

BOARDS

Research: Could Machine Learning Help Companies Select Better Board Directors?

by Isil Erel, Léa H. Stern, Chenhao Tan, and Michael S. Weisbach

APRIL 09, 2018



Innovating in Health Care

A newsletter bringing together
good management and health care.

Sign up now >

JENS MAGNUSSON/GETTY IMAGES

Ever since Adam Smith published *The Wealth of Nations* in 1776, observers have bemoaned boards of directors as being ineffective as both monitors and advisors of management. Because a CEO often effectively controls the director selection process, he will tend to choose directors who are unlikely to oppose him, and who are unlikely to provide the diverse perspectives necessary to maximize firm value. Institutional investors often are critical of CEOs' influence over boards and have made efforts to help companies improve their governance. Nonetheless, boards remain highly imperfect.

Could technology help? Advances in machine learning have led to innovations ranging from facial recognition software to self-driving cars. These techniques are rapidly changing many industries – could they also improve corporate governance?

To explore that question, we conducted a study of how machine learning might be used to select board directors, and how the selected directors might differ from those selected by management. Our intent is to demonstrate how a machine-learning model could potentially help investors by helping companies select better directors.

The first challenge with such a study is determining what makes a director “better” or “worse.” Most directors' actions occur in the privacy of the boardroom where they cannot be observed by outsiders. In addition, most of what directors do occurs within the structure of the board, so we cannot isolate their individual contributions.

Despite those complications, one clear measure of director performance is publicly available: the fraction of votes a director receives in annual shareholder re-elections. Although the CEO often influences the choice of the person *nominated* to the board and shareholders have virtually no control over the choice of directors, shareholders vote annually on their re-election. These votes reflect the support the director personally has from the shareholders and should, in theory, incorporate all publicly available information

about the director's performance. Our choice of performance measure is also motivated by the fact that the hiring decision for a corporate director is no different than any other hiring decision: it is fundamentally about predicting the individual's future performance. Since the mandate of the board is to represent shareholders' interests, shareholder votes stand out as a natural performance metric.

The second challenge we face is that we only have that measure of director performance for directors who are actually selected to join the board. Machine learning is all about prediction, but if we just try to predict how *selected* directors will fare in shareholder elections we are only looking at half of the problem. Ideally, we also want to predict how would-be directors who were not ultimately nominated would have done if they had the chance to join the board.

We address this issue by constructing a pool of potential directors for each board opening from those who, around that time, accept a directorship at a smaller nearby company. We assume that these individuals would have been attracted to a directorship at a larger, neighboring company. For the purposes of our study, we use the fraction of votes these individuals received at the company where they become director as our measure of their potential performance.

We trained a machine learning algorithm to predict directors' performance, using a dataset of large publicly traded U.S. corporations between 2000 and 2011. We used a machine learning method called gradient boosting, and then evaluated the results using a separate test dataset of directors who joined firms between 2012 to 2014 who the algorithm did not observe during this "training period." The algorithm was able to identify which directors were likely to be unpopular with shareholders. The directors that were actually hired but that our algorithm predicted would be unpopular with shareholders ended up faring much worse than other available candidates. In contrast, hired directors that our algorithm

predicted would do well indeed did better than other available candidates. (Our machine learning model performed substantially better than a standard econometric model such as ordinary least squares.)

The differences between the directors suggested by the algorithm and those actually selected by firms allow us to assess the features that are overrated in the director nomination process. We found that firms tend to choose directors who are much more likely to be male, have a large network, have a lot of board experience, currently serve on more boards, and have a finance background.

In a sense, the algorithm is telling us exactly what institutional shareholders have been saying for a long time: that directors who are not old friends of management and come from different backgrounds both do a better job in monitoring management and are often overlooked.

In light of our findings, it is worth asking: Why do real-world firms appoint directors who they could predict will be unpopular with shareholders? We think there are at least two possible reasons. First, it could be that CEOs *do not want* effective directors on their boards. Since the publication in 1932 of Adolph Berle and Gardiner Means' *The Modern Corporation and Private Property*, economists have argued that managers are able to maintain control over their firms by influencing the director selection process to ensure management-friendly boards.

Alternatively, it could be that because of behavioral biases, management *is not able* to select effective directors as well as an algorithm. In his book *Thinking, Fast and Slow*, Daniel Kahneman describes a long history of psychological research documenting that, in many circumstances, simple rules can lead to better outcomes than allowing individuals to have discretion over decisions. Machine-learning models, which are much more sophisticated

than the rules suggested by psychologists in their experiments, represent a potentially valuable way to operationalize the notion that rules rather than discretion can improve real world decision making.

How should our findings be applied in practice? The algorithms we present should be treated as “first pass” approaches; presumably more sophisticated models would predict director performance even better than the ones presented in this paper. In addition, our algorithms rely on publicly available data; if one had more detailed private data on director backgrounds, performance, etc., one could improve the algorithm’s accuracy even more. If algorithms such as these are used in practice in the future, as we suspect they will be, practitioners will undoubtedly have access to much better data than we have and should be able to predict director performance more accurately than we do in the paper.

Machine learning algorithms are not without their flaws. They are prone to bias, too, depending on the data they are fed and the outcomes they are optimizing for.

For the purpose of our study, though, it is clear that algorithms are not prone to agency conflicts and the biases that occur when boards and CEOs meet together to select new directors. Institutional investors are likely to find this attribute particularly appealing and to encourage boards to rely on an algorithm for director selections in the future. How well this approach to selecting directors will be received by management is an open question.

Isil Erel is the Fisher College of Business Distinguished Professor of Finance.

Léa H. Stern is an assistant professor of Finance and Business Economics at the Michael G. Foster School of Business at the University of Washington.

Chenhao Tan is an assistant professor at the Department of Computer Science and the Department of Information Science, by courtesy, at University of Colorado Boulder.

Michael S. Weisbach is the Ralph W. Kurtz Chair in Finance at Ohio State University, as well as a Research Associate of the National Bureau of Economic Research.

This article is about **BOARDS**

 FOLLOW THIS TOPIC


Related Topics: [TECHNOLOGY](#) | [DATA](#)

Comments

Leave a Comment

POST

0 COMMENTS

 **JOIN THE CONVERSATION**

POSTING GUIDELINES

We hope the conversations that take place on HBR.org will be energetic, constructive, and thought-provoking. To comment, readers must sign in or register. And to ensure the quality of the discussion, our moderating team will review all comments and may edit them for clarity, length, and relevance. Comments that are overly promotional, mean-spirited, or off-topic may be deleted per the moderators' judgment. All postings become the property of Harvard Business Publishing.