closer to the cutoff(small d_i) face lower costs. Employers cannot observe θ_i but observe credential attainment. They offer a wage premium $\Delta w > 0$ to students with honors.

Under full information about Δw, students obtain honors if and only if the benefit exceeds the cost:

$$e_i^* = 1 \iff \Delta w \geq \frac{d_i}{\theta_i} \tag{2}$$

This produces the standard Spence result: credentials induce sorting because only low-cost students (high productivity) find it optimal to invest (Spence, 1973). As a result, credential attainment is informative to employers.

3.2 Adding Information Frictions

We now relax the assumption that students know the true return Δw. Let $\tilde{\Delta w}_i$ denote student i's *perceived* wage premium, drawn from a distribution $G(\cdot)$ that may differ systematically from the true value. With imperfect information, students obtain honors if and only if:

$$e_i^* = 1 \iff \tilde{\Delta w}_i \geq \frac{d_i}{\theta_i} \tag{3}$$

Information frictions lead some students to underinvest in signaling. We define the undersignaling region as:

$$U = \left\{ i : \tilde{\Delta w}_i < \frac{d_i}{\theta_i} < \Delta w \right\} \tag{4}$$

Students in U would find credential attainment privately optimal under correct beliefs, but do not invest in it due to underestimated returns. Therefore, information frictions distort sorting: some high-productivity students fail to signal, reducing the informativeness of the credential.

3.3 Predictions

Consider an intervention that corrects beliefs, shifting $\tilde{\Delta w_i}$ toward Δw. The intervention affects behavior only for students in the undersignaling region U. This yields two testable predictions.

1. Near-threshold concentration. Treatment effects are concentrated among students close to the threshold.

The undersignaling region requires $c_i < \Delta w$, or equivalently $d_i < \theta_i \cdot \Delta w$. For students far from the threshold, this condition fails for most productivity levels—even high-θ_i types face costs exceeding the true return. For students near the threshold, the condition holds for a wider range of productivities.

2. Cost-based sorting. Among near-threshold students, treatment effects are larger for high-productivity (low-cost) types.

A student is in the undersignaling region if $\tilde{\Delta w_i} < d_i/\theta_i < \Delta w$. For high-$\theta_i$ students, the cost d_i/θ_i is low, so even moderate underestimation places them in U. For low-θ_i students, costs exceed the true return regardless of beliefs.

Prediction 2 directly tests the mechanism underlying signaling models (Spence, 1973). If credential attainment responds to information correction specifically among low-cost types, the sorting process operates through cost-based self-selection, the channel that makes credentials informative to employers (Spence, 1973; Stiglitz and Weiss, 1990).

3.4 Taking Stock

We finish the model discussion by summarizing the main insights. Our analysis indicates that information frictions lead to undersignaling among students for whom credential attainment is optimal, generating two sources of heterogeneity. First, treatment effects are concentrated among students close to the credential threshold, where signaling costs are low. Second, within this group, responses are larger for high-productivity (low-cost) students, who are most likely to underinvest in signaling due to underestimated returns. In the next section, we will empirically test these theoretical implications of the model in the data.

4 Data, Experimental Design, and Implementation

4.1 Subject Recruitment

We carried out the information experiment with students from Université Paris 1 Panthéon-Sorbonne, one of France's largest universities. Sorbonne offers Master's programs in diverse fields, including Economics, Law, Political Science, Business, Philosophy, Humanities, Mathematics, and the Arts. Programs are classified as either professional or research-oriented. Professional Master's prepare students for employment in firms and public administration; research Master's are intended for academic or analytical careers, including doctoral studies.

The survey was conducted within the University learning-management system (LMS) where students access courses, assignments, and activities. The LMS ensured institutional environment, secure single sign-on, and familiar navigation, as students are used to answer to surveys in this platform.

On March 17, we emailed invitations to 14,595 students enrolled in 2024–2025 master's programs (M1 and M2) across all disciplines, targeting first- and second-year students. The message described a study on labor-market prospects and linked directly to the LMS survey page (See Appendix Figures A.9 and A.11 for a screenshot of the invitation email and landing page). Participation was voluntary and survey questions were not incentivized, but to encourage participation, students were offered the option to enroll in a lottery after completing the survey, with five prizes of €200 Amazon gift cards.

The survey remained open through March 22. To encourage participation, we sent three reminder emails at two-day intervals to students who had not yet completed the survey. Access was restricted to eligible students via University credentials and log in. After consent, responses were linked to administrative records (final GPA and honors receipt) using University identifiers under an approved data-use agreement.

4.2 Survey Design

In this section, we summarize the main features of the survey.

Step 1 (Pre-Experimental background characteristics).

Measuring individual productivity (effort costs): First, we elicit a pre-treatment GPA–hours mapping ("You said your current GPA is X. What do you think your final GPA would be if you studied 1/3/5/10 more hours per week?") which we use to construct a local return on investment (ROI), which is the perceived marginal ΔGPA per extra study hour. We use these responses to construct students' perceived marginal returns to effort, that is, the expected change in GPA per additional study hour. We use these as proxies for individual productivity or, equivalently, the effort cost of credential attainment.

Measuring proximity to thresholds: Second, we measure characteristics relevant to how students should respond to credential returns according to signaling theory. We elicit: (a) students' distance to the High Honors threshold (their mid-semester GPA relative to 14/20); (b) prospective employer characteristics, how likely students expect to negotiate wages in their first job, and whether they are in fields (e.g., economics, business, management) where prior research documents larger honors premia (Astruc-Le Souder et al., 2024).

Additional pre-registered heterogeneity: Third, we collect additional pre-registered heterogeneity dimensions: gender, self-confidence, and for students in two-year master's programs, whether they are in their first or second year. Importantly, all heterogeneity measures were elicited before the information treatment to avoid endogenous responses.

Step 2 (Elicit prior beliefs). We then asked about students' prior beliefs on returns to graduating with High Honors (14/20 threshold). Specifically, we asked participants two questions on how they believe honors are considered by employers. First, with a more qualitative flavor, we asked whether they consider *mention Bien* (High Honors) to be important for employers during recruitment in a 1-4 likert scale (1-Very important, 2-Quite important, 3-Of little importance, 4-Not important). Second, with a more quantitative flavor, we asked about the expected percentage wage difference between otherwise similar graduates with and without High Honors. We give students a slider and ask the following question "*By how much does obtaining Bien mention (High Honors) affect first-year salary compared to* *an otherwise identical candidate without it? Enter a percentage between –50 and +50*"

Step 3 (Information-provision experiment). Our information treatment is designed to maximize credibility and relevance. We provide students with earnings premium information that speaks directly to their context: the estimate comes from a recent academic study (Astruc-Le Souder et al., 2024) conducted specifically with former master's students from their own institution. By receiving earnings data from previous cohorts of graduates at their university, rather than national statistics, we expect to increase the likelihood that students internalize and act on the information.

Treated students receive the following message: "*A recent academic study conducted specifically with former master's students from Université Paris 1 Panthéon-Sorbonne found that obtaining High Honors (mention bien) is associated with a 12.5% increase in cumulative first-year earnings after graduation, equivalent to approximately 1.5 additional months of salary.*" We randomize whether the student is shown this information with a probability of 50%. To avoid inference from the act of receiving information, randomization is made explicit: on the first screen, participants are told that some respondents will be randomly selected to receive the benchmark; on the next screen, they learn whether they were selected.

Step 4 (Elicit posterior beliefs). Immediately after the information screen, we re-elicit beliefs from all respondents covering the same objects as in Step 1 (expected honors premium, employer recognition). To avoid subjects inferring that re-elicitation signals an incorrect prior, we explicitly inform them that every participant is asked again regardless of initial answers.

Step 5 (Intention Outcomes). In the last module, we ask about intended behavior: willingness to display honors on the CV, willingness to exert effort, and planned weekly time allocation across study/work/rest (budget task).

To identify the causal effect of beliefs on behavior, it is crucial to observe priors and to introduce exogenous variation in posteriors. Under the information-provision experiment, under-estimators of the honors premium should revise beliefs upward and over-estimators downward, while accurate priors should exhibit little change. Whether students are willing to exert additional effort and, ultimately, responses in terms of realized outcomes will depend on the prior–posterior shift and other individual characteristics such as proximity to the 14/20 threshold and perceived returns to effor (GPA–hours mapping). We therefore embed the information shock within the survey and subsequently link all respondents to administrative outcomes (final GPA and honors attainment) once grades are released.

These questions form the core of the survey. The full questionnaire is available in Appendix B.

4.3 Data and Sample Selection

A total of 2,815 students accessed the survey, corresponding to a response rate of approximately 17%. We apply standard cleaning procedures to ensure data quality. First, we drop 507 students who did not complete the full survey. Next, we exclude 116 "speeders"—defined as students who spent fewer than 3 seconds (control group) or 5 seconds (treatment group) on the feedback page (i.e., below 5th percentile in on-screen time within group). We further remove 95 students who entered implausible prior beliefs (exactly 50 or –50), and 270 students with extreme priors (below 1 or above 20), corresponding to the bottom 8.5 and top 5 percentiles of the prior distribution. None of these exclusions are correlated with treatment assignment. Our final analysis sample consists of 1,827 students. We match survey responses to university administrative records to obtain final GPA and honors status, which is automatically awarded based on GPA cutoffs. Table 1 presents descriptive statistics.

4.4 Implementation

As previously mentioned, we timed the intervention to mid-semester. This timing was chosen to be after the first set of exams but still early enough for behavioral adjustments to affect final GPA. At this point, students already knew their first-semester GPA, which is essential for our survey because it anchors their expectations about final performance and the likelihood of reaching honors thresholds. Since final GPA still depends heavily on second-semester finals, students had enough time (around 13 weeks) to react to the treatment. We linked survey responses to each student's administrative record with final GPA and honors receipt. Importantly, our field experiment took place after the period in which students could choose their second-semester courses, avoiding the concern about strategical switching (e.g., students choosing more lenient courses after treatment).

Table 1: Descriptive Statistics and Balance Between Treatment Arms

	Treatment Arm		
	Control (1)	Treatment (2)	p-value test (3)
Female	64.341	66.509	0.416
	(1.887)	(1.877)	
Age	23.589	23.404	0.461
	(0.188)	(0.165)	
French	82.791	84.518	0.404
	(1.487)	(1.439)	
GPA	13.533	13.626	0.269
	(0.061)	(0.059)	
Internship	72.333	72.923	0.777
	(1.492)	(1.460)	
Part-time	52.889	48.652	0.070
	(1.665)	(1.643)	
Full-time	39.444	40.022	0.801
	(1.630)	(1.610)	
No experience	5.111	4.854	0.801
	(0.734)	(0.706)	
Knows Honors are based on GPA	86.889	86.084	0.615
	(1.126)	(1.137)	
Gap to Knows 14 Threshold	0.110	0.128	0.561
	(0.023)	(0.022)	
Prior: Honors matter to employers	29.556	27.832	0.416
	(1.522)	(1.473)	
Prior: Gap to Honors' Premium on Earnings	9.011	9.056	0.840
	(0.160)	(0.156)	
Answer Survey on First Day	43.556	41.748	0.435
	(1.654)	(1.621)	
Survey Time	12.145	13.071	0.382
	(0.721)	(0.776)	
Major: MGMT/Econ	31.000	31.823	0.705
	(1.543)	(1.531)	
Observations	900	927	

Notes: This table reports baseline covariate balance across treatment and control groups. Column (1) shows means for the control group, column (2) for the treatment group, and column (3) reports p-values from two-sided t-tests of equality of means. Standard errors are reported in parentheses

Survey perception measures require dealing with outliers in responses. Extreme answers can reflect misunderstanding, typos, or inattention rather than true beliefs, creating large but uninformative "information shocks" that bias estimates. Following standard practice, we exclude respondents with the most extreme prior misperceptions (Giaccobasso et al., 2025). In the baseline, we drop the bottom 5% and top 5% of prior misperceptions, removing 323 respondents, leaving 2,210 for analysis. This implies a net response rate of 1,827/14,495 ≈ 13%. Because exclusions use pre-treatment variables, internal validity is preserved. We also report results using different cutoffs and falsification tests.

4.5 Outcomes of Interest

As stated in the RCT pre-registration, our primary outcomes are academic performance and honors attainment. Academic performance is final GPA and honors attainment at pre-specified thresholds, measured from university records. To benchmark levels, we report control-group means for GPA and honors. We also include one question to serve as a secondary outcome: we ask each respondent's perceived causal effect of High Honors on their own wage. We ask the following: "Consider your personal case. By how many percent would your wage differ if you received High Honors versus if you did not?".

5 Results

5.1 Do students misperceive the returns to academic honors?

Signaling models of education rely on a critical assumption: workers on the supply side understand which credentials employers value and adjust their investments accordingly (Spence, 1973). This means that the awareness of the returns to a given educational credential is a necessary condition for a separating equilibrium to emerge (Stiglitz and Weiss, 1990). Yet despite the theoretical centrality of this assumption, empirical evidence on whether students perceive credential returns accurately remains limited. If students systematically underestimate the labor-market value of attainable credentials, high-ability students may not invest to obtain them, credentials lose informational content, and matching efficiency is affected. Testing for such supply-side information frictions is therefore key to understanding when and how educational credentials function as signals.

We begin by documenting students' knowledge about the French academic honors system and the returns associated with it. Figure 1 shows that students understand the "automatic" assignment of honors: approximately 80% correctly identify that honors are awarded purely on the basis of GPA at strict thresholds (Panel **a**), with 14/20 for High Honors (*mention bien*) and 16/20 for Highest Honors (*mention très bien*) (Panel **b**). This is expected since these rules have been in place since the 19th century, have not changed and are used since middle school until higher education. Yet students may not fully understand the labor market implications of these credentials. Previous evidence, not only in the context of education, shows that people misperceive simple and accessible statistics and underestimate the economic returns of education (Wiswall and Zafar, 2015; Jensen, 2010; Giaccobasso et al., 2025).

Figure 1: Students Knowledge and Beliefs about Honors

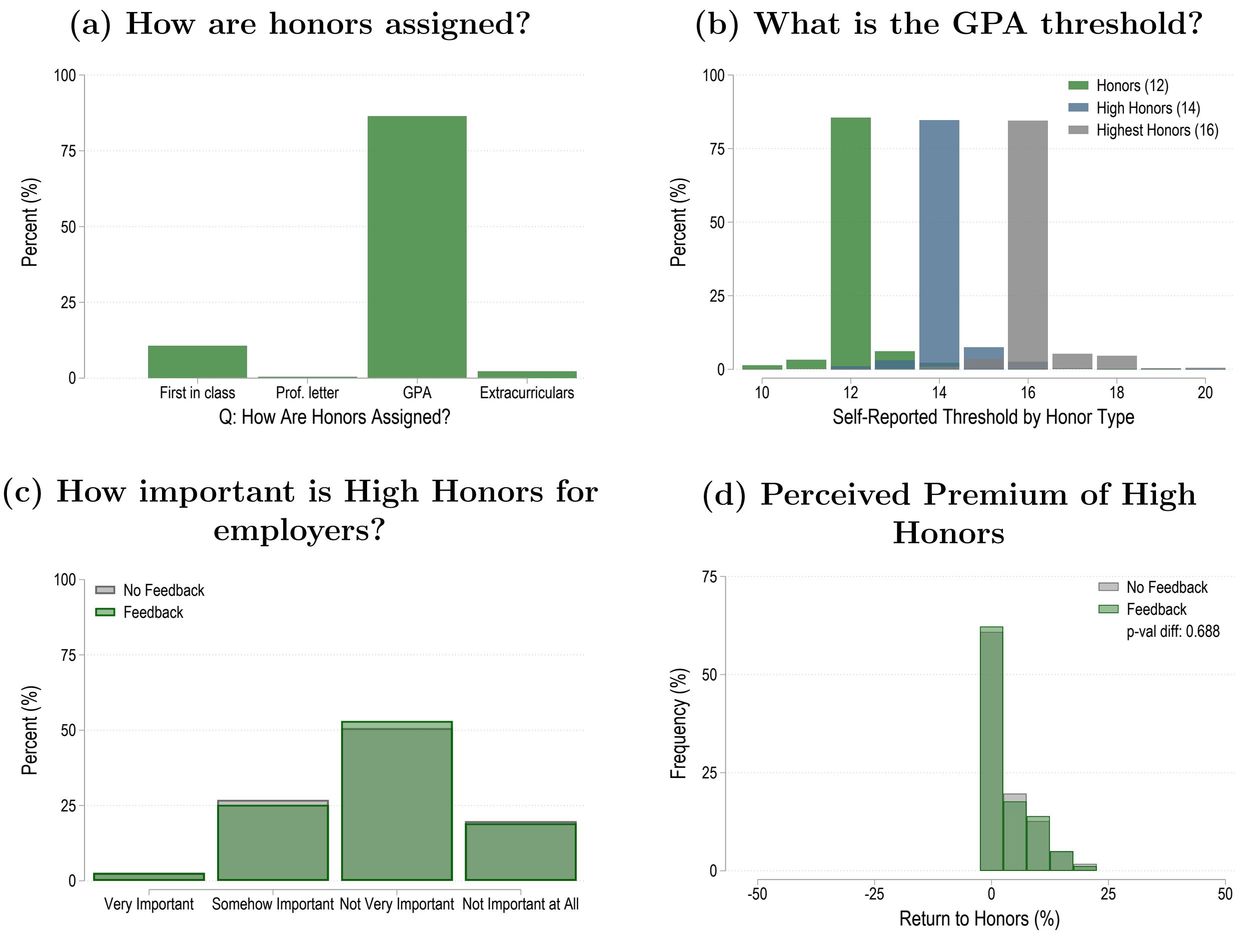

Notes: Panels (a)–(b) summarize knowledge of the honors assignment rule and GPA cutoffs. Panels (c)–(d) report qualitative and quantitative beliefs about the labor-market value of High Honors. In panel (d), 43% of respondents report a return within 1 percentage point of 0%.

Results show that the institutional knowledge does not translate into accurate beliefs about returns. When asked qualitatively how important High Honors are for employer recruiting, most students report "little importance" or "no importance" (Panel **c**). Nearly 70% believe High Honors carry zero wage premium, reporting a salary difference of 0% between otherwise identical candidates with and without High Honors (Panel **d**). This stands in sharp contrast to prior research documenting sizable earnings premia for honors in the French labor market (12.5%) (Astruc-Le Souder et al., 2024). In short, students know how to obtain the credential but not what it is worth, implying a substantial supply-side friction and confirming the scope for belief correction via information provision.

5.2 Belief Updating

We use the subscript i to index subjects. Let s_i^{prior} denote subject i's belief about the labor-market value of graduating with honors (relative to graduating without honors) before the information-provision stage. Let s_i^{feed} represent the feedback about the returns to honors that subject i can potentially receive in the experiment. Define T_i^S as a binary variable equal to 1 if subject i is selected to receive the information about the returns to honors, and 0 otherwise. Let s_i^{post} denote the posterior belief about the value of honors: s_i^{post} represents the perceived labor-market return to honors after the subject observes (or does not observe) the feedback.

An individual shown feedback will form her posterior belief s_i^{post} as the average of the prior belief s_i^{prior} and the feedback s_i^{feed}, weighted by a parameter α that captures the degree of learning. This parameter ranges from 0 (subjects ignore the feedback) to 1 (subjects fully adjust to the feedback) and depends on the relative precision of the prior compared to the precision of the feedback. This Bayesian updating model can be summarized by the following linear relationship:

$$s_i^{post} - s_i^{prior} = \alpha \cdot \left(s_i^{feed} - s_i^{prior} \right). \qquad (5)$$

Intuitively, Bayesian learning predicts that, when shown feedback, subjects who initially *overestimate* the signaling value of honors will revise their beliefs downward, while those who *underestimate* will revise their beliefs upward. Figure 2 illustrates this Bayesian learning model. In Figure 2(a) we show the mean error and mean absolute error binned by how far priors were from the true results. In Figure 2(b) we show a binned scatterplot: the x-axis corresponds to the gaps in prior beliefs ($s_i^{feed} - s_i^{prior}$), and the y-axis corresponds to belief updating ($s_i^{post} - s_i^{prior}$).

We find that students update their inaccurate beliefs when provided with accurate feedback. To measure belief updating, we use a simple Bayesian model that has been shown to accurately represent belief formation in other information-provision experiments on a wide range of topics (Giaccobasso et al., 2025; Cullen et al., 2022).

The green circles in Figure 2(b) correspond to the subjects shown feedback about the 12.5% higher earnings for students who graduated with High Honors. Consistent with significant updating, there is a strong relationship between the updated beliefs and prior gaps: an additional percentage point (pp) in the perception gap is associated with an actual revision of 0.603 pp higher.

Figure 2: Perceptions about the Honors' Earning Premium

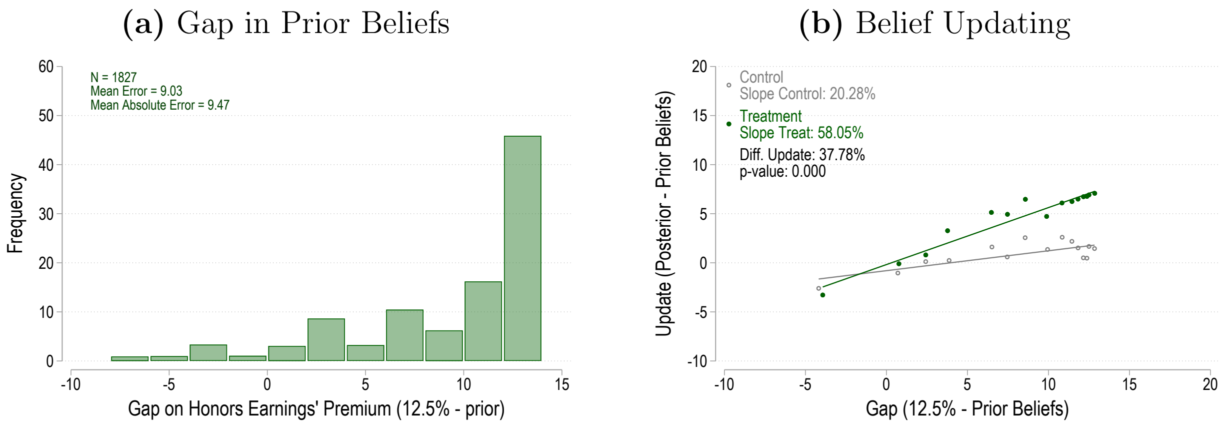

Notes: Panel (a) shows the gap in prior beliefs about the honors earnings premium. The x-axis reports the difference between the actual premium and respondents' prior beliefs in 10 percentage-point bins. The y-axis reports the percentage of survey respondents in each bin. The upper left corner reports the total number of observations, the mean error, and the mean absolute error. Panel (b) shows how respondents update their beliefs using a binned scatterplot (with ten bins based on deciles of the belief gap). The x-axis reports the difference between the actual premium and prior beliefs. The y-axis shows the change from prior to posterior belief (i.e., belief updating). Green circles (gray squares) represent the average update within each bin for the group that received (did not receive) feedback about the honors premium. Each line shows fitted values from separate OLS regressions of belief updating on the belief gap. Reported in the top-left are the slope coefficients, robust standard errors (in parentheses), p-value of the difference in slopes, and the number of observations.

The gray circles in Figure 2(b) correspond to subjects who do not receive feedback about honors. In the absence of feedback, these subjects should not update their beliefs. However, in practice, individuals may revise their beliefs in the direction of the feedback for spurious reasons even when they do not receive feedback. For example, respondents may reassess their answers or correct typos when asked a question a second time, leading to an answer closer to the truth.

The gray squares indicate a weak relationship between belief updating and prior gaps in the group that was not shown the feedback: an additional 1 pp in the prior gap is associated with an actual revision of 0.204 pp higher. This effect is statistically significant but economically very small. This result is consistent with other information-provision experiments that show evidence of spurious revisions (e.g.,Fuster et al. (2022) and Cullen et al. (2022)).

How do students update their beliefs? Our evidence suggests that learning among treated students reflects *how* they process information rather than simple exposure. Quantitative updating is stronger in math-heavy majors, consistent with students with better math knowledge feeling interpreting percentage returns and wage premia in an easier way. By contrast, qualitative learning, moving from “it doesn’t matter” to “it matters”, does not differ by math-intensity major. We interpret this as indicating that the basic update that “High Honors are valued by employers” is widespread even if precise magnitudes are harder to internalize. Moreover, we find little support for spillovers: the modest belief changes observed in the control group are most spurious rather than driven by cross-arm information being spread. Finally, learning does not vary systematically with baseline GPA, suggesting that proximity to performance cutoffs shapes *incentives* but not the cognitive uptake of information about returns.

Learning in the Control Group. We examine whether control group students learn about credential returns through spillovers or other channels. Figure A.3 shows that control students do not increase the share of “correct” answers (within 1 percentage point of the 12.5% premium we provided to treated students). Moreover, Figure A.4 indicates that control students do not change their qualitative beliefs about whether honors matter to employers. Finally, Figure A.5 shows that control students do not provide more precise answers on average. This pattern indicates no true learning in the control group.

But what explains the upward slope in control group updating? Students who initially report high priors (above 10%) appear to revise their beliefs downward when asked the same question again without receiving information. Figure A.6 shows that, when we exclude students with initial priors above 10%, the positive update in the control group disappears entirely. This indicates learning in the control group is most likely spurious, consistent with noisy belief reporting rather than information transmission, as seen in many other contexts (Giaccobasso et al., 2025; Cullen et al., 2022).

5.3 Average Treatment Effects on Intentions

Having established that students have updated their beliefs about the actual returns to honors, we estimate average treatment effects on our registered intention outcomes. Because most learning in our sample is more qualitative ("it matters") and numerical updating is uneven, we start with average treatment effects (ATEs) on preregistered intention outcomes.

Figure 3: Average Treatment Effects on Intentions

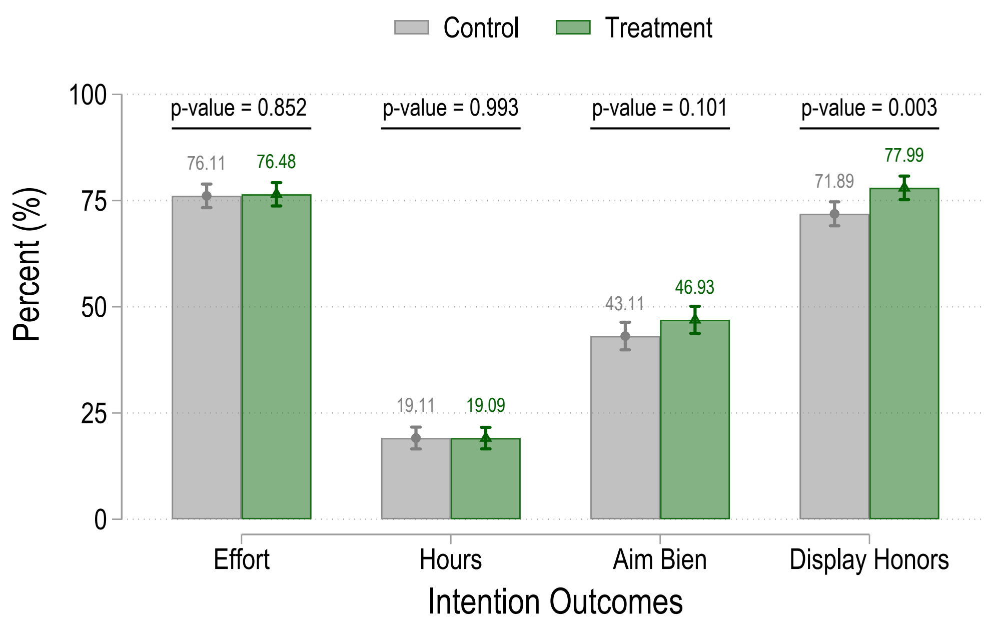

Notes: This figure illustrates the average treatment effects of the honors earnings premium information intervention. Panel (a) reports the mean of the posterior belief about the honors earnings premium by treatment status. Gray bars show the average posterior for the control group; green bars show the average posterior for the treated group. For each group, we report: (1) the raw difference in means between treated and control (Δ^C and Δ^F), (2) the difference in means conditional on baseline covariates ($\Delta^C|X_i$ and $\Delta^F|X_i$), (3) the p-value of the equality of means test within each group, and (4) the p-value for the difference in treatment effects across subgroups (i.e., a double-difference test). Panel (b) replicates this analysis using honors receipt as the dependent variable.

The results are presented in Figure 3, which reports both the raw mean difference between treatment and control groups, as well as conditional on a set of baseline control variables. We see that effort-related intentions, displayed in the left side of Figure 3 do not change with the information. Self-reported effort and weekly study hours are similar across treatment and control groups. The right panel of the figure shows honors-related intentions: treated students are more likely to target High Honors (14/20) and are 6 percentage points more likely to report that they would list the honor on their CV or LinkedIn. These first findings can be interpreted in two ways: (i) students update aspirations before adjusting costly inputs; or (ii) there is limited room to change study time by mid-semester. This finding aligns with a large body of literature documenting small educational choice responses to monetary incentives (Arcidiacono, 2004; Wiswall and Zafar, 2015).

5.4 Treatment Effect of Posterior Beliefs

To measure the causal effect of posterior beliefs on intentions and subsequent outcomes, we follow the econometric models used in other information-provision experiments (Giaccobasso et al., 2025; Cullen et al., 2022).

Let Y_i denote an outcome of interest for student i (e.g., hours of study, self-reported effort, or ΔGPA). Let s_i^{prior} be the student's prior belief about the signaling value of honors, s_i^{feed} the feedback provided by the experiment, and s_i^{post} the posterior belief. The treatment consists of random assignment to receive feedback, which shifts posterior beliefs independently of unobserved determinants of effort.

We estimate the following equation:

$$Y_i = \beta_0 + \beta_1 s_i^{post} + X_i \beta_X + \varepsilon_i, \tag{6}$$

where ε_i is the usual error term. The parameter of interest is β_1, which captures the effect of posterior beliefs on outcomes. However, since s_i^{post} may be endogenous, we estimate equation 6 using 2SLS, instrumenting s_i^{post} with the randomized treatment and the exogenous shift $(s_i^{feed} - s_i^{prior})$:

$$Y_i = \beta_0 + \beta_1 s_i^{post} + \gamma(s_i^{feed} - s_i^{prior}) + X_i \beta_X + \varepsilon_i. \tag{7}$$

The exclusion restriction is that treatment affects outcomes only through its impact on posterior beliefs. This setup allows us to isolate the causal channel from belief updating to effort and GPA.

5.5 Results: ATE and 2SLS on Academic Outcomes

Table 2 presents results on the impact of correcting information frictions on academic outcomes. Panel (a) reports effects on obtaining High Honors (*mention Bien*, while Panel (b) reports effects on percentage changes in GPA from mid-semester to year-end. Columns (1) to (3) provide intent-to-treat (ITT) estimates starting from the baseline specification and adding pre-registered controls and additional covariates. Columns (4) to (6) show two-stage least squares (2SLS) estimates in which we instrument posterior beliefs with treatment assignment.

Table 2: Average Treatment Effects and 2SLS on Hard Outcomes

	ATE			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
(a) Effects on High Honors (14/20) Attainment:						
Treatment/Posterior Effects	3.639	3.726*	3.920*	0.726	0.605	0.621
	(2.510)	(2.254)	(2.267)	(0.557)	(0.500)	(0.492)
Observations	1,443	1,443	1,443	1,443	1,443	1,443
Baseline Mean	33.10	33.10	33.10	33.10	33.10	33.10
(b) Effects on % Δ GPA						
Treatment/Posterior Effects	-0.159	-0.033	-0.029	-0.092	-0.072	-0.068
	(0.423)	(0.388)	(0.388)	(0.091)	(0.082)	(0.080)
Observations	1,347	1,347	1,347	1,347	1,347	1,347
Baseline Mean	1.90	1.90	1.90	1.90	1.90	1.90
Prereg. Controls	No	Yes	Yes	No	Yes	Yes
Addl. Controls	No	No	Yes	No	No	Yes

Notes: This table reports the effects of the information provision on (a) obtaining *mention Bien* (GPA $\geq$ 14) honors and (b) percentage changes in GPA. Columns (1)–(3) report average treatment effects (ATE) estimated using equation 6, while columns (4)–(6) report two-stage least squares (2SLS) estimates using equation 7, where posterior beliefs about the wage return to honors instrument for credential attainment. All specifications include cohort fixed effects. Columns labeled "Prereg. Controls" add the set of pre-registered baseline covariates, and columns labeled "Addl. Controls" further include academic and demographic controls. Standard errors are reported in parentheses. Baseline means correspond to the control group.

Columns (1) to (3) show modest positive ITT effects on honors attainment. When the full set of controls is included, the estimate corresponds to roughly a 4 percentage point increase relative to a baseline mean of 33%, or about a 12% increase compared with the control group. The estimates in columns (2) and (3) are statistically significant only at the 10% level, although they remain imprecisely estimated overall. By contrast, the 2SLS estimates in columns (4) to (6) are smaller and not statistically significant. Results for GPA growth in Panel (b) show no detectable effects: point estimates are close to zero and statistically insignificant across specifications, indicating that the treatment did not affect average academic performance.

The 2SLS estimates in columns (4) to (6) support the same conclusion. These specifications scale the reduced-form effect by the first-stage impact on beliefs, recovering the effect of a one percentage point increase in perceived returns. Again, the estimates are small and not statistically significant across all specifications. The stability of results across columns suggests that the null findings are not driven by omitted variable bias or specification choices.

These null average effects are consistent with previous evidence showing that information interventions shifts beliefs but often do not translate into behavioral changes in educational behavior (Oreopoulos and Petronijevic, 2019). Wiswall and Zafar (2015) show that providing major-specific earnings information shifts beliefs substantially but yields only small changes in stated major intentions. Arcidiacono (2004) document that students update beliefs about ability in response to grades, yet many persist in their initial choices.

This pattern is also consistent with the theoretical framework developed in Section 3, which predicts that average effects should be zero in our setting. The undersignaling region U, the set of students for whom belief correction changes behavior, is non-empty only for students with low signaling costs. For students far from the threshold or with low productivity, costs exceed the true return regardless of the change in beliefs. Therefore, the null average effect suggests that the information is relevant only for a specific subpopulation. We now test these predictions by examining heterogeneity in treatment effects.

5.6 Testing Prediction 1: Distance to Threshold

The theoretical framework in Section 3 highlighted that information frictions generate a "concentration" effect: belief correction shifts behavior only for students in the undersignaling region U, which requires signaling costs to fall below the actual honors return. In Section 5.5, we documented null average treatment effects, consistent with the prediction that most students lie outside this region. We now turn to directly testing Prediction 1, which states that treatment effects should concentrate among students close to the credential threshold, where distance costs are low.

To measure proximity to the threshold, we use students' mid-semester GPA relative to the 14/20 cutoff for High Honors. We split the sample into three equally sized terciles based on mid-semester GPA and re-estimate treatment effects separately for students close to versus far from the 14/20 High Honors threshold. We classify students as "close" if their midsemester GPA falls in the top tercile of the distribution (12.9–14.3) and "far" otherwise. Students close to the threshold face mechanically lower costs of credential attainment: a student at 13.7/20 needs only 0.3 additional GPA points to reach High Honors, whereas a student at 12/20 needs 2.0 points. Moreover, students near the 14 threshold may face less uncertainty about whether additional effort will secure the credential, making strategic optimization more credible. Our model predicts that belief correction should induce behavioral responses primarily in the former group.

Table 3 presents the results, reporting treatment effects separately for students close to and far from the threshold. A first observation is that the treatment effect for "far" students is essentially zero and not precisely estimated. The point estimate is slightly negative, and we can rule out positive effects for this group. This is consistent with the model: for students far from the threshold, even accurate beliefs about credential returns do not make signaling worthwhile because costs exceed the wage premium.

Table 3: Heterogeneity: Far and Close to the Threshold

	High Honors (14/20) Attainment		% Δ GPA	
	ATE	2SLS	ATE	2SLS
Close to 14	11.421***	1.875*	0.592	-0.061
	(4.419)	(0.991)	(0.734)	(0.148)
Far from 14	-0.116	0.179	-0.412	-0.082
	(2.792)	(0.611)	(0.460)	(0.098)
(Difference Close - Far)	11.537**	1.696	1.005	0.021
	(5.229)	(1.179)	(0.865)	(0.179)
Mean Outcome (Baseline)				
Close	35.09	35.09	1.62	1.62
Far	32.18	32.18	2.03	2.03
Observation	1,443	1,443	1,347	1,347

Notes: This table reports baseline covariate balance across treatment and control groups. Column (1) shows means for the control group, column (2) for the treatment group, and column (3) reports p-values from two-sided t-tests of equality of means. Standard errors are reported in parentheses

Students close to the threshold show a different response. The ITT estimate indicates that the information treatment increases honors attainment by roughly 11.5 percentage points, representing a 32% increase relative to the control baseline mean. The difference between near- and far-threshold effects is statistically significant at the 5% level. Figure 4(a) also shows the estimated treatment effects on High Honors (14/20) attainment by pre-experiment

GPA bins is only significant for students sitting below the threshold at midsemester. This concentration of effects among near-threshold students confirms Prediction 1.

Figure 4: ATE on High Honors and GPA Growth by Proximity to Threshold

(a) ATE on High Honors Attainment by Midsemester GPA

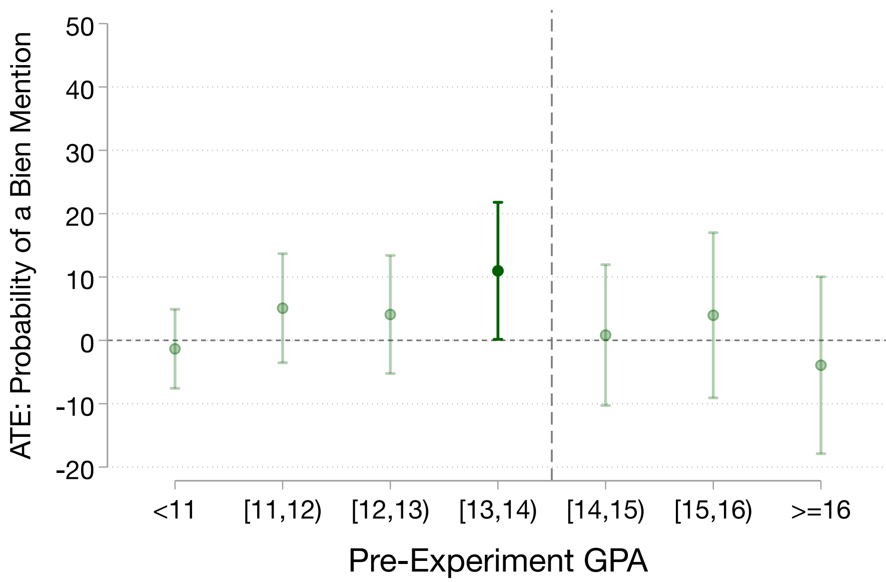

(b) Δ GPA: Far from Threshold

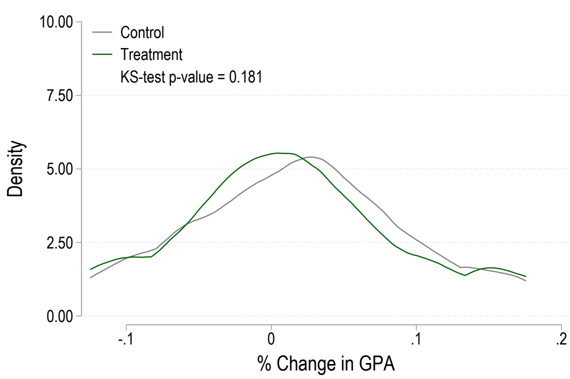

(c) Δ GPA: Close to Threshold

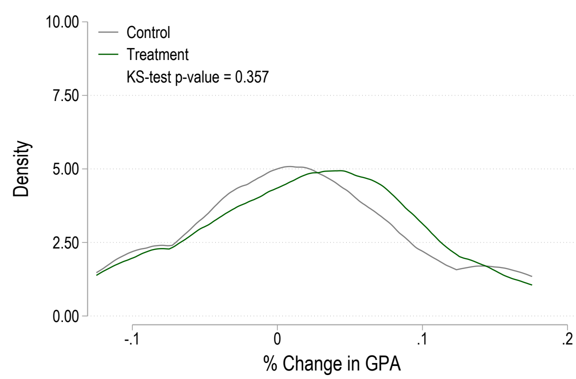

Notes: Panel (a) shows the estimated treatment effect on High Honors attainment (in percentage points) by midsemester GPA bin, with 95% confidence intervals. Panels (b) and (c) show kernel density estimates of GPA growth from midsemester to year-end, separately for treated and control students. "Far from Threshold" (panel b) denotes students in the bottom two terciles of midsemester GPA; "Close to Threshold" (panel c) denotes students in the top tercile. The Kolmogorov-Smirnov test evaluates the null that treatment and control distributions are drawn from the same population. All specifications in panel (a) include pre-registered controls and major fixed effects. Robust standard errors.

The GPA results are consistent with these findings but are less precisely estimated. Figure 4 displays non-parametric GPA growth distributions by treatment status. For students far from the threshold (panel b), the distributions are similar, with the control group slightly shifted to the right. For students close to the threshold (panel c), treated students display a modest rightward shift, with slightly more positive GPA growth, though the difference is not statistically significant. The corresponding estimates in Table 3 are positive but imprecise, matching the visual patterns.

This imprecision likely reflects the continuous nature of the GPA outcome. In contrast, the binary honors outcome aggregates this variation, with greater power and capturing the signal observed by employers. Taken as a whole, the evidence supports the model's first prediction that treatment effects are concentrated near the threshold. In the next section, we test Prediction 2 by examining whether these effects are driven by high-productivity, low-cost students.

5.7 Testing Prediction 2: Cost-Based Sorting

Prediction 2 states that among near-threshold students, treatment effects should concentrate among those with low effort costs: the high-productivity types for whom credential attainment is privately optimal once beliefs are corrected. In Section 5.6, we confirmed that treatment effects concentrate near the threshold. We now turn to directly testing the cost-based sorting mechanism that underlies signaling models.

To measure productivity, we use a key feature of our experiment. Before treatment, we elicited each student's perceived effort-GPA mapping by asking: "Your current GPA is X. What would it be if you studied 1 additional hour per week? What about 3 hours? 5 hours?" Students reporting large GPA gains from modest hour increases have high marginal productivity. In other words, they face low effort costs of credential acquisition. We define "high productivity" as students who report needing at most one additional hour per week to reach the 14/20 threshold. This measure captures the effort cost component of signaling costs, complementing the distance component analyzed above.

Table 4 presents treatment effects by proximity to threshold and productivity. The results show a pattern that is consistent with Prediction 2. First, among near-threshold students, treatment effects are driven by the high-productivity group. The ITT estimate indicates that correcting information frictions increases honors attainment by approximately 18 percentage points for high-productivity students close to the threshold, representing a 45% increase relative to the control mean. This effect is large and statistically significant. Figure 5 confirms In contrast, low-productivity students near the threshold show a positive but imprecisely estimated effect. Second, students far from the threshold show no treatment response regardless of productivity. The group of high-productivity students far from the threshold is very small and the point estimate is not statistically significant. For low-productivity students far from the threshold, the point estimate is also imprecisely estimated.

Table 4: Proximity and Productivity

	High Honors (14/20) Attainment		% Δ GPA	
	ATE	2SLS	ATE	2SLS
Close to 14, High Productivity	17.843**	2.464	1.320	-0.053
	(6.962)	(1.532)	(1.174)	(0.236)
Close to 14, Low Productivity	8.955	2.324	0.847	0.013
	(6.201)	(1.429)	(1.070)	(0.226)
Far from 14, High Productivity	13.217	-4.199	2.230	-0.065
	(17.142)	(5.506)	(4.269)	(0.507)
Far from 14, Low Productivity	3.388	0.632	-0.312	-0.135
	(2.858)	(0.678)	(0.807)	(0.176)
Mean Outcome (Baseline)				
Close, High ($\#T = 91$, $\#C = 93$)	39.81	39.81	1.78	1.78
Close, Low ($\#T = 94$, $\#C = 107$)	25.26	25.26	1.32	1.32
Far, High ($\#T = 10$, $\#C = 8$)	10.00	10.00	3.90	3.90
Far, Low ($\#T = 211$, $\#C = 203$)	8.37	8.37	4.24	4.24
Observations	870	870	817	817

Notes: This table reports baseline covariate balance across treatment and control groups. Column (1) shows means for the control group, column (2) for the treatment group, and column (3) reports p-values from two-sided t-tests of equality of means. Standard errors are reported in parentheses

As in the previous section, we also show non-parametric estimates of GPA changes for students close to the 14 threshold by productivity. Figure 5 shows an overlap of the distributions for low-productivity students (panel b). For high-productivity students (panel c), the pattern is substantially different: the treatment distribution shows a larger mass in the positive growth region, with treated students showing a shift to the right relative to controls. Although the Kolmogorov-Smirnov test lacks power to detect this difference at conventional levels given the cell size, the visual evidence is consistent with the regression estimates and suggests that high-productivity students near the threshold improved their academic performance in response to treatment. This evidence reinforces the main finding of our paper: correcting information frictions lead to behavioral responses specifically among low-cost types, consistent with the sorting mechanism underlying signaling models (Spence, 1973).

Taking stock. The concentration of effects among high-productivity students near the 14/20 threshold provides direct evidence for the separating equilibrium mechanism predicted

Figure 5: ATE on High Honors and GPA Growth by Proximity and Productivity

(a) ATE on High Honors Attainment by Productivity and Proximity

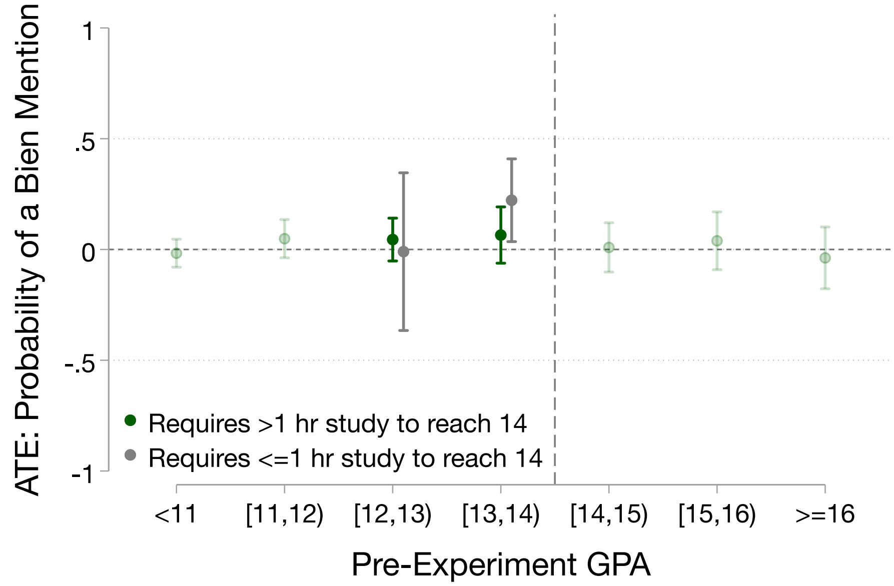

(b) Close to Threshold, Low Productivity

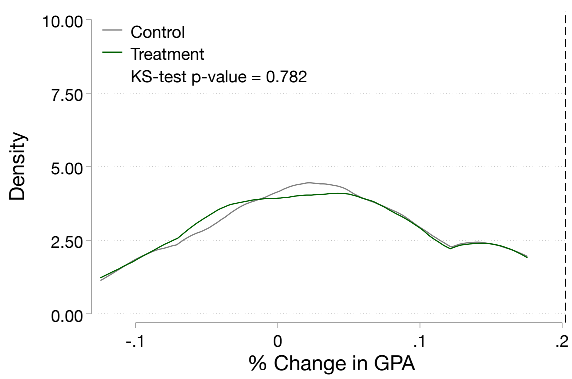

(c) Close to Threshold, High Productivity

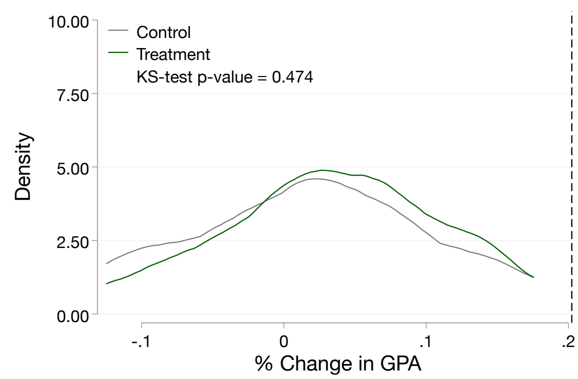

Notes: Panel (a) shows the estimated treatment effect on High Honors attainment (in percentage points) by midsemester GPA bin, split by low- and high-productivity. Panels (b) and (c) show histograms of percentage change in GPA. "Low Productivity" denotes students who report needing more than one additional study hour per week to reach the 14/20 threshold; "High Productivity" denotes students who report needing at most one hour. The sample is restricted to students in the top tercile of mid-semester GPA. The KS test evaluates the null hypothesis that distributions are drawn from the same underlying distribution.

by signaling theory. The canonical model assumes that high-productivity students face lower effort costs, making them the types for whom signaling is optimal (Spence, 1973). This assumption is difficult to test because productivity and costs are typically unobserved. Our design overcomes this challenge by measuring effort costs before treatment and introducing exogenous shifts in beliefs. Our results confirm that once information frictions are corrected, the sorting mechanism operates exactly as theory predicts: credential attainment responds among low-cost types, while high-cost types do not respond regardless of beliefs. Our field experiment provides the first direct evidence for the cost-based self-selection that makes credentials informative to employers.

6 Robustness Checks

In this section, we conduct a series of checks to show our results are robust to alternative definitions of proximity and productivity. Our central finding is that correcting information frictions induces credential attainment among students who are both near the threshold and face low effort costs. This result relies on two sample splits: proximity to the 14/20 cutoff and a pre-treatment productivity measure, which we now test the robustness of each. Moreover, we test whether our results reflect self-confidence rather than productivity and run placebo tests in the other two thresholds (12 and 16).

Alternative proximity definitions. A first concern is that the definition of "near threshold", which uses the top tercile of midsemester GPA, may be driving our results. To address this, we re-estimate treatment effects using four alternative definitions: a continuous measure of distance to the 14/20 cutoff interacted with treatment, and fixed bandwidths of 1, 0.5, and 0.25 GPA points below the threshold. Table A.1 reports the results. The point estimates are stable across definitions: 11.4 percentage points for the baseline tercile split, 12.1 within 1 GPA point, 10.8 within half a point, and 15.3 within 0.25 points. As the bandwidth narrows, estimates remain similar in magnitude but standard errors increase as expected given smaller cell sizes. The continuous specification confirms the pattern: the treatment effect declines with distance to the threshold, though the coefficient is imprecisely estimated.

Alternative productivity measures. A second potential concern is related to our productivity measure. We measure productivity based on a single elicitation point (whether one extra hour of study is enough to cross) may be sensitive to how productivity is defined. We construct four continuous productivity indices that progressively incorporate more of the elicited GPA–hours schedule. We first use the 1-hour response alone, then the average of the 1-hour and 3-hour responses, the average including the 5-hour response, and all four elicitation points. For each measure, we classify students above the 75th percentile as high-productivity and re-estimate the main specification. Figure A.7 reports the results. The pattern is stable and statistically significant across all definitions, except for including 10-hour responses, which is consistent with decreasing marginal returns to hours of study. Across different measures, the treatment effects on High Honors attainment concentrate among high-productivity students near the threshold, regardless of how productivity is measured.

Productivity or self-confidence? A related concern is that the productivity measure based on the GPA—hours mapping may capture self-confidence rather than effort costs. We pre-registered heterogeneity analyses based on three self-confidence questions[2], which we combine into a standardized index. In Table A.3, we start by showing that the correlation between elicited productivity and an index of self-confidence is very low (0.05). Next, we regress elicited productivity on the confidence index, controlling for midsemester GPA and additional covariates. We find that the confidence coefficient is not statistically different than zero. Finally, we use the control group to test whether elicited productivity or self-confidence predict actual GPA growth among near-threshold students. If students who report high productivity actually have higher GPA growth (even without the information shock), then our measure captures something real about their production function, not just optimism. We find that only the productivity coefficient is statistically significant, whereas the self-confidence index is not. This result confirms that our productivity measure captures real differences in the ability to convert effort into grades.

Placebo tests at thresholds 12 and 16. 4 We conduct a placebo tests at the 12/20 (Honors) and 16/20 (Highest Honors) thresholds. Our treatment specifically informed students about the earnings premium at the 14/20 cutoff. If the results reflect a threshold-specific response to information rather than a generic motivation effect, we should find no treatment effects at other thresholds. While Figure 4 shows average treatment effects by midsemester GPA bins, we complement this analysis with a placebo exercise that further splits students into high- and low-productivity groups, based on the elicited measure of hours required to

[2]See full questionnaire in Appendix B.

reach the 12 and 16 cutoffs. Figure A.8 reports the results. At the 12/20 threshold (panel a), treatment effects on Honors attainment are not statistically different from zero across all GPA bins, including for high-productivity students near the cutoff. The same pattern is observed near the 16/20 threshold (panel b), with large confidence intervals reflecting the small number of students near this cutoff. In sum, neither of these thresholds show the concentration of effects among near-threshold, high-productivity students that we document at 14/20.

7 Welfare Implications

We quantify the welfare effects of correcting information frictions using a sufficient statistics approach (Chetty, 2009). Let w denote the actual wage premium from High Honors, c_i the individual cost of obtaining it, and p the perceived return before treatment. Students with $c_i \in (p, w)$ would optimally obtain the credential but do not because they underestimate returns. Our treatment shifts perceived returns toward w, inducing those with $c_i < w$ to cross. The welfare gain is:

$$\Delta W = \int_p^w (w - c)\, dF(c) \approx N_{\text{switch}} \times (\bar{w} - \bar{c}_{\text{switch}}) \tag{8}$$

where N_{switch} is the number of treatment-induced switchers and $\bar{c}_{\text{switch}}$ their average effort cost.

Costs and benefits. We use the estimate from Astruc-Le Souder et al. (2024) that High Honors carries a 12.5% first-year earnings premium for French master's graduates. Each switcher gains $\bar{w}$ = €4,375.[3] Our ITT estimate of 17.8 percentage points in the Close × High cell ($N_T = 91$) implies approximately 16 treatment-induced switchers. On the cost side, high-productivity students report that one additional hour of weekly study suffices to cross the threshold. Over the 13 remaining weeks of the semester, at an opportunity cost of €12/hour, this implies $\bar{c}_{\text{switch}} = 1 \times 13 \times 12 =$ €156.[4].

[3]Astruc-Le Souder et al. (2024) estimate the premium using administrative records from master's graduates in business, economics, law, and political science. Table A.2 shows that these majors account for roughly 70% of our sample.

[4]Balance tests confirm that the share of high-productivity types, average weekly hours needed to reach the threshold ($\bar{h} = 3.3$), and average GPA distance ($\bar{d} \approx 0.5$) are all balanced across treatment and control. Hence, the net gain per switcher is €4,375 − €156 = €4,219, for an aggregate welfare gain of roughly €68,000

Decomposing the undersignaling region. Using the framework from Section 3, we classify each student within 1 GPA point of the threshold based on their elicited production function and pre-treatment beliefs. For each student, we interpolate the GPA–hours schedule to compute the euro cost of crossing and compare it to their perceived and true returns. This yields three groups (Table A.4 in the Appendix): students whose cost exceeds the true return (efficiently non-crossing)[5], students whose cost falls between their perceived and true return (undersignalers), and students whose perceived return exceeds their cost yet who do not cross (informed non-crossers).

This exercise leads to the following results. First, among high-productivity students near the threshold, 34.7% are classified as undersignalers under full information. This means that, if beliefs were perfectly updated, the crossing rate would rise from 40% to 75%. This striking result implies that more than one third of students who could efficiently cross the threshold fails to do so because of incorrect beliefs. Second, average signaling costs differ sharply across our productivity groups: €111 for high-productivity students vs. €677 for low-productivity students. This cost heterogeneity is at the basis of the Spence mechanism and what makes credentials informative when the sorting process works.

How much of this potential does our experiment recover? Our experiment recovers about half of this potential: the 17.8 percentage point treatment effect represents roughly half of the full-information benchmark (34.7). Two factors explain this gap. First, incomplete updating: students close about 60% of the gap between prior beliefs and the true return. If they had fully updated, another 9pp would have switched. This partial updating is consistent with Bayesian learning as students combine the treatment (the study we cite) with prior beliefs formed through previous experience (e.g., internships, peer networks), and updating rates of 40–70% are standard in information experiments (Wiswall and Zafar, 2015; Giaccobasso et al., 2025). Second, the remaining gap of 8pp likely reflects non-monetary effort costs that are not captured by our elicited production function. These students believe the return exceeds their measured cost, yet do not cross.

This decomposition has a direct policy implication because it shows both the potential and the limits of our information intervention. The 9 percentage points lost to incomplete updating represent the margin that better information provision could improve. For example,

[5] In practice, because the true return is large relative to the effort costs faced by students within 1 GPA point of the threshold, no student in this range has costs exceeding the true return.

through more credible or repeated information campaign. The remaining gap of 8 percentage points due to non-monetary factors, by contrast, would require interventions that reduce the non-monetary costs of effort (e.g., tutoring or tailored GPA distance alerts). In light of this, we believe our experimental estimate consists of a lower bound on the gains from correcting information frictions, but also a realistic upper bound on what a single information campaign can achieve.

Semi-elasticity. Prior beliefs about the honors premium average approximately 3.5%. Our treatment shifts posteriors to 8.5%, while control posteriors reach only 4.5%, implying a treatment-induced belief shift of $\Delta p = 4.0$ percentage points. The semi-elasticity of honors attainment with respect to perceived returns is $\varepsilon = 0.178/0.040 = 4.45$: a one percentage point increase in perceived returns raises honors attainment by 4.45 percentage points among near-threshold, high-productivity students. This statistic is portable to counterfactual policy evaluation: an intervention that shifts perceived returns by Δp would generate approximately $4.45 \times \Delta p \times N_{\text{eligible}}$ additional honors recipients.

Aggregate cost of undersignaling. To assess the potential national magnitude of these frictions, we present an extrapolation exercise. Each year, around 160,000 students graduate with a master's degree in France. Applying the share of undersignalers from our sample to this population, and assuming that the earnings premium and the distribution of costs and beliefs at Paris 1 are broadly representative of other French universities in these fields, we estimate that information frictions at the 14/20 threshold alone generate a welfare loss on the order of €20–€53 million per year. The lower bound restricts the calculation to high-productivity students, while the upper bound includes all near-threshold students whose measured costs fall below the true return. These estimates are conservative because they cover a single threshold, assume the earnings premium does not persist beyond the first year, and exclude potential undersignaling at the 12/20 and 16/20 cutoffs. Nevertheless, even under such assumptions, the cost of eliminating these frictions, which means incorporating credential return information into existing university communications, is negligible.

Resulting equilibrium. A plausible concern is whether widespread information provision would erode credential value. Our results suggest the opposite. Because effects concentrate among high-productivity students near the threshold, information correction improves the quality composition of the honors pool. In the Spence (1973) framework, removing frictions strengthens the separating equilibrium and the credential becomes more informative because the "high" types now sort into it. Despite the limitations of our survey, we find no evidence of wasteful effort among high-cost students, nor discouraging effects among students far from the threshold. While employers might in principle respond to aggregate shifts in honors rates, the magnitude of our effect is far too small to move the credential distribution.

8 Conclusion

This paper provides new causal evidence on whether correcting supply-side information frictions impacts the cost-based self-selection mechanism that makes educational credentials informative to employers. We combine a field experiment among master's students with pre-treatment measures of individual effort costs, allowing us to directly test the sorting mechanism at the core of signaling theory (Spence, 1973).

Standard signaling models assume individuals are fully informed, yet we document large gaps in students' beliefs about the returns to academic honors, even in a setting with transparent rules and a long-standing credential structure. When we correct these beliefs, High Honors attainment increases by 17.8 percentage points, but only among students who are both near the credential threshold and face low effort costs. This double selection is exactly the pattern predicted by Spence's separating equilibrium and, to our knowledge, the first direct evidence that educational signaling operates through cost-based self-selection.

The welfare analysis suggests that our intervention realizes about half of the potential gains. Under full information, roughly 35% of high-productivity near-threshold students would additionally obtain High Honors. Of the remaining gap, half reflects incomplete belief updating, while the other half comes from non-monetary effort costs. These findings indicate that, besides information provision, other policies to address effort barriers are needed. A back-of-the-envelope calculation suggests that, at the national level, these frictions generate welfare losses estimated between €20–€53 million each year.

Finally, our experiment captures short-run responses. Whether behavioral changes persist across cohorts, or whether equilibrium effects emerge as information spreads, remains an open question. Future work could extend this approach to other credentials and to longer-run outcomes such as job search behavior, wage negotiation, and career progression.

References

[1] Girum Abebe et al. “Anonymity or distance? Job search and labour market exclusion in a growing African city”. In: *The Review of Economic Studies* 88.3 (2021), pp. 1279–1310.

[2] Martin Abel, Rulof Burger, and Patrizio Piraino. “The value of reference letters: Experimental Evidence from South Africa”. In: *American Economic Journal: Applied Economics* 12.3 (2020), pp. 40–71.

[3] Joseph Altonji and Charles Pierret. “Employer Learning and Statistical Discrimination”. In: *Quarterly Journal of Economics* 116.1 (2001), pp. 313–350.

[4] Peter Arcidiacono. “Ability Sorting and the Returns to College Major”. In: *Journal of Econometrics* 121.1–2 (2004), pp. 343–375.

[5] Peter Arcidiacono et al. “Ex ante returns and occupational choice”. In: *Journal of Political Economy* 128.12 (2020), pp. 4475–4522.

[6] Mael Astruc-Le Souder, Olivier Bargain, and Gedeão Locks. *A Question of Honor? The Labor Market Advantage of Academic Signaling*. Working Paper. Bordeaux School of Economics, 2024.

[7] Felipe Barrera-Osorio and Hernando Bayona-Rodríguez. “Signaling or better human capital: Evidence from Colombia”. In: *Economics of Education Review* 70 (2019), pp. 20–34.

[8] Vittorio Bassi and Aisha Nansamba. “Screening and signalling non-cognitive skills: experimental evidence from Uganda”. In: *The Economic Journal* 132.642 (2022), pp. 471–511.

[9] Kelly Bedard. “Human Capital versus Signaling Models: University Access and High School Dropouts”. In: *Journal of Political Economy* 109.4 (Aug. 2001), pp. 749–775. DOI: 10.1086/322089. URL: https://ideas.repec.org/a/ucp/jpolec/v109y2001i4p749-775.html.

[10] Julian R Betts. “The impact of educational standards on the level and distribution of earnings”. In: *The American Economic Review* 88.1 (1998), pp. 266–275.

[11] Teodora Boneva and Christopher Rauh. “Socio-economic gaps in university enrollment: The role of perceived pecuniary and non-pecuniary returns”. In: (2017).

[12] David Card. “The Causal Effect of Education on Earnings”. In: *Handbook of Labor Economics* 3A (1999), pp. 1801–1863.

[13] CGF. *Équivalences des diplômes français aux USA*. French. PDF. Consulate General of France in Boston. 2015.

[14] Raj Chetty. “Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods”. In: *Annu. Rev. Econ.* 1.1 (2009), pp. 451–488.

[15] Damon Clark and Paco Martorell. “The signaling value of a high school diploma”. In: *Journal of Political Economy* 122.2 (2014), pp. 282–318.

[16] Zoë B. Cullen, Shengwu Li, and Ricardo Perez-Truglia. “What’s My Employee Worth? The Effects of Salary Benchmarking”. In: *Review of Economic Studies* (2022). Forthcoming.

[17] David J Deming et al. “The value of postsecondary credentials in the labor market: An experimental study”. In: *American Economic Review* 106.3 (2016), pp. 778–806.

[18] Laura Ehrmantraut, Pia Pinger, and Renske Stans. “The expected (signaling) value of higher education”. In: (2020).

[19] Andy Feng and Georg Graetz. “A question of degree: the effects of degree class on labor market outcomes”. In: *Economics of Education Review* 61 (2017), pp. 140–161.

[20] Andreas Fuster et al. “Expectations with endogenous information acquisition: An experimental investigation”. In: *Review of Economics and Statistics* 104.5 (2022), pp. 1059–1078.

[21] Matias Giaccobasso et al. “Where do my tax dollars go? tax morale effects of perceived government spending”. In: *American Economic Journal: Applied Economics* 17.4 (2025), pp. 223–259.

[22] Sara B Heller and Judd B Kessler. “Information frictions and skill signaling in the youth labor market”. In: *American Economic Journal: Economic Policy* 16.4 (2024), pp. 1–33.

[23] INA. *Master à l’université : diplôme, augmentation des effectifs, étudiants*. `https://www.ina.fr/ina-eclaire-actu/master-universite-diplome-augmentation-effectif-etudiant`. 2024.

[24] Robert Jensen. “The (perceived) returns to education and the demand for schooling”. In: *The Quarterly Journal of Economics* 125.2 (2010), pp. 515–548.

[25] Pauline Khoo and Ben Ost. “The effect of graduating with honors on earnings”. In: *Labour Economics* 55 (2018), pp. 149–162.

[26] Le Parisien. *Bac & Brevet : à quoi servent les mentions ?* Le Parisien Étudiant, published July 5, 2021. 2021. URL: https://www.leparisien.fr/etudiant/examens/bac/bac-brevet-pourquoi-distribue-t-on-des-mentions-RMLMNGGCF5MYFECFKC42KCPP4Y.php.

[27] Trang Nguyen. “Information, role models and perceived returns to education: Experimental evidence from Madagascar”. In: *Unpublished manuscript* 6.5 (2008).

[28] Philip Oreopoulos and Uros Petronijevic. *The Remarkable Unresponsiveness of College Students to Nudging and What We Can Learn from It*. NBER Working Paper No. 26059. 2019.

[29] Amanda Pallais. “Inefficient hiring in entry-level labor markets”. In: *American Economic Review* 104.11 (2014), pp. 3565–3599.

[30] John G Riley. *Silver signals: Twenty-five years of screening and signaling*. 2001.

[31] Michael Spence. “Job market signaling”. In: *The Quarterly Journal of Economics* 87.3 (1973), pp. 355–374.

[32] Joseph E Stiglitz and Andrew Weiss. *Sorting out the differences between signaling and screening models*. 1990.

[33] John H. Tyler, Richard J. Murnane, and John B. Willett. “Estimating the Labor Market Signaling Value of the GED”. In: *The Quarterly Journal of Economics* 115.2 (May 2000), pp. 431–468. DOI: 10.1162/003355300554818. URL: https://EconPapers.repec.org/RePEc:oup:qjecon:v:115:y:2000:i:2:p:431-468..

[34] Matthew Wiswall and Basit Zafar. “Determinants of College Major Choice: Identification Using an Information Experiment”. In: *The Review of Economic Studies* 82.2 (2015), pp. 791–824.

Online Appendix (For Online Publication Only)

Signaling Effects in Higher Education: Experimental Evidence from France

Locks, and Giacobasso

April 2, 2026

A Further details and results

A.1 Additional Tests for the Validity of the Experimental Design

Figure A.1: Distribution of Earning Premium Gap

(a) Honors Earning Premium Gap

(b) Honors Earning Premium Gap by *Close*/*Far* of Bien Threshold

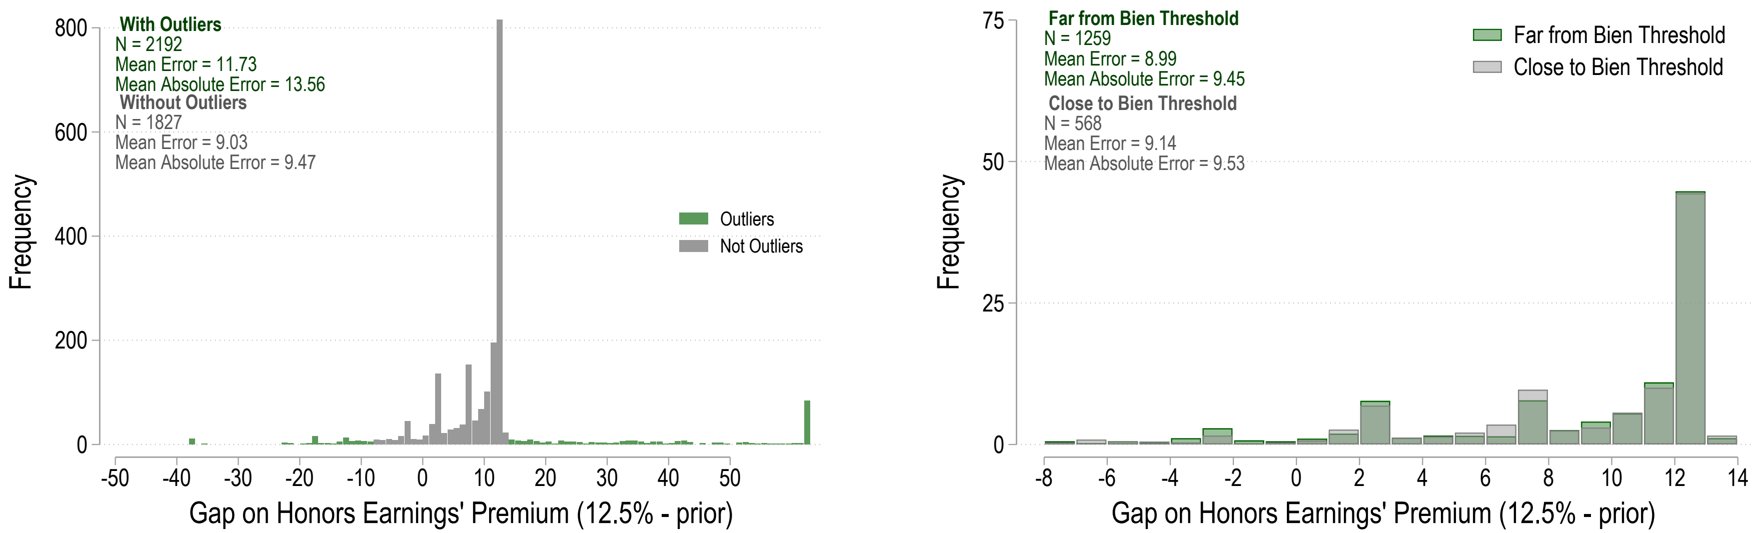

Notes: This figure shows the distribution of belief gaps regarding the honors earnings premium, defined as the difference between the actual premium and the respondent's prior belief. In Figure (a), gray bars represent the central 90% of the distribution, while green bars indicate the top and bottom 5%. The ranges shown in the figure use slightly different bins for visualization purposes. Figure (b) presents the distribution of honors premium belief gaps separately for students far from to the *mention Bien* threshold (GPA 14/20; green bars) and those close (gray bars), based on the main experimental sample ($N = 2{,}815$). The y-axis reports the share of respondents in each 10 percentage-point bin. In the upper left of each panel, we report the number of observations, the mean error, and the mean absolute error.

Figure A.2: Distribution of Gap by Treatment Status

(a) General Question, Analysis Sample   **(b)** General Question vs Specific (you) Question, Full Sample

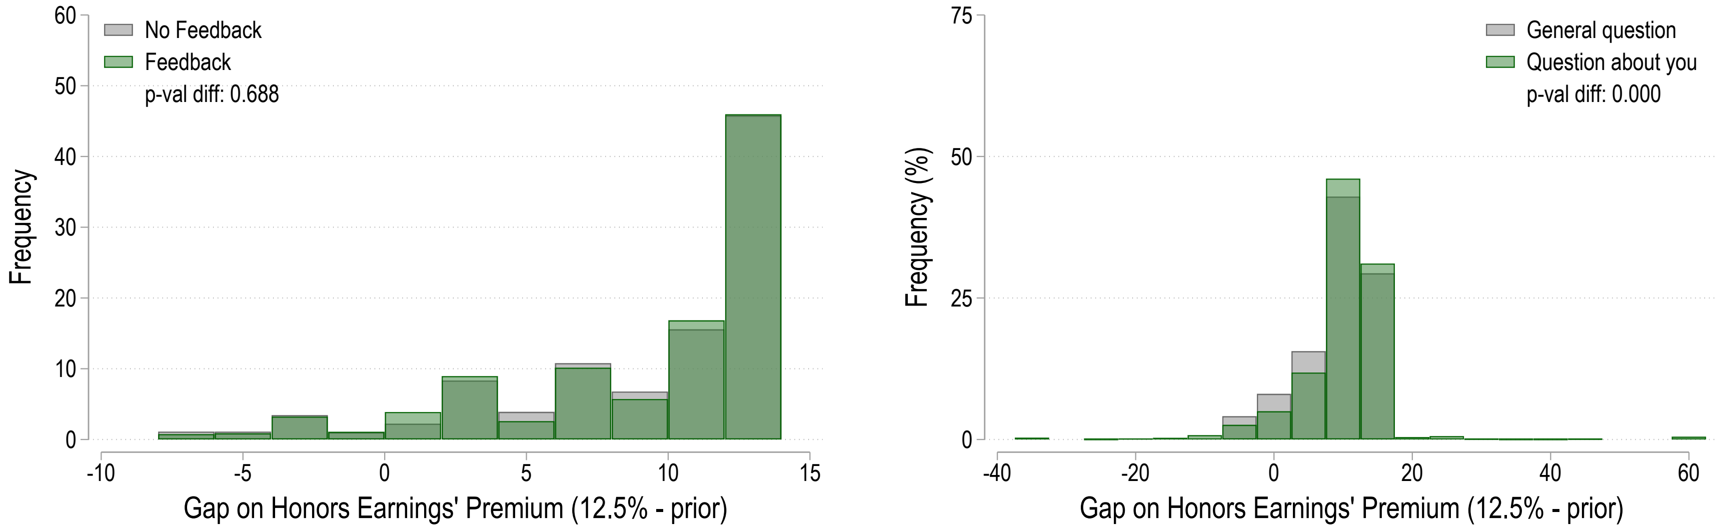

Notes: This figure shows the distribution of belief gaps regarding the honors earnings premium, defined as the difference between the benchmark premium (12.5%) and the respondent's prior belief. Figure (a) presents the distribution for the analysis sample, separately by treatment status. Figure (b) compares responses to a general question about the earnings premium for honors students versus a personalized question asking what respondents believe *they* would earn with honors.

Figure A.3: Share of Correct Answers (1pp) by Treatment Status

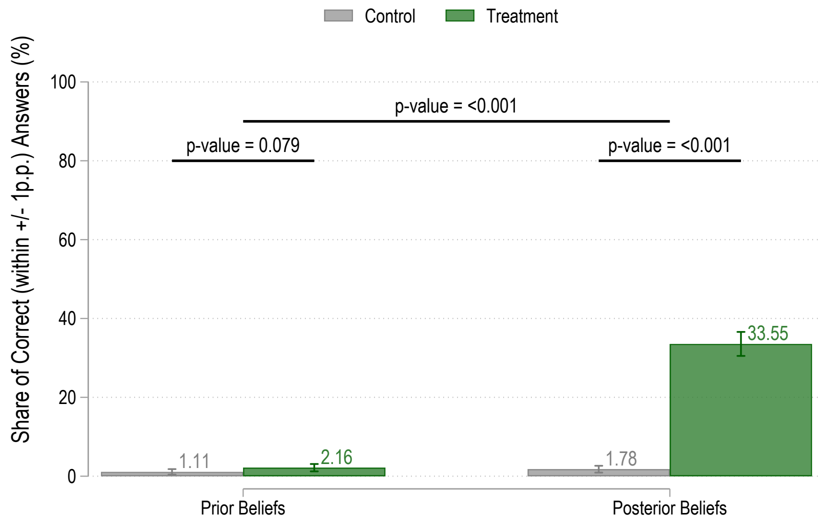

Notes: This figure shows the percentage of survey questions answered correctly (within ±1 percentage point of the truth) before and after the intervention, split by treatment group.

Figure A.4: Qualitative Beliefs About Whether Honors Matter to Employers

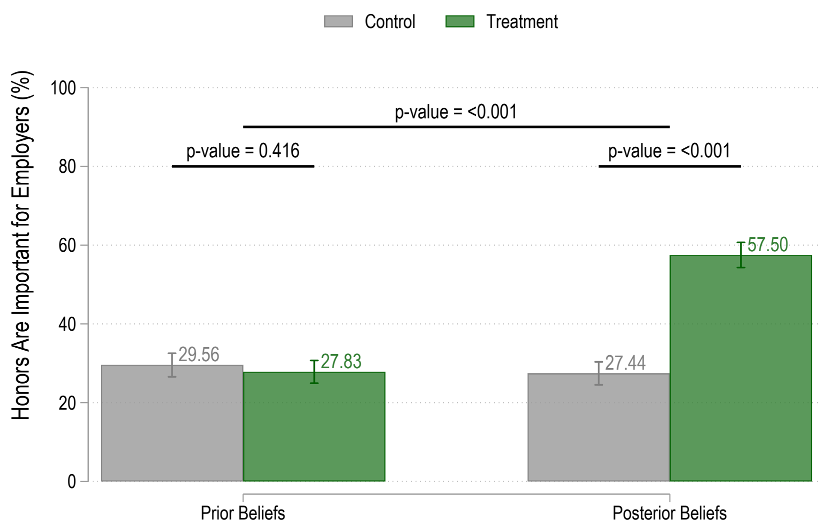

Notes: This figure shows beliefs about whether honors matter to employers, before and after treatment. For both panels, p-values correspond to tests of differences in means.

Figure A.5: Change in Absolute Prediction Error

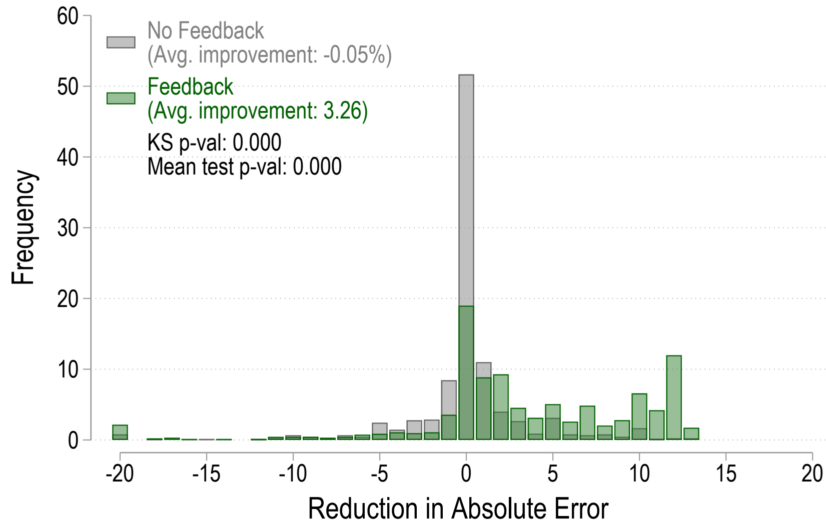

Notes: This figure displays the distribution of changes in absolute error between prior and posterior beliefs. The average reduction in absolute error is 3.26 percentage points in the treatment group.

Figure A.6: Learning by Treatment Status, Excluding High-Prior Students

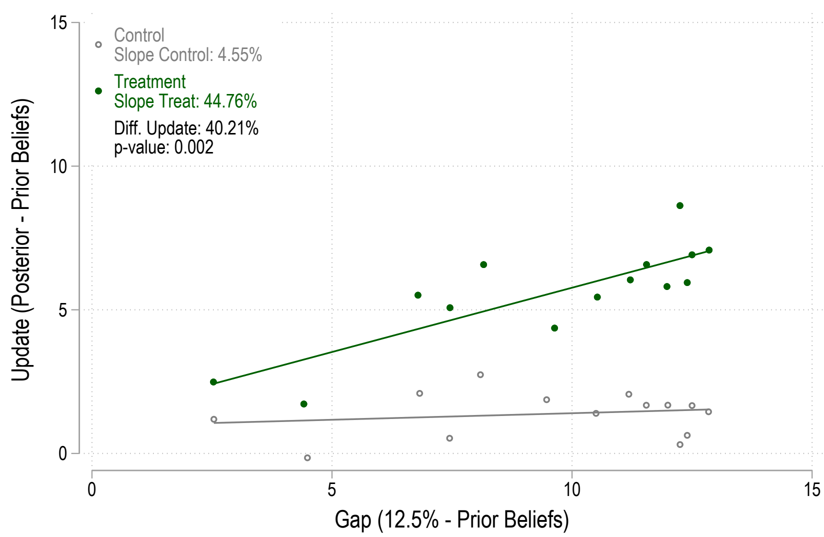

Notes: This figure shows belief updating as a function of the initial belief gap, excluding individuals whose priors exceed 10 percentage points (outliers).

Figure A.7: Alternative Productivity Measures

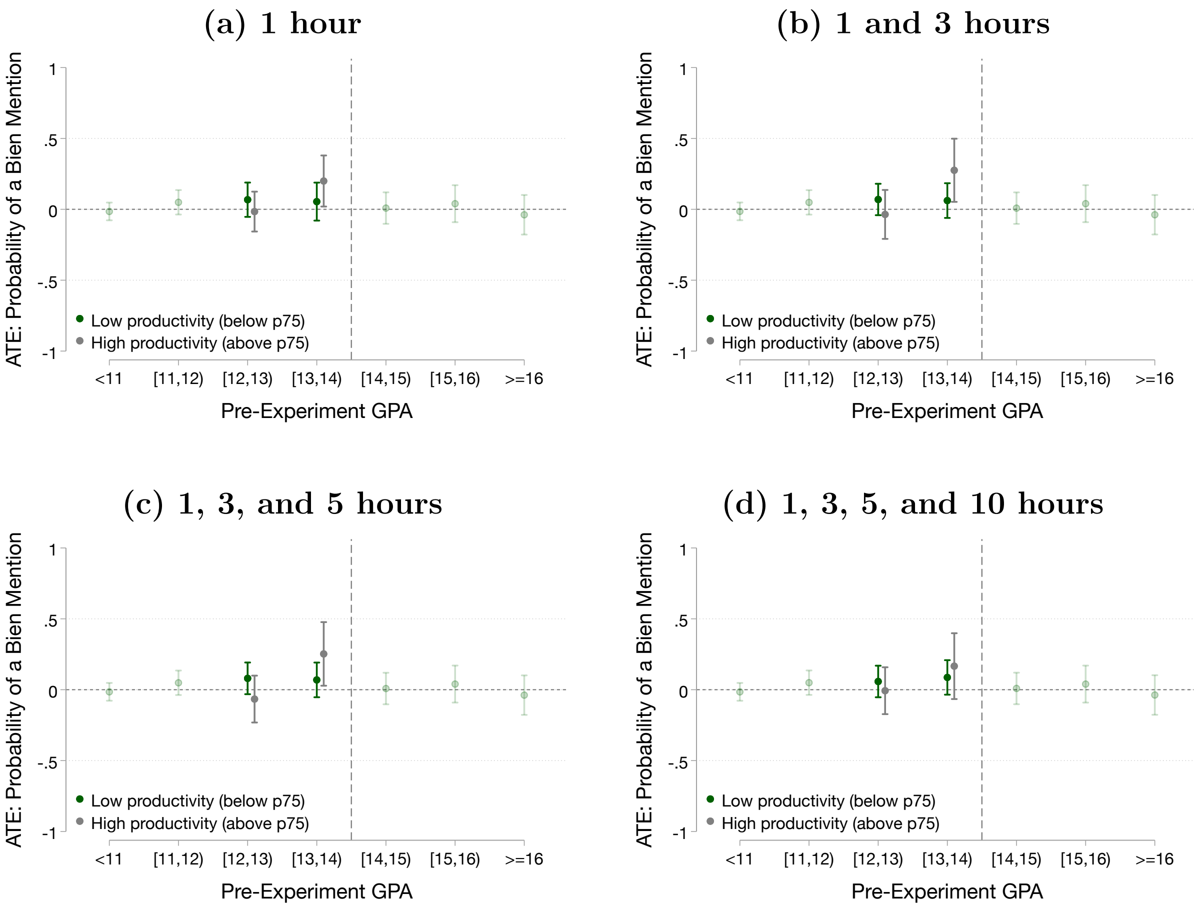

Notes: This figure shows the main heterogeneity analysis using alternative definitions of student productivity using the pre-treatment GPA–hours elicitation. Before the experiment, students reported their expected GPA under 1, 3, 5, and 10 additional weekly study hours. We compute the average GPA gain per hour: panel (a) uses only the 1-hour response; panel (b) averages gains per hour from the 1-hour and 3-hour responses; panel (c) averages gains per hour from the 1-hour, 3-hour, and 5-hour responses; panel (d) uses all four. In each panel, "High Productivity" denotes students above the 75th percentile of the relevant measure, while "Low Productivity" denotes those below. High-productivity status is restricted to students in GPA bins [12,13] and [13,14]. Each point represents the estimated treatment effect on High Honors attainment (Bien Mention) for the corresponding GPA bin, with 95% confidence intervals. All specifications include pre-registered controls. Robust standard errors.

Figure A.8: Robustness: Placebo tests in alternative thresholds (12 and 16)

(a) 12 (Honors) threshold   **(b)** 16 (Highest Honors) threshold

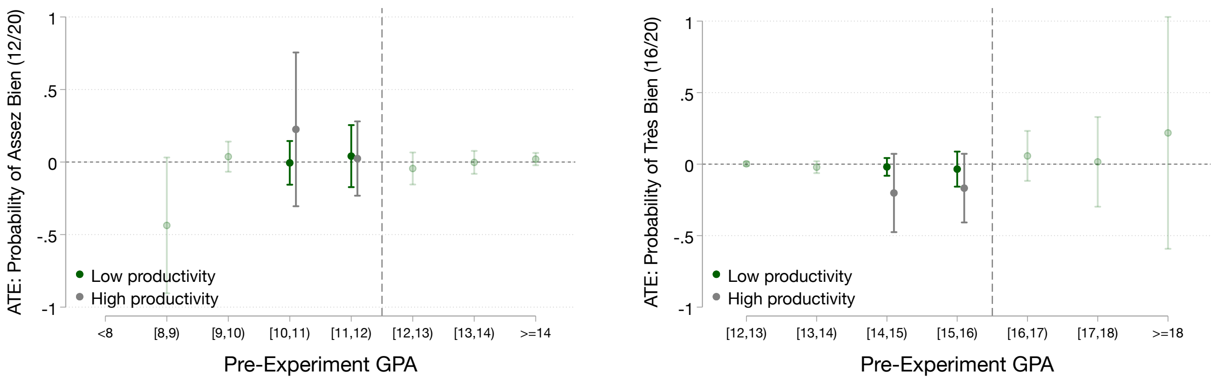

Notes: This figure replicates the main heterogeneity analysis at two placebo thresholds: 12/20 (Honors, panel a) and 16/20 (Highest Honors, panel b). Our treatment informed students about the earnings premium at the 14/20 threshold only. If the main result reflects a threshold-specific response to information rather than a generic motivation effect, we should find no treatment effects at other thresholds. Each point represents the estimated treatment effect on crossing the relevant threshold for the corresponding GPA bin, with 95% confidence intervals. "High productivity" denotes students who report needing at most one additional study hour per week to reach the relevant threshold. All specifications include pre-registered controls. Robust standard errors.

Table A.1: Robustness: Treatment Effects Using Alternative Proximity Definitions

	High Honors (14/20) Attainment
Treatment × Close (baseline: top tercile)	11.421***
	(4.419)
Treatment × Close (within 0.25 GPA point)	15.345*
	(8.637)
Treatment × Close (within 0.5 GPA point)	10.762*
	(6.466)
Treatment × Close (within 1 GPA point)	12.104**
	(4.814)
Controls	Yes
Observations	1,443

Notes: Each row reports results from a separate regression of High Honors attainment (in percentage points) on treatment interacted with proximity to the 14/20 threshold. Panel A interacts treatment with continuous distance to 14 (in GPA points). Panel B reports the treatment effect for students classified as under alternative definitions: the top tercile of midsemester GPA (baseline), within 1 GPA point of 14, within 0.5 GPA points, and within 0.25 GPA points. All specifications include pre-registered controls and major fixed effects. Sample restricted to students below 14/20 at midsemester. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.2: Distribution of Majors across Samples

Major	**Experiment Sample (%)**	**University Population (%)**
Law	26.7	29.5
Business & Management	16.4	23.4
Economics	14.5	7.0
Political Sc.	8.7	8.0
Arts & Humanities	9.5	10.0
Geography	6.6	4.5
History	6.1	5.3
Philosophy	3.7	5.3
Other / Unclassified	8.0	7.0
Observations	1,827	14,595

Notes: This table reports the distribution of students across majors in the experiment sample and in the university distribution. Source: Authors' calculation from University administrative records.

Table A.3: Robustness: Productivity and Self-Confidence

	(1)	(2)	(3)
Panel A: Does confidence predict elicited productivity?			
Correlation (productivity, confidence): 0.047			
Dep. var.: Elicited productivity (GPA gain per hour)			
Confidence index		0.038	
		(0.023)	
Observations		691	
Panel B: Predicting GPA growth in the control group			
Dep. var.: Δ GPA (control group, near-threshold students)			
Elicited productivity	0.024**		0.024**
	(0.011)		(0.011)
Confidence index		−0.001	−0.000
		(0.006)	(0.006)
GPA control	Yes	Yes	Yes
Additional controls	Yes	Yes	Yes
Observations	260	260	260

Notes: In this table, Panel A reports the unconditional correlation between elicited productivity and the confidence index, and coefficient of a regression of elicited productivity on the confidence index, controlling for midsemester GPA and additional covariates. The sample is restricted to near-threshold students. Panel B restricts the sample to untreated (control group) students near the 14/20 threshold and shows coefficients from a regression of Δ GPA (midsemester to year-end) on: elicited productivity only (column 1), the confidence index only (column 2), and both together (column 3). All specifications control for midsemester GPA, survey quality, job prospects, and knowledge about the honors system. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.4: Decomposition of Near-Threshold Students (GPA 13–14)

	Share (%)	Avg. cost (€)	Avg. perceived return (€)	Avg. net gain (€)
Panel A: High productivity ($\leq$ 1 hr/week to reach 14)				
Efficiently non-crossing	0.0	—	—	—
Undersignalers	34.7	112	-2	4,263
Informed non-crossers	65.3	110	2,062	4,265
Panel B: Low productivity ($>$ 1 hr/week to reach 14)				
Efficiently non-crossing	0.0	—	—	—
Undersignalers	56.2	736	111	3,639
Informed non-crossers	43.8	602	2,518	3,773

Notes: This table decomposes students with midsemester GPA between 13 and 14 — those within 1 GPA point of the High Honors threshold — based on their elicited production function and pre-treatment beliefs. Panel A restricts to high-productivity students (those reporting that one additional study hour per week suffices to reach 14/20); Panel B to low-productivity students. Individual signaling costs are computed by interpolating the elicited GPA–hours schedule, valued at €12/hour over 13 weeks. : cost exceeds the true return (€4,375). : cost falls below the true return but above the perceived return. : perceived return exceeds cost, yet the student did not cross. True return from Astruc-Le Souder et al. (2024). Perceived returns: $\widetilde{\Delta w}_i = \text{prior}_i \times$ €35,000.

Figure A.9: Invitation email (EN)

Dear Gedeao,

As part of a postdoctoral research project on future career prospects and the conditions of entering the job market for students in France, we invite you to take part in this survey. Conducted by a researcher affiliated with the CES (Centre for Economic Studies of the Sorbonne), this study aims to better understand the factors influencing career orientation, aspirations, and expectations of students in a constantly evolving socio-economic context. Your responses will help identify the main challenges you face and contribute to the development of public policies tailored to support the transition to the workforce.

The survey will only take 5 minutes. To thank you for your participation, **you will be entered into a lottery with 5 prizes of €200 (in the form of Amazon gift cards)**. Participation is voluntary and requires providing your email address. We noticed that you have not yet completed the questionnaire and remind you that it is still available should you wish to participate. **Please note that your responses will only be valid until Saturday (03/22).**

To participate, please click on the link below:
https://enquete.univ-paris1.fr/529934?token=cWQDihmkj639K9t&lang=fr

Notes: This is the email sent via the University system to 14,595 masters students enrolled in master programs in the 2024-2025 academic year.

Figure A.10: Invitation email (FR)

Cher(e) Gedeao,

Dans le cadre d'une recherche postdoctorale consacrée aux perspectives d'avenir et aux conditions d'insertion sur le marché du travail pour les étudiants en France, nous vous invitons à participer à cette enquête. Menée par un chercheur attaché au CES (Centre d'Économie de la Sorbonne), cette étude a pour but de mieux comprendre les facteurs qui influencent l'orientation professionnelle, les aspirations et les attentes des étudiants dans un contexte socio-économique en constante évolution. Vos réponses permettront d'identifier les principaux défis auxquels vous êtes confrontés, afin de formuler des recommandations visant à améliorer la transition vers le monde du travail et à soutenir le développement de politiques publiques adaptées.

L'enquête ne prendra que 5 minutes. Pour vous remercier de votre participation, **vous pourrez participer à une loterie avec 5 prix de 200 € (sous forme de cartes cadeaux Amazon)**. La participation est volontaire et nécessite de fournir votre adresse e-mail. Nous avons pris en compte que vous n'avez pas encore complété le questionnaire, et nous vous rappelons que celui-ci est toujours disponible si vous souhaitez participer. **N'oubliez pas que vos réponses seront valides jusqu'à samedi (22/03).**

Pour participer, veuillez cliquer sur le lien ci-dessous.
https://enquete.univ-paris1.fr/529934?token=cWQDihmkj639K9t&lang=fr

Cordialement,

Locks Ferreira, Gedeão - Chercheur en Économie (noreply@univ-paris1.fr)

Si vous ne souhaitez pas participer à ce questionnaire et ne souhaitez plus recevoir aucune invitation, veuillez cliquer sur le lien suivant :
https://enquete.univ-paris1.fr/optout/tokens/529934?token=cWQDihmkj639K9t&langcode=fr

Notes: This is the email sent via the University system to 14,595 masters students enrolled in master programs in the 2024-2025 academic year.

Figure A.11: First survey page

PANTHÉON SORBONNE
UNIVERSITÉ PARIS 1

Ce questionnaire n'est actuellement pas activé. Vous ne pourrez pas sauver vos réponses.

Enquête sur les Mentions

Bienvenu(e) dans l'enquête des étudiants de Paris 1 Panthéon-Sorbonne !

Nous vous invitons à participer à cette enquête conçue pour recueillir des informations sur la manière dont les étudiants en France perçoivent leur avenir sur le marché de l'emploi.

L'enquête ne prendra que 5 minutes, et pour vous remercier, elle permet la participation à une **loterie** à la fin de l'enquête **avec 5 prix de 200 € (cartes cadeaux Amazon) !** La participation est facultative et nécessite de fournir votre adresse e-mail à la fin de l'ênquete.

La participation à cette enquête entraîne le traitement de vos données personnelles conformément au RGPD. Pour plus de détails, vous pouvez consulter la mention d'information suivante ici.

Suivant

Notes: This is the first page after students clicked on the lik sent via the University system to 14,595 masters students enrolled in master programs in the 2024-2025 academic year.

B Full questionnaire (English Translation)

Welcome

Welcome to the Paris 1 Panthéon-Sorbonne student survey! We invite you to take part in this survey designed to gather information on how students in France view their prospects on the labor market.

The survey takes about 5 minutes. As a thank-you, you may enter a **lottery** at the end with **5 prizes of €200 (Amazon gift cards)**. Participation is optional and requires providing your email address at the end.

Participation involves processing your personal data in accordance with the GDPR. For details, see the information notice **here**. If you have questions, contact.

Consent:

- I agree to participate in this survey.
- I do not wish to participate.

C General Questions

In this section, we ask about your job-search expectations.

Target Employer Type (select all that apply)

- Self-employed
- Public sector (State, local, or hospital)
- Public company (La Poste, SNCF, EDF, France Télévisions, ...)
- Private company
- Non-profit or association
- Liberal profession / independent (law firm, notary office, ...)
- Private household

- I plan to continue studying (e.g., PhD)
- Other: ______________________________

Confidence Questions: Scale (1–4)

Please indicate how much you agree with the following statements on a 1–4 scale (1 = Not at all agree; 4 = Completely agree).

- I will perform well in a job in the sector I aim to work in.
- I am confident that I will succeed in the career I choose.
- I believe I have the necessary skills to perform well in my future professional tasks.

Salary Negotiation Atitude

In your view, how much room is there to negotiate the salary for the job you are aiming for after your master's?

- No room (fully fixed salary)
- Small room (minor adjustments possible)
- Negotiable (discussion possible)
- Large room (salary can vary significantly)

Importance of Pay

When choosing a job, how important is pay relative to other aspects (e.g., enjoyment, passion, work–life balance)?

- Not important
- Slightly important
- Fairly important
- Very important

D Studies, Performance, and Academic Honors

We now ask a few questions about your academic background.

Current Average (GPA 0–20)

Enter value: ____________

Follow-up Based on Stated GPA

Several numeric entries on the 0–20 scale:

- If your GPA is X, what would it be if you studied 1 more hour per week? ____________
- If you studied 5 more hours per week? ____________

We will now ask your opinion on the importance of academic honors (*Assez Bien*, *Bien*, *Très Bien*) obtained during the master's for career prospects.

How Are Honors Awarded?

Which statement best describes how academic honors are awarded to master's students in French universities?

- To students ranking near the top of their cohort
- Based on extracurricular activities
- To students whose overall average exceeds a given threshold
- Upon recommendation from some instructors

Honors Cutoffs (Enter values 10–20)

In France there are three honors: *Assez Bien*, *Bien*, *Très Bien*. Each is awarded when a student's overall average exceeds a given score.

- Assez Bien: ____________
- Bien: ____________
- Très Bien: ____________

E Honors (Information Module Setup)

In France, academic honors are awarded at these overall average thresholds: **Assez Bien $\geq$ 12/20**, **Bien $\geq$ 14/20**, **Très Bien $\geq$ 16/20**. We now focus on the honor *Bien* obtained during the master's and its impact on employment prospects in France.

Perceived Importance of *Bien* in Recruiting

- Very important
- Quite important
- Of little importance
- Not important

Perceived Effect on Starting Salary (General)

By how much does obtaining *Bien* affect first-year salary compared to an otherwise identical candidate without it? Enter a percentage between –50 and +50:

Perceived Effect on Own Salary

For you personally, by how much would obtaining *Bien* affect your own first-year salary? Enter a percentage between –50 and +50:

Confidence in Answer

- Very confident
- Rather confident
- Not very confident
- Not at all confident

F Experimental Module (Information Treatment)

A randomly selected group of participants will receive information related to the previous question. Continue to the next screen to learn whether you will receive this information.

Control Message

You have been randomly selected **not** to receive the information. The rest of the questionnaire will be available in 3 seconds. Please proceed.

Treatment Message

You have been randomly selected to receive the following information:

- A recent academic study of former master's students at the University Paris 1 Panthéon-Sorbonne found that obtaining the honor **Bien** is **associated with a 12.5% increase in cumulative income** during the first year after graduation.
- This is equivalent to **about 1.5 extra months of salary** during that first year.

The rest of the questionnaire will be available in 3 seconds.

G Post-Information Re-evaluation

We now give you the opportunity to revise some of your previous answers. This opportunity is offered to all participants regardless of their earlier responses.

Perceived Importance of *Bien* in Recruiting (Again)

- Very important
- Quite important
- Of little importance
- Not important

Perceived Effect on Starting Salary (General, Again)

Percentage between –50 and +50:

Perceived Effect on Own Salary (Again)

Percentage between –50 and +50:

Confidence in Answer (Again)

- Very confident
- Rather confident
- Not very confident
- Not at all confident

H Intentions

We would like to understand your study routines and the effort you plan to invest to reach your goals.

Planned Effort to Raise Average

- A lot of effort
- Some effort
- A little effort
- No effort

Allocation of 10 Extra Hours Next Week

Numbers must be between 0 and 10 and sum to 10.

- Studying / coursework: ____________
- Other activity: ____________

- Other activity: ____________

Planned Weekly Study Hours

On average, how many hours per week do you plan to devote to your studies (outside classes and extracurriculars) in the coming weeks? ____________ hours

I Final Questions

Do You Aim to Obtain an Honor? If Yes, Which One?

- Assez Bien (12/20)
- Bien (14/20)
- Très Bien (16/20)
- No

How Much Additional Effort is Needed?

- No additional effort
- A small additional effort
- A moderate additional effort
- A significant additional effort

Will You List *Bien* on Your CV or LinkedIn?

If you receive the honor *Bien* in your master's, how likely are you to list it on your CV or LinkedIn when applying for jobs?

- Very likely
- Rather likely
- Rather unlikely
- Very unlikely

Time Preferences Task

Now choose between two options. The first offers a lower payment today; the second offers a higher payment in the future.

- €300 today vs. €300 in one month
- €280 today vs. €300 in one month
- €250 today vs. €300 in one month
- €200 today vs. €300 in one month
- €150 today vs. €300 in one month
- €120 today vs. €300 in one month

Current Status

- Continuing studies (full-time student)
- Internship/apprenticeship
- Employed full-time
- Employed part-time
- Unemployed (looking for work)
- Other

Work Experience (Select All That Apply)

- Internship
- Part-time job
- Full-time job
- No work experience

J Lottery

Thank you for your time. You may now enter the prize draw for five €200 Amazon gift cards. To enter and/or receive the study results, please tick the box below if you agree to be contacted; otherwise, you can simply provide your email and close this window.