

GIVECAMPUS


White Paper

Leveraging Machine Learning to Predict Donor Behavior and Drive Better Outcomes



Presented with a
foreword by CASE

October 2023



GIVECAMPUS

©2023 GiveCampus

Original publication date: October 2023

All rights reserved. No part of this material protected by this copyright may be reproduced or used in any form, electronic or mechanical, including photocopying or recording, or by any information storage and retrieval system, without written permission from GiveCampus.

Limit of liability/disclaimer. While the publisher has used its best efforts in preparing this document, it makes no representations or warranties with respect to the accuracy or completeness of its contents. No liability or responsibility of any kind (to extent permitted by law), including responsibility for negligence, is accepted by GiveCampus, its servants, or its agents. All information gathered is believed to be correct at publication date. Neither the publisher nor the author is engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

Author

Michael Firn, Senior Data Scientist, GiveCampus

ACKNOWLEDGMENTS

We would like to thank the educational fundraisers we partner with for contributing to this research. Special thanks to our pilot participants who were and continue to be so generous with their time and feedback.

We also extend our appreciation to CASE for providing a thoughtful foreword to our findings.

About GiveCampus

GiveCampus is the world's leading digital fundraising platform for education. Trusted by more than 1,300 colleges, universities, K-12 schools, and millions of donors, our mission is to advance the quality, the affordability, and the accessibility of education. We provide software, services, and expertise that help schools raise more money, from more people, at a fraction of the cost of other fundraising methods. For more information, please visit go.givecampus.com.

Table of Contents

Foreword	04
Introduction	05
About our research	05
A small disclaimer	06
Why targeted outreach matters	07
Using machine learning to predict donor behavior	08
What we found and what it means for you	08
1. Who should you ask to make a recurring gift?	09
2. Who should you ask to upgrade their gift?	13
3. Who should you ask to serve as an advocate?	16
How fundraisers can prepare for what's next	19
Summary	19



Foreword

The Council for Advancement and Support of Education (CASE) is the world leader in helping advancement professionals in schools, colleges, and universities make data-informed decisions. Built upon the CASE Global Reporting Standards, the data and research provided through CASE Insights highlights both the deeper vision and understanding found through data and the actions you can take to expand your institution's impact.

This report from GiveCampus on the impact of using artificial intelligence to predict donor behavior is important and timely. While the largest, well-funded advancement teams around the world may have analytics capacity, this is not the norm. Most institutions are looking for more access to the artificial intelligence and insights these predictive models provide to help them use their data in deeper ways.

CASE has enjoyed long relationships with educational partners who are thought leaders and who support our members through direct service. We are partnering widely with our CASE Standards Champions, including GiveCampus, to ensure that we are promoting the same best practices, ethical principles, and reporting standards across the advancement profession. We are grateful to GiveCampus for their supportive partnership and for providing a contribution that moves the profession forward.

Cara Giacomini

VP of Data, Research and Technology
Council of Advancement and Support of Education

About CASE

CASE—the Council for the Council for Advancement and Support of Education—is a global, not-for-profit membership association with a vision to advance education to transform lives and society. CASE is the home for advancement professionals, inspiring, challenging, and equipping them to act effectively and with integrity to champion the success of their institutions. CASE defines the competencies and standards for the profession of advancement, and leads, and champions their dissemination and application for more than 97,000 advancement professionals at 3,100 member institutions in 80 countries.

Broad and growing communities of professionals gather under the global CASE umbrella. Currently, these professionals include individuals working in alumni relations, development and advancement services, communications, fundraising, government relations, and marketing.

Headquartered in Washington, D.C., CASE works across all continents from its regional offices in London, Singapore, and Mexico City to achieve a seamless experience for all its stakeholders, particularly its members, volunteers, and staff.

Introduction

The ability to predict donor behavior is the Holy Grail for fundraisers, and—with a little help from artificial intelligence (AI)—it may finally be within reach. GiveCampus has been training machine learning models to help anticipate donor behaviors and preferences so that fundraisers can better understand who may be most receptive to their mission and message, and then target their outreach accordingly.

In this white paper, we'll reveal what our research and modeling tells us so far about predicting donor behavior. We'll also suggest actionable ways educational fundraisers can start leveraging the data they already have to make the right asks to the right constituents at the right time.

About our research

GiveCampus employed a comprehensive predictive modeling methodology to arrive at the findings detailed in this white paper. Predictive modeling uses machine learning and data to forecast likely future outcomes based on historical and existing data. In this case, we leveraged five years of giving data from partner educational institutions to predict future donor behavior for the fiscal year beginning July 1, 2022 and ending June 30, 2023. A detailed accounting of our methodology follows.

Data Collection

In order to avoid leakage, features were built using data timestamped before the period in which labels were constructed. Specifically, labels were built during Fiscal Year 2023 (July 2022 through June 2023), while features were generated using information collected in the 5-year run-up to this window (July 2017 through June 2022). Data was anonymized and pooled across 295 GiveCampus partner institutions for which the full set of giving data was available stretching back to July 2017. From these giving histories a total of 125 features characterizing each constituent's philanthropic behavior were derived.

Data Preprocessing

To enhance model performance and prevent overfitting, principal component analysis (PCA) was applied to reduce dimensionality. This step ensured that the most salient features were retained while eliminating noise from the dataset.

Modeling

The data was randomly split into an 80% train set and a 20% test set to evaluate model performance. A grid search was conducted on the train set, employing Logistic Regression, Random Forest, and Gradient Boosting Trees algorithms. Hyperparameters for each algorithm were optimized through 4-fold cross-validation within the random 80% train set. The best-performing algorithm was selected based on cross-validation scores. Final performance was evaluated on the 20% test set to ensure that the model's predictive power could generalize to data not seen during training.

Performance Analysis

The models generated propensity scores ranging from 0 to 1, indicating the likelihood of each constituent to “convert” during the label window. To categorize constituents as predicted converters or non-converters, these propensity scores were thresholded based on the overall class balance in the training set. For example, if 5% of the training set converted, the top 5% of constituents in terms of prediction were designated as “predicted converters.” Precision was then measured within each bucket to determine the percentage of predicted converters that actually converted during the label window.

Interpreting the charts

This white paper includes two types of data visualizations designed to illustrate our findings:

- 1. Feature Data:** Feature data refers to the variables or attributes used to make predictions or classifications.
- 2. Performance Data:** Performance data refers to the information used to assess how well our model’s predictions matched the behavior actually observed among constituents.

Please note that percentages in charts were rounded to one decimal place to simplify visualizations.

A small disclaimer

Unlike predicting the weather, there’s no accurate forecast for fundraising ... yet. Truth be told, human nature is harder to predict than mother nature. But behavioral science does give us some clues as to why people do what they do.

And data modeling helps us simulate that behavior so we can make educated guesses (or in this case “machine-learned” predictions) about how individual constituents might respond to our solicitations. We can then use these predictions to prioritize fundraising efforts, tailor communication strategies, and allocate resources more effectively.

That said, even the best predictive models aren’t a hundred percent accurate—and that’s okay because philanthropic predictions don’t have to be perfect in order to be useful. As you’ll see in the pages that follow, **moving the predictability needle even a fraction of a percentage point can make a huge difference** in the number of donors you convert, advocates you recruit, dollars you raise, and more.

It’s worth noting too that predictive models are iterative and therefore improve over time. So, by continually collecting more data and retraining the model, we can achieve progressively more accurate results and better predictions as we go.

Why targeted outreach matters

Fundraising appeals are a powerful tool for advancement professionals, but now more than ever the battle for attention in the inbox—and mailbox—is fiercely competitive. It's increasingly difficult for your communications to stand out among the Amazons, Etsys, and Wayfairs of the world—to say nothing of the myriad subscription services and worthy nonprofits also angling for eyeballs.

Vying for your constituent's attention is also costly, so to be effective your outreach needs to be targeted. "Spray-and-pray" (i.e., sending as many communications to as many people as possible) is not the answer, and there are real costs to making the wrong ask to the wrong donor.

Mass communications not only eat up your time and marketing dollars, they run the risk of alienating your constituents. **It's one thing for your messages to go unnoticed, but quite another when potential donors simply unsubscribe.** Once you lose your audience, it's incredibly difficult to win them back—ask any marketer in any industry.

So who do you ask to give, advocate, or volunteer? To maximize the value of your appeals you need to find constituents who are likely to be receptive to them. This is a very fundamental and upstream problem: your giving experience might be top-notch, but it doesn't matter if you're not reaching the right people.

At GiveCampus, we believe that the best way to segment your communications is to identify those constituents who are already primed to convert—in other words, surface the people who are most inclined to say yes. We think we're in a unique position to tackle this, given the patterns we see across hundreds of schools and the volume of anonymized data that we can use to train our models.

Using machine learning to predict donor behavior

While there are a number of ways fundraisers already leverage data to anticipate donor behavior, predictive modeling can produce exponentially better results when done right. That's because when it comes to uncovering patterns in vast amounts of data, machines are simply better, faster, and more perceptive to changes over time than human beings.

Think of machine learning as a supercharged and automated version of the type of segmentation you already perform. For example, the strategy to target LYBUNT alumni who gave \$500+ last year is built on a simple "model" designed to predict who will be most receptive to your solicitation. Machine learning algorithms work in a similar way, but instead of analyzing a handful of criteria, predictive models analyze hundreds of variables and identify the precise combination likely to yield the best results.

Using machine learning, GiveCampus is training models to analyze donor data, identify patterns, and predict future behavior. The goal is to help schools better understand their constituents, tailor their fundraising efforts, and make more informed decisions to maximize their impact.

If educational fundraisers can identify constituents who are likely to be more receptive to their requests, they can:

- Boost the efficacy of their outreach and marketing efforts
- Reduce the negative consequences associated with a "spray-and-pray" approach to donor outreach—namely donor fatigue, email unsubscribe rates, and the cost of direct mail campaigns that target the wrong constituents
- Build more meaningful relationships with like-minded constituents who support their mission

What we found and what it means for you

For the purposes of this whitepaper, let's assume you're approaching a big giving day and working out the details of your marketing plan. You have set ambitious goals for your campaign, and your strategy for hitting them is three-fold:

1. Drive recurring gifts to help improve retention
2. Encourage existing donors to increase or upgrade the size of their gift to boost dollars raised
3. Leverage peer-to-peer advocacy to drive greater participation

While these are solid goals, proper segmentation will be key to achieving them. If you approach the wrong person with the wrong ask, or simply ask for too much, you may actually drive existing and potential donors away.

With this in mind, our data scientists trained machine learning models to effectively anticipate which constituents would be most likely to make a recurring gift, increase the size of their donation, or serve as an advocate. We detail our findings, and what they mean for you, below.

1. Who should you ask to make a recurring gift?

A recurring gift is an especially attractive donor vehicle because it represents sustainable revenue for fundraisers. Plus recurring gifts provide an easy, frictionless way for constituents to support initiatives they care about with small, regular, set-it-and-forget-it donations that get made automatically. It's easy for you and easy for your donors.

Although the initial 'ask price' is low the overall lifetime gift tends to be larger. In fact, the expected lifetime value of indefinite recurring gifts on the GiveCampus platform is nearly 10 times the value of a one-time donation.

While the benefit of soliciting recurring donations is clear, the universe of donors who are predisposed to give in this way is limited. At the baseline, only 1% of donors make a recurring gift. This seems like an incredibly small number, but 1% of all constituents in our database is actually a sizable cohort—and if you could zero in on only those most likely to commit to a recurring gift you could capture a substantial and reliable source of revenue for your institution.



What we found: The donors that our model identified as most likely to make a recurring gift actually went on to do so **140 times more often** than the donors who were identified as unlikely to make a recurring gift. Equipped with this kind of insight, it's easy to see how a school might focus fewer resources more effectively on recurring giving appeals that are hyper-targeted to this more inclined segment of constituents.

Performance Data

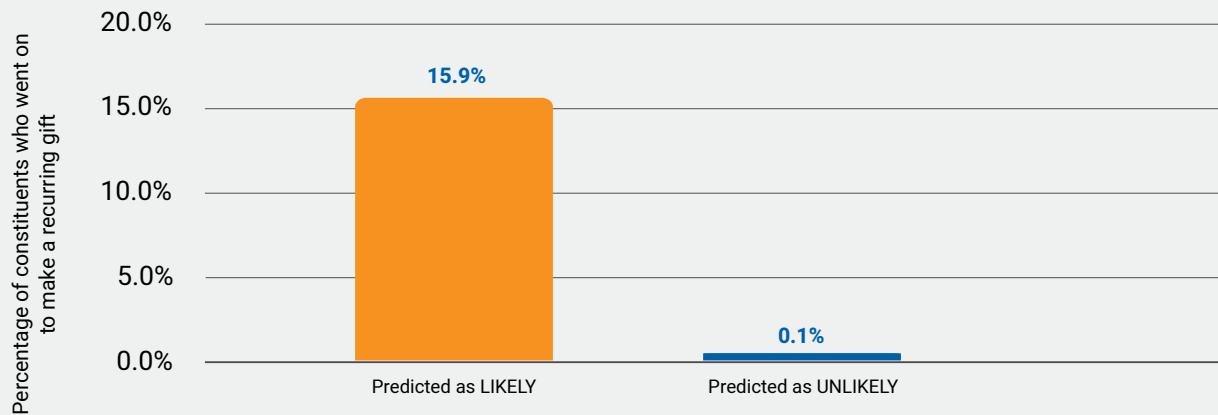


Figure 1. Likely to Make a Recurring Gift Performance Data

Predicting: Likelihood to make a recurring gift during Fiscal Year 2023

For: All constituents who don't already have an active recurring subscription

Still, our model was able to home in on this precise segment only after weighing, combining, and analyzing hundreds of variables. So, what can the average data-savvy fundraiser who doesn't have machine learning capabilities at their disposal do to get more targeted with their outreach? The short answer is that you can leverage the data that you already have to approximate what our machine model is doing, only on a human scale.



What this means for you: Start by identifying and critically thinking about which variables may be strong leading indicators of the behavior you're trying to predict. Using raw data that's readily available, like giving history for your donor base, you can begin to make some educated guesses about who is most likely to make a recurring gift.

For example, much can be gleaned from the frequency and size of donations a constituent makes consistently. One of the hundreds of feature sets that our model used to predict who was most likely to make a recurring gift included average gift size data for each constituent over a five-year period. If you have access to this type of data, you should be able to identify which donors have a history of making smaller, frequent gifts. These donors are more likely to embrace a subscription mindset and are therefore potentially more inclined to make a recurring donation.

Creating a targeted recurring gift appeal specifically for this segment is a far better strategy than tapping everyone in your CRM with the same ask. That’s because reaching out to the wrong constituents can have a negative impact. If they’re not inclined to give in this way they may tune you out completely—or worse—simply unsubscribe.

Key Takeaway: Little, successive gifts mean a lot. Donors with a history of making smaller, more frequent gifts may be more inclined to make a recurring donation.

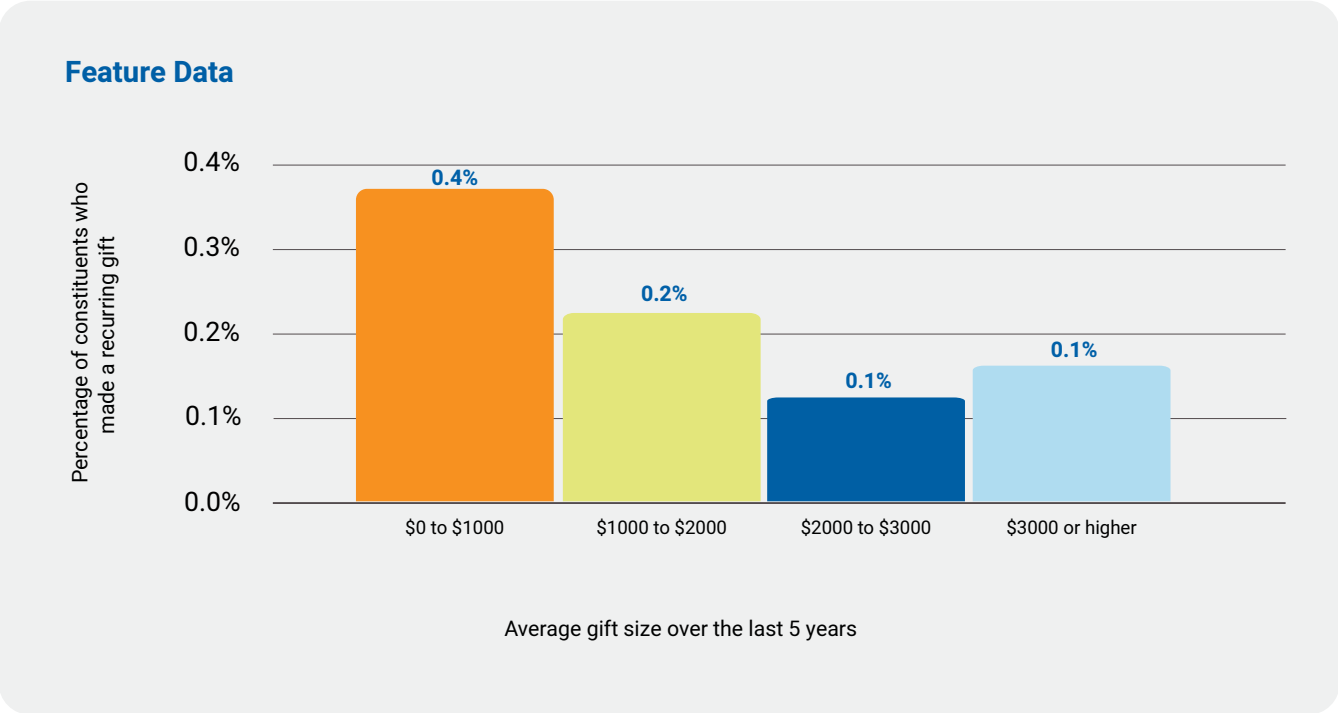


Figure 2. Likely to Make a Recurring Gift Feature Data
Tracking: Average gift size over a five-year period

But, for donors who are already making small regular gifts to your institution, the recurring model may be an attractive way to support your initiative, one that they'll appreciate and welcome. After all, the way they've given in the past is not unlike a subscription. By offering them the option to automate their gift, you're giving that cohort of loyal donors the opportunity to commit once for a set period of time—or perhaps indefinitely. And for those supporters who want to give on an ongoing basis there's really no easier way than making a recurring gift.

Fun Fact: Our annual recurring installment data shows that the most popular times to set up recurring gifts are in June at FYE and in December at CYE.



2. Who should you ask to upgrade their gift?

Encouraging those who donate consistently to increase their gift amount helps to deepen donor commitment and grow engagement over time. Of course in the short term, a fundraiser's ability to upgrade donor gifts also results in an immediate increase in revenue, which is the second goal of our hypothetical giving day.

So we trained our second predictive model to help us identify those donors who are most likely to say yes when asked to dig a little deeper into their pockets.



What we found: Our model looked at data on donors who had made at least 1 gift within the last 5 years, but had never given more than \$100 in a single transaction. We then built a model to predict the likelihood that each of these donors would go on to make a \$500+ gift within the coming year. Those donors that our model identified as likely to upgrade their gifts actually went on to do so **9.3 times more often** than those not identified as likely to upgrade.

Performance Data

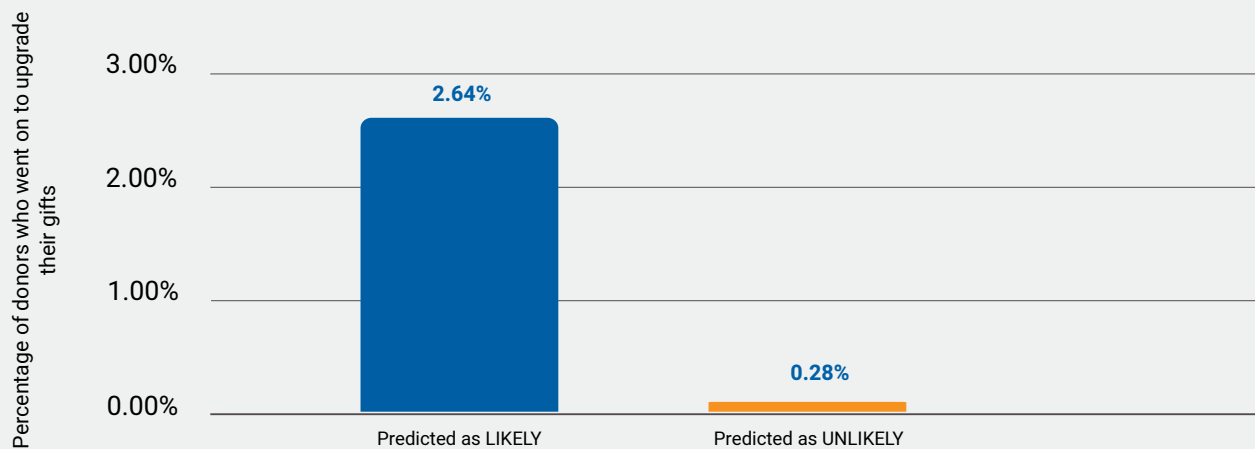


Figure 3. Likely to Upgrade Performance Data

Predicting: Likelihood to make a \$500 gift during Fiscal Year 2023

For: All constituents who have given in the past, but never more than \$100 in a single transaction



What this means for you: From the hundreds of data sets we used to train our predictive model, we can extrapolate a number of leading indicators that a donor may be both capable and inclined to increase the amount of their gift.

For example, donors who have consistently given more total dollars over successive years or have made higher-than-average gifts over successive years may be more receptive to an upgrade request. Also, as the feature data in Figure 4 below indicates, older donors are more inclined to make a larger gift when asked.

While these insights paint (in very broad strokes only) a portrait of our ideal upgrade candidate, the findings can also be used to identify who is most unlikely to be receptive to this ask. Knowing both who to target and who not to target is valuable data that you can use to inform your outreach strategy:

1. Create a segmented upgrade appeal targeting older donors with a history of making higher-than-average gifts. In this case, “higher than average” means donors who have made a gift of \$500 or more in the past 5 years.
2. Exclude or suppress younger donors who may not have the capacity to make a meaningful increase in contribution.



Feature Data

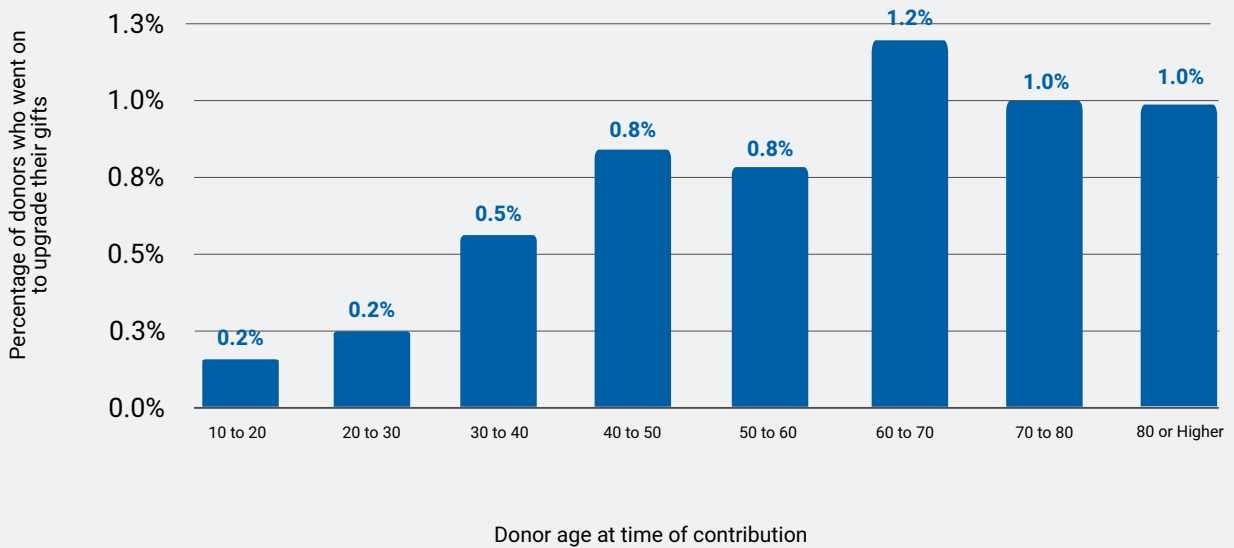


Figure 4. Likely to Upgrade Feature Data

Tracking: Donor age at contribution

Key Takeaway: Older donors who have a history of making higher-than-average gifts may be more receptive to an upgrade request.

3. Who should you ask to serve as a campaign advocate?

An advocate is someone who might share your campaign on social media or perhaps launch a match or challenge on behalf of your institution.

Peer-to-peer advocacy, our third giving day goal, plays a crucial role in the success of every online giving campaign. Advocates help to extend your reach, amplify your message, and grow your network of supporters. Sadly, it's no secret that the pool of people willing to advocate on your institution's behalf tends to skew small.

In fact, so few people enthusiastically raise their hand when asked, that school fundraisers inevitably end up asking those same few people to step up—over and over again. And understandably so. If someone has advocated before, chances are they'll do so again. Not surprisingly, our predictive model also arrived at this rather obvious conclusion.



What we found: At baseline, only .08% of people advocate. Of the people our model identified as likely to advocate, 21% went on to actually do so. In other words, **the constituents our model surfaced were 350 times more likely to advocate** than the average constituents. That said, this model included feature sets that tread familiar territory: namely folks who have a history of advocacy.

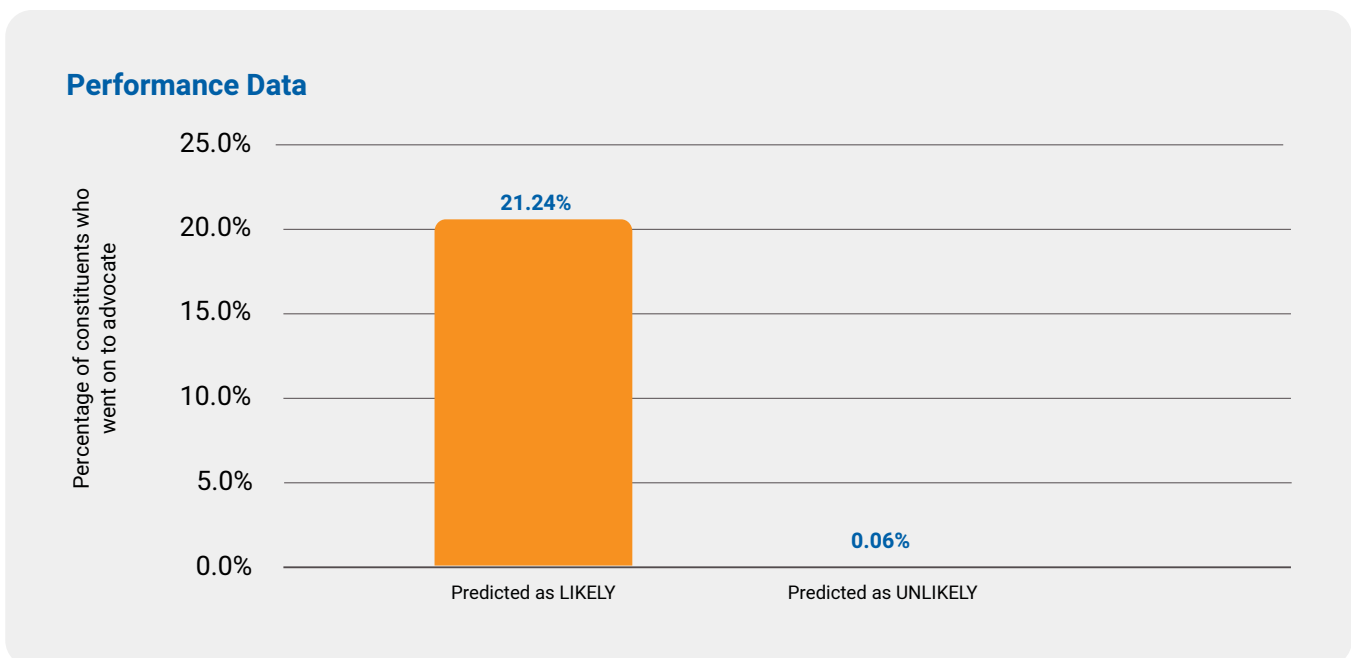


Figure 5. Advocacy Performance Data

Predicting: Likelihood to serve as a campaign advocate from Jan-1-23 to Jun-30-23
For: All constituents

Of course we all know that you can only go back to the same well so many times before it dries up—and one of the goals of predictive modeling is to prevent precisely that kind of fatigue by reaching out to the right people at the right time.

But how do you identify that rarest of breeds—the new advocate—someone who’s never advocated before? That’s what we set out to find with the following model.

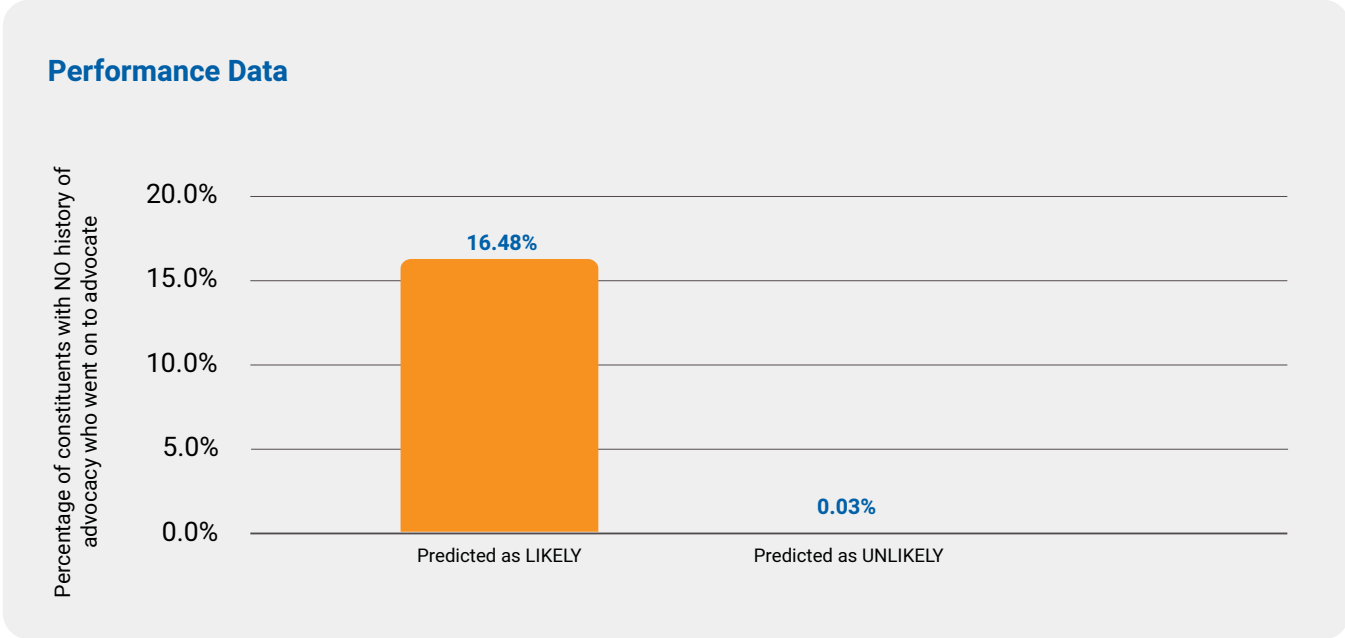


Figure 6. First-Time Advocacy Performance Data
Predicting: Likelihood to serve as a campaign advocate from Jan-1-23 to Jun-30-23
For: All constituents who have **never** served as a campaign advocate before

Here, we repeated the same exercise as above but removed constituents who had never advocated in the past. While the pool of potential advocates dropped from 21% to 16%, **that 16% still represents 540 times more advocates than those randomly selected.**



What this means for you: Because advocates tend to be your most enthusiastic supporters, they’re relatively easy to spot in the wild. Often they are among the very first donors to make a gift during a campaign. These supporters also feel personally invested, so much so that they keep tabs on your progress toward your goal and may even make additional gifts to help nudge your initiative across the finish line.

One of the many feature sets our predictive model included, that you could easily replicate, identified constituents who had made multiple gifts to the same social fundraising campaign.

Feature Data

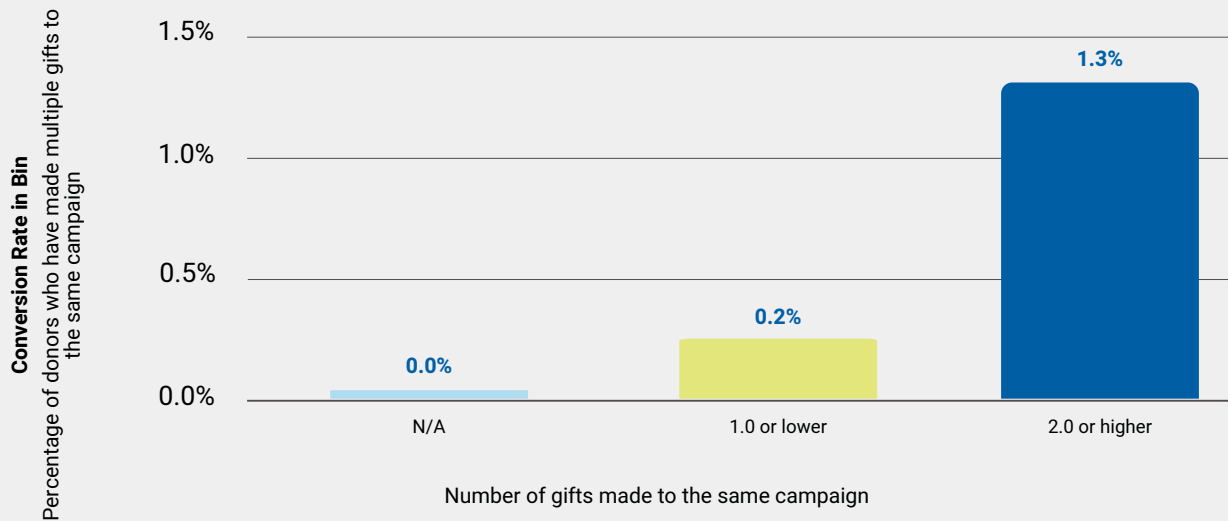


Figure 7. Advocacy Feature Data

Tracking: Maximum number of gifts made to any one social fundraising campaign over a five-year period

If you're looking for new advocates, check your giving data for people who have donated early and often to the same appeal. They may be your best bet—and even if they don't agree to advocate when requested, you should be speaking to this highly engaged group in a different way than you would speak to your new or unengaged prospects. So any extra legwork you do to identify these unique supporters is bound to ultimately yield better results.

Key Takeaway: If you're looking for new advocates, check your giving data for people who have donated early and often to the same appeal.

How fundraisers can prepare for what's next

Now is the time for advancement leaders at institutions of all sizes to begin to educate themselves about the increasingly vital and inevitable role data science will play in their school's future fundraising efforts.

Savvy fundraisers are already thinking about how they can cultivate a data-driven mindset at their organization—one that fosters curiosity about the patterns, trends, and factors that influence constituent behavior. We recommend you learn as much as you can now about how predictive modeling and machine learning works so you'll be ready to leverage new solutions and strategies as they become available.

Finally, stay up to date on fundraising technology and platforms as the space is evolving rapidly. As always, you can stay abreast of emerging trends by attending conferences and webinars and tapping your professional networks like [CASE](#) for peer insights. You might also consider partnering with data scientists and philanthropy experts who are pioneering this innovative new technology. GiveCampus is actively engaged in ongoing research with partner schools. **If you are interested in participating in one of our pilot studies, please email us at info@givecampus.com.**

Summary

Over the last 8 years, GiveCampus has grown the number of educational institutions we support to more than 1,300. That means we can now look at huge aggregate data sets and share meaningful insights that can help school fundraisers like you guide your work.

Overall, predictive modeling has the potential to significantly impact educational fundraising in both the short and long term. By leveraging data-led insights and personalization strategies, schools can enhance donor targeting, optimize campaigns, and nurture meaningful relationships with constituents, ultimately leading to increased fundraising success and sustained support for educational institutions.

GIVECAMPUS

To see GiveCampus in action,
please visit: info.givecampus.com/schedule

