

Response to “[An Improved Grade Cap Amendment](#)”

Alisha Holland, Gates Professor of Developing Societies
on behalf of the Subcommittee on Grading

April 22, 2026

I wanted to take a stab at explaining the thinking beyond the 20%+4 because I think it's actually a really elegant design that has some under-appreciated features.

First, we valued the fact that the + 4 handles the class-size problem without forcing us into crude categories. Princeton drew a hard line between lectures (35% cap on all A-range grades) and seminars (50% cap), based on the reasonable idea that smaller courses tend to attract more advanced students and allow for more individualized feedback. But that binary distinction creates awkward boundary cases — is a 25-person upper-level physics course a lecture or a seminar? Our formula handles this continuously. The +4 is essentially a constant term that has a larger proportional effect in smaller courses, which is exactly the behavior you want. Rather than a step function with an arbitrary cutoff, you get a smooth curve that produces sensible results at every class size.

Second, the +4 accounts for the reality that smaller courses have higher variance. In a class of 12, the difference between 2 and 4 outstanding students is huge in percentage terms but entirely plausible as normal year-to-year fluctuation. A flat percentage cap would force faculty in small courses to make agonizing distinctions driven more by noise than by meaningful differences in performance. The +4 provides a buffer that matters most precisely where the variance is highest.

Third, and relatedly, it separates the aspirational target from the hard constraint in a way that helps with implementation. We want faculty to internalize 20% as their benchmark. But a cap set precisely at 20% would create constant friction at the margins — the student at 89.7 versus the one at 89.4, where the distinction is being driven by an administrative rule rather than a genuine judgment about performance. The +4 gives faculty room to exercise judgment in those cases, which makes the policy feel workable rather than punitive. That matters a great deal for long-term compliance. A rule that faculty experience as rigid and arbitrary will face constant pressure for exceptions, which ultimately undermines the whole system.

Fourth—and I'll admit this is partly a matter of optics, but I think it's extremely important optics—the formula is good for our students. We were very attentive to the concern that Harvard students not be disadvantaged relative to peers at institutions without caps. The beauty of the 20%+4 formulation is that

externally, it reflects our shared commitment to rigor: a cap of 20%, down from roughly 60%. That's a striking headline number. But the +4 meaningfully softens the actual impact, particularly in the smaller courses where many of our students do their most advanced work. So our students get the reputational benefit of being graded under a system that signals high standards, without the cap being so severe that it puts them at a material disadvantage in graduate admissions or on the job market. I think this is an underappreciated feature of the design, which in my mind makes the 20%+4 formula far better for students overall than a flat 30%.

Finally, the +4 is doing something principled: it's acknowledging that as courses get smaller and more self-selected, the within-course distribution becomes a less reliable proxy for the broader standard. In a 200-person introductory course, the enrolled students approximate the university-wide population reasonably well, so a percentage cap works fine as a proxy. In a 10-person advanced seminar, the enrollment is heavily filtered, and forcing a strict percentage would confuse selectivity of enrollment with selectivity of grading.

You could imagine more sophisticated approaches — formulas that dynamically adjust the cap based on course level, prerequisites, historical enrollment patterns, or some measure of demonstrated prior student ability. These might do a better job of calibrating "extraordinary" to the right reference group in each individual course.

But there's a real tradeoff there. Grades serve two audiences. Internally, they communicate performance to students and faculty. Externally, they communicate something about a student's achievement to graduate admissions committees, employers, and the rest of the world. A more complex formula might improve internal fairness, but it would come at a serious cost to external legibility. If we can't explain our grading standard in a sentence — 20% plus 4 — then we've undermined one of the most important things grades do, which is signal clearly to people who have no access to the underlying mechanics. A system that's fairer in theory but opaque in practice doesn't actually serve our students well.