The Global Gender Gap in STEM Applications: Pipeline vs. Choice

Isaac Ahimbisibwe[a], Adam Altmejd[b], Georgy Artemov[c], Andrés Barrios-Fernández[d], Aspasia Bizopoulou[e], Martti Kaila[f], Jin-Tan Liu[g], Rigissa Megalokonomou[h], José Montalbán[i], Christopher Neilson[j,*], Sebastián Otero[k], Jintao Sun[l], Xiaoyang Ye[m]

[a]*Baylor University, USA*
[b]*SOFI, Stockholm University, Sweden*
[c]*University of Melbourne, Australia*
[d]*Universidad de los Andes, Chile*
[e]*VATT Institute for Economic Research, Finland*
[f]*University of Glasgow, UK*
[g]*National Taiwan University, Taiwan*
[h]*Monash University, Australia*
[i]*Swedish Institute for Social Research, Stockholm University, Sweden*
[j]*Yale University, USA*
[k]*Columbia University, USA*
[l]*Renmin University, China*
[m]*Amazon, USA*

Abstract

Women make up only 35% of global STEM graduates, a share unchanged for a decade. We separate gender differences in academic preparation (the pipeline) from gender differences in application decisions (the choice gap) across ten centralized university admissions systems, where eligibility is set by academic performance and assignment runs through variants of deferred acceptance, under which students have little incentive to misreport. Restricting attention to high-achieving students, for whom access constraints are least likely to bind, isolates application behavior from the most obvious barriers. The pipeline gap varies widely, from female disadvantage in Uganda to female advantage in Sweden, and broadly tracks labor-market gender parity. The choice gap does not: it is large and negative in every setting, and even among top scorers women are about 24 percentage points less likely than men to rank a STEM program first, a gap exceeding 20 points in eight of the ten. The pattern holds when STEM is counted anywhere on the list rather than only first. The gap is large enough that closing the pipeline gap entirely would still leave high-achieving STEM applicants majority male, between 57% and 75%, in every setting we study.

Keywords: gender inequality, STEM gender gap, centralized application platforms

1. Introduction

Despite decades of progress in educational attainment, women remain substantially underrepresented in science, technology, engineering, and mathematics (STEM). In 2024, they accounted for only 35% of STEM graduates worldwide, a share that has barely moved in a decade (UNESCO, 2024). Explanations for this persistence are typically grouped into two channels: a *pipeline* channel, gender differences in academic preparation and access to selective STEM programs (Card and Payne,

*Corresponding author. Email: christopher.neilson@yale.edu

2021, Aucejo and James, 2021, Humphries et al., 2023), and a *choice* channel, gender differences in field preferences conditional on eligibility (Zafar, 2013).

Field of study matters for later outcomes, and the evidence that this is causal rather than sorting is direct: exploiting admission cutoffs, Kirkebøen et al. (2016) and Hastings et al. (2013) show that the field a student enters changes later earnings substantially, with returns concentrated in the more selective programs, and admission to those programs also shapes who reaches top incomes and corporate leadership (Zimmerman, 2019, Altonji et al., 2016). Persistent gender gaps in entry to STEM are therefore relevant not only for equity but for efficiency, since losing talent there costs innovation, productivity and growth (Weinberger, 1999, Hoogendoorn et al., 2013, National Science Foundation, 2017, Hsieh et al., 2019).

Distinguishing pipeline from choice is empirically difficult, since it requires observing both program-specific eligibility and students' ranked application decisions. While Delaney and Devereux (2019, 2021) demonstrate the value of using application data to study gender differences in college applications and STEM choices, existing research has largely been confined to single-country settings, limiting our ability to assess whether these patterns generalize across distinct institutional and cultural contexts.

This paper addresses that challenge by leveraging administrative microdata from centralized admissions systems in ten settings across five continents: Australia, Brazil, Chile, China, Finland, Greece, Spain, Sweden, Taiwan, and Uganda. Despite large differences in population size, economic development, and gender norms, these systems share a key institutional structure: students submit ranked preferences over college–major combinations through coordinated platforms running deferred acceptance (DA) (Neilson, 2024), and are assigned to the highest-ranked option for which they are eligible. Eligibility depends almost exclusively on academic performance. This setting allows us to observe students' ranked choices while comparing women and men with similar admission prospects.

Within this setting, we focus on high-achieving students, defined as those in the top 10% of each setting's academic performance distribution. For this group, we examine two complementary measures of STEM choice: whether a student ranks a STEM program first, and whether a student lists one anywhere on the preference list. The top-10% cutoff is a reporting convention rather than a substantive assumption: Online Appendix B reproduces every result from the top 5% to the top 30% and across the full performance distribution.

The two measures capture different intensities of preference. Under DA, ranking a STEM program first means it was preferred to every other program the student could have entered, whereas listing STEM lower signals interest that other options outrank. Truthful ranking is optimal when lists are unconstrained (Roth, 1982), and although list-length limits vary widely across our settings, they do not order the gaps we find (Section 2).

Women are substantially underrepresented among high-achieving STEM applicants, and we decompose that underrepresentation into the pipeline and choice gaps defined in Section 3. Pipeline patterns vary markedly across settings and in sign, at every threshold we examine, yet women remain underrepresented even where they are a larger share of top students than men. The pipeline alone cannot account for what we observe.

By contrast, the choice gap is consistently large and negative. Among students in the top 10%, women are 23.7 percentage points less likely than men to rank a STEM program first, averaging across the ten settings, or 22.4 points pooling students; the gap exceeds 20 points in eight of the ten systems and holds under both measures everywhere. The central regularity is therefore not cross-country heterogeneity in the choice gap but its persistence, which extends to how the high-achieving group is defined: pooled across settings, the gap moves by less than one percentage point as that group widens six-fold, from the top 5% to the top 30%.

We also relate both gaps to cross-country gender parity in wages, as a proxy for norms in the domain most relevant to specialization decisions. The two behave very differently as parity rises: the pipeline gap narrows and can turn positive, whereas the choice gap barely moves, staying negative and sizable even in the most gender-equal settings. These correlations are descriptive, but they reinforce the main pattern—preparation varies strongly with context, STEM choices much less.

Our contribution is to provide, to our knowledge, the first cross-national decomposition of the STEM gender gap into a pipeline gap (access and preparedness) and a choice gap (application decisions), using harmonized application microdata. This evidence bridges two strands of research. A first strand disentangles pipeline and choice within single settings (Delaney and Devereux, 2019, Card and Payne, 2021), but its narrow scope limits external validity. A second strand, typically in the form of international reports (OECD, 2017, Encinas-Martín and Cherian, 2023, UNESCO, 2024), documents STEM gender gaps across education systems but cannot separate pipeline from choice, since it relies on enrollment outcomes that blend pipeline, choices, and admissions. We instead assemble a novel global dataset of ranked application lists, which records the programs students put themselves forward for rather than only where they ultimately enroll. Such a comparison has not previously been possible at this scale, because it requires harmonized application data across systems.

The point of working within DA platforms is comparative. In decentralized systems such as the United States, students choose where to apply under uncertainty about admission, so an observed gender difference in applications blends preferences with beliefs about admission chances, application costs, and the strategic construction of a portfolio. Coordinated DA platforms close much of that down: students rank programs rather than deciding where to apply, they have little incentive to misreport, and eligibility is a known function of academic performance rather than of a discretionary admissions review. What remains is a difference in what high-achieving women and men put at the top of their lists, in ten systems at once. That is the sense in which the design separates choice from pipeline: not that it identifies a preference parameter, but that it closes off the most common alternative accounts by construction.

Closing off misreporting is not the same as explaining the gap, and we do not claim otherwise. Two classes of explanation survive the design intact. Residual strategic incentives remain where lists are short or where applications precede scores, which we take up in Section 2. More importantly, a student who reports truthfully may still be reporting preferences formed under different information: women and men may hold different beliefs about what studying and working in STEM involves, may face or expect different returns to a STEM degree (Zafar, 2013, Wiswall and Zafar, 2018), or may search and learn less about STEM options before applying, so that the list they submit faithfully reflects a choice set they never fully explored (Agte et al., 2025). Our design speaks to the first class and not the second: we observe what students rank, not the beliefs that generated the ranking.

2. Data

This section outlines the institutional context and data for the ten settings we study: Australia, Brazil, Chile, China, Finland, Greece, Spain, Sweden, Taiwan, and Uganda.[1] Our data are nationwide in eight of the ten settings. In Australia we observe Victoria, a state holding 25% of the population and a strong proxy for the nation; in China, Ningxia, a more rural and less developed autonomous region holding 0.5% of the population.

Institutional coverage is broad across our settings, but what makes these systems suited to our question is how they allocate seats rather than how many institutions they include. In eight of the ten settings the centralized platform is the route into the most selective universities and their STEM programs, and where participation is partial the institutions outside it are typically less selective. Two settings depart from this. In Brazil the São Paulo state universities, which host some of the country's strongest STEM programs, run their own admissions, so the Brazilian sample covers the

[1] Online Appendix A provides additional details on each setting.

federal system rather than the full public sector. In Uganda the platform allocates only government-sponsored places, roughly a tenth of undergraduate seats at Makerere, the largest public university, with the remainder filled through privately sponsored admissions that each university runs itself, so the Ugandan sample covers the government-sponsored track rather than the selective sector as a whole. Elsewhere the same logic governs which students we observe: where selective seats are allocated centrally, a student who wants one must appear on the platform, so covering the top of the institutional distribution implies covering the top of the student distribution.[2]

The ten settings are far apart on every dimension that might be thought to govern the gaps we study: population runs from 5.5 million in Finland to 1.4 billion in China, GDP per capita from about $4,000 in Uganda to $66,000 in Taiwan, and the Gini index from 27.3 in Finland to 53.9 in Brazil. To characterize the gender norms most relevant to our question we use a wage parity index, defined as one minus the total gender wage gap, which ranges from 0.517 in Uganda to 0.965 in Greece; its construction is described in Section 4.4.[3]

Despite these differences, the systems share two features: they allocate most university seats through centralized platforms running variants of deferred acceptance, and eligibility is almost exclusively based on academic performance, typically high school grades and admission exam scores. Our data, drawn from the admissions agencies, include students' ranked applications, their performance records, and the program each is assigned to.[4]

Centralized admissions cover every university in Australia, Greece, and Sweden, and a subset elsewhere, from 36 of 38 in Finland down to 8 of 52 in Uganda and 132 of 2,448 in Brazil, whose 2,152 private institutions do not participate. Four settings charge no tuition among participating universities; financial aid is available in all ten.

Students apply to specific college–major combinations, choosing from 149 options a year in Uganda and more than 18,000 in China, and ranking at most two of them in Brazil against no numerical limit in Greece. All are DA mechanisms, which give students little incentive to misreport and so let us read the submitted list as an ordering over the programs a student was willing to enter. The value of this common structure is not that it makes the settings frictionless, but that it removes a class of explanations that would otherwise be hard to rule out.[5]

[2]The systems differ in how much of the wider cohort they capture, with Spain's entrance exam taken by 81% of a cohort and Uganda's pool smaller and more selected, but the choice gap is negative throughout the performance distribution in every setting, not only at the top; Online Appendix B reports the gap by score bin.

[3]Table A.I in Online Appendix A reports these and the admission-system characteristics summarized below.

[4]Because the periods covered differ across settings, Online Appendix D re-estimates on the cohort closest to 2018 in each; the results are very similar.

[5]Strategy-proofness weakens when the list is restricted (Fack et al., 2019), but the constraint has to bind before it

They also vary in scale, from 18.1 thousand yearly applicants in Uganda to 2.71 million in Brazil. Women outnumber men among applicants in eight of the ten settings, from 53.0% in Chile to 60.1% in Finland; the exceptions are Taiwan at 49.8% and Uganda at 43.9%. These patterns mirror findings for the United States (Goldin et al., 2006). The cohorts we observe span 2003–2012 in Greece at one end and 2021–2022 in Uganda at the other.

Two settings qualify the claim that eligibility runs through performance alone, and we state them here rather than leave them to the institutional appendix. Greek lists are unlimited in length but confined to two of five broad subject categories, so an applicant chooses a domain before ranking within it. Uganda gates STEM eligibility earlier and harder: a student whose A-level subject combination is non-STEM cannot apply to a STEM program at all, and that combination is fixed years before the application. In both settings, part of what we measure as choice is a track decision taken earlier, and in Uganda it is a decision that forecloses the option rather than merely discouraging it. This is a real limit on the interpretation, and it cuts in a knowable direction: Uganda has the smallest choice gap in our data and Greece sits mid-range, so neither setting is generating the pattern we report.[6]

The settings differ in other ways too—exam design, score scales, tested content—and how those differences bear on the two gaps is beyond what we attempt here. Our aim is to document within-system gender differences, and to ask whether a gap in STEM applications persists among students facing similar admission prospects.

3. Empirical strategy

Our empirical strategy exploits a defining feature of these centralized systems: because eligibility depends on academic performance, students with similar performance face similar admission probabilities. We therefore focus on high-achieving students, defined as those in the top 10% of their cohort on the mandatory components of each setting's performance measure (university admission exam, high school exit exam, or high school GPA), which lets us include all applicants regardless

can distort, and across our settings the caps do not order the gaps: counted in college–major cells they run from two in Brazil and 24 in China to a hundred in Taiwan, with no numerical limit in Greece (Online Appendix Table A.I). The smallest choice component in our sample is Uganda's, on a six-cell list, and the second smallest is Taiwan's, on a hundred-cell list, while Brazil's two-cell list, the tightest cap we observe, sits in the middle of the range. A mechanism operating through list length would have to reproduce that ordering, and we can see no account that does. Brazil also operates an iterative DA in which students revise across rounds, and the evidence on how well iterative designs sustain truthful reporting is mixed (Bó and Hakimov, 2020, 2022, Barahona et al., 2023).

[6] Online Appendix A gives both in full. A parallel restriction operates through prerequisites in several other settings, but only Greece and Uganda impose a categorical gate.

of intended field. These are the students most likely to gain admission to and succeed in selective STEM programs. The cutoff is a reporting convention: Online Appendix B repeats every exercise at the top 5%, 20%, and 30% and across the full distribution, and the choice gap is essentially unchanged. Thresholds are computed within cohort and system, so students are never ranked against a cohort other than their own.

We define STEM using ISCED-2013 two-digit fields of education and training,[7] grouping natural sciences, mathematics and statistics (05), information and communication technologies (06), and engineering, manufacturing and construction (07). Based on official national mappings, this is the finest partition we can construct consistently across all ten systems; finer within-STEM distinctions, such as separating engineering from natural sciences, are not available on a comparable basis everywhere.

Our analysis begins by characterizing the gender composition of high-achieving STEM applicants. We then decompose these gender differences by examining two gaps:

1. The pipeline gap: difference between women's and men's representation among students in the top 10% of each setting's academic performance distribution.

2. The choice gap: difference in the share of high-achieving women and men who rank a STEM program either as their top choice or at any rank of their preference list.[8]

These are not competing explanations of female underrepresentation but components of a single accounting identity. The number of high-achieving applicants to a field is the product of two margins: how many students of each gender reach the top group, and what share of them rank the field first. Because it is a product, taking logs makes the two margins additive, so the gender gap in applicants splits exactly into the two gaps defined above, with no residual.

Formally, let T_g denote the number of students of gender $g \in \{m, f\}$ in the top group and s_g the share of them who rank a given field first, so that $A_g = T_g\, s_g$. Taking logs of the ratio between genders gives

[7] See UNESCO Institute for Statistics, *International Standard Classification of Education (ISCED)*: `https://uis.unesco.org/en/topic/international-standard-classification-education-isced`. Business and law (code 04) and health (code 09) serve as comparison fields throughout; Online Appendix C reports those results in full.

[8] List-composition statistics, such as the STEM share of a list, are not comparable where list-length rules range from two programs (Brazil) to no limit (Greece).

$$\underbrace{\ln\left(\frac{A_m}{A_f}\right)}_{\text{gap among high-achieving applicants}} = \underbrace{\ln\left(\frac{T_m}{T_f}\right)}_{\text{pipeline component } P} + \underbrace{\ln\left(\frac{s_m}{s_f}\right)}_{\text{choice component } C} . \tag{1}$$

Note the orientation: the two gaps defined above are stated female minus male, so that a negative choice gap means women apply less often, whereas P and C are stated male relative to female, so that positive values favor men throughout the decomposition. The male share among high-achieving applicants is $\Lambda(P+C)$ with $\Lambda(z) = 1/(1+e^{-z})$, which yields the counterfactual we use below and in Online Appendix B: $\Lambda(C)$ is the male share that would obtain if the top group were gender balanced, and $\Lambda(P)$ the share if men and women applied at equal rates, each holding the other margin at its observed value.

Two properties organize what follows. First, the components are directly comparable, which the underlying percentage-point gaps are not: the pipeline gap is a difference in composition shares and the choice gap a difference in application rates. Second, P depends only on the composition of the top group, so it is common to every field within a setting and any cross-field difference is by construction attributable to C.[9] We conclude by examining whether these gaps correlate with cross-country wage parity, using a labor-market measure that is directly relevant for higher-education specialization decisions.

The choice component compares men and women within the top group rather than at a common score, so where men sit higher within it and application rates rise with performance, part of C reflects within-group composition. If composition were all that C measured, it should vanish as the group is narrowed. It does not vanish: recomputed on the top 5% it remains positive in all ten settings, from +0.28 to +1.08, and on the top 1% in all nine settings for which that group can be constructed, from +0.13 to +0.80. It does attenuate: across those nine the mean falls from +0.69 at the top 10% to +0.55 at the top 1%, and it is lower at the top 1% in eight of the nine. A composition contribution therefore cannot be ruled out. What the exercise establishes is narrower, and is what we rely on: composition cannot be the whole of C.

[9] Online Appendix E extends the identity to realized assignments and reports the additional component.

4. Results

4.1. Female representation among top STEM applicants

Figure 1 reports women's share among STEM applicants, defined as students who rank a STEM program first, twice for each system: among all applicants, and among students in the top 10% of the academic performance distribution. We focus on the first-choice definition because DA assigns students to the highest-ranked program for which they are eligible, so ranking STEM first is informative about the intensity of STEM preferences and bears directly on admissions outcomes. The any-choice definition, which counts a STEM program anywhere on the list, captures broader interest and separates students who do not apply to STEM at all from those who rank it lower.[10]

Women are a minority of first-choice STEM applicants in every system and in both samples, ranging among all applicants from 27.9% in Uganda to 38.8% in Brazil. Restricting attention to the top 10% raises the share in seven of the ten systems but closes the gap in none: women remain below 35% in six settings, lowest in Chile at roughly one quarter, and reach 41.6% to 45.3% in the other four. The any-choice definition raises the shares somewhat and leaves the pattern intact; Greece is the only setting in which women become a majority, at 53%.[11] Nor is the pattern specific to the top-10% threshold: the same six settings keep women below 35% at the top 5% and top 20% as well.[12]

The shares in Figure 1 combine the two forces of equation (1): the *pipeline gap*, women's representation among high-achieving students overall, and the *choice gap*, the propensity of high-achieving women to rank STEM relative to high-achieving men. We analyze each in turn.

4.2. The pipeline gap

Panel (a) of Figure 2 reports the *pipeline gap*, the female minus male share of students in the top 10% of the academic performance distribution; the column beside each setting's name gives women's share of that group.[13] It varies widely across our settings, and in sign: women are underrepresented among high-achieving students in Taiwan, Chile, and Uganda, and overrepresented in Australia, China, Greece, Spain, and Sweden; Brazil is balanced at every threshold and Finland changes sign

[10] We report the any-choice definition throughout Online Appendix B.

[11] These shares are the second stage of the funnel in Figure 4; Online Appendix B.8 reports the counterfactuals and Online Appendix B.6 the gaps at all three thresholds.

[12] Online Appendix B reports both definitions at all three thresholds.

[13] Online Appendix B.4 reports the same gap at the top 5% and top 20%, which we draw on below.

Figure 1: Women as a Share of STEM Applicants, All Applicants and Top 10%

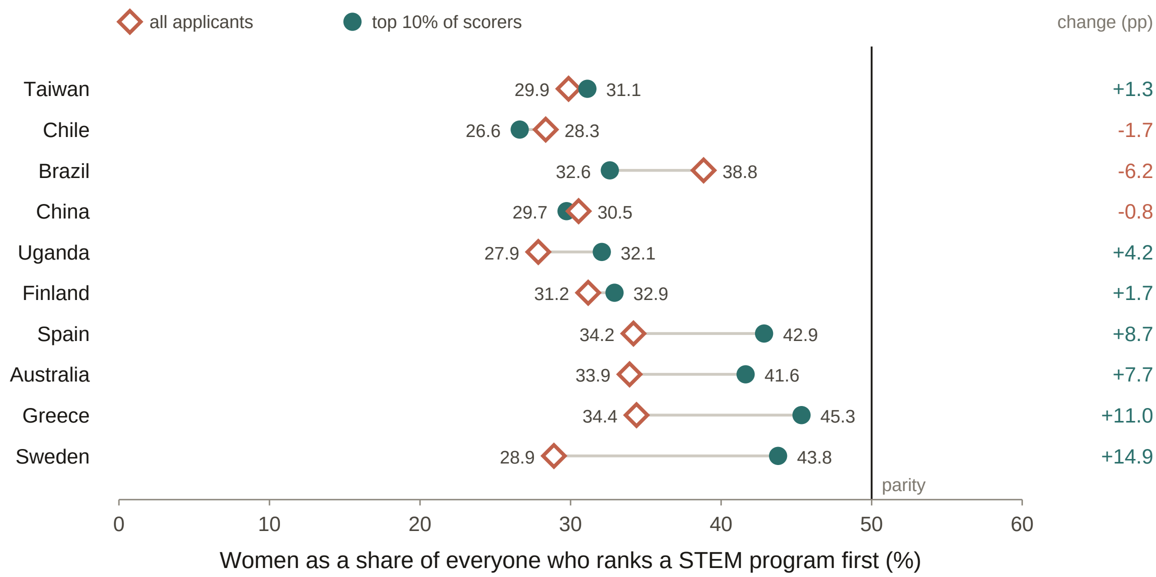

Notes: Women as a percentage of everyone ranking a STEM program first, among all applicants (diamonds) and among the top 10% of the academic performance distribution (circles). The right-hand column gives the change, in percentage points. Samples in Online Appendix Table A.II.

with the threshold. Uganda and Sweden mark the extremes, at 20 to 25 percentage points below parity and 31 to 33 above.

The pipeline gap therefore cannot account for the pattern in Figure 1. In every setting, and under either definition, women's share among high-achieving STEM applicants falls below their share among high-achieving students overall, including in the settings where they are the majority of that group. Factors beyond academic performance must be at work.

4.3. The choice gap

Following the analytical framework established by Delaney and Devereux (2019), we examine the *choice gap*, the extent to which high-achieving women and men differ in their likelihood of ranking a STEM program first. Panel (b) of Figure 2 reports the female-minus-male choice gap among students in the top 10% of the academic performance distribution, drawn on the same scale as the pipeline gap in panel (a). The column beside each setting's name gives the share of top-10% women who rank STEM first, from which men's rate follows by subtracting the gap.[14]

[14]Online Appendix B reports the same gap at the top 5% and the top 20%.

Figure 2: The Pipeline Gap and the Choice Gap among Top 10% Students

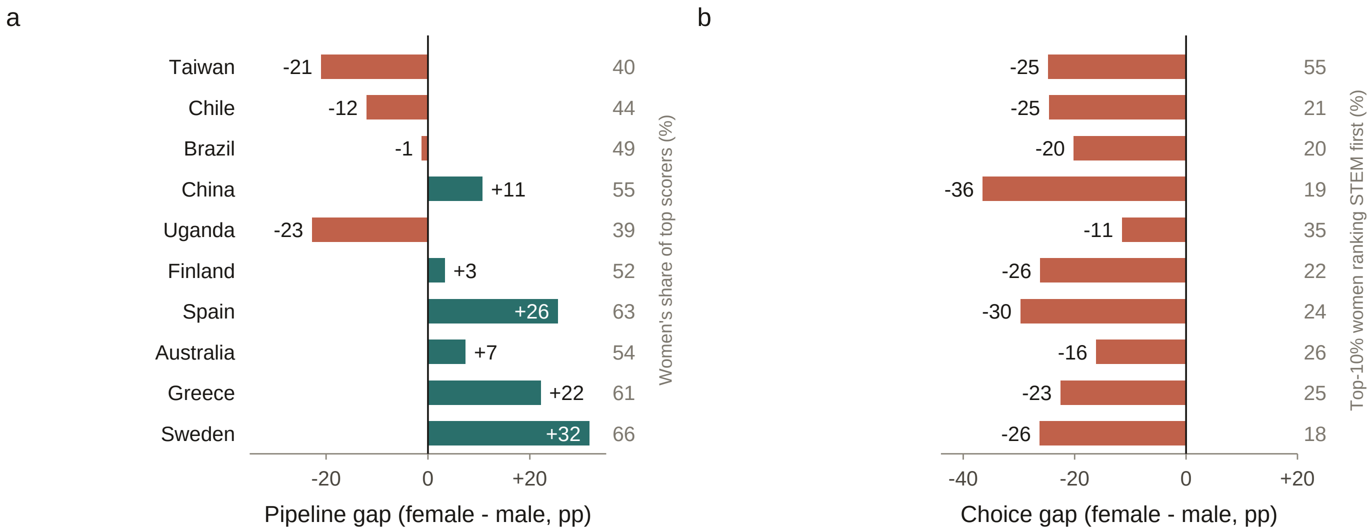

Notes: Female minus male, percentage points, top 10% of the academic performance distribution: (a) the pipeline gap, (b) the choice gap. Grey columns give the female level. The two are differences in different quantities and do not sum; Online Appendix E gives the components that do. For Chile, panel (b) conditions on the top-decile students who submit an application list, while panel (a), Table B.I and the components in Online Appendix C use the full top-decile totals; that convention moves Chile's choice gap from −23.1 to the −24.6 plotted here and leaves every other setting unchanged. Samples and denominators in Online Appendix Table A.II.

Two features support interpreting this gap as a difference in choice rather than access. First, admission depends on academic performance and submitted rankings, and the performance measure on which we rank students is constructed identically for women and men in every setting. Admission rules themselves are not everywhere sex-neutral, and Uganda is the clearest case: female applicants receive an additional 1.5 points on the official admission score, the largest public university reserves 40% of STEM places for women, and the published guide lists separate male and female STEM cutoffs (Online Appendix A.1). The choice gap is measured on submitted rankings rather than on admission outcomes, so such rules can reach it only through anticipation. In Uganda's case, anticipating a female advantage would push the gap toward zero, which is where Uganda sits, with the smallest choice component in our sample. Second, the gap is not an artifact of strategic behavior under limited lists: the list-length caps do not order the gaps we find (Section 2), and the gap persists across the full performance distribution, where such constraints bind hardest.[15] Single-country evidence documents similar patterns within centralized systems (Bordón et al., 2020).

[15]Online Appendix B reports the gap by score bin.

In contrast to the wide cross-country variation in the *pipeline gap*, the choice gap is negative in every setting we study. Focusing on first-choice rankings, it exceeds 20 percentage points in most systems at every threshold, and in eight of the ten at the top 10%, ranging from 12 to 17 points in Uganda to 33 to 39 points in China.

The gap is not an artifact of the first-choice measure. Counting STEM applications in any position of the preference list, which captures broader interest, it is smaller on average across settings, though marginally larger when students are pooled, and remains negative in every setting at every threshold, exceeding 20 percentage points in seven of the ten at the top 10%.[16]

The gap is equally insensitive to where the high-achieving group is drawn. Pooling the ten settings, the first-choice gap is −21.5, −22.4, −22.4 and −21.9 percentage points at the top 5%, 10%, 20% and 30%: a range of less than one point while the group widens six-fold, and negative in all forty setting-by-threshold cells. Individual settings move in both directions and largely offset, which is why the pooled figure is so flat. The concern that a gap measured in the extreme tail reflects selection into that tail therefore finds no support.[17]

The pipeline component is common to every field within a setting (Section 3), so any difference across fields is attributable to the choice component alone. Table 1 reports the components for all three fields, setting by setting. STEM and the two comparison fields, health and business and law, behave very differently. For STEM the component is positive in all ten settings, ranging from +0.29 in Uganda to +1.07 in China. For health it is negative in eight of ten, from −0.67 in Chile to −0.06 in Uganda: conditional on being high-achieving, women are *more* likely than men to rank health first. Under the any-choice measure that reversal is universal, holding in all ten settings, and the two first-choice exceptions, Greece (+0.08) and Taiwan (+0.04), correspond to gaps of less than two percentage points. For business and law the component carries no systematic sign at all, positive in four of the nine settings that report it and running from −0.88 in Uganda to +0.64 in Australia.

The contrast is sharpest in the counterfactual $\Lambda(C)$ of Section 3, the male share among high-achieving applicants to a field that would obtain with a gender-balanced top group. For STEM it lies between 0.57 and 0.75; for health, between 0.34 and 0.52. The ranges do not overlap. Closing the pipeline gap entirely would leave high-achieving STEM applicants majority male in every setting

[16] Further details are reported in Online Appendix B.

[17] Table B.I in Online Appendix B.1 reports these figures setting by setting. The margin that does move systematically with the threshold is admission rather than application: the gap among students admitted to STEM narrows from 18.2 to 14.2 percentage points between the top 5% and the top 30%, as cutoffs bind on a larger share of a wider group.

we observe, and health applicants majority female in most.[18]

Table 1: Decomposition of the gender gap among high-achieving applicants and assignees, by field

	Pipeline	Choice component C			Assignment	Balanced-pipeline
Setting	P	STEM	Health	Business/Law	V_{STEM}	male share $\Lambda(C_{\mathrm{STEM}})$
Australia	-0.15	+0.48	-0.25	+0.64	-0.05	0.62
Brazil	+0.02	+0.70	-0.60	+0.06	+0.05	0.67
Chile	+0.24	+0.77	-0.67	+0.12	-0.04	0.68
China	-0.21	+1.07	-0.10	-0.74	-0.17	0.75
Finland	-0.07	+0.78	-0.52	+0.40	-0.10	0.69
Greece	-0.45	+0.64	+0.08	–	+0.01	0.65
Spain	-0.52	+0.81	-0.46	-0.30	-0.08	0.69
Sweden	-0.66	+0.91	-0.40	-0.11	+0.06	0.71
Taiwan	+0.43	+0.37	+0.04	-0.83	+0.08	0.59
Uganda	+0.46	+0.29	-0.06	-0.88	+0.28	0.57
Positive in	4/10	10/10	2/10	4/9	5/10	

Notes: High-achieving students are those in the top decile of each setting's academic performance distribution. P is the pipeline component and C the choice component of (1); V is the assignment component of (E.1), which extends the identity to realized assignments. Positive values favor men, and P is common to all fields within a setting by construction. $\Lambda(C)$ is the male share among high-achieving STEM applicants that would obtain if the pipeline were balanced. Online Appendix E derives both identities and reports V for the other fields. Greece does not report business/law counts. Rates here use each setting's full top-decile totals as denominators, so that the components satisfy the identity exactly; Chile's panel (b) of Figure 2 instead conditions on students who submit an application, which is why Chile's C and its plotted gap are not directly comparable.

Within each setting this comparison holds constant the applicant pool, the performance threshold, the admission mechanism and any list-length constraint, so it also rules out the reading that the gap is a general feature of selective, high-return fields: business and law shows no systematic pattern, and health reverses.

4.4. The pipeline gap, the choice gap, and gender wage parity

Gender norms and labor-market conditions are often cited as potential drivers of gender differences in educational outcomes (Akerlof and Kranton, 2000, Guiso et al., 2008, Bertrand, 2020). We therefore relate the *pipeline* and *choice gaps* to a wage parity index, defined as one minus the total gender wage gap, so values closer to one imply greater parity.[19] Wage gaps come from the International Labour Organization (International Labour Organization, 2026), except for Australia and Taiwan, taken from official sources (Workplace Gender Equality Agency, 2026, Directorate-General of Budget, Accounting and Statistics (DGBAS), Taiwan, 2016), and China, from Zhang and Dong (2025).

The female share among students in the top 10% of the academic performance distribution rises with wage parity; in many of the countries with higher parity female shares exceed male shares, so

[18]Online Appendix C gives the underlying probabilities and gaps by country and threshold.
[19]Online Appendix G reports both relationships graphically.

the *pipeline gap* narrows and can even turn positive. The *choice gap* does not follow: it remains large and negative in every system, including those with the highest wage parity.

The decomposition in Section 3 sharpens this picture. Expressed as the components of (1), the two margins diverge as wage parity rises: the pipeline component P falls sharply (slope -2.25), while the choice component C shows no comparable decline and if anything rises (slope $+1.02$). The two are negatively correlated across settings (-0.65): in systems where women are better represented among high achievers, the women who are there are *less* likely to rank STEM first, not more. Figure 3 plots them against one another, one point per setting.[20]

This is consistent with the pattern documented by Stoet and Geary (2018) and debated since (Stoet and Geary, 2020, Richardson et al., 2020, Breda et al., 2020, Falk and Hermle, 2018), and it arises here within a single harmonized design rather than across heterogeneous international assessments, and without recourse to a composite gender-equality index. What the decomposition adds is that the paradox is a statement about one margin and not the other: aggregate measures combine P and C, which respond to parity with opposite signs, so the aggregate can appear flat even where one component moves systematically. We report this as description, not mechanism. With ten settings these slopes are imprecise, wage parity is one of many things that covary across these countries, and nothing here identifies what generates either component.

Wage parity, which is closely associated with prevailing gender norms, is therefore strongly correlated with the pipeline gap but not with the choice gap: even in relatively progressive settings, high-achieving women remain considerably less likely than similarly talented men to rank a STEM program first. The evidence is descriptive, but the persistence of this gap across such diverse contexts suggests that an important part of female underrepresentation in STEM is driven by factors common across settings.

4.5. From applications to placements

The gaps survive the assignment stage. Centralized allocation lets us observe not only what students rank but where they are placed. Figure 4 follows high achievers through four successive margins: pooled across settings, women are 51.2% of the top decile, 38.6% of those ranking STEM anywhere, 34.0% of those ranking it first, and 34.7% of those assigned to a STEM program. The two application margins account for essentially all of the shortfall, and the larger is whether STEM

[20]The slopes are $t = -3.7$ and $t = +2.0$ respectively. Online Appendix G reports how each slope responds to the choice of units and to dropping settings one at a time.

Figure 3: Pipeline and Choice Components of the STEM Gender Gap

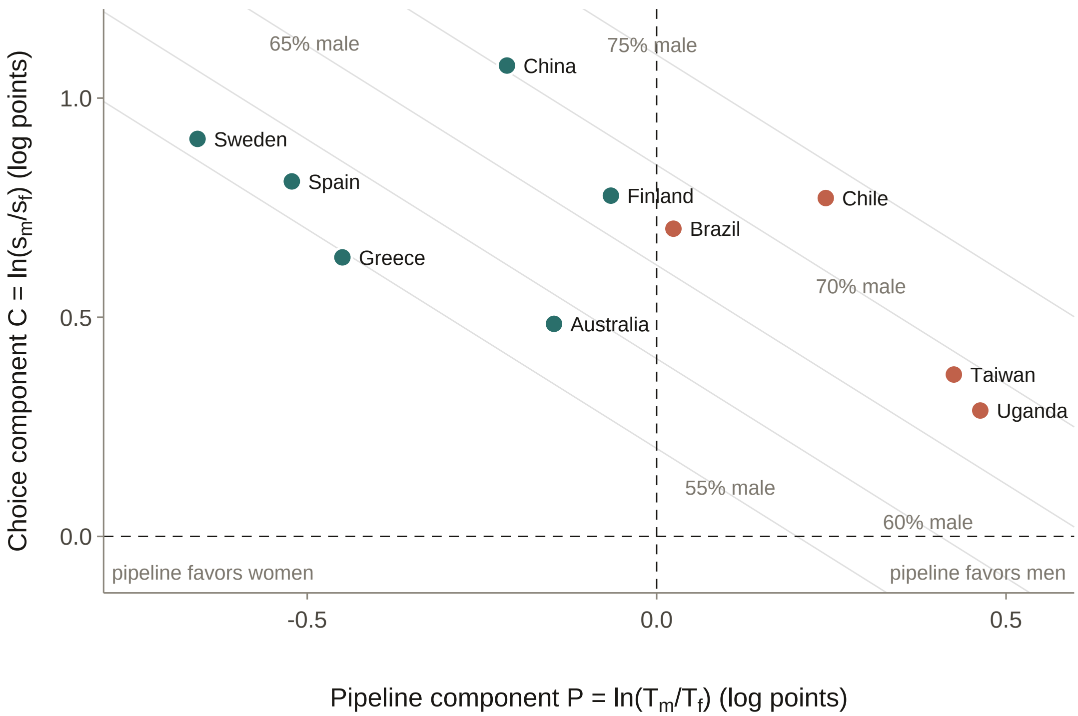

Notes: One point per setting, top decile. Axes are the components of (1), both positive where men are favored; grey diagonals are contours of constant male share among high-achieving STEM applicants. Online Appendix E derives them. Samples in Online Appendix Table A.II.

appears on the list at all rather than where it is ranked. The assignment stage adds a component smaller than the choice component in nine of the ten settings and comparable in the tenth, carrying no systematic sign and averaging +0.005.[21] Two qualifications apply: the comparison is one of aggregate composition and does not speak to any individual admission decision, and we observe the program a student is assigned to rather than the one she enrolls in.

What generates a choice gap this stable is beyond what these data can settle: admissions records observe choices, not the information, expectations and preferences that produce them.[22] Two clarifications matter. First, the stability of the gap is not evidence of innate or biological differences in preferences: application choices are formed within social, cultural and institutional environments,

[21]Online Appendix F sets out the assignment component formally; Online Appendix B.8 reports the counterfactuals implied by closing each margin in turn.

[22]The project's online materials set out the four families of explanation the literature offers, and what would be needed to distinguish them.

Figure 4: Female share of high achievers at each stage of the STEM funnel, top decile

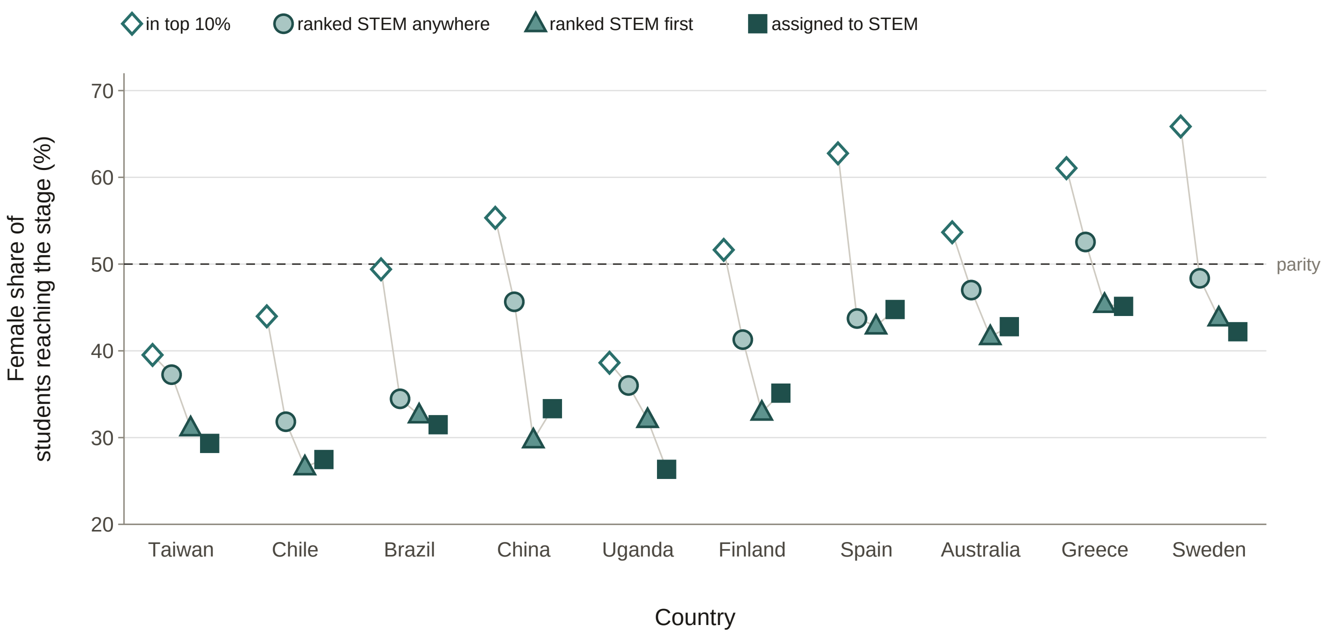

Notes: Female share of the students reaching each successive stage, top decile. These are compositions of the group reaching a stage, not the application-rate gaps reported elsewhere. Online Appendix B.8 reports the counts and the counterfactuals. Samples in Online Appendix Table A.II.

and our data cannot identify the channels through which those environments operate. Second, while the persistence of the gap suggests its drivers are at least partly shared, we cannot rule out that different combinations of context-specific forces produce similarly sized gaps in each setting.

5. Discussion and conclusion

The gender disparities we document in university applications matter for both equity and efficiency. Returns to higher education vary by field and are especially high in STEM, so women's underrepresentation there sustains labor-market gaps, and a better gender balance across fields would improve the allocation of talent (Hsieh et al., 2019, National Science Foundation, 2017, Weinberger, 1999, Hoogendoorn et al., 2013).

Our main contribution is to decompose women's underrepresentation in STEM into a pipeline gap and a choice gap. The pipeline gap varies substantially across settings—women make up under 40% of high achievers in Uganda and over 60% in Sweden—and that variation is a feature of the settings rather than of the cutoff, surviving at every threshold from the top 5% to the top 30%. It cannot on its own explain women's underrepresentation in STEM.

The choice gap behaves quite differently: it is large and negative in every setting, under both measures, and at every threshold we examine. Among top-10% students, women are 23.7 percentage points less likely than similarly talented men to rank STEM first, averaging across the ten settings, or 22.4 points pooling students.

Closing performance gaps is therefore not sufficient on its own: even where women are a clear majority of high achievers, as in Sweden, they remain a minority of STEM applicants. Prior research points to gendered preferences over non-pecuniary attributes such as stability and flexibility (Wiswall and Zafar, 2018) and to differences in program tastes (Zafar, 2013), and finds that pipeline factors alone cannot account for persistent underrepresentation (Patnaik et al., 2021). What we add is that these differences hold across societies with widely varying levels of development and gender norms, in contexts as different as Chile, Greece, Sweden, and Taiwan.

Understanding the mechanisms behind this stable choice gap remains a central challenge. It may reflect differences in how men and women value job attributes, expectations of discrimination, family considerations, identity and belonging, or self-efficacy. Our data cannot distinguish among these, nor rule out that different combinations of them produce gaps of similar size in each setting. What the evidence does establish is that the gap survives in settings that differ in almost every respect one might expect to matter, and that closing it is a separate task from closing the performance gap—an objective with implications not only for equity but also for realizing the efficiency gains from a more inclusive allocation of talent.

Finally, we contribute a market-design perspective. Deferred acceptance gives students little incentive to misreport, so the rankings these systems collect are unusually informative about how applicants order the programs open to them, and focusing on high-achieving students strips out the most obvious differences in access. That does not make the settings frictionless; a gap of this size is itself evidence that something is operating on women's applications. What the design buys is the ability to set aside the more mechanical explanations by construction rather than by assumption: strategic misreporting, binding list constraints, and differential admission chances at a common cutoff.

Read that way, the contribution is an existence result. In ten admissions systems that differ in almost every respect one might expect to matter, high-achieving women and men still rank STEM differently, and by a strikingly similar amount. This does not identify why, and we have been explicit that it cannot: a truthful ranking can still express beliefs that are inaccurate, or a choice set that was never fully searched. What it does establish is that whatever produces the difference is not confined to one country's institutions, labor market, or norms, which is what makes it worth pursuing with data that observe beliefs and search as well as choices.

Understanding when and how these differences emerge remains a key agenda for future work. As digital application platforms expand, the same approach can be used to study not only gender disparities but also racial, socioeconomic, and other inequalities in higher education at scale.

References

Agte, P., Allende, C., Kapor, A., Neilson, C.A., Ochoa, F., 2025. Search and Biased Beliefs in Education Markets. NBER Working Paper 32670. National Bureau of Economic Research. doi:10.3386/w32670. issued July 2024, revised October 2025.

Akerlof, G.A., Kranton, R.E., 2000. Economics and identity. The Quarterly Journal of Economics 115, 715–753.

Altonji, J.G., Arcidiacono, P., Maurel, A., 2016. The analysis of field choice in college and graduate school, in: Handbook of the Economics of Education. Elsevier, pp. 305–396. doi:10.1016/b978-0-444-63459-7.00007-5.

Aucejo, E., James, J., 2021. The path to college education: The role of math and verbal skills. Journal of Political Economy 129, 2905–2946.

Barahona, N., Dobbin, C., Otero, S., 2023. Affirmative action in centralized college admissions systems. Working Paper .

Bertrand, M., 2020. Gender in the twenty-first century. AEA Papers and Proceedings 110, 1–24. doi:10.1257/pandp.20201126.

Bó, I., Hakimov, R., 2020. Iterative Versus Standard Deferred Acceptance: Experimental Evidence. The Economic Journal 130, 356–392. URL: https://doi.org/10.1093/ej/uez036, doi:10.1093/ej/uez036.

Bó, I., Hakimov, R., 2022. The iterative deferred acceptance mechanism. Games and Economic Behavior 135, 411–433.

Bordón, P., Canals, C., Mizala, A., 2020. The gender gap in college major choice in chile. Economics of Education Review 77, 102011.

Breda, T., Jouini, E., Napp, C., Thébault, G., 2020. Gender stereotypes can explain the gender-equality paradox. Proceedings of the National Academy of Sciences 117, 31063–31069.

Card, D., Payne, A.A., 2021. High school choices and the gender gap in stem. Economic Inquiry 59, 9–28.

Delaney, J.M., Devereux, P.J., 2019. Understanding gender differences in STEM: Evidence from college applications. Economics of Education Review 72, 219–238. doi:10.1016/j.econedurev.2019.06.002.

Delaney, J.M., Devereux, P.J., 2021. Gender differences in college applications: Aspiration and risk management. Economics of Education Review 80, 102077. doi:10.1016/j.econedurev.2020.102077.

Directorate-General of Budget, Accounting and Statistics (DGBAS), Taiwan, 2016. Gender at a Glance: Gender Overall Earnings Gap (GOEG). Official Taiwan statistical release. URL: https://eng.stat.gov.tw/News_Content.aspx?n=2561&s=223911. gOEG estimate reported by DGBAS; accessed June 2026.

Encinas-Martín, M., Cherian, M., 2023. Gender, education and skills: The persistence of gender gaps in education and skills. OECD Skills Studies , 1–54.

Fack, G., Grenet, J., He, Y., 2019. Beyond truth-telling: Preference estimation with centralized school choice and college admissions. American Economic Review 109, 1486–1529.

Falk, A., Hermle, J., 2018. Relationship of gender differences in preferences to economic development and gender equality. Science 362, eaas9899.

Goldin, C., Katz, L.F., Kuziemko, I., 2006. The homecoming of american college women: The reversal of the college gender gap. Journal of Economic Perspectives 20, 133–156.

Guiso, L., Monte, F., Sapienza, P., Zingales, L., 2008. Culture, gender, and math. Science 320, 1164–1165.

Hastings, J.S., Neilson, C.A., Zimmerman, S.D., 2013. Are some degrees worth more than others? Evidence from college admission cutoffs in Chile. Technical Report 19241. National Bureau of Economic Research. doi:10.3386/w19241.

Hoogendoorn, S., Oosterbeek, H., Van Praag, M., 2013. The impact of gender diversity on the performance of business teams: Evidence from a field experiment. Management science 59, 1514–1528.

Hsieh, C.T., Hurst, E., Jones, C.I., Klenow, P.J., 2019. The allocation of talent and us economic growth. Econometrica 87, 1439–1474.

Humphries, J.E., Joensen, J.S., Veramendi, G.F., 2023. Complementarities in high school and college investments. Available at SSRN .

International Labour Organization, 2026. ILOSTAT Indicator EAR_GGAP_OCU_RT_A: Gender wage gap by occupation skill level. ILOSTAT data portal. URL: https://rplumber.ilo.org/data/indicator/?id=EAR_GGAP_OCU_RT_A. latest available year by country; accessed June 2026.

Kirkebøen, L.J., Leuven, E., Mogstad, M., 2016. Field of study, earnings, and self-selection. The Quarterly Journal of Economics 131, 1057–1111. doi:10.1093/qje/qjw019.

National Science Foundation, 2017. Women, minorities, and persons with disabilities in science and engineering. National Science Foundation and National Center for Science and Engineering Statistics Special Report NSF.

Neilson, C., 2024. The Rise of Coordinated Choice and Assignment Systems in Education Markets Around the World. Working Paper. The World Bank. Washington, DC. URL: https://thedocs.worldbank.org/en/doc/fa6ab76559c14f38def82e2cedf3b446-0050022024/the-rise-of-coordinated-choice-and-assignment-systems-in-education-markets-around-the-world.

OECD, 2017. The Pursuit of Gender Equality-An Uphill Battle. OECD publishing.

Patnaik, A., Wiswall, M., Zafar, B., 2021. College majors. The Routledge handbook of the economics of education , 415–457.

Richardson, S.S., Reiches, M.W., Bruch, J., Boulicault, M., Noll, N.E., Shattuck-Heidorn, H., 2020. Is there a gender-equality paradox in science, technology, engineering, and math (stem)? commentary on the study by stoet and geary (2018). Psychological Science 31, 338–341.

Roth, A.E., 1982. The economics of matching: Stability and incentives. Mathematics of Operations Research 7, 617–628.

Stoet, G., Geary, D.C., 2018. The gender-equality paradox in science, technology, engineering, and mathematics education. Psychological Science 29, 581–593.

Stoet, G., Geary, D.C., 2020. Corrigendum: The gender-equality paradox in science, technology, engineering, and mathematics education. Psychological Science 31, 110–111.

UNESCO, 2024. Global education monitoring report 2024, gender report: technology on her terms. Global Education Monitoring Report Team.

Weinberger, C.J., 1999. Mathematical college majors and the gender gap in wages. Industrial Relations: A Journal of Economy and Society 38, 407–413.

Wiswall, M., Zafar, B., 2018. Preference for the workplace, investment in human capital, and gender. The Quarterly Journal of Economics 133, 457–507.

Workplace Gender Equality Agency, 2026. ABS gender pay gap data. Australian Government official statistics page. URL: https://www.wgea.gov.au/data-statistics/ABS-gender-pay-gap-data. refers to November 2025 estimate released in 2026; accessed June 2026.

Zafar, B., 2013. College major choice and the gender gap. Journal of Human Resources 48, 545–595.

Zhang, W., Dong, X., 2025. Gender wage gap in early careers: Evidence from university graduates in china. Review of Development Economics 29, 1677–1692. URL: https://doi.org/10.1111/rode.13181, doi:10.1111/rode.13181.

Zimmerman, S.D., 2019. Elite colleges and upward mobility to top jobs and top incomes. American Economic Review 109, 1–47. doi:10.1257/aer.20171019.

Funding

Andrés Barrios-Fernández acknowledges partial support from ANID through FONDECYT grant 11230169, the Spencer Foundation through grant 10039719, and the YJS Foundation. Martti Kaila acknowledges partial support from the OP Foundation through grant 20230175. Georgy Artemov acknowledges support from the Australian Research Council grant DP160101350.

Acknowledgments

We thank the agencies in charge of university admissions in each country for granting us access to the administrative data we use in this project. Josefina Muñoz-Ávila provided excellent research assistance.

The Global Gender Gap in STEM Applications: Pipeline vs. Choice

Online Appendix

Contents

Appendix A. Institutional Details

This section documents the ten admission systems and the data we draw from each. Table A.I collects the setting, university-system and admission-system characteristics summarized in Section 2. Table A.II gives the administrative source, cohorts, sample size, and gender composition of the top decile—the group the main-text analysis is based on. Table A.III then sets the ten systems side by side on the features that bear most directly on how we read a submitted list: the score students are ranked on, whether they know that score when they apply, how places are assigned, and whether any admission rule differs by sex.

Fuller per-setting descriptions (each system's application calendar, the institutions inside and outside the centralized platform, and how the analysis sample was constructed) are available in the project's online materials, which have room for detail this appendix does not.

Table A.I: Institutional Characteristics

	Australia (1)	Brazil (2)	Chile (3)	China (4)	Finland (5)	Greece (6)	Spain (7)	Sweden (8)	Taiwan (9)	Uganda (10)
Panel A										
Setting Characteristics										
Population (millions)	24.6	209.5	18.7	1,402.8	5.5	10.7	46.8	10.2	24.0	40.1
GDP per capita (thousands)	56	21	31	24	56	41	50	62	66	4
Human Development Index	0.937	0.764	0.849	0.755	0.937	0.881	0.905	0.943	0.925	0.534
Gini index	33.7	53.9	44.4	38.5	27.3	32.9	34.7	30.0	34.2	42.8
Wage parity index (1− wage gap)	0.885	0.900	0.837	0.786	0.832	0.965	0.941	0.901	0.671	0.517
Panel B										
University System										
N of Institutions	21/21	132/2,448	34/60	1,252/2,663	36/38	41/41	50/86	41/41	65/71	8/52
Tuition fees	Yes	No	Yes	Yes	No	No	Yes	No	Yes	Yes
Financial aid	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Panel C										
Admission System										
Options available (yearly avg.)	1,078	6,310	1,423	18,671	1,458	609	2,169	15,374	1,659	149
Max. apps	12	2	10	24	6	No limit	12	20	100	6
N of applicants in a year (avg.)	38.8	2,711.0	249.7	52.4	69.6	68.0	171.6	108.2	72.1	18.1
N of admitted students in STEM (avg./year)	7.9	65.0	30.7	18.9	8.1	17.9	32.6	13.6	27.7	0.7
Female share (apps)	55.8%	57.1%	53.0%	55.5%	60.1%	55.8%	56.8%	57.4%	49.8%	43.9%
Data coverage	2009/10	2016	2012/18	2018	2016/20	2003/12	2020/24	2017	2010	2021/22

Notes: Panel A describes each setting, panel B its university system, and panel C its admission system. Australia is Victoria and China is Ningxia. "N of Institutions" gives those in the centralized system over the total. "Max. apps" counts college–major combinations; in China this is the first undergraduate batch, where a student ranks four institutions with six majors in each, and the cap applies within a batch rather than across the cycle. Applicant and admitted-student counts are yearly averages, in thousands; the applicant row counts the score holders who form each setting's performance distribution, which in Chile is DEMRE's registered candidates, about half of whom go on to submit an application. Performance groups are formed within setting and cohort; Table A.III gives each setting's score measure, and the setting-specific notes that follow record where the groups cannot be filled exactly. Sources and the wage parity index are described in Section 4.4.

Appendix A.1. Setting-specific notes

Uganda. Admission rules here are the most explicitly sex-conditional in our sample, and they favor women. Female applicants have received an additional 1.5 points on the official weighted score since 1990. Makerere University, the largest public university, admits students to STEM programs on a 40:60 female-to-male ratio from the 2020 academic year, and Mbarara, Gulu, Busitema and Soroti apply 30:70 ratios for the underrepresented sex; the published admissions guide correspondingly lists separate male and female cutoff points for STEM programs, while non-STEM programs carry a single cutoff. None of this enters the performance measure we use. We rank Ugandan applicants on an index we construct ourselves from raw national-examination grades, which is identical in construction for women and men and carries no adjustment; the official PUJAB weighted score, which does carry the adjustment, is program-specific and is not used here. Because applications are submitted before results are released, these rules can reach the choice gap only through anticipation,

Table A.II: Data sources and composition of the top 10% of the score distribution

				Students in the top 10%			
Country	Administrative source	Cohorts	N	Women	Men	Total	% Female
Australia	Victorian Tertiary Admissions Centre	2009–2010	77,648	4,158	3,590	7,748	53.7
Brazil	SISU platform / ENEM records	2016	2,711,049	133,899	137,150	271,049	49.4
Chile	DEMRE	2012–2018	1,747,925	76,534	97,486	174,020	44.0
China	Ningxia Department of Education	2018	52,427	2,879	2,324	5,203	55.3
Finland	Ministry of Education (HAREK, AMKOREK)	2016–2020	347,904	17,859	16,726	34,585	51.6
Greece	Ministry of Education	2003–2012	680,035	41,523	26,484	68,007	61.1
Spain	Ministry of Universities	2020–2024	857,870	53,783	31,908	85,691	62.8
Sweden	Swedish Council for Higher Education	2017	108,237	7,085	3,673	10,758	65.9
Taiwan	Joint Board of College Recruitment Comm.	2010	72,067	2,766	4,232	6,998	39.5
Uganda	Public University Joint Admissions Board	2021–2022	36,247	1,231	1,956	3,187	38.6

Notes: Counts are the number of students in the top decile of each setting's academic performance distribution, from the harmonized administrative tables supplied by each country team. Deciles are defined over the full distribution of score holders in each setting, and N reports that full distribution (deciles 1–10, female + male). Australia corresponds to the state of Victoria and China to the Ningxia Autonomous Region. The top decile contains close to but not exactly 10% of score holders in Uganda (8.8%) and Taiwan (9.7%), because coarse score scales generate ties at the cutoff; Uganda's score is the integer sum of two national exam grades and takes only 17 distinct values. The final column reports women as a percentage of the top-decile total, the pipeline measure plotted in Figure 2. Unless a figure note states otherwise, these country samples, cohort years, and sample sizes apply to all figures in the main text and the Online Appendix; Online Appendix D instead uses the cohort closest to 2018 in each country.

and anticipating an advantage would if anything push Uganda's gap toward zero—which is where Uganda sits, with the smallest choice component in our sample.

The performance measure is also unusually coarse. The index takes only seventeen distinct values, so cell sizes are uneven and ties mean the groups cannot be filled exactly: the top decile covers 8.8% of the cohort and the top 20% covers 14.7%, against 10% and 20% elsewhere. A true top-1% cutoff cannot be identified on a seventeen-value score, so where we report that group for Uganda it is the top score group, 547 students, or 1.5% of applicants. Interpolating between adjacent thresholds moves Uganda's gaps by under 0.03 log points.

Chile. Denominators differ here because the exam can be taken without submitting an application. Panel (b) of Figure 2 and the corresponding appendix figures condition on top-decile students who submit an application list, while Table B.I, Table 1 and the decomposition components use the full top-decile totals reported in Table A.II. The convention moves Chile's top-10% choice gap from −23.1 to −24.6 percentage points and leaves every other setting unchanged.

Australia. More than one candidate ranking variant is available. We use the Australian Tertiary Admissions Rank throughout, the rank the clearinghouse itself publishes and allocates on; a variant ranking students on English and mathematics results only is not comparable with the measures used in the other settings and is not used here.

Greece. Field coverage is incomplete: counts for business and law are available only for the top 5%, so Greece is absent from those columns of Table 1 and is marked n/a in the top-10% and top-20% panels of the corresponding figures. Its STEM and health counts are complete.

Table A.III: Admission Systems at a Glance

Setting	Performance measure	Scores known when applying?	Assignment mechanism	Sex-conditional admission rules
Australia (Victoria)	ATAR, a percentile rank built from scaled subject scores	Yes	Clearinghouse; offer rounds from January	—
Brazil (federal)	ENEM score, weighted by program	Yes	Iterative DA; provisional cutoffs updated daily	—
Chile	High-school GPA and PSU sections, weighted by program	Yes	Deferred acceptance	—
China (Ningxia)	Gaokao total, within province and track	Yes	Parallel-choice DA within each batch	—
Finland	Matriculation grades and/or entrance exam	Partly; most applicants already hold a completed certificate	DA re-run over the results period	—
Greece	Panhellenic examination score	Yes	Score-ordered allocation along declared preferences	—
Spain	Bachillerato GPA and EBAU, weighted by program	Yes	Score-ordered regional assignment	—
Sweden	Upper-secondary GPA and/or scholastic aptitude test	No; the list closes before results are released	Rank-ordered assignment in two rounds	—
Taiwan	AST subject scores, weighted by program	Yes	Program-proposing DA	—
Uganda	Index of national-examination grades, constructed by us	No; applications precede results	Merit-order selection	+1.5 points for female applicants; 40:60 STEM ratio at Makerere; separate published STEM cutoffs

Notes: One row per setting, describing the system as it operated in the cohorts we observe (Table A.II); several of these systems have since been reformed. “Performance measure” is the score on which we rank students to form performance groups; in every setting it is constructed identically for women and men. “Scores known when applying” records whether a student’s examination result is available at the moment the ranked list is submitted, which governs how much scope there is for admission-risk considerations to shape the list. “—” in the final column means we document no sex-conditional rule for that setting; we have not audited every program-level entry route in every system, and the entry for Uganda is the one we establish from the admissions guide and ministry sources. Institution counts, list-length caps, fees and cohort coverage are in Table A.I. Fuller per-setting descriptions, including application calendars and data construction, are available in the project’s online materials.

Appendix B. Pipeline and STEM choice gaps along the academic performance distribution

The main text defines high-achieving students as those in the top 10% of each setting's academic performance distribution. That focus is motivated by the fact that these students are most likely to gain admission to and succeed in selective STEM programs, i.e., the ones associated with high earnings, and that they are the least affected by strategic concerns in settings that limit the number of programs students can list. This section establishes that the pooled gap is stable to the cutoff, and that where individual settings do move across thresholds they offset one another rather than accumulate. We vary the threshold used to define high-achieving students and the definition of a STEM applicant, and we extend the analysis to students below the top 20%, so that every result in the main text can be read off the full performance distribution rather than a single group.

Appendix B.1. The choice gap is stable across thresholds

Table B.I reports the choice gap at four thresholds, from the top 5% to the top 30%, under both definitions of a STEM applicant. It is the compact statement of this appendix's main point. Pooling the ten settings, the first-choice gap moves from −21.5 to −22.4 percentage points and back to −21.9 as the group widens six-fold, a range of 0.9 percentage points; the any-choice gap varies by 0.6. The gap is negative in all forty setting-by-threshold cells under each definition.

Two features of the table are worth drawing out. First, the direction in which a setting's gap responds to the threshold is not uniform. Between the top 5% and the top 30% the first-choice gap widens in Taiwan (from −20.7 to −31.9 percentage points) and Australia (from −14.7 to −21.1), and narrows in Brazil, Chile, Finland and Spain; moving from the top 10% to the top 30% it narrows in five of the ten settings. These movements offset rather than accumulate, so the pooled gap is nearly flat. The top-10% column is the joint largest of the four, not one end of a trend. Second, the settings whose gaps move most across thresholds (Taiwan, Greece, Australia) are not the settings with the largest gaps, so the ordering across settings is largely preserved.

Table B.I: The STEM choice gap at alternative performance thresholds

Setting	Top 5%	Top 10%	Top 20%	Top 30%
Panel A. Ranks a STEM program first				
Australia	-14.7	-16.1	-18.9	-21.1
Brazil	-19.7	-20.2	-19.5	-18.9
Chile	-23.1	-23.1	-22.8	-22.3
China	-32.8	-36.5	-38.0	-38.5
Finland	-25.3	-26.2	-25.8	-24.6
Greece	-17.8	-22.5	-24.5	-24.3
Spain	-30.0	-29.7	-29.7	-29.1
Sweden	-23.5	-26.3	-26.0	-25.5
Taiwan	-20.7	-24.7	-28.8	-31.9
Uganda	-12.8	-11.5	-15.6	-17.4
Pooled	-21.5	-22.4	-22.4	-21.9
Range			0.9 percentage points	
Panel B. Ranks a STEM program anywhere on the list				
Australia	-17.0	-17.5	-19.3	-21.1
Brazil	-22.3	-22.7	-22.1	-21.5
Chile	-22.4	-22.7	-22.8	-22.2
China	-28.9	-26.4	-26.1	-25.8
Finland	-20.7	-21.9	-23.8	-24.0
Greece	-23.5	-25.0	-25.7	-25.7
Spain	-29.3	-29.0	-28.8	-28.3
Sweden	-25.1	-27.3	-26.8	-26.8
Taiwan	-6.1	-8.7	-11.0	-13.9
Uganda	-7.4	-9.4	-13.3	-16.4
Pooled	-22.3	-22.8	-22.9	-22.5
Range			0.6 percentage points	

Notes: Female-minus-male difference in the probability of applying to a STEM program, in percentage points, among students at or above each performance threshold. Negative values indicate that women are less likely than men to apply. *Pooled* counts all students in the ten settings together, and *Range* is the spread of the pooled gap across the four thresholds. The main text reports the top 10% column.

Appendix B.2. The choice gap at the top 1%

Table B.I asks what happens as the high-achieving group is made *wider*. Figure B.I asks the opposite question, and reports the top 1% alongside the top 5% and the top 10%. The group can be formed in nine of the ten settings; Greece is the exception, since its data reach us by decile and top 5%.

The sign is unchanged: the gap is negative in all nine settings at the top 1%, as it is at every other threshold we can construct. The magnitude is not. Averaged over the nine settings the gap falls from -23.8 percentage points at the top 10% to -18.2 at the top 1%, and it is smaller at the top 1% in seven of the nine, with Chile and Spain the exceptions. Two settings account for much of the movement: China falls from -36.5 to -17.0 and Taiwan from -24.7 to -11.5.

We report this rather than lean on it. The top 1% is a group of a few hundred students in several settings—524 in China, 547 in Uganda, 695 in Taiwan, 798 in Australia—and the two settings that move most are among the smallest, so some of the attenuation is likely noise. What the figure establishes is the same thing Table B.I establishes at the other end: the gap is a feature of the upper distribution rather than of any particular cutoff, and no threshold we can construct makes it disappear or change sign.

Figure B.I: The STEM Choice Gap at the Top 1%, 5%, and 10%

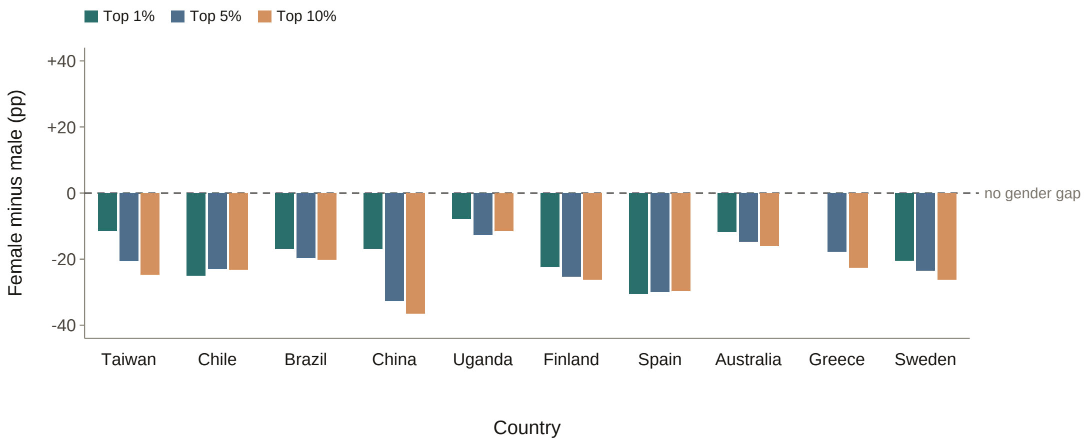

Notes: Female-minus-male difference in the probability of ranking a STEM program first, in percentage points, at three performance thresholds. Greece is omitted from the top 1% because its data are reported by decile and top 5%, so the group cannot be formed. Uganda's top-1% bar is its top score group, 1.5% of applicants, because its seventeen-value admission score admits no exact 1% cutoff (Online Appendix A.1). Denominators are all students in the score group throughout, the convention of Table B.I, with which this figure reconciles exactly; it therefore differs from Figures B.VI and B.VIII, where Chile conditions on applicants. Samples in Online Appendix Table A.II.

Appendix B.3. Gender shares among STEM applicants at alternative performance thresholds

Figure 1 in the main text reports women's share among first-choice STEM applicants, for all applicants and for students in the top 10% of the academic performance distribution. Figure B.II reports the same share at the top 5% and the top 20%, under both the first-choice and the any-choice definitions.

Figure B.II: Gender Shares among STEM Applicants at Alternative Performance Thresholds

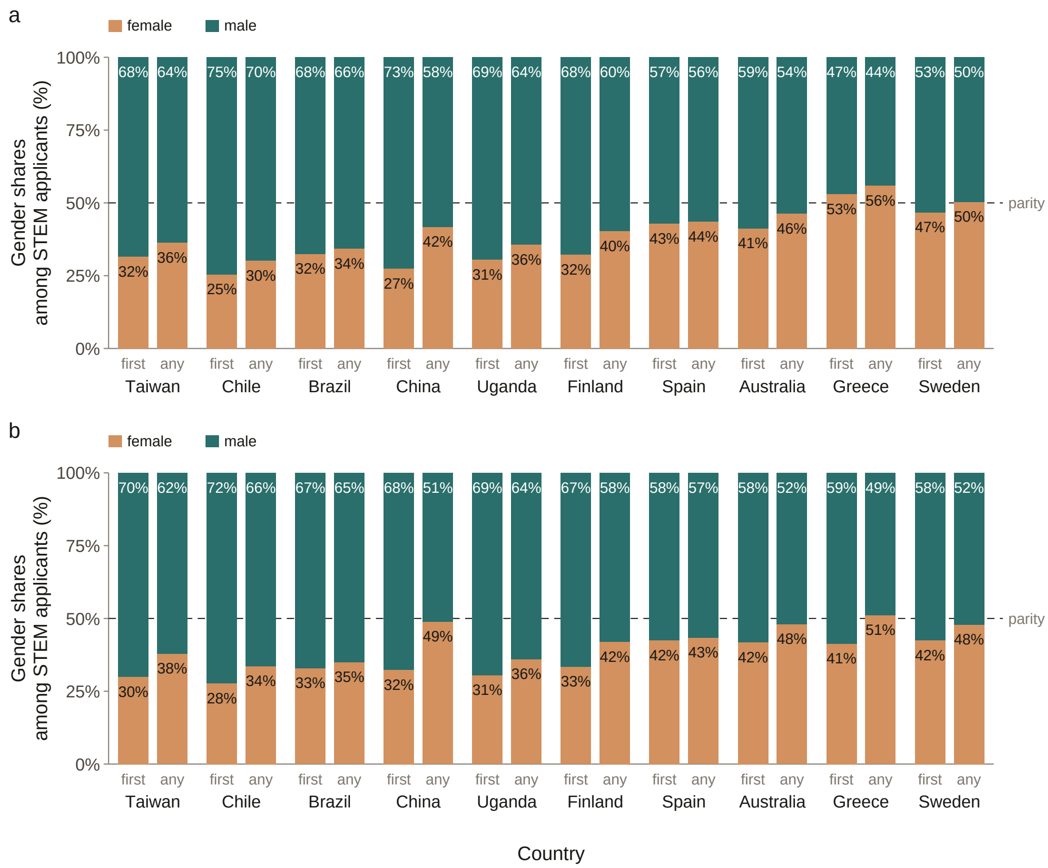

Notes: Female (lower) and male (upper) shares among STEM applicants, at the top 5% (a) and the top 20% (b); each bar spans 0–100 and the pair sums to 100 by construction. Within each country the left bar ranks STEM first, the right ranks it anywhere. Samples in Online Appendix Table A.II.

The main-text conclusion is unchanged at either threshold. Under the first-choice definition, the same six settings—Taiwan, Chile, Brazil, China, Uganda, and Finland—keep women below 35% of high-achieving STEM applicants, and Greece is the only case in which women exceed 50%, at 53.0% among top-5% students, falling to 41.4% at the top 20%. Under the any-choice definition, female shares rise, but the ordering across settings is largely preserved. The largest increases again occur in China, where the share rises by 14.3 percentage points at the top 5% and by 16.6 at the top 20%, remaining below 50% in both cases. Greece is the only setting where women are a majority under this measure at both thresholds, at 55.9% and 51.2%; Sweden reaches 50.3% among top-5% students but falls back to 47.8% at the top 20%.

Appendix B.4. The pipeline gap at alternative performance thresholds

Figure 2 in the main text reports the gender composition of the top 10% of the academic performance distribution and the associated pipeline gap. Figures B.III and B.IV repeat the exercise at the top 5% and the top 20%, the first reporting the shares by sex and the second the gap.

The main-text conclusion is unchanged at either threshold. Women remain underrepresented among high-achieving students in Taiwan, Chile, and Uganda, and overrepresented in Australia, China, Greece, Spain, and Sweden, at both thresholds. Brazil stays balanced, with the two shares within about one percentage point of each other everywhere. Finland is the only setting whose gap changes sign across thresholds, moving from −0.6 percentage points among top-5% students to +9.3 at the top 20%; Australia widens over the same range, from +4.9 to +11.8. Uganda records the largest negative gap at both thresholds (−24.7 and −20.4 percentage points) and Sweden the largest positive one (+30.7 and +32.6).

Figure B.III: Gender Shares among Top Students at Alternative Performance Thresholds

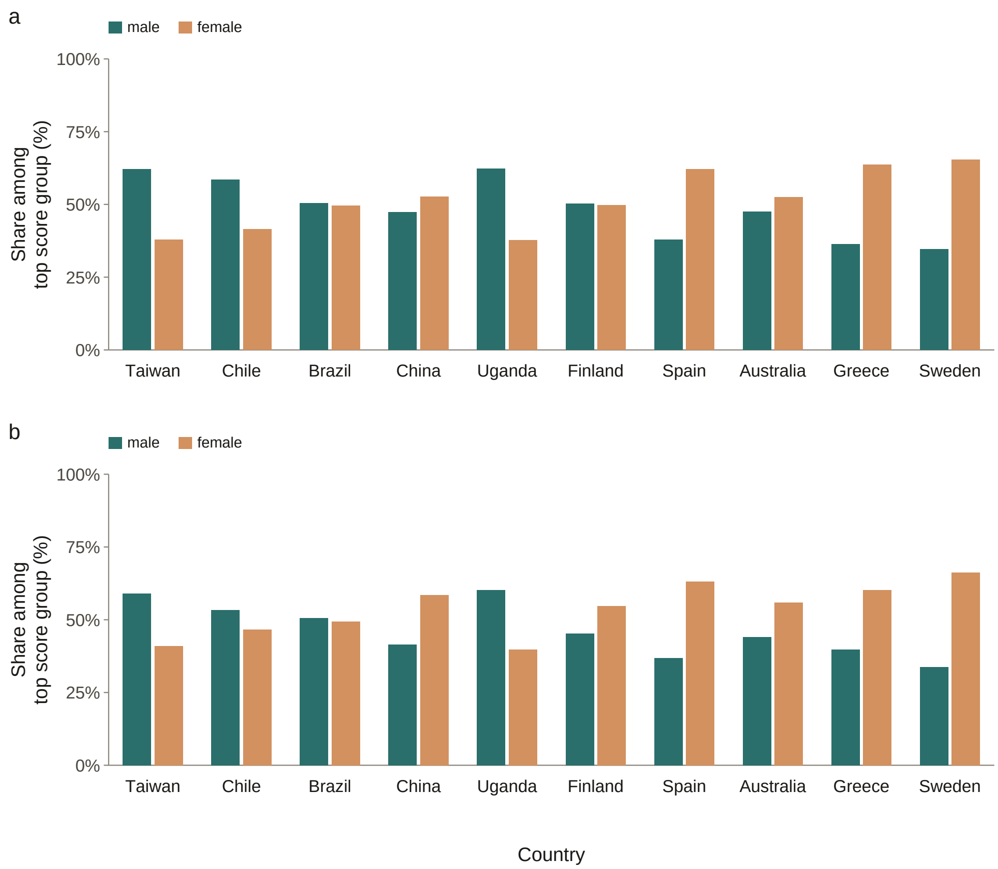

Notes: Shares of each performance group that are male and female, at the top 5% (a) and the top 20% (b). The female-minus-male pipeline gaps are reported in Figure B.IV. Samples in Online Appendix Table A.II.

Figure B.IV: The Pipeline Gap at Alternative Performance Thresholds

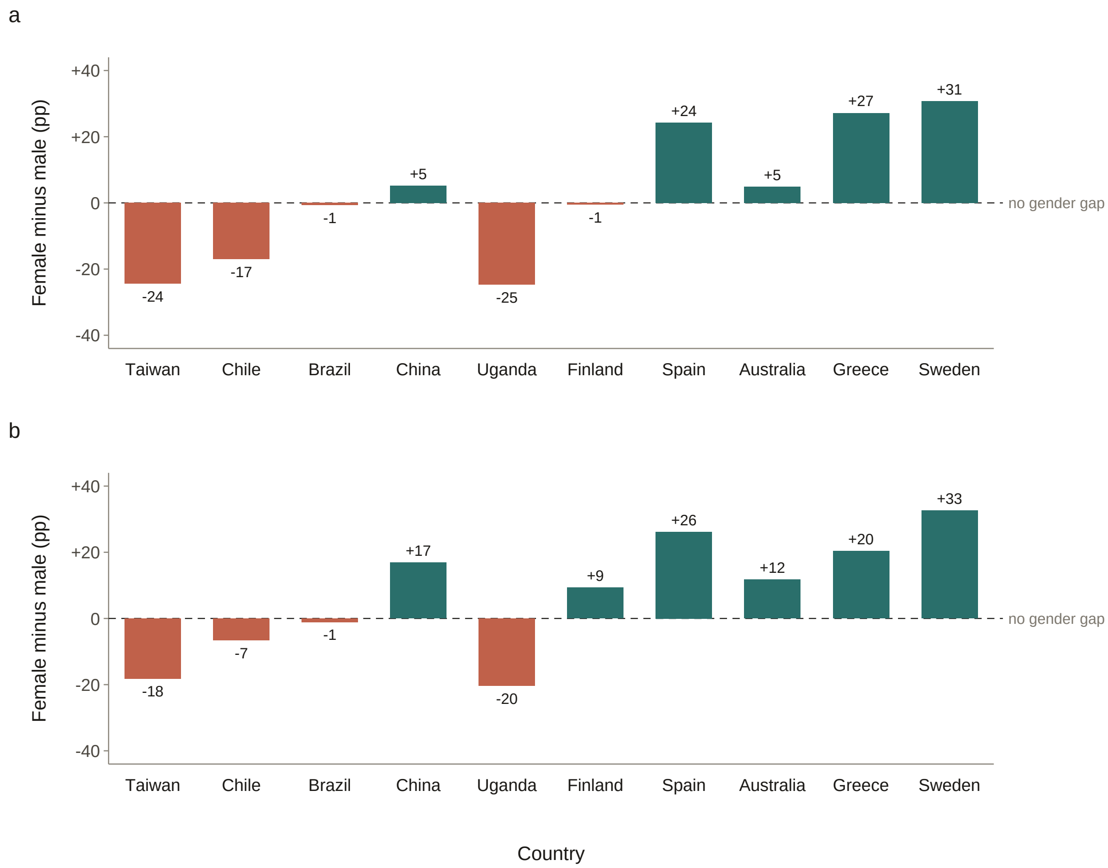

Notes: Female-minus-male pipeline gap, at the top 5% (a) and the top 20% (b). The underlying shares by sex are reported in Figure B.III. Samples in Online Appendix Table A.II.

Appendix B.5. The first-choice STEM gap at alternative performance thresholds

Figure 2 in the main text reports the share of top-10% women who rank a STEM program first and the associated gender gap. Figures B.V and B.VI repeat the exercise at the top 5% and the top 20%, the first reporting the probabilities by sex and the second the gap.

The main-text conclusion is unchanged at either threshold. The gap is negative in all ten settings at both thresholds, and it exceeds 20 percentage points in six of the ten among top-5% students and in seven at the top 20%. Uganda records the smallest gap at both thresholds (−12.8 and −15.6 percentage points) and China the largest (−32.8 and −38.0). The ordering across settings is broadly stable; the clearest exception is Taiwan, where the gap widens from −20.7 percentage points among top-5% students to −28.8 at the top 20%, moving it from the sixth to the third largest.

Figure B.V: STEM First-Choice Probabilities by Sex at Alternative Performance Thresholds

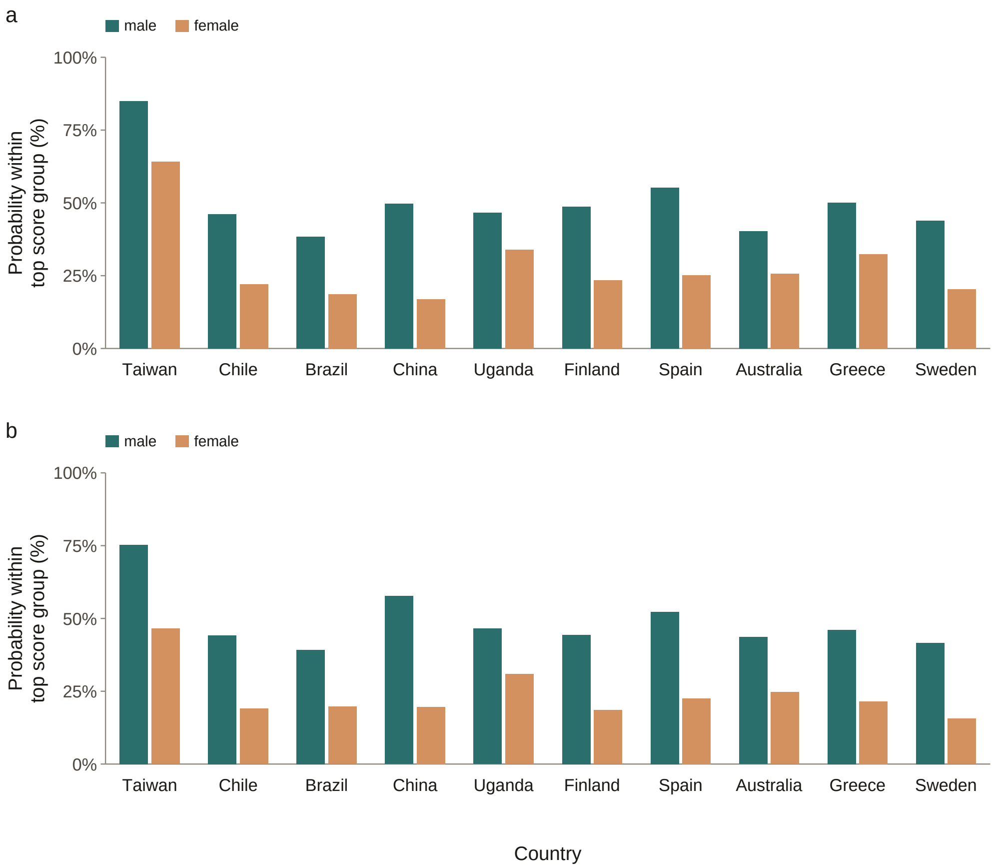

Notes: Probability of ranking a STEM program first, by sex, at the top 5% (a) and the top 20% (b). Chile conditions on applicants. The female-minus-male gaps are reported in Figure B.VI. Samples in Online Appendix Table A.II.

Figure B.VI: The STEM First-Choice Gender Gap at Alternative Performance Thresholds

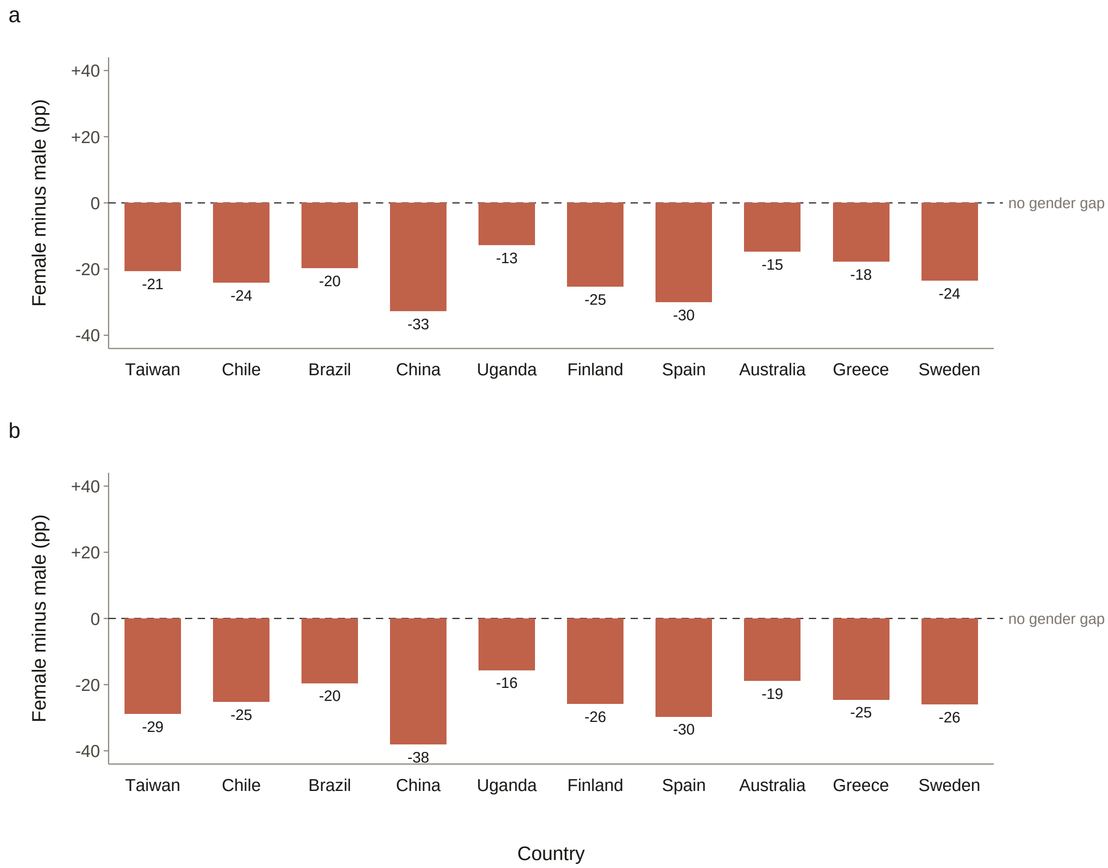

Notes: Female-minus-male gap in the probability of ranking a STEM program first, at the top 5% (a) and the top 20% (b). Chile conditions on applicants. The underlying probabilities by sex are reported in Figure B.V. Samples in Online Appendix Table A.II.

Appendix B.6. The choice gap under an any-choice definition of a STEM applicant

The main text measures the choice gap using students who rank a STEM program first, the definition most informative about the intensity of STEM preferences. Figures B.VII and B.VIII repeat the exercise under an any-choice definition, which counts a student as a STEM applicant whenever a STEM program appears anywhere on the application list, at all three performance thresholds: the first reports the probabilities by sex, the second the female-minus-male gap.

The main-text conclusion is unchanged. The gap is negative in all ten settings at every threshold. Averaged across settings it is smaller than under the first-choice measure, though pooling students it is marginally larger (-22.8 against -22.4 at the top 10%), and it is larger under the any-choice definition in Australia, Brazil, Greece and Sweden. Under either aggregation it remains substantial: it exceeds 20 percentage points in seven of the ten settings at each of the three thresholds. The three smallest gaps are found in Taiwan, Uganda, and Australia at all three thresholds, with Taiwan the smallest throughout (-6.1, -8.7, and -11.0 percentage points). No setting changes its position markedly across thresholds.

Figure B.VII: Any-STEM Application Probabilities by Sex at Alternative Performance Thresholds

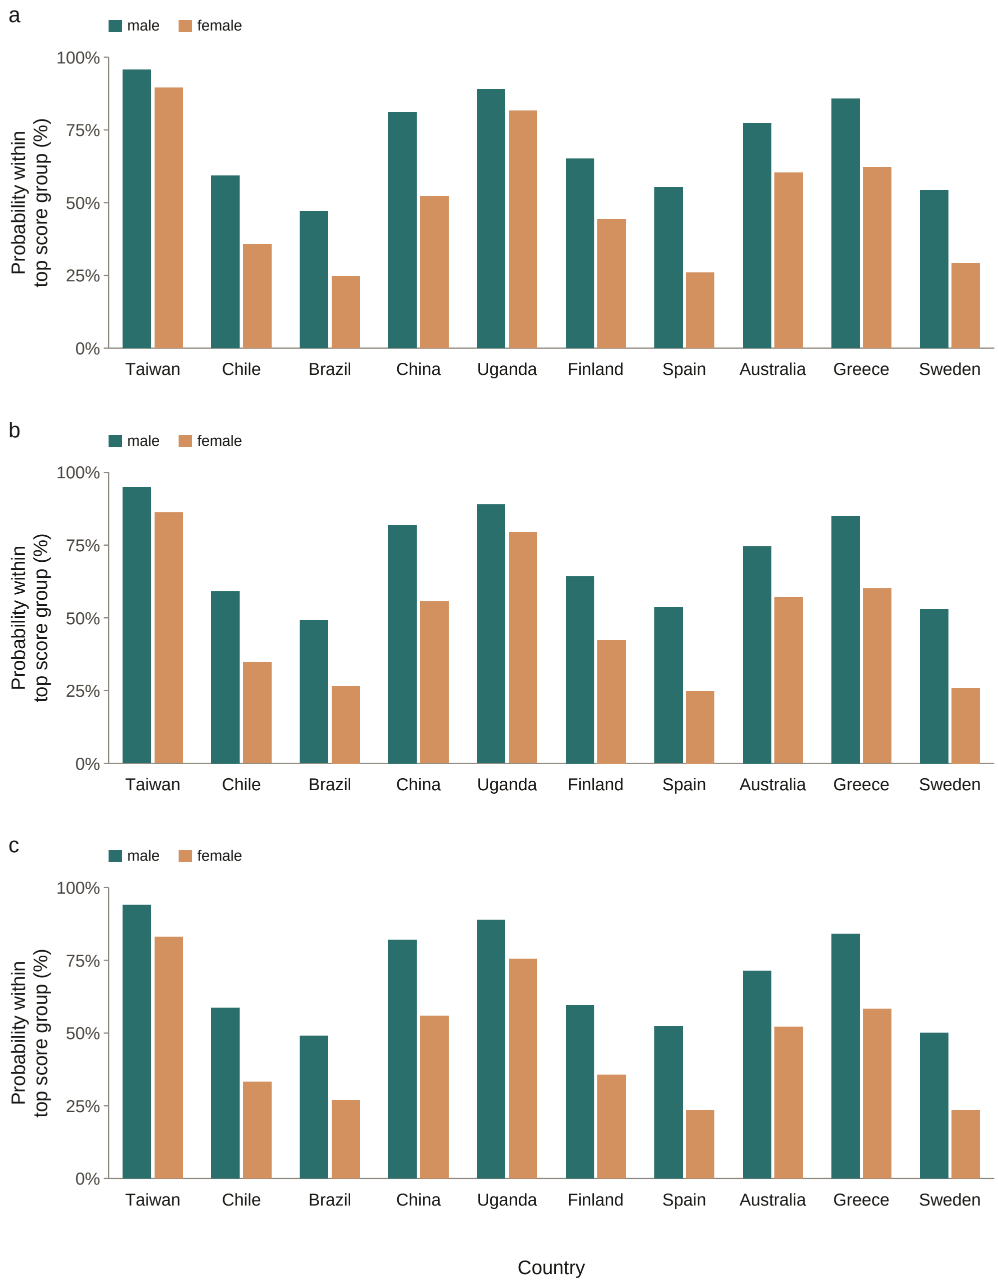

Notes: Probability of listing a STEM program anywhere, by sex, at the top 5% (a), top 10% (b) and top 20% (c). Chile conditions on applicants. The female-minus-male gaps are reported in Figure B.VIII. Samples in Online Appendix Table A.II.

Figure B.VIII: The Any-STEM Gender Gap at Alternative Performance Thresholds

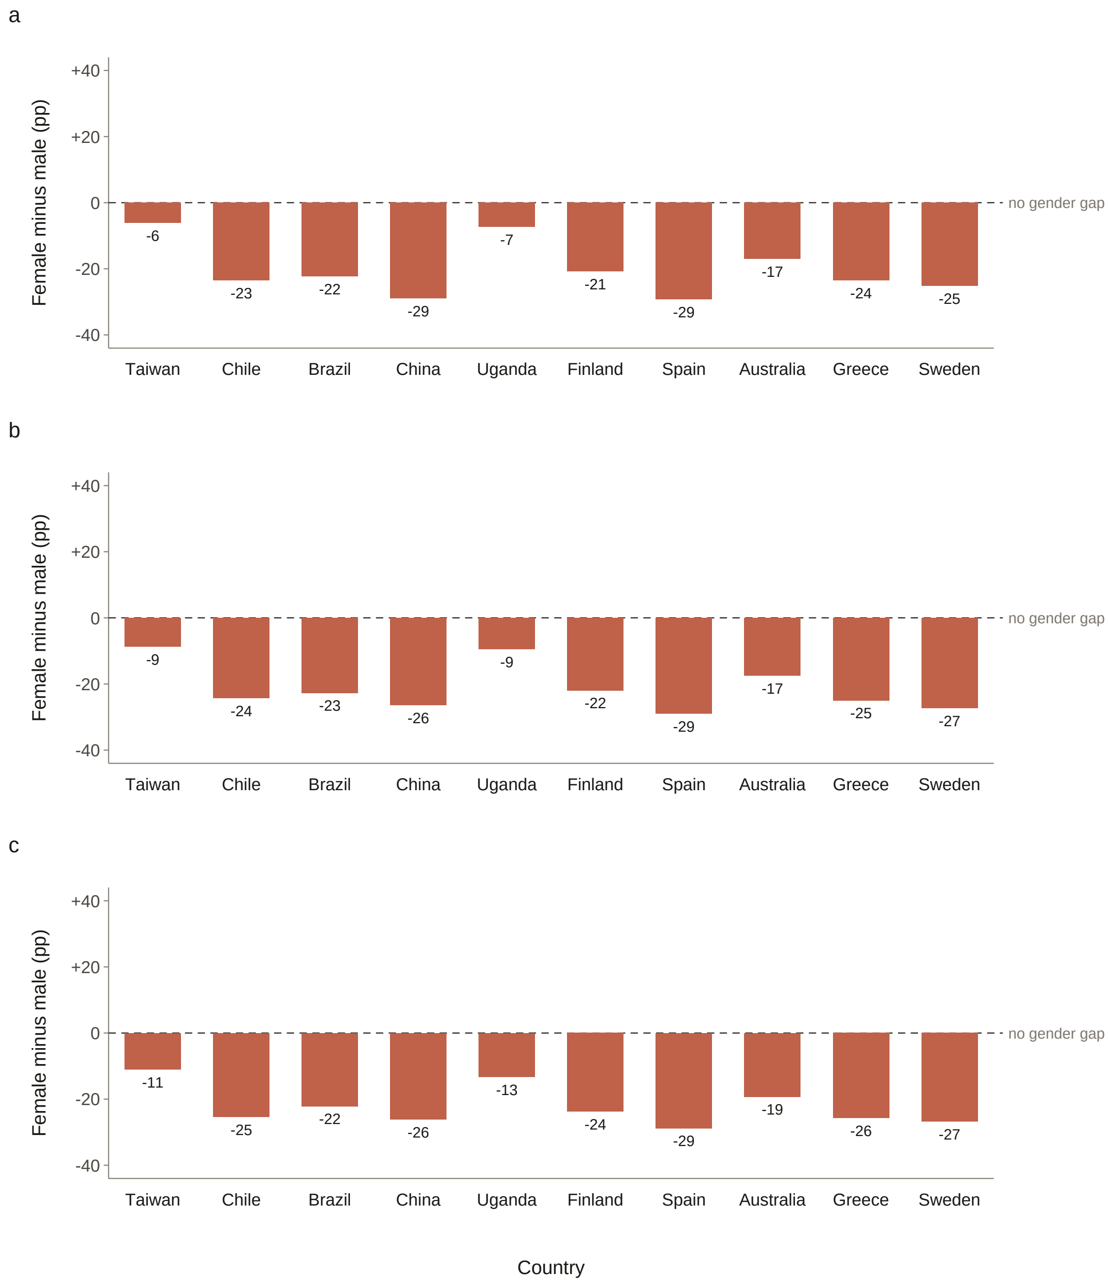

Notes: Female-minus-male gap in the probability of listing a STEM program anywhere, at the top 5% (a), top 10% (b) and top 20% (c). Chile conditions on applicants. The underlying probabilities by sex are reported in Figure B.VII. Samples in Online Appendix Table A.II.

Appendix B.7. Across the full score distribution

To extend the analysis below the top quintile, we pool deciles 1–5 into a single "bottom 50%" group and report deciles 6–10 separately. This preserves a decile-by-decile view in the upper tail, where admission to selective STEM programs is most plausible, while reducing noise in the lower half of the distribution.

The two margins behave differently across these bins, in the same way they behave differently across settings. Female shares vary in direction: in Brazil, Chile and Finland female representation tends to decline with academic performance, while in Greece and Sweden it tends to increase, consistent with the main-text evidence that pipeline differences vary by context. The choice gap does not vary in direction. The female-minus-male gap in STEM application probabilities is negative in every country and in every bin: all 120 country-by-bin cells—ten settings, six score bins, and two definitions of a STEM application—are negative. Its magnitude varies with performance in some settings, but the sign is uniform. The cross-country regularity is not that pipeline differences look similar, but that conditional STEM choice differences are persistently negative for women.[23]

Appendix B.8. From applications to placements: counterfactuals

The main text reports that the assignment component V is small and unsigned, and Figure 4 there follows high achievers through the four successive margins. Table B.II gives the setting-by-setting counterfactuals behind that statement.

The assignment margin moves the female share by less than two percentage points in seven of the ten settings and is positive in five. The largest exception is Uganda, at -5.7 percentage points; Uganda also has the lowest ratio of STEM assignments to STEM first preferences in our sample, 0.49 for men and 0.37 for women, so a larger share of first preferences goes unmet there than elsewhere. Uganda is also the setting whose admission rules differ most clearly by sex (Online Appendix A.1); elsewhere a non-zero assignment margin reflects which programs each group ranks within STEM and how far down the list they are placed, rather than differential treatment at a common cutoff.

Table B.II sets the two counterfactuals implied by (E.1) side by side. Giving high-achieving women the STEM rates of high-achieving men raises the female share among STEM entrants in all ten settings, in most cases substantially: from 42.2% to 65.9% in Sweden and from 31.5% to 49.4% in Brazil. Balancing the top decile while leaving both rates as observed raises the share in only four settings. In the other six—Australia, China, Finland, Greece, Spain and Sweden—it *lowers* it, by 14.7 percentage points in Sweden, 12.3 in Spain and 10.7 in Greece. The reason is arithmetic rather than behavioral: women already outnumber men among high achievers in those settings, and that surplus is currently offsetting part of the choice gap, so removing it removes the offset as well. Neither column predicts how a margin would respond to an intervention.

[23]The three figures behind these statements—female shares by bin, and the first-choice and any-choice gaps by bin—are reported in the project's online materials, together with the underlying counts.

Table B.II: Counterfactual female shares among high-achieving students assigned to STEM

Setting	Observed	Counterfactual		Change from
		Balanced pipeline	Common choice rates	balancing
Uganda	26.3	36.2	38.6	+9.9
Chile	27.5	32.5	44.0	+5.1
Taiwan	29.3	38.8	39.5	+9.5
Brazil	31.5	32.0	49.4	+0.5
China	33.3	28.8	55.3	-4.6
Finland	35.1	33.6	51.6	-1.5
Sweden	42.2	27.5	65.9	-14.7
Australia	42.8	39.2	53.7	-3.6
Spain	44.8	32.5	62.8	-12.3
Greece	45.1	34.4	61.1	-10.7

Notes: All entries are the percentage of high-achieving students assigned to a STEM program who are women, where high-achieving means the top decile of each setting's academic performance distribution. *Observed* is the realized share. *Balanced pipeline* holds each gender's rate of reaching a STEM assignment at its observed value and sets the top decile to equal numbers of women and men; it equals $1 - \Lambda(C + V)$ in the notation of (E.1) (Online Appendix E). *Common choice rates* holds the observed composition of the top decile and gives women men's rate, so the share equals women's share of the top decile. The final column is the balanced-pipeline share minus the observed share; a negative entry means that balancing the pipeline, on its own, would lower the female share among STEM entrants, because women already outnumber men in that setting's top decile. Both columns are accounting exercises that reallocate observed rates and are not predictions of how either margin would respond to an intervention.

Appendix C. The choice gap in other selective degrees

This section reports the same choice-gap exercises as in the main body of the paper, but for other selective fields that are also associated with high labor-market returns, namely business-law and health. Results cover students in the top 5%, top 10%, and top 20% of each country's academic performance distribution. The cross-field comparison is discussed in Section 3 and in the Results section, where Table 1 reports the decomposition components for all three fields on a common scale, setting by setting. The figures below give the underlying first-choice probabilities and gaps by country and threshold, reporting the probabilities by sex and the female-minus-male gap in turn for each field.[24]

[24] The corresponding any-choice figures for both fields are reported in the project's online materials; the any-choice pattern is summarized in the Results section, and the components in Table 1 are computed on first-choice rankings.

Appendix C.1. Business and law

Figure C.I: Business/Law First-Choice Probabilities by Sex at Alternative Performance Thresholds

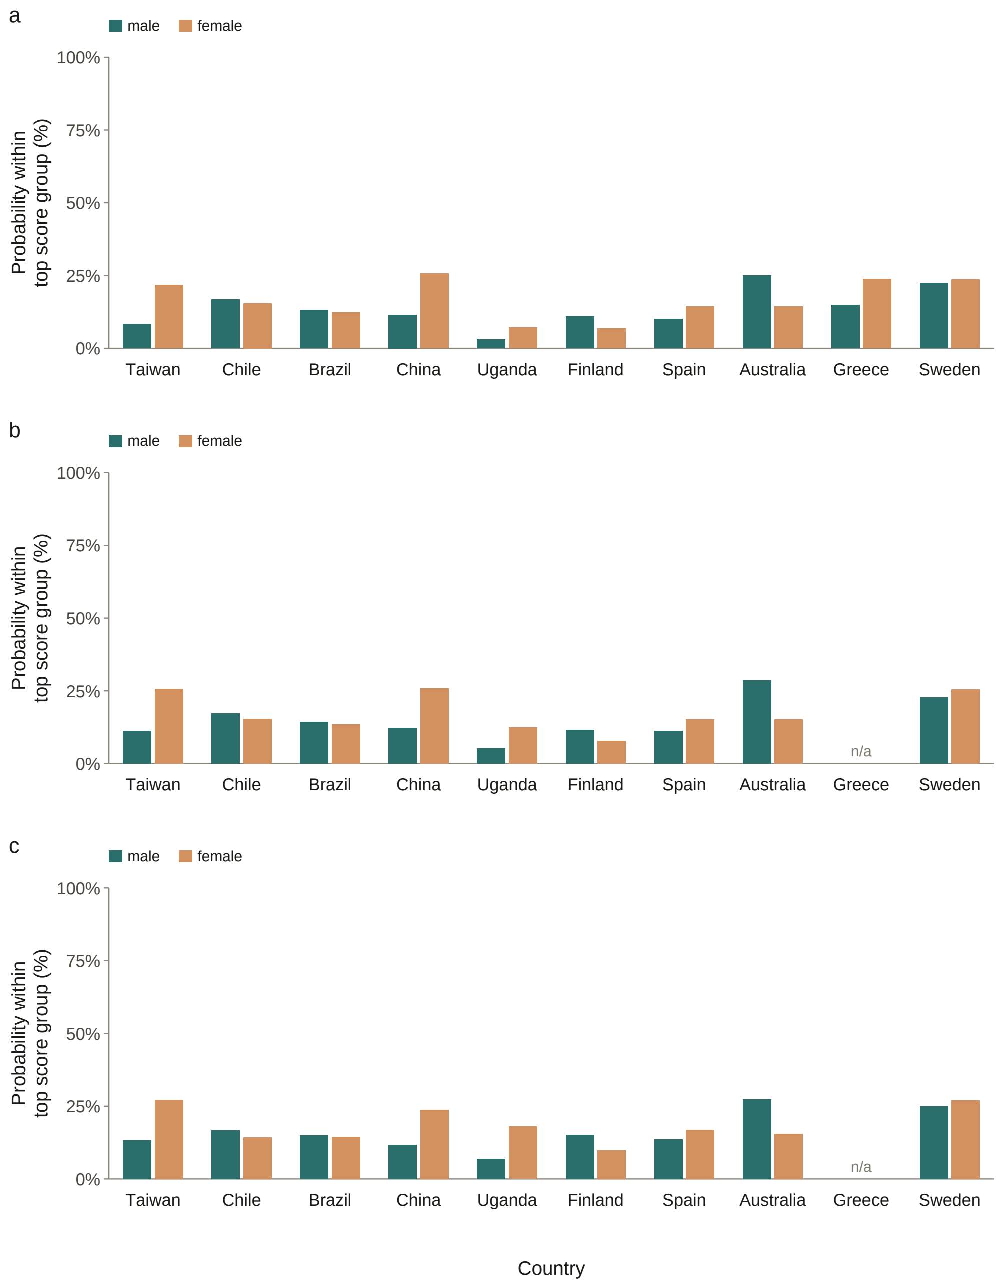

Notes: Probability of ranking a Business/Law program first, by sex, at the top 5% (a), top 10% (b) and top 20% (c). Greece is unavailable at the top 10% and top 20% and is marked n/a. The female-minus-male gaps are reported in Figure C.II. Samples in Online Appendix Table A.II.

Figure C.II: The Business/Law First-Choice Gender Gap at Alternative Performance Thresholds

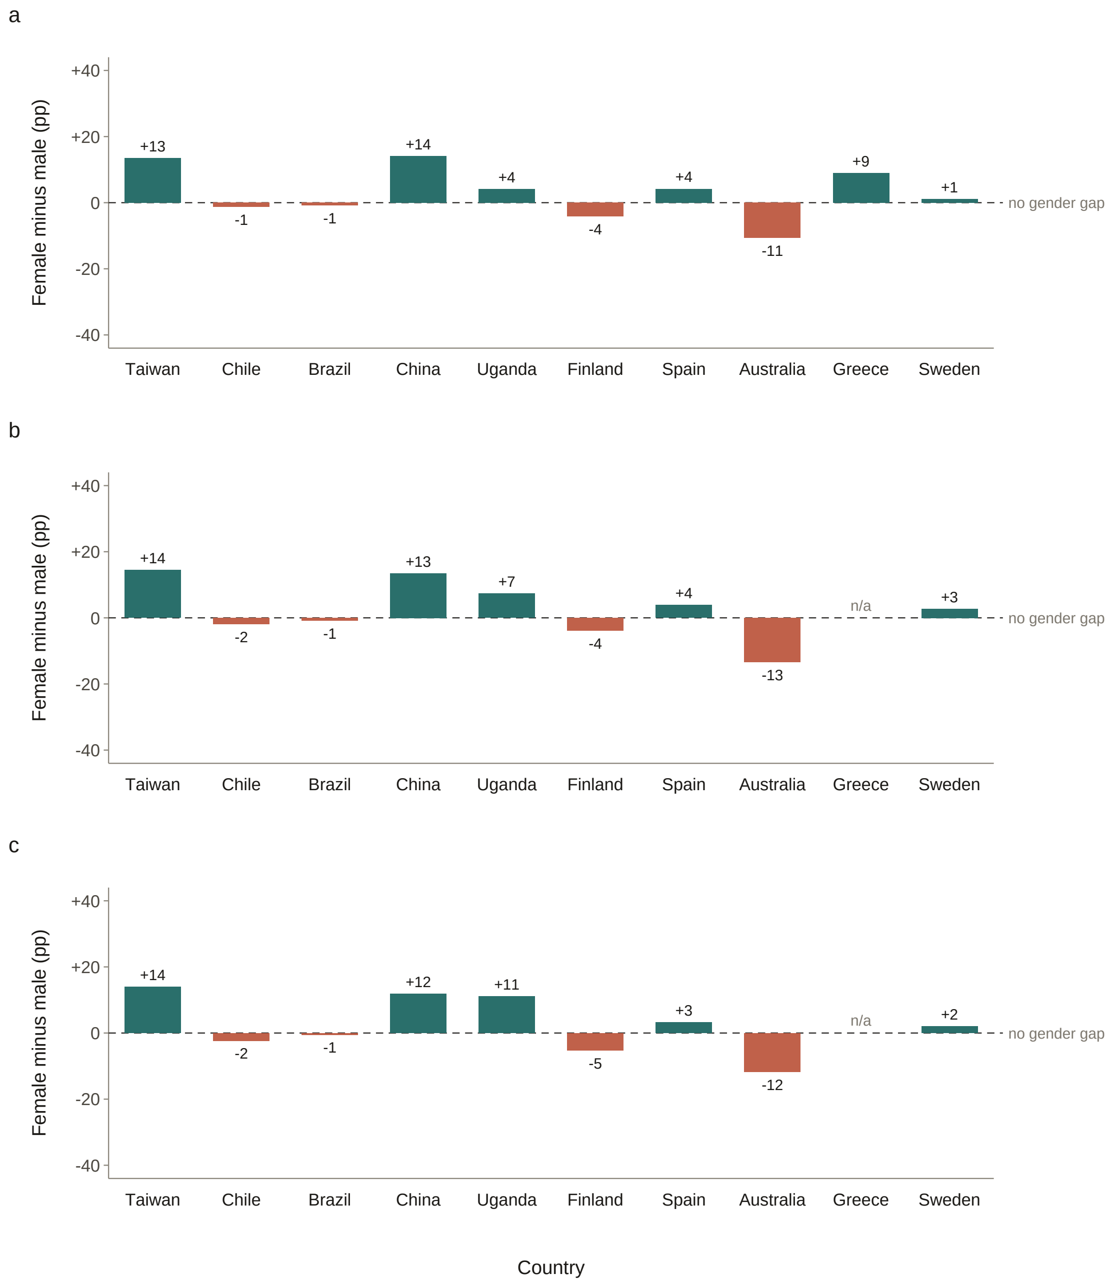

Notes: Female-minus-male gap in the probability of ranking a Business/Law program first, at the top 5% (a), top 10% (b) and top 20% (c). Greece is unavailable at the top 10% and top 20% and is marked n/a. The underlying probabilities by sex are reported in Figure C.I. Samples in Online Appendix Table A.II.

Appendix C.2. Health

Figure C.III: Health First-Choice Probabilities by Sex at Alternative Performance Thresholds

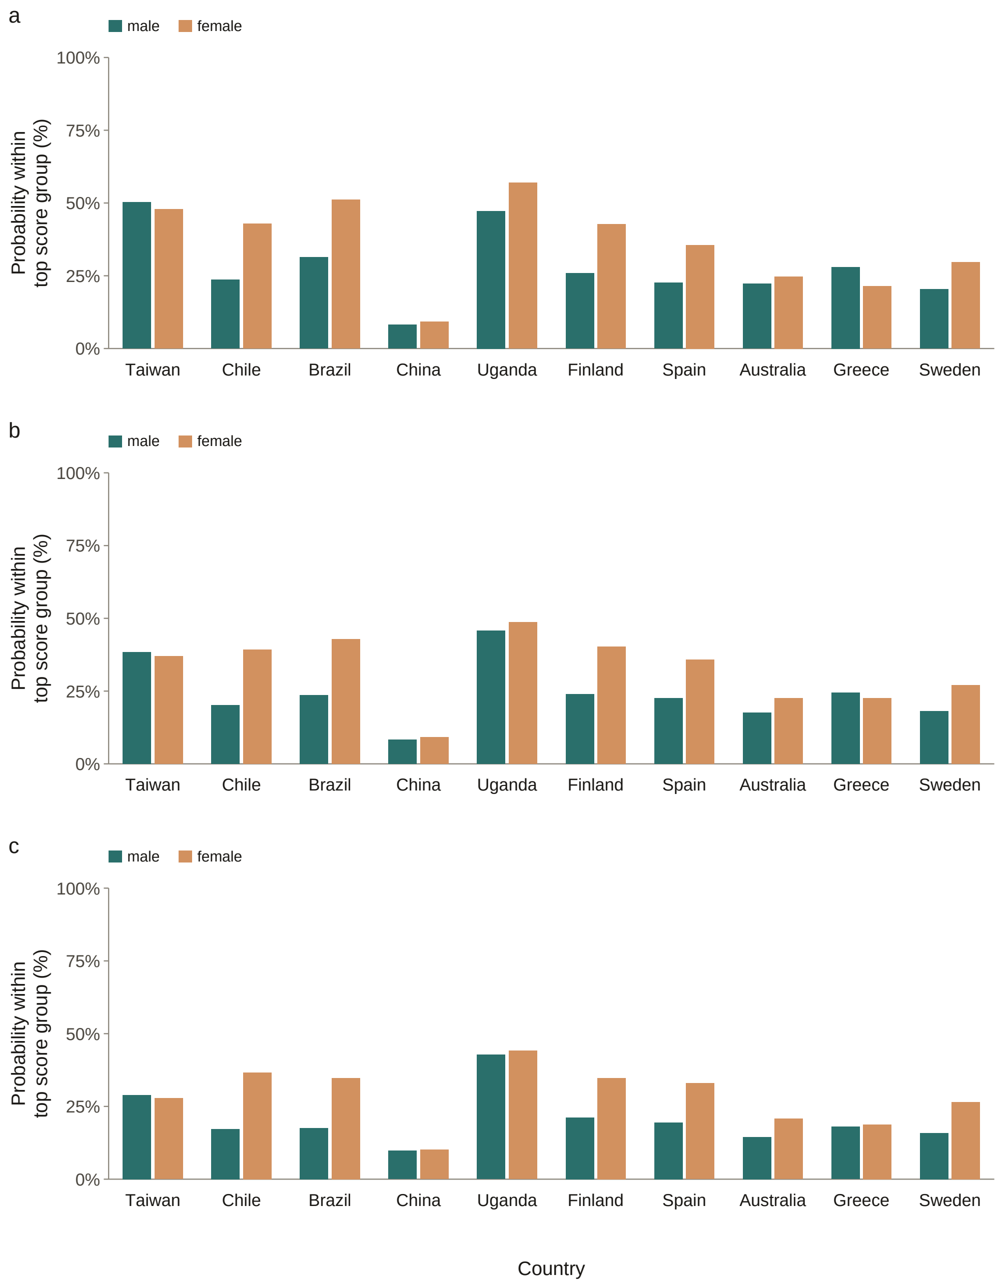

Notes: Probability of ranking a Health program first, by sex, at the top 5% (a), top 10% (b) and top 20% (c). The female-minus-male gaps are reported in Figure C.IV. Samples in Online Appendix Table A.II.

Figure C.IV: The Health First-Choice Gender Gap at Alternative Performance Thresholds

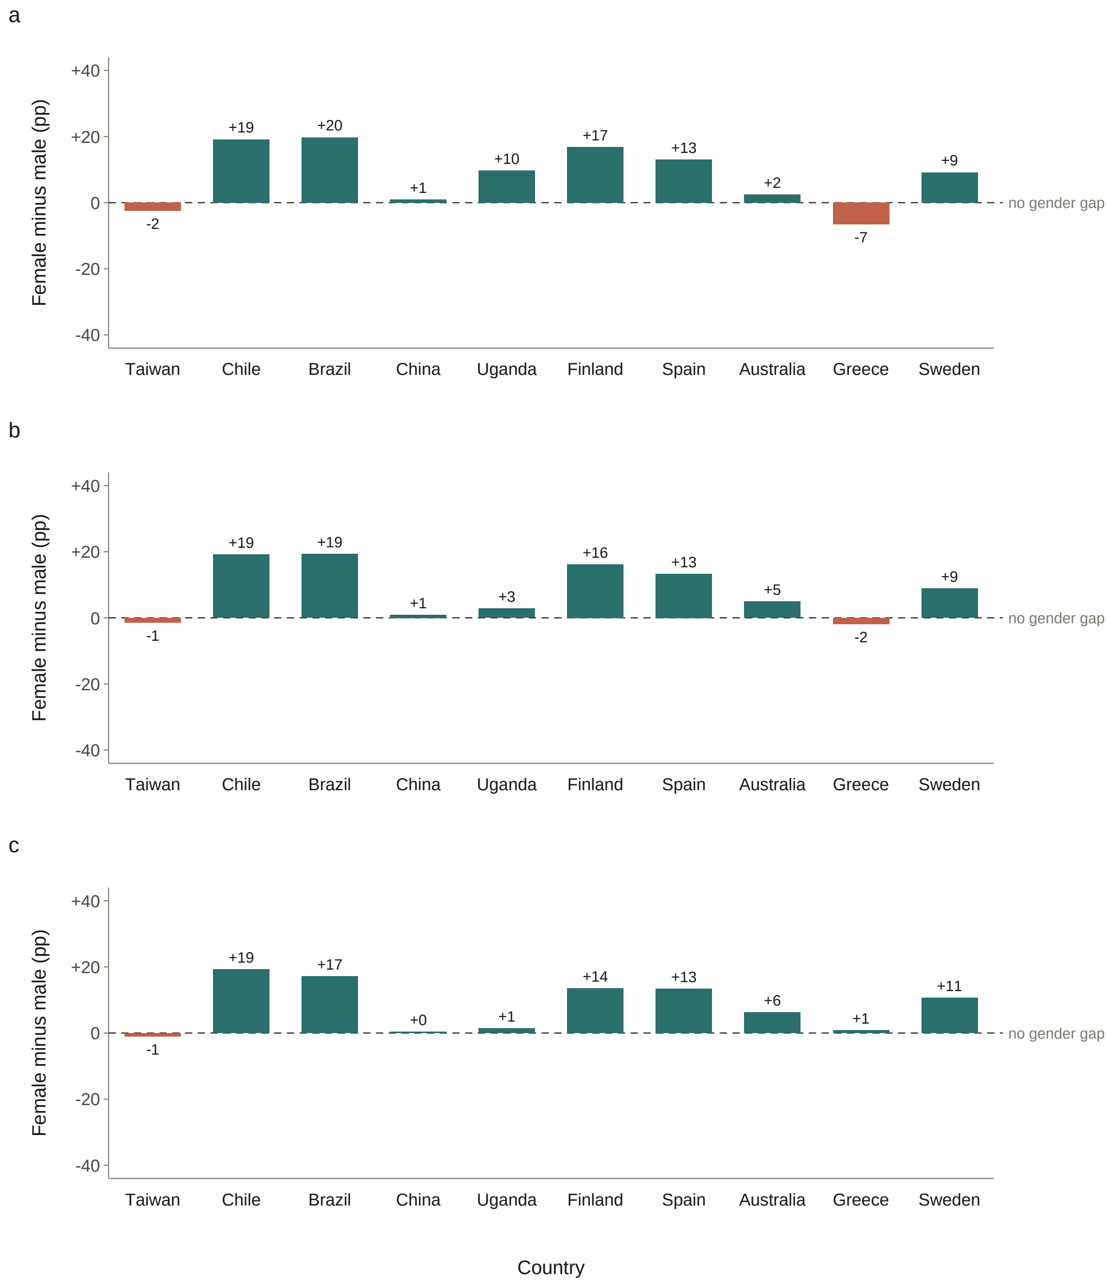

Notes: Female-minus-male gap in the probability of ranking a Health program first, at the top 5% (a), top 10% (b) and top 20% (c). The underlying probabilities by sex are reported in Figure C.III. Samples in Online Appendix Table A.II.

Appendix D. Pipeline and STEM choice gaps in cohorts closest to 2018

Appendix D.1. Sample and time coverage

This appendix evaluates whether the main findings are sensitive to cohort composition by re-estimating top-10% pipeline and choice gaps using cohorts as close as possible to 2018 in each setting. The objective is to make the time period as comparable as possible across countries, although this is not feasible everywhere given differences in the underlying source files.

The country-year selections used are: Australia (2010), Brazil (2016), Chile (2018), China (2018), Finland (2018), Greece (2003–2012 pooled), Spain (2020), Sweden (2017), Taiwan (2010), and Uganda (2021/2022).

Relative to the main-text samples, we move closer to 2018 in Australia (from 2009–2010 pooled to 2010), Chile (from 2012–2018 pooled to 2018), Finland (from 2016–2020 pooled to 2018), and Spain (from 2020–2024 pooled to 2020). In the remaining settings we use the same period as in the main text because no closer harmonized cohort is available in the current country files. In several cases this is still reasonably close to 2018: China is exactly 2018, Sweden is 2017, and Brazil is 2016.

Appendix D.2. Pipeline gap in recent cohorts

In Figure D.I, the first series shows the top-10% pipeline gap for cohorts closest to 2018. As in the main text, the pipeline gap continues to vary substantially across countries and changes sign across settings.

Appendix D.3. Choice gaps in recent cohorts

The second and third series in Figure D.I report the first-choice and any-choice STEM gaps in recent cohorts. The first-choice STEM gap remains negative and sizable in all countries, and the any-choice STEM gap is also negative in all countries, with magnitudes generally close to the first-choice benchmark. Overall, restricting to cohorts closest to 2018 leaves the core empirical pattern unchanged.

Figure D.I: Top-10% Pipeline and Choice Gaps in Cohorts Closest to 2018

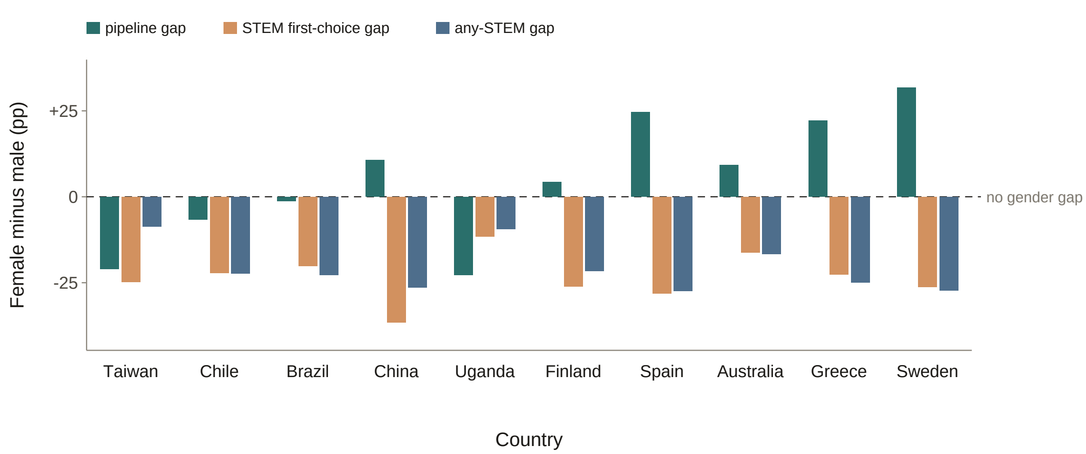

Notes: Female-minus-male gaps for the top 10%, using the cohort closest to 2018 in each setting. Online Appendix D.1 lists the country-years used.

Appendix E. The pipeline–choice decomposition

The main text states the identity in equation (1): at a fixed performance threshold the pipeline component P and the choice component C sum exactly to the gender gap among high-achieving applicants to a field. This appendix extends it to realized assignments and reports the additional component.

Extending the identity to assignment. Equation (1) stops at the application stage. Because each of our settings places students through a centralized assignment mechanism, we also observe where the same students end up, and the identity extends by one further factor. Let M_g denote the number of students of gender g in the top group assigned to the field, and let $v_g = M_g/A_g$. Then

$$\underbrace{\ln\left(\frac{M_m}{M_f}\right)}_{\text{gap among high-achieving assignees}} = P + C + \underbrace{\ln\left(\frac{v_m}{v_f}\right)}_{\text{assignment component } V}, \tag{E.1}$$

so that the male share among high-achieving students assigned to the field is $\Lambda(P+C+V)$. The third factor is a ratio of counts rather than a conditional probability: a student may be assigned to a field she did not rank first, and v_g exceeds one for at least one gender in four of our settings. V therefore combines two channels, differential admission among those who ranked the field first and differential entry from lower positions in the submitted list.

In our data V is small and carries no systematic sign. Across the ten settings it ranges from -0.17 to $+0.28$ log points, lies below 0.10 in absolute value in eight, and is smaller in magnitude than C in nine settings and comparable in the tenth (Uganda, 0.28 against 0.29); its mean is $+0.005$. The assignment component is similarly small in the other fields, ranging over $[-0.21, +0.06]$ for health and $[-0.46, +0.12]$ for business/law. Its magnitude is stable across thresholds: recomputing V on the top 20% moves no setting by more than 0.12 log points. The gap among high-achieving applicants therefore survives to assignment almost unchanged. We stress what this does and does not say. It is a statement about the aggregate composition of assigned cohorts, not about any individual admission decision: it does not establish that the mechanism treats otherwise identical men and women alike, only that the composition it produces is close to the one already implied by the lists students submit.

We report P, C and V in log points rather than as shares of the total. The components frequently carry opposite signs in our data, and ratios of the form $P/(P+C)$ are not well behaved in that case. We emphasize that (1) and (E.1) are accounting identities and carry no causal content: they partition an observed gap into measured margins, and are silent on what generates any of them.

Appendix F. From applications to placements

The main text reports that the gaps documented among applicants survive the assignment stage, and Figure 4 there gives the funnel. Online Appendix E states the extension of (1) to realized assignments and defines the assignment component V in (E.1). This section records what that component does in our data.

V is small and carries no systematic sign: $|V|$ is at most 0.10 in eight of the ten settings, smaller in magnitude than C in nine of the ten settings and comparable in the tenth (Uganda, 0.28 against 0.29), and averages +0.005. Because V is small, the counterfactual in the final column of Table 1 carries over from applicants to entrants: with a gender-balanced top group, the male share among high-achieving students assigned to STEM would lie between 0.61 and 0.73. Uganda operates admission rules that differ by sex and favor women (Online Appendix A.1); its assignment component nonetheless runs against women, which these data cannot explain and we do not try to. Elsewhere a non-zero V reflects which programs each group ranks within STEM and how far down the list they are placed, rather than differential treatment at a common cutoff. Online Appendix B.8 reports the counterfactuals implied by (E.1) setting by setting. Two qualifications apply throughout: the comparison is one of aggregate composition and does not speak to any individual admission decision, and we observe the program a student is assigned to rather than the one she enrolls in.

Appendix G. The pipeline gap, the choice gap, and gender wage parity

This section reports the cross-country evidence behind the discussion in Section 4.4 of the main text. Figure G.I relates the *pipeline* and *choice gaps* among students in the top 10% of the academic performance distribution to the wage parity index, defined as one minus the total gender wage gap, so values closer to one imply greater parity: panel (a) reports the pipeline gap and panel (b) the choice gap. Wage gaps come from the International Labour Organization (International Labour Organization, 2026), except for Australia and Taiwan, taken from official sources (Workplace Gender Equality Agency, 2026, Directorate-General of Budget, Accounting and Statistics (DGBAS), Taiwan, 2016), and China, from Zhang and Dong (2025).

Figure G.I: Pipeline and Choice Gaps vs Wage Parity Index

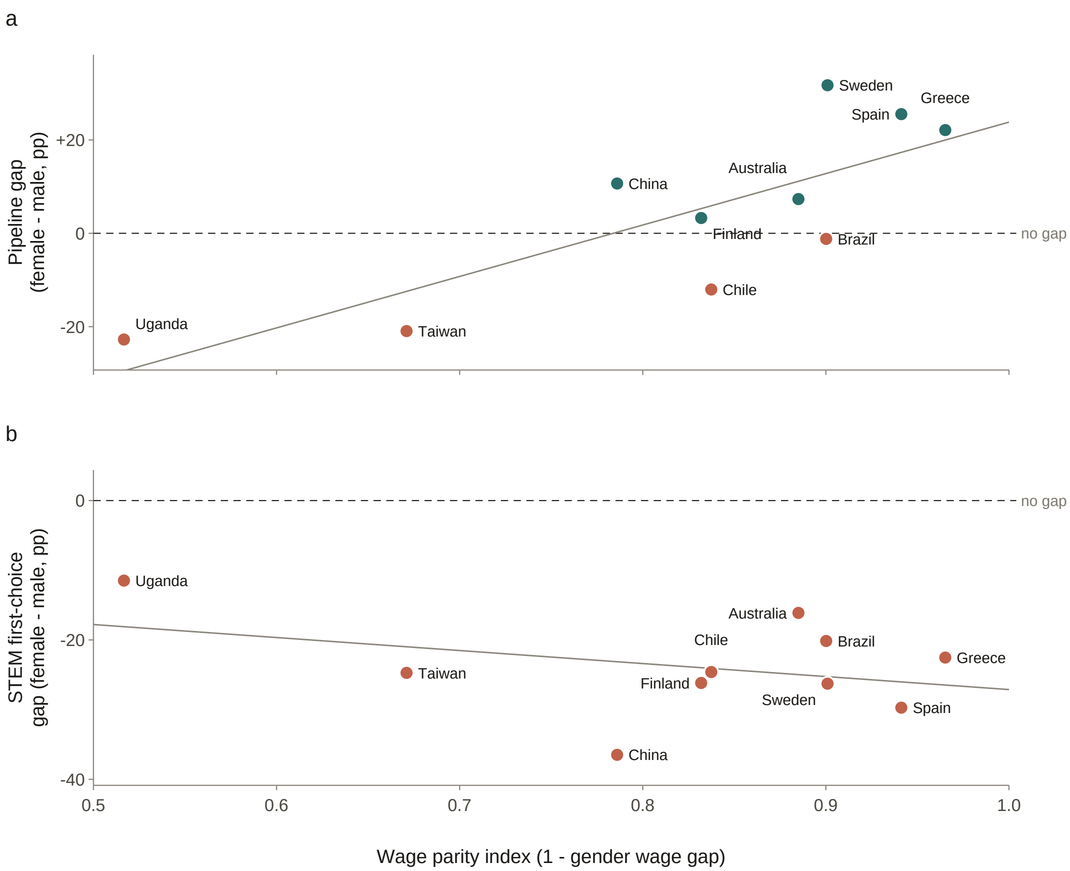

Notes: Both gaps are female minus male, percentage points, top 10% of the academic performance distribution: (a) the pipeline gap, (b) the choice gap. The line is an unweighted least-squares fit. Samples in Online Appendix Table A.II.

Appendix G.1. Sensitivity of the two slopes

Section 4.4 rests on the contrast between the two margins rather than on the precision of either slope, and the two are not equally robust. The pipeline relationship is estimated equally precisely in either scale ($|t| = 3.70$ in log points against 3.71 in percentage points). The choice relationship is not: excluding settings one at a time, the slope on C stays between $+0.58$ and $+1.22$ and never changes sign, whereas the corresponding percentage-point slope changes sign when a single setting is dropped. The precision of the slope on C, though not its sign, rests on Uganda, the lowest-parity setting in the sample; excluding Uganda leaves the sign intact but takes the statistic to $t = 0.7$. We therefore read the positive slope on C as evidence against a decline rather than as evidence of an increase, and rest the section's claim on the divergence between the margins, which survives every one of these variations.