The AIR Professional File

Fall 2025 Volume

Supporting quality data and decisions for higher education.

PREFACE

This issue of *Professional File* addresses two pressing challenges for IR professionals in higher education: declining survey response rates and the complexities of predicting student success.

In *A Rising Tide Lifts All Boats: Assessing and Addressing Survey Nonresponse*, Sanjeewa Karunaratne and John Nugent examine the persistent issue of survey nonresponse. Drawing on national surveys of institutional research directors and case studies, they explore whether response rates vary across subgroups and analyze patterns in the representation of different demographic groups. They hypothesize that as overall response rates increase, gaps between sample and population proportions in gender and race will narrow, resulting in a more representative sample. Their findings support this hypothesis: disparities shrink as response rates rise, particularly once overall participation approaches 50%. The authors also offer practical guidance for improving response rates through targeted outreach, thoughtful survey design, and coordinated campus engagement, as well as strategies for weighting survey data to strengthen representativeness.

Shifting from survey methodology to predictive analytics, *Identifying At-Risk Course Combinations for Freshman Students Using Market Basket Analysis*, by Fikrewold Bitew and Khoi To, investigates how institutional data can forecast student success. Using data from 7,466 first-time, full-time freshmen across five cohorts at The University of Texas at San Antonio, they analyzed 15,698 course enrollment records with failing grades. Their analysis identified seven course combinations that, when taken together, significantly increase the likelihood of student attrition—a perspective not possible when courses are examined individually. Identifying these high-risk course combinations enables targeted interventions: academic advisors can use the insights to help students select courses more strategically, while universities can implement supplemental instruction or tutoring programs for these specific combinations. Beyond academic outcomes, this analysis also has financial implications, as failing courses can extend time to graduation and increase costs for students, families, and public funding. The application of market basket analysis in this context highlights the value of data-driven approaches in higher education and lays the groundwork for further research.

Whether by improving survey participation or uncovering patterns that predict student outcomes, these studies underscore the importance of investing in both high-quality data and its thoughtful use to guide institutional strategy and enhance student success.

Iryna Muse
Editor

Editors

Iryna Muse
Editor
University of Alabama System

Leah Ewing Ross
Managing Editor
Association for Institutional Research

ISSN 2155-7535

Table of Contents

A Rising Tide Lifts All Boats: Assessing and Addressing Survey Nonresponse Results from Two National Surveys and Institutional Case Studies

Sanjeewa Karunaratne

John D. Nugent

About the Authors

Sanjeewa Karunaratne is Associate Director of Institutional Research and Planning at Connecticut College, New London, Connecticut. John D. Nugent is Director of Institutional Research and Planning at Connecticut College, New London, Connecticut.

Abstract

Survey nonresponse poses a significant challenge to the validity and reliability of research findings. In higher education, nonresponse warrants special attention because higher education institutions often make decisions based on empirical research. Nonresponse has been increasing in surveys of general populations, and survey researchers across various social science disciplines have observed a steady decline in survey participation over time. This trend has also impacted higher education research, as we can see in the reductions in response rates in national surveys, such as the Cooperative Institutional Research Program surveys and the National Survey of Student Engagement. We report data from two national surveys of directors of institutional research that we conducted, and we highlight demographic disparities in response rates. We demonstrate that improvement in response rates and representativeness reduces variations between sample and survey populations by gender and race. In the same way that "a rising tide lifts all boats," an improving survey response rate lifts the representativeness and thus the value of the results. We bolster our national survey findings with case studies of surveys of first-year and senior students at Connecticut College in New London, Connecticut. We find that disparities in gender and race are very low when the overall response rate reaches 50% or more of all students. We offer actionable solutions for researchers to enhance response rates and reduce nonresponse. Finally, in an institutional case study, we demonstrate how these actionable solutions significantly increased response rates and improved sample representativeness. Our key finding is that male and non-White participants have lower response rates overall, compared with female and White participants, but that, with concerted efforts within a college campus, male and non-White participants' response rates can be improved.

The AIR Professional File, Fall 2025
Article 180

https://doi.org/10.34315/apf1802025

1. INTRODUCTION

Survey research is a cornerstone of data collection in higher education. However, nonresponse, defined here as the failure to obtain information from selected participants, can introduce bias, reduce statistical power, and compromise the generalizability of findings to the study population. Despite its pervasive nature, researchers often underestimate or inadequately address nonresponse issues. Response rate issues in surveys at higher education institutions, where the target population is concentrated and can easily be contacted through marketing and other efforts, are particularly intriguing.

Sample bias occurs because we gather a data sample from the population, and the makeup of the sample may differ from—and be underrepresentative or overrepresentative of—the population in key demographics. To address this possible sample bias, it is crucial to understand the nature of nonresponse and its variations. Two key questions emerge: Who aren't we regularly hearing from in our surveys? Who are we often hearing from in our surveys? Identifying and addressing any gaps and understanding overrepresentation are fundamental to improving data quality and ensuring representative findings.

To explore these issues, we surveyed directors of institutional research at national liberal arts colleges and national universities regarding their survey practices and response rates. Data from liberal arts colleges and universities were combined for the analysis in this research paper in order to improve subgroup sample sizes. This methodological choice enhances the statistical power and precision of response-rate comparisons among demographic groups, such as gender and race.

In today's increasingly data-driven world, higher education institutions rely heavily on survey research to guide decision-making. Data collected from surveys influence curricular offerings, policies, resource allocation, physical space, and program development. However, when survey nonresponse leads to incomplete datasets, institutional leaders risk making decisions based on inaccurate, nonrepresentative, or overrepresentative samples. Addressing survey nonresponse ensures that institutional decisions will be grounded in high-quality data, thereby improving student experiences, instructional delivery, and administrative efficiency.

Survey nonresponse is a growing concern, particularly in an era of declining participation rates. Studies indicate that response rates in telephone and online surveys have dropped over the years. As a result, the literature on survey methods is full of recommendations and best practices to improve survey response rates (Dillman, 2000; Heberlein & Baumgartner, 1978). It is almost always better to spend resources on increasing the response rate than to spend resources on surveying a large group.

2. LITERATURE REVIEW

There are limited recent studies on survey response rates. Nonresponse has been increasing in surveys of the general population (De Leeuw & Heer, 2002; Smith, 1995; Steeh, 1981), and survey researchers across various social science disciplines, both in the United States and in other countries, have observed a steady decline in survey participation over time (Brick & Williams, 2013; National Research Council, 2013). This trend has also impacted higher education research. Dey (1997) highlighted a significant drop in response rates for the American Council on Education and the Cooperative Institutional Research Program (CIRP) senior follow-up surveys; those rates fell from 60% in the 1960s to just 21% in 1991.

2.1. Patterns in Nonresponse

Moreover, research focusing on student populations has identified certain demographic patterns in survey response rates. For instance, female students typically participate at higher rates than male students (Crawford et al., 2001; Dey, 1997; Hutchinson et al., 1987; National Center for Education Statistics, 2002; National Survey of Student Engagement [NSSE], 2003; Sax et al., 2003). Additionally, White students are generally more likely to respond to surveys compared to non-White students (Dey, 1997; National Center for Education Statistics, 2002), leading to potential representation gaps in unweighted data. Dey also found that academic factors are related to survey participation—high school GPA and students' self-assessed academic abilities were among the strongest predictors of survey participation. In their 1987 study, Hutchinson et al. investigated response bias in mail surveys targeting first-year college students. They compared the attitudes and demographic characteristics of students who responded to the survey with those who did not. The study found notable differences between respondents and nonrespondents, which indicates the presence of response bias. Similar to Dey, Hutchinson et al. found that late respondents had lower GPAs than either early respondents or the overall student population, suggesting that academic performance may influence the timing and likelihood of survey participation. In addition, Hutchinson et al. suggest that students with higher GPAs and greater confidence in their academic skills are more likely to respond to surveys.

Porter and Whitcomb (2005) analyze the impact of demographics, engagement, and personality on student survey nonresponse, providing a comprehensive understanding of factors influencing participation. Using the CIRP first-year student survey administered by the Higher Education Research Institute (HERI), Porter and Whitcomb found that female students and students with high GPAs are more likely to take part in surveys, consistent with Dey (1997) and Hutchinson et al. (1987), while students receiving financial aid are less likely to complete surveys. The addition of CIRP panel survey data provides several significant predictors of survey response to the model. For example, there is a positive relationship between social engagement and survey participation, with socially engaged students being more likely to take part in surveys than their less-engaged peers.

2.2. Response Rate and Sample Representativeness

Though variation exists in defining nonresponse bias, most view it as a function of the response rate and nonresponse effect; that effect is defined as how much respondents and nonrespondents differ on survey variables of interest (Keeter et al., 2000). In other words, low response rates may or may not lead to nonresponse bias, because answers to survey items might not differ substantially between respondents and nonrespondents. Similarly, Massey and Tourangeau (2013) argue that a high rate of nonresponse increases the potential for biased estimates, but does not necessarily bias an estimate.

Using NSSE data, Fosnacht et al. (2020) calculate the difference in means between the full and simulated samples of NSSE data using both the response rate and the respondent count. Then they correlate the standardized mean difference between the sample and the population mean of the metrics, such as the level of academic challenge and active and collaborative learning, computed in the NSSE questionnaire. Fosnacht et al. found that, even with response rates as low as 5% to 10%, survey results still provide reliable estimates, particularly when sample sizes exceed 500 respondents. Thus, lower response rates do not necessarily introduce significant bias, provided that the sample is diverse and representative of the population. Their study challenges the traditional assumption that a higher response rate always ensures better data quality. Fosnacht et al. suggest that researchers should focus on survey design and representativeness rather than strictly aiming for higher response rates.

2.3. Procedural Considerations

Research by Höfler et al. (2005) highlights the effectiveness of weighting techniques to adjust for nonresponse and dropout, providing a methodological approach to mitigating these issues. They argue that, by using these weights, researchers can better approximate the characteristics of the original sample, thereby enhancing the generalizability of the results. Furthermore, Little (1993) discusses the importance of post-stratification data weighting in correcting nonresponse, and emphasizes how model-based adjustments can enhance the representativeness of survey results. Using data from the CIRP follow-up survey, Dey (1997) applied a weighting procedure described by Astin and Molm (1972) to correct for nonresponse. The findings indicated that this weighting method was highly effective in reducing underrepresentativeness in univariate distributions, ensuring that the survey results more accurately reflected the target population. However, the advantage of using this weighting procedure was less clear when adjusting correlation and regression analyses.

Nair et al. (2008) explore declining survey response rates in higher education and argue that student engagement plays a critical role in improving participation. The authors examine how institutional strategies, survey design, and student perceptions influence response rates. Nair et al. suggest that, by fostering stronger connections between institutions and students, response rates can improve, leading to data that are more representative of the overall student population. They also highlight the role of clear communication regarding the significance of survey participation and the implementation of targeted engagement strategies to encourage student involvement.

Porter et al. (2004a and 2004b) and Coates (2006) suggest that students are more likely to participate in evaluation surveys if they believe that their feedback makes a meaningful contribution. This means that students not only believe that their voice was crucial in providing valuable information, but also that they believe that the institution will act on their feedback.

3. RESEARCH QUESTION: DOES A RISING TIDE LIFT ALL BOATS?

We explore whether there are subgroup differences in response rates, and whether improving the response rates across the board leads disparities between sample and population proportions to converge or diverge. The literature review has shown that student engagement is a factor in survey response; this means that, if students are engaged, we will see a higher response rate in surveys. If, for example, engaged students are disproportionately represented by certain demographics (e.g., White or female students), then our efforts to improve overall survey response rates might backfire. However, we hypothesize that, as we improve overall response rates, the disparities between sample and population proportions in gender and race will converge and the sample will become more representative of its underlying population.

4. METHODOLOGY

Before initiating this study, we conducted a limited literature review to find out what had already been done regarding response rates in higher education. We were primarily inspired by research conducted by Porter and Whitcomb in the early 2000s (Porter & Whitcomb, 2005). Armed with the information from our limited literature review, in January 2023 we conducted a series of personal interviews with several colleagues and experts in higher education and the private sector to evaluate the feasibility of our study. We asked these interviewees seven questions over the phone and followed up with one more question via email (Appendix 1). The input from these personal interviews was incorporated into the questionnaire. Once the questionnaire had been drafted, it was piloted with a sample from the same colleagues. Our sample of colleagues included Alexander C. Yin, Jeffrey E. Luoma, Michael Whitcomb, Patrick Glaser, Susan Canon, and Viola Simpson. We applied for and received Connecticut College's Institutional Review Board (IRB) approval for our liberal arts survey, and we applied for and received an IRB exemption from further review under the Code of Federal Regulations for the national universities survey. The Connecticut College IRB number assigned to this study is 2022.23.20.

Identical online surveys were administered to directors of institutional offices at liberal arts colleges (April 2023 through May 2023) and national universities (May 2024 through June 2024) via email. The survey asked about general survey practices and requested disaggregated response rates on a major survey they had administered in the preceding academic year. The survey questionnaire is available in Appendix 2.

An invitation to take the survey and three reminders were sent to individuals on the liberal arts contact list. After the final reminder, a postcard embedded with a QR code for the survey was mailed to non-completers of the liberal arts survey. For the national universities survey, an invitation and five reminders were emailed. A $15 Amazon gift card was offered to all completers in both surveys. We gratefully acknowledge the Connecticut College Center for the Critical Study of Race and Ethnicity for awarding us two grants to cover the cost of these survey incentives.

We used the 2023 edition of the U.S. News & World Report dataset and its classification to obtain the liberal arts contact list for the spring 2023 survey, and the 2024 edition of the same dataset to obtain the national universities contact list for the 2024 survey. The U.S. News & World Report rankings of institutional data were merged with internal databases to include the contact information of directors of institutional research offices. We manually searched online and updated missing or outdated contact information. The liberal arts contact list had a coverage of 92.0% and the national universities contact list had a coverage of 86.7%.

The surveys yielded 70 responses from 185 liberal arts colleges (38.9%) and 89 responses from 361 national universities (24.7%), for a total of 159 respondents. The surveys underrepresented liberal arts colleges and national universities that are ranked 151 or below in the U.S. News & World Report rankings, as seen in Figures 1 and 2.

Figure 1. Liberal Arts Gap Between Contact List and Finished Respondents (%)

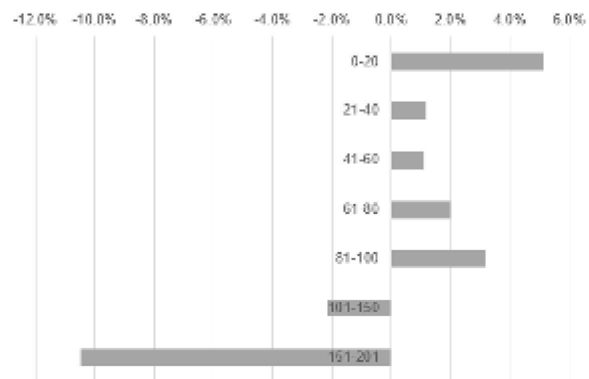

Figure 2. National Universities Gap Between Contact List and Finished Respondents (%)

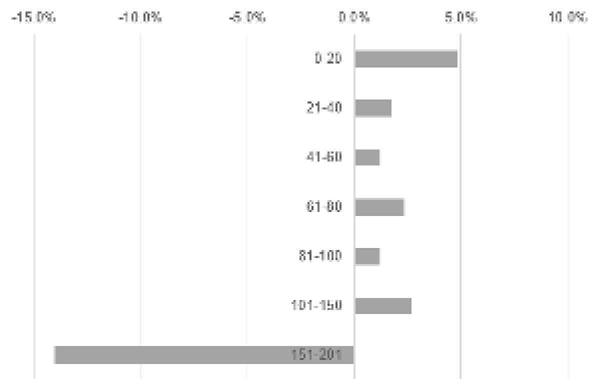

5. SURVEY FINDINGS

Comparing the response rate from male student and female student differentials between population and sample (the mean male response rate difference between population and sample when paired with the mean female response rate difference between population and sample), as reported by our survey respondents, we found a statistically significant difference with a strong effect size (Cohen's d = 0.727) using a paired sample t-test.

Fewer men and more women are represented in the sample than in the population. (Figure 3 shows more men represented in the population than in the sample, and Figure 4 shows fewer women represented in the population than in the sample as a proportion.) In addition, both figures show response rate differentials converging around a 50% overall response rate, indicating that the disparities in gender and race shrink as response rates increase beyond 50%.

Figure 3. Differentials by Overall Response Rate: Male Respondents

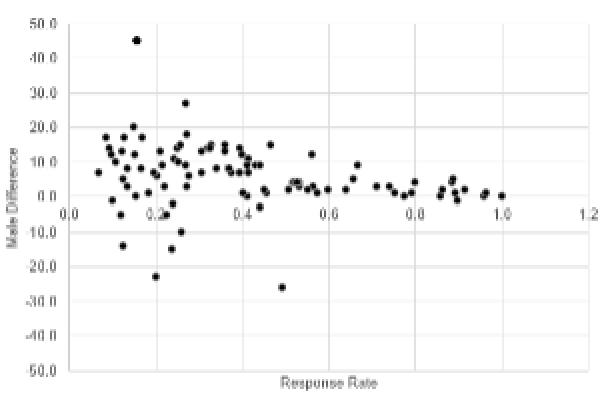

Figure 4. Differentials by Overall Response Rate: Female Respondents

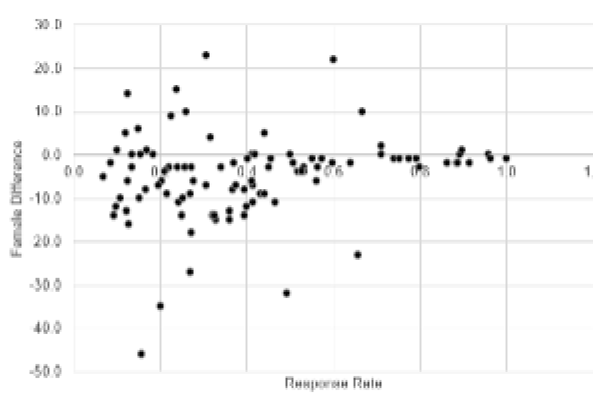

Figure 4a. Paired Samples *T*-Test

		Paired Differences							Significance	
					the Difference					
		Mean	Std. Deviation	Std. Error Mean	Lower	Upper	t	df	One-Sided p	Two-Sided p
Pair 1	Male_Diff - Female_Diff	11.26316	15.49482	1.58974	8.10670	14.41961	7.085	94	0.000	0.000

Comparing White and Hispanic/Latino response rate differentials between population and sample (the mean White response rate difference between population and sample when paired with the mean Hispanic/Latino response rate difference between population and sample), as reported by the respondents, we found a statistically significant difference with a weak effect size (Cohen's *d* = –0.248) using a paired samples *t*-test. In other words, fewer Hispanic/Latino respondents and more White respondents are represented in the sample than in the population. (Figure 5 shows more Hispanic/Latino respondents represented in the population than in the sample, and Figure 6 shows fewer White respondents represented in the population than in the sample as a proportion.) In addition, both figures show response rate differentials converging around a 50% overall response rate, except for a few outliers, indicating that the disparities in Hispanic/Latino and White respondents also shrink as response rates increase beyond 50%.

Figure 5. Hispanic/Latino Differentials by Overall Response Rate

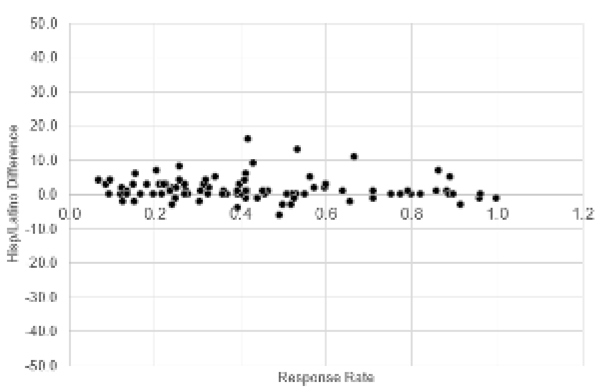

Figure 6. White Differentials by Overall Response Rate

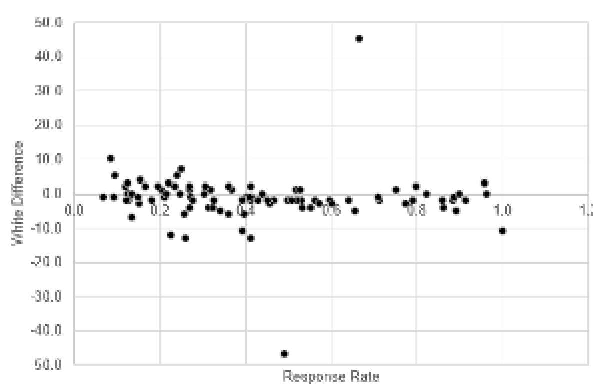

Figure 6a. Paired Samples *T*-Test

		Paired Differences							Significance	
					the Difference					
		Mean	Std. Deviation	Std. Error Mean	Lower	Upper	t	df	One-Sided p	Two-Sided p
Pair 1	White_Diff - Hisp_Latino_Diff	-2.28571	9.22952	0.96752	-4.20786	-0.36357	-2.362	90	0.010	0.020

Comparing White and Black response rate differentials between population and sample, as reported by the respondents (the mean White response rate difference between population and sample when paired with the mean Black response rate difference between population and sample), we found a statistically significant difference with a weak effect size (Cohen's d = –0.210) using a paired sample t-test. In other words, fewer Black respondents and more White respondents are represented in the sample than in the population. (Figure 7 shows more Black respondents represented in the population than in the sample, and Figure 6 shows fewer White respondents represented in the population than in the sample as a proportion.) In addition, Figures 7 and 6 show response rate differentials converging around a 50% overall response rate, except for a few outliers; this indicates that the disparities in Black and White respondents' response rates diminish thereafter.

Figure 7. Black Differentials by Overall Response Rate

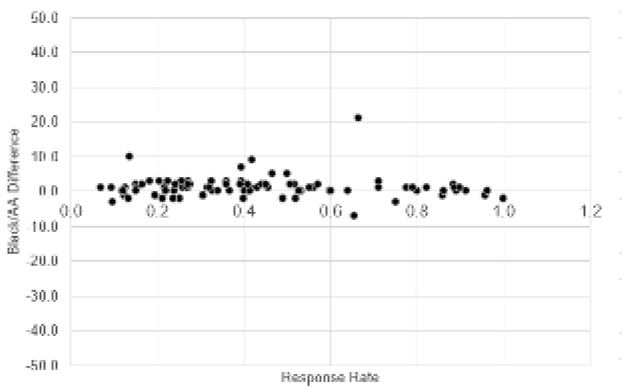

Figure 7a. Paired Samples *T*-Test

		Paired Differences							Significance	
					the Difference					
		Mean	Std. Deviation	Std. Error Mean	Lower	Upper	t	df	One-Sided p	Two-Sided p
Pair 1	White_Diff - Black_AA_Diff	-1.93407	9.20363	0.96480	-3.85081	-0.01732	-2.005	90	0.024	0.048

6. CASE STUDIES CONDUCTED AT CONNECTICUT COLLEGE

The major finding of our research is that increasing survey response rates leads to greater representativeness of the respondent pool as compared to the population. This was not a foregone conclusion, since it is possible that institutional researchers' efforts to improve response rates could simply increase the response rates of one or a few student subgroups rather than increasing them across the board. For example, female students tend to respond to surveys at higher rates than male students, and efforts to improve response rates could potentially simply run up this differential even higher. However, we find that this was not the case; as shown in Figures 8, 9, 10, and 11, we see that, as response rates climb, the gender gap in response rates drops between the sample and population.

Figures 8, 9, 10, and 11 pertain to Connecticut College case studies we conducted when we monitored survey response rates when administering the NSSE, our annual first-year and senior surveys in 2024, and our senior survey in 2025. (Note that these figures contain unweighted percentages.)

Figure 8. Male % Difference: Sample and Population

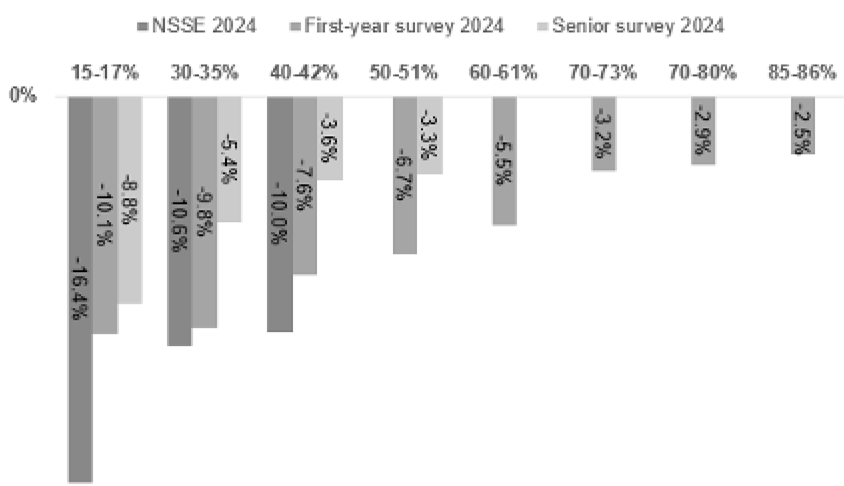

Figure 9. NSSE 2024 Male/Non-White % Difference: Sample and Population

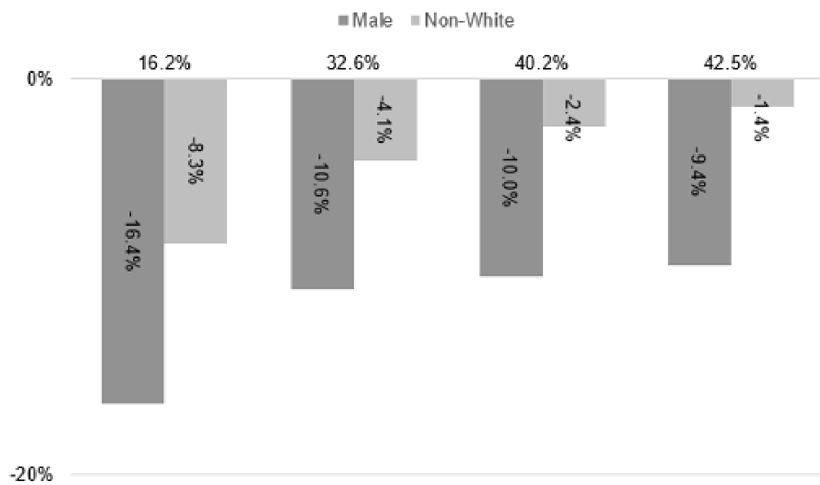

Figure 10. Senior Survey 2024 Male/Non-White % Difference: Sample and Population

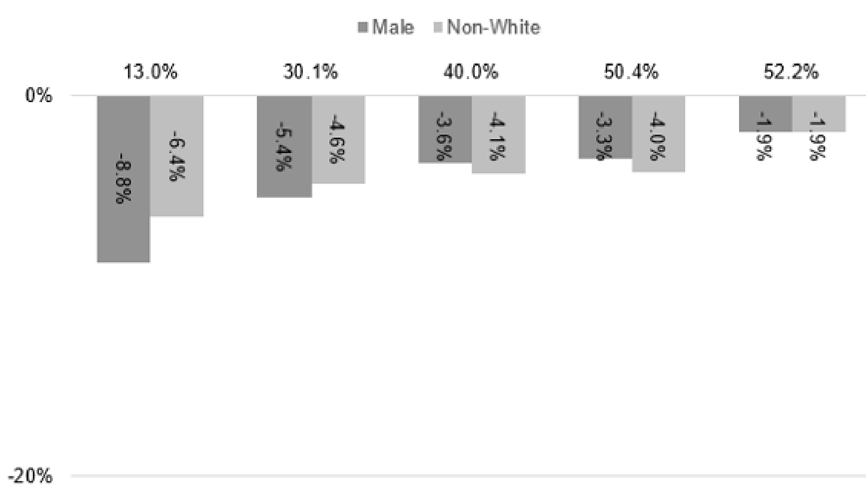

Figure 11. Senior Survey 2025 Male/Non-White % Difference: Sample and Population

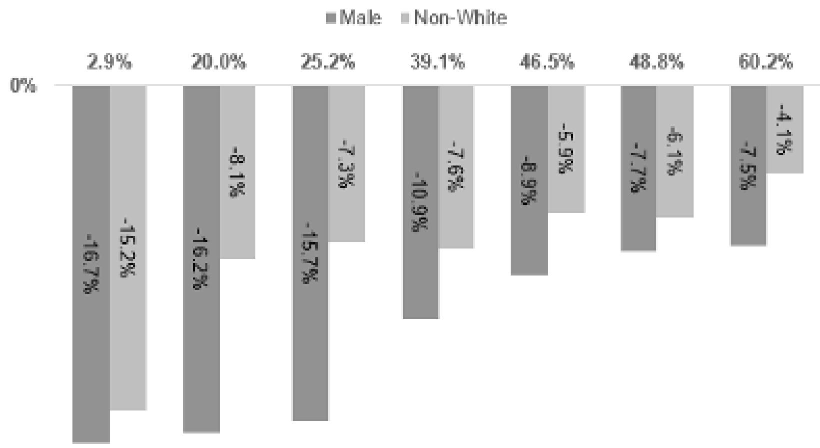

As we can observe in the four figures above, at higher response rates, especially when the rate surpasses 50%, the subgroup gap between sample and population declines to manageable levels, less than or closer to 5%, except for the male students. Thus, there may be no need for a researcher to have additional weighting procedures in place to correct this imbalance. How, then, can a researcher achieve a high response rate?

7. ACTIONABLE SOLUTIONS FOR ADDRESSING NONRESPONSE

In our literature review, in-person interviews, the two in-house national surveys of directors of institutional research conducted in 2023 and 2024, and in our institutional case studies, we gathered insights about the actionable solutions that institutional research offices can deploy when conducting surveys on campus. A summary of these strategies follows.

7.1. Survey Administration

- Administering major national surveys such as NSSE, CIRP, Noel-Levitz Student Satisfaction Inventory, and Consortium on Financing Higher Education (COFHE) and Higher Education Data Sharing (HEDS) consortium surveys on a rotating basis
- Managing survey scheduling through coordination with other offices on campus to avoid overlap
- Reducing the number of internal surveys overall
- Targeting outreach to student organizations, residential assistants, and athletics staff to get their help with recruiting respondents
- Monitoring response rates and using adaptive strategies during the survey administration period

7.2. Survey Design Improvements

- Making surveys shorter and reducing completion time
- Sending out survey prenotifications

- Communicating survey best practices to on-campus researchers
- Consolidating similar internal surveys
- Addressing participants by their names in email invitations and reminders
- Asking people in respected/high-profile roles on campus—college leaders, faculty, student support office staff, advisors, coaches, etc.—to announce the survey and encourage its completion

7.3. Approval Process

- Implementing survey preapproval processes, such as IRB approval to prioritize in-house surveys
- Limiting survey software access based on institutional needs and priorities to avoid survey proliferation

7.4. Survey Distribution Methods

- Administering surveys in person, such as during orientation or at tables set up in the student center
- Posting survey links in student learning management systems such as Moodle to increase visibility
- Posting on the institution's social media feeds, advertising in student newspaper, putting announcements on campus TV screens or information kiosks, and posting flyers and table tents with QR codes for accessing the survey in conspicuous campus areas such as dining halls
- Considering the use of text messaging to students with survey reminders or shorter text surveys with fewer questions
- Implementing sampling techniques to ensure representative survey responses and to mitigate survey fatigue
- Using oversampling methods to capture underrepresented populations or targeted outreach to specific student groups
- Conducting pilot studies to assess the effectiveness of the survey or of each question before full deployment
- Explaining how the data will be used; pointing to concrete changes that have happened as a result of past survey responses; citing survey results in reports; writing articles about survey results for the student newspaper

7.5. Practical Application of These Strategies

Connecticut College's annual senior survey was administered with a mix of online and in-person approaches until 2018. No special efforts were undertaken between 2019 and 2022. Some of the actionable solutions mentioned above were adopted in 2023, and we saw a 23.7% gain in response rate. That rate dropped by 7.9 percentage points in 2024, yet still passed the 50% threshold. In 2025 the response rate bounced back to 60.2%.

Figure 12. Senior Survey Response Rate, 2014 through 2025

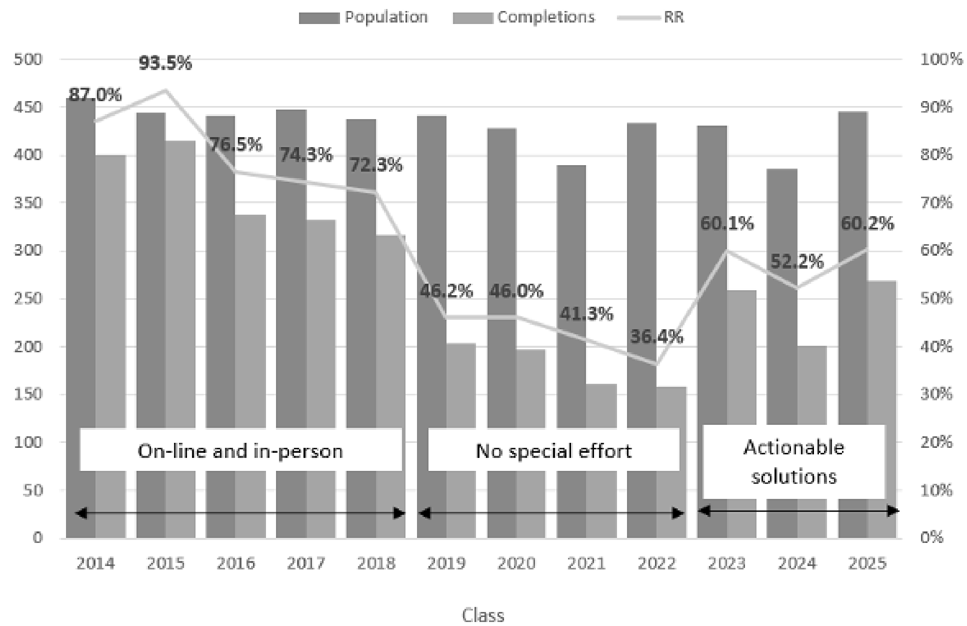

7.6. Comparison of 2022 and 2025 Senior Survey Respondents by Gender and Race

The improvement of our overall response rate by 23.8 percentage points resulted in a 3.4 percentage-point gain in response rates for male respondents and a 3.3 percentage-point gain in response rates for non-White respondents. Although the gain among male students seems low compared with the gain in the overall response rate, enticing male students to participate in surveys is a challenge.

In our surveys of institutional research directors, when we asked respondents to name one student subgroup they think is hard to reach, out of 109 respondents, 34 (31.2%) named male students. A more targeted approach, such as outreach to athletics or concentrated marketing efforts to male students, may be recommended to improve male students' participation in surveys.

Figure 13. Senior Survey 2022, Senior Survey 2025

	Senior Survey 2022	Senior Survey 2025
Overall Response Rate	36.4%	60.2
Male	-10.9%	-7.5%
Non-White	-7.4%	-4.1%

Three of the actionable solutions—incentives, sampling, and data weighting—require greater attention and are described in detail below, along with survey findings. These are high-impact practices in survey research. Incentives generate a cost for the institutional research office, even though they are effective in improving the response rate. Sampling and data weighting require preparation and technical expertise in the institutional research office and political will at the institution to adopt them.

8. INCENTIVES

Survey researchers widely use incentives to encourage survey participation. In higher education, the type and value of the incentives differ from institution to institution. Some institutions have used creative incentives to improve survey participation. In our literature review, in-person interviews, and national surveys, we gathered insights about the use of incentives when conducting surveys on campus. Notable suggestions from respondents included the following:

- Offer well-chosen incentives (one big prize, small prizes to every completer, nonmonetary prizes, a prize each week, early respondent prizes, etc.).
- Offer cash or gift card incentives, which tend to be more effective than nonmonetary incentives.
- Offer prepaid incentives, which may yield higher response rates than post-survey ones.

Our survey of institutional research directors included an open-ended question asking respondents whether they had offered an incentive for the survey they reported on, and, if so, the specific type of incentive. However, because this question was open ended, we were unable to systematically assess preferences between offering smaller incentives to a larger pool of respondents versus larger incentives to a smaller group, or any other specific preference that might be better than the other.

Close to a majority of survey participants (49.7%) have offered an incentive. Among those who offered an incentive, two-thirds (67.1%) reported offering monetary gift cards. This question was followed by another item assessing respondents' opinions about the impact of their incentives. Porter and Whitcomb (2003) found that 81.3% of researchers who used incentives perceived at least somewhat increased response rates. Consistent with these findings, we found that a large portion (70.9%) of the respondents who offered an incentive believe that incentives have a moderately or slightly positive impact on the response rate of their surveys.

Figure 14. What Impact Do You Think This Incentive May Have Had on Your Overall Response Rate?

	N	%
Moderately or slightly positive	56	70.9
Neither positive nor negative	15	19.0%
Extremely positive	8	10.1%
Total	79	100.0%

9. SAMPLING

The objective of surveys is to make inferences about a population from the information contained in the sample that has been selected from that population. The inferences take the form of population mean or population proportion (Scheaffer et al., 2006). Sampling is a crucial aspect of survey methodology, and influences the representativeness and accuracy of findings. While it is common, particularly at smaller institutions, to send survey invitations to the entire student body, this may not be necessary or desirable. Proper sampling techniques to yield an invitation list of, say, half the student population can mitigate nonresponse and improve data quality. Strategies such as stratified sampling help ensure that diverse populations are adequately represented in survey results, reducing variability between the survey sample and the underlying population, and thereby effectively reducing sampling error (Bock et al., 2018). Sampling on college campuses can also reduce survey fatigue, meaning that, if researchers contact samples of students rather than the entire population, the probability of repeatedly reaching any individual student declines. Sampling helps optimize the use of incentives by increasing the probability of success of getting an incentive in a random drawing. Since researchers have more leverage to follow up with a sample of students than with the entire student population, it might encourage harder-to-reach students to participate in the survey. If a researcher communicates effectively that a student has been "selected" to be included in the sample, it might produce an extra motivation for the student to participate in the survey.

In their 2004 studies, Porter et al. (2004a, 2004b) explored the phenomenon of survey fatigue among college students, focusing on how administering multiple surveys within a single academic year affects response rates. The researchers conducted two experiments at a selective liberal arts college to assess the impact of multiple survey requests on student participation and found that administering multiple surveys within a single academic year significantly reduces response rates due to survey fatigue. They observed that closely timed surveys and those perceived as less relevant had the most substantial drop in participation. However, a small group of students consistently responded, suggesting that some individuals are less affected by fatigue.

As researchers in higher education experience this phenomenon today, they may want to consider sampling smaller groups of students when administering surveys. For example, if the aim is to collect 100 responses for a reliable estimate and your population is 2,000, we propose sampling 300 students, assuming a 35% response rate. The remaining 1,700 unsampled students could be used for other research projects, thus reducing survey fatigue and potentially increasing response rates for all surveys being administered on campus.

9.1. Survey Findings About Sampling

Although there seems to be great potential for the use of sampling in higher education, our two national surveys found that 47.8% of respondents rarely or never used sampling at their institutions. As we have observed in our practice, there can be political considerations of using samples in college campus surveys, depending on the objective of the study.

9.2. Political Considerations in Sampling

Despite being a well-established statistical technique, sampling remains underused and is often viewed with skepticism in higher education. Higher education's strong commitment to equity and inclusion presents a unique challenge when justifying the exclusion of individuals from data collection, even when probability-based sampling methods ensure that every member of the population has a fair and known chance of selection. Given the politically engaged nature of higher education institutions, the concerns about fairness, representation, and potential biases in sampling methodologies have led many colleges and universities to favor full-population surveys over scientifically rigorous sampling strategies.

One major factor contributing to this skepticism is the perception that political influences may shape how samples are drawn, which groups are surveyed, and which perspectives are either emphasized or underrepresented. In an environment where diversity, equity, and inclusion are central values, stakeholders—ranging from students and faculty to administrators—often question whether sampling techniques can truly provide an accurate and fair representation of the entire population. This concern has resulted in a preference for comprehensive surveys that include all members of the student body, faculty, or staff, ensuring that no voices are inadvertently excluded. However, as noted above, this approach comes with drawbacks, including survey fatigue, reduced response rates, and inefficient use of resources.

To address these concerns, institutions must prioritize transparency, fairness, and methodological rigor in their sampling strategies. Clearly communicating how samples are selected, ensuring randomness in participant selection, and using stratified sampling techniques to reflect the diversity of the population can enhance credibility and trust in the results. Moreover, establishing independent oversight committees to review sampling methodologies can further mitigate concerns about bias.

In fact, when researchers send out a survey to students, the data they gather is *always* from a sample of the student body, and typically is influenced by self-selection rather than by inclusion in a scientific sample, as discussed above. By embracing well-designed sampling methods, higher education institutions can strike a balance between efficiency and inclusivity. When properly implemented, sampling not only preserves statistical integrity, but also upholds the core values of fairness and representation in decision-making processes.

10. POST-STRATIFICATION DATA WEIGHTING

When survey results appear not to be fully representative of the population, researchers can use weighting techniques to bring the results into closer alignment with what they would have been with a more robust response rate. Post-stratification data weighting is a vital statistical technique that enhances the accuracy and representativeness of survey results by adjusting for demographic imbalances that result from nonresponse. By applying weighting adjustments, researchers ensure that underrepresented groups are properly accounted for, thereby improving the reliability and validity of survey findings.

Data weighting is a widely used and essential statistical technique in election polling. The University of Connecticut's former polling center, the Center for Survey Research and Analysis (CSRA), weighted its polling data by gender, age, level of education, and race. As a policy, CSRA refrained from weighting data by party affiliation (CSRA, 2006). Even though data weighting can be a complex process, in the absence of data weighting, pollsters might not be able to predict election outcomes within the margin of error (Enten, 2024).

10.1. Survey Findings About Data Weighting

Despite its potential benefits, data weighting remains underused in higher education, with 59.7% of respondents reporting in our survey that they rarely or never apply weighting to their survey data.

10.2. Political Considerations in Data Weighting

Similar to when applying survey sampling techniques, data weighting may carry a stigma in a higher education institution, posing an appearance of data manipulation if information about the weighting procedure is somewhat lacking, or if it is overly complex.

There is nothing unethical about weighting data as long as such weighting is done objectively and for good reasons, and it is fully disclosed in the methodology section of reports and presentations. The weighted results change the number of survey respondents because their objective is to reduce imbalance in the final sample and improve the overall quality of the results. Above all, weighting is not a technique for repairing any deficiencies in the sample. If it is used wisely, weighting can improve the accuracy of sample estimates (Dorofeev & Grant, 2006). The researcher must first strive to get a representative sample with a reasonably good response rate, then apply weights to adjust minor imbalances in the sample. Transparency and consistency in the application of weighting techniques, along with clear explanations in results reporting, are essential for fostering trust and acceptance of weighted survey data.

One of the most widely used surveys in higher education is NSSE, administered by the Center for Postsecondary Research at the School of Education at Indiana University–Bloomington. The NSSE researchers provide a clear and concise explanation of their weighting methodology. The final Statistical Package for the Social Sciences (SPSS) dataset produced for institutions includes two weight variables: Weight 2, adjusted by institutional size for comparisons between institutions (to ensure proportional representation of institutions within selected peer groups), and Weight 1, adjusted by enrollment status (part time/full time) and gender. When an institution downloads its NSSE data, it must first apply Weight 1 to ensure accurate year-to-year comparisons of its own data (NSSE, n.d.).

11. LIMITATIONS OF THIS STUDY

The target population for our survey was directors of institutional research, a hard-to-reach yet naturally responsive group. As a result of the group's responsiveness, the survey achieved a reasonably strong response rate. Some of the survey's key questions were highly technical, requiring respondents to compile data on gender and race distributions within survey samples and corresponding target populations, as well as overall response rates from surveys conducted in the preceding academic year. Our study relies on these self-reported data.

To accommodate institutional differences, respondents were asked to reference data from a national survey that they had conducted in the preceding academic year. These surveys include NSSE, CIRP, and Noel-Levitz Student Satisfaction Inventory surveys, as well as surveys from the COFHE and the HEDS consortium. However, our study does not account for variations among these surveys or their timing. Additionally, institutional research directors use different techniques to improve response rates, such as mandatory participation, in-person administration, and incentives. The data presented in this paper do not adjust for variations in survey administration methods across institutions.

Moreover, engagement patterns identified by previous studies may be unique to each institution and could influence response rates. For example, institutions with an engaged student body may naturally yield higher response rates and vice versa, a factor that is not accounted for in this study. Finally, the data do not incorporate institutional characteristics such as graduation or retention rates. These characteristics may further influence response rate patterns.

The response rate readings for our internal surveys (NSSE, first-year, and senior) were taken only in limited instances during the assessment of survey progress. More-frequent response rate readings would have provided consistent ranges to showcase the decline in subgroup response rates (x-axis).

The marketing materials in our internal surveys included a QR code, allowing respondents to participate anonymously. However, we requested an email address for the drawing of the incentive in an effort to validate the responses and guard against ballot stuffing. Though we were able to find out gender and race information for most of our QR respondents, some respondents remained anonymous.

12. DISCUSSION

Research finds that socially engaged students are more likely to respond to surveys in higher education. It is also likely that the same group of engaged students is responding to most of the surveys distributed on a campus. In addition to taxing the time of the same set of socially engaged students, declining response rates in surveys present a significant challenge, requiring innovative strategies and a willingness to step beyond traditional methods to mitigate the impact. A combination of effort, creativity, collaboration, and strategic timing can enhance response rates. We found that, when response rates surpass 50%, the marginal differences in subgroup rates decline, and samples become more representative of their target population in gender and race.

The act of surveying itself can also serve as a valuable tool for student engagement. Researchers can involve students in piloting surveys, reviewing questionnaires, suggesting attractive incentives, or participating in focus groups, thus fostering a sense of inclusion and ownership. Additionally, engaging faculty, staff, and students in discussions about research findings and sharing insights through student newspapers or campus events can generate more interest and participation in future surveys.

To further encourage participation, researchers can offer innovative and meaningful incentives, such as a refund of the cap and gown purchase price for graduating seniors, course credits, parking passes, game tickets, a better position in a housing lottery, or early course registration, thus ensuring that students feel that their time has been appreciated and valued. And, by using technical strategies—such as sampling, monitoring response rates in real time, and adjusting outreach efforts accordingly—

researchers can improve participation and reduce survey fatigue.

Once all feasible methods to increase response rates have been exhausted, researchers can apply post-stratification weighting to make marginal adjustments, ensuring that the final sample accurately represents the broader student population based on key demographic characteristics. Through a thoughtful and multifaceted approach, researchers can enhance survey effectiveness while strengthening student engagement in the research process.

Critical components of our approach are highlighting underrepresented groups, developing strategies to improve their participation, and emphasizing the role of sampling and post-stratification data weighting based on key demographics such as gender and race. These techniques serve as solutions for addressing nonresponse and overrepresentation, reducing survey fatigue, and improving the overall efficiency of data collection in higher education. By ensuring that survey samples more accurately reflect the broader student population, institutions can derive more-reliable insights while minimizing the burden of excessive surveying.

Ultimately, our work aims to refine survey methodologies in higher education, and to ensure that data collection efforts are both equitable and methodologically sound. We provide a framework for institutions to enhance survey representation, improve response rates, and make data-driven decisions with greater confidence.

APPENDIX 1. IN-PERSON INTERVIEW QUESTIONNAIRE

1| In your experience, what type of student populations may have historical low response rates?

2| What type of community engagements such as posting flyers, outreach to athletes, faculty, student organizations etc., if any, have you undertaken to improve response rate in student surveys?

3| Have you observed an increase in response rate due to these efforts?

4| Do you typically give incentives to participate in a survey? If so, what kind of incentives?

5| Have you used a sample of students rather than the population, performed data weighting, or used any other technique to reduce nonresponse?

6| There can be a number of surveys conducted at a college campus in a given year. In a typical year, how many student surveys do you conduct?

7| Have you done things like limit number of student surveys by having an approval process like IRB, cut down or consolidate surveys, or anything else to reduce survey fatigue?

Follow up via email.

The study we are about to commence is attempting to gauge and propose best practices to reduce nonresponse in undergraduate student surveys. If you were to conduct this study, what are the three questions you may be asking in your survey?

APPENDIX 2. ASSESSING AND ADDRESSING SURVEY NONRESPONSE IN NATIONAL UNIVERSITIES

Start of Block: Block 1

IQ1 Dear ${m://FirstName},

The Office of Institutional Research and Planning at Connecticut College is conducting a national study on assessing and addressing nonresponse in undergraduate surveys. As a part of the study, we are contacting institutional research offices at national universities to participate in a brief survey.

The survey will ask for details about your administration of a national survey of undergraduates in the most recent academic year, 2022-2023. Please have those survey materials available before you begin our survey so you can answer all the questions about the nature of your respondents.

The survey will take about 10 minutes of your time. It has been approved by the Connecticut College Institutional Review Board. Your participation in the survey is voluntary, and your responses will be kept confidential. Any contact information collected will be removed from the data we analyze and report on. You must be 18 years or older to participate in this study.

Should you have any questions or concerns about this survey, please feel free to contact Dr. John Nugent (jdnug@conncoll.edu), Sanjeewa Karunaratne (skarunara@conncoll.edu), or Dr. Jason Nier (irb@conncoll.edu).

Thank you.

End of Block: Block 1

Start of Block: Block 2

Q1 Below are some of the most common undergraduate surveys conducted at higher education institutions and administered by external organizations or internally by the institution. Please select the most recent survey conducted at your institution in the academic year 2022-2023 (between July 1, 2022, and June 30, 2023). If you administered more than one of these surveys, choose just one to provide details on -- the one for which you have the most complete information about respondents. *(Please note: This is a required question.)*

- ○ National Survey of Student Engagement (NSSE) (1)
- ○ HERI-CIRP Freshman Survey (TFS) (2)
- ○ Noel-Levitz Student Satisfaction Inventory (NSL) (3)
- ○ HEDS Student Satisfaction Survey (HEDS) (4)
- ○ Senior survey or graduating student survey (internal) (8)
- ○ First-year or freshman survey (internal) (15)
- ○ COFHE Senior Survey (Graduating Seniors) (16)
- ○ COFHE New Student Survey (New Fall Students) (17)
- ○ Other (please specify) (11) ____

End of Block: Block 2

Start of Block: Block 3

IQ2 *The following questions are used to compute overall response rate in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}.*

Q2 How many survey responses were recorded at your institution in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? *(Please enter a valid number.)*

____

Q3 What was the size of the eligible undergraduate population who received the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? *(Please enter a valid number.)*

____

End of Block: Block 3

Start of Block: Block 4

IQ3 *The following questions ask about demographic information collected in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}.*

Q4 Was any of the following demographic information collected in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? *(Please select all that apply.)*

- ○ race/ethnicity (1)
- ○ gender (2)

Display This Question:

If Was any of the following demographic information collected in the ... = gender

Q5 What was the ${Q4/ChoiceDescription/2} breakdown of the survey respondent population in ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? *(Please move the slider to your selection.)*

	Survey respondent population (%) 0 10 20 30 40 50 60 70 80 90 100
Male ()	
Female ()	
Other (non-binary) ()	

Display This Question:

If Was any of the following demographic information collected in the ... = gender

Q6 What was the ${Q4/ChoiceDescription/2} breakdown of the eligible undergraduate population in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? *(Please move the slider to your selection.)*

	Eligible student population (%) 0 10 20 30 40 50 60 70 80 90 100
Male ()	
Female ()	
Other (non-binary) ()	

Display This Question:

If Was any of the following demographic information collected in the ... = race/ethnicity

Q7 What was the ${Q4/ChoiceDescription/1} breakdown of the <u>survey respondent population</u> in ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? *(Please note: Race categories are based on IPEDS race/ethnicity definitions; move the slider to your selection.)*

	Survey respondent population (%) 0 10 20 30 40 50 60 70 80 90 100
Hispanic or Latino ()	
American Indian or Alaska Native ()	
Asian ()	
Black or African American ()	
Native Hawaiian or Other Pacific Islander ()	
White, non-Hispanic ()	
Nonresident ()	
Other non-citizens ()	
Two or more races ()	
Race/ethnicity unknown ()	

Display This Question:

If Was any of the following demographic information collected in the ... = race/ethnicity

Q8 What was the ${Q4/ChoiceDescription/1} breakdown of the eligible undergraduate population in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? *(Please note: Race categories are based on IPEDS race/ethnicity definitions; move the slider to your selection.)*

	Eligible student population (%) 0 10 20 30 40 50 60 70 80 90 100
Hispanic or Latino ()	
American Indian or Alaska Native ()	
Asian ()	
Black or African American ()	
Native Hawaiian or Other Pacific Islander ()	
White, non-Hispanic ()	
Nonresident ()	
Other non-citizens ()	
Two or more races ()	
Race/ethnicity unknown ()	

End of Block: Block 4

Start of Block: Block 5

IQ4 *The following set of questions are related to the administration of the ${Q1/ChoiceGroup/SelectedChoicesTextEntry} in the academic year 2022-2023 (between July 1, 2022 and July 30, 2023).*

Q11 How many reminders, excluding the invitation, were sent during the administration of the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}? (The reminders include any reminder notification sent to students, either by the external organization or by the institution.)

- None (1)
- One (2)
- Two (3)
- Three (4)
- Four (5)
- Five (6)
- Six or more (7)

Q12 When administering the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}, did you do any of the following? *(Please select all that apply.)*

- [ ] Posted flyers on campus (1)
- [ ] Communicated with faculty to spread the word (2)
- [ ] Posted digital notices on campus website (3)
- [ ] Posted digital messages on campus TV screens (4)
- [ ] Posted digital messages on student learning management system (e.g., Moodle) (5)

- ▢ Posted digital messages on student portal (e.g., Banner) (6)
- ▢ Contacted student organizations or clubs (7)
- ▢ Contacted departments or divisions (8)
- ▢ Posted on your institution's social media (e.g., Twitter, Instagram, or Facebook) (9)
- ▢ Sent pre-notification emails to students about the upcoming survey (10)
- ▢ Communicated with athletic coaches to spread the word (11)
- ▢ Contacted resident hall assistants or staff to spread the word (12)
- ▢ Sent survey reminders via text messages (15)
- ▢ Spread information about the survey by word-of-mouth (16)
- ▢ Something else (please describe) (13) ____

End of Block: Block 5

Start of Block: Block 6

IQ5 *Incentives are often offered to encourage students to participate in surveys. The following questions are related to the incentives offered in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry} in the academic year 2022-2023 (between July 1, 2022 and June 30, 2023).*

Q13 Was an incentive to participate offered in ${Q1/ChoiceGroup/SelectedChoicesTextEntry}?

- ◯ Yes (1)
- ◯ No (2)

Display This Question:

If Was an incentive to participate offered in ${q://QID1/ChoiceGroup/SelectedChoicesTextEntry}? = Yes

Q14 If an incentive was offered in the ${Q1/ChoiceGroup/SelectedChoicesTextEntry}, what was the incentive?

____

____

____

____

____

Display This Question:

If Was an incentive to participate offered in ${q://QID1/ChoiceGroup/SelectedChoicesTextEntry}? = Yes

Q15 What impact do you think this incentive may have had on your ${Q1/ChoiceGroup/SelectedChoicesTextEntry} overall response rate?

- ◯ Extremely positive (1)
- ◯ Moderately positive (2)
- ◯ Slightly positive (3)
- ◯ Neither positive nor negative (4)
- ◯ Slightly negative (5)
- ◯ Moderately negative (6)
- ◯ Extremely negative (7)

End of Block: Block 6

Start of Block: Block 7

Q16 Suppose you were to provide an incentive to participants in your next undergraduate survey. Which of the following incentives is most similar to what you would offer?

- ◯ A drawing for one iPad, Samsung Galaxy, or similar device (1)
- ◯ A drawing for one $300 gift card (Amazon, Visa, or similar) (4)
- ◯ A drawing for ten $30 gift cards (Amazon, Visa, or similar) (5)
- ◯ A drawing for thirty $10 gift cards (Amazon, Visa, or similar) (6)
- ◯ A non-monetary incentive (e.g., game tickets for a college sports event) (7)
- ◯ A pre-paid or contingency incentive prior to administering the survey (10)
- ◯ I would not offer an incentive (8)
- ◯ Other (9) ____

End of Block: Block 7

Start of Block: Block 8

IQ6 *The following questions are related to ALL undergraduate surveys conducted by your office in the academic year 2022-2023 (between July 1, 2022 and June 30, 2023), including the surveys administered by external organizations such as NSSE, HEDS, HERI, or Noel-Levitz, and the internal surveys administered by your office.*

Q17 How many surveys that were administered or coordinated by your office in the academic year 2022-2023 were sent to at least 10% of the undergraduate population?

- None (1)
- 1-5 (7)
- 6-10 (8)
- 11-15 (9)
- 16-20 (10)
- 21-30 (11)
- More than 30 (12)

End of Block: Block 8

Start of Block: Block 9

IQ7 *The level of student engagement at an institution can have an impact on survey response rates. The following questions are related to overall student engagement at your institution.*

Q18 Based on your best estimate, how many on-campus clubs or organizations are available to the students to participate?

- None (9)
- 1-20 (6)
- 21-40 (2)
- 41-60 (3)
- 61-80 (4)
- 81-100 (5)
- More than 100 (7)

Q19 Based on your best knowledge, how engaged are the students with the surrounding community through on-campus organizations? *(Please move the slider to your selection.)*

	Not at all engaged					Engaged					Highly engaged
	0	1	2	3	4	5	6	7	8	9	10
Community outreach ()											

Q20 How big a role do intercollegiate sports play at your institution? (Please move the slider to your selection.)

	Small role					Big Role					Very big role
	0	1	2	3	4	5	6	7	8	9	10
Intercollegiate sports ()											

Q21 How often does your institution organize events for students that address important social, economic, or political issues, for example, the war in Ukraine?

- Always (5)
- Frequently (6)
- Sometimes (7)
- Rarely (8)
- Never (9)

Q22 Does your institution have a Greek system (social fraternities and sororities)?

- Yes (1)
- No (4)

End of Block: Block 9

Start of Block: Block 10

IQ8 *These questions are related to your general survey practices of conducting undergraduate student surveys at your institution over the years.*

Q23 College campuses may try to reduce the survey "load" on their students. Have you done any of the following when administering or coordinating student surveys over the years? *(Please select all that apply.)*

- ◯ Managed the scheduling of the surveys through survey coordination with other offices, or similar process (1)
- ◯ Allowed only some offices or departments to send out surveys (2)
- ◯ Had a survey pre-approval process such as a form or IRB approval process (3)
- ◯ Managed access to survey software such as Survey Monkey or Qualtrics (4)
- ◯ Combined or consolidated similar internal surveys (5)
- ◯ Made surveys that are shorter in length (6)
- ◯ Communicated survey best practices to on-campus researchers (9)
- ◯ Made surveys that take less time to complete (10)
- ◯ Cut down the number of internal surveys (11)
- ◯ Administered external surveys such as NSSE, CIRP, NSL, or HEDS in rotational basis (e.g., every other year) (12)
- ◯ Something else (please describe) (8) ____

Display This Question:

If College campuses may try to reduce the survey "load" on their students. Have you done any of the... = Administered external surveys such as NSSE, CIRP, NSL, or HEDS in rotational basis (e.g., every other year)

Q24 How often do you administer ${Q1/ChoiceGroup/SelectedChoicesTextEntry}

- ◯ Annually (2)
- ◯ Every other year (3)
- ◯ Once in every three years (4)
- ◯ Once in every four years (5)
- ◯ Once in every five years (6)

Q25 In your practice over the years, how often have you surveyed a sample of students rather than the entire student population in a survey?

- Always (1)
- Frequently (2)
- Sometimes (4)
- Rarely (6)
- Never (7)

Display This Question:

If In your practice over the years, how often have you surveyed a sample of students rather than the... = Never

Q26 If you have never used a sample of students rather than the entire student population, what is the rationale for it?

____

____

____

____

____

Q27 In your practice over the years, how often have you had the chance to perform data weighting to make the survey respondent population look similar to your student population?

- ○ Always (1)
- ○ Frequently (2)
- ○ Sometimes (4)
- ○ Rarely (6)
- ○ Never (7)

Q28 In administering surveys over the years, please name one student subgroup or population that you think is harder to reach or typically has lower response rates compared with other student populations.

____

____

____

____

____

End of Block: Block 10

Start of Block: Block 11

Q29 What do you consider to be a high response rate on an undergraduate survey at your institution?

- ○ 0-20% (1)
- ○ 21-25% (2)
- ○ 26-30% (3)
- ○ 31-35% (4)
- ○ 36-40% (5)
- ○ 41-45% (6)
- ○ 46-50% (7)
- ○ More than 50% (8)

Display This Question:

If What do you consider to be a high response rate on an undergraduate survey at your institution? = 0-20%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 21-25%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 26-30%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 31-35%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 36-40%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 41-45%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 46-50%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = More than 50%

Q30 How important it is for you to achieve the ${Q29/ChoiceGroup/SelectedChoices} response rate in an undergraduate survey?

- ◯ Extremely important (1)
- ◯ Very important (2)
- ◯ Important (3)
- ◯ Not so important (4)
- ◯ Not at all important (5)

Display This Question:

If What do you consider to be a high response rate on an undergraduate survey at your institution? = 0-20%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 21-25%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 26-30%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 31-35%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 36-40%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 41-45%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = 46-50%

Or What do you consider to be a high response rate on an undergraduate survey at your institution? = More than 50%

Q31 How often do you achieve the ${Q29/ChoiceGroup/SelectedChoices} response rate in an undergraduate survey?

○ Always (1)

○ Frequently (2)

○ Sometimes (3)

○ Rarely (4)

○ Never (5)

End of Block: Block 11

Start of Block: Block 12

Q32 Regarding undergraduate surveys you have been involved in administering over the years, please name one practice you have identified that has helped improve student response rates.

____

____

____

____

____

End of Block: Block 12

Start of Block: Block 13

*

Q33 Please enter your college email address for confirmation. (Required question)

____

REFERENCES

Astin, A. W., & Molm, L. D. (1972). *Correcting for nonresponse bias in follow-up surveys.* Unpublished. American Council on Education.

Bock, D. E., Velleman, P. F., & De Veaux, R. D. (2018). *Stats: Modeling the world* (5th ed.). Pearson.

Brick, J. M., & Williams, D. (2013). Explaining rising nonresponse rates in cross-sectional surveys. *The ANNALS of the American Academy of Political and Social Science, 645*(1), 36–59. https://doi.org/10.1177/0002716212456834

Center for Survey Research and Analysis (CSRA). (2006). Weighting methodology. Unpublished manuscript. University of Connecticut.

Coates, H. (2006). *Student engagement in campus-based and online education: University connections.* Taylor & Francis.

Crawford, S. D., Couper, M. P., & Lamias, M. J. (2001). Web surveys: Perceptions of burden. *Social Science Computer Review, 19*(2), 146–162.

De Leeuw, E., & Heer, W. (2002). Trends in household survey nonresponse: A longitudinal and international comparison. In R. M. Groves, D. A. Dillman, J. L. Eltinge, & R. J. A. Little, (Eds.), *Survey nonresponse.* John Wiley & Sons.

Dey, E. L. (1997). Working with low survey response rates: The efficacy of weighting adjustments. *Research in Higher Education, 38*(2), 215–227.

Dillman, D. A. (2000). *Mail and internet surveys: The tailored design method* (2nd ed.). Wiley.

Dorofeev, S., & Grant, P. (2006). *Statistics for real-life sample surveys: Non-simple-random samples and weighted data.* Cambridge University Press.

Enten, H. (2024, November 1). Did the polls miss again because nobody answered their phones? *The Washington Post.* https://www.washingtonpost.com/business/2024/11/01/polls-accurate-nobody-answers-phones/

Fosnacht, K., Sarraf, S., Howe, E., Peck, L. K. (2020). How important are high response rates for college surveys? *The Review of Higher Education, 40*(2), 245–265. https://prod.wp.cdn.aws.wfu.edu/sites/202/2020/02/Fosnacht-2017-how-important-are-high-response-rates-for-college-surveys-3.pdf

Heberlein, T. A., & Baumgartner, R. (1978). Factors affecting response rates to mailed questionnaires: A quantitative analysis of the published literature. *American Sociological Review, 43*(4), 447–462.

Höfler, M., Pfister, H., Lieb, R., & Wittchen, H.-U. (2005). The use of weights to account for non-response and drop-out. *Social Psychiatry and Psychiatric Epidemiology, 40*(4), 291–299.

Hutchinson, S. R., Tollefson, N., & Wigington, H. (1987). Response bias in college student surveys. *Research in Higher Education, 27*(1), 61–73.

Keeter, S., Miller, C., Kohut, A., Groves, R. M., & Presser, S. (2000). Consequences of reducing nonresponse in a national telephone survey. *Public Opinion Quarterly, 64*(2), 125–148.

Little, R. J. A. (1993). Post-stratification: A modeler's perspective. *Journal of the American Statistical Association, 88*(423), 1001–1012.

Massey, D. S., & Tourangeau, R. (2013). Where do we go from here? Nonresponse and social measurement. *The ANNALS of the American Academy of Political and Social Science, 645*(1), 222–236.

Nair, C. S., Adams, P., & Mertova, P. (2008) Student engagement: The key to improving survey response rates. *Quality in Higher Education, 14*(3), 225-232. https://doi.org/10.1080/13538320802507505

National Center for Education Statistics (NCES). (2002). *The condition of education 2002.* U.S. Department of Education. https://nces.ed.gov/pubs2002/2002025.pdf

National Research Council (NRC). (2013). *Nonresponse in social science surveys: A research agenda.* The National Academies Press.

National Survey of Student Engagement (NSSE). (2003). *NSSE 2003 annual report.* Indiana University Center for Postsecondary Research.

National Survey of Student Engagement (NSSE). (n.d.). *Weighting: An explanation of weighting in NSSE reports.* Indiana University Center for Postsecondary Research. https://nsse.indiana.edu/nsse/reports-data/weighting.html

Porter, S. R., & Whitcomb, M. E. (2003). The impact of lottery incentives on student survey response rates. *Research in Higher Education, 44*(4), 389–407.

Porter, S. R., & Whitcomb, M. E. (2005). Non-response in student surveys: The role of demographics, engagement, and personality. *Research in Higher Education, 46*(2), 127–152.

Porter, S. R., Whitcomb, M. E., & Weitzer, W. H. (2004a). Multiple surveys of students and survey fatigue. *New Directions for Institutional Research, 121,* 63–73.

Porter, S. R., Whitcomb, M., E., & Weitzer, W. H. (2004b). Multiple surveys of students and survey fatigue. In S. R. Porter (Ed.), *Overcoming survey research problems.* Jossey-Bass.

Sax, L. J., Gilmartin, S. K., & Bryant, A. N. (2003). Assessing response rates and nonresponse bias in web and paper surveys. *Research in Higher Education, 44*(4), 409–432.

Scheaffer, R. L., Mendenhall, W. III, & Ott, R. L. (2006). *Elementary survey sampling* (6th ed.). Thomson Brooks/Cole.

Smith, T. W. (1995). Trends in nonresponse rates. *International Journal of Public Opinion Research, 7,* 157–171.

Steeh, C. G. (1981). Trends in nonresponse rates, 1952–1979. *Public Opinion Quarterly, 59,* 66–77.

Identifying At-Risk Course Combinations for Freshman Students Using Market Basket Analysis

Fikrewold Bitew and Khoi To

About the Authors

Fikrewold Bitew, PhD, and Khoi To, PhD, have a combined 25 years of experience in the field of institutional research. They currently work in the Office of Institutional Research and Analysis at the University of Texas at San Antonio. They are interested in applying data mining, artificial intelligence and machine learning, and statistical modeling to support decision-making at various levels.

Abstract

As institutions increasingly focus on improving student retention and graduation rates, understanding factors that influence student success has become crucial. Course failures in students' first terms have been shown to have strong associations with student retention. While extensive research has analyzed failure rates of individual courses, this study advances the field by using market basket analysis to identify courses that, when taken during the same semester, lead to high probabilities of failure. These insights are not available when courses are analyzed individually.

Using data from 7,466 first-time, full-time freshman students across five cohorts (Fall 2018–Fall 2022) at The University of Texas at San Antonio, we applied the Apriori algorithm to analyze 15,698 course enrollment records with failing grades (F, D–, D, D+). Our findings identified seven high-risk course combinations, with mathematics courses (MAT 1053, MAT 1073) appearing in four combinations and a writing course (WRC 1013) appearing in five combinations. Chi-square analyses revealed that students taking both courses in these combinations had significantly lower first-term retention rates (69.3%–81.0% vs. 83.0%–86.5%) and first-year retention rates (34.6%–54.0% vs. 54.6%–59.6%) compared to students taking only one course from the same combination.

The AIR Professional File, Fall 2025
Article 181
https://doi.org/10.34315/apf1812025

These findings provide actionable insights for academic advisors and curriculum designers to implement targeted intervention strategies, to enhance course scheduling guidance, and to develop support systems for high-risk course combinations. The study demonstrates the value of data-driven approaches in higher education and establishes a methodological framework that other institutions can replicate to improve student success outcomes.

Keywords: freshman students, at-risk course combinations, market basket analysis (MBA)

INTRODUCTION

Student success is a critical measure of institutional effectiveness in higher education, with retention and graduation rates serving as key indicators. However, despite substantial resources dedicated to supporting students, many of them still face challenges in achieving academic success. Understanding the factors that influence student success or failure is, therefore, crucial for developing strategies to improve academic outcomes. A key factor influencing student success is course failures in students' first semester, which research has identified as a strong predictor of academic performance (Slim et al., 2016). From that standpoint, being able to identify two or more courses that would pose a high risk of failure when taken together can help institutions design better support systems and interventions, enhance academic advising, and ultimately improve retention and graduation rates.

Much work has been done in analyzing failure rates of single courses. This study will take it one step farther by using market basket analysis (MBA) to identify courses that, if taken together, would lead to the high probability that a student will fail one or both courses. This data-mining technique, traditionally used in retail to identify associations between products that are often purchased together, offers a unique approach to examining these factors in an educational setting. By applying MBA to academic records, we can uncover patterns and associations among courses taken by students who encounter academic difficulties. The findings will provide actionable recommendations for academic advisors and curriculum designers, and will inform targeted intervention strategies, ultimately contributing to improved student performance.

LITERATURE REVIEW

The concept of "student success" encompasses multiple dimensions depending on stakeholder perspectives (Ezarik, 2023). It could include retention and persistence rates, degree completion, post-graduation outcomes such as employment or graduate school enrollment, or the acquisition of specific skills such as communication and critical thinking. Within this analysis, we focus on retention as our primary metric of student success; we define "retention" as the rate of return among students from term to term or from year to year (Soika, 2021).

Research has shown that certain courses, often referred to as "high-risk courses" or "bottleneck courses," have higher rates of failure and can significantly impact student retention and eventual graduation. A study by the Colorado State University Office of Institutional Research, Planning and Effectiveness (2012) found a strong negative association between unsuccessful course attempts and retention for first-time, full-time freshman students.

Similarly, a study by Michaels and Milner (2021) revealed significant gaps in retention and graduation among students receiving different numbers of D/F grades in the first term. On average, students saw an 8.5% retention differential with one D/F grade in the first term and a 35% retention differential with more than one D/F grade, as compared to students with no D/F grade in the first term. Although the data were not causal, early grades were clearly correlated with student success.

A recent study conducted by the University of Wyoming's Office of Institutional Analysis (Zong & Koller, 2023) also pointed out that students taking fewer high-risk courses are more likely to retain after the first year. The study identified high-risk courses as those with a pass rate of less than 80% over 5 years. These courses were predominantly in science, technology, engineering, and mathematics (STEM) fields, with quantitative reasoning and physical sciences being particularly challenging courses for students.

Recent advances in business intelligence have enabled researchers to apply data mining techniques to analyze student course-taking patterns. MBA has emerged as a pioneering approach to identify course combinations associated with high failure rates. This data mining technique, originally developed for retail to identify product associations, has been effectively adapted for educational settings (Papadogiannis et al., 2024; Romero & Ventura, 2020).

This approach provides insights that are not observable if individual courses were examined in isolation. For instance, students who enroll in multiple high-risk courses in their first semester are more likely to struggle academically and therefore are more vulnerable to attrition. Additionally, the analysis can reveal hidden patterns, such as the impact of taking certain elective courses alongside core courses (McKinsey & Company, 2023).

Safour et al. (2024) applied MBA to identify associations between courses in a university setting, finding that certain course combinations were more likely to result in student dropouts. Similarly, González et al. (2008) used MBA to explore the relationship between course sequences and student performance, revealing that some course sequences were strongly associated with lower grades.

Bautista (2005) applied MBA to analyze first-year engineering student course combinations, identifying the top 20 combinations with high failure rates. The study highlighted significant financial implications of course failures, since unsuccessful attempts result in extended enrollment periods, creating additional financial burdens for students and their families, and creating pressure on government funding sources in public institutions.

RESEARCH QUESTIONS AND SIGNIFICANCE OF THE STUDY

Given the demonstrated impact of course failures on student retention, this study analyzes data from first-time, full-time freshman students at The University of Texas at San Antonio (UTSA), using MBA to explore first-semester course-taking patterns. Two primary research questions guide this investigation:

Research Question 1: Among the courses taken by first-time, full-time freshman students at UTSA in their first semester, which course combinations, when taken together, lead to high probabilities of failure?

Research Question 2: Are there significant relationships between first-term retention and first-year retention when comparing two groups of students—those taking only one course from a high-risk combination versus those taking both courses from the same combination?

Findings from these research questions contribute to understanding how certain course combinations lead to high failure rates and adversely affect student retention. This knowledge enables university administrators, academic advisors, and curriculum planners to develop effective strategies for teaching methodology, curriculum design, and targeted student support systems.

DATA AND METHODS

This study applies the Knowledge Discovery in Databases (KDD) approach that was developed by Usama Fayyad, Gregory Piatetsky-Shapiro, and Padhraic Smyth (see Figure 1). This framework is recognized as being one of the most reliable research methodologies for academic purposes, since it was designed to identify hidden patterns, unseen trends, and correlations in databases to inform future decision-making (Bautista, 2005).

Figure 1. Knowledge Discovery in Databases (KDD) Approach in Data Mining

The Knowledge Discovery in Databases (KDD) approach was systematically applied to analyze student course data at UTSA, with each step detailed below in this section, as well as in Results and Discussion.

Data Selection

The study uses anonymized data from UTSA, covering course enrollments and final grades of five first-time, full-time freshman cohorts (Fall 2018–Fall 2022). Courses with final grades indicating failure (F, D–, D, D+) were retained for the analysis.

We use Oracle SQL to retrieve data from the university's student database (DMARTPROD), extracting student-course records needed for the analysis. The dataset has a total of 15,698 records and was structured in a transactional layout, with each row representing a student-course enrollment ("a transaction") in each given term. This layout is also known as the "long format."

The dataset was exported as an Excel file and imported into Python for processing and analysis. The following describes each variable in the dataset, with a representative sample presented in Figure 2.

- COHORT_TERM: Identifies the student cohort at UTSA, representing the student's first enrollment term (201910, 202010, 202110, 202210, and 202310, corresponding to Fall 2018, Fall 2019, Fall 2020, Fall 2021, and Fall 2022, respectively).
- COHORT_PIDM: Unique student identifiers (anonymized to protect privacy).
- COHORT_IPED: Binary indicator of cohort status (1 = Roadrunner, 0 = CAP). "Roadrunner" refers to fully enrolled UTSA students (represented by the university mascot). "CAP" refers to students in the University of Texas at Austin Coordinated Admission Program (CAP) who begin their studies at UTSA. All students in this dataset are Roadrunners.
- COHORT_FT: Binary indicator of full-time enrollment status (1 = full time, 0 = part time). All students in this dataset are full time.
- COHORT_TRANSFER: Indicates student type (1 = new transfer, null = first timer). All students in this dataset are first timers.
- COURSE: Lists courses taken during the first semester that resulted in final grades of F, D–, D, or D+. Each row represents one course per student, so students with multiple failing grades appear in multiple rows.
- GRDE_CODE_FINAL: Final grade received (F, D–, D, or D+) for the corresponding course.
- CRS_INDEX: Sequential counter of failing courses per student in the first semester (reference only, not used in analysis). For example, a student with four failing courses would have index values 1–4 across their respective rows.
- FIRST_SPRING: Spring term enrollment indicator showing whether the student returned for the subsequent Spring semester.
- RETAINED_FIRST_SPRING: Binary indicator derived from FIRST_SPRING (1 = retained for Spring semester). Used in chi-square tests to examine relationships between course combinations and first-term retention.
- SECOND_FALL: Fall term enrollment indicator showing whether the student returned for the second Fall semester (first-year retention).
- RETAINED_SECOND_FALL: Binary indicator derived from SECOND_FALL (1 = retained for second Fall semester). Used in chi-square tests to examine relationships between course combinations and first-year retention.

Figure 2. Sample of Dataset Queried Out from The University of Texas at San Antonio's Student Database (DMARTPROD)

COHORT_TERM	COHORT_PIDM	COHORT_IPED	COHORT_FT	COHORT_ TRANSFER	COURSE	GRDE_CODE _FINAL	CRS_INDEX	FIRST _SPRING	RETAINED_FIRST _SPRING	SECOND_ FALL	RETAINED_SECOND _FALL
201910	1612364	1	1		AIS 1203	F	1	201920	1		0
201910	1612364	1	1		MAT 1214	D	2	201920	1		0
201910	1612364	1	1		ME 1403	F	3	201920	1		0
201910	1612364	1	1		WRC 1023	F	4	201920	1		0
201910	1612375	1	1		MAT 1073	D	1	201920	1	202010	1
201910	1612615	1	1		ANT 2063	D+	1	201920	1		0
201910	1612615	1	1		HIS 1043	D-	2	201920	1		0
201910	1612615	1	1		MAT 1073	F	3	201920	1		0
201910	1613347	1	1		CLA 2323	F	1	201920	1		0
201910	1613347	1	1		SOC 2013	D	2	201920	1		0
201910	1613349	1	1		HIS 1053	F	1	201920	1	202010	1
201910	1613558	1	1		CHE 1121	D	1	201920	1	202010	1
201910	1613558	1	1		ECO 2023	F	2	201920	1	202010	1
201910	1613558	1	1		MAT 1093	F	3	201920	1	202010	1
201910	1613564	1	1		GLA 1013	D	1	201920	1	202010	1
201910	1613577	1	1		CRJ 1113	D	1	201920	1		0
201910	1613577	1	1		MAT 1073	F	2	201920	1		0
201910	1613577	1	1		POL 1213	F	3	201920	1		0
201910	1613577	1	1		WRC 1013	D+	4	201920	1		0

Data Pre-processing

From this point forward, we utilize Python and Jupyter Notebook, taking advantage of their robust data manipulation, analysis, and visualization capabilities in data mining.

Relevant packages are imported into Python, and the Excel data file is loaded and set up for pre-processing.

```
# Load relevant packages into Python
import pandas as pd
import numpy as np
from mlxtend.frequent_patterns import apriori, association_rules
from mlxtend.preprocessing import TransactionEncoder
import matplotlib.pyplot as plt
import seaborn as sns
```

```
# Load the data from Excel file
df = pd.read_excel('student_course_data.xlsx')
print(f"Data shape: {df.shape}")
print(f"Columns: {list(df.columns)}")
```

Data quality assessment revealed no missing values or duplicate records. All variable data types were verified as appropriate for the intended analyses.

```python
# Check for missing values
print("Missing values per column")
print(df.isnull().sum())

# Check for duplicate records
print(f"\nDuplicate records: {df.duplicated().sum()}")

# Display basic statistics
print("\nDataset overview:")
print(df.info())

# Display first few rows
print("\nFirst 5 rows of the dataset")
df.head()
```

Following is a screenshot of the first few rows of the dataset loaded in Python.

COHORT_TERM	COHORT_PIDM	COHORT_IPED	COHORT_FT	COHORT_TRANSFER	COURSE	GRDE_CODE_FINAL	CRS_INDEX	FIRST_SPRING	RETAINED_FIRST_SPRING
201510	[redacted]	1	1	NaN	MAT 1073	D	1	201520.0	1
201510	[redacted]	1	1	NaN	CS 1711	F	2	NaN	0
201510	[redacted]	1	1	NaN	PHY 1943	F	4	NaN	0
201510	[redacted]	1	1	NaN	CS 1713	F	3	NaN	0
201510	[redacted]	1	1	NaN	WRC 1023	F	5	NaN	0

Table 1 presents a summary of the data. Several key patterns emerge from these data as described below.

1| Distribution of grades: F grades are consistently the most common low grade across all semesters, comprising 9,623 instances (61.3% of all low grades). D grades are the second-most common (4,348 or 27.7%), while D+ (815 or 5.2%) and D– (912 or 5.8%) occur less frequently.

2| Pandemic impact: The data suggest possible effects of the COVID-19 pandemic, with Fall 2020 showing a notable increase in F grades (2,134) while having fewer D-range grades, potentially indicating greater academic challenges during this period.

3| Recent developments: The most recent data (Fall 2022) show some improvement, with fewer F grades compared to the previous year while still maintaining relatively high counts of D– grades.

Table 1. Descriptive Summary of the Dataset in Figure 2

Cohort	Final Grade				Total
	F	D–	D	D+	
Fall 2018	1,657	107	1,126	99	2,989
Fall 2019	1,494	66	1,019	87	2,666
Fall 2020	2,134	102	845	73	3,154
Fall 2021	2,415	310	679	295	3,699
Fall 2022	1,923	327	679	261	3,190
Total	**9,623**	**912**	**4,348**	**815**	**15,698**

Data Transformation

Before we perform MBA, we must convert the dataset into a format compatible with the Apriori algorithm. Specifically, the data require transformation from their current transactional layout (long format) into a tabular structure (wide format) comprising Boolean values. A "True" value indicates student enrollment in a course, while a "False" value indicates non-enrollment.

The first step in data transformation involves removing variables that are not relevant to MBA. Several columns were excluded, including cohort identifiers, full-time status, retention indicators, and grade information; the remainder focused exclusively on student-course relationships. The analysis retains only student identifiers (COHORT_PIDM) and courses taken (COURSE).

```
# Select relevant columns for transformation
columns_to_keep = ['COHORT_PIDM', 'COURSE']
df_transformed = df[columns_to_keep].copy()

# Display the structure of the filtered data
print("Filtered dataset structure")
print(df_transformed.head())
print(f"\nShape after filtering: {df_transformed.shape}")
print(f"Unique students: {df_transformed['COHORT_PIDM'].nunique()}")
print(f"Unique courses: {df_transformed['COURSE'].nunique()}")
```

The following screenshot displays a sample of the filtered dataset.

COHORT_PIDM	COURSE
[redacted]	CRJ 1113
	POL 1013
	AIS 1203
	ES 2013
	HIS 1053
	IS 1403
	MAT 1053
	FIN 3013
	ANT 2043
	ES 1121
	ES 1123
	HIS 2533
	MAT 1073
	WRC 1023
	MAT 1023
	PSY 1013

The next step involves transforming the data into a tabular format suitable for the Apriori algorithm, where each row represents a student and each column represents a course. The values in this transformed dataset are either "True" or "False." A "True" value indicates that a student takes a specific course, and a "False" value means the student does not take that course.

```
# Create a pivot table to transform from long to wide format
basket = df_transformed.groupby(['COHORT_PIDM', 'COURSE']) .size().unstack().
fillna(0)

# Convert to boolean values (True/False)
def encode_units(x):
    if x <= 0:
        return False
    if x >= 1:
        return True
basket_sets = basket.applymap(encode_units)
print("Transformed dataset structure:")
print(f"Shape: {basket_sets.shape}")
print("Sample of transformed data:")
print(basket_sets.head())
```

After transformation, the dataset contains 7,466 rows, each representing a unique student, and 373 columns, each representing a distinct course. Since full-time students typically enroll in four courses per semester, each student row contains four "True" values among the 373 available columns, with the remaining columns showing "False" values. In the MBA analogy, students represent "baskets" and courses represent "items," hence the term "market basket analysis." The transformed dataset is now ready for analysis (Figure 3).

Figure 3. Sample of Transformed Dataset Ready for Market Basket Analysis

	AAS 2013	AAS 2113	AAS 3013	AAS 3123	ACC 2003	ACC 2013	ACC 2033	ACC 3123	AHC 1113	AHC 1123	...	STA 3313	SWK 1013	UCS 2033	UTE 1111
0	False	False	False	False	False	False	False	False	False	False	...	False	False	False	False
1	False	False	False	False	False	False	False	False	False	False	...	False	False	False	False
2	False	False	False	False	False	False	False	False	False	False	...	False	False	False	False
3	False	False	False	False	False	False	False	False	False	False	...	False	False	False	False
4	False	False	False	False	False	False	False	False	False	False	...	False	False	False	False
...	...	...	...	...	...	...	...	...	...	...	... ...	...	...	...	...

Data Mining

There are several algorithms in MBA: they include Apriori, artificial immune system (AIS), sort and merge scan (SETM), and Frequent Pattern–Growth (FP–Growth). In this study, we use Apriori, a well-established and commonly used algorithm in MBA, to identify frequent itemset and association rules within large datasets. The Apriori algorithm was selected for its simplicity and ease of implementation. The algorithm's straightforward approach to generating candidate itemset and applying support and confidence thresholds facilitates both implementation and interpretation. However, Apriori has computational limitations, and requires multiple dataset scans that can demand significant memory and processing resources, particularly with large datasets containing numerous frequent patterns.

Figure 4 illustrates the steps in Apriori algorithm process. The following describes each step from initial data processing through identification of significant item combinations:

Figure 4. Steps in Apriori Algorithm Process

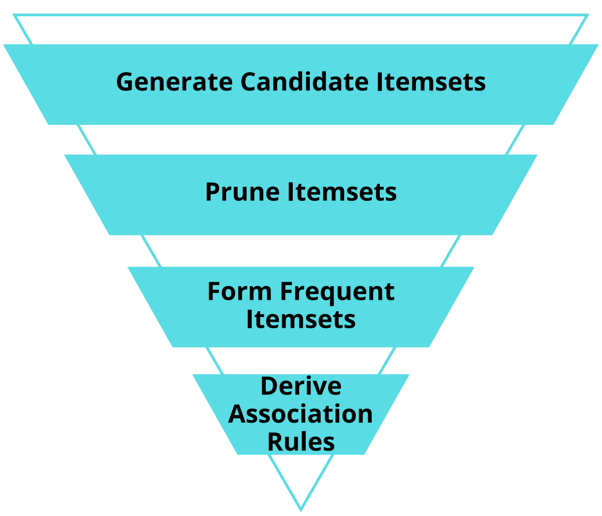

Source: Mwiti 2025.

- **Step 1: Generate candidate itemset:** The algorithm identifies individual items and counts their occurrences to determine frequent items. The fundamental principle states that, if an itemset is frequent, all its subsets must also be frequent. This assumption reduces the number of itemsets requiring evaluation, and improves algorithmic efficiency.
- **Step 2: Prune itemsets:** Itemsets occurring below the minimum support threshold are eliminated from further consideration.
- **Step 3: Form frequent itemsets:** The algorithm creates larger itemsets by combining frequent smaller itemsets, iterating until no additional frequent itemsets can be identified.
- **Step 4: Derive association rules:** The algorithm extracts meaningful association rules based on support, confidence, and lift values to identify significant relationships among items.

In this analysis, itemsets are defined as combinations of courses taken by students; our goal is to identify rules that can predict course failure based on these combinations.

Metrics Used to Evaluate Association Rules Among Items

The Apriori algorithm uses four key metrics—support, confidence, lift, and Zhang's metric—to identify meaningful association rules among items. These metrics are essential for selecting significant association rules, particularly when analysis generates numerous potential relationships (Alangari & Alturki, 2020). Each metric is detailed next (Derouiche, 2024):

SUPPORT

Definition: Support measures the proportion of course combinations in the dataset that contain a particular course set.

Interpretation: High support suggests that the course set occurs frequently in the dataset. The higher the support, the more frequently the itemset occurs. Rules with a high support are preferred since they are likely to be applicable to a large number of future course combinations. For instance, if the course set {Math, Physics} appears in 20 out of 100 course combinations, the support would be 0.20 or 20%.

Formula:

$$\textit{Support } (A \rightarrow B) = \frac{\textit{Number of transactions containing both A and B}}{\textit{Total number of transactions}}$$

Confidence

Definition: Confidence measures the likelihood that Course B is taken when Course A is taken.

Interpretation: High confidence indicates a strong association between Courses A and B. The higher the confidence, the greater the likelihood that Course B (the right-hand side) will be taken when Course A (the left-hand side) is taken. For example, if 15 out of 20 course combinations that contain mathematics also contain physics, the confidence is 0.75 or 75%.

Formula:

$$\textit{Confidence } (A \rightarrow B) = \frac{\textit{Support}(A \cup B)}{\textit{Support}(A)}$$

LIFT

Definition: Lift measures how much more likely that Course B is taken when Course A is taken, compared to the likelihood of taking Course B independently.

Interpretation: A lift value greater than 1 indicates a positive association between Courses A and B, suggesting they are more likely to be taken together than independently. For instance, a lift of 1.5 means that enrolling in mathematics increases the likelihood of enrolling in physics by 50%.

Formula:

$$\textit{Lift}(A \rightarrow B) = \frac{\textit{Confidence}(A \rightarrow B)}{\textit{Support}(B)}$$

ZHANG'S METRIC

Zhang's metric offers a valuable alternative measure for evaluating association rules in MBA. Unlike lift, which can sometimes overemphasize rare item relationships, Zhang's metric provides a more-balanced evaluation by considering both the presence and the absence of items. This metric offers a normalized measure (–1 to 1), where positive values indicate positive correlation, zero indicates independence, and negative values indicate negative correlation.

Zhang's metric is calculated using the following formula:

$$\textit{Zhang's Metric}(A \rightarrow B) = \frac{\textit{Confidence}(A \rightarrow B) - \textit{Support}(B)}{1 - \textit{Support}(B)}$$

RESULTS AND DISCUSSION

Identifying Association Rules for Course Failure Patterns

With the dataset properly formatted, the Apriori algorithm was applied to identify frequent course combinations and their association rules. The minimum support threshold was set at 0.01 (1%) to capture meaningful patterns while avoiding overly restrictive filtering. This threshold requires course combinations to appear in at least 1% of total student records, equivalent to 75 students (7,466 × 1%). The minimum confidence was set at 0.2 (20%), indicating that, when a student fails one course in a combination, there is at least a 20% probability that the student will fail the associated course.

These moderate threshold values were selected considering the large course catalog (373 total courses) and the extensive number of possible combinations. Higher thresholds would eliminate potentially meaningful patterns due to the natural diversity in student course selections. Effective association rule mining requires careful threshold selection based on domain knowledge, data characteristics, and iterative refinement to identify actionable insights.

The analysis specifications targeted two-course combinations (length = 2) with minimum support of 0.01 (support > = 0.01) and minimum confidence of 0.2 (metric = "confidence," min_threshold = 0.2).

```
# Filter frequent itemsets by length & support
apriori_model_colnames = apriori_model_colnames[
    (apriori_model_colnames['length'] == 2) &
    (apriori_model_colnames['support'] >= 0.01)
]

# Generate association rules
rules = mlxtend.frequent_patterns.association_rules(
    apriori_model_colnames,
    metric="confidence",
    min_threshold=0.2,
    support_only=False
)
```

The model identifies the top seven course combinations that are associated with failure (F, D–, D, D+ grades), if taken together. Table 2 shows these course combinations and their performance metrics of support, confidence, lift, and Zhang's. The course combinations and the interplay among their performance metrics are discussed below.

Table 2. Course Combinations Leading to Failure for First-Time, Full-Time Freshman Students

Antecedents	Consequents	Support	Confidence	Lift	Zhang's
BIO 1404	MAT 1073	0.023	0.460	3.080	0.711
AIS 1203	WRC 1013	0.020	0.315	1.695	0.437
CHE 1073	MAT 1073	0.030	0.285	1.908	0.532
HIS 1053	WRC 1013	0.017	0.277	1.488	0.350
HIS 1043	WRC 1013	0.015	0.276	1.481	0.343
MAT 1053	WRC 1013	0.014	0.238	1.282	0.234
MAT 1073	WRC 1013	0.035	0.235	1.261	0.243

For reference, the following are the full course titles for the identified combinations:

- AIS 1203–Academic Introduction and Strategies
- BIO 1404–Biosciences I
- CHE 1073–Basic Chemistry
- MAT 1053–Mathematics for Business
- MAT 1073–Algebra for Scientists and Engineers
- HIS 1043–United States History: Pre–Columbus to Civil War Era
- HIS 1053–United States History: Civil War Era to Present
- WRC 1013–Freshman Composition I

The support metric indicates the frequency of co-occurrence, with approximate values ranging from 0.014 to 0.035 across the top seven rules. While these support values may appear low, they are significant in educational context, where failing multiple courses when taken together is relatively uncommon. The confidence metric, whose approximate values range from 0.235 to 0.460, represents the conditional probability of failing the consequent course, given failure in the antecedent course.

The lift metric provides a complementary perspective on association strength. Lift values exceeding 1 (ranging from 1.261 to 3.080) confirm positive association. Zhang's metric also offers useful insight into the performance of the association rules. For the top seven rules identified, the values of Zhang's metric range from 0.711 to 0.243, indicating positive association within each course combination.

- Rule 1 (BIO 1404 and MAT 1073) is obtained under a high degree of confidence and lift. It signifies that those students who struggled in BIO 1404 also struggled in MAT 1073.
- Rule 2 (AIS 1203 and WRC 1013) also has high degree of confidence and lift. It illustrates that students who struggled in AIS 1203 also had difficulty in WRC 1013.
- Rule 3 (CHE 1073 and MAT 1073) is selected because it has a high degree of support. It shows that these two courses were taken together frequently; this rule is applicable to many future course combinations.
- Rules 4 and 5 (HIS 1043/HIS 1053 and WRC 1013) are based on a high degree of confidence and support. Those rules indicate association of student performance in HIS 1043/HIS 1053 and student performance in WRC 1013.
- Rules 6 and 7 (MAT 1053/MAT 1073 and WRC 1013) are obtained based on a high degree of confidence and support. Those rules show that students who struggled in MAT 1053/MAT 1073 also struggled in WRC 1013.

The analysis reveals critical patterns: mathematics courses (MAT 1053/MAT 1073) appear in four of the seven high-risk combinations, while the writing course (WRC 1013) appears in five combinations. This frequency suggests that targeted intervention strategies should prioritize mathematics and writing support, since these subjects frequently co-occur with failures in other disciplines. These cross-disciplinary failure patterns would remain undetected through individual course analysis.

A Closer Look at the Performance Metrics of the Association Rules for Course Failure Patterns

The scatter plot in Figure 5 illustrates a fundamental tradeoff that exists in association rule mining. Confidence and support are two primary metrics used to evaluate the strength and relevance of association rules; their relationship reveals important insights about the underlying data patterns.

Figure 5. Plot of Confidence and Support Levels of Association Rules

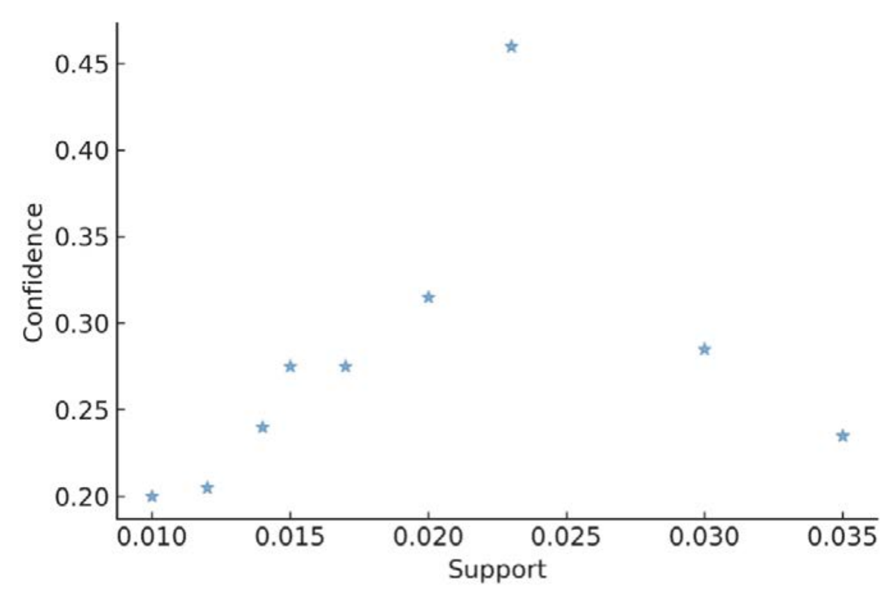

Looking specifically at Figure 5, we can observe several key characteristics.

- **Distribution pattern:** The points are scattered across different support and confidence values, with no rules appearing in the upper-right quadrant (high support and high confidence). Instead, we see a cluster of points with approximate support values ranging from 0.010 to 0.035 and approximate confidence values ranging from 0.200 and 0.460.
- **Highest confidence rule:** The rule with the highest confidence (approximately 0.460) has a moderate support value of about 0.023. This suggests that, while this rule is highly reliable when it applies, it affects only a moderate number of students.
- **Highest support rules:** The rules with the highest support values (around 0.030 to 0.035) have relatively lower confidence values (around 0.235), indicating these combinations occur more frequently but are less reliable as predictors.
- **Moderate confidence–support balance:** Several rules fall in the middle range, with confidence values between 0.240 and 0.320 and support values between 0.015 and 0.020, representing potentially valuable insights that balance frequency and reliability.

The absence of rules in the upper-right quadrant of the plot reflects a natural constraint in most real-world datasets, particularly in educational contexts. This absence of rules occurs for several interconnected reasons.

- **Statistical dilution effect:** Courses with high support (frequently failed courses) naturally appear across many different student records and combinations. This widespread distribution means they co-occur with many different courses, and not just with specific ones. Consequently, the conditional probability (confidence) of one high-support course leading to another specific course tends to decrease.
- **Specificity versus generality tradeoff:** Association rules that capture highly specific relationships (high confidence) typically involve more-distinctive or more-specialized course combinations that fewer students take, resulting in lower support. Conversely, rules with high support often capture patterns that are more general, and that apply to many students but are less predictive of specific outcomes.
- **Course selection diversity:** Students follow diverse academic paths based on their majors, interests, and requirements. This diversity naturally limits the number of students taking identical course combinations, creating a ceiling effect on support values for highly predictive rules.

In educational data mining specifically, the above analysis of the performance metrics suggests that the most actionable insights often come from rules with moderate support and confidence, representing course combinations that both affect several students and demonstrate reliable predictive power. For academic advisors and curriculum designers, these middle-ground rules often identify the most promising targets for intervention, since they balance relevance (affecting enough students to matter) with predictive accuracy (reliably identifying problematic course combinations). As explained at the beginning of Results and Discussion, the minimum support and minimum confidence are set with moderate values since there are many individual courses (373 courses in total) that students can choose from, and a significant number of possible combinations that can be made of these courses. Additionally, choosing the appropriate support and confidence values requires domain knowledge, understanding of the data at hand, experimentation, and iterative refinement to find values that yield actionable insights.

In the next phase of the analysis, we analyze the relationships between the course combinations identified above with student retention rates. Specifically, we will examine whether there are statistically significant relationships with first-term retention and first-year retention between two groups of students: those taking only one course from a course combination compared with those taking both courses from the same combination.

Chi-square tests of independence are conducted for the seven course combinations discussed previously, and all show significant relationships between two groups of students and retention rates. The results of combinations BIO 1404–MAT 1073, CHE 1073–MAT 1073, and HIS 1053–WRC 1013 are presented next.

BIO 1404–MAT 1073

First-Term Retention

From Table 3 we see that 1,023 students take either BIO 1404 or MAT 1073, and 174 students take both courses. The first-term retention rates of the two groups are 83.4% and 81.0%, respectively.

Table 3. BIO 1404–MAT 1073, First-Term Retention

BIO 1404–MAT 1073	Retained after first term		Total	% Retained
	No	Yes		
Taking either course	170	853	1,023	83.4%
Taking both courses	33	141	174	81.0%

A chi-square test is performed to examine the relationship between the two groups of students in Table 3 and first-term retention rates. There is significant relationship between the two variables, with X^2 (1, N = 1,197) = 5.0, $p < .05$. Students taking both courses are less likely to retain after the first term than are students taking either course.

First-Year Retention

First-year retention rates of the two groups of students are 55.7% and 54.0%, respectively (Table 4).

Table 4. BIO 1404–MAT 1073, First-Year Retention

BIO 1404–MAT 1073	Retained after first year		Total	% Retained
	No	Yes		
Taking either course	453	570	1,023	55.7%
Taking both courses	80	94	174	54.0%

A chi-square test is performed to examine the relationship between the two groups of students in Table 4 and first-year retention rates. There is significant relationship between the two variables, with X^2 (1, N = 1,197) = 156.5, $p < .001$. Students taking both courses are less likely to retain after the first year than are students taking either course.

CHE 1073–MAT 1073

First-Year Retention

Table 5 shows that there are 1,452 students taking either CHE 1073 or MAT 1073, and that 222 students are taking both courses. The retention rates of the two groups are 86.5% and 71.6%, respectively.

Table 5. CHE 1073–MAT 1073, First-Term Retention

CHE 1073–MAT 1073	Retained after first term		Total	% Retained
	No	Yes		
Taking either course	196	1,256	1,452	86.5%
Taking both courses	63	159	222	71.6%

A chi-square test is performed to examine the relationship between the two groups of students in Table 5 and first-term retention rates. There is significant relationship between the two variables, with X^2 (1, N = 1,674) = 32.6, p < .001. Students taking both courses are less likely to retain after the first term than are students taking either course.

First-Year Retention

First-year retention rates of the two groups of students are 59.6% and 39.6%, respectively (Table 6).

Table 6. CHE 1073–MAT 1073, First-Year Retention

CHE 1073–MAT 1073	Retained after first year		Total	% Retained
	No	Yes		
Taking either course	586	866	1,452	59.6%
Taking both courses	134	88	222	39.6%

A chi-square test is performed to examine the relationship between the two groups of students in Table 6 and first-year retention rates. There is significant relationship between the two variables, with X^2 (1, N = 1,674) = 31.4, p < .001. Students taking both courses are less likely to retain after the first year than are students taking either course.

HIS 1053–WRC 1013

First-Year Retention

There are 1,597 students taking either HIS 1053 or WRC 1013, and 127 students taking both courses. The retention rates of the two groups are 83.0% and 69.3%, respectively (Table 7).

Table 7. HIS 1053–WRC 1013, First-Term Retention

HIS 1053–WRC 1013	Retained after first term		Total	% Retained
	No	Yes		
Taking either course	272	1,325	1,597	83.0%
Taking both courses	39	88	127	69.3%

A chi-square test is performed to examine the relationship between the two groups of students in Table 7 and first-term retention rates. There is significant relationship between the two variables, with X^2 (1, N = 1,724) = 18.5, $p < .001$. Students taking both courses are less likely to retain after the first term than are students taking either course.

First-Year Retention

First-year retention rates of the two groups of students are 54.6% and 34.6%, respectively (Table 8).

Table 8. HIS 1053–WRC 1013, First-Year Retention

HIS 1053–WRC 1013	Retained after first year		Total	% Retained
	No	Yes		
Taking either course	725	872	1,597	54.6%
Taking both courses	83	44	127	34.6%

A chi-square test is performed to examine the relationship between the two groups of students in Table 8 and first-year retention rates. There is significant relationship between the two variables, with X^2 (1, N = 1, 724) = 262.9, $p < .001$. Students taking both courses are less likely to retain after the first year than are students taking either course.

The bivariate analyses above show that the seven course combinations have statistically significant relationship with retention rates. Students taking both courses in a course combination are less likely to retain after the first term and after the first year than are students taking only one course from the same course combination.

SUMMARY AND IMPLICATIONS

The findings from this study provide valuable insights into the course combinations that are negatively correlated with retention of first-time, full-time freshman students. We identified seven pairs of courses that, when taken together, significantly increase the likelihood of student attrition. This information is crucial for university administrators, academic advisors, and curriculum planners aiming to improve student retention rates. Such insights would not have been available if courses were analyzed individually rather than in combination.

The identification of the high-risk course combinations above allows for targeted interventions. Academic advisors can use this information to guide first-time, full-time freshman students in selecting their courses more strategically. For example, advisors might recommend against taking CHE 1073 and MAT 1073 in the same semester, or they might suggest additional support resources for students enrolled in these courses.

These findings also have implications for curriculum design and institutional policy. Universities could consider restructuring the timing and prerequisites of high-risk courses to reduce the likelihood of students encountering these problematic combinations. Additionally, supplemental instruction or tutoring programs could be developed specifically for the identified high-risk course pairs.

The analysis of high-risk course combinations also offers practical financial implications to consider. Failing a course would mean an additional semester or year enrolled, hence increasing financial responsibilities for students, their families, and the government. Being able to identify high-risk course combinations would help university administrators design strategies to support students taking these courses and help relieve the financial burden of the parties involved.

The application of MBA has proven to be a valuable tool in identifying course combinations that are associated with high rates of failure. Understanding these associations provides colleges and universities with insights to develop effective strategies that support student success. From that perspective, this study highlights the importance of data-driven approaches in higher education and sets the stage for further research in this area.

LIMITATIONS OF THE STUDY AND FUTURE RESEARCH

Limitations of the Study

The study acknowledges limitations on data generalizability and temporal changes. On the one hand, the findings in this analysis are based on data from a single institution (UTSA) and might not be readily generalizable to other contexts. On the other hand, the data are from Fall 2018–Fall 2022 cohorts. Therefore, changes in curriculum and academic policies after the study period could affect subsequent analyses.

An additional limitation involves the exclusion of demographic variables (race/ethnicity, gender, age groups) from the current analysis due to the study's broad scope. Future research could apply MBA within specific demographic subgroups to identify population-specific patterns, as demonstrated by Çiçekli and Kabasakal (2021).

MBA provides valuable pattern identification but has inherent limitations. Association rules identify correlations rather than causal relationships, requiring careful interpretation when developing interventions. Additionally, the algorithm's multiple database scans demand substantial computational resources, particularly with large datasets containing numerous frequent patterns.

Future Research

Future research should aim to replicate this study across multiple institutions to validate the results. Additionally, future studies could explore the underlying reasons why these specific course combinations lead to higher failure rates, such as course content difficulty, teaching methods, or student preparedness. Another study that can be considered to follow up from this one is to explore how these course combinations could be associated with student attrition. This can be done by treating student attrition as the response variable and course combinations among the predictors.

REFERENCES

Alangari, N., & Alturki, R. (2020). Association rule mining in higher education: A case study of computer science students. In R. Mehmood, S. See, I. Katib, & I. Chlamtac (Eds.), *Smart infrastructure and applications* (pp. 311–328). EAI/Springer Innovations in Communication and Computing. https://doi.org/10.1007/978-3-030-13705-2_13

Bautista, R. M. (2005). Clustering failed courses of engineering students using association rule mining. *Journal of Theoretical and Applied Information Technology, 96*(4), 875–886. https://www.jatit.org/volumes/Vol96No4/3Vol96No4.pdf

Çiçekli, U. G., & Kabasakal, I. (2021). Market basket analysis of basket data with demographics: A case study in e-retailing. *Alphanumeric Journal, 9*(1), 1–12. https://doi.org/10.17093/alphanumeric.752505

Colorado State University Office of Institutional Research, Planning and Effectiveness. (2012). *Unsuccessful course completion and student success.* https://irpe-reports.colostate.edu/pdf/ResearchBriefs/Unsuccessful_Courses_Student_Success.pdf

Derouiche, M. (2024). *Association rule mining for MBA.* https://www.kaggle.com/code/mohammedderouiche/association-rule-mining-for-mba

Ezarik, M. (2023). *Defining, measuring and unifying approaches to student success.* https://www.insidehighered.com/news/student-success/life-after-college/2023/11/22/how-college-campus-professionals-speak-about

González, J., Romero, C., & Ventura, S. (2008). Data mining in course management systems: Moodle case study and tutorial. *Computers & Education, 51*(1), 368–384. https://doi.org/10.1016/j.compedu.2007.05.016

McKinsey & Company. (2023). *Strengthening higher education outcomes through partnerships, alliances, and mergers.* https://www.mckinsey.com/industries/education/our-insights/higher-ed-is-consolidating-transforming-the-sector

Michaels, K., & Milner, J. (2021). *Powered by publics learning memo: The big ten academic alliance cluster exploring foundational course DFW rates, equity gaps, and progress to degree.* https://www.aplu.org/wp-content/uploads/powered-by-publics-learning-memo-the-big-ten-academic-alliance-cluster.pdf

Mwiti, D. (2025). *Apriori algorithm explained: A step-by-step guide with Python implementation.* https://www.datacamp.com/tutorial/apriori-algorithm

Papadogiannis, I., Wallace, M., & Karountzou, G. (2024). Educational data mining: A foundational overview. *Encyclopedia* 2024, *4*(4), 1644–1664. https://doi.org/10.3390/encyclopedia4040108

Romero, C., & Ventura, S. (2020). Educational data mining and learning analytics: An updated survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 10*(3), e1355. https://doi.org/10.1002/widm.1355

Safour, H., Essgaer, M., Alshareef, A. (2024). Unraveling academic failure: Examining the influence of course correlations on student performance through association rules algorithms. *2024 International Conference on Computer and Applications (ICCA)*, Cairo, Egypt. https://doi.org/10.1109/ICCA62237.2024.10927837

Slim, A., Heileman, G., Al-Doroubi, W., & Abdallah, C. (2016). The impact of course enrollment sequences on student success. *2016 IEEE 30th International Conference on Advanced Information Networking and Applications (AINA)*, Crans-Montana, Switzerland. https://doi.org/10.1109/AINA.2016.140

Soika, B. (2021). *What is student success? New insight into a complex question.* https://rossier.usc.edu/news-insights/news/what-student-success-new-insight-complex-question

Zong, C., & Koller, S. (2023). *The relationship between high-risk courses and Fall-to-Fall retention of first-time full-time students at UW* [University of Wyoming]. https://www.uwyo.edu/oia/_files/research-reports/high-risk-course-research-report-oct-2023.pdf

ABOUT THE AIR PROFESSIONAL FILE

The AIR Professional File features journal-length articles grounded in relevant literature that synthesize current issues, present new processes or models, or share practical applications related to institutional research. All submissions are peer-reviewed. For more information about AIR Publications, including *The AIR Professional File* and instructions for authors, visit www.airweb.org/publications.

ABOUT AIR

The Association for Institutional Research (AIR) empowers higher education professionals to leverage data, analytics, information, and evidence to make decisions and take actions that benefit students and institutions and improve higher education. For over 50 years, AIR has been committed to building and sustaining strong data-informed decision cultures within higher education institutions.

Association for Institutional Research
Christine M. Keller, Executive Director and CEO
www.airweb.org
ISSN 2155-7535

THANK YOU

Authors who submit materials to AIR Publications receive feedback from members of the AIR community. AIR expresses sincere appreciation for the individuals who serve as peer reviewers. The following individuals reviewed manuscripts recently submitted to *The AIR Professional File*.

Tracey Baham
Eddie Bolden
Mike Bolen
Mary Ann Coughlin
Soohong Kim
Michael Littrell
Amirtha Prakash
Stephen Terry
David Vintinner
Osundwa Wanjera
Xilin Zhang
Linli Zhou