

# Predictive Analytics in Higher Education: The Promises and Challenges of Using Machine Learning to Improve Student Success

---

Kelli Bird

## About the Author

Kelli Bird is a Research Assistant Professor at the University of Virginia's School of Education and Human Development.

## Acknowledgments

I am grateful to Ben Castleman, whose continued collaboration and mentorship has been instrumental in my exploration of this topic. I am also grateful for the long-standing partnerships with the Virginia Community College System and Bloomberg Philanthropies that have enabled us to do this work.

## Abstract

Colleges are increasingly turning to predictive analytics to identify “at-risk” students in order to target additional supports. While recent research demonstrates that the types of prediction models in use are reasonably accurate at identifying students who will eventually succeed or not, there are several other considerations for the successful and sustained implementation of these strategies. In this article, I discuss the potential challenges to using risk modeling in higher education and suggest next steps for research and practice.

# INTRODUCTION

With persistently low retention and graduation rates at many colleges and universities, higher education administrators are increasingly looking for innovative ways to improve student success outcomes. As a result, predictive analytics are increasingly pervasive in higher education (Ekowo & Palmer, 2016). The most common and arguably the most impactful application of predictive analytics is to use a prediction model to identify students who are at risk of doing poorly in a course or of leaving college without completing, and to intervene with these students early before they are too far off track.<sup>1</sup> For instance, more than half of colleges and universities report using “statistical modeling to predict the likelihood of an incoming student persisting to degree completion” (Ruffalo Noel Levitz, 2021, p. 22). Once the at-risk students have been identified by the prediction model, then faculty or staff proactively reach out to these students with offers of additional supports, such as academic advising or tutoring. While these types of resources are typically available to students upon request (though perhaps at limited capacity), many students do not take advantage of them. Since colleges do not typically have the resources to provide all students with these extended supports—at the median community college, academic advisors are responsible for 2,000 students (Carlstrom & Miller, 2013)—the goal of predictive analytics is for colleges to efficiently target the resources to students who need the resources to succeed. I will refer to this application of predictive analytics as “risk modeling and resource targeting” throughout this article.

To administrators who have been searching for solutions to improve student success, risk modeling and resource targeting are tempting solutions. Because colleges often lack the analytic capacity to implement these methods, private industry has stepped in with solutions, and those solutions are now a \$500 million industry. Roughly a third of colleges and universities have bought predictive analytics products, with each institution spending approximately \$300,000 per year (Barshay & Aslanian, 2019). Despite this investment, however, there is no rigorous evidence to show that these methods (either proprietary or in-house applications developed by colleges themselves) are successful at improving student outcomes.<sup>2</sup> What’s more, there are concerns that racially biased algorithms or poorly executed messaging could exacerbate, instead of mitigating, existing inequities (Acosta, 2020; Angwin et al., 2016; Burke, 2020; Engler, 2021). In this article, I will discuss the promises of predictive analytics in higher education, the challenges of predictive analytics (human vs. machine), obstacles to effective implementation, and recommendations for next steps for research and practice.

## PROMISES OF PREDICTIVE ANALYTICS IN HIGHER EDUCATION

While the current research is lacking in rigorous evaluations of the impact of risk modeling and resource targeting on student success, an increasing body of literature demonstrates that algorithms can achieve relatively high levels of accuracy at

1. Colleges also use predictive analytics for enrollment management purposes, such as identifying high-target students for recruitment or offering generous financial aid packages. These enrollment management practices are designed to bolster the quality of a colleges’ incoming class. In this article, I choose to focus on predictive analytic applications designed to support at-risk students.

2. Still, there are several anecdotes to suggest that current applications risk modeling and resource targeting are leading to improved student outcomes. Most notably is Georgia State University (GSU), which reports an 8-percentage-point increase in its graduation rate since implementing EAB’s predictive analytics products. This implementation accompanied several other changes at the university, however (Swaak, 2022).

predicting student success. For a recent cohort of high school seniors, my colleagues and I compared the accuracy of a relatively simple logistic regression model with the students' professional college advisors at predicting the students' college enrollment outcomes (Akmanchi et al., 2023). We found that the logistic model is at least as accurate as the advisors for students who interacted with the advisors up to eight times. This is true even though advisors likely had much more pertinent information about the students' college search, such as the names of colleges where they had been admitted. In a separate line of work, my colleagues and I found that incorporating behavioral trace data from online learning management systems can significantly improve the prediction accuracy for new students—which is the population with the lowest retention rates and thus those for whom predictions could be most important (Bird et al., 2022). In recent University of Oregon applications, a more advanced machine learning (ML) algorithm (XGBoost) is roughly three times better at identifying at-risk students than relying on students' high school GPAs alone (Greenstein et al., 2023).

## CHALLENGES OF PREDICTIVE ANALYTICS: HUMAN VS. MACHINE

There are many challenges to successfully deploying risk modeling and resource targeting in higher education. However, as the research I briefly discuss above demonstrates, the main challenge will *not* be whether the algorithms (i.e., machine) are able to identify at-risk students better and more efficiently than humans. Instead, most of the challenges surround the question of how humans will use what

the output the machines provide. A quote from Pedro Domingos highlights this tension: "It's not man versus machine; it's man with machine versus man without. Data and intuition are like horse and rider, and you don't try to outrun a horse; you ride it." For humans to harness the machine effectively, it is important to remember two important distinctions. First, much like a horse and rider, the human and machine have different objectives when it comes to predicting which students are at risk. Humans (administrators, policymakers, researchers, etc.) have complex objectives of increasing student success, improving equity, and ensuring the longevity of the colleges and universities. The machine's objective is much simpler: to make the best predictions possible using the information provided. Second, the human and machine have different responsibilities. The humans have the responsibility to rely on context when building the prediction models, since there are many subjective decisions to be made regarding sample construction, outcome specification, and predictors to include. Humans must also investigate potential biases within models, which I will discuss below. Once the predictions have been made and at-risk students have been flagged, the machine's job is done, but the human's job is not: people must decide how to communicate to at-risk students and which additional supports to provide. This is no simple undertaking, and requires significant engagement with colleges' faculty and advising staff. Allison Calhoun-Brown at GSU highlights the importance of the human work: "The innovation is not the technology. The innovation is the change that accompanies the technology" (Calhoun-Brown quoted in Swaak, 2022). In other words, if we want to improve student success outcomes, it is not a question of if we use predictive analytics, but instead *how* we use it.

## OBSTACLES TO EFFECTIVE IMPLEMENTATION

One of the biggest obstacles that colleges face in implementing predictive analytics is effectively communicating to students (Acosta, 2020). You could imagine someone drafting this message: “Kelli, an algorithm flagged you as someone likely to fail English 101. Work hard to improve your grade.” This message raises several concerns. A recipient might be concerned about their data privacy: How is the college using their personal data to determine their likelihood of failing? This type of messaging could also reinforce stereotype threats of not being “good enough” or “college material,” and being labeled as likely to fail could become a self-fulfilling prophecy. Perhaps this message would be more appropriate: “Hi Kelli, this is Professor Smith. I noticed you’ve been interacting less frequently than some of your classmates. Let’s set up a time to talk about how you’re doing in the class.” This message puts more of a human touch on the outreach, does not lead with the idea of failure, and provides a concrete next step on which the student can act. My colleagues and I are currently working with social psychologists to design effective messaging for an upcoming pilot program, which I describe below. Simply getting the communication right is not sufficient, however. Several recent low-touch nudge interventions with behaviorally informed messaging failed to improve student outcomes (e.g., Bird, Castleman, Denning, et al., 2021), so it is also imperative for students to be connected to the right supports to meet their needs.

Another barrier to successfully implementing risk modeling and resource targeting is achieving buy-in from faculty and staff. Among colleges and universities using statistical modeling to predict

graduation, fewer than one-third of administrators thought it was a very effective strategy at improving student success (Ruffalo Noel Levitz, 2021). One of the reasons that faculty may distrust predictive analytics is their black box nature. Many prediction models in use are from third-party for-profit vendors; their proprietary nature means that institutions have little understanding of what goes on under the hood. A recent GAO report specifically calls out these higher education models as needing more scrutiny from both their consumers and from regulators (Bauman, 2022).

Humans also may find it difficult to incorporate risk modeling due to the impersonal nature of the machine. Prediction models inherently rely on information from a large historical sample and generate predictions to optimize the accuracy for the group as a whole, as opposed to considering potential nuance in a particular individual’s circumstance. In a recent pilot where my colleagues and I collaborated with a community college to improve transfer outcomes for their students, we incorporated an algorithm that generated personalized course recommendations that accounted for the probability that the student would succeed in the course. Despite significant collaboration on how the algorithm would select the courses to recommend, the advisors still changed roughly one out of three courses the algorithm had identified before communicating the recommendations to students.

Finally, many are concerned about the potential negative impacts of algorithmic bias to exacerbate, instead of mitigate, existing inequities. These concerns are not unfounded: several studies have found the existence of algorithmic bias in higher education prediction models (e.g., Baker &

Hawn, 2021; Yu et al., 2020).<sup>3</sup> When my colleagues and I investigated algorithmic bias in two models predicting course completion and degree completion among community college students, we find evidence of meaningful bias (Bird et al., 2023). Specifically, we find that the calibration bias present in the models would lead to roughly 20% fewer at-risk Black students receiving additional supports, compared with a simulated unbiased model.<sup>4</sup> Our exploration suggests that this bias is driven not by the inclusion of race or socioeconomic information as model predictors, but instead by success being inherently more difficult to predict for Black students. This result may reflect structural racism in K–12 education systems where many Black have access to fewer advantages. Specifically, model predictors based on past performance reflect that unequal circumstances would not be as powerful to predict a disadvantaged student's full potential. The algorithmic bias is particularly prevalent among new students for whom there is very little baseline information, suggesting that additional pre-matriculation data collection could mitigate bias in this case. We also find that the amount of algorithmic bias—and the strategies for mitigating the bias—can vary substantially across models; it is therefore imperative to address bias on a case-by-case basis.<sup>5</sup>

## RECOMMENDATIONS FOR NEXT STEPS FOR RESEARCH AND PRACTICE

First and foremost, we need rigorous evaluations of different strategies that incorporate predictive

analytics. My colleagues and I are planning a pilot program that we will evaluate through a randomized control trial, with three experimental conditions: (1) control (i.e., business as usual); (2) early-term predictions, in which community college instructors will be informed which of their students a prediction model flagged as being at risk, with the instructors receiving training in how best to communicate with those students; and (3) early-term predictions plus additional embedded course supports. We include the third condition recognizing that community college instructors likely face meaningful constraints in the additional supports they can provide students on their own. While randomized control trials are the gold standard of research, they are not the only rigorous method. For institutions interested in evaluating their predictive analytic applications, there are many researchers, including me, who would be happy to collaborate on designing a quasi-experimental study.

Another important topic for future research is to better understand which point(s) in the distribution of predicted risk would be most effective and efficient for intensive resource targeting. While students are typically lumped into categories based on their risk (e.g., two categories: at risk or on track; three categories: green, yellow, or red), the raw model output is a continuous predicted risk score ranging from zero to one. An immediate thought may be to target the students at highest risk, meaning those least likely to succeed. However, it might be quite difficult to get these students to engage with additional supports, and they may not have a high likelihood of success even when they are targeted. So perhaps students at a more moderate

3. Algorithmic bias has been found in other predictive analytic applications outside higher education, including criminal justice and health care (Angwin et al., 2016; Obermeyer et al., 2019).

4. Calibration bias occurs when, conditional on predicted risk score, subgroups have different actual success rates. In our application, this means that, at a particular point in the distribution of predicted risk scores, Black students have a higher success rate than White students.

5. Our related work also suggests that small changes in modeling decisions (e.g., choosing logistic regression versus XGBoost as the prediction model) can significantly change the sorting of students within the risk score distribution, and therefore have the potential to significantly alter which students would receive additional supports (Bird, Castleman, Mabel, et al., 2021).

risk level, or students just at the margin of success, would be a more appropriate targeting strategy. It is not clear where in the distribution of risk we would expect to see the most bang for the buck in terms of resources moving students from failure to success; thus future research could significantly improve the cost-effectiveness of risk modeling and resource targeting. It is important to note that the answer to this question will almost certainly be context-dependent. For example, at more-selective colleges with higher persistence and graduation rates, the best strategy would likely target those with the highest risk scores; at broad or open access institutions, however, there is a much wider range of students who could benefit from additional resources. Institutional research (IR)/institutional effectiveness (IE) professionals who are involved in institution assessment are positioned well to contribute important context of student success that would not only inform the design of student success supports tied to the risk models, but also estimate the institution's return on investment of these additional resources.

I also believe that ML has the potential to improve how we structure the targeted students supports. Struggling students have a variety of different needs that may be inhibiting their success: lack of academic preparedness, financial constraints, inflexible schedules, unfamiliarity with administrative processes, and so on. So how do we connect students to the right supports that they need? ML methods commonly used in the private sector such as market basket analysis (Aguinis et al., 2013) have a lot of potential to inform this question, although it would require colleges to invest in the collection of student support usage data. IR officials who are involved in campus-wide data governance could help colleagues think about data collection, management, and analytic uses of this and other student data,

including the integration of this data collection into existing learning management systems or student success platforms that already track several other student behaviors (e.g., Blackboard).

Finally, it is imperative for us as an education research community to develop standards for ethical considerations relevant to these applications. Researchers and policymakers are increasingly recognizing the need for transparency and student rights with regard to artificial intelligence (AI) in education (e.g., Holmes & Tuomi, 2022; U.S. Department of Education, 2023), though additional considerations should be given to the technical aspects of algorithmic bias. There are many metrics that could be used to determine whether a model is generating fair predictions, and the choice of metric is critical since they can be at odds with each other (Kleinberg et al., 2016). In the paper I describe above (Bird et al., 2023), my colleagues and I chose to focus on calibration bias because we thought the most important type of bias in this application would be at-risk students from underrepresented or minoritized groups who are less likely to receive additional supports, compared to at-risk students from majority groups. However, this metric is less appropriate for an application where at-risk students are counseled out of college majors that are associated with the highest earnings (e.g., Barshay & Aslanian, 2019). We also need to develop standards for what level algorithmic bias is acceptable since reducing bias often leads to decreases in overall model accuracy, and it may not be feasible to achieve zero bias while still maintaining a high-performing model.

At this time predictive analytics has shown its promise at efficiently identifying at-risk students; with the possible inclusion of more-detailed data from learning management systems, these

predictions will only improve (Bird et al., 2022). Still, there is much important work to be done to both unlock its full potential and to ensure its safe use. Before risk modeling practices and applications that use predictive analytics become too ingrained in our colleges and universities, it is critical that we use the momentum fueled by the various discussions I mention above to ensure a fruitful future for predictive analytics in higher education.

## REFERENCES

- Acosta, A. (2020). *How you say it matters: Communicating predictive analytics findings to students*. Report, New America. <https://www.newamerica.org/education-policy/reports/how-you-say-it-matters/>
- Aguinis, H., Forcum, L. E., & Joo, H. (2013). Using market basket analysis in management research. *Journal of Management*, 39(7), 1799–1824. <https://journals.sagepub.com/doi/10.1177/0149206312466147>
- Akmanchi, S., Bird, K. A., & Castleman, B. L. (2023). Human versus machine: Do college advisors outperform a machine-learning algorithm in predicting student enrollment? EdWorkingPaper 23-699, Annenberg Institute at Brown University. <https://doi.org/10.26300/gadf-ey53>
- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. *International Journal of Artificial Intelligence in Education*, 32, 1052–1092. <https://doi.org/10.1007/s40593-021-00285-9>
- Barshay J., & Aslanian S. (Hosts). (2019, August 6). Under a watchful eye: Colleges are using big data to track students in an effort to boost graduation rates, but it comes at a cost. [Audio podcast episode]. The Educate Podcast, [apmreports.org. https://www.apmreports.org/story/2019/08/06/college-data-tracking-students-graduation](https://www.apmreports.org/story/2019/08/06/college-data-tracking-students-graduation)
- Bauman, D. (2022, June 3). Congress should scrutinize higher ed's use of predictive analytics, watchdog says. *Chronicle of Higher Education*. [https://www.chronicle.com/article/congress-should-scrutinize-higher-eds-use-of-predictive-analytics-watchdog-says?cid2=gen\\_login\\_refresh](https://www.chronicle.com/article/congress-should-scrutinize-higher-eds-use-of-predictive-analytics-watchdog-says?cid2=gen_login_refresh)
- Bird, K. A., Castleman, B. L., Denning, J. T., Goodman, J., Lambertson, C., & Rosinger, K. O. (2021). Nudging at scale: Experimental evidence from FAFSA completion campaigns. *Journal of Economic Behavior & Organization* 183(March), 105–128. <https://www.sciencedirect.com/science/article/abs/pii/S0167268120304819?via%3Dihub>
- Bird, K. A., Castleman, B. L., Mabel, Z., & Song, Y. (2021). Bringing transparency to predictive analytics: A systematic comparison of predictive modeling methods in higher education. *AERA Open*, 7. <https://journals.sagepub.com/doi/10.1177/23328584211037630>
- Bird, K. A., Castleman, B. L., & Song, Y. (2023). Are algorithms biased in education? Exploring racial bias in predicting community college student success. EdWorkingPaper 23-717, Annenberg Institute at Brown University. <https://www.edworkingpapers.com/ai23-717>

- Bird, K. A., Castleman, B. L., Song, Y., & Yu, R. (2022). Is Big Data better? LMS data and predictive analytic performance in postsecondary education. EdWorkingPaper 22-647. Annenberg Institute at Brown University. <https://www.edworkingpapers.com/ai22-647>
- Burke, L. (2020, December 13). The death and life of an admissions algorithm. *Inside Higher Ed*. <https://www.insidehighered.com/admissions/article/2020/12/14/u-texas-will-stop-using-controversial-algorithm-evaluate-phd>
- Carlstrom, A. H., & Miller, M. A. (Eds.). (2013). 2011 NACADA national survey of academic advising. Monograph No. 25, Global Community for Academic Advising. <https://nacada.ksu.edu/Resources/Clearinghouse/View-Articles/2011-NACADA-National-Survey.aspx>
- Ekowo, M., & Palmer, I. (2016). The promise and peril of predictive analytics in higher education: A landscape analysis. Policy paper, New America. <https://www.newamerica.org/education-policy/policy-papers/promise-and-peril-predictive-analytics-higher-education/>
- Engler, A. (2021). Enrollment algorithms are contributing to the crises of higher education. Part of series "AI Governance." The Artificial Intelligence and Emerging Technology (AIET) Initiative, Brookings Institution. <https://reachhighscholars.org/Articles/Enrollment%20Algorithms%20-%20Brookings%20Institution%209.14.21.pdf>
- Greenstein, N., Crider-Phillips, G., Matese, C., & Choo, S-W. (2023). Predicting risk earlier: Machine learning to support success and combat inequity. Academic data analytics briefing document. Office of the Provost, University of Oregon. <https://provost.uoregon.edu/analytics/student-success>
- Holmes, W., & Tuomi, I. (2022). State of the art and practice in AI in education. *European Journal of Education*, 57, 542-570. <https://onlinelibrary.wiley.com/doi/10.1111/ejed.12533>
- Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. Cornell University. <https://doi.org/10.48550/arXiv.1609.05807>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453. <https://www.science.org/doi/10.1126/science.aax2342>
- Ruffalo Noel Levitz. (2021). *2021 effective practices for student success, retention, and completion report*. Ruffalo Noel Levitz. <https://www.ruffalonl.com/papers-research-higher-education-fundraising/student-retention-practices-report/>
- Swaak, T. (2022, August 26). How higher ed is trying to improve student performance with data. PBS News Hour. <https://www.pbs.org/newshour/education/how-higher-ed-is-trying-to-improve-student-performance-with-data>
- U.S. Department of Education. (2023). Artificial intelligence and future of teaching and learning: Insights and recommendations. Office of Educational Technology. <https://tech.ed.gov/ai-future-of-teaching-and-learning/>
- Yu, R., Li, Q., Fischer, C., Doroudi, S., & Xu, D. (2020). Towards accurate and fair prediction of college success: Evaluating different sources of student data. In A. N. Rafferty, J. Whitehill, V. Cavalli-Sforza, & C. Romero (Eds.), *Proceedings of the 13th International Conference on Educational Data Mining* (pp. 292-301). <https://files.eric.ed.gov/fulltext/ED608066.pdf>