First-Generation Students in European Higher Education: Access and Completion

What the Category Conceals and Why It Still Matters

Vikki Boliver & Rebecca Thornton

Vikki Boliver, Durham University
Rebecca Thornton, Durham University

First-generation students are the most studied under-represented group in European higher education and the least consistently defined. This article examines the definitions in use, the evidence about access and completion, and the institutional measures that make a difference. It argues that the category conceals as much as it reveals — combining students whose disadvantage is financial, cultural and informational in different proportions — and that the measures which work are those that address the informational deficit rather than the financial one alone.

Contents **Page**

1. A Category and Its Definitions **14**

2. Access **14**
2.1 What the Evidence Shows 14
2.2 The School System 14

3. Completion **15**
3.1 The Completion Gap 15
3.2 What Institutions Observe 15

4. Measures That Work **16**
4.1 Informational Measures 16
4.2 Financial Measures 16
4.3 Belonging 16

5. The Institutional Question **17**
5.1 Who Owns the Issue 17
5.2 The Data Question 17

6. Policy Implications **17**

7. A Comparable Indicator: What the Question Must Do **18**

8. Cost and the Sequence of Investment **18**

9. The Institution's Own Data **19**

10. The Student's Own Account **20**

11. What the Institution Gains **20**

11. Conclusions **21**

1. A Category and Its Definitions

One label, four definitions

The term first-generation student describes a student whose parents did not attend higher education. The description is simple and the definitions in use are not: some studies count a student as first-generation if neither parent completed a degree, others if neither attended at all, others if no member of the household has a qualification.

The differences matter for measurement. A system that counts students whose parents attended but did not complete will report a larger first-generation population than one that counts only those whose parents never entered, and the two populations have different needs.

Definition	Population size	Typical finding
Neither parent attended	Smaller	Largest attainment gap
Neither parent completed	Larger	Moderate gap
No household qualification	Variable	Depends on migration
First in family to attend	Variable	Cultural rather than financial

Table 1 Definitions in use and their effects

2. Access

2.1 What the Evidence Shows

The gap widens with selectivity

Across European systems, first-generation students are under-represented in higher education and more strongly under-represented in the most selective institutions and programmes. The pattern is consistent and its size varies with the structure of the school system and with the selectivity of the entry route.

The mechanism is informational as much as financial. A student whose parents attended university has access to informal knowledge about which programmes lead where, how an application is assessed and what an institution expects; a student without that knowledge relies on formal information, which is frequently late and generic.

2.2 The School System

The second mechanism is the structure of secondary education. Systems that separate students early into academic and vocational tracks produce larger gaps, because the decision about the track is taken at an

age when family background has its strongest influence and is difficult to reverse.

Factor	Mechanism	Effect
Early tracking	Decision at an early age	Large and persistent gap
Selective institutions	Information and preparation	Gap concentrated at the top
Fee levels	Price of participation	Gap widens where fees are high
Admission tests	Preparation and coaching	Gap widens with test weight
Late specialisation	More time to decide	Smaller gap

Table 2 Structural factors and their effect on the gap

3. Completion

3.1 The Completion Gap

The evidence on completion is weaker than the evidence on access and it points in a consistent direction: first-generation students complete at lower rates, and the difference is largest in the first year.

Preparation, money, belonging

The explanations fall into three categories. The academic explanation points to preparation: a student from a school that did not prepare her for the demands of a programme struggles in the first semester. The financial explanation points to work: a student who must earn to study has less time. The cultural explanation points to belonging: a student who does not recognise herself in the institution's assumptions about what a student is may withdraw.

The three are not alternatives and they require different responses. A programme that addresses preparation and not belonging will improve grades and lose students; one that addresses belonging and not money will retain students who cannot afford to stay.

3.2 What Institutions Observe

Institutions observe the completion gap and rarely explain it. The data available are records of enrolment, progression and award, and they do not contain the information that would identify the mechanism: hours of paid work, the reason for a withdrawal, the experience of the first semester.



Figure 1 Where the completion gap is produced

The second stage is where most of the loss occurs and where the institution has the most leverage, because the first semester is the period in which a student decides whether she belongs.

4. Measures That Work

4.1 Informational Measures

Cheapest and most effective

The measures with the strongest evidence are informational. They include outreach in schools with low participation, a named contact for applicants who have no family experience of higher education, and a pre-entry programme that introduces the conventions of academic work before the first assessment.

The third is the most effective and the most expensive, and its cost is concentrated in the period before the academic year, when institutions have capacity.

4.2 Financial Measures

Financial measures are necessary and insufficient. A grant that permits full-time study removes the need to work; a grant that does not removes nothing. The evidence suggests that the level matters more than the targeting, because a small targeted grant does not change the decision that a student makes about how many hours to work.

4.3 Belonging

The measures that address belonging are the least expensive and the least used. A cohort that begins together, a mentor who is herself a first-generation graduate, a seminar in which the conventions of the discipline are made explicit rather than assumed: each of these costs little and each addresses the mechanism that the cultural explanation identifies.

5. The Institutional Question

5.1 Who Owns the Issue

Distributed responsibility is no responsibility

The first institutional question is ownership. In most institutions responsibility for first-generation students is distributed between an outreach office, a student services unit and, informally, individual academics who take an interest. The distribution produces activity without a strategy.

The institutions that do best have assigned the issue to a named person with a budget and a report to the governing body, and the assignment is what converts a set of activities into a policy.

5.2 The Data Question

The second is data. An institution that does not know which of its students are first-generation cannot evaluate a measure, and the collection of the information raises a question of privacy that is legitimate and answerable: the data can be collected for the purpose of support, held separately from the academic record and reported only in aggregate.

6. Policy Implications

Three implications follow.

The first is that the definition should be stated and used consistently, because a category that changes its meaning between reports cannot support a policy.

The second is that access and completion should be measured together, with the difference between them treated as a measure of institutional capacity.

The third is that the informational measures should be funded before the financial ones are increased, because a student who does not apply cannot be supported by a grant.

7. A Comparable Indicator: What the Question Must Do

The category is defined by a question, and the question is rarely designed with the comparison in mind. A survey that asks whether a parent attended university produces a different population from one that asks whether a parent holds a degree, and the two populations differ not only in size but in the needs that policy is meant to address.

A definition is a question with a reference period

Four properties determine whether an indicator travels. The first is the object of the question: attendance, completion or qualification at a given level. The second is the household: one parent or both, and whether a step-parent or a guardian counts. The third is the reference period: the situation when the student was fourteen, which is the convention in several surveys, or the situation at the time of the interview, which is easier to answer and less stable.

The fourth is the treatment of non-response. A student who does not know her parents' education is not a missing value to be discarded; she is frequently a student whose parents are absent, deceased or educated abroad, and the three cases have nothing in common except the difficulty of the answer.

The last choice is the one that institutions make least deliberately. A gap measured against the student body and a gap measured against the children of graduates are different quantities, and a report that does not state which of them it uses allows a system to select the smaller one when the news is bad. The published indicator should name its baseline in the same sentence as its value.

8. Cost and the Sequence of Investment

The measures that institutions describe as widening participation differ by an order of magnitude in cost, and the difference is not reflected in the order in which they are adopted. Information and advice are cheap. Mentoring and pre-entry programmes are moderate. Grants, fee reductions and residential accommodation are expensive. Institutions adopt them in the reverse order.

The reason is visibility rather than ignorance. A scholarship is announced at a ceremony and appears in a press release; a change in the advice given to applicants by a school liaison officer appears nowhere, and its effect is a difference in applications that nobody observes. An institution that is accountable for what it spends and not for what it achieves will spend on the visible instrument.

The cheapest measure is the one that changes an application

The sequence that the evidence supports is therefore not the sequence that institutions follow, and the divergence has a cost that is rarely counted: a scholarship programme that consumes the whole budget produces a smaller number of graduates than a combination of advice and preparation would produce with the same money.

There is a case for the expensive instruments as well, and it is an argument about adequacy rather than about access. A student who has been admitted and cannot afford to remain enrolled is not helped by better advice at the point of application, and a system that spends everything on outreach and nothing on maintenance has moved the loss from the first year to the second. The correct statement is that the measures are sequential: information determines who applies, preparation determines who is admitted, and support determines who completes.

9. The Institution's Own Data

A national indicator describes a system and an institution needs a number it can act upon, which is why the first measure in most universities is the construction of a local definition. The definition is usually assembled in the admissions office, used for the annual report and never checked against the register that records outcomes.

Three sources exist and they are rarely joined. The application form records what the applicant states about her parents. The registry records the programme, the results and the date of withdrawal. The student survey records the experience, the hours of work and the money. A report that draws on one of the three describes a population; a report that joins them describes a problem.

The join is technically simple and administratively delicate, because it combines an application record, which the applicant supplied for one purpose, with a performance record held for another. The delicate part is not the computing; it is the statement that the institution makes to the applicant about what it will do with the answer.

A target on access can be met at the gate

A local target is the second measure and it is the one that determines whether the data are used. A target expressed as a share of the student body is easy to meet by admitting more students and hard to meet by retaining them; a target expressed as the difference in completion between groups cannot be met by admission at all, and it directs the institution to the only measure that matters.

The third measure is publication. An institution that publishes the difference in completion between first-generation students and the rest has made a commitment that its own staff can see, and the visibility is

the mechanism. The publication need not be a league table and it should not be a comparison between institutions, which measures the intake rather than the teaching.

The argument for local data is not that national indicators are wrong. It is that a national indicator tells a university that the problem exists in its country and a local indicator tells it where in the institution the problem is produced, which is the only form of the information that a programme director can use.

10. The Student's Own Account

The literature on first-generation students is written mostly from institutional data, and the accounts that students give of their own experience describe a mechanism the data cannot show: the cost of not knowing how the institution works.

A student whose parents attended university arrives with a set of assumptions that she has never had to articulate. She knows that a lecturer has office hours and that a question is welcome in them; that a reading list is a starting point; that a low grade in the first term is not a verdict; that a tutorial is a place to speak badly and be corrected. None of this is taught, and all of it is assumed.

The hidden curriculum is the difference, and it is teachable

A student without that inheritance learns the same things, and she learns them by making mistakes that are recorded. The recorded mistakes are what the statistics describe as a completion gap, and the description is accurate and it locates the cause in the wrong place: the student did not lack ability, she lacked an instruction manual that nobody had written.

Three measures address it and none of them is expensive. A first-term module that explains how the institution works, taught by a member of staff and assessed so that it is taken seriously. A named adviser for the first year whose role is to answer the questions a student is embarrassed to ask. And a published account, on the programme page, of what a week looks like, what a tutorial is, and what happens when an essay is late.

11. What the Institution Gains

The argument for supporting first-generation students is usually made in the language of equity, and there is a second argument that is rarely

made and that is more persuasive to a governing board: the institution that teaches students who do not arrive prepared learns how to teach.

A programme that has never had to explain its own conventions has never examined them, and the examination is where the improvement in teaching comes from. The module that explains how a discipline argues, the tutorial that teaches a student to speak, the feedback that states what would improve the work: these are the practices that a diverse intake forces an institution to develop, and they improve the education of every student who arrives.

The evidence for this is in the institutions that changed their intake and their practice together. They report that the changes designed for the new students were adopted by the staff as improvements rather than as accommodations, and that the strongest resistance came from those who assumed the old intake would return. It did not, and the practice that was designed for the students who arrived is now the practice of the department.

11. Conclusions

First-generation status is a real disadvantage and a heterogeneous category. It combines a financial deficit, an informational deficit and a cultural distance in proportions that differ between students and between systems, and a measure that addresses one of the three will produce a partial result that is frequently reported as a failure.

Three conclusions follow. The informational measures are the cheapest and the most effective and they should come first. Support must be adequate rather than targeted at a symbolic level. And institutions must own the issue with a name, a budget and a report, because a responsibility distributed between three offices is a responsibility that nobody holds.

References

[1] Boliver, V. (2017). Widening participation and fair access. Oxford Review of Education, 43(3), 331–347.

[2] Thornton, R. (2019). First-generation students and completion. Durham: Durham University.

[3] Spiegler, T., & Bednarek, A. (2013). First-generation students in Europe. Brussels: European University Association.

[4] Orr, D., & Hovdhaugen, E. (2014). Completion and dropout in Europe. Oslo: NIFU.

[5] Tinto, V. (2017). Reflections on student retention. In The student engagement handbook. Bingley: Emerald.

Biography:

Vikki Boliver is professor of sociology at Durham University, where she works on social inequalities in access to higher education and on fair admissions.

Rebecca Thornton is a researcher at Durham University working on first-generation students, completion and the evaluation of widening participation measures.