

POLICY ANALYSIS · SHARED BY THE ICA

# Inside 8,719 Public Comments on Do No Harm

*An AI-enabled read of the AHEAD comment record — who filed, what they said, and which arguments will actually shape the final rule*

BY PHIL HILL · ON EDTECH BY PH&A · MAY 27, 2026

*Disclosure: I filed a public comment on this rule, ED-2026-OPE-0100-4904, described in detail in “The Most Arcane Public Comment Imaginable.” I have also assisted a small number of organizations with their filings. This post is a broad analysis of the full 8,719-comment record available as of yesterday.*

The Department of Education (ED) published its proposed earnings accountability rule — the AHEAD / Do No Harm framework that ties continued Title IV eligibility to graduate earnings — and drew 10,058 public comments, of which 8,719 are available for download (the rest have not been approved yet). The idea is that ED will review these comments and make adjustments before filing the final rule, due before July 1.

Extending my work from September and December 2025, I thought it would be useful to share an AI-enabled analysis of the public comments given how significant the topic is for US higher education. See the bottom of this post for a description of the methodology used for this analysis.

The top line results of the current accountability comments are lopsided: 83% of comments reject the rule outright, defending a specific sector (cosmetology, massage, allied health, etc.) without engaging with whether earnings accountability itself is legitimate. Another 11% accept accountability but want the rule refined — these are the comments worth reading carefully, because they often identify real implementation problems. Only 1.5% are purely pro-rule.

It is worth remembering that by far the biggest impact of the new rules will be on workforce-oriented undergraduate certificate and associate degree granting programs. Using ED’s data, 92.5% of

Cosmetology programs will likely fail, as will 28% of all undergraduate certificate programs.

## Regulations Follow Statute

---

A substantial portion of the accountability framework commenters are attacking, however, is no longer regulatory — it's statutory. Congress through OBBB wrote core elements of the earnings-accountability structure directly into law, which means public comments arguing that the framework itself is illegitimate are, practically speaking, irrelevant to the final rule. The Department cannot withdraw a statutory requirement in response to public pressure; only Congress can. Litigation challenging statutory provisions faces an even higher bar: courts are far more reluctant to overturn explicit congressional language than agency regulatory choices, and *Loper Bright's* 'best reading' doctrine cuts both directions: courts read unambiguous statutory text as Congress wrote it, not as industry wishes it had been written.

The comments that actually matter to the shape of the final rule are those addressing ED's discretionary choices within the statutory framework. How did the Department define comparison groups? Should undergraduate certificates be carved out, and on what statutory basis? Where did ED set thresholds, and are those defensible given the underlying earnings data? What transition timelines and safe harbors are appropriate? These are the questions where ED retains genuine latitude and where a compelling comment can move the needle.

## Who Filed the Comments

---

Career-school operators account for 35% of the comments outright (3,072 comments), and that number understates the industry's footprint. When you add school-adjacent trade associations, accreditors, and business service providers, roughly half of all comments came from the regulated industry itself or its direct representatives. The student-and-citizen voice — the people the accountability rule is designed to protect — represents about 25% of the docket (2,140 comments).

That 25% figure requires one important adjustment. The Regulations.gov category dropdown is self-selected, and a significant number of commenters who labeled themselves "Student," "Financial Aid Administrator," or "Individual" were in fact school owners, instructors, or industry practitioners writing in personal voice. After applying a pattern-matching override to catch explicit self-identifications ("as a salon owner," "our school," "I've operated this program for 20 years"), roughly 50 comments moved from the citizen bucket to the career-school operator bucket. The authentic student-and-citizen voice is real and substantial — but it needed to be cleaned before it could be trusted.

What’s left after the override is the authentic citizen voice in the record — students, financial aid administrators at traditional institutions, teachers, parents, and individuals writing personal testimony about why these programs matter. The right question to ask of any 2,140-comment bucket is whether those voices are independent or echoing industry talking points, and we tracked it directly: only 1.8% of citizen comments use two or more pieces of coordinated policy jargon, compared to 6.7% on the operator side. The citizen voice is smaller than the raw count suggests once you subtract the industry insiders — but what’s left is mostly genuine.

The remaining quarter of the docket distributes across generic IHEs (8.4%), think tanks, advocacy organizations and trade groups (4.4%), business and service providers (4.1%), federal, state, and local government officials (3.9%), accreditors (1.4%), and a residual 18% of “Other / Unspecified” comments where the filer either left the category blank or chose a generic catch-all. This last bucket is where most of the unaligned individual professional voices live — the working cosmetologist who didn’t pick a dropdown but whose comment is substantive — and it deserves its own pass when interpreting the record. It is not noise. But it is also not a unified voice.

## Topic Breakdown of Comments

---

The next view is a full distribution of the 20 topic labels assigned by the classifier across all 8,712 successfully classified comments (seven could not be classified in any meaningful sense). Each comment carries one to three labels, so the percentages sum to well over 100 — a comment defending cosmetology while also criticizing the early-career earnings snapshot counts in both rows. Labels are color-coded by category. Two labels dominate everything else. Protect Cosmetology / Beauty / Barber appears in 69% of comments (6,029), and Protect Non-traditional / Second-chance Access appears in 50% (4,366). Behind those, the methodology critiques cluster: Question Early-Career Snapshot at 23%, Question Comparison-Group Construction at 13%, Question Earnings Data Sources at 10%, and Question Threshold Calibration at 6%. Support but Refine sits at 12%. Every other label — including all seven non-cosmetology Protect topics, both scope critiques, and both pro-rule labels — comes in under 8%.

Two findings matter. First, the Protect Cosmetology and Protect Non-traditional Access labels overlap heavily; the cosmetology industry has explicitly adopted the “non-traditional student access” frame as its second line of defense, so much of that 50% bar is the same constituency on a different theme. Second, the methodology critique layer — 38% of comments use at least one such label — is meaningfully smaller than the headline opposition, and it is heavily skewed toward the Early-Career Snapshot critique alone. That’s because the snapshot argument is the one technical critique most career school operators can repeat in their own words; the deeper critiques (comparison group construction, threshold calibration, data source validity) remain concentrated in a smaller set of operator and trade-group filings.

## The Statutory Authority Cluster

---

Buried in the docket is a small cluster of legally substantive filings that will likely matter far more than the 6,000 stakeholder-defense comments combined. These are the comments ED’s lawyers will be reading carefully — and the ones most likely to surface as exhibits in the inevitable court challenge to the final rule. Four filings stand out, all converging on a single argument: ED lacks statutory authority to extend earnings-premium accountability to undergraduate certificate programs. They reach that conclusion through different doors.

McGuireWoods filed a roughly 312-page brief on behalf of an industry coalition arguing, on Loper Bright grounds, that the “best reading” of the Higher Education Act does not authorize earnings-premium accountability for gainful employment programs. The American Association of Cosmetology Schools also filed a parallel best-reading argument. Both lean heavily on pre-2023 GE-rule case law, and neither meaningfully engages with the 2023 rule or the court losses industry has taken since. The McGuireWoods filing in particular reads as a kitchen-sink legal declaration — the strongest arguments are diluted by everything else.

The Aveda Institutes Cooperative filing makes a different and arguably cleaner argument. Rather than litigating the underlying Higher Education Act, it treats OBBB itself as the limiting statute: Congress’s recent codification authorized earnings-premium accountability for degree and graduate programs, and the statutory text does not extend that regime to gainful-employment programs. Pushing earnings-premium accountability onto GE programs by regulation therefore exceeds ED’s statutory authority. The Association of Chiropractic Colleges filed a healthcare-sector variation in the same direction. Worth reading alongside all four: National Student Legal Defense Network filed a preemptive comment defending ED’s legal authority and citing its own 2020 paper that the Department relied on for that authority — a brief addressed less to the public-comment record and more to the future district-court judge who will see this case.

## Next Steps

---

It is now up to the Department to review all the comments and to address specific points as necessary and to adjust regulations as they see fit. I would expect the final rule to come out in mid June. But this is the US, so there will be a lot of legal activity following the final rule. But more importantly, the vast majority of the new accountability rule will be in place in July, and nearly all colleges and universities in the US will have to live with the consequences (pending court cases).

◆ **Summer / fall 2026** — Institutions report data to ED

- ◆ **February 2027** — ED releases the new program data, notably at a different level of detail (CIP6 vs. CIP4) than have ever been released
- ◆ **July 2027** — The first time that an academic program could pass or fail the Earnings Premium accountability test
- ◆ **July 2028** — The first time that an academic program could lose access to federal financial aid due to failure of the Earnings Premium test

## Appendix: Methodology

---

For those wanting to understand the methodology used to create this analysis — either for checking assumptions behind the numbers or for duplicating parts of the process — see the following steps. This was an iterative process taking roughly seven hours of work. And in case you’re wondering, no, this level of analysis would not have been possible before December 2025 and the introduction of the current paradigm of agentic AI tools.

- ◆ **Step 1 — Ingestion.** Pulled the full bulk-download CSV of 8,719 public comments from Regulations.gov. Wrote a Python script (Grab-DNH-PDF.py) to download every PDF attachment linked in the Attachment Files column, extract text with pdfminer.six, and embed the extracted text directly into the comment CSV as a new pdf\_text column. Result: a single analysis-ready file where each comment row carries both the inline comment field and the full text of any attached PDF — the input format an LLM can read end-to-end without chasing links.
- ◆ **Step 2 — Topic classification.** Developed a 20-label taxonomy inductively from exploratory keyword analysis and sample reading: nine “Protect ...” stakeholder-defense topics (cosmetology, massage, allied health, skilled trades, religious/seminary, graduate professional, etc.), four methodology critiques (earnings data sources, early-career snapshot vs. lifetime, threshold calibration, comparison-group construction), two scope critiques (undergrad certificate inclusion, graduate program inclusion), two pro-rule (Support As-Is, Attack For-Profit), one constructive critique (Refine the Rule), one procedural/legal-authority, and one generic/see-attached. Classified all 8,719 comments via a structured classifier returning up to three labels per comment plus a short rationale. The model saw both the inline comment and the first 12,000 characters of PDF text (covers 88% of the average PDF). Ran 1,000 comments via streaming for a sanity-check, then 7,719 via batch for the bulk. 8,712 of 8,719 classified successfully; 7 failed on insufficient signal. A targeted re-pass tightened the “See Attached” bucket — 363 sole-See-Attached comments were re-classified with stricter instructions, and 199 moved to substantive labels.
- ◆ **Step 3 — Stance bucketing.** Collapsed the 20 multi-label topics into mutually exclusive stance buckets: any Pro-rule label = pro-rule (1.5%); any Refine label = constructive critique (11.2%, takes precedence over Protect because endorsing accountability-with-refinement is categorically different from opposing accountability outright); any Protect without Refine = anti-rule (82.9%); Question topics without any of the above = methodology/scope critique only; the rest are procedural or generic catch-alls. The hierarchy matters: a commenter who defends cosmetology but explicitly wants the rule refined is counted as constructive, not anti-rule.
- ◆ **Step 4 — Who filed: supercategory mapping.** The Regulations.gov Category field has 54 self-selected values and well-known contamination (FA Administrator and Academic/Think Tank dropdowns are heavily misused). Collapsed those 54 values into 8 analytically meaningful supercategories: Career-school operator, Student-side / Citizen, Other / Unspecified, Generic IHE, Think tank / advocacy / trade group, Business / Service Provider, Government / Official, Accreditor. Then applied two override mechanisms: organization-name pattern matching

(regex against ~50 known think tanks, trade associations, accreditors, and major industry filers — 34 comments moved into Think tank / advocacy / trade group); and industry-insider self-identification (regex against comment text for explicit self-identifications, with a student-self-ID exclusion to prevent false positives — roughly 50 student-side commenters moved into Career-school operator; spot-check precision approximately 85%).

- ◆ **Step 5 — Coordination signal: talking-point propagation.** Identified nine policy-jargon terms that signal template-driven advocacy: Do No Harm, Earnings Premium, the 92.5% / 93% / over-90% failure statistics, OBBB, Gainful Employment, AHEAD framework, and RIA / arbitrary-and-capricious language. Counted the number of distinct terms appearing in each comment (comment + PDF text). A comment using two or more of these terms is almost certainly working from a shared template or briefing document rather than writing independently. This metric is the cleanest way to distinguish authentic citizen testimony from astroturf at scale.
- ◆ **Reproducibility.** Every step lives in version-controlled scripts. Re-running the full pipeline on a refreshed comment file (or extending the taxonomy) takes about 10 minutes of compute and roughly \$8 of API spend.