An Improved Grade Cap Amendment

Matthew D. Schwartz
Department of Physics, Harvard University

April 22, 2026

Abstract

Grade inflation at Harvard has let down our students, who can no longer be proud of their A's, and damaged our institutional reputation. The report of the Subcommittee on Grading identified a cap on the number of A grades as the most practical mechanism to address grade inflation. Their proposed formula is that the number of A grades in a class of N students be no greater than $0.2N + 4$. The 0.2 quantifies "extraordinary distinction" as being within the top 20% of Harvard undergraduates, and the +4 is meant to give a buffer to small classes where students may perform their best work. Unfortunately, the proposed formula will not actually solve the problem. The flaw is that the flat constant (+4) is far too generous to small classes, which is where grade distinctions are most useful, and puts all the corrections in larger classes. It lowers the number of A's overall, but does not address the real concern the committee emphasized: grade compression. Concretely, it reduces the A rate from 77% to only 71% in classes with $N \leq 12$ students while cutting the A rate from 61% to 27% in larger classes. So any student taking an advanced class can still expect an A, regardless of their concentration or how exceptional they are. Instead, this amendment proposes an alternative formula for the cap: at most $0.2N + 0.6\sqrt{N}$ A's. This formula is simply the 95% confidence upper bound on the number of extraordinary students expected in a class of N. It produces nearly the same net A rate as the committee's formula (31.0% vs. 32.3%), but is statistically honest, scales appropriately with N, and addresses grade compression where it is most critical. It provides the right amount of flexibility to match historical trends in A rates in small and large classes; it allows instructors in small classes to be generous with A's, within reason, but not without limit; it forces all instructors to bear the responsibility for fighting grade inflation. And most importantly, it allows Harvard to say publicly that we are not only making it look like we are fighting grade inflation, but are actually fighting it.

1. What grades are for

The Student Handbook states that all grades represent varying levels of mastery, but the A grade is special. The A grade indicates "extraordinary distinction." The Committee on Grading lamented that without "a recognized standard for 'extraordinary distinction', instructors are left unmoored in their grading." It then noted that "Mastery is an absolute measure whereas distinction is a relative measure" and proceeded to recommend a concrete interpretation of the Handbook's ambiguous wording. Grades A− and below can be awarded on an absolute scale: a student with an A− has mastered the subject, a B+ grade is not quite there, and so on. An A grade, on the other hand, representing extraordinary distinction, is to mean extraordinary relative to other Harvard students. More precisely, extraordinary means in the top 20% of each class, and Harvard should aim to have 20% A's awarded on average. Considering that A is the only relative grade, the logical recommendation was then to cap the number of A grades. The proposed formula is

$$N_A \leq 0.2N + 4, \tag{1}$$

where N_A is the number of A grades in a class with N students. The 0.2 in this formula is the 20% target and the +4 is an allowance for small classes to award relatively more A's. The consequence of this formula is that overall, even if instructors maximize the number of possible A's, the net number of A's among undergraduates at Harvard should not exceed 32%. Compared to current rates in excess of 60%, this would be a successful counter to grade inflation.

This amendment does not present any revision of the committee's main findings. The cap on A's seems like an excellent and easily implementable way to fight grade inflation while staying within the guidelines of the Student Handbook. Instead it takes strong exception to the formula for the cap. The main point of contention is the +4. This is added by hand to the 20% baseline with the justification "Because small courses attract advanced and highly motivated students, we recommend allocating an additional four A grades to each class to raise the effective cap for smaller courses." There is merit to this claim. However, all students have some concentration, so they all take advanced classes which are generally small. Yes, a student is excellent in their field compared to the Harvard population at large, but that's true of all students! For a student to be in the top 20% in an intro class, they must be extraordinary relative to other students in intro classes. For a student to be in the top 20% in advanced classes, they must be extraordinary at advanced work, relative to other advanced students.

The consequences of the committee's formula are dire. One salient statistic is that currently students in smallish classes, with no more than 12 students, get around 77% A's on average, while students in larger classes, above 12 students, get around 61% A's. The cap formula $0.2N + 4$ would lead to 71% A's in the same smaller classes and 27% A's in the larger ones. Thus the implication of the committee's formula is that we do not consider grade inflation in small classes to be a problem worth correcting. Having north of 70% A's in small classes is okay, because these students are advanced and highly motivated. The main purpose of this amendment is to contest this claim.

Why is it important to curb grade inflation in advanced classes? One reason is that it is precisely these classes for which the grade is most important and where grade compression is most damaging. When students apply to graduate school, the admissions committee will look at their transcript and almost entirely focus on the advanced classes in their concentration. If those classes are giving A's to 70% of students, then the grade is not providing any information about how exceptional a student is in their field, beyond what is already in the transcript. In addition, there are unintended consequences: students may try to pick small courses they are not interested in or advanced courses they are not prepared for to improve their grades. A likely outcome is a bimodal system, with fierce competition in large classes and no competition at all in small classes. These outcomes benefit no one.

So what is the alternative? This amendment proposes a formula for the cap that treats all students and all courses fairly:

$$\text{cap}(N) \;=\; \left\lceil\, 0.2\,N + 0.6\sqrt{N}\,\right\rceil, \tag{2}$$

where $\lceil x \rceil$ means the smallest integer greater than or equal to x (one can write this alternatively as $N_A \leq 0.2\,N + 1 + 0.6\sqrt{N}$). This formula is the unique solution to a basic statistical problem. If extraordinary means top 20% relative to the cohort of Harvard undergraduates, then how many

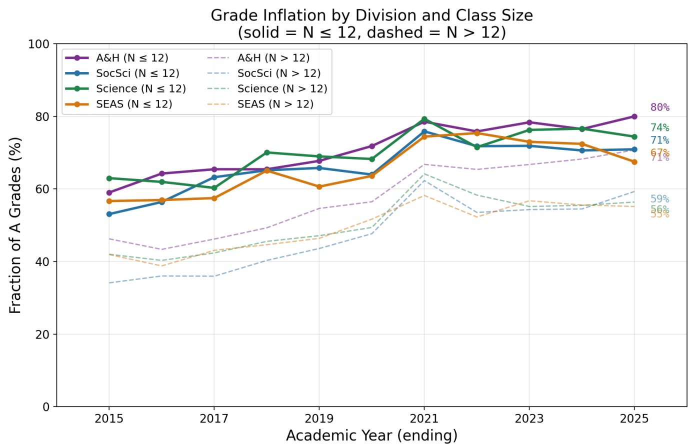

Figure 1: Share of A grades by division *and* class-size bin, 2014–15 through 2024–25. Solid lines: small classes ($N \leq 12$). Dashed lines: large classes ($N > 12$). In every division, small classes give A's at a substantially higher rate than large ones; both bins have inflated together. Arts & Humanities is the most inflated division in both bins, reaching 80% in small classes and 71% in large classes in 2024–25.

extraordinary students should we expect in a class of N students? The answer is, on average $0.2N$. If fewer than 20% of these enroll in a class, fewer A's can be given, but what if there's an upward fluctuation? The $\sqrt{N}$ factor handles these fluctuations. The formula given provides enough A's for all deserving students 95% of the time.

The remainder of this proposal compares quantitatively and qualitatively the committee's proposed formula and the replacement. It provides a detailed derivation of the main formula in Appendix A and a table of cap values under each rule in Appendix B. It also looks at the history of FAS grading to show that the +4 formula is not consistent with historical trends, while the main formula is, and provides a discussion and response to various questions about the proposal.

2. A lesson from history

One way to check whether a formula for the grade cap is sensible, and whether it sensibly fights inflation, is to ask how it affects grade distributions across class sizes and across divisions. In the 2023 *Report on Grading* it was noted that in 2002–2003, there were 30% A's given on average in Arts and Humanities classes compared to 24% in SEAS, 25% in Science and 21% in Social Science, while in 2020–2021 there were 73% A's given in Arts and Humanities compared to 60% in SEAS, 65% in Science and 65% in Social Science. This shows that grade inflation has increased across all divisions but perhaps more so in Arts and Humanities. One can then ask: is the disparity based on a cultural difference across divisions, or is it simply because Arts and Humanities has more small classes? The answer is the latter.

In Figure 1 we show the fraction of A's awarded in each division separated into different class size bins ($N \leq 12$ and $N > 12$) over time. This figure uses actual grade data from 2014–2025 extracted from the FAS Grade Distribution Dashboard. It indicates that grading is in fact remarkably consistent across divisions, that grade inflation has occurred in all class sizes at nearly the same rate, and that the disparity in A rates can be attributed almost wholly to class size differences. Harvard instructors consistently award 10–20% more A's in smaller classes than larger ones, and

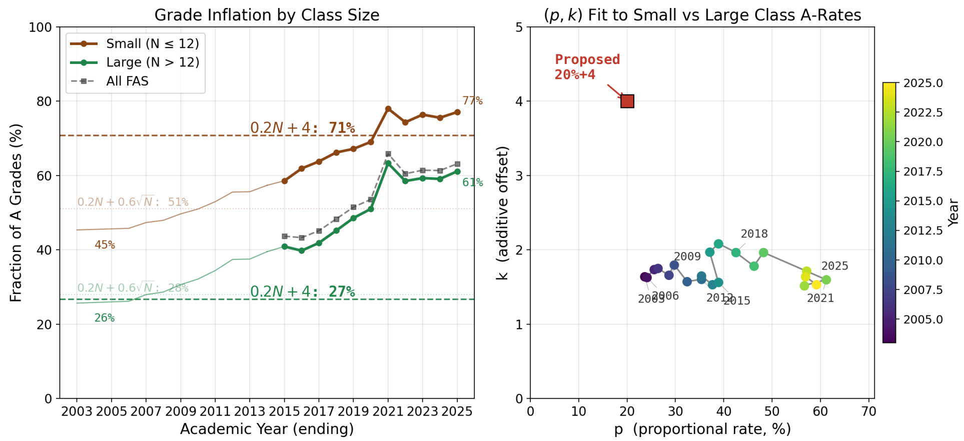

Figure 2: **Left:** Share of A grades for small ($N \leq 12$) and large ($N > 12$) classes, FAS-wide, 2003–2025. Dashed horizontal lines mark the enrollment-weighted caps each rule imposes on each bin using the 2024–25 class-size distribution: committee's rule (71% small, 27% large) and proposed statistical rule (51% small, 28% large). The committee's +4 cap on small classes sits essentially on top of the current rate. **Right:** Best-fit (p, k) in $pN + k$ that reproduces the small- and large-class A-shares at each year, applied to the 2024–25 class-size distribution. k stays in the range 1.5–2.0 throughout the decade while p rises from 39% to 59%; the committee's proposed $(20\%, 4)$, marked in red, is an outlier in both coordinates.

have done so for at least a decade.[1]

One conclusion from Figure 1 is that Harvard instructors seem to have a shared sense that smaller classes deserve more A's, at a rate consistent across divisions. This validates the committee's recommendation that the formula allot more A grades to smaller classes. Moreover, it gives a quantitative anchor as to how many grades students in small classes deserve relative to students in larger classes. We can then ask whether the $0.2N + 4$ formula preserves this relative grade distribution or distorts it. To answer this question, we first pool the divisions. The result is shown on the left of Figure 2, which demonstrates a 20% difference in grading between small and large classes that is consistent over time (and among divisions). In this figure we have also extrapolated back to 2003 using the anchor points on division disparity from the 2023 report and normalizing to the total A rate, which is also given.

We can then ask: if the historical differences are well-modeled by a formula $pN + k$, what would be the appropriate values of p and k? For each year we solve for the values of p and k that reproduce the observed A rates in small and large classes. The result is in the right panel of Figure 2. The conclusion is clear: the baseline rate p of A grades has increased from roughly 20% in 2003 to 60% today, while the offset k has stayed between 1 and 2. The committee's proposal of $k = 4$ is simply

[1] One can extend the curves farther back in time. Although the FAS Grade Distribution Dashboard data is not available from before 2014, the *Report on Grading*'s numbers from 2002–2003 confirm that the class-size disparity has persisted for at least 20 years. Dean Claybaugh's 2025 *Update on Grading and Workload* shows that the average GPA at Harvard has risen nearly linearly since 1985, from 3.17 to 3.75. The history before that is striking: average grades were essentially flat for the better part of a century (mean GPA 2.46 in 1889, 2.55 in 1950, 2.7 in 1963), then jumped sharply during the Vietnam War, when faculty were reportedly reluctant to fail students into the draft (mean GPA 3.0 by 1967), and have climbed steadily ever since. Concern about inflation is not new either: an 1894 Harvard report—issued just eleven years after letter grades were introduced at Harvard in 1883—was already lamenting that "Grades A and B are sometimes given too readily."

Table 1: Fraction of A's split by class size. This table uses 2024–2025 enrollment data; grade fractions and bin sizes include only undergraduates. The predicted A rates under the two formulas assume that instructors will assign the maximum allowed number of A grades.

Class size	courses	ugrads	%A actual	%A under $\lfloor 0.2N + 4 \rfloor$	%A under $\lceil 0.2N + 0.6\sqrt{N} \rceil$
Small ($N \leq 12$)	1,319	7,201	77.1%	70.8%	51.1%
Large ($N > 12$)	912	49,367	61.1%	26.7%	28.1%
All FAS	2,231	56,568	63.1%	32.3%	31.0%

not justified by the data. It is the outlier shown in red on the right. Taking $k = 1$ is more defensible. But with $k = 1$ the lower bound formula $\lfloor 0.2N + 1 \rfloor$ is the same as an upper bound $\lceil 0.2N \rceil$ with no offset at all. This would actually be a sensible formula. However, it does not allow for sufficient variance: a statistic not shown in the net A rates in these plots.

We can directly see how the committee's cap formula fails to address grade inflation by comparing its allowed A's in small and large classes to the historical grades. The formula leads to 71% A's in small classes and 27% in large classes, shown as the dashed horizontal lines in the left panel of Figure 2. This brings small-class A rates back to 2020 rates and knocks large classes back to 2004. The 71% cap for small classes is essentially on top of the current 77% A rate, so it does not meaningfully roll back grading in small classes at all. Harvard cannot honestly claim it is curbing grade inflation if it continues to allow upwards of 70% A grades in small classes.

In contrast, the statistically-justified formula $0.2N + 0.6\sqrt{N}$ gives a cap of 51% A's in small classes and 28% in large classes. The small-class level corresponds to circa 2010 rates and the large-class level to circa 2006. These rates are shown as faint dotted lines in the left panel of Figure 2. One can still object that 51% A's in small classes is too high. But it is a meaningful rollback from 77%, and it is consistent with the historical trend of giving more A's in small classes. It also preserves the relative distribution of A's across class sizes, which is something the committee's formula fails to do.

Table 1 gives the detailed breakdown. Two features of the table stand out. First, at the college level the two rules give nearly identical FAS-wide A rates (32.3% under $0.2N + 4$, 31.0% under the proposed rule), so any worry that one rule is dramatically more permissive than the other is misplaced. Second, where the two rules differ is almost entirely in small classes. In the large-class group the committee's cap (26.7%) and the proposed cap (28.1%) are within two percentage points of each other and both well below the current 61.1% A rate. A breakdown into finer class size bins is given in Appendix B.

3. Considerations

Here we collect some concerns about the amendment, with responses.

Isn't only allowing 2 A's in a class of 3 cruel? Who is the odd one out? Giving two A− grades and one A grade in a class of 3 is perfectly consistent with A− being mastery and A being extraordinary distinction. When these students apply for graduate school, it is likely that only one will get in. When applying for jobs, there is often only one position. If two students are exceptional, an instructor can give two A's. That 3 exceptional students end up in the same small seminar of 3, by chance alone, should only happen once in 125 years.

That being said, classes of 3 are so small that the within-class distinction may not be meaningful. In that case, the instructor is free to move to SAT/UNS. In fact, seeing a larger number of ultra-small classes move to SAT/UNS is a sensible way to handle these unusual learning environments. Grades and the GPA are meant to be statistical averages, so if a class type is not statistically meaningful it should be excluded from the average.

More importantly, the perspective of this question is backwards. It is not that the student receiving the A− is being singled out as a loser, but rather the student receiving an A being singled out as extraordinary. Fighting inflation means replacing the shame of failure with pride of success.

The statistical model treats each student as independent. Is that realistic? No. Students in a class are not independent draws: they may have prepared together, share a concentration, or all have learned from the same extraordinary instructor. In fact, to be extraordinary almost certainly means working and learning with your peers, helping them to perform at their best as you do. But if we want extraordinary to mean top 20%, this must mean top 20% even including these correlations.

Correlations are difficult to measure and include. Any attempt to add a term to the cap formula to model them would be controversial and complicated. A simple formula, like either the committee's or the amended one, has the benefit of transparency. The best defense that the amended formula allows for an appropriate level of correlation and the committee's formula does not is in Figure 2. Both formulas allow for more A grades in smaller classes where correlations are higher, but the committee's formula abandons any attempt to rein in grade inflation at all in these classes, while the amended formula does so in a principled and justified way.

In 3,000 courses a year, 5% is 150 classes where the cap is too tight by chance. Isn't that too many? This reasoning assumes that grading would naturally be balanced around 20%. There is no reason to believe instructors would grade that way. Grade inflation is the problem that instructors, if left unconstrained, would give well more than 20% A's. This means that students who are on the true A/A− boundary would almost always be pushed up. That many students will get many more A's than they deserve, even with the statistical model, far overcompensates for the rare case when the slack in the formula can't cover the fluctuation.

You assume every instructor will award the maximum number of A's the cap allows. Many classes already give fewer; aren't the projected rates inflated? The policy question is what the rule should *permit*, independent of what current practice happens to be. The actual consequence of imposing a cap are hard to predict, but a natural expectation is that the cap would become the new norm over time. Instructors who currently undershoot have a clean Schelling point to drift toward, and students may expect the maximally permitted number of A's.

If we instead assume instructors will give no more A's than they do now, even if permitted, the overall A rate comes out to 31% under the committee's $0.2N + 4$ cap and 30% under the proposed rule, so these rates are essentially unchanged (from 32% and 31% respectively). In the $N \leq 12$ bin the committee's rule would give 62% (instead of 71%) and the proposed rule would give 48% (instead of 51%). The improved stability across bin size under varied assumptions of grading behavior is additional support for the amended rule.

Why ceiling rather than floor? Because we want to be 95% confident that the cap is *at least* as large as the number of extraordinary students in the class. Rounding up gives an integer cap at or above $0.2N + 0.6\sqrt{N}$, so the 95% guarantee holds (and may in fact be exceeded, giving higher confidence at some class sizes). Rounding down gives a cap below that value, and can push the confidence below 95%. One can also write the proposed formula as $\lfloor 0.2N + 1 + 0.6\sqrt{N} \rfloor$ if that for some reason seems more palatable.

The 20% baseline is arbitrary. It is. This was the committee's recommendation. Here is a defense. 20% is sufficient to lower the median grade from A to A−. It thereby significantly reduces the grade *compression* problem without making the A so rare that it is unmotivating. The strongest critique of a 20% target is that it only reduces the mean GPA mildly, from 3.77 to 3.66. The recent report of the Yale Committee on Trust in Higher Education suggested that a mean GPA of 3.0 should be the goal. No cap on A's alone can produce a 3.0 mean — that requires the full grade range to come back into use, which would only happen if instructors willingly and actively correct their own grading. A first step in that direction is for all instructors to acclimate to A no longer being the default grade in any class, and the $0.2N + 0.6\sqrt{N}$ cap forces exactly that adjustment.

The 95% confidence threshold is itself arbitrary. Every threshold is a choice. 95% is a widely understood and accepted standard in statistics. A tighter 99% cap would give more A's at every class size; a looser 90% cap would give fewer.

Freshman seminars and expository writing. Freshman seminars are excluded from Table 1 because they are graded SAT/UNS. Expository writing courses are included; if any of those are also SAT/UNS the effect on the totals is small.

Is the class-size distribution stable? A natural worry is that the cap's impact depends sensitively on how class sizes are distributed today, and that any formula calibrated to 2025 might drift out of calibration as courses shift. The enrollment data allow a direct check. Restricting to letter-graded lecture-style courses and aggregating discussion sections under a shared instructor into single units, the distribution of class sizes across Fall terms since 2013 is essentially stationary: the median course has held at 11 or 12 students for more than a decade, and the share of courses with more than 50 students has drifted up by a single percentage point. Whatever is driving the observed growth in A grade share, it is not a changing denominator.

4. Conclusions

Restoring the meaning of the A is a collective act. The faculty must commit together to a rubric in which the A again means extraordinary distinction and A− indicates mastery. The committee proposes that the most practical approach to correcting grade inflation is to cap the A's and aim for a nominal 20% A rate. The refinement to the committee's proposal in this amendment is not about the goal or the strategy but about the details of the mechanism. The committee's formula for the A grade cap $0.2N + 4$ leaves the most inflated segment, classes with $N \leq 12$ students, nearly untouched and asks students taking larger classes to bear the entire burden of the rollback. The proposed amended formula $0.2N + 0.6\sqrt{N}$ corrects this.

A faculty member with an unusually strong small class still has substantial room to grade generously under the amended rule. A class of five retains a cap of three A's (three times the baseline); a class of ten retains four. What is removed is the ability to declare every student in a seminar of five extraordinary, which happens by chance for a given class less than once in a thousand years and should not be a routine outcome of a grading rule. Critically, every instructor in every class with more than 2 students must give some A− grades. Every instructor must participate in fighting grade inflation: it is a collective problem that requires a collective solution.

Taking a step back, it is worth keeping sight of what fighting grade inflation is really about. Harvard is the most visible undergraduate institution in the country, the one most associated with grade inflation, and the one most equipped to lead grading reform if it chooses to. Harvard then has the unique opportunity to lead by example and restore the meaning of the A grade. The choice in front of the faculty is binary. Either A is meant to pick out extraordinary work relative to other Harvard students, in which case the faculty has to commit to that meaning at every level of the curriculum, in every department, in seminars and lectures alike, or A is not meant to pick out anything in particular. Fixing grade inflation requires the faculty to stand behind a single sensible approach, setting the right example, for the right reasons.

A. Derivation of the statistical cap

The origin of the $0.6\sqrt{N}$ factor in Eq. (2) is the statistics of discrete fluctuations. We imagine that the N students in a class are drawn from an appropriate pool (all Harvard undergraduates, all English concentrators, etc.), in which a fraction p (= 20%) of the students are extraordinary. What is the probability that exactly m of the N students in the class are extraordinary? Number the students 1 to N, say by the order of enrollment. The chance that the first m to enroll are extraordinary and the remaining $N-m$ are not is $p^m(1-p)^{N-m}$. Allowing that the m extraordinary students can be any subset leads to the **binomial distribution**

$$B_N(m) \;=\; p^m\,(1-p)^{N-m}\,\frac{N!}{m!\,(N-m)!}. \tag{3}$$

The combinatorial factor ${}_NC_m = \binom{N}{m} = \frac{N!}{m!\,(N-m)!}$ counts the number of ways to choose which m of the N students are the extraordinary ones.

Using the binomial distribution with $p = 0.2$ gives the following probabilities of m extraordinary students in a class of size N:

	$m=0$	$m=1$	$m=2$	$m=3$	$m=4$	$m=5$
$N=1$	80.00%	20.00%	—	—	—	—
$N=2$	64.00%	32.00%	4.00%	—	—	—
$N=3$	51.20%	38.40%	9.60%	0.80%	—	—
$N=4$	40.96%	40.96%	15.36%	2.56%	0.16%	—
$N=5$	32.77%	40.96%	20.48%	5.12%	0.64%	0.03%

This shows how rare it is to have many extraordinary students in one class. In a class of 5, around 1/3 of the time there will be no extraordinary students at all, and only once out of every 3000 times the class is offered might one expect all 5 to deserve A's.

The mean and standard deviation of the binomial distribution are

$$\langle m\rangle = \sum_{m=0}^{\infty} m B_N(m) = pN, \qquad \sigma(m) = \sqrt{\langle m^2\rangle - \langle m\rangle^2} = \sqrt{N\,p\,(1-p)}.$$

For $p = 0.2$ the standard deviation simplifies to $\sigma = 0.4\sqrt{N}$. The absolute fluctuation σ grows like $\sqrt{N}$; the relative fluctuation σ/N shrinks like $1/\sqrt{N}$. In a seminar of 5, the standard deviation is about 18% of the class. In a lecture of 500, it is about 2%. Large classes are numerically stable; small classes are not.

Suppose we had a hard cap: no more than 20% A's ever. If the underlying rate really is $p = 0.2$, the probability of exceeding a strict 20% cap by at least one student hovers around 25–35% across class sizes from 1 to 20, and by at least two students around 10–20%. To improve on the hard cap, we ask what the cap should look like so that we are at least 95% confident of having enough A's for the number of extraordinary students in the class. That is, we want

$$\Pr(m > \text{cap}) \;=\; \sum_{m\,=\,\text{cap}+1}^{N} B_N(m) \;\leq\; 0.05.$$

There is no closed-form expression for the cap as a function of N, but there is a good closed-form approximation that holds in the limit of large N. For N large enough, the binomial distribution is well approximated by a normal (Gaussian) curve with the same mean and variance:

$$B_N(m) \;\approx\; \frac{1}{\sqrt{2\pi}\,\sigma}\,\exp\!\left[-\frac{(m-\mu)^2}{2\sigma^2}\right], \qquad \mu = pN, \;\; \sigma = \sqrt{Np(1-p)}.$$

The sum becomes a Gaussian integral, and after the change of variable $z = (m-\mu)/\sigma$ the inequality becomes

$$\int_{z^*}^{\infty} \frac{1}{\sqrt{2\pi}}\, e^{-z^2/2}\, dz \;\leq\; 0.05, \qquad z^* \;=\; \frac{\text{cap} - \mu}{\sigma}.$$

The smallest z^* that satisfies this is $z^* \approx 1.645$, computed numerically. Taking the cap to be $\mu + 1.645\sigma$ and rounding up to the nearest integer gives the large-N rule

$$\text{cap}_{\text{large}\,N}(N) \;=\; \left\lceil pN \,+\, 1.645\sqrt{Np(1-p)} \right\rceil \;=\; \left\lceil 0.2N + 0.66\sqrt{N} \right\rceil \qquad (p = 0.2). \tag{4}$$

This agrees with the exact binomial condition to within one A grade at every class size. However, because the binomial is discrete and the Gaussian is continuous, the two do not agree exactly at small N, and the Gaussian approximation is consistently generous there: the closed form with $0.66\sqrt{N}$ overshoots the exact binomial, giving an extra A in more than half the class sizes in the range $N = 1$–50. We can *tune* the coefficient of the $\sqrt{N}$ term to improve the formula at small class sizes, where an extra A matters most. If we replace 0.66 by 0.60, the formula

$$\boxed{\text{cap}(N) \;=\; \left\lceil 0.2N \,+\, 0.6\sqrt{N} \right\rceil} \tag{5}$$

matches the exact binomial at *every* class size from $N = 3$ through $N = 48$ except at $N = 14$ and $N = 29$ where it gives an extra A. The first undershoot—where the formula gives one A fewer than the strict 95% confidence rule—is at $N = 49$, and further undershoots happen at a sprinkling of class sizes thereafter.

We adopt Eq. (5) as the amended rule. The reasoning is that for the purposes of reining in grade inflation, small classes are where correct calibration matters most: they are where $0.2N + 4$ is most generous, where the intuitive appeal of "everyone is extraordinary" is strongest, and where a cap that is even one or two A's off in either direction changes the grading experience in a visible way. At large N, by contrast, the difference between the 0.60 and 0.66 coefficients for $N \geq 49$ amounts to only about one fewer A than the 95% confidence fluctuation allows.

B. Lookup table of A-grade caps

Table 2: A-grade cap as a function of class size N (undergraduate enrollments only). The "90% conf." column gives the central 90% interval of Binomial$(N, 0.2)$, from the 5th to the 95th percentile: the number of extraordinary students one would see in 90% of classes of that size if the true rate is $p = 0.2$. The $\lfloor 0.2N + 4 \rfloor$ column is the committee's proposed rule; the cap(N) column is the amended statistical rule $\lceil 0.2N + 0.6\sqrt{N}\,\rceil$.

N	90% conf.	$\lfloor .2N{+}4 \rfloor$	cap(N)	N	90% conf.	$\lfloor .2N{+}4 \rfloor$	cap(N)	N	90% conf.	$\lfloor .2N{+}4 \rfloor$	cap(N)
1	[0–1]	1	1	26	[2–9]	9	9	60	[7–17]	16	17
2	[0–1]	2	2	27	[2–9]	9	9	70	[9–20]	18	20
3	[0–2]	3	2	28	[2–9]	9	9	80	[10–22]	20	22
4	[0–2]	4	2	29	[2–9]	9	10	90	[12–24]	22	24
5	[0–3]	5	3	30	[3–10]	10	10	100	[14–27]	24	26
6	[0–3]	5	3	31	[3–10]	10	10	110	[15–29]	26	29
7	[0–3]	5	3	32	[3–10]	10	10	120	[17–31]	28	31
8	[0–4]	5	4	33	[3–11]	10	11	130	[19–34]	30	33
9	[0–4]	5	4	34	[3–11]	10	11	140	[20–36]	32	36
10	[0–4]	6	4	35	[3–11]	11	11	150	[22–38]	34	38
11	[0–5]	6	5	36	[3–11]	11	11	160	[24–40]	36	40
12	[0–5]	6	5	37	[4–12]	11	12	170	[26–43]	38	42
13	[0–5]	6	5	38	[4–12]	11	12	180	[27–45]	40	45
14	[1–5]	6	6	39	[4–12]	11	12	190	[29–47]	42	47
15	[1–6]	7	6	40	[4–12]	12	12	200	[31–49]	44	49
16	[1–6]	7	6	41	[4–13]	12	13	210	[33–52]	46	51
17	[1–6]	7	6	42	[4–13]	12	13	220	[34–54]	48	53
18	[1–7]	7	7	43	[4–13]	12	13	230	[36–56]	50	56
19	[1–7]	7	7	44	[5–13]	12	13	240	[38–58]	52	58
20	[1–7]	8	7	45	[5–14]	13	14	250	[40–61]	54	60
21	[1–7]	8	7	46	[5–14]	13	14	260	[42–63]	56	62
22	[2–8]	8	8	47	[5–14]	13	14	270	[43–65]	58	64
23	[2–8]	8	8	48	[5–14]	13	14	280	[45–67]	60	67
24	[2–8]	8	8	49	[5–15]	13	14	290	[47–69]	62	69
25	[2–8]	9	8	50	[6–15]	14	15	300	[49–72]	64	71

C. Percent A grades by division and class size

The three tables below use a common color scale: the shade of each cell is directly proportional to the percentage shown, so the tables can be read against one another at a glance. All three are computed from AY 2024–25 course-level grade data restricted to undergraduate letter-graded enrollments. Each course is assigned to a size bin based on its *undergraduate* enrollment only, matching the binning used in Table 1. This differs from the public FAS grading dashboard, which bins by total class size including graduate and cross-registered students. The difference is not large, and we have verified that the dashboard numbers are reproduced with the dashboard class-size definition. Within each (division, size) cell the reported percentage is enrollment-weighted across courses (total A's / total undergraduate enrollments). The "All" column is the same weighted average taken over every course in the bin. Table 3 shows the observed share of A grades, and Tables 4 and 5 show the maximum share that would be allowed under the committee's proposed cap $0.2N + 4$ and the amended statistical cap $0.2N + 0.6\sqrt{N}$, respectively.

Table 3: Observed percent of undergraduate grades awarded at the A level, by division and class size, AY 2024–25. The public FAS grading dashboard reports the same underlying A counts but bins each course by total class size rather than undergraduate enrollment, so the numbers here are a few points different from what appears on the dashboard.

Class Size	Gen. Ed.	Arts & Hum.	SEAS	Science	Soc. Sci.	All
Small (≤ 12)	—	80.4%	67.5%	73.5%	70.9%	77.1%
Medium (13–30)	45.1%	76.2%	64.2%	69.2%	68.6%	71.3%
Large (31–50)	54.0%	74.9%	59.1%	68.5%	73.5%	68.5%
Extra Large (51–150)	66.3%	75.4%	55.0%	63.4%	59.5%	63.7%
Huge (≥ 151)	66.8%	47.3%	50.1%	42.8%	53.0%	52.1%

Table 4: Maximum percent of A grades allowed under the committee's proposed cap $\lfloor 0.2N + 4 \rfloor$, enrollment-weighted within each (division, class size) cell.

Class Size	Gen. Ed.	Arts & Hum.	SEAS	Science	Soc. Sci.	All
Small (≤ 12)	—	72.0%	69.0%	69.3%	68.2%	70.8%
Medium (13–30)	35.2%	40.0%	37.4%	38.0%	39.2%	39.0%
Large (31–50)	29.1%	29.1%	29.2%	29.0%	29.6%	29.2%
Extra Large (51–150)	24.5%	25.0%	24.7%	24.3%	24.3%	24.6%
Huge (≥ 151)	21.4%	20.6%	21.6%	21.4%	21.3%	21.3%

Table 5: Maximum percent of A grades allowed under the amended statistical cap $\lceil 0.2N + 0.6\sqrt{N} \rceil$, computed the same way.

Class Size	Gen. Ed.	Arts & Hum.	SEAS	Science	Soc. Sci.	All
Small (≤ 12)	—	51.7%	49.6%	50.2%	50.3%	51.1%
Medium (13–30)	33.8%	36.8%	36.0%	36.1%	36.7%	36.5%
Large (31–50)	30.6%	31.0%	30.9%	30.8%	31.2%	30.9%
Extra Large (51–150)	27.3%	27.7%	27.4%	27.0%	27.0%	27.3%
Huge (≥ 151)	23.7%	22.4%	24.2%	23.8%	23.6%	23.6%

Acknowledgements

The course-level grade data used in this proposal come from the FAS Grade Distribution Dashboard. The FAS-wide and per-division A-shares prior to 2014–15 are taken from the 2025 FAS Report on Grading (overall FAS by year) and the 2023 FAS Report on Grading (per-division anchors for 2003 and 2021). Class-size distributions used in the appendix cap tables come from the Fall 2025 course enrollment data pulled from the FAS Registrar's archive. Analysis of the data was performed with Claude Opus 4.6, but the amendment itself was written entirely by a human.