2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #50033

Developing an Outcome-Based Attendance Policy for Introductory Computing Courses

Dr. Chris Bourke, University of Nebraska-Lincoln

Developing an Outcome-Based Attendance Policy for Introductory Computing Courses

Numerous studies have demonstrated a positive correlation between attendance and grades. Given the generally lower outcomes in engineering and computing courses, an effective policy that improves attendance can potentially have a positive effect on student outcomes. However, traditional policies are either too lenient and ineffective or overly stringent and inflexible. In this study we present an alternative: an *Outcome-Based Attendance Policy* which requires and grades attendance but allows students who do not attend the opportunity to earn back attendance points by demonstrating competence in the material, as evidenced by achieving a certain threshold grade in all other assessments. This approach is designed to give students agency in how they engage with the course and material while emphasizing the importance of participation. It shifts the focus to learning *outcomes* rather than the *process*.

We document our experience in developing and administering this policy for laboratories in introductory computing courses. We present and analyze data from several years of implementing this policy and share our observations and best practices for those considering adopting or revising attendance policies in their own courses.

1 Introduction

Regular attendance in courses is vital to academic success. This has been demonstrated by large scale studies in general [1, 2] and corroborated by smaller studies involving engineering courses [3] and specifically in computing [4–6]. Educators and administrators at all levels have tried to improve student attendance using a variety of methods: incentivizing attendance (carrots) and/or establishing punitive measures for non-attendance (sticks).

Institutions may provide only broad guidelines or policies with respect to attendance and "excused absences," giving faculty discretion in applying them. Faculty may have a variety of viewpoints. The traditionalist viewpoint is to have no course attendance policy. From this perspective, students are viewed as adults and expected to be present and to make responsible decisions regarding their academic success. Having an explicit (graded or otherwise) attendance policy is seen as antithetical to instilling responsibility and robs students of their autonomy. Moreover, a student shouldn't "earn" points by simply being there. Attendance is the bare minimum effort and should not be rewarded.

A more modern viewpoint takes a more paternalistic approach. College students, especially freshmen, are coming from a K–12 environment where attendance is a legal obligation and are suddenly thrown into a hands-off environment. They cannot be expected to make such an abrupt transition and take full responsibility for their learning. Some guardrails are necessary. Instructors with this perspective may require attendance either directly (as a portion of the grade) or indirectly (grading components based on "participation," active learning activities, or "pop" quizzes as a proxy for attendance). They may view regular attendance as necessary for engaging with the material or building a community.

In a post-COVID world, rigid graded attendance policies have been actively discouraged, as forcing sick students to attend poses serious health risks. Furini et al. [7] observe that in a post-COVID world students now have a large variety of alternative learning methods. Mandatory attendance "negates many of the benefits students have found in online learning, such as cost reduction, management of complex personal schedules, and personalized learning." Instead, the focus has shifted to *flexible* policies that encourage attendance but allow broad exceptions, such as a set number of no-questions-asked absences and alternatives like make-up sessions or exercises. Each approach has its benefits, costs, and administrative burdens (especially for instructors).

Several recent studies have called for more flexible attendance policies. Furini et al. [7] suggest that "finding a balance between tradition and innovation, might be a reasonable approach to meet the expectations and needs of contemporary students." Indeed, others have called for more flexible or balanced attendance policies. When studying ways to build community in the classroom Minnes et al. [8] observed that a "successful attendance policy occupies a middle ground between providing the incentive students need to attend and being overly coercive to the point that students who truly do not wish (or need) to come to class do so begrudgingly. We found that attendance policies that rewarded students for attending class, but that provide alternative ways to earn these rewards, were most successful in creating an engaged classroom community." Malhotra et al. [9] employed a mandatory attendance policy along with flexible (but short) assignment extensions and observed "that while flexibility is crucial, it must be balanced with structure and proactive support."

In answer to these authors, we present a novel *Outcome-Based Attendance Policy*. This policy assigns a modest portion of a student's final grade (10%) to weekly lab attendance but allows a non-attending student to earn the points back by demonstrating competence in the material. This policy focuses on the end result (learning outcomes), rather than the process (attendance), and gives students flexibility and agency in deciding how they engage with the material while still emphasizing the importance of attending. Secondarily, the policy was meant to increase overall attendance. Given the strong positive correlation between attendance and performance the aim is to improve student outcomes.

In section 2 we discuss related work and research on attendance and academic success. In section 3 we describe our outcome-based attendance policy in detail. In section 4 we report and analyze data from administering this policy in three introductory computing courses. We discuss and reflect on these results in section 5.

2 Related Work

At the K–12 level, students are *legally* required to be in school, but many schools face high rates of "chronic" absenteeism (truant in $\geq 10\%$ of classes), which is primarily a consequence of economic factors such as health conditions, lack of safe and reliable transportation, and food and housing insecurity. These factors disproportionately affect vulnerable and marginalized student populations [10]. Undergraduates are not immune to these issues, but reasons for absenteeism at the college level are more commonly due to "boredom, illness, and interference with other coursework or social life" [11], or a general lack of interest [12].

Many studies have established that regular attendance is a cofactor and has a positive correlation with grades and academic success in undergraduate education [13–15]. A large meta-analysis [1] found that attendance and class grades/GPA were strongly correlated, making attendance one of the strongest predictors of student success. In fact, the results showed that "class attendance explains large amounts of unique variance in college grades because of its relative independence from SAT scores and high school GPA and weak relationship with student characteristics such as conscientiousness and motivation." However, the same study found that mandatory attendance policies have only a small positive impact on grade averages.

Some studies have shown that having a graded attendance policy significantly increases attendance [14]. Some older studies indicate that punitive policies (point penalties) or incentives (bonus points) [16] can improve attendance, but the bulk of more recent research indicates otherwise. St. Clair [17] found results to be mixed at best, and that compulsory attendance policies can actually be counterproductive. Moore [18] found that penalties did not improve rates of attendance but did have a negative impact on passing rates. In a separate study [19], he found mandatory attendance policies were not even necessary to positively impact grade outcomes. Simply reiterating the importance of attending class resulted in higher grades.

Students recognize the importance of attendance but value their autonomy. In one study [20], students were given a choice to have attendance mandatory or optional. Most opted for the mandatory policy, which led to an attendance rate of 88%. Another study [21] offered students the opportunity to have attendance count (or not count) toward their grade and found that nearly universally (90%), students opted in to this scheme as a "pre-commitment" device. The key difference in these approaches is that students are given the autonomy to determine their own motivating factors, which has been shown to be important in general [22] and specifically in education [23].

3 Outcome-Based Attendance Policy

Our outcome-based attendance policy was originally developed in the context of a Computer Science I (CS1) course. Prior to 2022, this course had no attendance policy. Attendance was tracked informally (head counts) and observed to be very low, averaging about 30-40% each week by midsemester. It was clear from informal observations that students who attended benefited, but we were hesitant to adopt an attendance policy because we held to the aforementioned traditionalist viewpoint.

Nevertheless, in order to improve student outcomes, we decided to devise and adopt an attendance policy. The policy was piloted in fall 2022 and was successful enough that it was revised and has remained in place since. In spring 2025, the policy was expanded to our CS2 course. In the following subsections, we provide course background and context, describe the policy, and explain how we administer it.

3.1 Course Setup

This is a typical CS1 course [24] that serves primarily freshman computing majors at a large midwest R1 institution. For the majority of students, this is their first computing course. Though the course is offered every semester, we focus on "on-schedule" fall offerings. Delivered over 15 weeks, the course is organized into 12 weekly modules (covering variables, conditionals, loops, . . . , recursion, searching & sorting).

In addition to the traditional lecture, there are *two* weekly lab sessions. The first lab session requires students to complete small programming tasks with substantial starter code and is due the same day. The second is a "hack" session in which students are given a larger programming exercise and it is due the following week. Both are organized using peer programming, where students may work in pairs, and near-peer Learning Assistants (LAs) facilitate pairing if necessary. LAs also assist with learning and provide help during both sessions. Collectively, they hold an additional 30–40 open help hours per week.

Assessment for each module consists of a pre-module graded reading [25], the lab, the hack, and attendance points. Attendance counts for about 15% of the module grade (or 10% of the final grade). Each module is weighted equally. Students are allowed to utilize a filtered AI tutor chatbot [26] but no other AI usage is allowed. About 30% of the final grade is based on exams which are taken individually. Exam problems are closely aligned with previous exercises but in a new context.

3.2 Policy

Students are graded based on whether or not they attend each weekly lab and hack session. Attendance points are awarded only if they are present for the full session and, in the judgment of the LAs, put forth a full effort. Students may not be more than 5 minutes late, leave early, work on non-course material, or remain idle.

However, if a student does not earn the full points, they may still earn them back provided they achieve at least 75% of the remaining credit on the reading, lab, and hack. The idea is that if a student can demonstrate their ability to complete the material without attending, they are effectively exempt from attendance. For those unable to attend one or two sessions, it provides an alternative opportunity to succeed.

We call this an *outcome-based attendance policy* because it refocuses emphasis on learning outcomes rather than the process (attending). It gives students agency in how they approach the course while still emphasizing the importance of regular attendance.

3.3 Logistics & Administration

The outcome-based attendance policy is fully explained to students during the course introduction and in the syllabus. We emphasize the importance of attendance and engagement and strongly advise students to attend during the first few weeks to acclimate to the course. In collaboration with professional advisors, we perform regular early grade checks. For underperforming students, we reemphasize the importance of regular attendance.

During each lab session, LAs take attendance as part of the startup process (they ask for names and "Do you need help getting started?") and enter records into our Learning Management System (LMS) gradebook. Students who miss one or more sessions are initially given zeros for sessions they did not attend. Once all grading and assessment is complete for the module, attendance is updated by running a script that interfaces with our LMS API. The script applies the 75% rule and updates the attendance points accordingly, along with a documented message explaining why the points were changed or *not* changed.

The initial attendance data collection is somewhat manual but necessary, as it involves a judgment call by the LAs. Otherwise, the administrative burden is effectively zero, as the rest of the process is nearly fully automated.

3.4 Adoption in CS2

Following the successful adoption of this policy in CS1 over several years, it was adopted in our CS2 course in spring 2025. It had not been adopted earlier because the thinking was that students in CS2 should be fully transitioned into college life and responsible for their own educational success. CS2 represents a greater challenge than CS1, and some of the guardrails in CS1 should be removed. However, we observed a substantial population of students who still needed these guardrails.

Unlike CS1, our CS2 course has only one weekly lab session, but it is twice as long (about 2 hours) and is not based on modules. Instead, attendance points are awarded only if the student attends or fully completes the lab on their own.

4 Results

During the initial pilot of this policy, the process was more manual and data was not formally collected. Data was collected in fall 2023 and fall 2024 after the administration was mostly automated. The data only includes students who completed the course (withdrawals are omitted, constituting about 5% of the initial enrollment). We summarize the week-to-week breakdown of students in Figures 1 and 2 according to the following categories.

- **Attended** - students who initially attended and received credit (but who would have received credit anyway by earning at least 75% on the rest of the module assessment).
- **No Attend/Credit** - students who did not initially attend but received credit by earning at least 75% on the rest of the module assessment.
- **Attend/No Credit** - students who initially attended and received attendance points but who would *not* have received credit because they earned less than 75% on the rest of the module assessment.
- **Neither** - students who did not attend and did not receive credit because they earned less than 75% on the rest of the module assessment.

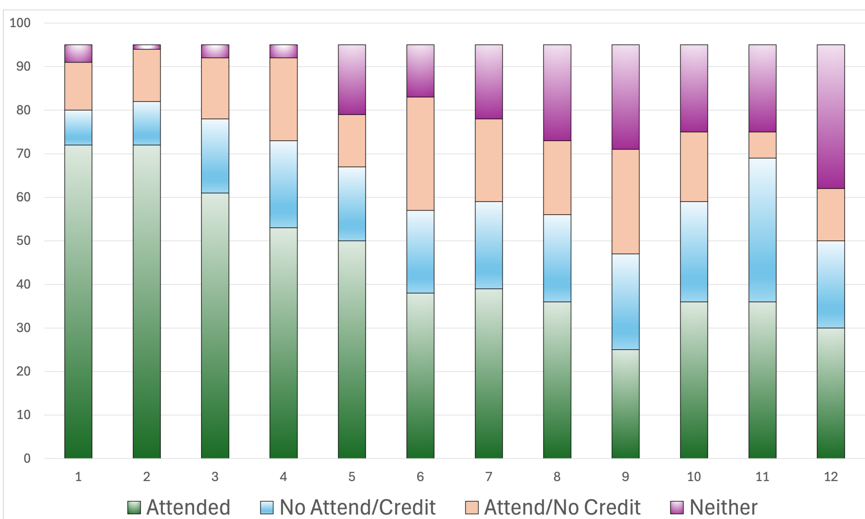

Figure 1: Week-to-week outcomes for CS1, fall 2023; total enrollment was 95.

Week-to-week students in each of the four categories do not necessarily represent the same individuals. A student could choose to attend one week and then not the next, switching back and forth.

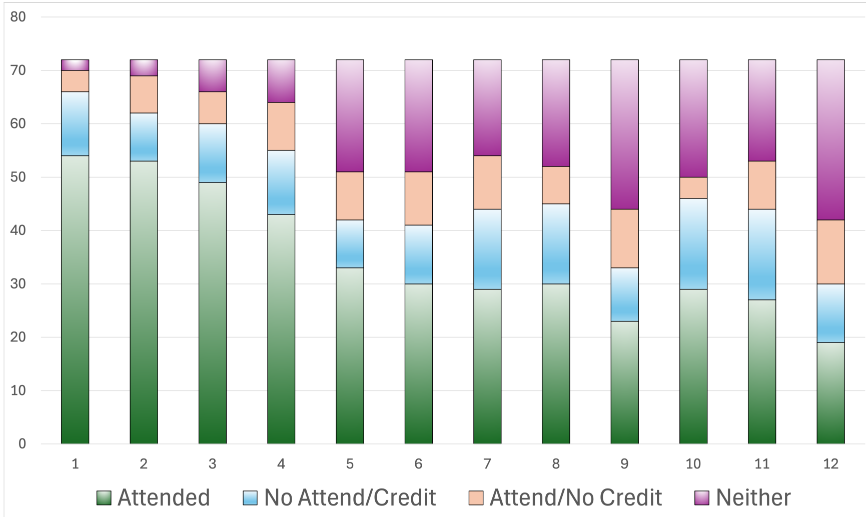

Figure 2: Week-to-week outcomes for CS1, fall 2024; total enrollment was 72.

The general trend for both years was virtually the same. There was an initially high rate of attendance of about 80% (**Attended** and **Attend/No Credit**) which gradually decreased week after week. By the end of the course, only about 35–40% were still attending.

The population of students who chose not to attend but still earned the points back by taking advantage of this policy was fairly consistent (15–20%) throughout the semester, starting even in the first week. Likewise, there was a consistent population of students (10–15%) who attended and earned credit for attending but did not meet the 75% threshold on the rest of the assessment. This does not mean that these students necessarily failed the course. Instead, this category was more volatile and contained different students week-to-week. The most negative trend was the proportion of students who neither attended nor reached the 75% threshold (**Neither**). This group was initially small (but not insignificant) at 5%, which jumped to 25–35% by week 5 and remained consistent until the end of the semester. Week 5 coincides with our first exam and represents students who didn't do well and subsequently "checked out." Ultimately, these students either failed (18% overall received a D/F) or did not perform very well (15% overall received a C).

4.1 CS2 Results

Data was also collected in our spring 2025 CS2 course and can be found in Figure 3. There were 14 labs delivered in 15 weeks. Bad weather caused one lab to be cancelled and we omit it from the data (students were unconditionally granted the attendance points).

The trends in CS2 are quite different. In the first week, attendance was about 75%, which quickly declined over the first three weeks and stabilized around 20% on average. Simultaneously, the proportion of students who did not attend but still completed the lab (and received attendance credit as a result) grew. However, the proportion who attended but did not receive credit was insignificant (and nonexistent for the majority of weeks). There were a few students in the **Neither** category, which varied slightly week-to-week but averaged only 5%.

The general trend likely differs from our CS1 experience for several reasons. First, assessment for CS2 is based primarily on a substantial programming project that had only an indirect relation

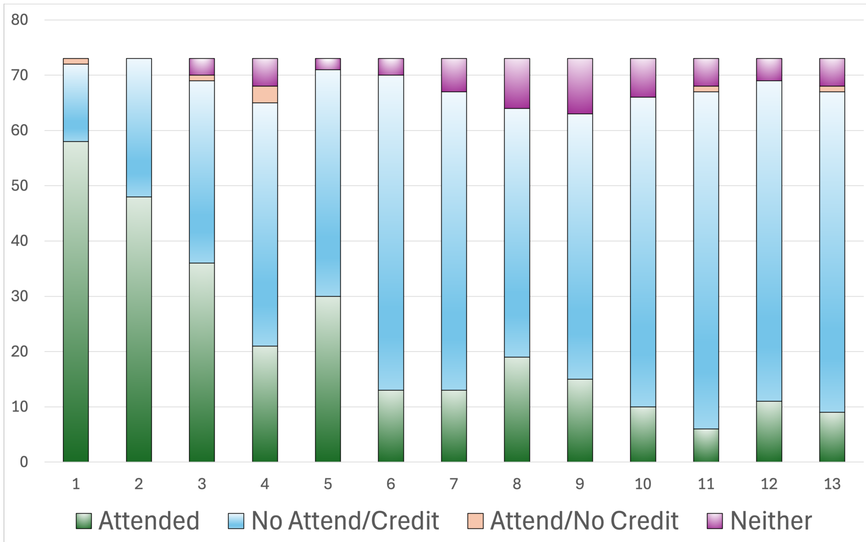

Figure 3: Week-to-week outcomes for CS2, spring 2025; total enrollment was 74.

to lab attendance. Second, lab attendance points represented only about 5% of the final grade. Finally, CS2 students are (slightly) more mature, and very low-performing students have already been filtered out in CS1.

4.2 *Correlation Analysis*

We were interested in confirming whether attendance rates correlated with performance (as measured by final grade percentage), as previous research has indicated. We combined both years of CS1 data ($n = 156$ after retake duplicates were removed) and ran a Shapiro-Wilks test to confirm that final grades were not normally distributed ($W = 0.8475$, $p = 1.916 \times 10^{-11}$). In addition, since there were only 12 lab/hack sessions, attendance rates were essentially discrete values (potentially violating linearity and/or normality). Therefore, we chose to apply nonparametric statistical tests for correlation.

A Spearman's ρ test found that attendance rates and final grades were statistically significantly correlated ($\rho = 0.4658$, $p = 8.923 \times 10^{-10}$, $\alpha = 0.01$). The interpretation is that this was a moderate (though nearly strong) positive correlation [27].

Table 1: Spearman analysis for each quartile and overall.

Quartile	Size	ρ	p-value	significant?
Q4	75	0.2214	0.0563	False
Q3	29	-0.321	0.0894	False
Q2	25	0.1781	0.3942	False
Q1	27	0.1031	0.6089	False
Overall	156	0.4658	< 0.001	True

We were also interested whether this correlation held for certain groups of students. We stratified students into quartiles based on their attendance rates: Q1 attended [0, 25]% of the time, Q2 (25, 50]%, Q3 (50, 75]%, and Q4 (75, 100]%. We explored other stratifications (median split, quantiles, etc.) but settled on quartiles because the bottom three quartiles had roughly the same number of students.

We applied Spearman's ρ to each quartile and found that none of them were statistically significantly correlated (see Table 1). This may be due to reduced sample sizes and restricted range within quartiles, which can diminish correlation strength and statistical power. Our interpretation is that, though there is a significant correlation overall, other factors do not hold uniformly across each quartile.

To better illustrate the relationship, we graphed attendance rates against final grades as a scatter plot (see Figure 4). In addition to a linear regression line, we've graphed a locally weighted scatterplot smoothing (LOWESS), which is a nonparametric regression method that fits smooth curves through data without assuming a specific functional form (like linearity). The LOWESS curve suggests that there is less of a relationship between attendance and final grades for those who **do not attend** (Q1, Q2) and a greater one for those who do (Q3, Q4).

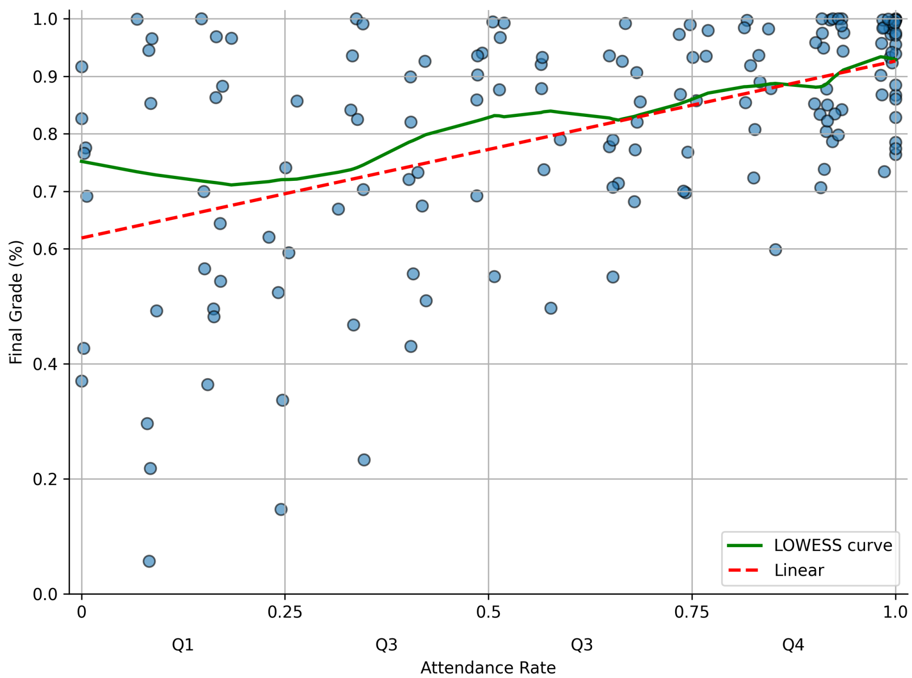

Figure 4: Attendance Rates vs. Final Grades.

4.3 Attendance vs. Grades

To determine if there was any significant difference in final grades among quartiles, we applied a Kruskal–Wallis test [28]. The results were statistically significant ($H(3) = 33.5973$, $p < 0.001$), which means there was a statistical difference between *at least* two quartiles. To determine which quartile(s) differed, we applied Dunn's test [29], a nonparametric post hoc analysis (see Table 2a).

The results indicate that no "near quartile" pair showed any significant difference (Qn vs. Q$n+1$), but students *more* than one quartile apart had significantly different final grades. This suggests a threshold effect. There is little difference in final grades between a student who (say) attends only 15% versus 35% (or even 60% vs. 90%), but there *is* a significant difference when attendance rates differ drastically; for example, a student attending only 40% of the time compared to one attending 80% of the time.

Table 2: Dunn's Post-hoc Tests (p-values with Bonferroni correction; statistically significant values are in **bold**, $\alpha = 0.01$).

(a) Attendance vs. Grades

Group	z	p
Q1 vs Q2	1.489926	0.817462
Q1 vs Q3	4.072193	**0.000279**
Q1 vs Q4	5.212110	**0.000001**
Q2 vs Q3	2.445175	0.086869
Q2 vs Q4	3.703481	**0.001276**
Q3 vs Q4	1.513149	0.781451

(b) Grades vs. Attendance

Group	z	p
Q1 vs Q2	3.823268	**0.000790**
Q1 vs Q3	4.466788	**0.000048**
Q1 vs Q4	5.726069	**0.000001**
Q2 vs Q3	0.643520	1.000000
Q2 vs Q4	1.902801	0.342399
Q3 vs Q4	4.764156×10^5	1.000000

4.3.1 Threshold Effect

To further explore a possible threshold effect, we ran a logistic regression (maximum likelihood estimation) to estimate the attendance rate at which students have a specified predicted probability of earning a failing (D/F) grade. We initially ran the model for a threshold of 50%. It indicated that attendance significantly predicted the likelihood of earning a failing grade ($\beta = -0.040$, Odds Ratio = 0.961, $p < 0.001$). The model implied that an attendance rate of only 17.9% (95% CI $[3.9, 32.0]$) corresponded to a 50% threshold. That is, those who attended *at least* 18% of the time had a better than 50% chance of passing the course. We also ran the model for a threshold of 25% and found that the corresponding attendance rate was 45.4% (95% CI $[11.1, 79.6]$) with the same model coefficients. That is, those who attended less than 45% of the time had a 25% probability of failing the course.

A graph of the logistic curve and the two probability thresholds can be found in Figure 5. Both of these thresholds have interpretive value. The 50% threshold represents the point of maximum ambiguity and indicates that even moderate attendance (more than 18%) is associated with passing the course. The 25% threshold provides an early-warning indicator and identifies a greater number of students who may be at risk of failing. Policy-wise, instructors may be reassured to know that only extreme non-attendance is likely to lead to failing grades. At the same time, this suggests that students who do not attend half of the time early in the semester should be targeted for early interventions and more communication on the value of attendance.

4.4 Grades vs. Attendance

Finally, we conducted an exploratory data analysis of how attendance rates differed among low- and high-performing students. We performed the same Kruskal-Wallis and Dunn tests but split the data into quartiles based on final grades (Q1 are the lowest performing 25%, mostly Ds and Fs; Q4 is the top scoring 25%, mostly As), so each quartile had uniform size ($n = 39$). The Kruskal-Wallis results were statistically significant ($H(3) = 36.4890$, $p < 0.001$), indicating a statistical difference between *at least* two quartiles. Dunn's test results are in Table 2b.

We found that Q1 (the lowest-performing quartile) was statistically significantly different from the other groups with respect to attendance. However, there were no significant differences among the other quartiles. This means that the lowest-performing students differed in their attendance patterns from other students. When combined with the other data, the difference in Q1's attendance pattern

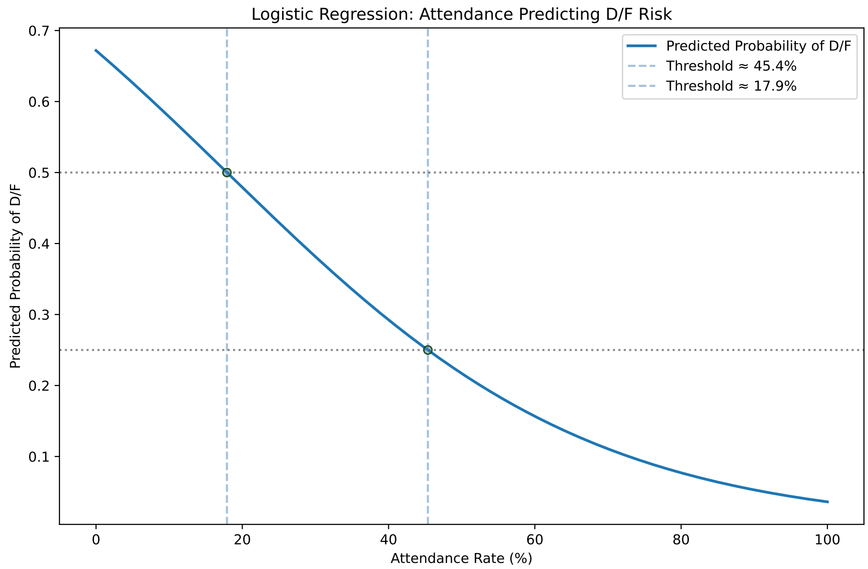

Figure 5: Logistic Regression Curve. Attendance rate as a predictor of earning a failing grade.

is clear: they did not attend regularly. The upper three quartiles showed no observable difference, likely an artifact of clustering. Q4 and Q3 were essentially all A students, while Q2 were mostly B students. This reinforces the threshold analysis: poor attendance is an indicator of overall lower achievement. However, the threshold is somewhat low. Once a student's attendance rate is high enough, attendance patterns become less predictive.

5 Discussion

Our first observation is that the policy had a positive effect on improving attendance. Prior to the policy, attendance rates were only 30–40% (stabilizing by mid-semester). With this policy about 70% had attendance rates of more than 50% while a majority had attendance rates of 70–100%. In the following sections we discuss the data further.

5.1 Attendance as a Predictor of Success

Our results indicate that there is a (rather low) threshold of attendance rate that is a good predictor of performance. In order to understand student attendance patterns better, we examined our scatter plot and identified several distinct groups (see Figure 6).

Group A ($n_A = 13$) corresponds to students who did not regularly attend and failed the course. Group D ($n_D = 20$) includes students who attended sporadically and also performed poorly. Together, these represent the lowest final grade quartile—the group whose attendance is a good predictor of performance. The median observed attendance rate of these two groups is about 37.5%. This is significantly higher than the regression model indicated (which is expected as the regression model is predictive while this data is descriptive after the fact). This may be a better threshold for instructors setting policies: students who attend less than one third of lab sessions are likely to fail the course. We also observe that students who regularly attended but still failed the course (lower right) *did not exist*!

Group B ($n_B = 14$) corresponds to students who also did not attend regularly but nevertheless

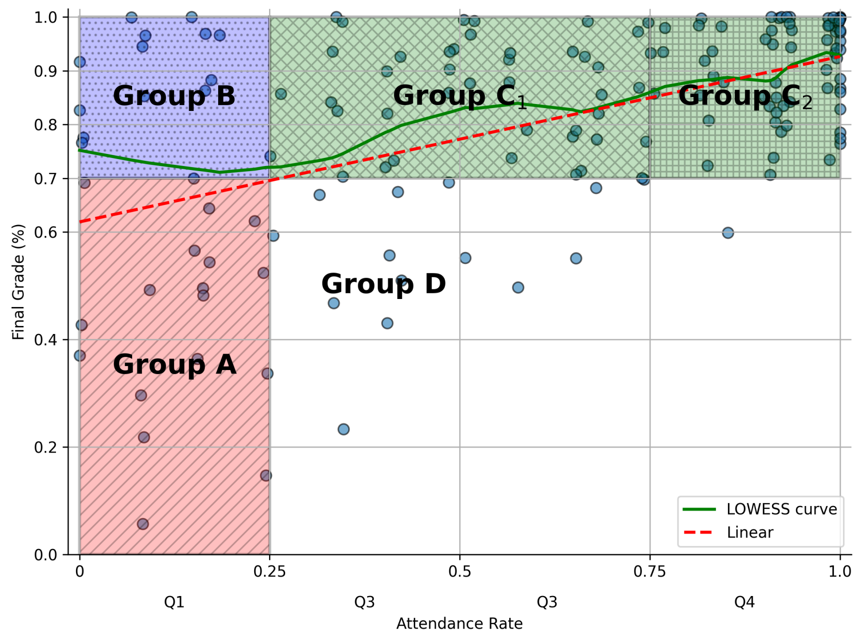

Figure 6: Identified Groups

passed (or did very well) in the course. This is the previously identified "independent" group of students who are able to formulate their own academic success strategies. Prior to this study, we estimated this group to be about 20–30% of students. However, the data show that this is not the case; this group makes up less than half our original estimate.

Group C was the largest ($n_C = 109$), so we broke it into two subgroups: C_1 ($n_{C_1} = 37$), representing those who succeeded in the course but only attended sporadically, and C_2 ($n_{C_2} = 72$), representing students who attended regularly *and* succeeded in the course. C_2 is the largest group and includes nearly half of all students. Though only weakly correlated—and not a strong predictor for high-performing students—this result gives a clear, actionable message that instructors can give to students: the most **typical** outcome for a student who succeeds is regular attendance, and the most typical outcome for students who regularly attend is success.

5.2 *Comparison With Prior Research*

Our CS1 data are consistent with our informal observations prior to this policy. We had previously identified three distinct groups: (1) those who never attend but do not need to, (2) those who do not attend regardless of any incentives or rewards, and (3) those who do attend but have varying outcomes. Prior to this data, we estimated the proportions of these groups to be roughly even (one-third each). However, the data shows that the actual proportions are more nuanced. The first group is actually only 15–20%—about half of what we initially thought. The second group *is not initially significant* at all. Only after a substantial portion of the assessment has been completed (about 35% of the course) does this group emerge and remain significant. This group includes students whose grade is no longer recoverable and who have essentially given up, as well as a smaller portion who are merely treading water.

Comparing with prior research, our own data confirms that attendance and academic success are moderate-to-strongly correlated. Our data also confirms that students will (generally) choose to attend when given the option. About 70% did so overall, either sporadically (C_1) or regularly (C_2, which was the majority).

5.3 Philosophy

Our basic philosophy with this policy is to prioritize *flexibility* and to provide students with autonomy. Students generally understand that attendance is important but tend to resist rigid policies, whether reward-based or punitive. Students believe they can formulate their own strategies for achieving academic success. This policy responds by placing the burden of demonstration on the student.

Traditional graded attendance is recorded on the day-of or through a proxy such as in-class activities or quizzes. In this sense, there is an immediate reward, regardless of whether the student has successfully learned the material. Our approach is fundamentally different: it provides a post hoc verification of whether the student successfully formulated their own approach to academic success. There is still the immediate reward (points) for attending or the consequence (initial zero) for not attending, but only after the student has *demonstrated* some degree of mastery do they have that consequence removed and earn the reward for their strategy. This gives students agency in their learning while making topic mastery the primary focus.

A traditional attendance policy risks being overly paternalistic, while the absence of any attendance policy represents the opposite extreme. Although college students are legally adults, many are still in a developmental transitional phase, particularly those enrolled in introductory courses. This policy attempts to balance student autonomy with structured accountability, offering flexibility while preserving mechanisms that help guide students toward success.

5.4 Benefits & Best Practices

Other than student success, the policy has had additional benefits. During labs, there are fewer students overall, which means that LA contact time and other resources are more readily available to those that choose to attend. We would also argue that it is more equitable and aligns with Universal Design [30] by re-centering responsibility and agency for learning with the student. The policy makes it far more flexible for non-traditional students and those who face irregular schedules (due to health, work, lack of reliable child care, etc.). It provides a consistent and fair logistical solution for students with excused absences for military service, athletics, student organizations, religious observances, etc.

The policy has minimal administrative overhead (after the initial investment in developing scripts). There are some best practices that we recommend based on our experience. First, the policy needs to be clearly and consistently communicated, not only in the syllabus but also in class early and reiterated throughout the semester. For example, at the end of each module, we report aggregate data to students on how many attended and how many earned points, etc. on our course message board. The scripts that update attendance points also generate a message to the student explaining why the points were updated (or not) and encourage attendance going forward if appropriate. We clearly articulate the purpose and benefits of attending by reporting historical data.

5.5 Limitations

Our study is not making substantial generalization claims as it is limited to one institution. Our study may also be subject to several other limitations. We excluded withdrawal students due to lack of complete data. However, withdrawals only accounted for less than 5% of the total enrollment and would not have affected the analysis substantially. Further, attendance points were awarded based on subjective LA judgments regarding student engagement. Finally, though we found significant correlations, this does not necessarily mean that attendance leads to success—correlation is not causation.

6 Conclusion & Future Work

We have presented our experience in developing our *outcome-based attendance policy* and have demonstrated that it was successful in improving attendance while providing students with a level of flexibility and autonomy. When given agency in how they approach their learning, *most* students will make the right choices.

There are several avenues for further research and additional data analysis. We did not consider trends of individual students over the course of the semester–only those of the class as a whole. Does a student stop attending because they have fallen below a certain grade threshold (and have given up)? If so, does that make attendance less correlated or less predictive? Can we characterize students by their trends in attendance (ex: some may initially not attend and then show up later, some may attend initially but then stop, and others may be less consistent)?

Another interesting analysis would involve gathering student perspectives on the policy. In course evaluations, several students expressed appreciation for the flexibility of the course and the attendance policy in particular. However, we did not specifically solicit feedback on the policy. Likewise, LAs also expressed similar positive attitudes, but we did not systematically gather their feedback. Future work would involve explicitly gathering and analyzing such feedback.

References

[1] M. Credé, S. G. Roch, and U. M. Kieszczynska, "Class attendance in college: A meta-analytic review of the relationship of class attendance with grades and student characteristics," *Review of Educational Research*, vol. 80, no. 2, pp. 272—295, 2010.

[2] M. Credé, S. G. Roch, and U. M. Kieszczynka, "Class attendance in college: A meta-analytic review of the relationship of class attendance with grades and student characteristics," *Review of Educational Research*, vol. 80, no. 2, pp. 272–295, 2010. [Online]. Available: https://doi.org/10.3102/0034654310362998

[3] J. Durfee, W. Loendorf, D. Richter, T. Geyer, and D. Munson, "A formal research study on correlating student attendance to student success," 06 2012.

[4] J.-F. J. Yao and T.-M. Chiang, "Correlation between class attendance and grade," *J. Comput. Sci. Coll.*, vol. 27, no. 2, p. 142–147, Dec. 2011.

[5] M. Urban-Lurain and D. Weinshank, "Attendance and outcomes in a large, collaborative learning, performance assessment course," in *Proceedings of the Annual Meeting of the American Educational Research Association (AERA)*, 2000.

[6] W. C. Hsu and S. W. Plunkett, "Attendance and grades in learning programming classes," in *Proceedings of the Australasian Computer Science Week Multiconference*, ser. ACSW '16. New York, NY, USA: Association for Computing Machinery, 2016. [Online]. Available: https://doi.org/10.1145/2843043.2843061

[7] M. Furini, O. Gaggi, V. Ghini, S. Mirri, and M. Montangero, "The university learning experience: A dual approach to understanding study habits and attendance patterns," in *Proceedings of the 2024 International Conference on Information Technology for Social Good*, ser. GoodIT '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 408–415. [Online]. Available: https://doi.org/10.1145/3677525.3678690

[8] M. Minnes, C. Alvarado, and L. Porter, "Lightweight techniques to support students in large classes," in *Proceedings of the 49th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 122–127. [Online]. Available: https://doi.org/10.1145/3159450.3159601

[9] V. Malhotra and D. Garcia, "Steering student behavior and performance toward success with mastery learning through policy optimization," in *SIGCSE Virtual 2024*, ser. SIGCSE Virtual 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 144–150. [Online]. Available: https://doi.org/10.1145/3649165.3690109

[10] Hanover Research, "Strategies for improving student attendance: Policies and practices," https://www.ksde.gov/Portals/0/CSAS/CSAS%20Home/Graduation%20and%20Schools%20of%20Choice/Strategies-for-Improving-Student-Attendance-Policies-and-Practices.pdf, Tech. Rep., July 2022.

[11] M. L. Van Blerkom, "Class attendance in undergraduate courses," *The Journal of Psychology*, vol. 126, pp. 487–494, 1992.

[12] M. H. Launius, "College student attendance: Attitudes and academic performance," *College Student Journal*, vol. 31, pp. 86–92, 1997.

[13] S. E. Gump, "The cost of cutting class: Attendance as a predictor of success," *College Teaching*, vol. 53, no. 1, pp. 21–26, 2005. [Online]. Available: https://doi.org/10.3200/CTCH.53.1.21-26

[14] T. Chenneville and C. Jordan, "Impact of attendance policies on course attendance among college students," *Journal of the Scholarship of Teaching and Learning*, vol. 8, no. 3, pp. 29–35, October 2008.

[15] V. Kassarnig, A. Bjerre-Nielsen, E. Mones, S. Lehmann, and D. D. Lassen, "Class attendance, peer similarity, and academic performance in a large field study," *PlOS One*, vol. 12, no. 11, 2017.

[16] T. M. Hancock, "Effects of mandatory attendance on student performance," *College Student Journal*, vol. 28, no. 3, pp. 326–329, 1994.

[17] K. St. Clair, "A case against compulsory class attendance policies in higher education," *Innovative Higher Education*, vol. 23, pp. 171—180, 1999.

[18] R. Moore, “Attendance: Are penalties more effective than rewards?” *Journal of Developmental Education*, vol. 29, January 2005.

[19] ——, “Attendance and performance: How important is it for students to attend class?” *Journal of College Science Teaching*, vol. 32, no. 6, 2003.

[20] J. Gerald and B. Brady, “Time to make your mandatory-attendance policy optional?” *The Chronicle of Higher Education*, January 2019. [Online]. Available: https://www.chronicle.com/article/Time-to-Make-Your/245455

[21] S. Cullen and D. Oppenheimer, “Choosing to learn: The importance of student autonomy in higher education,” *Science Advances*, vol. 10, no. 29, p. eado6759, 2024. [Online]. Available: https://www.science.org/doi/abs/10.1126/sciadv.ado6759

[22] R. M. Ryan and E. L. Deci, “Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being.” *American psychologist*, vol. 55, no. 1, p. 68, 2000.

[23] R. M. Ryan, J. Reeve, H. Kaplan, L. Matos, and S. H. Cheon, “Education as flourishing: Self-determination theory in schools as they are and as they might be,” *The Oxford Handbook of Self-Determination Theory*, 2023.

[24] M. Hertz, “What do “CS1” and “CS2” mean? investigating differences in the early courses,” in *Proceedings of the 41st ACM Technical Symposium on Computer Science Education*, ser. SIGCSE ’10. New York, NY, USA: Association for Computing Machinery, 2010, p. 199–203. [Online]. Available: https://doi.org/10.1145/1734263.1734335

[25] C. Gordon, R. Lysecky, and F. Vahid, “The rise of the zylab program auto-grader in introductory CS courses,” zyBooks, Tech. Rep., November 2020, updated May 2023.

[26] R. Liu, C. Zenke, C. Liu, A. Holmes, P. Thornton, and D. J. Malan, “Teaching cs50 with ai: Leveraging generative artificial intelligence in computer science education,” in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 2*, ser. SIGCSE 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 1927. [Online]. Available: https://doi.org/10.1145/3626253.3635427

[27] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1988.

[28] W. H. Kruskal and W. A. Wallis, “Use of ranks in one-criterion variance analysis,” *Journal of the American Statistical Association*, vol. 47, no. 260, pp. 583–621, 1952.

[29] O. J. Dunn, “Multiple comparisons using rank sums,” *Technometrics*, vol. 6, no. 3, pp. 241–252, 1964.

[30] A. Meyer, D. H. Rose, and D. Gordon, *Universal Design for Learning: Theory and Practice*, 1st ed. Wakefield, MA: CAST Professional Publishing, 2014.