

Social Bee Lab


DISCUSSION PAPER SERIES

SBL-DP-00002

Closing class gaps with simple nudges: Experimental evidence on opportunity allocation

Manuel Munoz

July 2025



Closing class gaps with simple nudges: Experimental evidence on opportunity allocation

Manuel Munoz^{* †}

Abstract

In this paper I evaluate if an attentional nudge can reduce socio-economic bias in the allocation of opportunities? I partnered with an organization to run a field experiment involving 528 nominators and 6,098 potential candidates. Nominators were invited to propose candidates for an international training program; among those they are connected to in the intra-organizational network. They were either asked to focus on performance or additionally nudged not to overlook talented candidates from disadvantaged backgrounds. The nudge did not discourage nominators' propensity to make a recommendation. Instead, it raised the share of low-SES nominees from 23% to 31%, fully closing the representation gap (as 1 out of 3 candidates are low-SES). This equity shift occurred without lowering candidate quality or affecting nominations of other groups (e.g., women). In a follow-up study, I explore mechanisms and found additional evidence that the salience nudge shifts attention, which drove low-SES nominations. Because nudges cost virtually nothing, preserve gatekeepers' autonomy, and can fit within existing workflows, they offer a scalable path to address raising problems of class inequality in the allocation of opportunities within organizations.

Keywords: Social class, Attention biases, Educational Access, Field Experiment, Social Networks

JEL Classification: C93, D83, I24, J71

^{*}New York University Abu Dhabi, Abu Dhabi, United Arab Emirates, and Universidad Autonoma de Bucaramanga, Bucaramanga, Colombia (e-mail: manumunoz@nyu.edu).

[†]I am grateful to Michelle Acampora, Ernesto Reuben, Pedro Rey-Biel, Sharlane Scheepers and Robert Stuber for helpful comments. I thank Jhon Diaz and Claudia Gomez from UNAB for their invaluable institutional support. The experiments reported in this paper obtained ethical approval at UNAB and LISER and were pre-registered at AEA RCT Registry under numbers AEARCTR-0014135 and AEARCTR-0015473. I used OpenAI's ChatGPT for light copy-editing.

1 Introduction

Affirmative action policies are commonly used to address disparities and widen access to valuable opportunities. However, these policies usually impose quotas or strict selection rules, which limit the autonomy of those responsible for selection processes (Beaman et al. 2018; Bolte et al. 2024; Alesina et al. 2024). Such restrictions frequently lead to resistance from decision-makers (Ahern and Dittmar 2012), potentially undermining the long-term effectiveness of affirmative action interventions, and sometimes reversing gains achieved for disadvantaged groups.

This issue is particularly relevant within organizational settings, where opportunities for career advancement, such as promotions, mentoring, or specialized training, are typically allocated through internal nomination processes (Diaz et al. 2025). Gatekeepers, such as managers, team leaders or department heads, often nominate individuals who are most visible to them, frequently favoring candidates of higher social status (Fernandez et al. 2000).¹ This occurs even when equally qualified but less visible, lower-status candidates are also available (Shukla 2022). The resulting selection biases reinforce existing inequalities in career trajectories and limit the diversity of organizational talent pools (Granovetter 1974; Topa 2001; Dustmann et al. 2016). Thus, evaluating interventions that shift gatekeepers' attention toward equally capable but less visible candidates, is essential to improving equity in opportunity allocation (see e.g., Link et al. 2025).

In this paper, I present results from a natural field experiment designed to test whether a nudge can encourage gatekeepers to consider qualified candidates from low-status backgrounds when nominating individuals for a valuable opportunity. Unlike traditional affirmative action approaches, the nudge does not alter eligibility criteria (Thaler and Sunstein 2009) and the message it sends does not label nominated candidates as quotas (Leibbrandt and List 2018). Instead, it aims to shift gatekeepers' attention toward capable candidates who otherwise may be easily overlooked due to their lower-status backgrounds. This is of particular importance given the raising problems of class discrimination that have been identified within organizations (Stansbury and Rodriguez

¹ Evidence shows that evaluators rely on cues linked to status (e.g., accent, leisure tastes, social networks) as proxies for “cultural fit” within organizations. Cultural fit, in turn, operates as a powerful screening device, making high-status candidates more salient and systematically favored in selection processes (see e.g., Rivera 2012; Friedman and Laurison 2020; White-Lewis 2020).

2025; Shukla 2022)

To evaluate if a salience nudge can effectively address such selection biases within organizations requires overcoming various empirical challenges. It requires (I) identifying opportunities of meaningful value that can only be accessed through nominations from gatekeepers; (II) tracing an intra-organizational network that contains the history of interactions between gatekeepers and potential nominees; (III) obtaining reliable and detailed measures of both social status and candidate quality; and (IV) randomly varying the communication sent to gatekeepers, to embed a salience nudge for some but not all, as part of the organizational workflow.

To address these empirical requirements, which are rarely satisfied simultaneously, I carried out a field experiment as part of a collaborative partnership with a university in Colombia. Colombia is an ideal setting to study class discrimination, as the country has a stratification system implemented by the central government, which assigns every household to one of six strata on the basis of objectively verifiable dwelling and neighborhood characteristics. As such, it provides an explicit social class measure that allows me to identify how status affects the allocation of opportunities within organizations.²

In this setting, I launched an international training program for students of the local partner university. Access to the training program was obtained only through nominations from faculty (i.e., requirement I). Faculty members at the local university were invited to nominate one student each for the training program, provided they had previously taught the nominee at least once. This allows me to use course enrollment registries to map the connections and history of interactions between nominators and potential candidates: the intra-organizational network (i.e., requirement II). The partnership further gave me access to extensive administrative data covering 6,098 students and 528 faculty members, including indicators of socioeconomic status, gender, and academic performance (i.e., requirement III). Faculty were invited by the university administrators to nominate students for the training program, as part of an internal selection process. They were randomly assigned to either a control group (CTRL), receiving instructions to nominate candidates by focusing solely on their academic performance, or a treated group (SALIENCE) that additionally received a nudge explicitly reminding nominators about candidates from low-status groups (i.e., requirement IV).

² For details on Colombia's unique stratification system see Section 2.1.

This design allows me to evaluate the efficacy of the salience nudge in the field by measuring three outcome variables. First, *backlash*: does the nudge produce resistance from gatekeepers by lowering the overall nomination rate? Second, *equity shift*: does the nudge raise the representativeness of low-status candidates compared to how representative they are in the control group? Third, *quality trade-offs*: does any equity gain from a higher representation of low-status nominations come at the cost of lower candidate performance or weaker participation in the training program?

The main findings of the experiment are as follows: First, there is no backlash. The salience nudge did not discourage invited nominators from participating in the selection process. Nomination rates are indistinguishable between conditions: 48% in the CTRL group and 50% in the SALIENCE group. This supports the conjecture that such a low-touch prompt does not trigger resistance from decision-makers.

Second, there is an equity shift. At the local university, low-status candidates (i.e., low-SES students) constituted approximately one-third of the undergraduate body, but only received 23% of nominations in the CTRL. With the introduction of the SALIENCE nudge, this proportion increased to 31%, effectively closing the class representation gap. The nudge altered the demographic composition of nominated candidates and corrected selection biases, which remains robust when controlling for faculty network characteristics.

Third, there are no quality trade-offs. Nominated candidates in both the CTRL and SALIENCE groups ranked at the top of the performance distribution (i.e., GPA), and mean performance scores were statistically indistinguishable between conditions. Also, take-up and completion rates of the training program were equally high, between conditions. Thus, confirming that the observed shift in selection represents a reallocation of opportunities rather than a weakening of candidate standards.

These core results persist when I control for the history of interactions between nominators and candidates, including frequency of prior interactions, elapsed time since shared encounters, and candidate visibility. This further supports the conjecture that the nudge itself rather than the network structure drives the change.

The findings from the field experiment suggest that the nudge operated by shifting attention to low-SES candidates, who are otherwise less visible regardless of their share in faculty networks. However, the observed shift in nominations could be explained to

some extent through alternative channels that cannot be ruled out conclusively with the design and data from the main field experiment. Rather than simply heightening awareness of low-SES candidates, the nudge might have operated by signaling the presence of biases, leading to a reaction from faculty to correct or prove the bias wrong. It may also have done it by signaling institutional expectations. Faculty might then have nominated more low-status students not because of heightened salience, but as compliance to what they interpreted as a directive from university administrators.

To better disentangle these mechanisms and address the identified limitations, I designed and implemented a second study three years later, employing a similar nomination framework. This time, I used a different opportunity that has fewer requirements and no ambiguity about its value for low-SES candidates of all socio-economic backgrounds.³ The opportunity was a student excellence award, which involved a certificate and cash prizes for candidates who win the award. In this case, nominators were randomly assigned to one of four conditions. As in the first study, the CTRL invites nominators to focus on candidate quality as the main selection criterion. A SALIENCE nudge, resembling that of the original experiment, reminds nominators to consider low-status candidates. A DEBIASING nudge encourages reflection on unconscious biases against low-status candidates. Finally, an INSTITUTIONAL nudge that appeals to the university's mission of diversity and inclusion.

The experiment of the second study replicates the main result and sharpens its interpretation. The SALIENCE nudge, which is strictly a reminder, produces the largest increase in low-SES nominations, relative to the CTRL group. The nudge that asks faculty to reflect on unconscious biases has a smaller effect, and the nudge that cites the university's diversity policy has almost none. The pattern is consistent with a shift in attention rather than with moral persuasion or institutional pressure. Also, as in the first experiment, nominee performance scores are above average and indistinguishable between conditions. Thus, confirming that these types of nudges shape equity without hurting candidate quality.

The experimental findings of my work speak to three lines of research.

First, my study contributes to the work on *selection bias remediation*. Existing strate-

³ See a discussion of potential limitations of the training program as a *class-neutral* opportunity at the end of Section 3.4.

gies in the literature generally fall into either coercive mechanisms like quotas imposed on decision-makers (Alesina et al. 2024; Bolte et al. 2024; Beaman et al. 2018), or persuasive mechanisms like equal opportunity statements directed towards candidates (Leibbrandt and List 2018). While quotas can improve representation, they limit the autonomy of decision-makers and may induce resistance.⁴ Persuasive interventions do not restrict agency, but they can deter candidates of underrepresented groups by triggering stigma, given that people dislike feeling like quotas themselves (Bertrand and Mullainathan 2004; Agan and Starr 2018; Bohnet 2016).

To address these challenges, I propose and test a third pathway: a salience nudge directed at gatekeepers, which preserves autonomy while redirecting attention, by activating awareness of overlooked qualified candidates. The findings of my study show that such a nudge eliminates class gaps in the allocation of opportunities without backlash or performance trade-offs. This occurs because the nudge mitigates cognitive neglect of existing talent rather than altering selection criteria (Castilla 2008; Goldin and Rouse 2000).

Second, this paper contributes to the growing literature *integrating cognitive psychology into economic decision-making*, particularly through salience theory (Bordalo et al. 2012, 2016). The theory highlights that attention is context-dependent and drawn to attributes that stand out (see e.g., Link et al. 2025). This causes people to overweight salient features while neglecting less visible ones, when making judgements or decisions.⁵ To counteract such biases, existing research has tested interventions to help recalibrate attention by either reframing choice sets (Chetty et al. 2009; Thaler and Bernartzi 2004), highlighting typically neglected information (Bollinger et al. 2011; Allcott and Taubinsky 2015), or exposing decision-makers to counter-stereotypical examples (Goldin and Rouse 2000; Alesina et al. 2024).

My study contributes to this line of inquiry by showing how salience nudges can effectively mitigate selection biases in a realistic organizational environment, demonstrating that subtle and timely shifts in attention can lead to meaningful reductions in inequal-

⁴ Quotas may even impose quality trade-offs when the required level of representation (the size of the quota) is larger than the pool of available high-performing candidates from the targeted group (Ahern and Dittmar 2012; Niederle et al. 2013).

⁵ For example, in financial choices, investors fixate on extreme past returns (see e.g., Frydman and Wang 2020); in hiring, stereotypes may dominate if group traits appear statistically exaggerated (see e.g., Bertrand and Mullainathan 2004).

ity (see e.g., Arslan et al. 2025, for the case of hiring). When the salience nudge is introduced, it redirects attention toward overlooked candidates activating their *latent* availability (Small and Sukhu 2016).⁶ This suggests that representation may be necessary but not sufficient, and availability translates into opportunity only when paired with attentional shifts.

Third, this paper contributes to the growing body of evidence showing that social class operates as a distinct and understudied cause of inequality, one that rivals, and sometimes exceeds, gender and race gaps (Friedman and Laurison 2020). This is specially prominent in organizational settings where individuals from low-SES backgrounds are less likely to get promotions or raises (Rivera 2012; Stansbury and Rodriguez 2025) even when their credentials match those of high-status peers (Zimmerman 2019; Michelman et al. 2022; Shukla 2022).

My paper documents the presence of socio-economic biases in the allocation of opportunities within a university setting. It also provides field-experimental evidence that a salience nudge can heighten decision-makers' awareness of overlooked, low-SES candidates, which significantly improves their visibility. My work provides causal evidence of how nudges can narrow class gaps in the allocation of opportunities within organizations. It also shows how such nudges do not impose unintended spillovers on members of other groups, such as women.

The rest of the paper proceeds as follows. Section 2 describes the experimental design and hypotheses. Section 3 presents the main results, and Section 4 reports the follow-up study on mechanisms. Section 5 concludes.

⁶ The cognitive representation of networks is an important complementary line of research, which shows that latent (or potential) ties are contacts who exist in an individual's network but remain cognitively *off-line* until something makes them come to mind. The activation of contacts is highly sensitive to attentional cues (Simpson et al. 2011), and salience manipulations could change which connections people remember (Brashears and Quintane 2015). In my setting, this would mean that heightening the visibility of low-SES candidates can lead to real opportunity access (Small 2013; Small and Sukhu 2016; Menon et al. 2024).

2 The experiment

2.1 Setting of the study

This project is the result of a collaboration partnership with a university in Colombia (the local partner). Colombia is an ideal setting to study class discrimination as the country has an explicit social-stratification scheme assigned by the central government since the early 1990s. Every household receives a number from 1 (poorest) to 6 (wealthiest) that is printed on utility bills and determines the tariffs paid for water, electricity, and other public services. It also operates as a redistributive tax where lower strata receive subsidies that are financed by higher strata (see [Bogliacino et al. 2018](#)).

The system provides a granular and objective measure of socio-economic status that is known to virtually everyone in the country. For example, students routinely list their stratum on administrative forms, and the number is frequently used in public debate. Beyond its fiscal and economic role, the *estrato* has deep symbolic content. Evidence shows that Colombians can identify another person's stratum with considerable accuracy using cues such as accent, attire, and residential address ([García-Sánchez et al. 2018](#)). As a result, the scheme not only shapes access to subsidized goods but also signals social position in everyday interactions. The pervasive salience of strata means that class distinctions are quickly inferred on campus, in the workplace, and in social networks, reinforcing status hierarchies even when economic stakes are minimal ([Cardenas et al. 2021](#)).⁷

The stratification system also has important implications for the socio-demographic composition of the student body at universities, which differs greatly between private and public institutions. Public universities are almost exclusively for low income students because tuition fees are a function of family social strata, which means that those in lower strata pay very little and those in higher strata would pay substantial fees. In private universities there is no price discrimination, but among the private there are two types: elite and non-elite. Private elite universities are mostly for students from high income families, as they charge very high tuition fees (see e.g., [Londono-Velez 2022](#)).

⁷ India's caste categories provide the closest administrative analogue to this system, yet Colombia's *estrato* scheme cuts across ethnic boundaries (see e.g., [Shukla 2022](#); [Oh 2023](#)).

Private non-elite universities reach students from all social classes as their prices are intermediate, so their diversity levels are highest among the universities in the country. The local partner university is private but not elite, which makes it an ideal partner for my study.⁸

In 2022, at the time of the study, the local partner employed 528 faculty members and enrolled 6,098 undergraduate students from a diverse set of social groups and backgrounds (see [Cardenas et al. 2021](#)). One out of every three students (31%) comes from low-SES backgrounds.⁹ This provides a fitting environment to analyze selection-based mechanisms for expanding access to valuable opportunities, such as the international training program offered. Low-SES students outperform their high-SES peers academically by 0.15 standard deviations (GPA: 0.10 vs. -0.05, $p < 0.001$) but expect longer job searches, lower starting salaries on their first job (wages), and perceive their social connections made in college to be less effective in helping secure employment after graduation (for details see Appendix [D.1](#)). These discrepant expectations may lead to associations of low institutional fit despite objective academic success, making low-SES students a focal group for interventions aimed at widening access to high-return opportunities.

2.2 Features of the training program

The training program was offered as an opportunity for career progression. It aimed to strengthen self-efficacy to help students better set and pursue goals (see e.g., [Milkman 2021](#)).¹⁰ This program was an exclusive opportunity accessible only through faculty nominations. It was not announced publicly and students were not informed about it, so they had no way to request or influence faculty nominations. This allows me to observe

⁸ Universities in Colombia regularly report the average strata of their student population. Recall strata goes from 1 for those with the lowest income to 6 for those most affluent. Private elite universities have an average strata above 4 with the highest case being 5.4, public universities have average strata below 2.5 with the lowest being 1.2. Private non-elite universities have an average strata ranging between 2.5 and 3.5. The partner university reported having an average strata of 3.3. See <https://www.universidad.edu.co/de-mayor-a-menor-ies-colombianas-segun-el-estrato-socioeconomico-de-sus-estudiantes/>.

⁹ In Colombia, socioeconomic status can be generally divided into three categories: *Low* groups strata 1 and 2, *Middle* groups strata 3 and 4, and *High* groups strata 5 and 6. My focus in this project is on the allocation of opportunities to low-SES candidates, which I also refer at times as low-status candidates.

¹⁰ Self-efficacy refers to an individual's belief in their ability to accomplish tasks or goals ([Bandura 1978](#)).

a gatekeeping decision with real stakes in the field.¹¹ Its content was intentionally broad, benefiting students across academic programs, years of study, and demographic groups.¹²

To motivate nominations, faculty members were informed that the program was offered at no cost to participants and consisted of nine pre-recorded, 30-minute video sessions. Students could complete the sessions online at their own pace, allowing flexibility to accommodate varied schedules, so that people with jobs or dependents (e.g., low-SES students) would not be perceived as less fitting. On-campus labs and free Wi-Fi were available for students without reliable internet or personal devices.

Nominators were aware that participants who finished the program would receive an international certificate of completion, a credential valued in job, internship, and scholarship applications (see e.g., [Athey and Palikot 2024](#)).¹³ Also, participants who finished the entire program were entered into a lottery to win one of two iPads.

2.3 Treatments and sample

Faculty members were invited to nominate students for the international training program, as part of an internal selection process.¹⁴ This is a big advantage of the partnership, as the field study is embedded within an institutional initiative. In fact, the invitation email was sent by the Office of International Relations of the local university, an office that regularly offers international opportunities to students, making this

¹¹ At the time of this experiment, the only way to access the program was through faculty nominations. After this experiment concluded, other students were invited to the training program through different channels, as reported in [Munoz \(2024\)](#).

¹² Unlike specialized programs that focus on cognitive skills, such as coding or advanced mathematics (see e.g., [Carlana and Fort 2022](#)), this program emphasized broad, transferable skills to ensure its relevance across diverse fields and social groups.

¹³ The entire training program was sponsored also by a global American university. That is why the certificate of completion is *international*, which has high value for those completing the program.

¹⁴ This experiment aligns with the definition of a natural field experiment as outlined by [Harrison and List \(2004\)](#). A natural field experiment involves participants who are unaware they are part of a study, with experimental conditions closely mirroring real-world settings. In this case, faculty were invited to nominate students as part of an internal institutional process, with no indication that the nomination decisions were part of an experimental study. For details see Appendix [E](#).

process fit naturally within the organizational workflow.¹⁵ The invitation emphasized the program's prestige and outlined its tangible benefits, including an internationally recognized certificate of completion (see Appendix A.1 for the full invitation text). Faculty members could recommend any student as long as that student had taken at least one course with the nominator.¹⁶ To test the impact of the salience nudge, I randomly assigned nominators to one of two conditions:

- CTRL condition: Recipients read the following instruction in their invitation message:

We ask you to respond to this message with the full name and student code of the person you want to recommend, considering that this person has the academic performance to benefit from this great opportunity.

- SALIENCE condition: Recipients saw the same text plus one sentence designed to make low-status candidates salient:

We ask you to respond to this message with the full name and student code of the person you want to recommend, considering that this person has the academic performance to benefit from this great opportunity. For your recommendation, please focus on the academic performance of the students and do not exclude anyone because of their socioeconomic background, gender or any other demographic characteristic.

Nominators in both cases were asked to select candidates based on academic performance. On top of that, the salience nudge was designed to encourage more equitable nominations by prompting nominators to consider candidates from diverse social groups, without compromising such academic performance.¹⁷

¹⁵ Note that the Office of International Relations has no authority over faculty affairs, which aims to prevent the invitation from being perceived as a directive that would restrict faculty's autonomy. For a comment of this potential limitation see the discussion at the end of Section 3.4.

¹⁶ By requiring that nominated students had taken at least one course with the faculty member recommending them, I can use course registries to construct weighted faculty-student ties based on shared courses. This results in an enrollment network that allows me to conduct a detailed analysis of the relationship between network attributes (e.g., tie strength, student availability) and nomination choices, as described in Section 2.4.

¹⁷ When designing the nudge, university administrators raised their concern that calling instructors' attention to students from disadvantaged socioeconomic backgrounds could unintentionally reduce nominations of other groups, especially women. To avoid potential spillovers, they required the nudge to explicitly remind faculty not to exclude by gender. The formatting of the nudge allows me to test for spillovers across social groups.

Table 1 Balance table on treatment assignment

This table compares individual demographics and network characteristics between those nominators assigned to the CTRL and the SALIENCE condition. p-values indicate significant differences.

	CTRL I	SALIENCE II	p-value III
<i>Socio-Demographics</i>			
Nominator is female	0.40 (0.49)	0.45 (0.49)	0.222
Nominator is married	0.51 (0.50)	0.50 (0.50)	0.856
Nominator has permanent contract	0.58 (0.49)	0.54 (0.49)	0.265
Nominator holds graduate degree	0.87 (0.33)	0.84 (0.36)	0.247
Nominator's years at university	9.30 (7.79)	8.30 (7.32)	0.131
<i>Networks</i>			
Average number of ties (degree)	199 (164)	176 (142)	0.092
Frequency of contact with ties	1.60 (0.62)	1.51 (0.51)	0.046
Share of ties with female students	0.56 (0.17)	0.55 (0.17)	0.455
Share of ties with low-SES students	0.32 (0.17)	0.32 (0.15)	0.738
Average GPA of ties	0.79 (0.03)	0.79 (0.03)	0.602
Number of observations	258	270	

Of the 528 invited faculty members, 258 (49%) were assigned to the CTRL condition and 270 (51%) to the SALIENCE condition. I randomized faculty by stratifying key demographics and job characteristics, including gender, marital status, tenure (i.e., permanent contract), education degree (i.e., graduate or not) and time of employment. Table 1 reports that randomization was balanced across demographics. It also shows post-randomization network measures and confirms that most of these variables were balanced, although faculty members in the CTRL group were connected to slightly more students, i.e. average number of ties, and have slightly more courses with them, i.e. frequency of contact (for details see Table B-2 in Appendix B).

Once the invitation was sent, nominations remained open for three weeks. After that window closed, the Office of International Relations contacted each nominee to invite them to enroll in the training program. This is the first time students were made aware of the offered opportunity.

2.4 Data

The collaboration agreement with the local partner university gives me access to work with the following datasets for this project:

Course-enrollment network data. This data links 528 faculty to 6,098 students through course enrollment registries. There are 99,272 links in the resulting network, which are weighted by the number of shared courses (including measures of class size and time when the course took place). The network defines each nominator’s choice set and the strength of every nominator-candidate connection.

Administrative data. This data provides demographics (gender, SES), academic performance (GPA), and faculty records to qualify each candidate. It also allows me to construct availability measures, such as a nominator’s share of low-SES or female students in her network.

Experimental data. For every faculty-student pair, I have data from the field experiment indicating whether a nomination occurred and the nominator’s treatment assignment (CTRL or SALIENCE).

Training-program data. These are records of take-up and completion for all nominees, which serve as post-selection quality measures.

By integrating these datasets, I have a unique panel that provides the structure needed to test the effect of the salience nudge on backlash, equity shift, and quality trade-offs.

2.5 Hypotheses

Next, I give a brief description of the predictions for the experiment, summarized in three hypotheses. The interested reader can find a formal model of gatekeeper’s nomination choices and the impact of the salience nudge in Appendix C, from which the hypotheses are derived.

The first prediction focuses on potential backlash or reduced participation due to the intervention. There is backlash if the share of faculty who submit a recommendation falls in the SALIENCE condition relative to the CTRL.

Hypothesis 1 *If the nudge triggers backlash, then nomination rates will be lower in the treated group than in the control.*

Notably, if the salience nudge does not alter the structure or criteria of nominations, it should not impose additional psychological or administrative costs, and therefore, nomination rates should remain unaffected.

The second prediction focuses on the equity shift due to the salience nudge. By increasing cognitive salience of low-status candidates, the nudge is conjectured to reduce selection biases and to increase their representation among nominees.

Hypothesis 2 *If the nudge promotes an equity shift, then the selection bias against low-status candidates will be more attenuated in the treated group than in the control.*

Hypothesis 2 implies that the share of low-status nominees is predicted to match closer their availability in the network in SALIENCE than in the CTRL.

The third prediction focuses on potential trade-offs in nominee quality. I assess this with two proxies, each nominee's GPA at the time of nomination and their subsequent take-up and completion of the training program.

Hypothesis 3 *If the nudge imposes quality trade-offs, then candidate quality will be lower in the treated group than in the control.*

Hypothesis 3 is motivated by the conjecture that as the salience nudge does not change the valuation of candidate quality, gatekeepers are expected to continue prioritizing performance equally across conditions.

3 Results

In this results section, I first assess whether the salience nudge affects the overall nomination rate (Hypothesis 1), then assess the predicted equity shift in the share of low-status nominees (Hypothesis 2), and examine whether there are trade-offs on candidate quality (Hypothesis 3). I close this section with a network analysis of what determines the selection choices nominators make. For this section, I run difference-in-

means tests and regression specifications. I report two-sided p-values from the analysis in the text and provide regression tables in Appendix B.

3.1 Backlash: who makes a nomination?

The first step is to test whether the salience nudge discourages nominators from participating in the selection process. I invited 528 nominators, out of which 259 (49%) recommended a candidate for the training program. Figure 1A illustrates that about one of every two nominators recommends a candidate for the opportunity across conditions: 48% in CTRL and 50% in SALIENCE (two-sided proportion test, $p = 0.536$). The evidence does not support Hypothesis 1, as the nudge does not change the likelihood that invited nominators recommend a candidate for the training program. That is, I find no evidence of backlash.

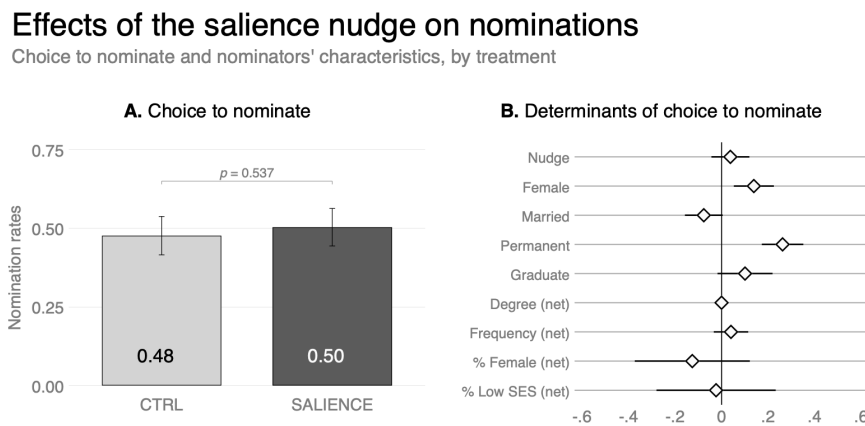


Figure 1 Effects of the salience nudge on nominations.

Panel A compares nomination rates across treatments. **Panel B** displays differences in the likelihood of making a nomination using faculty characteristics as covariates, with diamonds showing point estimates and bars representing 95% confidence intervals.

I also check which faculty attributes are determinant of the choice to nominate (for details see Table B-2 in Appendix B). Figure 1B illustrates that female professors, those on permanent contracts (i.e., the equivalent of tenure at the local university), and faculty holding a graduate degree are all significantly more likely to recommend a student for the training program, whereas married faculty are modestly less inclined to do so. By contrast, neither the size (i.e., Degree (net)) nor the demographic make-up of a faculty's student network (% Female (net) or % low SES (net)) meaningfully affects

participation.¹⁸

The main result from this section is summarized below:

Result 1 *The salience nudge does not reduce faculty participation in the nomination process.*

3.2 Equity shift: who gets the nomination?

Next, I test whether the nudge shifts who receives a nomination, as predicted by Hypothesis 2. For each faculty member, I compare the share of low-SES nominees with the share of low-SES candidates available in the nominators' networks. When the nominee share falls below the network share, I label the gap as *inbreeding bias*, meaning nominators favor higher-status candidates even after accounting for group availability (see e.g., Currarini et al. 2009, 2010).¹⁹ As a control, I carry out the same analysis for female students, which allows me to test for unintended effects on other social groups.

Social Class. Figure 2A shows that nominators in the CTRL condition nominate low-SES candidates 23% of the time even though they account for 31% of the nominators' networks. The 8 p.p. gap is statistically significant (two-sided t -test, $p = 0.029$) and is suggestive of inbreeding bias in the CTRL condition.

Under SALIENCE the share of low-SES nominations rises to 31%, matching their share in the networks (two-sided t -test, $p = 0.848$). This result aligns with Hypothesis 2: when low-status groups are made salient, the bias is attenuated and the nomination rate for low-SES candidates moves up. In this case, it goes up to the level of its availability.²⁰

¹⁸ Note that gender, contract status, and graduate-degree attainment are the very covariates I used to stratify the randomization into treatments, so their predictive power does not compromise causal inference. This confirms that the design balanced these key characteristics across conditions while letting them operate as first-stage predictors of nomination behavior.

¹⁹ Group availability at the university is such that low-SES students comprise roughly 30 percent of the population, so some inequality is mechanical (i.e., *baseline imbalance*). Inbreeding bias refers to any further shortfall that appears after I adjust for the availability of low-status students in each faculty member's networks. That is why it is most natural to test differences in the share of nominations against the share of low-SES available in the networks of faculty nominators.

²⁰ For this analysis, I use two-sided t -tests as the structure of the data collapses network measures into averages. In Section 3.4, I carry out a more rigorous analysis accounting for heterogeneity in faculty networks using a multinomial mixed logit model (McFadden and Train 2000), which leads to consistent results.

Effects of the salience nudge on low-SES nominations

Nominations vs. availability of low-SES candidates, by treatment




Figure 2 Effects of the salience nudge on low-SES nominations.

Panel A compares the share of low-SES students who receive a nomination with the share available in faculty networks. **Panel B** displays treatment differences in the likelihood of nominating a low-SES candidate using faculty characteristics as covariates. **Panel C** plots the relationship between low-SES availability and the probability of nominating a low-SES student. Numbers next to the plotted lines represent the count of faculty members in each bin, by treatment.

On the right-hand side (see Panel B), the figure reports results from a linear regression analysis testing what drives the nomination of low-SES candidates, by treatment (for details see Table B-4 in Appendix B). Availability (the percentage of low-SES candidates in the network) raises the likelihood of selecting a low-SES candidate, yet it is statistically significant only in SALIENCE. When the nudge attenuates the bias, nominators respond more strongly to the presence of low-SES candidates in their networks, but without the nudge they do not. The nudge remains effective even if I restrict the sample to low-SES students who are also top performers, defined as having a GPA above

median (for details see Appendix D.2).²¹

Panel C further illustrates the result on availability. The figure orders faculty into four bins of low-SES availability and plots the corresponding nomination rates.²² In the CTRL group the line is flat, which is aligned with a selection bias that is not affected by availability alone. Being exposed to higher shares of low-status candidates does not appear to be enough to correct the bias. In SALIENCE, the line tilts upward with availability. Nominators who are connected to more low-SES candidates in their networks react more strongly to receiving the nudge. This suggests that low-SES representation may be necessary but not sufficient and availability translates into opportunity only when paired with attentional shifts.

Gender. Figure 3A shows that women account for 55% of the candidates in faculty networks and receive 62% of nominations in CTRL and 61% in SALIENCE (two-sided tests, $p = 0.097$ and $p = 0.163$). This supports the conjecture that there is no negative bias against female candidates in this setting, and there are no unintended spillovers from the nudge. Figure 3B confirms that availability drives the selection of female candidates: the coefficient on the nominators' share of females in their networks is positive and statistically significant in both treatments, and the two coefficients are virtually identical.²³ This is further confirmed in Figure 3C. It shows that in each bin the nomination rate for women rises as a function of availability, and the slopes coincide across the two conditions. That is, the salience nudge leaves this already proportional

²¹ Socioeconomic background is not part of the administrative records at the local university for faculty, although it is for students. Despite potential concerns, this is unlikely to threaten the interpretation of the experiment. Evidence suggests that as organizational norms that equate *fit* with high-status are expected to be internalized, class-based penalties are likely to emerge even when evaluators come from mixed backgrounds (Bisin and Verdier 2011; Turco 2010; Castilla 2015; Castilla and Ranganathan 2020; Shukla 2022).

²² I group faculty into four intervals of the distribution of low-SES availability: B1 contains faculty members with a share of low-SES candidates in their network that lies in the (0%, 25%] range, B2 those in (25%, 50%], B3 in (50%, 75%], and B4 in (75%, 100%].

²³ Notably, there is no evidence of homophily in selection for female candidates. Although female faculty members are more likely to make a nomination they are not more likely to nominate a woman. This may suggest that if homophily is not a main factor in faculty nominations, the lack of information on faculty SES should not affect the interpretation of the experimental findings.

Effects of the salience nudge on female nominations

Nominations vs. availability of female candidates, by treatment




Figure 3 effects of the salience nudge on female nominations.

Panel A compares the share of female students who receive a nomination with the share available in faculty networks. **Panel B** displays treatment differences in the likelihood of nominating a female candidate using faculty characteristics as covariates. **Panel C** plots the relationship between female availability and the probability of nominating a female student. Numbers next to the plotted lines represent the count of faculty members in each bin, by treatment.

pattern unchanged.²⁴ This is summarized in the result below:

Result 2 *Without the nudge, nominations underrepresent low-SES candidates. The salience nudge raises their share to match their availability in the nominators' networks, eliminating the class representation gap. Female nominations already match availability, and the nudge does not impose negative spillovers but leaves that pattern unchanged.*

²⁴ Because faculty could interact outside the study, some cross-treatment communication is possible. One way to check for contamination in this setting is to look at duplicate nominations, which should not occur if there is coordination among nominators (see e.g., Bolte et al. 2024). Of the 259 total nominations, 86 percent nominate unique students. The remaining 14 percent are duplicate nominations, and among those students 58 percent are nominated by faculty in both the CTRL and the SALIENCE groups rather than within the same group. This suggests that it is unlikely for spillovers to bias the main estimates.

3.3 Quality trade-offs: how good are those nominated?

After showing that the salience nudge widens the allocation of opportunities for low-SES candidates while leaving participation rates unchanged, I now turn to candidate quality. To evaluate potential quality trade-offs, I use two metrics: academic performance and program participation.

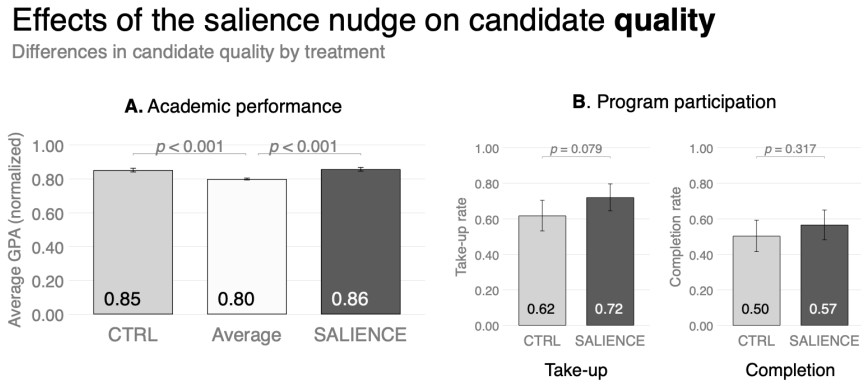


Figure 4 effects of the salience nudge on candidate quality.

Panel A compares the normalized GPA of students who receive a nomination with the average GPA of all students in faculty networks. **Panel B** displays treatment differences in take-up and completion rates of the training program.

Academic performance. Figure 4A shows that nominees outperform the student body in both treatments. Nominees in both CTRL and SALIENCE conditions hold an average normalized GPA that is about 6 p.p. above the student population ($p < 0.001$).

Program participation. Figure 4B reports how candidates engage with the training program after being nominated. In CTRL, 62% of nominees start the program and 50% finish it. In SALIENCE, those figures go up to 72% and 57%, respectively, but the differences are only marginally significant for take-up ($p = 0.079$) and not for completion ($p = 0.317$).²⁵

Attenuating the bias shifts attention toward overlooked candidates, yet the evidence does not support Hypotheses 3, as the salience nudge leaves candidate quality unaffected. The findings on candidate quality are summarized in the result below:

²⁵ The actual rate of take-up and completion after being nominated by a faculty member is high, as the average take-up when top performing students are invited directly is about 40% as reported in Munoz (2024). For a descriptive comparison see Appendix D.3.

Result 3 *The salience nudge raises representation for low-SES nominations without lowering candidate quality.*

3.4 Determinants of selection choices

The uniqueness of my dataset offers a rare opportunity to analyze gatekeepers' nomination choices through detailed network-level information. Specifically, the administrative data allow me to map every faculty-student interaction at a granular level, resulting in a comprehensive network with 47,993 distinct course-level ties.²⁶ This granular structure makes it possible to explicitly measure how network characteristics shape faculty nominations, to disentangle structural exposure from attention biases.

Table 2 Multinomial Mixed Logit estimates of nomination determinants

The table reports multinomial mixed logit regression results examining the determinants of faculty nomination choices. Columns I, IV, VII include the full network of nominators and candidates, while Columns II, V, VIII and III, VI, IX display results for the CTRL and SALIENCE groups, respectively. The models estimate the likelihood of a student being nominated based on their socio-demographic characteristics (gender, social class), academic performance (standardized GPA), and the strength of the faculty-student interactions. For the network interaction I use three specifications: tie strength (columns I-III), ties strength weighted by class size (columns IV-VI) and tie strength weighted by recency (columns VII-IX). Standard errors are shown in parentheses. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	Tie strength			Class size			Time decay		
	<i>Pooled</i> I	CTRL II	SALIENCE III	<i>Pooled</i> IV	CTRL V	SALIENCE VI	<i>Pooled</i> VII	CTRL VIII	SALIENCE IX
<i>Female</i>	-0.063 (0.143)	-0.022 (0.214)	-0.081 (0.212)	0.036 (0.222)	0.167 (0.255)	-0.138 (0.359)	0.114 (0.212)	0.191 (0.261)	0.024 (0.297)
<i>low-SES</i>	-0.977* (0.578)	-1.222* (0.725)	-0.271 (0.504)	-0.392 (0.385)	-0.598* (0.319)	-0.130 (0.344)	-0.357 (0.339)	-0.633* (0.332)	-0.259 (0.389)
<i>GPA (std)</i>	1.551** (0.122)	1.519** (0.161)	1.637** (0.226)	1.415** (0.140)	1.407** (0.173)	1.544** (0.270)	1.345** (0.150)	1.365** (0.168)	1.418** (0.260)
<i>Network interaction</i>	-0.272 (0.172)	0.074 (0.191)	-0.703* (0.340)	0.706** (0.060)	0.686** (0.045)	0.712** (0.049)	0.595** (0.030)	0.588** (0.040)	0.606** (0.041)
# Observations	47993	24530	23463	47993	24530	23463	47993	24530	23463
# Individuals	259	123	136	259	123	136	259	123	136
Chi-test	162.79	95.90	57.55	223.41	246.37	229.40	414.45	236.67	213.50

To effectively benefit from the richness of these dyadic data, I run a multinomial mixed logit model (McFadden and Train 2000) and report results in Table 2. This modeling approach is particularly suitable because it accounts for heterogeneity across nominators' choice sets, as each faculty member only selects among the students they have actually taught. In doing so, I incorporate three unique network measures to

²⁶ This is about 50% of the 99,272 ties in the complete network, as here I restrict the analysis to faculty who actually nominated someone for the training program.

identify the determinants of a faculty selection choices: (1) tie strength, capturing the frequency of past interactions between faculty and students, (2) class size, reflecting the visibility of candidates to nominators,²⁷ and (3) time decay, accounting for how recent interactions might weigh more heavily in the nominators' selection choice.²⁸ This integrated approach allows me to estimate the marginal effects of network structure on nomination decisions, separating the role of candidate visibility and exposure from potential behavioral biases. This is consistent with the narrative-driven heterogeneity documented by Andre et al. (2025), which shows that what people attend to depends on the information (e.g., narratives, networks) available to them. Columns I, IV, and VII pool the two treatments together; subsequent columns split the sample by treatment (CTRL and SALIENCE).

The analysis confirms that the history of interactions between nominators and candidates and the visibility this implies matter for selection. Smaller classes and more recent contact raise the chances of being nominated. Yet these structural factors do not explain the bias away. In fact, the impact of status is inline with the predictions. In the CTRL (see columns II, V and VIII) the low-SES coefficient is negative and significant, replicating the inbreeding bias already identified in Section 3.2. In SALIENCE (see columns III, VI, and IX) that coefficient shrinks and loses significance, consistent with the idea that the bias is attenuated by the nudge.²⁹ As expected, candidate quality is central in selection. Across all nine specifications the GPA coefficient is large, positive, and highly significant, confirming that candidate quality drives nomination choices and the nudge imposes no trade-offs. These observations are summarized in the following result:

Result 4 Academic performance drives nominations, network exposure amplifies visibility, and the salience nudge removes the residual low-SES bias without affecting quality.

²⁷ This tie measure is generated as $\frac{1}{n}$, where n is the number of students in a course. This gives more weight to connections with students in smaller classrooms.

²⁸ This tie measure is generated as 0.9^t , where $t=0$ if the course takes place in the current term, 1 if it took place the previous term, and so on. This gives more weight to connections with students in more recent courses.

²⁹ Female status never enters significantly in either condition, echoing the earlier evidence that there are no gender biases in this setting, and the nudge cannot lower what is already absent.

Discussion

The main results from this study suggest that the salience nudge addresses selection biases by redirecting gatekeepers' attention toward low-status candidates. However, the observed shift in nominations could be explained to some extent through alternative channels that cannot be ruled out conclusively with the design and data from the main field experiment.

One alternative interpretation is directly connected to the nature of the opportunity itself. The international training program required sustained engagement over multiple sessions, and was explicitly focused on enhancing self-efficacy; skills potentially perceived as particularly beneficial or fitting to high-SES students. Faculty might have considered low-SES students less likely to complete such a demanding opportunity. If this were the case, the nudge might not only have shifted attention but also provided nominators with additional information or implicit reassurance that the program was equally suitable and relevant for low-status students.

A second alternative interpretation relates to the psychological effect of the nudge itself, irrespective of the offered opportunity. Rather than simply heightening awareness of low-SES candidates, the nudge might have operated by signaling the presence of biases, leading to a reaction from faculty to correct or prove the bias wrong. It may also have done it by signaling institutional expectations. Faculty might have nominated more low-status students not because of heightened salience, but as compliance to what they interpreted as a directive from university administrators.

To disentangle these mechanisms and address the identified limitations, I designed and implemented a second study, which I report in the next section.

4 Mechanisms experiment

4.1 Setting and opportunity

To isolate the mechanisms underlying the observed effect of the salience nudge, I designed and conducted a follow-up experiment three years after the initial study. This second experiment took place at the same partner university, maintaining the organi-

zational context and nomination-based structure of the first study. However, this time the opportunity offered was a *Student Excellence Award*. Unlike the original training program this award required no sustained commitment. Instead, it provided immediate recognition and tangible benefits: selected students received a formal certificate as well as a cash prize, which carry significant professional and material value.

This simpler, low-commitment award allows me to rule out concerns specific to the nature of the training program, particularly faculty perceptions about the appropriateness or viability of low-status students' participation. Because the Excellence Award imposed no ongoing demands, nominators had little reason to question students' ability to engage fully, thus creating an ideal setting to clarify whether the nudge's effectiveness arose purely from shifting attention or from the provision of information about the opportunity itself.

I am also able to test directly the alternative mechanisms that might explain the effectiveness of the salience nudge observed in the original experiment. To disentangle whether the observed equity shift was driven by attention shifts, faculty debiasing, or institutional expectations, I randomly varied the messages embedded within the invitation to nominate. I designed a multi-armed experiment using three distinct nudges as well as a control. One is linked to the original SALIENCE nudge, one explicitly encourages faculty to reflect on unconscious biases (DEBIASING), and one emphasizes the university's commitment to diversity (INSTITUTIONAL). By comparing nomination patterns of these treatments to a CTRL, the new design allows me to better identify which channel is most influential on gatekeepers' nomination choices. Finally, I designed the nudges to exclusively target socio-economic background (i.e., gender was not made salient), to focus on the relevant status category for this setting.³⁰

4.2 Treatments and sample

At the moment of conducting the second field experiment (2025), there was an increase in the number of employed faculty members and the share of enrolled low-SES students compared to the first study. There were 664 faculty members working at the local

³⁰ After seeing that there were no spillovers from the nudge on female students in the first study, administrators from the local university accepted to focus exclusively on socioeconomic background in the follow-up study.

university. About 40% of the students were low-SES, while the share of female students remained at about 55%.

All faculty members received an invitation email from the Office of International Relations asking them to nominate a student for an *Excellence Award*, based on academic performance (GPA).³¹ To test the differential impact of the nudges, I randomly assigned nominators to one of four conditions (for details see Table B-8 in Appendix B). The following statements display the differences in the messages embedded within the invitation to nominate (see Appendix A for the full invitation text):

CTRL condition:

Select a student based on their academic performance (GPA).

SALIENCE condition:

Select a student based on their academic performance (GPA). Remember to consider all the students you know, including those who come from disadvantaged socioeconomic backgrounds, so as not to overlook anyone who deserves this recognition.

DEBIASING condition:

Select a student based on their academic performance (GPA). Note that, according to scientific research, sometimes we may overlook students from disadvantaged socioeconomic backgrounds due to unconscious biases. We invite you to reflect on this when making your choice, so as not to overlook anyone who deserves this recognition.

INSTITUTIONAL condition:

Select a student based on their academic performance (GPA). The university has a commitment to equity and inclusion of students from disadvantaged socioeconomic backgrounds, ensuring that everyone has the opportunity to be recognized. We invite you to strengthen this commitment with your nomination, so as not to overlook anyone who deserves this recognition.

³¹ To encourage participation of nominators, this time the university offered twenty \$100 bonuses, which would be raffled among nominators. Also, the nomination process involved two steps as reported in Appendix A.2. In step 1 all faculty received an identical invitation email to fill-out a short online survey and nominate a student. Once entered the survey, they were presented with the instructions depending on their treatment, which allows me to properly measure the impact of the nudge as those entering the survey are already showing interest in the nomination process (see e.g., Leibbrandt and List 2018).

4.3 Results

Here, I summarize the main findings of the second field experiment, looking at the effectivity of each message relative to the CTRL group.³² I focus on the same outcome measures as before: backlash, equity shift and quality trade-offs. For simplicity, I present results pooling the three treatments into a single NUDGE category when there are no differences between conditions and the CTRL. When there are, I report the disaggregated outcomes.³³

First, I find no evidence of backlash. Overall participation was similar to the first experiment: 302 of the 664 (46%) invited nominators showed interest in the excellence award process. Out of these, 252 made a nomination (83%). Figure 5A illustrates that nomination rates are statistically indistinguishable across conditions ($p = 0.488$), reinforcing the evidence that low-touch nudges do not trigger backlash.

Second, there is an equity shift. In the CTRL group, low-SES students account for 26% of nominees, resembling the nomination rate observed in the first experiment. Figure 5B reports separately the effect of each treatment, and shows that the SALIENCE nudge nearly doubles the share of low-SES nominees to 47% ($p = 0.018$). The DEBIASING nudge yields a more modest increase to 42% ($p = 0.079$), while the INSTITUTIONAL nudge raises low-SES nominations only to 33% ($p = 0.366$).³⁴

Third, there are no quality trade-offs. Academic performance among nominees was high and statistically indistinguishable between treatments. This is illustrated in Figure 5C. Thus, the equity gains generated by the nudge came at no observable cost to candidate quality. Also, consistent with the previous findings (even though this time they were not mentioned explicitly in any of the nudges), female candidates were not affected by the prompts. This mirrors the gender-neutral pattern observed in the first

³² For this study, the data-sharing agreement with the local partner did not include course enrollment registries (unlike study 1). For the analysis, I focus on *between-treatment* comparisons. I report regression tables in Appendix B.

³³ I pre-registered for this study, that I would analyze the pooled effect of using a message as well as the separate evaluation of each nudge against the control condition (see a discussion in Appendix E). The interested reader can find all disaggregated results in Appendix D.4.

³⁴ The salience nudge used in the mechanisms experiment singled out *only* low-SES status, while that of the first study also mentioned gender. It is possible that focusing on socioeconomic background alone might help explain the larger shift in low-SES nominations in Study 2. However, this is only tentative and not fully conclusive as the two experiments also differed in the nature of the opportunity and in the baseline share of low-SES students: about 40% in 2025 vs. 31% in 2022.

Effects of the nudges on nominations and quality

Selection either polling all nudges or by type of nudge

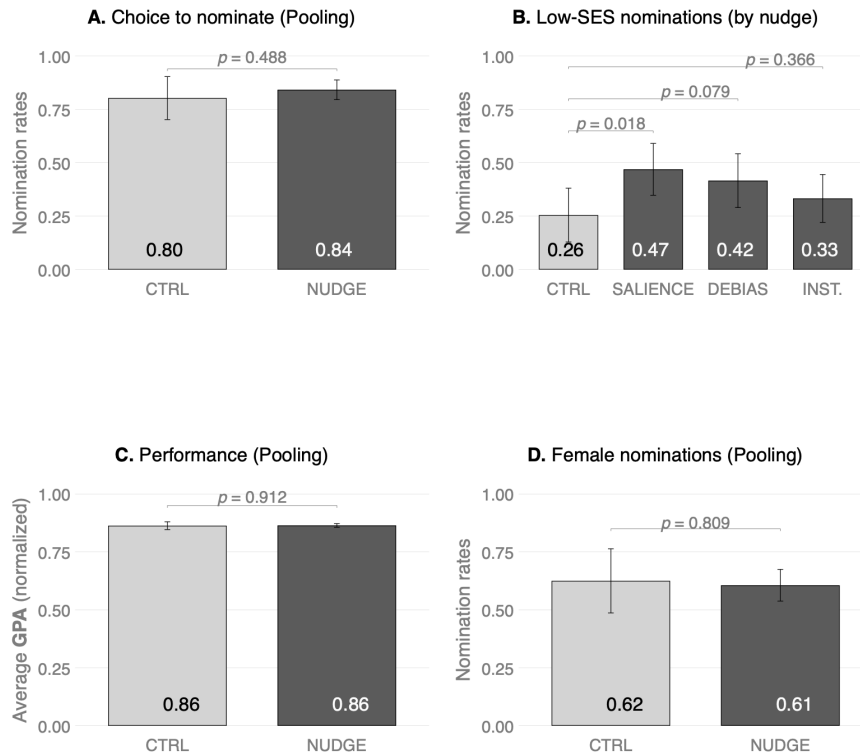


Figure 5 Mechanism experiment: Effects of the nudges on nominations and quality.

Panel A compares nomination rates between CTRL and the pooled NUDGE condition. **Panel B** compares the share of low-SES students who receive a nomination focusing on differences between each nudge and the CTRL. **Panel C** compares the normalized GPA of students who receive a nomination between CTRL and the pooled NUDGE condition. **Panel D** compares the share of female students who receive a nomination between CTRL and the pooled NUDGE condition.

experiment and confirms that nudging attention toward a disadvantaged group does not impose spillovers on others (see Figure 5D).

This second study, using a new opportunity, confirms that a salience nudge delivers the largest and most cost-effective equity shift. Encouraging self-reflection on bias helps, but to a lesser extent, perhaps because it demands greater cognitive effort or induces defensiveness (as it explicitly suggests to nominators that they may be biased). Invoking institutional values results in a modest shift. Combined, the two experiments demonstrate that a simple SALIENCE nudge that brings overlooked candidates to mind delivers large equity gains at no cost to quality, whereas more forceful moral or insti-

tutional messages appear to be less effective.³⁵ This is summarized in the following result:

Result 5 *The mechanisms experiment further shows that the salience nudge significantly increases low-SES nominations without backlash or quality trade-offs. Alternative nudges designed to trigger debiasing or institutional compliance have smaller or negligible effects, underscoring attention as the most effective mechanism.*

5 Conclusions

This paper asks whether a salience nudge, delivered as a reminder at the moment of choice, can steer gatekeepers toward fairer nominations for a valuable opportunity, without lowering standards or triggering resistance. The evidence of my study supports this conjecture. In a field experiment at a Colombian university, a salience nudge that drew attention to low-status candidates closed the class nomination gap for low-SES individuals, left overall nomination rates intact, and preserved both the academic performance and program engagement of those selected.

A follow-up experiment strengthened the case for the attentional explanation. When the nudge made low-SES students more salient, the equity gain was large. When the message instead urged faculty to reflect on implicit biases, the gain was not as big. When it cited the university's diversity policy, the gain was much smaller. These differences suggest that shifting attention to overlooked candidates moves behavior most, while appeals to moral reflection or institutional duty may not be as effective.

For policy, the implication is direct. In organizational settings where access to opportunities depends on informal nominations, salience nudges can be embedded into their calls to prompt a fairer allocation of opportunities without resistance. Implementation is cheap, scalability is high, and the absence of backlash makes the approach politically attractive.

Two limitations qualify the findings in this paper, which can be promising avenues for future research. First, the experiment took place in a setting where nominators (faculty)

³⁵ I also test and find that spillovers from the 2022 study do not explain the findings in this experiment. In Appendix D.5, I report regression analysis of the impact of being treated with the salience nudge in Study 1 on selection in Study 2 and find no evidence for this.

had little reason to expect future favors from the students they nominated, so reciprocal motives were minimal. In some organizational settings, a nomination may carry implicit expectations of reciprocal benefits (see e.g., [Caria et al. 2023](#)). The salience nudge might interact differently with anticipated reciprocity, although it is unclear if it would amplify or reduce its effectiveness. Second, despite faculty nominating high-quality candidates without backlash, only about half of those invited participated in the nomination process in both experiments. Thus, increasing nominator engagement remains an open challenge. A follow-up study could allow nominators to submit multiple candidates, as providing multiple-choice opportunities has been shown to enhance diversity (see e.g., [Chang et al. 2020](#)). However, this approach may introduce quality trade-offs by incentivizing the inclusion of less-qualified candidates to fill additional slots (see e.g., [Ahern and Dittmar 2012](#)). Addressing these limitations opens valuable avenues for future research on optimizing the effectiveness of salience interventions in selection processes.

Even so, the core take-away is clear. A light-touch nudge that simply widens the spotlight of attention towards low-status groups can reduce a persistent form of inequality, align opportunity with performance, and do so at virtually no cost. Salience nudges deserve a prominent place in the repertoire of managers and policymakers committed to fair and inclusive selection practices.

References

- Agan, A. and Starr, S. (2018). Ban the box, criminal records, and racial discrimination: A field experiment. The Quarterly Journal of Economics, 133:191–235.
- Ahern, K. R. and Dittmar, A. K. (2012). The changing of the boards: The impact on firm valuation of mandated female board representation. The Quarterly Journal of Economics, 127:137–197.
- Alesina, A., Carlana, M., La Ferrara, E., and Pinotti, P. (2024). Revealing stereotypes: Evidence from immigrants in schools. American Economic Review, 114:1916–1948.
- Allcott, H. and Taubinsky, D. (2015). Evaluating behaviorally motivated policy: Experi-

- mental evidence from the lightbulb market. American Economic Review, 105:2501–2538.
- Andre, P., Haaland, I., Roth, C., Wiederholt, M., and Wohlfart, J. (2025). Narratives about the macroeconomy. Review of Economic Studies.
- Arslan, C., Chang, E. H., Chilazi, S., Bohnet, I., and Hauser, O. P. (2025). Behaviorally designed training leads to more diverse hiring. Science, 387(6732):364–366.
- Athey, S. and Palikot, E. (2024). The value of non-traditional credentials in the labor market. Working paper, arXiv:2405.00247:1–n.
- Bandura, A. (1978). Self-efficacy: Toward a unifying theory of behavioral change. Advances in Behaviour Research and Therapy, 1:139–161.
- Beaman, L., Keleher, N., and Magruder, J. (2018). Do job networks disadvantage women? evidence from a recruitment experiment in malawi. Journal of Labor Economics, 36:121–157.
- Bertrand, M. and Mullainathan, S. (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. American Economic Review, 94:991–1013.
- Bisin, A. and Verdier, T. (2011). The economics of cultural transmission and socialization. In Benhabib, J., Bisin, A., and Jackson, M. O., editors, Handbook of Social Economics, volume 1, pages 339–416. North-Holland.
- Bogliacino, F., Jiménez Lozano, L., and Reyes, D. (2018). Socioeconomic stratification and stereotyping: lab-in-the-field evidence from colombia. International Review of Economics, 65:77–118.
- Bohnet, I. (2016). What works. Gender equality by design. Harvard University Press.
- Bollinger, B., Leslie, P., and Sorensen, A. T. (2011). Calorie posting in chain restaurants. American Economic Journal: Economic Policy, 3:91–128.
- Bolte, L., Immorlica, N., and Jackson, M. O. (2024). The role of referrals in immobility, inequality, and inefficiency in labor markets. Journal of Labor Economics, X:x–x.

- Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2016). Stereotypes. The Quarterly Journal of Economics, 131:1753–1794.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience theory of choice under risk. The Quarterly Journal of Economics, 127:1243–1285.
- Brashears, M. E. and Quintane, E. (2015). The microstructures of network recall: How social networks are encoded and represented in human memory. Social Networks, 41:113–126.
- Cardenas, J. C., Fergusson, L., and Garcia Villegas, M. (2021). La quinta puerta: De cómo la educación en Colombia agudiza las desigualdades en lugar de remediarlas. Ariel Colombia.
- Caria, S., Franklin, S., and Witte, M. (2023). Searching with friends. Journal of Labor Economics, 41:887–922.
- Carlana, M. and Fort, M. (2022). Hacking gender stereotypes: Girls’ participation in coding clubs. AEA Papers and Proceedings, 112:583–587.
- Castilla, E. J. (2008). Gender, race, and meritocracy in organizational careers. American Journal of Sociology, 113:1479–1526.
- Castilla, E. J. (2015). Accounting for the gap: A firm study manipulating organizational accountability and transparency in pay decisions. Organization Science, 26:311–333.
- Castilla, E. J. and Ranganathan, A. (2020). The production of merit: How managers understand and apply merit in the workplace. Organization Science, 31:909–935.
- Chang, E. H., Kirgios, E. L., Rai, A., and Milkman, K. L. (2020). The isolated choice effect and its implications for gender diversity in organizations. Management Science, 66:2752–2761.
- Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. American Economic Review, 99:1145–1177.
- Currarini, S., Jackson, M. O., and Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. Econometrica, 77:1003–1045.

- Currarini, S., Jackson, M. O., and Pin, P. (2010). Identifying the roles of race-based choice and chance in high school friendship network formation. Proceedings of the National Academy of Sciences, 107:4857–4861.
- Diaz, B. S., Neyra Nazarrett, A., Ramirez, J., Sadun, R., and Tamayo, J. A. (2025). Training within firms. NBER Working Paper, 33670:1–44.
- Dustmann, C., Glitz, A., Schönberg, U., and Brucker, H. (2016). Referral-based job search networks. Review of Economic Studies, 83:514–546.
- Fernandez, R. M., Castilla, E. J., and Moore, P. (2000). Social capital at work: Networks and employment at a phone center. American Journal of Sociology, 105:1288–1356.
- Friedman, S. and Laurison, D. (2020). The Class Ceiling: Why It Pays to Be Privileged. Policy Press, Bristol, UK.
- Frydman, C. and Wang, B. (2020). The impact of salience on investor behavior: Evidence from a natural experiment. Journal of Finance, 75:229–276.
- García-Sánchez, E., Willis, G. B., Rodríguez-Bailón, R., García-Castro, J. D., Palacio-Sañudo, J., Polo, J., and Rentería-Pérez, E. (2018). Perceptions of economic inequality in colombian daily life: More than unequal distribution of economic resources. Frontiers in Psychology, 9:1660.
- Goldin, C. and Rouse, C. (2000). Orchestrating impartiality: The impact of ‘blind’ auditions on female musicians. American Economic Review, 90:715–741.
- Granovetter, M. (1974). Getting a Job: A Study of Contacts and Careers. Harvard University Press.
- Harrison, G. W. and List, J. A. (2004). Field experiments. Journal of Economic Literature, 42:1009–1055.
- Leibbrandt, A. and List, J. (2018). Do equal employment opportunity statements backfire? evidence from a natural field experiment on job-entry decisions. NBER Working Paper, 25035:1–40.
- Link, S., Peichl, A., Pfäuti, O., Roth, C., and Wohlfart, J. (2025). Attention to the macroeconomy. CEBI Working Paper Series.

- Londono-Velez, J. (2022). The impact of diversity on perceptions of income distribution and preferences for redistribution. Journal of Public Economics, 214:1–29.
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, Frontiers in Econometrics, pages 105–142. Academic Press, New York.
- McFadden, D. and Train, K. (2000). Mixed mnl models for discrete response. Journal of Applied Econometrics, 15:447–470.
- Menon, T., Shea, C. T., and Smith, E. B. (2024). Mobilization capacity: Tracing the path from having networks to capturing resources. Research in Organizational Behavior, 44:100210.
- Michelman, V., Price, J., and Zimmerman, S. D. (2022). Old boys’ clubs and upward mobility among the educational elite. The Quarterly Journal of Economics, 137:845–909.
- Milkman, K. (2021). How to change: The Science of Getting from Where You Are to Where You Want to Be. Penguin Random House LLC, New York.
- Munoz, M. (2024). How publicly targeting by group identities impacts take-up of educational opportunities. Working Paper.
- Niederle, M., Segal, C., and Vesterlund, L. (2013). How costly is diversity? affirmative action in light of gender differences in competitiveness. Management Science, 59:1–16.
- Oh, S. (2023). Does identity affect labor supply. American Economic Review, 113:2055–2083.
- Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. American Sociological Review, 77:999–1022.
- Shukla, S. (2022). Making the elite: Top jobs, disparities, and solutions. International Finance Discussion Paper, 1331r1.
- Simpson, B., Markovsky, B., and Steketee, M. (2011). Power and the perception of social networks. Social Networks, 33:166–171.

- Small, M. L. (2013). Weak ties and the core discussion network: Why people regularly discuss important matters with unimportant alters. Social Networks, 35:470–483.
- Small, M. L. and Sukhu, C. (2016). Because they were there: Access, deliberation, and the mobilization of networks for support. Social Networks, 47:73–84.
- Stansbury, A. and Rodriguez, K. (2025). The class gap in career progression: Evidence from us academia. Working Paper.
- Thaler, R. H. and Benartzi, S. (2004). Save more tomorrow™: Using behavioral economics to increase employee saving. Journal of Political Economy, 112:164–187.
- Thaler, R. H. and Sunstein, C. R. (2009). Nudge: Improving Decisions About Health, Wealth, and Happiness. Penguin, New York.
- Topa, G. (2001). Social interactions, local spillovers and unemployment. Review of Economic Studies, 68:261–295.
- Turco, C. J. (2010). Cultural foundations of tokenism: Evidence from the leveraged buyout industry. American Sociological Review, 75:894–913.
- White-Lewis, D. K. (2020). The facade of fit in faculty search processes. The Journal of Higher Education, 91:833–857.
- Zimmerman, S. D. (2019). Elite colleges and upward mobility to top jobs and top incomes. American Economic Review, 109:1–47.

Online Appendix:

Closing class gaps with simple nudges: Experimental evidence on opportunity allocation

Manuel Munoz

A Invitations to nominate

Appendix A reproduces every message that faculty received in each of the two field experiments. Section A.1 shows the email invitation used in Study 1 (the training-program experiment). Section A.2 shows the two-step sequence used in Study 2 (the excellence-award experiment).

A.1 Nominations for the training program (Study 1)

The invitation message below was sent to all faculty members via institutional email. The original email was sent in Spanish by the Office of International Relations of the local university. I include below the English version (translated by the author).

Dear Professor [Name],

The [Local University] has a collaboration agreement with the [International University], a global university of re-known quality. As part of the agreement, professors from the [International University] will teach a training program to help students at [Local University] acquire and further develop the necessary tools to achieve their goals and increase their chances of personal and professional success (you can see details of the program at the end of this message).

This great opportunity provides multiple benefits. First, being able to learn from ex-

cellent professors. Second, by completing the program participants will receive an attendance certificate from the [International University]. These types of credentials can have a very large impact in a student's CV and open doors for future jobs or scholarships. In addition, at the end of the program there will be a lottery of various last-generation iPads among those who complete the program, with the aim of giving students a tool that may help them in their academic endeavors.

As part of the agreement, the [International University] is inviting each faculty member at [Local University], including you, to recommend a student, who will automatically have a spot in the training program. You can recommend students from any program, but please only recommend students that you have taught at least once.

We ask you to respond to this message with the full name and student code of the person you want to recommend, considering that this person has the academic performance to benefit from this great opportunity.

SALIENCE condition: For your recommendation, please focus on the academic performance of the students and do not exclude anyone because of their socioeconomic background, gender or any other demographic characteristic.

Additional information about the program:

- Name: "How to Change: Scientific Tools to Achieve Your Personal and Professional Goals."
- Instructor(s): The program will be taught by professors from high international standing from the [International University].
- Language: Spanish.
- Duration: 9 online sessions, half an hour each. All sessions are independent and students will be able to complete them at their own pace. So, they will not have any conflicts of scheduling with other academic activities.
- Start: Recommended students will be contacted at the beginning of the second semester of [Year of intervention]. The program will start on [Start date].
- Costs: Free course for recommended students.

- Benefits: An international certificate of attendance. Also, students will participate in the lottery of various iPads.
- Deadline: Please make your recommendation before [Deadline date].

We appreciate your collaboration in making a recommendation, so your recommended student can benefit from the opportunities in this program.

Sincerely,

[Signature]

Head / Office of International Relations

A.2 Nominations for the student excellence award (Study 2)

Section A.2 explains the two steps involved in Study 2. First, every faculty member received the same email inviting them to nominate a student for the *Excellence award*. The message did not include any treatment information. Only after respondents entered the survey and progressed to the final screen, immediately before submitting their nomination, they were displayed one of four nudges. This 2-step process ensured that (a) the decision to open the survey would not be influenced by treatment wording, and (b) any treatment effects are attributable solely to the single sentence displayed at the point of choice, not to differential recruitment or demand effects introduced earlier in the process.

Step 1: Invitation email

Dear Professor [Name],

We are writing to invite you to take part in the 2025 [Internal name] Survey.

This short survey takes about **7 minutes** to complete and is part of a university-wide initiative to help identify candidates for the **2025 EXCELLENCE AWARD**, a prize the [local University] is launching this year in alliance with [international University].

Every faculty member has the opportunity **to nominate a student** by completing the survey. As a token of our appreciation for your participation, we will raffle **10 vouchers worth 300,000 pesos** each.

To participate and nominate a student for the award, please click the link below:

Link: [link]

Sincerely,

[Signature]

Head / Office of International Relations

Step 2: Nomination survey

You have been invited to take part in an initiative that gives you the opportunity to **nominate a student for the 2025 EXCELLENCE AWARD**.

Each faculty member can nominate one student. The student you nominate may be informed that you have recommended them for this award.

The entire survey takes **approximately 7 minutes** to complete. As a token of appreciation for your participation, **10 vouchers of 300,000 pesos** will be raffled.

On the following screens we will ask you some questions about your role as an instructor and, at the end, you will be able to nominate a student for the award.

Before continuing, please read the consent form carefully so that you understand how we handle your data and responses.

[Page break]

The survey first included some demographic and attitudinal questions, for which I do not have data. At the end, it presented the nomination item, which I report below:

Please write the full name and major/program of the student you wish to nominate.

CTRL condition:

Select a student based on their academic performance (GPA).

SALIENCE condition:

Select a student based on their academic performance (GPA). Remember to consider all the students you know, including those who come from disadvantaged socioeconomic backgrounds, so as not to overlook anyone who deserves this recognition.

DEBIASING condition:

Select a student based on their academic performance (GPA). Note that, according to scientific research, sometimes we may overlook students from disadvantaged socioeconomic backgrounds due to unconscious biases. We invite you to reflect on this when making your choice, so as not to overlook anyone who deserves this recognition.

INSTITUTIONAL condition:

Select a student based on their academic performance (GPA). The university has a commitment to equity and inclusion of students from disadvantaged socioeconomic backgrounds, ensuring that everyone has the opportunity to be recognized. We invite you to strengthen this commitment with your nomination, so as not to overlook anyone who deserves this recognition.

B Regression Tables

This appendix provides detailed regression results supporting the analysis reported in the main text.

Table B-1 GPA and Students' Expectations by SES

This table presents regression results for GPA outcomes and students' expectations across different specifications. Columns I–III examine GPA differences by social class using a dataset of 6,098 students, where 36% are low class. Columns IV–VI use a dataset of 1,924 survey respondents, where 26% are low class. Column IV reports students' expected time to find their first job after college, Column V their expected salary for the first job, and Column VI their beliefs about the number of peers from their program who found a job through social connections. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV	V	VI
Low-SES	0.011*** (0.002)	0.009*** (0.003)	0.009*** (0.003)	0.019** (0.010)	-0.045*** (0.010)	-0.019* (0.011)
Program		0.000*** (0.000)	0.000*** (0.000)	-0.000 (0.000)	-0.000 (0.000)	0.001** (0.000)
Semester			-0.003*** (0.001)	0.000 (0.002)	-0.017*** (0.002)	-0.007*** (0.002)
Constant	0.794*** (0.001)	0.786*** (0.003)	0.805*** (0.004)	0.468*** (0.012)	0.441*** (0.015)	0.244*** (0.015)
# Observations	6098	2923	2923	1921	1921	1921
R ²	0.005	0.010	0.024	0.003	0.052	0.011

Table B-2 Balance Randomization into Treatments (Study 1)

This table presents regression results testing baseline balance on five demographic and career characteristics. Column I reports the difference in the proportion of female faculty between the treatment and control groups; Column II the difference in the proportion of married faculty; Column III the difference in the share holding a graduate degree; Column IV the difference in the share employed full-time (i.e., tenure); and Column V the difference in average years the faculty member has been at the university. In every specification the coefficient on the variable SALIENCE measures the mean difference between treated and control faculty. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV	V
SALIENCE	0.053 (0.043)	-0.008 (0.044)	-0.035 (0.030)	-0.048 (0.043)	-0.995 (0.658)
Constant	0.399*** (0.031)	0.512*** (0.031)	0.876*** (0.021)	0.585*** (0.031)	9.302*** (0.485)
Observations	528	528	528	528	528
R ²	0.003	0.000	0.003	0.002	0.004

Table B-3 Balance Post-Randomization - Network measures

This table presents regression results for five baseline network composition and performance measures. These are post-randomization measures as they were not used to stratify treatment assignment. Column I examines the number of connections in the network (i.e., degree); Column II the average frequency of contact, counting the number of courses taught to a given connection; Column III the share of female student in the faculty member's network; Column IV the share of low-SES students; and Column V the average GPA of students. In each column, the coefficient on *SALIENCE* captures the mean difference between treated and control faculty. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV	V
<i>SALIENCE</i>	-22.588 (13.427)	-0.099* (0.050)	-0.011 (0.015)	-0.005 (0.014)	0.002 (0.003)
Constant	199.566*** (10.267)	1.610*** (0.039)	0.560*** (0.011)	0.320*** (0.011)	0.793*** (0.002)
Observations	528	528	528	528	528
R^2	0.005	0.008	0.001	0.000	0.001

Table B-4 Determinants of Nominating Low-SES or Female Students

This table reports linear-probability regressions of the likelihood that a faculty members nominate a low-SES (Columns I-II) or female (Columns III-IV) candidate. Columns I and III restrict the sample to the *CTRL* condition, while Columns II and IV restrict the sample to the *SALIENCE* condition. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV
<i>Socio-demographics</i>				
Faculty is female	-0.018 (0.068)	0.011 (0.058)	0.135 (0.090)	-0.078 (0.081)
Faculty is married	0.119* (0.058)	-0.010 (0.061)	0.076 (0.082)	-0.130 (0.078)
Permanent contract	0.012 (0.060)	-0.165* (0.069)	0.022 (0.092)	-0.089 (0.089)
Graduate degree	0.111* (0.056)	0.046 (0.097)	-0.076 (0.300)	-0.150 (0.132)
<i>Networks</i>				
Average number of ties	0.003 (0.043)	0.119 (0.064)	0.086 (0.074)	-0.127* (0.063)
Share of female ties	0.258 (0.156)	0.149 (0.196)	1.315*** (0.197)	1.270*** (0.156)
Share of low-SES ties	0.486 (0.252)	0.796*** (0.210)	0.202 (0.206)	0.249 (0.217)
Constant	-0.366* (0.159)	-0.265 (0.142)	-0.492 (0.350)	0.250 (0.210)
# Observations	123	136	123	136
R^2	0.117	0.211	0.244	0.254

Table B-5 Differences in GPA among Nominated Students (Study 1)

This table presents regression results for GPA differences among nominated students across five specifications. The independent variable, *SALIENCE* nudge, refers to the faculty's treatment status (*CTRL* is the omitted category). Column I pools all nominated students, Column II focuses on low-SES students, Column III on high-SES students, Column IV on female students, and Column V on male students. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV	V
<i>SALIENCE</i>	0.005 (0.008)	0.008 (0.013)	0.003 (0.010)	0.006 (0.009)	0.004 (0.014)
Constant	0.852*** (0.006)	0.857*** (0.009)	0.850*** (0.007)	0.858*** (0.007)	0.842*** (0.009)
# Observations	259	70	189	159	100
R ²	0.002	0.005	0.000	0.002	0.001

Table B-6 Differences in Program Take-up among Nominated Students

This table presents regression results for the probability of completing the first session of the training program (take-up) across five specifications. Column I pools all nominated students, Column II focuses on low-SES students, Column III on high-SES students, Column IV on female students, and Column V on male students. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV	V
<i>SALIENCE</i>	0.103 (0.059)	0.202 (0.112)	0.060 (0.070)	0.115 (0.074)	0.084 (0.097)
Constant	0.618*** (0.044)	0.607*** (0.094)	0.621*** (0.050)	0.632*** (0.056)	0.596*** (0.072)
# Observations	259	70	189	159	100
R ²	0.012	0.050	0.004	0.016	0.008

Table B-7 Differences in Program Completion among Nominated Students

This table presents regression results for the likelihood of completing all nine sessions of the training program (completion) across five specifications. Column I pools all referred students, Column II focuses on low-SES students, Column III on high-SES students, Column IV on female students, and Column V on male students. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV	V
<i>SALIENCE</i>	0.062 (0.062)	0.119 (0.122)	0.037 (0.073)	0.040 (0.079)	0.098 (0.101)
Constant	0.504*** (0.045)	0.500*** (0.096)	0.505*** (0.052)	0.526*** (0.058)	0.468*** (0.074)
# Observations	259	70	189	159	100
R ²	0.004	0.014	0.001	0.002	0.010

Table B-8 Baseline Balance Across Treatment Arms (Study 2)

This table reports regressions testing whether random assignment in Study 2 is balanced on five pre-treatment faculty characteristics. Column I shows differences in the share of female faculty; Column II differences in the share who are married; Column III differences in the share holding a permanent (tenure-track) contract; Column IV differences in the probability the faculty member received the SALIENCE nudge in Study 1; and Column V differences in the probability the faculty member made a nomination in Study 1. In every specification, the coefficients on SALIENCE, DEBIAS, and INSTITUTIONAL represent mean differences relative to the CTRL condition. The bottom panel gives p -values from F -tests for equality of coefficients across pairs of treatment. Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	I	II	III	IV	V
SALIENCE	-0.010 (0.055)	0.073 (0.048)	0.004 (0.054)	-0.004 (0.089)	-0.003 (0.091)
DEBIAS	-0.006 (0.054)	0.014 (0.046)	-0.001 (0.054)	0.021 (0.088)	-0.006 (0.091)
INSTITUTIONAL	-0.001 (0.055)	0.083 (0.048)	0.000 (0.054)	0.045 (0.087)	0.004 (0.091)
Constant	0.420*** (0.039)	0.222*** (0.033)	0.599*** (0.039)	0.914*** (0.063)	0.864*** (0.065)
<i>Inter-treatment comparisons (p-value)</i>					
SALIENCE vs. DEBIAS	0.933	0.227	0.929	0.771	0.969
SALIENCE vs. INSTITUTIONAL	0.861	0.840	0.947	0.576	0.940
DEBIAS vs. INSTITUTIONAL	0.927	0.157	0.983	0.786	0.909
# Observations	664	664	664	664	664
R^2	0.000	0.007	0.000	0.001	0.000

Table B-9 Determinants of making a nomination

This table presents linear-probability regressions for two outcomes. Columns I–II model the probability that a faculty member *participates*, i.e. fills out the nomination survey. Columns III–IV model the probability that a faculty member actually *submits a nomination*, conditional on participating in the survey. Columns I and III include only treatment indicators (SALIENCE, DEBIAS, INSTITUTIONAL, where the omitted category is the CTRL condition). Columns II and IV add faculty covariates: gender, tenure status, prior-study treatment (Old SALIENCE), and the *helper* indicator (if the faculty nominated a student in Study 1). The bottom panel reports p -values from pair-wise F -tests that compare treatment coefficients within each specification (Salience vs. Debias, Salience vs. Institutional, Debias vs. Institutional). Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV
SALIENCE	0.105 (0.055)	0.105* (0.052)	0.047 (0.065)	0.055 (0.063)
DEBIAS	0.091 (0.054)	0.093 (0.052)	0.007 (0.068)	0.004 (0.066)
INSTITUTIONAL	0.114* (0.054)	0.116* (0.052)	0.063 (0.064)	0.074 (0.063)
Female		0.044 (0.037)		−0.073 (0.044)
Tenure		0.258*** (0.039)		0.161** (0.056)
Old SALIENCE		−0.050 (0.036)		−0.019 (0.039)
Helper		0.142*** (0.034)		0.047 (0.043)
Constant	0.377*** (0.038)	0.127* (0.051)	0.803*** (0.051)	0.685*** (0.081)
<i>Inter-treatment comparisons (p-value)</i>				
SALIENCE vs. DEBIAS	0.792	0.823	0.506	0.377
SALIENCE vs. INSTITUTIONAL	0.869	0.832	0.774	0.734
DEBIAS vs. INSTITUTIONAL	0.667	0.659	0.341	0.236
# Observations	664	664	302	302
R^2	0.008	0.106	0.005	0.058

Table B-10 Determinants of Nominating a Low-SES Student

The dependent variable equals 1 if the candidate nominated by a faculty member is from a low-SES background. Columns I–II collapse the three nudges into a single NUDGE indicator (Column I without, Column II with faculty covariates). Columns III–IV include separate indicators for SALIENCE, DEBIAS, and INSTITUTIONAL; the omitted category is CTRL. Robust standard errors are in parentheses. Columns II and IV add faculty covariates: gender, tenure status, prior-study treatment (Old SALIENCE), and the *helper* indicator (if the faculty nominated a student in Study 1). The bottom panel reports p -values from pair-wise F -tests comparing treatment coefficients (only defined for the disaggregated specifications). Statistical significance: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

	I	II	III	IV
NUDGE (Pooling)	0.150* (0.073)	0.141 (0.073)		
SALIENCE			0.214* (0.089)	0.206* (0.090)
DEBIAS			0.161 (0.091)	0.157 (0.091)
INSTITUTIONAL			0.078 (0.086)	0.065 (0.086)
Female		0.069 (0.063)		0.080 (0.063)
Tenure		−0.057 (0.077)		−0.062 (0.077)
Old SALIENCE		0.019 (0.058)		0.015 (0.058)
Helper		−0.064 (0.061)		−0.055 (0.062)
Constant	0.255*** (0.064)	0.323** (0.107)	0.255*** (0.064)	0.317** (0.107)
<i>Inter-treatment comparisons (p-value; Cols III–IV)</i>				
SALIENCE vs. DEBIAS	—	—	0.553	0.591
SALIENCE vs. INSTITUTIONAL	—	—	0.107	0.100
DEBIAS vs. INSTITUTIONAL	—	—	0.334	0.282
# Observations	242	242	242	242
R^2	0.015	0.028	0.026	0.040

Table B-11 Determinants of Nominating a Female Student

The dependent variable equals 1 if the candidate nominated by a faculty member is female. Columns I–II pool the three nudges into a single NUDGE indicator (Column I without, Column II with covariates). Columns III–IV include separate indicators for SALIENCE, DEBIAS, and INSTITUTIONAL; CTRL is the omitted category. Robust standard errors are in parentheses. Columns II and IV add faculty covariates: gender, tenure status, prior-study treatment (Old SALIENCE), and the *helper* indicator (if the faculty nominated a student in Study 1). The bottom panel shows p -values from pair-wise F -tests comparing treatment coefficients (defined for Columns III–IV only). Statistical significance: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

	I	II	III	IV
NUDGE (Pooling)	−0.019 (0.078)	−0.028 (0.078)		
SALIENCE			−0.064 (0.094)	−0.073 (0.092)
DEBIAS			0.004 (0.094)	−0.005 (0.095)
INSTITUTIONAL			0.004 (0.091)	−0.005 (0.092)
Female		−0.088 (0.064)		−0.090 (0.064)
Tenure		−0.120 (0.073)		−0.120 (0.073)
Old SALIENCE		−0.067 (0.060)		−0.065 (0.060)
Helper		0.089 (0.062)		0.086 (0.062)
Constant	0.625*** (0.070)	0.735*** (0.105)	0.625*** (0.070)	0.737*** (0.105)
<i>Inter-treatment comparisons (p-value; Cols III–IV)</i>				
SALIENCE vs. DEBIAS	—	—	0.434	0.433
SALIENCE vs. INSTITUTIONAL	—	—	0.423	0.422
DEBIAS vs. INSTITUTIONAL	—	—	0.996	1.000
# Observations	246	246	246	246
R^2	0.000	0.024	0.004	0.028

Table B-12 Differences in GPA among Nominated Students (Study 2)

The dependent variable is the nominee's GPA (4-point scale, centered on the faculty-wide mean). Columns I–II collapse the three nudges into a single NUDGE indicator, Column I without, Column II with faculty covariates. Columns III–IV include separate indicators for SALIENCE, DEBIAS, and INSTITUTIONAL; the omitted category is CTRL. Robust standard errors are shown in parentheses. Columns II and IV add faculty covariates: gender, tenure status, prior-study treatment (Old SALIENCE), and the *helper* indicator (if the faculty nominated a student in Study 1). The bottom panel reports p -values from pair-wise F -tests comparing treatment coefficients in the disaggregated specifications (Salience vs. Debias, Salience vs. Institutional, Debias vs. Institutional). Statistical significance is denoted by * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

	I	II	III	IV
NUDGE (Pooling)	0.005 (0.047)	0.018 (0.047)		
SALIENCE			−0.006 (0.057)	0.006 (0.057)
DEBIAS			0.022 (0.054)	0.040 (0.053)
INSTITUTIONAL			0.001 (0.054)	0.010 (0.055)
Female		0.023 (0.037)		0.025 (0.037)
Tenure		0.032 (0.046)		0.031 (0.046)
Old SALIENCE		0.072* (0.036)		0.073* (0.036)
Helper		−0.014 (0.037)		−0.014 (0.037)
Constant	4.314*** (0.043)	4.215*** (0.064)	4.314*** (0.043)	4.215*** (0.064)
<i>Inter-treatment comparisons (p-value; Cols III–IV)</i>				
SALIENCE vs. DEBIAS	—	—	0.580	0.494
SALIENCE vs. INSTITUTIONAL	—	—	0.892	0.938
DEBIAS vs. INSTITUTIONAL	—	—	0.653	0.529
# Observations	242	242	242	242
R^2	0.000	0.032	0.001	0.034

C Theoretical model

In this section, I present the theoretical framework that guides my empirical analysis. Drawing on random-utility discrete-choice theory (e.g. [McFadden 1973](#)) and salience theory ([Bordalo et al. 2012, 2016](#)), I formalize how organizational gatekeepers nominate candidates and predict how a light-touch salience nudge shifts those nominations. I use the multinomial-logit specification because it is the workhorse model for settings in which each decision-maker selects exactly one option from a finite set, and because it lets me model the nudge as a simple attenuation of the bias parameter.

Building on this framework, I model an organizational gatekeeping environment in which each *gatekeeper* (e.g., a faculty member) i nominates one candidate j from the set of candidates to whom they are connected, J_i . Formally, $J_i = L_i \cup H_i$, where L_i (H_i) is the subset of i 's connections to candidates who belong to the low-status (high-status) group.

Each candidate j is characterized by two observable attributes: quality and status. I denote the candidate's quality by $q_j \in \mathbb{R}$, a continuous measure such as GPA. I denote group membership with a binary status indicator $s_j \in \{0, 1\}$, where $s_j = 1$ when the candidate belongs to the low-status group, and $s_j = 0$ otherwise.

Baseline utility and nomination probabilities

In the absence of any intervention (CTRL condition), I model the utility that gatekeeper i gets from nominating candidate j as:

$$U_{ij} = \alpha q_j - \beta s_j + \varepsilon_{ij}, \quad \alpha > 0, \beta \geq 0. \quad (1)$$

The parameter α scales the marginal value of candidate quality, while β captures the disutility of selecting a low-status candidate ($s_j = 1$), reflecting bias or inattention.¹

I treat the idiosyncratic shocks ε_{ij} , which capture candidate-specific factors such as prior interactions with the gatekeeper, as independent and identically distributed across all

¹ When $\beta = 0$, the gatekeeper is unbiased ex-ante, so any salience nudge leaves choices unaffected; see [Remark 1](#).

(i, j) pairs. Under this assumption, the probability that gatekeeper i nominates candidate j takes the standard multinomial-logit form:

$$P_{ij}^{\text{CTRL}} = \frac{\exp(\alpha q_j - \beta s_j)}{\sum_{k \in J_i} \exp(\alpha q_k - \beta s_k)}. \quad (2)$$

A salience nudge

The experimental treatment (SALIENCE condition) introduces a nudge that makes low-status candidates *cognitively salient*. Following [Bordalo et al. \(2012, 2016\)](#), I model salience as an attenuation of the bias parameter: the baseline penalty β is reduced to $\beta - \delta$, with $\delta \in [0, \beta]$.² The utility a gatekeeper gets after receiving the SALIENCE nudge is:

$$U_{ij}^{\text{SALIENCE}} = \alpha q_j - (\beta - \delta) s_j + \varepsilon_{ij}. \quad (3)$$

The corresponding nomination probability, such that gatekeeper i nominates a candidate j after receiving the nudge is:

$$P_{ij}^{\text{SALIENCE}} = \frac{\exp(\alpha q_j - (\beta - \delta) s_j)}{\sum_{k \in J_i} \exp(\alpha q_k - (\beta - \delta) s_k)}. \quad (4)$$

After establishing the choice probabilities for both conditions, I now aggregate these probabilities across candidates and gatekeepers to derive the model’s testable predictions, which I summarize in the next subsection.

Gatekeeper-level and aggregate shares

To move from individual choice probabilities to a quantity I can compare across experimental treatments, I proceed in two steps.

² I treat δ as an attenuation factor for the bias, such that $0 \leq \delta \leq \beta$. However, if a nudge were strong enough ($\delta > \beta$) the sign would flip, turning a bias *against* low-status candidates into a bonus *in their favor*. There is some indicative evidence of this case in Study 2, where nomination rates are larger in magnitude than the share of low-SES students in the population (at the time of the study): 47% vs. 40%.

Step 1: Gatekeeper-level probability. For every gatekeeper i I group the logit weights into two totals:

$$A_i = \sum_{j \in L_i} \exp(\alpha q_j - \beta), \quad B_i = \sum_{j \in H_i} \exp(\alpha q_j).$$

where A_i is the total weight the logit rule assigns to *low-status* candidates, and B_i is the total weight on *high-status* candidates³.

Using these two totals, the chance that i nominates a low-status candidate is:

$$\pi_i^{\text{CTRL}} = \frac{A_i}{A_i + B_i}, \quad (5)$$

$$\pi_i^{\text{SALIENCE}} = \frac{e^\delta A_i}{e^\delta A_i + B_i}. \quad (6)$$

Step 2: Aggregate share. Because each gatekeeper nominates exactly one candidate, the expected proportion of low-status nominees in *treatment* $t \in \{\text{CTRL}, \text{SALIENCE}\}$ is simply the average of the gatekeeper-level probabilities:

$$\text{Share}_{\text{Low}}^t = \frac{1}{N} \sum_{i=1}^N \pi_i^t.$$

In words, π_i^t is gatekeeper i 's own chance of nominating a low-status candidate, and averaging those chances across all N gatekeepers turns individual decisions into the model's ex-ante prediction for the overall share of low-status nominees under treatment t .

Theoretical predictions

Equations (5)–(6) show that the SALIENCE nudge multiplies every low-status logit weight by the common factor $e^\delta > 1$ while leaving high-status weights unchanged. The next results translate that mechanical shift into testable implications.

³ For an individual candidate j , the baseline (CTRL) logit weight is $\exp(\alpha q_j - \beta s_j)$. Under SALIENCE that weight is multiplied by the common factor e^δ whenever j is low-status ($s_j = 1$). Summing over $j \in L_i$ and $j \in H_i$ results in A_i and B_i , respectively; the ratio $A_i/(A_i + B_i)$ is gatekeeper i 's probability of nominating a low-status candidate.

Lemma 1 (Gatekeeper-level equity) *If gatekeeper i faces at least one low- and one high-status candidate ($|L_i|, |H_i| \geq 1$), then $\pi_i^{\text{SALIENCE}} > \pi_i^{\text{CTRL}}$ for every $\delta > 0$.*

Proof. Let $A = A_i$ and $B = B_i$ and note from (6) that $\pi_i^{\text{SALIENCE}} = e^\delta A / (e^\delta A + B)$. Differentiating with respect to δ yields $\partial \pi_i^{\text{SALIENCE}} / \partial \delta = AB e^\delta / (e^\delta A + B)^2 > 0$

Hence π_i^{SALIENCE} is strictly increasing in δ , and the inequality follows for all $\delta > 0$. \square

Lemma 2 (Aggregate equity) *Because $\text{Share}_{\text{Low}}^t = N^{-1} \sum_i \pi_i^t$ is a simple average, Lemma 1 implies*

$$\text{Share}_{\text{Low}}^{\text{SALIENCE}} > \text{Share}_{\text{Low}}^{\text{CTRL}} \quad \text{whenever } \delta > 0.$$

Lemma 3 (Nomination-rate neutrality) *The overall nomination rate is identical in the two treatments:*

$$\text{NominationRate}^{\text{SALIENCE}} = \text{NominationRate}^{\text{CTRL}}.$$

Proof. Each gatekeeper is asked to nominate exactly one candidate, regardless of treatment. Summing over all N gatekeepers therefore produces the same total number of nominations in both arms by construction. \square

Lemma 4 (Quality neutrality) *Let q_j denote candidate j 's observable quality. Then*

$$\mathbb{E}[q \mid \text{SALIENCE}] = \mathbb{E}[q \mid \text{CTRL}].$$

Proof. For gatekeeper i , $\mathbb{E}_i[q \mid t] = \sum_{j \in J_i} P_{ij}^t q_j$, where $t \in \{\text{CTRL}, \text{SALIENCE}\}$. Splitting the sum by status yields

$$\mathbb{E}_i[q \mid t] = \frac{\sum_{k \in H_i} e^{\alpha q_k} q_k + \sum_{k \in L_i} e^{\alpha q_k - \beta + \mathbf{1}_{\{t=\text{SALIENCE}\}} \delta} q_k}{\sum_{k \in H_i} e^{\alpha q_k} + \sum_{k \in L_i} e^{\alpha q_k - \beta + \mathbf{1}_{\{t=\text{SALIENCE}\}} \delta}}.$$

Moving from CTRL to SALIENCE multiplies every low-status term in both numerator and denominator by e^δ , which cancels out. Consequently,

$$\mathbb{E}_i[q \mid \text{SALIENCE}] = \mathbb{E}_i[q \mid \text{CTRL}] \quad \text{for every } i.$$

Averaging over gatekeepers preserves the equality. \square

Remark 1 (No underlying bias) *If $\beta = 0$ the penalty on low-status candidates is already zero, so the nudge has no leverage: Equations (5)–(6) give $\pi_i^{\text{SALIENCE}} = \pi_i^{\text{CTRL}}$ and Lemma 2 collapses to an equality. The intervention can raise low-status representation only when a positive bias ($\beta > 0$) is present. This is the conjectured case for female candidates.*

Taken together, the Lemmas predict that a light-touch salience nudge will mitigate selection biases: it increases the representation of low-status candidates while leaving both the number of nominations and the average quality of nominees unchanged. I summarize these Lemmas into three testable hypotheses in Section 2.5 of the main manuscript.

D Additional results

D.1 Performance and expectations by social group

Figure C-1 illustrates performance and labor market expectations by SES and gender. Low-SES students outperform their high-SES peers academically by 0.15 standard deviations (GPA: 0.10 vs. -0.05, $p < 0.001$) but expect longer job searches, lower starting salaries on their first job (wages), and perceive their social connections made in college to be less effective in helping secure employment after graduation. Similarly, female students achieve higher GPAs than males with a difference of 0.24 standard deviations (standardized GPA: 0.10 vs. -0.14, $p < 0.001$), but report systematically lower expectations for their labor market outcomes.⁴




Figure C-1 Performance and expectations by social groups.

Panel A shows standardized GPA differences by socio-economic status and gender. **Panel B** presents students' expectations related to job search duration (Time), anticipated wages in their first job, and the utility of university connections in helping secure employment. Diamonds represent estimates for low-SES students; circles for female students. Horizontal lines show 95% confidence intervals.

⁴ These job expectations were measured through a university-wide survey ($n=1924$), using the following questions: "How long do you expect it would take you to get your first job after graduation? [<6 months, 6-12, 12-24, >24 months]", "How many minimum wages do you expect to earn in your first job after graduation? [1-10]", and "Out of 100 peers who graduated from your program, how many found a job through the social connections they made at university? [0-100]".

D.2 Top performers

Here I report selection outcomes when I restrict the sample to top performers. That is, candidates with GPA above median.

Figure C-2A shows that top-performing low-SES students make up about one-third of all top performers (32%), but in the CTRL group faculty nominate them only 20% of the time ($p = 0.001$). In SALIENCE, the share of top-performing low-SES nominees rises to 26%, and the remaining gap is no longer statistically distinguishable ($p = 0.129$). Figure C-2B displays the coefficient plot of a regression analysis with top performers, which confirms the same results found before: after the nudge, nomination probabilities respond much more to how available these top low-SES students are in a faculty member's network. Even with performance held constant at the top of the grade distribution, gatekeepers continue to overlook disadvantaged talent, yet a brief reminder that makes this talent salient significantly improves their chances of being chosen.

Effects of the salience nudge on low-SES nominations

Nominations vs. availability of low-SES candidates, by treatment



Figure C-2 Effects of the salience nudge on top low-SES nominations.

Panel A compares the share of top performing low-SES students who receive a nomination with their share in faculty networks. **Panel B** displays treatment differences in the likelihood of nominating a top performing low-SES candidate using faculty characteristics as covariates.

D.3 Effectivity of nominations

Now I report results from an exploratory analysis of two groups who entered the training program at the same time, yet were invited through different channels. The first group (i.e., *Nominated*) pools all nominated students in CTRL and SALIENCE, who were

recommended by faculty in Study 1. The second group (i.e., *Direct*) are students who were invited directly in a follow-up study, as reported in Munoz (2024).

Nominated vs. direct invitation: Program participation

Participation (un)conditional on take-up, by targeting strategy




Figure C-3 Nominated vs. Direct invitation: Program participation.

Panel A reports the share of students who have completed each of the nine sessions of the training program, contrasting those nominated (pooling CTRL and SALINECE) with those invited directly. **Panel B** repeats the progression but keeps only those who finished session 1, i.e., conditional on take-up.

Figure C-3A reports completion levels for each of the 9 sessions of the training program, separately for those Nominated and those who received a Direct invitation. The line graph shows that receiving an invitation as a consequence of a nomination leads to significantly higher levels of completion than when invited directly. In fact, 67% of nominated students started the program (i.e., take-up) while only 41% of those with a direct invitation did ($p < 0.001$). Figure C-3B shows, however, that once students cross the threshold of the first session the curves converge. Conditional on take-up, there are no meaningful differences in completion between groups.

Some caveats in the interpretation are important given the invitations (and enrollment process) were different. As such, the comparison should be read as exploratory rather than causal. First, nominated students were explicitly told that a professor had recommended them, whereas the direct-invite cohort learnt that they had been chosen on the basis of strong academic performance. Second, every participant needed a faculty endorsement for registration, but those nominated received it without requesting it while those directly invited had to ask for it. Although evidence in Munoz (2024) suggests this does not impact take-up, the comparison between groups is not fully clean.

In short, faculty nomination appears to operate mainly on the extensive margin: they persuade more students to enroll, but it does not, on its own, help them persevere once

they have begun, compared to a group of students that were motivated to start with a direct invitation.

D.4 Results on excellence award by nudge

In this appendix, I report pooled and disaggregated outcomes of the different nudges tested in the mechanisms experiment of Study 2.

Effects of the nudges on **nomination** choices

Selection by Type of Prompts




Figure C-4 Effects of nudges on nomination choices.

Panel A compares nomination rates between CTRL and the pooled NUDGE condition. **Panel B** compares nomination rates focusing on differences between each nudge and the CTRL.

Effects of the nudges on **low-class** nominations

Selection by type of nudge




Figure C-5 Effects of nudges on low-SES nominations.

Panel A compares the share of low-SES students who receive a nomination between CTRL and the pooled NUDGE condition. **Panel B** compares the share of low-SES students who receive a nomination focusing on differences between each nudge and the CTRL.

Effects of the nudges on **female** nominations

Selection by type of prompt




Figure C-6 Effects of nudges on female nominations.

Panel A compares the share of female students who receive a nomination between CTRL and the pooled NUDGE condition. **Panel B** compares the share of female students who receive a nomination focusing on differences between each nudge and the CTRL.

Effects of the nudges on **GPA**

Differences in academic merit by type of prompt




Figure C-7 effects of the salience nudge on candidate quality.

Panel A compares the normalized GPA of students who receive a nomination between CTRL and the pooled NUDGE condition. **Panel B** compares the normalized GPA of students who receive a nomination focusing on differences between each nudge and the CTRL.

D.5 Spillovers between experiments

Here I test whether there are experimental spillovers. That is, whether being treated in study 1 (i.e., SALIENCE condition) impacts the choice to nominate a low-SES candidate three years later, in Study 2. Among the 252 faculty members who submitted a nomination in Study 2, 100 had also taken part in Study 1, providing a natural test for spillover effects from the earlier intervention.

Figure C-8 illustrates coefficient plots of different linear-probability models, where the sample is progressively broadened across four specifications. The first case focuses only

Experimental spillovers

Effect of being treated in Study 1 on the nomination of low-SES candidates in Study 2




Figure C-8 Experimental spillovers.

The figure displays coefficient estimates and 95 percent confidence intervals for the effect of prior assignment to the SALIENCE nudge (Study 1) on the probability of nominating a low-SES candidate in Study 2, across four sample specifications.

on those assigned to the CTRL in Study 2, to separate prior from current nudging. The second model restricts the analysis to faculty in the current SALIENCE condition, to test the effect of being treated twice. The third case pools all NUDGES together (i.e., excludes the CTRL). Finally, I look at all four treatments together. Across all four specifications there is no evidence of experimental spillovers.

E Details and deviations from the pre-registration

This section details the pre-registration process for the two studies, as well as any deviations or clarifications.

E.1 Study 1 - Nominations for training program

Pre-registration timeline and updates

Initial pre-registration The faculty referral experiment was initially pre-registered in As-Predicted.org under identifier #99285. At the time, the data-sharing agreement with the local partner university did not include access to the full network data (i.e., course enrollment registries) or complementary administrative datasets. The pre-registration was limited to administrative data on faculty members and the referred students.

Revised pre-registration After finalizing the new data-sharing agreement with the local partner but before accessing the network data, I updated the pre-registration to include the new dataset. This updated pre-registration was submitted to the AEA RCT Registry under identifier AEARCTR-0014135. The revised pre-registration reflects the inclusion of network and additional administrative data in the analysis.

Sample size adjustments

The original pre-registration estimated a sample size of approximately 600 faculty members, based on information provided by the university administration. This estimate included individuals who held the status of “faculty” but were primarily engaged in administrative roles and did not interact with students.

For the experiment, the final sample was restricted to faculty members who actively teach and interact with students. This adjustment resulted in an effective sample size of 528 faculty members.

E.2 Study 2 - Nominations for excellence awards

This experiment was pre-registered in the AEA RCT Registry under identifier AEARCTR-0015473. The intention of the analysis was exploratory, as different competing mechanisms were being evaluated: salience, debias and institutional. As such, I pre-registered two specific comparisons: pooling and separate.

E.3 Ethics

All studies reported in this paper were approved by the Ethics Review Board at Universidad Autónoma de Bucaramanga (UNAB), the local partner university. Faculty invited to the field experiments were not informed they were part of a study. As the initiative is implemented as part of an opportunity offered by the Office of International Relations of the local partner. There is, however, a consent form put in place at the local partner university, in which all students and faculty members at the beginning of their employment or their studies are informed that they will be part of research projects and that their administrative data can be used and shared with third-parties for research purposes. Both faculty and students can rescind their consent at any time. At the moment of the study none of those involved in any of the two studies had done so.

At the time of Study 2, the author was also affiliated with the *Luxembourg Institute of Socio-Economic Research* (LISER). In line with LISER's research-ethics policy, the complete protocol for Study 2, together with a full, ex-post revision of Study 1, was submitted to LISER's Ethics Review Panel and received approval on 7 March 2025 (Reference REC/2025/138.IPES). Because the panel required the retrospective assessment of Study 1, its decision explicitly covers both studies reported here, providing a unified ethical clearance in addition to the approval already granted by UNAB.