Developing a Single Tool for Assessing Student Retention Interventions

Jenny Zhu
Nathan Dickmeyer
Institutional Research & Assessment
LaGuardia Community College
May 31, 2016

History
At the request of the college's provost several years ago, we used the information in our data warehouse to predict which students were most likely to drop out. This effort informed us about which data elements were most useful for predicting college persistence. Nevertheless, we realized that these techniques had limited accuracy at the individual student level. The characterization "likely to drop out" becomes increasingly inaccurate as the size of target population increases and we go down the probability scale from a high probability of dropping out to a middling one. Most students with the 50 highest probabilities of dropping out will drop out, but a smaller proportion of the 500 highest will drop out.

In the development of regression equations to predict college success, we tried a number of dependent variables, including a semi-continuous outcome that scored dropping out after one semester as a zero and retention to graduation as one with longer retention lengths and early transfer to another college scored in between. Of particular interest, however, was two-semester (fall-to-fall) return rate. From one fall to the next, the college loses more than one-third of its students. Dropping out is reasonably likely, and predicting who will drop out is not like finding a needle in a haystack. This dependent variable also allowed us to use stepwise logistic regression, resulting in each degree student enrolled in the baseline fall semester having a probability score for return to the next fall.[1]

We then realized that, while we were always somewhat wrong when we selected a group for an intervention (assuming that they *all* were likely to drop out), we could average the probabilities for all members of the group and get a predicted proportion who would drop out, based on historical data. That is, for a selection of 100 students we could use historical data to generate a probability of two-semester return for each and, thus, an average return rate for the whole group. This opened the door for us to assess any intervention that had an impact during an academic year. All students in the intervention would have a probability of two-semester return based on a historical analysis of student characteristics available in our data warehouse.

[1] Note: as a convention in our office, we use the term "retention rate" to pertain to the return rate of a cohort and the term "return rate" to pertain to any mixed group of students, not all of whom may be at the same point in their academic career. Retention rates are more comparable than return rates because the mix of students involved in a return rate (for example, the number of new as opposed to continuing students), which can vary when a cohort is not used, can affect the rate of return. In the use of return rates in this paper, however, comparability is preserved by comparing a prediction to an actual for the same mixed group.

Using the average of these probabilities, the intervention group as a whole had a predicted return rate. If the intervention was successful, this group of students should be retained at a rate higher than the predicted rate.

Setting a target to indicate the success of an intervention

Thus, we had a way of assessing any intervention in which degree students were enrolled in the fall by averaging the predicted return probabilities of the students in the baseline fall semester and comparing that prediction with the actual result two semesters later. At the start of the intervention, we could take the predicted average and add an amount to give us assurance of statistical significance and call this a target for the intervention. In the next fall semester we would then count the number of returning students and add to that the number who had graduated at the end of the previous fall and spring semesters. If the proportion of those who returned or graduated was higher than the target, then we could declare the program successful.

We decided to build the prediction using return data from the previous fall semester. This gave us over 16,000 individual records. In every case we had data on the beginning state of the student, like cumulative GPA, and on the ending state, either retained (enrolled in or graduated before the second fall) or not retained. (Note, this means that early transfer and stopping out during the follow-up fall are defined as failures on the part of the intervention. Not all interventions may want to define success this way.)

Table 1 shows the results of the regression and how it might be translated into a probability of two-semester return for a fictional student. Fifteen variables proved to contribute to the probability. Fall 2013 data were used to predict the return (or graduation) of fall 2014 students to fall 2015. The fit of the model to the historical data can be described as good, but not at all perfect (P > chi-square .18).

Factors affecting probability of return/grad Fall 2014 to Fall 2015	Percentage	Sample	Joe's
(Based on actual data from Fall 2012, Spring 2013 and Fall 2013)	Point Impact	Student Joe	Prediction
Starting point	-28.13		-28.13
Each year of age	0.646	20	12.92
Each credit earned	1.02	12	12.24
Each point of cumulative GPA	20.22	2	40.44
Each one percent of WU grades out of total grades received	-1.147	0	0
Being female	6.57	0	0
Not completing developmental math requirement	-16.81	0	0
Being full-time	42.42	1	42.42
Registering early	11.37	0	0
Being a continuing student	-10.66	1	-10.66
*Being a new student**	49.42	0	0
Being a new transfer student	24.02	0	0
Being an AA degree student	2.72	0	0
Being an AS degree student	13.79	1	13.79
Being an AAS degree student	-6.04	0	0
Not being on a student visa	-18.87	0	0
(Joe's predicted probability of returning Fall 2015)			**83.02%**
**Comparing a new student with a continuing student with a 2.00 GPA and 12 earned credits, the new student is missing 52.68 percentage points and already is behind more than 10 percentage points.*			

Table 1

Nevertheless, the chi-square test only shows how well the model predicts individual performance in the historical data. Our concern is on how well the model predicts future group performance.

For all 15,935 fall 2014 degree students, the model predicted return/graduation at 61.6%. Actual to fall 2015 was 64.3%. There are two possible causes of the variance: 1) differences between modeling data and scored data; or 2) college interventions that improved upon predictions. That is, either the relationship between, for example, GPA and return was weak and changed over time, or, something the college did changed the relationship between GPA and return.

The latter is most likely true. The college's three-year graduation rate for full-time, first-time students, for example, jumped from the fall 2011 cohort to the fall 2012 cohort from 16.3% to 20.0%. The fall 2012 cohort had the benefit of the 2014-15 academic year, the year of the prediction for the fall 2014 students in the model. A jump in graduation rates indicates improved programs administered by the college, causing the variance from model predictions noted above.

Use

These return targets were then used to uniformly assess various programs (and other group designations) at the college. Figure 1 shows some of the results.

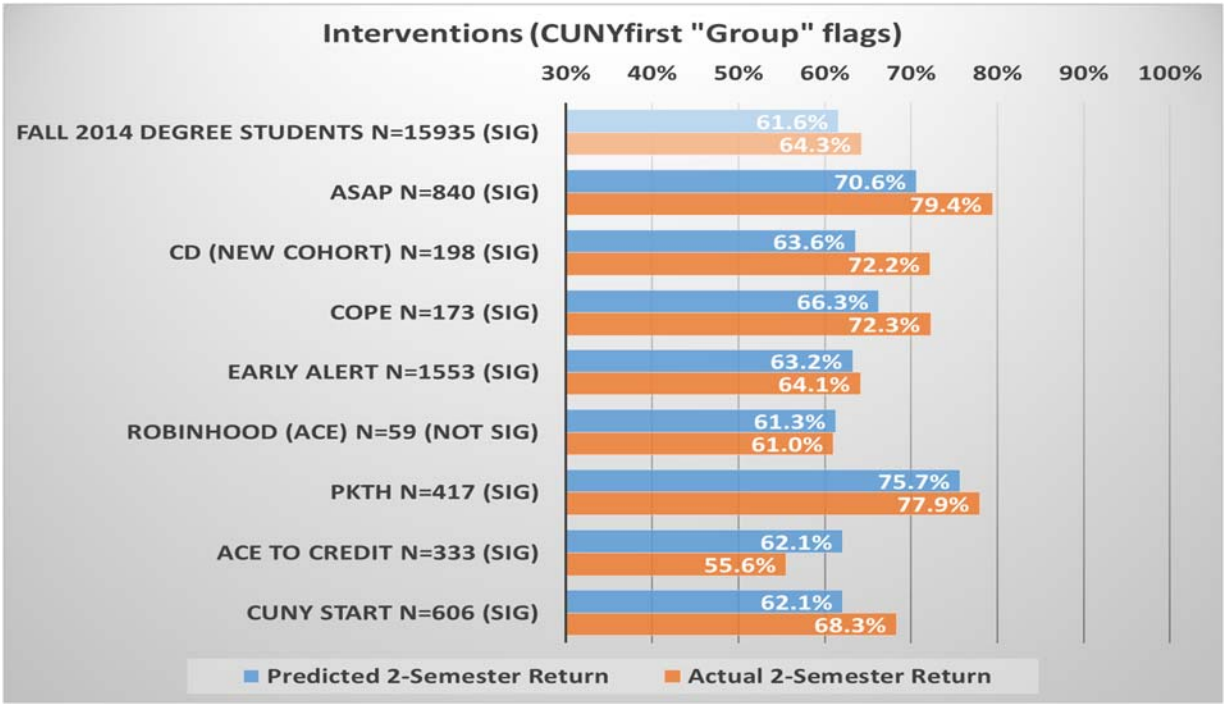

Figure 1

Figure 1 shows that the ASAP and CD program performed above target, while the “ACE to Credit” program did not. Other programs did not show as much improvement as the college as a whole, including Robinhood and the honors society, PKTH. Note that the students in the PKTH group already had high return probabilities, making improvement difficult.

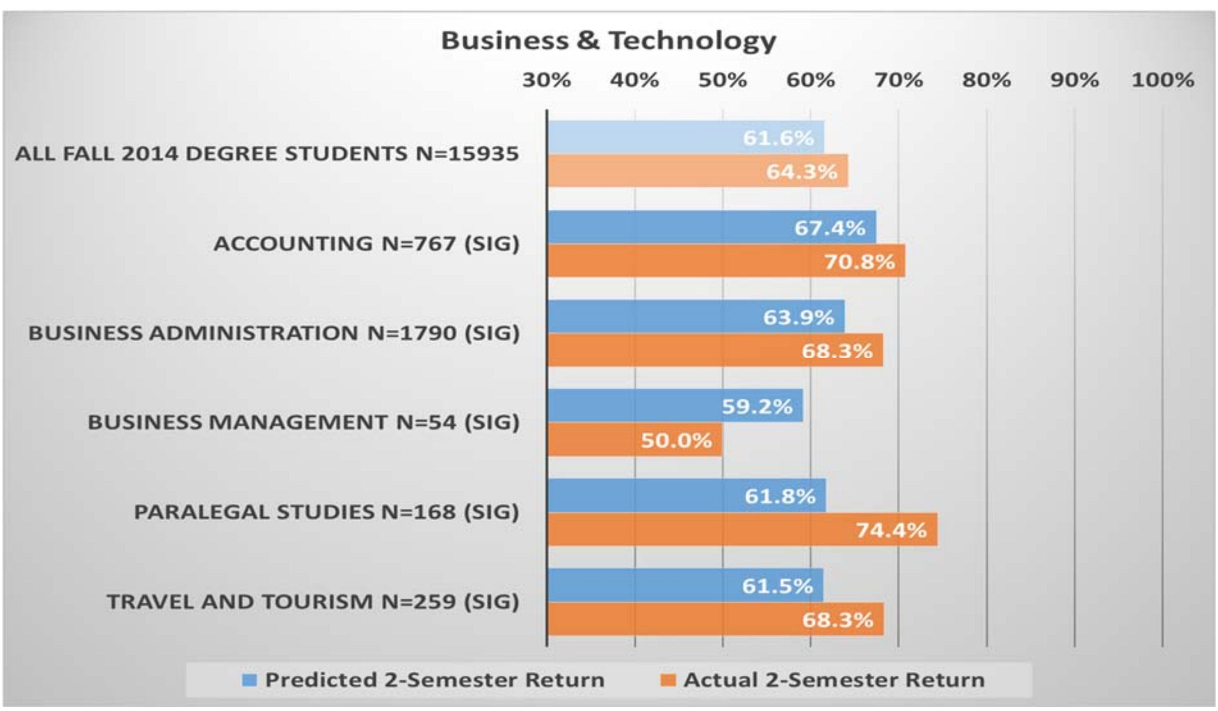

Figure 2

Figure 2 shows the same technique used to assess students grouped by advising teams (usually majors) for one of the college's divisions. Note, for example, that a major being phased out (Business Management) shows very poor results against the prediction, while Paralegal Studies shows very good results.

2015-16

For the current year, we increased the historical data to five semesters, yielding over 65,000 records. The two-semester return data on students from fall 2012, spring 2013, fall 2013, spring 2014, and fall 2014 were used to predict return of fall 2015 students to fall 2016. A somewhat different probability equation was derived. The fall 2015 prediction coefficients are compared to the fall 2014 in Table 2.

	Fall 2014 Factors	**Fall 2015 Factors**
Factors	**Percentage Point Impact (3 semesters data)**	**Percentage Point Impact (5 semesters data)**
Starting point	-28.13	-6.49
Each year of age	0.646	0
Each credit earned	1.02	0.91
Each point of cumulative GPA	20.22	22.26
Each one percent of attempted equated credits earned	0	0.1695
Each one percent of WU grades out of total grades received	-1.147	-1.28
Being female	6.57	8.58
Not completing developmental math requirement	-16.81	-7.42
Being full-time	42.42	36.33
Registering early	11.37	12.72
Being a continuing student	-10.66	-10.73
Being a new student	49.42	36.68
Being a new transfer student	24.02	31.37
Being an AA degree student	2.72	3.12
Being an AS degree student	13.79	12.19
Being an AAS degree student	-6.04	-3.41
Not receiving financial aid	0	-3.25
Not being on a student visa	-18.87	-23.24

Table 2

Note that age showed up in the 2014 predictive model, but not in the 2015 model, while the percentage of credits earned and receiving financial aid show up as significant in 2015 and not in 2014. There may be some degree of correlation among the three variables, causing some of them to drop out, when the others prove to improve the fit better. In this case, whichever combination works best would seem to be justified. Fortunately, none of the other 2015 variables switch signs or are completely out of range of the 2014 variables. This adds to our confidence in the validity of the model.

Summary

LaGuardia College now has a method for assessing the two-semester return impact on any group of students included in an intervention program that has an impact during the academic year. As soon as the target population is identified, the IR&A office (the data are confidential, so the data are circulated only on a need-to-know basis to advisors) can quickly calculate a target number of students who must be retained through to the next fall to demonstrate program effectiveness.

A percentage improvement over the target score allows comparison among various programs. In general, the college is finding that resources used with low-risk students may have more impact when used on high-risk students, where the opportunities for turning around students are greater.

In continuing this line of research, we have proposed a number of refinements for future work:

1) Developing similar formulas for spring to spring return.
2) Separating the probability calculations between new and continuing students to allow the variables to act fully independently on the two very different populations.
3) Testing more behavioral variables besides time of registration, perhaps including timing of first log-in to LaGuardia online services, counts of office visits, and previous absence rates.