Selecting Directors Using Machine Learning

Finance Working Paper N° 605/2019
September 2020

Isil Erel
Ohio State University, NBER, and ECGI

Léa Stern
University of Washington

Chenhao Tan
University of Colorado, Boulder

Michael S. Weisbach
Ohio State University, NBER and ECGI

ECGI Working Paper Series in Finance

Selecting Directors Using Machine Learning

Working Paper N° 605/2019
September 2020

Isil Erel
Léa Stern
Chenhao Tan
Michael S. Weisbach

We thank Renée Adams and Reena Aggarwal (who graciously shared data), Lucian Bebchuk, Philip Bond, Lisa Cook, Ran Duchin, Daniel Ferreira (discussant), Fabrizio Ferri, Itay Goldstein, Shan Ge, Jarrad Harford, Ben Hermalin, Joan MacLeod Heminway, Joshua Lee (discussant), Nadya Malenko (discussant), Jordan Nickerson (discussant) Miriam Schwartz-Ziv, Anil Shivdasani, Tracy Yue Wang (discussant), Ayako Yasuda, Luigi Zingales (discussant), two anonymous referees, and conference and seminar participants at London School of Economics, Michigan State-UIUC Virtual Finance Seminar, North Carolina, Nanyang Tech, National University of Singapore, Northeastern, Ohio State, SMU, Tennessee, Utah, Washington, 2017 Pacific Northwest Finance Conference, 2017 WAPFIN Conference at NYU Stern, 2017 NABE TEC Conference, 2018 University of Miami-AFFECT conference, 2018 Drexel Corporate Governance Conference, 2018 ICWSM BOD workshop, 2018 NBER Economics of AI Conference, 2018 Wash. U. Olin Corporate Finance Conference, 2019 AFA Annual Meetings, 2019 NBER Big Data Conference, 2019 Conference on Emerging Technologies in Accounting and Financial Economics at USC, 2019 Wine Country Finance Conference and Machine Learning Day Conference at Georgetown. Special thanks to Ronan Le Bras of the Paul Allen Institute for AI for providing invaluable help throughout the project.

Abstract

Can algorithms assist firms in their decisions on nominating corporate directors? Directors predicted to do poorly by algorithms indeed do poorly compared to a realistic pool of candidates in out-of-sample tests. Predictably bad directors are more likely to be male, accumulate more directorships, and have larger networks than the directors the algorithm would recommend in their place. Companies with weaker governance structures are more likely to nominate them. Machine learning holds promise for understanding the process by which governance structures are chosen and has potential to help real-world firms improve their governance.

Keywords: Corporate Governance, Boards of Directors, Machine Learning

JEL Classifications: C10, C45, G30, M12, M51

Isil Erel
Professor of Finance
Ohio State University, Department of Finance
832 Fisher Hall, 2100 Neil Avenue
Columbus, OH 43210-1144, United States
e-mail: erel@fisher.osu.edu

Léa Stern
Assistant Professor of Finance and Business Economics
University of Washington, Foster School of Business
429 Paccar Hall
Seattle, WA 98195, United States
phone: +1 206 543 5365
e-mail: leastern@uw.edu

Chenhao Tan
Assistant Professor of Computer Science
University of Colorado, Department of Computer Science
ECES 118A
Boulder, CO, 80309, United States
e-mail: Chenhao.Tan@Colorado.edu

Michael S. Weisbach*
Professor of Finance
Ohio State University
700 Fisher Hall, 2100 Neil Avenue
Columbus, OH 43210-1144, United States
phone: +1 (614) 292 2418
e-mail: weisbach.2@osu.edu

*Corresponding Author

THE OHIO STATE UNIVERSITY
FISHER COLLEGE OF BUSINESS

Fisher College of Business
Working Paper Series

Charles A. Dice Center for Research in Financial Economics

Selecting Directors Using Machine Learning

Isil Erel,
The Ohio State University, NBER, and ECGI

Léa H. Stern,
University of Washington

Chenhao Tan,
University of Colorado, Boulder

Michel S. Weisbach,
The Ohio State University, NBER and ECGI

Dice Center WP 2018-05
Fisher College of Business WP 2018-03-005

September 1, 2020

Selecting Directors Using Machine Learning

ISIL EREL

LÉA STERN

CHENHAO TAN

MICHAEL WEISBACH*

September 1, 2020

Abstract

Can algorithms assist firms in their decisions on nominating corporate directors? Directors predicted to do poorly by algorithms indeed do poorly compared to a realistic pool of candidates in out-of-sample tests. Predictably bad directors are more likely to be male, accumulate more directorships, and have larger networks than the directors the algorithm would recommend in their place. Companies with weaker governance structures are more likely to nominate them. Machine learning holds promise for understanding the process by which governance structures are chosen and has potential to help real-world firms improve their governance.

(JEL C10, C45, G30, M12, M51)

---

* Stern: Foster School of Business, University of Washington (leastern@uw.edu); Erel: Fisher College of Business, Ohio State University, NBER, and ECGI (erel@fisher.osu.edu); Tan: Department of Computer Science, University of Colorado, Boulder (chenhao.tan@colorado.edu); Weisbach: Fisher College of Business, Ohio State University, NBER, and ECGI (weisbach.2@osu.edu). We thank Renée Adams and Reena Aggarwal (who graciously shared data), Lucian Bebchuk, Philip Bond, Lisa Cook, Ran Duchin, Daniel Ferreira (discussant), Fabrizio Ferri, Itay Goldstein, Shan Ge, Jarrad Harford, Ben Hermalin, Joan MacLeod Heminway, Joshua Lee (discussant), Nadya Malenko (discussant), Jordan Nickerson (discussant) Miriam Schwartz-Ziv, Anil Shivdasani, Tracy Yue Wang (discussant), Ayako Yasuda, Luigi Zingales (discussant), two anonymous referees, and conference and seminar participants at London School of Economics, Michigan State-UIUC Virtual Finance Seminar, North Carolina, Nanyang Tech, National University of Singapore, Northeastern, Ohio State, SMU, Tennessee, Utah, Washington, 2017 Pacific Northwest Finance Conference, 2017 WAPFIN Conference at NYU Stern, 2017 NABE TEC Conference, 2018 University of Miami-AFFECT conference, 2018 Drexel Corporate Governance Conference, 2018 ICWSM BOD workshop, 2018 NBER Economics of AI Conference, 2018 Wash. U. Olin Corporate Finance Conference, 2019 AFA Annual Meetings, 2019 NBER Big Data Conference, 2019 Conference on Emerging Technologies in Accounting and Financial Economics at USC, 2019 Wine Country Finance Conference and Machine Learning Day Conference at Georgetown. Special thanks to Ronan Le Bras of the Paul Allen Institute for AI for providing invaluable help throughout the project.

1. Introduction

A company's board of directors is legally responsible for managing the company. In principle, the board of directors reports to the shareholders and represents their interests. In practice, however, there is much variation in director quality and the extent to which they serve shareholders' interests.[1]

Many of the concerns about boards come from the director selection process, which has been a source of debate since at least Berle and Means (1932).[2] The selection process for directors is one of the most important yet least studied questions in corporate governance. Despite the checks and balances built into a public corporation's governance system, the CEO often controls the selection of new directors.[3] In practice, appointed directors are almost always supporters of the CEO and his policies. Aside from occasional proxy contests, shareholders have virtually no control over the choice of the directors whose mandate is to represent their interests.

In this paper, we consider a potential alternative approach to selecting directors: one that uses algorithms that rely on data on firms, current board members, and the attributes of potential directors, to identify the quality of directors being considered for a given firm's board. Because boards must make predictions for how potential nominees would perform, the selection of directors is essentially a prediction problem. And while "traditional" econometrics is typically designed for estimating structural parameters and drawing causal inferences, machine learning algorithms are substantially better at making predictions. They are designed to maximize out of sample predictive accuracy by avoiding overfitting and by not being constrained by specific parametric assumptions or restrained in the number of covariates.[4] Because there is no one size fits all "good" governance, many covariates potentially matter and interact in non-linear ways

[1] See Hermalin and Weisbach (2003), Adams, Hermalin and Weisbach (2010), and Adams (2017) for surveys.
[2] Berle and Means (1932) wrote: "Control will tend to be in the hands of those who select the proxy committee and by whom the election of directors for the ensuing period will be made. Since the proxy committee is appointed by the existing management, the latter can virtually dictate their own successors" (p. 87). Hermalin and Weisbach (1998) present a formal model of this process in which boards vary in their independence from the CEO in equilibrium.
[3] See Shivdasani and Yermack (1999) and Kramarz and Thesmar (2013) for anecdotal evidence suggesting that the CEO typically holds a veto power over the choice of directors. See also Cai, Nguyen, and Walkling (2017), who document that more complex firms and firms in more competitive environments are more likely to appoint directors who are connected to the CEO or the existing board.
[4] See Athey and Imbens (2017) and Mullainathan and Spiess (2017).

when predicting director performance. This is what motivates the use of machine learning algorithms in this paper. Algorithms identify complex patterns and structure in the data by sifting through a large set of covariates and combining them in nonlinear and interactive ways to identify functional forms that reliably predict outcomes. We then use the predictions to provide insights into the decision-making process that governs the selection of corporate directors.

We construct a large database of publicly traded U.S. firms and independent directors appointed between 2000 and 2014. We build several machine learning algorithms designed to predict director performance using director, board, and firm level data available to the nominating committee at the time of the nominating decision. We compare the algorithms' selections of directors to the ones actually chosen by firms. The discrepancies between firms' actual choices of directors and the choices based on the predictions from our algorithms allow us to characterize which individual features are overrated by decision makers. In addition, by characterizing firms that tend to nominate directors with predictably poor performance, our analysis speaks to the role of governance structures in the selection of directors. As such, the algorithms' predictions can provide insights into the decision-making process that governs the selection of corporate directors.

A crucial element of any algorithm designed to select valuable independent directors is a process for assessing a director's performance in a particular firm. The task of measuring the performance of an individual director is challenging. Directors generally act collectively on the board and it is usually impossible for a researcher to ascertain the actions of any director. Hart and Zingales (2017) emphasize that directors' fiduciary duty is to represent the interests of the firm's shareholders. Their popularity among shareholders is thus a natural metric for evaluating them. For that reason, our main measure of director performance is based on levels of shareholder support in annual director re-elections. This is an individualized, market-based measure of performance which captures some aspects of investors' preferences. We task the algorithm, with predicting the excess support relative to the slate, as well as a dummy variable that captures low absolute shareholder support – i.e., strong dissent against a director.

We employ several machine learning algorithms to predict the performance of any potential director at any particular company, taking into consideration who is currently sitting on the board. Using our sample of public firms, we train each algorithm on a training set (directors appointed between 2000 and 2011), and then compare the predictions to the observed out-of-sample data using a test set (directors appointed between 2012 and 2014).

We find that these algorithms make accurate out-of-sample predictions of the distribution of outcomes, whether predicting the level of shareholder support, the excess support relative to the slate, strong dissent against a director by shareholders, or director turnover. The directors the algorithms predicted would do poorly did worse on average than the directors the algorithms predicted would do well. In comparison, the directors predicted to do poorly by an Akaike Information Criterion (AIC)-selected OLS model do not actually have worse performance out of sample than those it predicted would do well. Machine learning algorithms, by letting the data speak about the underlying relationships among a large set of candidate predictors, end up fitting the data much better and consequently do better at predicting future outcomes out of sample.

We ensure that the out of sample predictive accuracy of the algorithm is not dependent on one particular measure of director performance by considering multiple measures of director performance. For example, we use strong dissent against a director as an alternative performance measure. Also, a director who leaves soon after being appointed is likely to have been a poor choice (Ferreira et al., 2017, Bates et al., 2016). Therefore, an additional director performance measure we use in our predictions is whether the newly appointed director leaves the board within two years of appointment. We show that our algorithms do well in predicting this director-turnover measure as well. Finally, we consider the model's ability to predict firm profitability and announcement returns of director appointments. We find that the algorithm's predictions of shareholder votes to re-elect directors are also strongly related to firm profitability and announcement returns around director appointments.

We only observe director performance for directors who were actually nominated to the board but do not observe them for potential candidates who were not nominated. This *"selective labels"* problem of observing director performance only for directors who were actually selected is a common issue in prediction problems (see Kleinberg et al., 2017). Consequently, we are only able to evaluate the algorithm's predictive ability for nominated directors. If decision makers consider features that are not observable to our algorithm in their nominating decisions of directors, the distribution of outcomes in the set with observed labels (nominated directors) could differ from that in the set with missing labels (not nominated directors), even if they share exactly the same observable characteristics. In other words, if boards are skilled at using unobservables in their nominating decisions, nominated directors could have higher expected performance than otherwise similar (based on observables) passed-over directors.

To address this selective-labels problem, for each board appointment in our test set, we construct a realistic pool of potential candidates: directors who joined the board of a smaller neighboring company within a year. Presumably these potential candidates would have found the opportunity to be on the board of a larger nearby company attractive, since directorships at larger companies tend to be better paying and more prestigious than directorships at smaller companies. They also signaled that they were available and willing to travel to this specific location for board meetings. Although we do not observe the performance (*i.e.* the *label*) of those potential candidates at the focal firm (this is the essence of the selective labels problem), the design of our candidate pools allows us to observe what we refer to as their *"quasi-label"*: their performance on the board they effectively joined. Importantly, quasi-labels do not need to be *perfect* substitutes to labels in our procedure to assess the algorithms' predictions when labels are missing and the decision maker relies on unobservables.

We find that directors the algorithm predicted would perform poorly (well) do perform poorly (well) when compared to potential available alternatives. Directors in the bottom decile of predicted performance rank at the 23rd percentile in the distribution of quasi-labels. In contrast, those in the top decile rank at the 80th percentile. We find that OLS models are unable to predict ex ante who will perform well compared to alternatives and who will not.

One of the differences between machine learning algorithms and traditional econometric modeling is that machine learning algorithms do not provide an easy formula that can be used to infer the influence of any particular independent variable on performance. However, in recent years, *explainable AI*, has become an important and growing strand in the machine learning literature, focusing on improving model interpretability (e.g. Lundberg and Lee, 2016 and Ribeiro, Singh and Guestrin, 2016). We also employ these state-of-the-art techniques to gain insights into our machine learning algorithm and quantify the contribution of each feature to predicting director performance. In addition, while machine learning models do not generate estimates of the underlying structural parameters of a model, we can use the algorithm's predictions to understand the features that are overvalued and undervalued by firms in the director selection process. Relative to algorithm-selected directors, management-selected directors who receive predictably low shareholder approval are more likely to be male, have larger networks, and sit on more boards. These attributes characterize the stereotypical director in most large companies. A plausible interpretation of our results is that firms that nominate predictably unpopular directors tend to be subject to homophily when choosing directors, while the algorithm suggests that adding diversity would be a better idea.

Finally, to help understand the process determining the nomination of directors, we attempt to distinguish between potential explanations (error or agency) for why firms regularly make poor decisions when they hire directors. We find that firms that nominate directors who were predictably poor choices have worse governance structures. This evidence is consistent with agency conflicts distorting nominating decisions and, therefore, with what we refer to as the "cocooning" view of director selection.

Several papers in the recent economics and finance literature have used machine learning techniques. In a seminal paper, Kleinberg et al. (2017) study judges' bail decisions and show that machine predictions could significantly reduce crime. The authors establish a useful road map for properly utilizing machine learning within an economic framework. Machine learning is quickly being adopted as a new methodology in the asset pricing literature (e.g. Rossi, 2018, Ke et al., 2019, Abis, 2018) and Bubb et al., 2018) and in microstructure (Easley, Lopez de Prado, O'Hara and Zhang, 2020) Corporate-finance applications are developing. For example, Li et al. (2020) measure corporate culture using word embedding on earnings call transcripts.

Our paper uses machine learning to add to the literature on corporate governance. We contribute to the literature on the selection of directors (Smith, 1776, Berle and Means, 1932, Hermalin and Weisbach, 1998, Kramarz and Thesmar, 2013, Coles, Daniel and Naveen, 2014, Cai, Nguyen and Walkling, 2017) by showing that the quality of director hiring decisions is related to firms' governance structure. Our evidence suggests that firms that select predictably bad directors adhere to a "cocooning" view of director selection. We also contribute to the literature on director elections (Cai, Garner and Walkling, 2009, Fischer et al. 2009, Iliev, Lins, Miller, and Roth, 2015, Fos, Li and Tsoutsoura 2017, Aggarwal, Dahiya and Prabhala, 2017, Ertimur, Ferri and Oesch, 2017) by confirming that the information contained in shareholder votes in director elections is relevant and predicts firm profitability.

This paper is the first, to our knowledge, to apply supervised machine learning to improve our understanding of corporate governance. We also introduce a new empirical strategy, based on quasi-labels, as a way to address the selective labels problem that is pervasive in social science prediction problems. Machine learning tools have the potential to help answer many unanswered questions in the social sciences, both by academics wishing to understand the way the world actually works, and by practitioners and policy makers wishing to make better real-world decisions. In terms of boards of directors, an algorithmic decision aid could allow firms to choose better among existing candidates, without stripping decision makers of their judgement. We emphasize strongly that algorithms hold promise to *complement* rather than substitute human judgement. As such, we expect the economic value of board decisions to *increase* with the use of algorithmic decision aids (Autor, 2015 and Agrawal, Gans, and Goldfarb, 2017). In addition, algorithmic decision aids could help firms identify alternative choices of potential directors, thereby opening up board seats to a broader set of candidates with more diverse backgrounds and experiences, who would have otherwise been overlooked.[5]

2. Constructing a Sample on which Algorithms Can Select Directors

2.1. Measuring Director Quality

[5] We thank Oren Etzioni of the Paul Allen Institute for AI for pointing out this benefit of our approach.

2.1.1 Measuring Director Quality through Re-Election Results

An essential part of designing the algorithm is specifying a measure of director performance as the basis for which directors are selected. Our analysis focuses on the relative shareholder support that directors receive in annual director re-elections as a market-based measure of *individual* directors' performance. Our main outcome variable is *excess votes*: the average level of shareholder support over the first three years of director tenure, adjusted by the average support for the entire slate of directors up for re-election on that board that year.[6] Our results are qualitatively unchanged if we task the algorithms with predicting the absolute level of shareholder support, *i.e.* if we do not subtract the average for the slate.

One potential concern with using shareholder support as our measure of director performance is that in the vast majority of cases, all directors receive an overwhelming majority of the votes, with mean shareholder support usually around 95%.[7] There is almost no variation in the *outcome* of the re-elections. Nonetheless, variation among winning votes does appear to reflect differences in directors' quality. Cai et al. (2009), Fischer et al. (2009), and Iliev et al. (2015) find that vote totals predict stock price reactions to subsequent turnover. In addition, vote totals are negatively related to CEO turnover, board turnover, management compensation levels, as well as the probabilities of removing poison pills and classified boards.

Moreover, director re-elections appear to impact a firm's real activities, even if the elections are not contested and the nominated directors end up being re-elected. Fos et al. (2018) find that when directors are closer to getting re-elected, they are more likely to fire CEOs, presumably to persuade shareholders that they are being more diligent. Aggarwal et al. (2017) suggest that directors with low relative support are

[6] The distribution of shareholder support does not change over the first few years of a director's tenure. We obtain similar results using shareholder support at year one, year two or year three instead of using the average over the first three years.

[7] The literature on director re-elections is large, including Boone, Field, and Karpoff (2007), Linck, Netter, and Yang (2008), Cai, Garner and Walkling (2009), Linck, Netter, and Yang (2009), Fischer et al. (2009), Coles, Daniel and Naveen (2014), Iliev, Lins, Miller, and Roth (2015), Aggarwal, Dahiya and Prabhala (2017), Ertimur, Ferri and Oesch (2017), Cai, Nguyen and Walkling (2017), Fedaseyeu, Linck, and Wagner (2017), Fos, Li and Tsoutsoura (2018).

more likely to leave the board, and if they stay, tend to move to less prominent positions. And Ertimur et al. (2018) find that when votes are withheld from directors, boards explicitly attempt to address shareholders' concerns.

Shareholder support could reflect arbitrary recommendations by proxy advisors such as ISS. Ertimur et al. (2018) report that since 2003, large institutional investors take an active role in developing the guidelines that are the basis of ISS recommendations, which, as such, reflect its clients' aggregated preferences. Aggarwal, Erel and Starks (2016) confirm this result, documenting that institutional investors and proxy advisors pay attention to the changing opinions of their beneficiaries and shareholders. However, institutional investors do *not* follow proxy advisors' recommendations blindly. Aggarwal et al. (2016) find that shareholders are less likely to follow the recommendations of either management or proxy advisory firms as shareholders are forming their own views due to changes in public opinion. Iliev and Lowry (2014) show that institutional investors with larger size of ownership tend to vote more independently from ISS recommendations.

We therefore repeat our tests by focusing on a subsample of firms with larger-than-median (26%) ownership by the top-5 institutional owners and our results are unchanged. Using detailed voting data from 2003-2017, Heath et al. (2019) show that when ISS recommends voting against management, index (active) funds vote with management 54% (42%) of the time. This recent stream of the literature strongly suggests that shareholder votes are not simply the reflection of an arbitrary recommendation issued by proxy advisors.

Overall, the literature finds that the level of shareholder support does reflect perceptions of director quality, that directors do care about these perceptions, and take actions to influence them. We test whether algorithms can pick up variations in these perceptions of director quality despite the fact that most directors receive extremely high support.

In addition to using shareholder votes as a measure of director performance, we also train the algorithm to predict dissent, which is an indicator variable equal to one if a director receives low (less than 90%) support.

2.1.2 Measuring Director Quality through Turnover

Director turnover is a measure of director-firm match quality used in the literature (see e.g., Ferreira, Ginglinger, Laguna, and Skalli, 2017). Therefore, we use whether a new director leaves within two years of his or her appointment as an alternative measure of director quality.

2.2. Sample Selection

To evaluate the performance of an algorithm to select directors, we must gather a sample in which we can observe the attributes of firms and boards, and also measure the performance of directors. Because of these requirements, we focus on a sample of boards from large, publicly-traded, U.S. firms with an average market capitalization of $6.6 billion. We identify 41,015 new independent directors appointed to 4,887 unique corporate boards of these firms between 2000 and 2014 using *BoardEx*, which is our main data source for director and board-level characteristics.

We obtain data on the level of shareholder support for individual directors from *ISS Voting Analytics* and focus on directors appointed during our sample period. Because of the possibility of factors that lead all directors in one firm to receive higher average votes from directors in other firms, we rely on most specifications on a measure of *excess votes.* To construct this measure, we start with the number of votes in favor over all votes cast (yes, no, withheld). We then subtract the average for the slate of directors up for re-election on that board and take the average of this variable over the first three years of tenure.[8] Our sample contains the voting outcome, i.e. *excess votes*, for 24,054 new director appointments.[9]

2.3. Summary Statistics

Table 1 presents summary statistics for the four measures of director performance: average shareholder support (mean total votes), excess votes, dissent, and turnover. As previously documented in the literature on uncontested director elections, the overall level of shareholder support is typically very high. Given that

---

[8] Some firms have “staggered” boards that are elected for three-year terms. For these firms, instead of averaging the support over the first three elections of a director’s tenure, we use the support in the one election covering the three-year period.

[9] All the results reported below are similar when we use shareholder support not adjusting for the average support of the other directors at the firm.

the mean level of support is .948 and the median is .975 (with a standard deviation of .07), a voting outcome below 90% is a relatively poor outcome for a director. Although shareholder support in uncontested elections is typically very high, shareholders do on occasion oppose newly nominated directors (see figure in online appendix). We also present the means for strong dissent against a director – i.e., a dummy variable that takes on a value of one if the director gets less than 90% of shareholder support – as well as a dummy for director turnover within two years of appointment. Over the sample period, dissent is 12% while director turnover is about 11%, on average.

Table 2 illustrates that the frequency of shareholder discontent varies by director and board characteristics. For example, the fraction "poor outcomes", representing the bottom 10% of the sample in terms of excess votes, is 10.6% for male directors and 7.9% for female directors.[10] Similarly, busy directors (serving on three or more boards) experience low shareholder support more frequently than non-busy directors.

However, theory provides little guidance regarding the particular variables and functional forms of the relation between the various director, board and firm characteristics and the performance of directors. For example, we do not know whether we should expect female busy directors with a Ph.D. serving on the large board of a small firm in the pharmaceutical industry to receive higher or lower shareholder support on average than a male director who serves on a single small board of a large manufacturing corporation. The problem increases in complexity when many more covariates are likely to matter. For this reason, we wish to utilize an estimation procedure that does not impose the specific form for the relationship between potential explanatory variables. This logic is the reason why machine learning algorithms are successful at prediction. They are designed for the many problems, such as predicting which directors will be successful in a given firm, for which theory is silent on the appropriate functional form between the explanatory variables and the outcome to predict (Athey, 2017).

[10] The pattern is similar if we define a poor outcome as having shareholder support is below 80%.

3. Evaluating Machine Learning Predictions of Director Performance

3.1. Model Specification

We employ machine learning algorithms that predict the performance of potential directors. The algorithms use a set of observable director, board, and firm features that are available to the nominating committee at the time of the nominating decision. The algorithms are commonly used in the supervised machine learning literature: *lasso*, *ridge*, *neural networks* and *gradient boosting trees (XGBoost)*.[11] We first train each algorithm on the 2000-2011 portion of our sample containing 18,476 new independent director appointments, of which 12,815 are unique directors, at 2,407 firms. Training involves having the algorithm determine which combinations of variables best predict future performance.[12] We evaluate the models' out-of-sample predictions on the held out 2012-2014 portion of our sample containing 5,578 new director appointments, of which 4,019 are unique directors, at 569 firms. We compare those out-of-sample predictions to those from an AIC-selected OLS model. All comparisons are based on predictions for the 2012-2014 subsample of director appointments, which does not overlap with the 2000-2011 subsample on which the algorithms are trained.

The optimal way to choose the size of the training and test sets depends on the signal-to noise ratio in the data and the training sample size. Therefore, it is difficult to establish a general rule on how much training data is enough. For very large datasets, a 90-10% split can be done, though 70-30% or 80-20% splits are typically used in practice. We use an 80-20% split but our results do not depend on these alternative ways of splitting the dataset into training and test set. Since the data is inherently temporal, we only consider data splits where there is no future information leaked to the past. In addition to using 2011 as the cutoff, we also experimented with 2010, the results are the same. In the Internet Appendix, we report our main results using 2011-2014 as the test sample.[13]

[11] *XGBoost* provides an efficient implementation of the *Gradient Boosting Trees* algorithm. It is known for generating excellent predictions on a variety of problems, and was the most often used algorithm among the winning solutions in the 2015 machine learning Kaggle competition. See the Internet Appendix IA1 for a description of the way each algorithm is structured and for technical references.
[12] The algorithms rely on a regularizer that balances out in-sample fit and out-of-sample overfitting.
[13] See Hastie et al. (2009) for a discussion of methodological issues involved in choosing training and testing sets.

3.2. Predictions of Director Performance

A way to evaluate the quality of a model predicting performance is to compare whether actual performance increases with predicted performance. Table 3 summarizes the ability of the machine learning models, once trained on the earlier portion of the sample, to predict director success in the later part. Table 3 indicates that average observed shareholder support increases across model-predicted performance percentiles for each machine learning model. In contrast, in the OLS model, there is no relation between predicted and actual director performance.

Among the machine learning algorithms, *XGBoost* performs best at predicting the subsequent success of directors using both excess votes and total votes as measures of director performance. Directors predicted to be in the bottom percentile as predicted by *XGBoost* have an average *observed* excess shareholder support of -3.1%, whereas the average observed excess support is 1.2% for directors in the top percentile of predicted performance.

Figure 1 presents the average observed level of shareholder support for directors across the ten deciles of predicted performance for OLS and for the machine learning algorithms in the 2012-14 test period. The figure documents that the mean shareholder support for a director is an increasing function of the predicted one for all the machine learning algorithms, but not for the OLS model. The difference in the predictive ability of various models illustrates the difference between standard econometric approaches and machine learning. OLS fits the data well in sample but poorly out of sample. In contrast, machine learning algorithms are specifically designed to predict well out of sample.

The inability of the OLS model to predict director performance could potentially occur because the particular model we picked was not well specified. For this reason, in Appendix Table IA1, we present various OLS specifications that include director, firm, and board-level variables that have been used in the prior literature. We also include industry and time fixed effects in various specifications. It is important for us to use only the *ex-ante* variables that would be available to the nominating committee when they pick new directors. To compare the out-of-sample predictable power of these specifications, we calculate the Akaike Information Criterion (AIC) for each model and report it in the last row. The AIC provides each specification's out-of-sample prediction error and allows us to compare the relative quality of OLS models presented. The OLS model used in Table 3 corresponds to Model (4) of the Appendix Table IA1, which gives the lowest AIC error.

The fact that machine learning models do better than OLS at predicting director performance out of sample is consistent with the arguments of Athey and Imbens (2017) and Mullainathan and Spiess (2017), who emphasize that machine learning algorithms are superior for prediction problems such as this one. Machine learning models can handle a large set of covariates without imposing strong parametric assumptions avoid overfitting. By letting the data speak about the underlying relationships among the variables, the algorithms end up fitting the data better and consequently do better at predicting outcomes out of sample.

3.3. Excluding Poorly Performing Firms

A possible concern with this analysis is that the relation between predicted performance and subsequent performance could occur only because of poorly performing firms. A poorly performing firm would likely be less attractive to a director, so it could be that only low ability directors are attracted to poorly performing firms, even if the firms are relatively large and otherwise prestigious. Because of their low ability, these directors would tend to perform worse. It is also possible that shareholders express their discontent only when firms perform poorly.

For this reason, we repeat our analyses omitting firms that experience negative abnormal returns in the year prior to the nomination. We find similar results to those reported above when we exclude poorly performing firms from the sample. Therefore, it does not appear that the relation between subsequent performance and predicted performance compared to alternative potential directors is driven by poorly performing firms with disgruntled shareholders.

4. Comparing Directors who were Appointed to Potential Alternative Choices

4.1 Designing the Quasi-Labels Procedure to Evaluate the Algorithm's Performance

The results so far suggest that directors identified by our algorithm as likely to have low (high) future shareholder support, are in fact on average more likely to have low (high) support in subsequent elections. The algorithm predicts the distribution of outcomes. Accurate out of sample predictions, however, are not sufficient to imply that algorithms could assist firms in their nominating decisions of corporate directors. Specifically, there are two important and related challenges in assessing whether the algorithmic predictions can actually lead to better outcomes. First, we can only observe how well our algorithm's predictions do for directors who are actually appointed to that position. This issue is known as the *selective labels* problem. Second, when deciding on their choice of directors, decision makers presumably take factors into account that are not observable to the algorithm.

Therefore, it is possible that the directors who were nominated could differ from alternative possible choices in terms of unobservables even if they are observationally equivalent from an outsider's perspective. In particular, new directors could have been chosen because they have a set of skills that are valuable to the firm, or because they have a personal relationship with the CEO or existing directors. A firm could also have decided not to nominate a candidate based on some characteristics unobservable to the algorithm that would make this candidate a poor choice. We cannot observe these factors, yet they could lead to different average outcomes for nominated vs. not nominated individuals, even if both are identical on the basis of observable characteristics.

To formalize these concepts, we develop a framework in the spirit of Kleinberg et al. (2017) and present it in the online appendix. Our empirical strategy to address these concerns involves the design of a pool of realistic potential candidates for each vacant board position. Using candidate pools, we can evaluate the algorithm's predictions for directors who firms actually nominated (focal directors). In cases where our algorithm predicted low performance for directors, we are interested in whether there were plausible alternatives available and how these alternative candidates would have performed. We compare how the focal director performed relative to potential alternatives.

Each new board appointment in the test set is associated with a candidate pool, comprised of directors who, within one year of the appointment, joined the board of a smaller neighboring firm.[14] By revealed preference, we know that these directors were available to join a board at that time and were willing to travel to that specific location for board meetings. We restrict the pool of potential candidates to directors who joined a *smaller* neighboring firm since the prestige and remuneration of being a director tends to increase with company size (see Masulis and Mobbs, 2014). There are on average 147 candidates in a candidate pool. Importantly, our results are similar regardless of the way in which we construct pools of candidate directors. For example, our results are robust to further restricting the set of candidates in candidate pools to directors who joined the board of a firm in the same industry or by relaxing the smaller firm constraint (see Appendix Figure IA7).

To generate predictions of the quality of potential director candidates at the focal firm, the algorithms use the board and firm characteristics as well as the committee assignments of the appointment at the focal firm with the individual potential candidate's features. We do not observe the performance of these potential candidates at the focal firm (the selective labels problem). We do, however, observe the director's performance on the board he or she joined. This performance is an informative signal that serves as a substitute for a direct measure of performance, which we refer to as a potential director's *"quasi-label"*. This quasi-label provides an indication of how a potential candidate would have performed on the focal board.

We design a procedure to assess the algorithm's performance based on quasi-labels, and provide a schematic representation in Figure 2. We first rank all directors in our test set according to their algorithm-predicted performance. For all directors in the bottom decile of predicted performance, we consider their associated candidate pool and rank candidates according to their predicted performance *on the focal board*. We consider promising candidates the algorithm would have recommended (top decile) and rank these promising candidates according to their quasi-labels. The question we consider is: how does the observed

[14] A neighboring firm is defined as a firm whose headquarters is within 100 miles of the focal firm's headquarters. The average distance with the focal board is 35 miles (median distance is 26 miles).

performance of hired directors predicted to do poorly compare to the performance of promising available alternative candidates?

If the observed performance of the nominated director ranks high in the distribution of quasi-labels, then she ended up doing well relative to available alternatives even if our algorithm had predicted this particular director would do poorly. The focal board could have relied on unobservables in the nomination process, and the high rank in the distribution of quasi-labels would suggest that unobservables were used as *signal*. On the other hand, if the observed performance ranks low in the distribution of quasi-labels, then our algorithm would have identified *ex ante* that this director would perform poorly, and relative to alternatives, she indeed did perform poorly. This pattern would suggest that any unobservables used in the nomination decision process were not a signal of performance, but either noise, bias, or related to agency problems.

4.2 Quasi-Labels Results

Table 4 presents the median rank in the distribution of quasi-labels for directors in the bottom and top deciles of predicted performance for several machine learning algorithms, as well as for the OLS model. For all machine learning models, nominated directors predicted to do well performed noticeably better than available alternative candidates, while nominated directors predicted to do poorly performed worse than available alternative candidates. *XGBoost* and *Lasso* appear to be the preferred algorithms as they can best discriminate the directors who will do well from those who will not. In the rest of the paper, we focus on results with *XGBoost* to simplify the discussion. The results using *Lasso* are very similar to those with *XGBoost*, and are presented in the Internet Appendix Figure IA6.

The median director predicted by the *XGBoost* algorithm to be in the *bottom* decile of performance ranks at the 23rd percentile in the distribution of quasi-labels. The median director predicted to be in the *top* decile ranks at the 80th percentile in the distribution of quasi-labels. In contrast, the predictions from the OLS model are uninformative about subsequent performance relative to available alternative candidates.

Figure 3 shows that the mean and median rank (percentile) in the distribution of quasi-labels increases across deciles of predicted performance. The observed performance of hired directors in the test set is compared to the performance of *all* potential candidates in their respective candidate pool. These results suggest that machine learning models can be helpful in predicting whether an individual will be successful as a director in a particular firm.

4.3 Interpreting the Quasi-Labels Results

An underlying assumption of the quasi-labels procedure is that the *distribution* of quasi-labels for a candidate pool – i.e, candidates' performance on the boards they actually join – mirrors the distribution of performance these available (but passed-over) candidates would have had, had they served on the focal board. Matches between boards and directors are likely driven by unobservables and directors are chosen with the intent of maximizing the fit between directors and firms. If unobservable characteristics influence firms' decisions *not* to select some directors, the endogenous nature of the board-director match could potentially lead to systematically inflated quasi-labels. By revealed preference, the performance of available candidates would be lower on the focal board than on the board they effectively joined. The quasi-label could be greater than the missing label (the unobserved level of support a potential director would have received on the focal board) if quasi-labels are inflated due to the endogenous board-director match.

In other words, quasi-labels are not perfect substitutes for missing labels. The assumption we are making in their use is that the difference between the quasi-label and the missing label is not systematically negatively correlated with predicted performance. As long as this assumption holds, the *distribution* of quasi-labels is useful: it can be used to approximate the distribution of labels we would have observed for passed-over directors at the focal firm.[15] The selected director's rank in the distribution of quasi-labels indicates her performance relative to how other available candidates would have performed. The quasi-labels procedure allows us to evaluate the algorithm's predictive ability in the presence of selective labels.

[15] Suppose the error was systematically decreasing as predicted performance increases. For example, suppose it was positive for focal directors predicted to do poorly, and negative for those predicted to do well. Quasi-labels would systematically overestimate the performance a potential candidate would receive on the focal board for focal directors predicted to do poorly but would systematically underestimate it for focal directors predicted to do well. In this example, a finding that directors in the bottom (top) decile of predicted performance rank low (high) relative to alternative candidates could potentially be partly driven by this positive (negative) difference between the true labels and the quasi-labels.

There potentially exist many settings in which quasi-labels could be used to assess algorithms' predictions if they represent a plausible substitute for missing labels. Under the assumption that the difference between the unobserved missing label and its quasi-label does not vary systematically across the distribution of predicted outcomes, the quasi-label procedure can offer a useful approach in various contexts.[16]

5. Alternative Measures of Performance?

An important concern is the extent to which, being trained to predict *excess votes*, the algorithm is actually predicting a director's performance, or merely predicting which directors will be popular with shareholders. Hart and Zingales' (2017) argument implies that a director's performance is definitionally equal to her popularity with shareholders. However, others would disagree, claiming that a director's performance is her impact on a firm's profitability regardless of what shareholders say about her. One reason why this issue is a concern is that many institutional shareholders decide on their votes through recommendations of shareholder services companies such as ISS. ISS introduced guidelines in the latter part of our training period. For example, explicit guidelines to support proposals aimed at increasing female board representation were introduced in 2010.[17] Moreover, as discussed before, we find similar results when we focus on a subsample of firms with larger-than-median ownership by the top-5 institutional owners (Iliev and Lowry, 2014).

5.1. Predicting Abnormal Returns on Director Announcements

The recent literature on routine director re-elections does find that votes capture the performance of directors. We confirm that this pattern occurs in our data as well. We compare the cumulative abnormal returns (CARs) around the announcement of director appointments in our test set for directors predicted to do well to those for directors predicted to do poorly.[18]

[16] Algorithmic predictions of loan performance could be an example. Quasi-labels for denied loans could be the loan performance for the firm (or individual) offered a similar loan by a different institution.
[17] Our training sample covers data from 2000-2011. Less than 20% of appointments in our training set take place when ISS had those specific guidelines in place.
[18] We collect announcement dates from BoardEx, CapitalIQ and Lexis-Nexis.

Table 5 reports the mean CARs using a (-1; +1) window around announcements. The same pattern emerges using longer windows as well. Using *XGBoost* to predict excess votes, we find that the mean CAR for directors predicted to do poorly (decile 1) in our test set is -1.94% whereas it is +0.75% for directors predicted to do well (decile 10). The difference is statistically significant at the 1% level. Directors predicted to be unpopular also tend to be viewed by the market as worse directors. We also used the algorithm to predict announcement CARs using a smaller sample for which announcement dates are available, with similar results.

5.2. Predicting Future Profitability

Next, we train an *XGBoost* algorithm to predict a measure of firm profitability, *EBITDA/Total Assets*, three years post appointment. We then sort directors in our test set into deciles based on predicted profitability. We report the actual profitability as well as the shareholder support in the first two rows of Table 6.[19] The model trained to predict profitability in the subsequent period does predict future profitability well. The actual profits for the firms sorted into deciles based on expected profits increase monotonically, with average profits increasing with the model's expectation of profitability.

Firms that nominated directors in the bottom decile of predicted performance have an average profitability of -49.8% while those in the top decile have an average of 20.5%. What is perhaps more surprising is that even though the model is trained to predict profitability, it can also predict future shareholder support. Directors predicted to be in the bottom decile of profitability have shareholder support of 94% three years subsequent to the model's training, and directors predicted to be in the top decile have shareholder support of 96%. The difference between the two is statistically significantly different from zero at the 1% level. The model trained on profitability also does reasonably well at predicting *excess votes*. The average *excess votes* is -.0004 for directors in the bottom decile of predicted profitability and it is .004 for those in the top decile (the p-value of the difference is 6.68%).

[19] The correlation of EBITDA/Total Assets with the shareholder support measure is 0.12 (p-value: 0.000).

When the model is trained using profitability, the pattern of predictions is similar. Directors joining firms the algorithm predicts will have lower (higher) profitability three years after the appointment receive lower (higher) shareholder support over their first three years of tenure. In addition, for the algorithm trained on shareholder support that we discussed above, we consider whether it can also predict future profitability in addition to future shareholder support. We break the sample into deciles based on the algorithm's predictions of *excess votes* and *total votes*, and present average observed *excess votes*, *total votes*, as well as the average profitability for each decile. We present these averages in the bottom four rows of Table 6.

As discussed above, *XGBoost* is successful in predicting future shareholder support (i.e. total votes) and excess votes: average shareholder support in the lowest decile is 92% (-1.3% for excess votes), compared to 97.7% in the top decile (1.1% for excess votes). In addition, it also predicts future profitability. Firms that nominated directors in the bottom decile of predicted shareholder support have an average profitability of -0.3%, whereas firms that nominated directors in the top decile of predicted shareholder support have an average profitability of 10%. When *XGBoost* predicts excess votes, the average profitability of firms in the bottom decile is -0.6% and it is 11.1% for the top decile (Figure 4). This finding has two important and related implications: first, shareholder votes do appear to be closely related to firm performance, thereby supporting their use as a metric to evaluate director performance and second, nominating directors on the recommendation of an algorithm trained to predict shareholder votes would not come at the expense of poor firm performance.[20]

5.3. Predicting Strong Dissent against a New Director

It is possible that only when dissent is particularly strong is it a useful signal of director performance. With this notion in mind, we create a dummy variable that equals one if there are more than 10% dissenting votes in their reelection (within the next three years) from shareholders, and use it as an alternative measure

[20] This result alleviates concerns related to the omitted payoff bias articulated in Kleinberg et al. (2017), which in our setting refers to the concern that the decision-maker could have alternative objectives other than satisfying shareholders when making the nominating decision.

of poor director quality. Since the mean shareholder support in our sample is 95% for a new director in his/her reelection, we consider a support level of less than least 90% to be a particularly a bad outcome.

In Figure 5, we repeat our exercise of predicting director performance using the training set of 2000-2011 and comparing it with the actual performance in the test set of 2012-2014, this time using whether the director gets less than 90% shareholder support as our measure of performance. We find that only 1.3% of directors in the bottom decile of predicted dissent (i.e. directors the algorithm predicted had the lowest chance of receiving less than 90% support) end up with strong dissent from shareholders. This ratio increases to 23% for directors in the top decile of predicted dissent.

5.4. Predicting Director Turnover

Director turnover is argued to be a powerful measure of director-firm match quality – i.e., how a director performs on a given board of a given firm (see e.g., Ferreira, Ginglinger, Laguna, and Skalli, 2017). Therefore, we train the *XGBoost* algorithm to predict whether a director will leave within two years following her appointment. Figure 6 documents the way in which the *XGBoost* algorithm's predictions compare with the actual director turnover. Specifically, the figure shows the average observed director turnover within two years of appointment across the ten deciles of *XGBoost* predicted turnover in the 2012-2014 test period. There is clearly a monotonically increasing relationship between the mean fraction of directors who left within two years in the test sample and the algorithm's predictions, as we observed in Figure 1 for excess votes. The difference in mean turnover between the bottom and top deciles is large in magnitude: while less than 2% of directors in the bottom decile leave within two years, this fraction increases to 43% in the top decile. The unconditional mean is about 11%.

Overall, the estimates using these alternative measures of director performance are consistent with the ones presented above using *excess votes*. Regardless of the measure we use, machine learning algorithms can predict the distribution of future director performance.

6. Characteristics that Affect Director Performance

One of the differences between machine learning algorithms and traditional econometric modeling is that machine learning algorithms do not provide an easy formula that can be used to infer the influence of any particular independent variable on performance. To take prediction-based actions, however, it is helpful to understand what led a model to make a specific prediction. In recent years, *explainable AI*, an important and growing strand in the machine learning literature, has focused on improving model interpretability in an attempt to make black-box models more transparent (e.g. Lundberg and Lee, 2016 and Ribeiro, Singh and Guestrin, 2016). In this section, we employ state of the art techniques to gain insights into our machine learning model. We use SHAP values (Shapley Additive exPlanations) which are derived from Shapley values (Shapley, 1951) and were introduced by Lundberg and Lee (2016) as a way to reverse-engineer a model's output. Then, we explore the characteristics of directors that are overrated and underrated in the nomination process and design tests to understand why some firms would hire predictably bad directors. We conclude the section by discussing some potential reasons.

6.1 Factors affecting Predicted Director Quality

We would like to use the algorithm's predictions to learn more about the decision-making process that governs the nomination of corporate directors. For that purpose, we study the importance of characteristics that our algorithm uses to predict a given director's performance when she joins a given board. We use SHAP values to quantify the contribution of each feature to predicting director performance. SHAP values capture the marginal contribution of each feature. They can be computed for each observation, i.e. each individual prediction. This improves the model's transparency (*local interpretability*) as it shows which variables were instrumental in generating a specific prediction. SHAP values for each feature can be averaged across observations in order to improve the *global interpretability* of the model. This produces a ranking of variables to understand which contribute the most to the predictions, across observations.

6.1.1 Global interpretability

Figure 7 reports the top ten attributes that contribute the most to *XGBoost's* predictions of director performance according to their SHAP values. Features in red (blue) are positively (negatively) correlated with the variable the algorithm predicts. Panel A uses excess votes as the performance measure. Therefore, in this panel, the variables presented in red are the ones that contribute positively to predicting director success, while the variables presented in blue contribute negatively to success. For example, being a member of the compensation committee *decreases* the excess vote prediction by 0.22% on average while being on the audit committee *increases* it on average by 0.16%. The standard deviation of incumbent directors' time on board (which could proxy for groupthink) as well as the total number of boards the director sits on currently and sat on in the past *decrease* their predicted performance on the current board.[21]

In Panel B, we calculate SHAP values for attributes when using an alternative measure of director performance: the dummy variable indicating dissent. This variable is also based on votes but focuses on a relatively extreme outcome, identifying bad directors that get less than 90% of the votes. Similar attributes are identified as contributing to dissent predictions.

In Panel C, we switch our performance measure to director turnover. Since this measure is based on a director's possibly forced decision to leave, different variables, mostly individual characteristics of these directors, appear to explain turnover. One exception is the classified board, which has a negative effect on turnover. Having a classified board by definition should lead to lower director turnover. The top three attributes predictive of leaving the board within two years are being the chairman, being an entrepreneur, and having graduated from an Ivy League university.

SHAP values do not enable a causal interpretation. They quantify features' contribution to the model based on correlations. However, they provide a helpful basis for expanding our qualitative understanding of how machine learning models generate predictions.

6.1.2 Local interpretability

SHAP values can be computed for individual observations. For example, the three panels in Figure 8 report the attributes that contributed the most to three randomly selected observations, one for each measure of director performance. For each, the "model output value" is the model's prediction for this specific observation ($\hat{y}$), while the "base value" is the mean prediction across all observations. Features in red (blue)

[21] These attributes appear to be consistent across models. For example, we find an 80% overlap with the top-ten Lasso-selected covariates (with the same direction of the effect for each attribute).

increase (decrease) the individual prediction $\hat{y}$. The value of the attribute is reported (i.e., the X) for each observation. The arrow's length for each attribute corresponds to its SHAP value (i.e. longer arrows represent more important attributes for this observation). The difference between the model output value and the base value is equal to the sum across all attributes of all the SHAP values for a given observation.

We provide an example of a typical director using the example in Panel A. The most important variable affecting the model's prediction in this case was that this incoming director sat on a total of eleven unlisted boards, which pushed the predicted *excess votes* downwards. However, the fact that the director is joining the audit committee but not the compensation committee pushed the prediction on performance upwards. The standard deviation of incumbent directors' time on the board is six years (two years above the sample average), which contributed to increasing predicted performance $\hat{y}$. The incoming director is currently sitting on two listed boards and 76% of the focal firm's stock is owned by institutions. Both attributes increased the prediction $\hat{y}$.

Interestingly, some attributes with a given value sometimes push $\hat{y}$ up and sometimes down (e.g. classified board). This speaks to the importance of interactions and non-linearities in making predictions of director performance and to the fact that there is no one-size fits all "good" corporate governance. What may be a "bad" attribute for one firm or director, may be a positive one for another. This supports our rationale for using machine learning in the context of predicting director quality.

6.1.3 Multivariate comparisons

Next, we estimate OLS models to evaluate the marginal impact of the director, board, and firm variables on predicted performance. We present the estimates in Table IA2. Director variables related to predicted subsequent shareholder support are gender, a dummy variable that indicates whether the director is "busy," and the number of listed boards a director serves on. In particular, the algorithm suggests that male directors and directors who are on at least three boards ("busy" directors) tend to receive less support from shareholders. This pattern could reflect the commonly stated concern of shareholders that directors are too often the same people, are on many boards but do not monitor to the extent that shareholders would like (see for example Biggs (1996)). Consistently, network size has a significantly negative coefficient as well.

Board level variables that are significantly related to the predicted shareholder support are the size of the board, the average tenure of incumbent board members, and the average number of independent directors. These variables again are likely to reflect the independence of the board from management. Firm-level variables that appear to be associated with subsequent performance are size (total assets), operating performance, and whether the firm pays dividends.

The SHAP values presented above identified current or past service on multiple boards as one of the key features contributing to predictions of directors' performance in a negative way as well. This pattern could potentially reflect that these directors are often too "busy" to be diligent advisors and monitors.[22]

6.2. Overvalued Director Characteristics

In this section we consider the individual characteristics of *predictably bad directors*, i.e., directors who the model predicts will do poorly but are chosen anyway, and subsequently do poorly. These characteristics can help us identify the individual director features that tend to be overvalued by firms when they select new directors. To do so, we identify directors who were nominated but were of predictably low quality and we compare them to those directors the algorithm would have preferred for that specific board position.[23] The patterns of discrepancies between these two groups recognize the types of directors that tend to be overvalued in the nomination process.

In Table 7, we report characteristics of these *predictably bad directors* who were nominated, were predicted to do poorly, and indeed did perform poorly. Compared to promising candidates identified by the algorithm, predictably unpopular directors are on average more likely to be male, have a larger professional network and more current and past directorships.[24]

[22] Fich and Shivdasani (2006) present evidence suggesting that a director being overly busy can meaningfully affect their monitoring of management.
[23] Predictions for candidates assume the same committee assignments as the nominated director. We find very similar results for all alternative specifications mentioned in previous sections, including when using Lasso to generate the predictions.
[24] We find very similar results for all alternative specifications mentioned in previous sections, including when using Lasso to generate the predictions.

These results highlight the features that are likely overrated by management when nominating directors. They are consistent with the view that directors tend to come from an "old boys club", in which men who have sat on a lot of boards are chosen to be directors. The underlying reason for this pattern, however, is not clear. As suggested by the literature on boards going back to Smith (1776) and Berle and Means (1932), managers and existing directors could implicitly collude to nominate new directors unlikely to rock the boat and upset the rents managers and existing directors receive from their current positions. Alternatively, a long literature in psychology dating to Meehl (1954) and highlighted in Kahneman (2011) has found that even simple algorithms can outperform interviews by trained professionals at predicting subsequent performance in a number of contexts. It is possible that managers and boards could be attempting to find value-maximizing directors but because of behavioral biases, underperform the algorithms we present.

6.3. Why do Firms Pick Predictably Bad Directors?

A key issue in corporate governance is why firms regularly choose directors who can be predicted to do a poor job. Our model provides a way to address this issue by allowing us to examine the types of firms that are most likely to hire predictably bad directors. Table 8 provides estimates of probit equations that estimate the likelihood that a firm hires a predictably bad director.

One possibility is that these choices reflect agency problems. If so, we expect firms with worse overall governance to be more likely to hire predictably worse directors. If boards are simply "mispredicting" future performance, we would not expect governance quality to affect the probability of hiring predictably bad directors. As a measure of governance quality, we use Bebchuk et al.'s (2009) Entrenchment Index (E-index). The index is based on six governance attributes and is constructed so that it increases as the firm-level governance gets worse. The estimates in Table 8 indicate that there is a statistically-significant positive relationship between a firm's E-index and its likelihood of selecting predictably bad directors. This finding suggests that bad choices of directors are a consequence of a firm's overall poor corporate governance.[25]

---

[25] Confirming this interpretation of the results is the fact that governance variables are not significant when examining the likelihood of selecting a director who received low support but this was not predicted by the algorithm (unpredictably bad directors).

Next, we use the percentage of *co-opted* directors, measured as the fraction of the board directors appointed after the CEO assumed office. Coles et al. (2014) provide results suggesting that board monitoring decreases as co-option increases. Therefore, we predict that there should be a positive relation between the number of co-opted directors and the likelihood of a board selecting a predictably bad director. We do find a positive and significant relation between predictably bad directors and % co-opted boards. This finding is consistent with the view that CEOs tend to "cocoon" themselves with boards of directors that are less likely to monitor him/her effectively.

Finally, we study whether the fraction of independent directors on a board affects the choice of new directors. With this idea in mind, Congress passed the Sarbanes Oxley Act in 2012, requiring exchange-listed firms to have a majority of independent directors. Consistent with the agency view, the estimates in Table 8 imply that the likelihood of a firm selecting a predictably bad director decreases with the fraction of independent directors on the board.

Overall, the results in Table 8 suggest that firms that nominate directors who were predictable poor choices have worse governance structures. Agency conflicts appear to distort nominating decisions and CEOs like to be "cocooned," and not exposed to directors who are more likely monitor them.

7. Summary and Discussion

We develop machine learning algorithms that could potentially help firms choose directors for their boards. In developing these machine learning algorithms, we contribute to our understanding of corporate governance, specifically boards of directors, in at least four ways. First, we evaluate whether it is possible to construct an algorithm that can predict whether a particular individual will be successful as a director in a particular firm. Second, we compare alternative approaches to forecasting director performance; in particular, how traditional econometric approaches compare to newer machine learning techniques. Third, we identify the characteristics that tend to be associated with effective directors. Finally, we use the selections from the algorithms as benchmarks to understand the process through which directors are actually chosen and identify the types of individuals who are more likely to be chosen as directors *counter* to the interests of shareholders. In line with the "cocooning" view of director selection, we show that firms that select predictably bad directors tend to have worse governance structures.

There are a number of methodological issues we must address before we can design such an algorithm. We must be able to measure the performance of a director to predict which potential directors will be of highest quality. Our main measure of director performance comes from the level of support a director receives from shareholders when they vote relative to other directors at the same firm. This vote-based performance measure is an *individual* measure which reflects the support the director personally has from the shareholders she represents and which should incorporate all publicly available information about her performance. We also use firm profitability, the firm's abnormal returns at the time of the announcement of a director's appointment, a dummy variable indicating whether more than 10% of shareholders opposed the director as measures of director performance, and a dummy variable on whether the director leaves within two years of appointment.

Using publicly available data on firm, board, and director characteristics, our algorithm can predict the success of directors using any of these measures. The fact that the machine learning models outperform econometric approaches is consistent with the arguments of Athey and Imbens (2017) and Mullainathan and Spiess (2017): machine learning is a valuable approach for prediction problems in the social sciences.

There is an additional methodological issue we need to address before we can conclude that algorithms can help us understand the director nomination process. We observe the predictive accuracy of our algorithm only for directors who were nominated. We design a quasi-labels procedure which exploits the fraction of votes plausible candidates received at the company whose board they joined as an indication of their performance. We find that directors the algorithm predicted would do poorly (well) indeed do poorly (well) when compared to realistic alternatives.

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. Comparing predictably bad directors to promising candidates suggested by the algorithm, it appears that firms choose directors who are more likely to be male, have a large network, and have many past and current directorships. In a sense, the algorithm is saying 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 are more likely to monitor management. In addition, less connected directors potentially provide different and potentially more useful opinions about policy. For example, TIAA-CREF (now TIAA) has had a corporate governance policy aimed in large part at diversifying boards of directors since the 1990s for this reason (see Biggs (1996) and Carleton et al. (1998)).[26]

Our finding on the predictability of which directors will or will not be popular with shareholders has important implications for corporate governance. Observers since Smith (1776) and Berle and Means (1932) have been concerned about whether managers intentionally select boards that maximize their own interests rather than those of the shareholders. The psychology literature started by Meehl (1954) has found that due to behavioral biases, even simple algorithms can outperform humans in deciding on personnel decisions. We believe that machine learning algorithms, with their powerful predictive ability, present an opportunity for firms to improve their selection process.

A natural question concerns the applicability of algorithms such as the ones we developed in practice. We view our work as a "first pass" approach, aimed at bringing the predictive power of machine learning tools to the issue of director selection. More sophisticated models with richer data would undoubtedly predict individual director performance better than the models presented here. 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 this paper. An important benefit of algorithms is that they are not prone to the agency conflicts that occur when boards and CEOs together select new directors.

[26] Similarly, Glenn Kelman, the CEO of RedFin, recently wrote: "Redfin has recently completed a search for new board directors, […] and we had to change our process, soliciting many different sources for candidates rather than relying exclusively on board members' connections. If you don't pay attention to diversity, you'll end up hiring people who are nearest at hand, who have had similar jobs for decades before. This is how society replicates itself from generation to generation, in a process that seems completely innocuous to those who aren't the ones shut out." https://www.redfin.com/blog/2016/11/how-to-triple-the-number-of-women-appointed-to-boards-in-three-years.html

Institutional investors are likely to find the algorithm's independence from agency conflicts particularly appealing and are likely to use their influence to encourage boards to rely on an algorithmic decision aid such as the one presented here for director selections in the future. An important advantage of an algorithm over the way in which directors have been chosen historically is that "algorithms can overcome the harmful effects of cognitive biases" (Sunstein, 2018). Rivera (2012) studies the hiring practices of top investment banks, consulting and law firms and concludes that recruiters overvalue personal fit which is not necessarily a function of expected performance. In the context of lower skill workers, Hoffman et al. (2017) find that managers who hire against test recommendations end up with worse average hires. Cowgill (2018) shows that the job-screening algorithm at a software company prefers "nontraditional" candidates. Our results suggest that the same idea applies to the nominating of corporate directors. Including algorithmic input to limit (but not strip) discretion and reliance on soft information in these decisions could help minimize agency problems, and thus lead to a modified rank ordering of candidates that could in turn lead to better directors than the current process.

On the other hand, if the algorithm omits attributes of potential directors that are valuable to management, such as specialized knowledge of an industry or government connections, then it potentially could lead to suboptimal solutions. This is why tools built on algorithms are likely in practice to be valuable aids in decision-making, but not substitutes for human judgement. Humans and machines both have limits and make different kinds of mistakes, i.e. they tend to have uncorrelated errors. Achieving the right balance in the division of labor between humans and machines to take advantage of their relative strengths is key.[27]

In this paper, we use 21st century technology to confirm an observation that dates back over two hundred years: the board selection process leads to directors who are often those nearest at hand and are not necessarily the best choices to serve shareholders' interests. This technology can, however, in addition to confirming this observation, provide us with the tools to change it. By providing a prediction of performance for *any* potential candidate, a machine learning algorithm could expand the set of potential

[27] The issues around the consequences of AI-based decisions are exposed in grounded discussions in Agrawal, Gans and Goldfarb (2018)

directors and identify individuals with the skills necessary to become successful directors, who would have otherwise been overlooked. We expect that in the not too distant future, algorithms will fundamentally change the way corporate governance structures are chosen, and that shareholders will be the beneficiaries.

References

Abis, Simona (2017) Man vs. Machine: Quantitative and Discretionary Equity Management, working paper.

Adams, R. B (2017) Boards and the Directors Who Sit on Them. Chapter 6 of *Handbook of the Economics of Corporate Governance,* edited by B. E. Hermalin and M. S. Weisbach, Elsevier.

Adams, R., B. Hermalin and M. Weisbach (2010) The Role of Boards of Directors in Corporate Governance: A Conceptual Framework and Survey. *Journal of Economic Literature* 48, 58–107.

Alexander, C., M. Chen, D. Seppi, and C. Spatt (2010) Interim News and the Role of Proxy Voting Advice. *Review of Financial Studies* 23: 4419-4454.

Aggarwal, R., S. Dahiya and N. Prabhala (2017) The Power of Shareholder Votes: Evidence from Uncontested Director Elections. *Georgetown McDonough School of Business Research Paper*.

Aggarwal, R., I. Erel, and L. Starks (2016) Influence of Public Opinion on Investor Voting and Proxy Advisors, Working Paper, Ohio State University.

Agrawal, A., J. Gans and A. Goldfarb (2018) Prediction Machines. *Harvard Business Review Press*, Boston, Massachussetts.

Athey, S. (2017) Beyond Predictions: Using Big Data for Policy Problems. *Science* 355: 483-385.

Athey, S. and G.W. Imbens (2017) The State of Applied Econometrics: Causality and Policy Evaluation, *Journal of Economic Perspectives*, 31, 3-32.

Autor, D. (2015) Why Are There Still So Many Jobs? The History and Future of Workplace Automation, *Journal of Economic Perspectives*, 29, 3-30.

Bebchuk, L, A. Cohen, and A. Ferrell, 2009, What Matters in Corporate Governance? *The Review of Financial Studies* 22, 783–827.

Biggs, John (1996) Corporate Governance Assessment: A TIAA-CREF Initiative, *Directors Monthly,* 20 (10) 1-8.

Boone, A., L. Field, and J. Karpoff (2007) The Determinants of Corporate Board Size and Composition: An Empirical Analysis, Journal of Financial Economics *8*5, 66-101.

Cai, J., J. Garner and R. Walkling (2009) Electing Directors. *The Journal of Finance,* 64(5): 2389-2421.

Cai, J., T. Nguyen, and R. Walkling (2017) Director Appointments – It is Who You Know, Drexel University Working Paper.

Carleton, Willard T., James M. Nelson and Michael S. Weisbach (1998) The Influence of Institutions on Corporate Governance through Private Negotiations: Evidence from TIAA-CREF, *The Journal of Finance*, 53, 1335-1362.

Chen, T. and C. Guestrin (2016) XGBoost: A Scalable Tree Boosting System. KDD '16 Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining: 785-794.

Coles, J., N. Daniel and L. Naveen (2008) Boards: Does One Size Fit All? *Journal of Financial Economics*, 87, 329-356.

Coles, J.L., N. Daniel and L. Naveen (2014) Co-opted Boards. *Review of Financial Studies* 27: 1751-1796.

Coles, J., N. Daniel and L. Naveen (2015) Board Groupthink. *Working Paper*.

Cowgill, B. (2018) Automating Judgement and Decision-making: Theory and Evidence from Resumé Screening. *Working Paper*.

Daines, R., I. Gow, and D. Larcker (2010) Rating the ratings: How Good are Commercial Governance Ratings? *Journal of Financial Economics* 98: 439-461.

Easley, D., M. Lopez de Prado, M. O'Hara and Z. Zhang (2020) Microstructure in the Machine Age, *The Review of Financial Studies*, forthcoming.

Edmans, A. (2012) The Link Between Job Satisfaction and Firm Value, With Implications for Corporate Social Responsibility. *Academy of Management Perspectives*: 26, 1-19.

Ertimur, Y. F. Ferri, and D. Oesch (2013) Shareholder Votes and Proxy Advisors: Evidence from Say on Pay. *Journal of Accounting Research* 51 (5): 951-996.

Ertimur, Y. F. Ferri, and D. Oesch (2018) Understanding Uncontested Director Elections. *Management Science* 64(7): 3400-3420.
Fedaseyeu, V., J. Linck, and H. Wagner (2017) Do Qualifications Matter? New Evidence on Board Functions and Director Compensation, *Journal of Corporate Finance, forthcoming.*
Ferreira, D., E. Ginglinger, M. Laguna, and J. Skalli, (2017) Closing the Gap: Gender Quotas and Corporate Board Composition *ECGI - Finance Working Paper*
Fich, E. and A. Shivdasani (2006) Are Busy Boards Effective Monitors? *The Journal of Finance* 61: 689-724.
Field, L., M. Souther, and A. Yore (2017) Does diversity pay in the boardroom? University of Delaware Working Paper.
Fischer, P., J. Gramlich, B. Miller and H. White (2009) Investor Perceptions of Board Performance: Evidence from Uncontested Director Elections. *Journal of Accounting and Economics* 48: 172–189.
Fos, Vyacheslav, Kai Li, and Margarita Tsoutsoura (2018) Do Director Elections Matter? *Review of Financial Studies* 31: 1499-1531.
Hart O. and L. Zingales (2017) Companies Should Maximize Shareholder Welfare Not Market Value. *Journal of Law, Finance, and Accounting*, 2: 247–274.
Hastie, T., R. Tibshirani, and J. Friedman, (2009) *The Elements of Statistical Learning*, Springer Series in Statistics.
Hastie, T., R. Tibsharani and M. Wainwright (2015) Statistical Learning with Sparsity: The Lasso and Generalizations. *Chapman and Hall/CRC.*
Heath, D., D. Macciocchi, R. Michaely and M. Ringgenberg (2019) Passive Investors are Passive Monitors. Working Paper.
Hermalin, B. E., and M. S. Weisbach (1998) Endogenously Chosen Boards of Directors and Their Monitoring of the CEO. *American Economic Review*, 88(1): 96–118.
Hermalin, B. E., and M. S. Weisbach (2003) Boards of Directors as an Endogenously Determined Institution: A Survey of the Economic Literature. *Federal Reserve Bank of New York Economic Policy Review*, 9(1): 7–26.
Hoberg, G. and G. Phillips (2010) Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis. *Review of Financial Studies* 23 (10), 3773-3811.
Hoberg, G. and G. Phillips (2016) Text-Based Network Industries and Endogenous Product Differentiation. *Journal of Political Economy* 124 (5) 1423-1465.
Hoffman, M., L. B. Kahn and D. Li (2017) Discretion in Hiring. *The Quarterly Journal of Economics. qjx042, https://doi.org/10.1093/qje/qjx042.*
Iliev, P., K. Lins, D. Miller and L. Roth (2015) Shareholder Voting and Corporate Governance Around the World. *The Review of Financial Studies* 28(8): 2167–2202.
Iliev, P. and M. Lowry (2014) Are Mutual Funds Active Voters? *Review of Financial Studies* 28: 446-485.
Kahneman, D. (2011) *Thinking, Fast and Slow*, Farrar, Strauss and Giroux: New York.
Ke, Zheng, B. Kelly and D. Xiu (2019) NBER Working Paper No. 26186.
Kleinberg, J., H. Lakkaraju, J. Leskovec, J. Ludwig, S. Mullainathan (2017) Human Decisions and Machine Predictions. *The Quarterly Journal of Economics* 133 (1) 237–293.
Kleinberg J., Liang A. and S. Mullainathan (2017) The Theory Is Predictive, but Is It Complete? An Application to Human Perception of Randomness. *Working Paper*.
Kleinberg, J., J. Ludwig, S. Mullainathan and Z. Obermeyer (2015) Prediction Policy Problems. *American Economic Review Papers & Proceedings* 105(5): *491–495.*
Kleinberg, J., J. Ludwig, S. Mullainathan and C. Sunstein (2019) Discrimination in the Age of Algorithms. NBER Working Paper No. 25548.
Kramarz, F. and D. Thesmar (2013) Social Networks in the Boardroom. *Journal of the European Economic Association* 11(4): 780-807.
Larcker, D., A. McCall, and G. Ormazabal (2015) Outsourcing Shareholder Voting to Proxy Advisory Firms. *Journal of Law and Economics* 58, 173-204.

Li, K., M. Feng, R. Shen, and X. Yan (2020) Measuring Corporate Culture Using Machine Learning. *The Review of Financial Studies* (forthcoming).
Linck, J, J. Netter, and T. Yang (2008) Determinants of Board Structure, Journal of Financial Economics 87, 308-328.
Linck, J, J. Netter, and T. Yang (2009) The effects and unintended consequences of the Sarbanes-Oxley Act on the supply and demand for directors, Review of Financial Studies 8, 3287–3328.
Lundberg, S, S. Lee (2017) A Unified Approach to Interpreting Model Predictions, Advances in Neural Information Processing Systems 30.
Malenko, N., and Y. Shen (2016) The Role of Proxy Advisory Firms: Evidence from a Regression-Discontinuity Design, *The Review of Financial Studies* 29: 3394–3427.
Masulis, R. and S. Mobbs (2014) Independent Director Incentives: Where do Talented Directors Spend their Limited Time and Energy, *Journal of Financial Economics* 111, 406-429.
Meehl, P. E. (1954) *Clinical vs. Statistical Prediction: A Theoretical Analysis and a Review of the Evidence*, University of Minnesota Press: Minneapolis.
Miller, A. (2018) Want Less-Biased Decisions? Use Algorithms. *Harvard Business Review.*
Mitchell, T. (1997) Machine Learning. *McGraw Hill*.
Mullainathan, S. and J. Spiess (2017) Machine learning: An Applied Econometric Approach, *Journal of Economic Perspectives*, 31, 4: 87-106.
Pearl J. and D. Mackensie (2018) The Book of Why. *Basic Books.*
Peters, R. H. and L. A. Taylor (2017) Intangible capital and the investment-q relation. *Journal of Financial Economics* 123 (2): 251-272.
Ribeiro, M, S. Singh and C. Guestrin (2016) “Why Should I Trust You?” Explaining the Predictions of Any Classifier, *KDD* 2016, SanFrancisco, CA.
Rivera, L. (2012) Hiring as Cultural Matching: The Case of Elite Professional Service Firms, *American Sociological Review*, 77: 999-1022.
Rossi, A. (2018) Predicting Stock Market Returns with Machine Learning, working paper.
Shapley, L. S. (1951) Notes on the n-Person Game -- II: The Value of an n-Person Game. Santa Monica, Calif.: RAND Corporation.
Shivdasani, A. and D. Yermack (1999) CEO Involvement in the Selection of New Board Members: An Empirical Analysis, *The Journal of Finance*, 54, 1829-1853.
Shmueli, G. (2010) To Explain or to Predict? *Statistical Science* 25(3): 289-310.
Smith, A. (1776) *An Inquiry into the Nature and Causes of the Wealth of Nations*, Indianapolis, Liberty Press.
Sunstein, C. (2018) Algorithms, Correcting Biases. Working paper.
Suresh, H. and J. V. Guttag (2019) A Framework for Understanding Unintended Consequences of Machine Learning. arXiv:1901.10002. In the Proceedings of *Association for the Advancement of Artificial Intelligence* (AAAI) *2019.*
Yermack, D. (1996) Higher Market Valuation of Firms with a Small Board of Directors, *Journal of Financial Economics* 40: 185-211.
Zhao Jieyu, T. Wang, M. Yatskar, V. Ordonez, K. Chang (2017) Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints, *EMNLP* 2017.

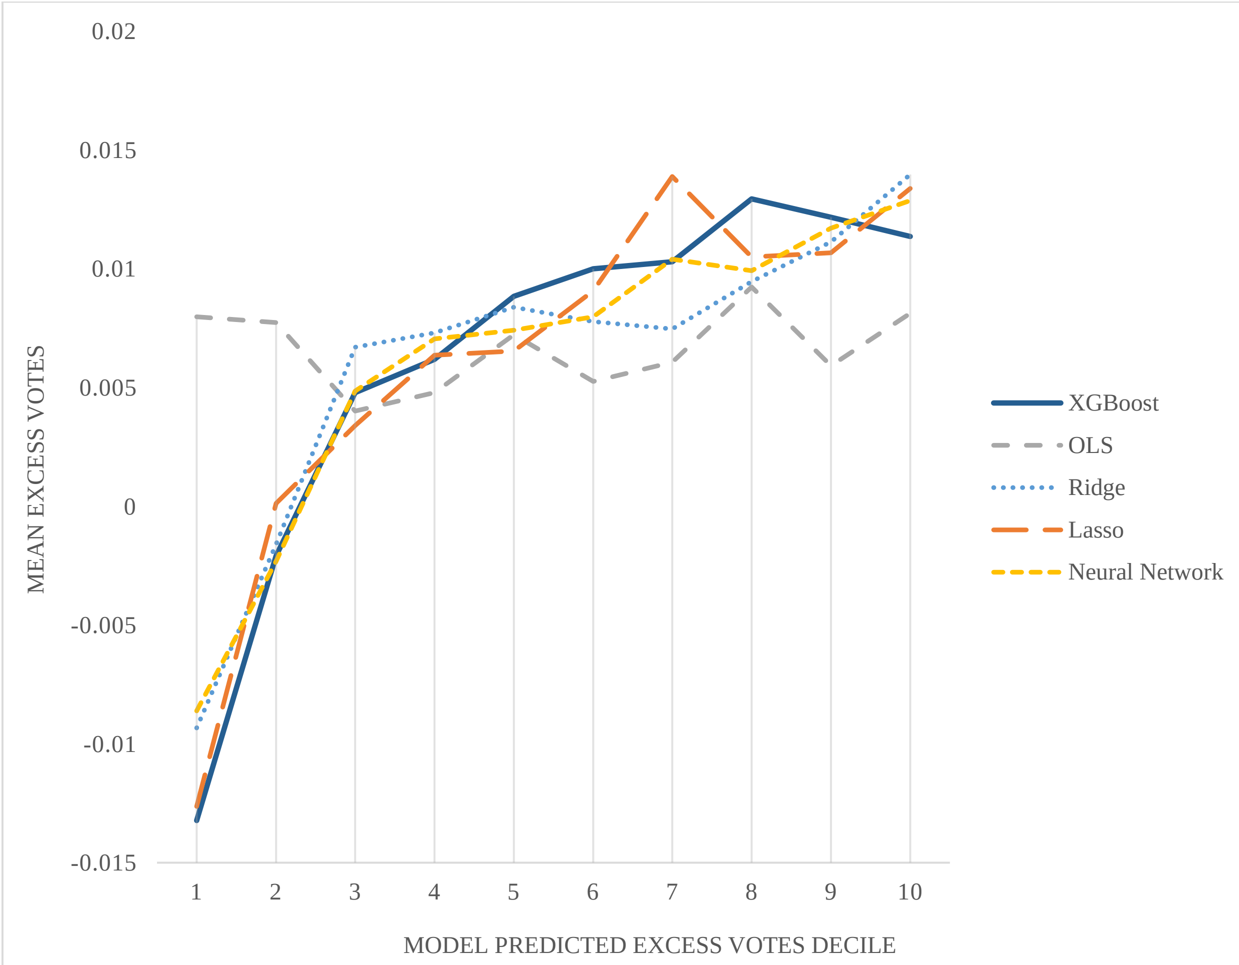

FIGURE 1: MEAN OBSERVED *EXCESS VOTES* VS. PREDICTED *EXCESS VOTES*

This figure shows the average observed level of excess shareholder support for directors across the ten deciles of predicted performance for OLS and ML models in the 2012-14 test set. To compute *excess votes*, we first compute the fraction of votes in favor of a given director over all votes cast for the director. Next, we subtract the average of that variable for the slate of directors up for reelection that year on the focal board. Finally, we take the average of this relative vote measure over the first three years of the new director's tenure.

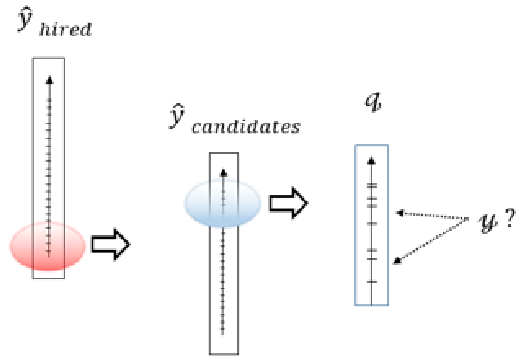

FIGURE 2: ASSESSING THE ALGORITHM'S PREDICTIONS USING QUASI-LABELS

This figure shows the procedure to evaluate our algorithmic predictions using quasi-labels. We rank all hired directors in our test set according to their predicted performance ($\hat{y}_{hired}$). The bottom decile represents directors who were predicted to receive low shareholder approval. For each of these hired directors, whom our algorithm predicted would be unpopular, we consider their associated candidate pool and rank candidates in this candidate pool according to their predicted performance on the focal board ($\hat{y}_{candidates}$). We retain the top decile of candidates, who are the most promising candidates based on our algorithms' predictions. We then re-rank these promising candidates according to their quasi-labels q, i.e. their performance on the board they actually joined. We then to compare the observed performance of the hired director on the focal board (y) to the quasi-labels of promising candidates.

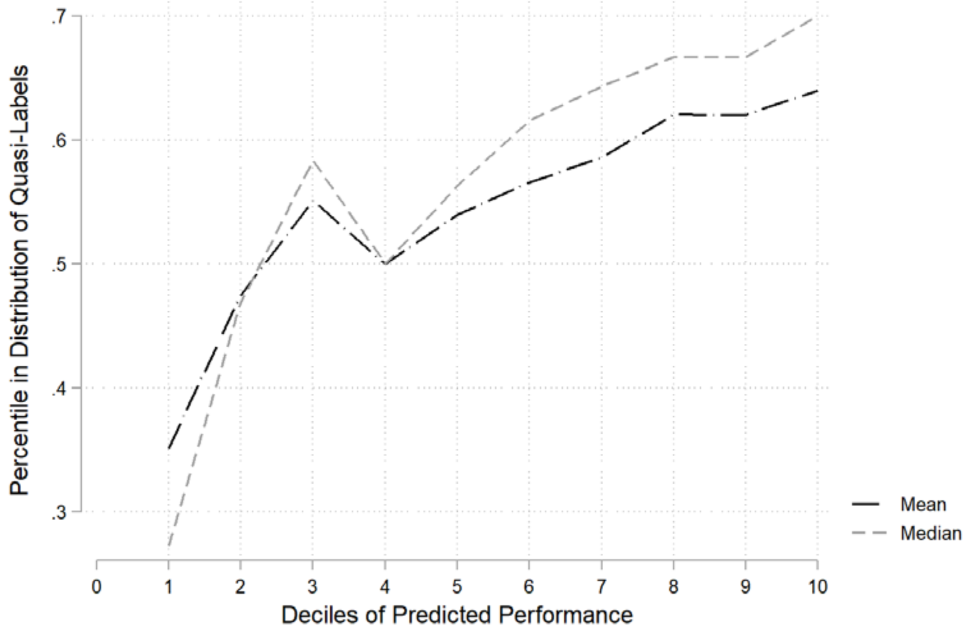

FIGURE 3: MEAN AND MEDIAN RANK IN QUASI-LABEL DISTRIBUTION ACROSS DECILES OF PREDICTED PERFORMANCE

This figure shows the mean and median rank in the distribution of quasi-labels for directors in each of the ten deciles of *XGBoost*-predicted performance (*Excess votes*). The observed performance of nominated directors in our test set is compared to the quasi-labels of *all* potential candidates in their respective candidate pool: Each new board appointment in the test set is associated with a candidate pool, comprised of directors who, within one year of the appointment, joined the board of a smaller neighboring firm

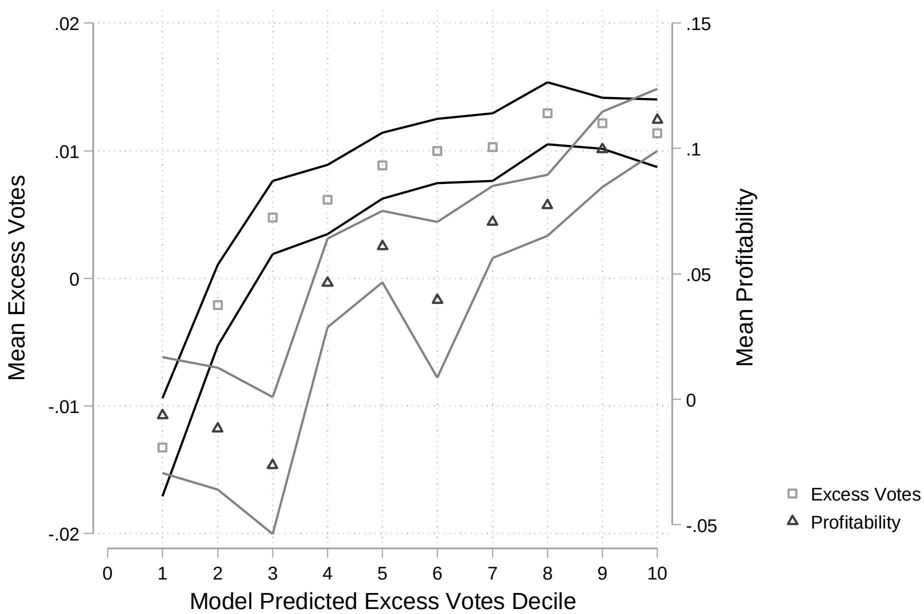

FIGURE 4: OBSERVED PERFORMANCE (EXCESS VOTES AND PROFITABILITY) ACROSS DECILES OF PREDICTED PERFORMANCE

This figure reports the actual mean *Excess Votes* (left y-axis) and mean firm profitability (right y-axis), with their respective 95% confidence interval, for each decile of *XGBoost*-predicted director performance (Excess Votes) for directors in the test set.

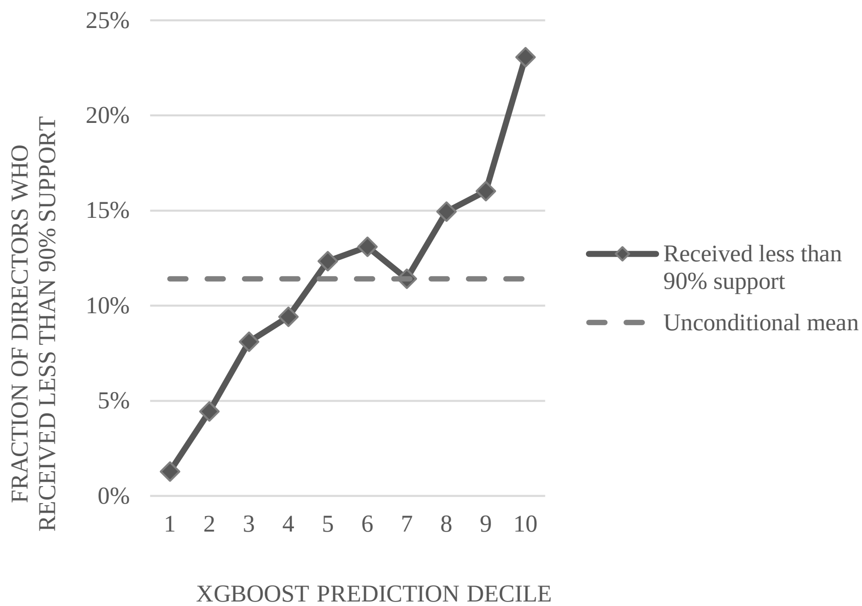

FIGURE 5: MEAN OBSERVED *DISSENT* VS. PREDICTED *DISSENT*

This figure shows the average observed level of dissent against a director by shareholder - i.e., fraction of new directors who received less than 90% shareholder support in their reelection within the next three years - across the ten deciles of *XGBoost*-predicted dissent in the 2012-14 test set. Directors in decile one (ten) were predicted by *XGBoost* as least (most) likely to face strong decent. Unconditional mean dissent is about 11%.

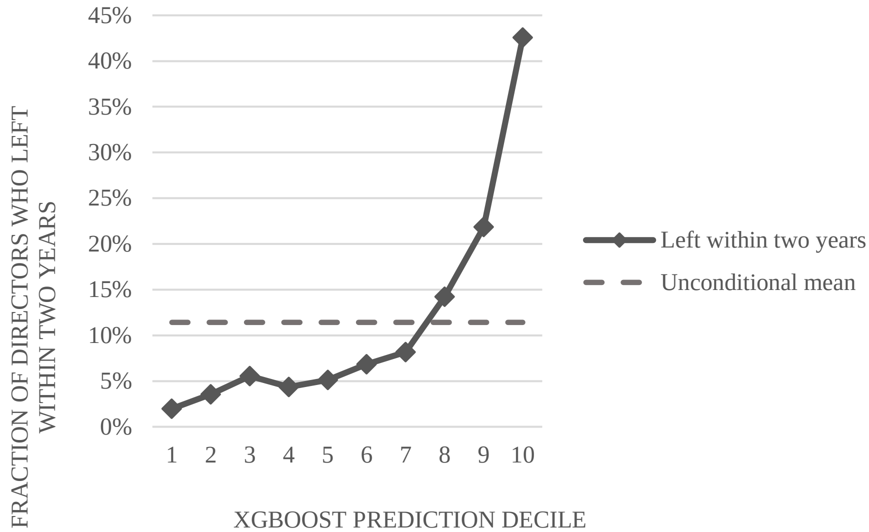

FIGURE 6: MEAN OBSERVED *DIRECTOR TURNOVER* VS. PREDICTED *TURNOVER*

This figure shows the average observed level of director turnover - i.e., fraction of new directors who left within the next two years - across the ten deciles of *XGBoost*-predicted turnover in the 2012-14 test set. Unconditional mean turnover is about 11%.

Panel A: Excess Votes as a Measure of (Better) Director Performance

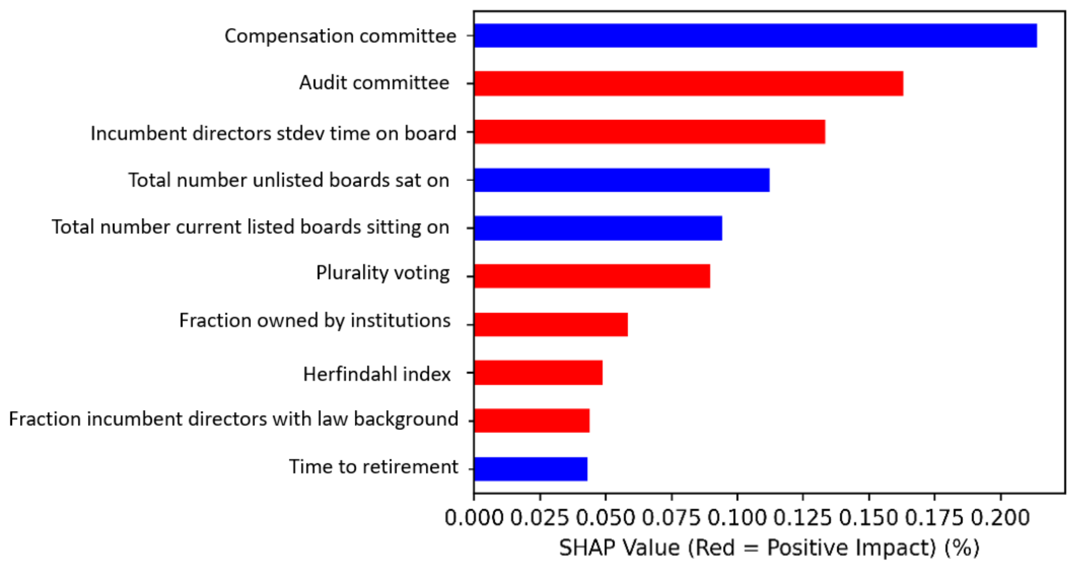

Panel B: (Larger than 10%) Dissent as a Measure of (Worse) Director Performance

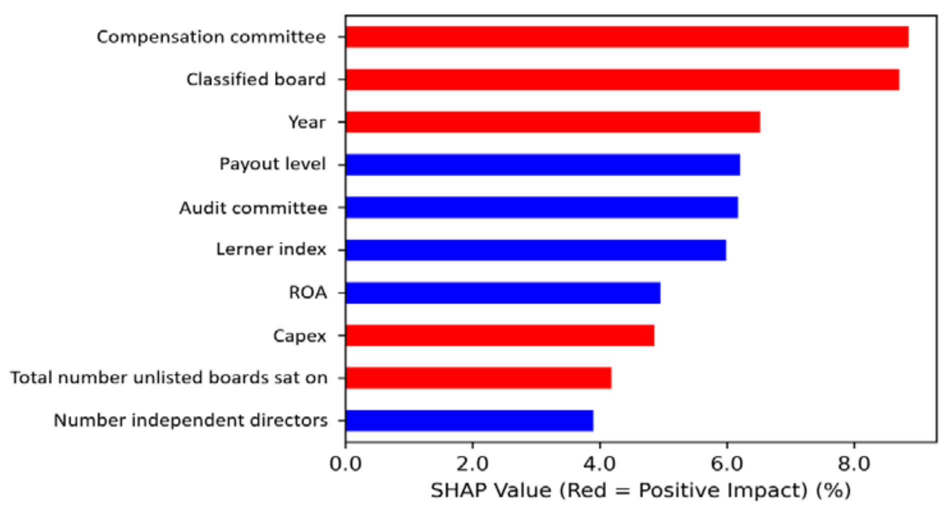

Panel C: Director Turnover as a Measure of (Worse) Director Performance

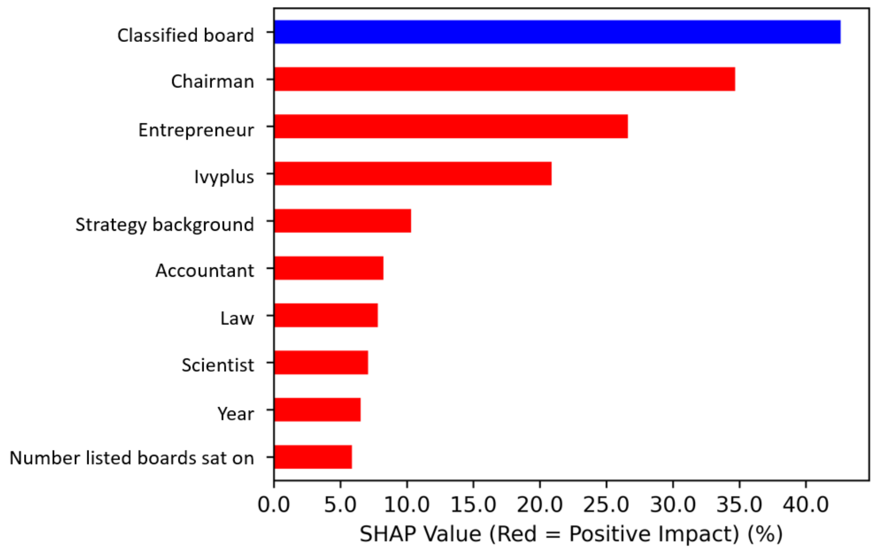

FIGURE 7: VARIABLE IMPORTANCE PLOT

This figure presents the SHAP values for the top ten characteristics in terms of variable importance in predicting director performance. We use the *XGBoost* algorithm in predictions. Variables are ranked in decreasing order of importance. Panel A uses excess votes, Panel B uses larger than 10% dissent, and Panel C uses director turnover within the next two years of appointment as a measure of director performance. While higher excess votes represent better performance, higher dissent or turnover represent worse performance. Features in red (blue) are positively (negatively) correlated with the variable the algorithm predicts.

Panel A: Predicting Excess votes

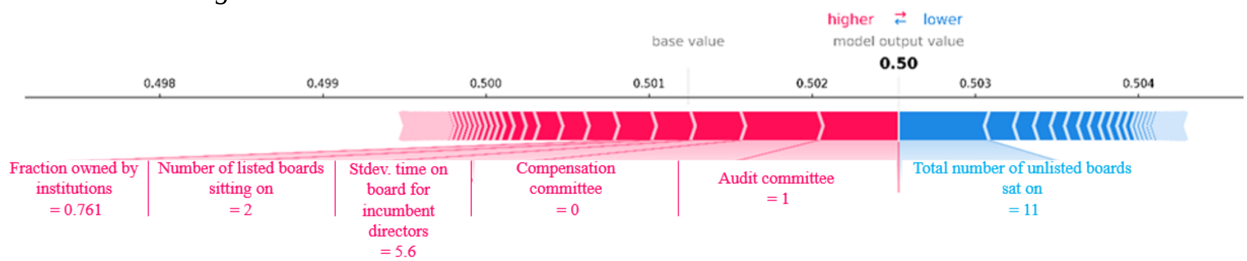

Panel B: Predicting Dissent

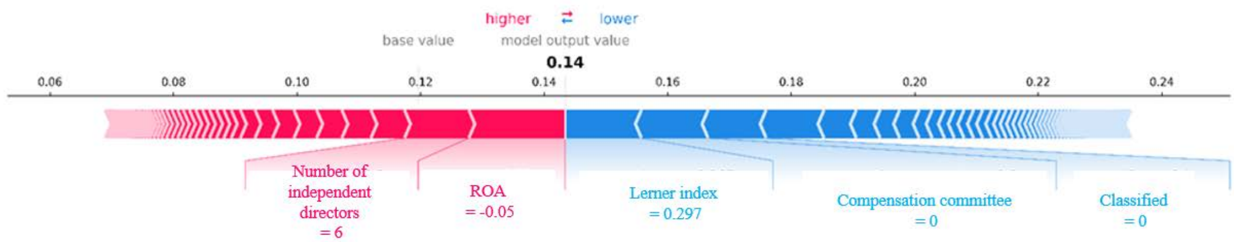

Panel C: Predicting Turnover

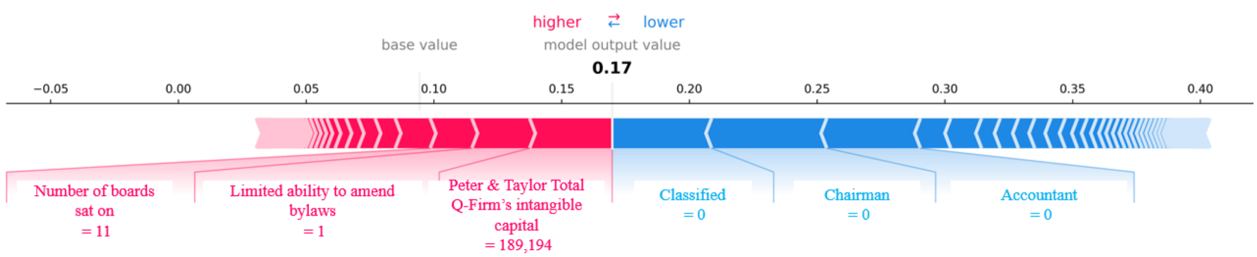

FIGURE 8: INDIVIDUAL SHAP VALUES FOR LOCAL INTERPRETABILITY

This figure shows the variables that contributed the most to the *XGBoost*-generated prediction for three random individual observations. Panel A is for an observation when the model predicts excess votes, Panel B when it predicts dissent and Panel C when it predicts turnover. The "model output value" is the model's prediction for this specific observation ($\hat{y}$). The "base value" is the mean prediction across all observations. Features in red (blue) increase (decrease) the individual prediction $\hat{y}$. The value of the attribute is reported for each observation. The arrow's length for each attribute corresponds to its SHAP value (i.e. longer arrows represent more important attributes for this observation). The difference between the model output value and the base value is equal to the sum across all attributes of all the SHAP values for a given observation.

	n	Total Votes	Excess Votes	Strong Dissent	Turnover
2000	331	0.950	0.001	0.127	0.132
2001	772	0.944	0.000	0.134	0.131
2002	1,057	0.946	0.002	0.118	0.086
2003	1,774	0.951	0.006	0.087	0.081
2004	2,019	0.953	0.007	0.086	0.104
2005	1,893	0.948	0.005	0.103	0.108
2006	1,789	0.941	0.005	0.129	0.099
2007	1,942	0.940	0.005	0.160	0.095
2008	1,691	0.944	0.007	0.155	0.112
2009	1,541	0.948	0.007	0.116	0.111
2010	1,842	0.948	0.004	0.136	0.114
2011	1,825	0.954	0.004	0.118	0.114
2012	1,862	0.952	0.005	0.111	0.113
2013	2,148	0.948	0.003	0.131	0.126
2014	1,568	0.959	0.006	0.094	0.098
	24,054	0.948	0.004	0.120	0.108

TABLE 1: DIRECTOR PERFORMANCE MEASURES SUMMARY STATISTICS

This table presents the *mean* for total and excess shareholder support over time, as well as for strong dissent against a director and director turnover. Shareholder support is defined as the fraction of votes in favor of a given director over all votes cast for the director's reelection within three years of her tenure. Dissent is equal to one for a director whose average total votes is less than 90% within three years of her tenure. To compute *Excess Votes*, we subtract the average of that variable for the slate of directors up for reelection that year on the focal board. We take the average of this relative vote measure over the first three years of the new director's tenure. Turnover is equal to one if the director leaves the board within two years of being appointed. The data is from ISS Voting Analytics and BoardEx.

	Full sample	yes	no	Difference p-value
Director level				
Male	0.102	0.106	0.079	0.000
Foreign	0.101	0.115	0.100	0.138
Qualifications > median	0.102	0.094	0.106	0.005
Network size > median	0.102	0.108	0.096	0.002
Generation BBB	0.101	0.093	0.118	0.000
Generation X	0.101	0.151	0.096	0.000
Busy director	0.102	0.145	0.090	0.000
Finance background	0.102	0.106	0.101	0.328
Board level				
Fraction male > median	0.102	0.116	0.091	0.000
Board size > median	0.102	0.089	0.114	0.000
Nationality mix > median	0.102	0.108	0.100	0.064
Attrition rate > median	0.098	0.106	0.086	0.000

TABLE 2: AVERAGE FRACTION OF POOR OUTCOME

This table presents the average fraction of "poor outcome" for various director-level and board-level characteristics. A director is considered to experience a poor outcome if her excess votes is < -2%. Poor outcomes represent 10% of the sample.

	Predicted Percentile of Excess Votes	Average Observed Performance for Directors in a Given Percentile of Predicted Performance as Predicted by: OLS	XGBoost	Ridge	Lasso	Neural Network
Directors predicted to perform **poorly**	1%	0.030	-0.031	-0.017	-0.026	-0.017
	5%	-0.002	-0.012	-0.011	-0.018	-0.011
	10%	0.007	0.001	-0.004	0.002	-0.003
Directors predicted to perform **well**	90%	0.004	0.008	0.012	0.004	0.013
	95%	-0.004	0.016	0.017	0.010	0.011
	100%	0.006	0.012	0.012	0.018	0.014

TABLE 3: OLS VS. MACHINE LEARNING TO PREDICT DIRECTOR PERFORMANCE

This table reports the average observed level of excess shareholder support over the first three years of a new director's tenure for directors who were ranked by their predicted level of shareholder support by an OLS model and several machine learning algorithms (*XGBoost*, Ridge, Lasso and Neural Network). Shareholder support is defined as the fraction of votes in favor of a given director over all votes cast for the director's reelection within three years of her tenure. To compute *Excess Votes,* we subtract the average of that variable for the slate of directors up for reelection that year on the focal board. Then we take the average of this relative vote measure over the first three years of the new director's tenure.

	Median percentile of observed performance in the distribution of quasi-labels (candidate pools)				
	OLS	XGBoost	Ridge	Lasso	Neural Network
Bottom decile of predicted performance	66th	23rd	33rd	25th	35th
Top decile of predicted performance	70th	80th	80th	84th	77th

TABLE 4: EVALUATING THE PREDICTIONS USING QUASI-LABELS

This table reports how nominated directors rank in the distribution of quasi-labels of their candidate pool. For each nominated director in our test set, we construct a pool of potential candidates who could have been considered for the position. Those candidates are directors who accepted to serve on the board of a smaller nearby company within a year before or after the nominated director was appointed. The quasi-label for each of these candidates is how she performed on the competing board she chose to sit on. The first (second) row shows the median percentile of observed performance in the distribution of quasi-labels for directors the model predicted to be in the bottom (top) decile of predicted performance. Each column presents the results from a different model.

	N	**Mean**	**Median**
Directors in Decile 1 of predicted performance (excess votes)	292	-1.94%	-0.64%
Directors in Decile 10 of predicted performance (excess votes)	575	0.75%	0.34%
Difference in means (p-value)		**0.0043**	

TABLE 5: CUMULATIVE ABNORMAL RETURNS AROUND APPOINTMENT ANNOUNCEMENTS

This table reports the mean and median cumulative abnormal returns for directors predicted to do poorly and for directors predicted to do well. Directors predicted to do poorly (well) are directors in decile 1 (decile 10) of predicted performance (excess votes) as predicted by the *XGBoost* algorithm. Results are shown for appointments in the test set only. The cumulative abnormal returns reported are computed using a (-1; +1) window.

		1	2	3	4	5	6	7	8	9	10	Decile 10 - 1 *p-value*
Algorithm trained on **profitability**	Average observed profitability	-0.498	-0.064	-0.017	0.017	0.078	0.083	0.113	0.114	0.144	0.205	0.0000
	Average observed shareholder support	0.942	0.946	0.956	0.937	0.957	0.961	0.953	0.954	0.960	0.961	0.0002
	Average observed excess votes	-0.0004	0.002	0.006	0.002	0.006	0.004	0.003	0.005	0.006	0.004	0.0668
Algorithm trained on **excess votes**	Average observed profitability	-0.006	-0.012	-0.026	0.046	0.061	0.040	0.071	0.077	0.100	0.111	0.0000
	Average observed excess votes	-0.013	-0.002	0.005	0.006	0.009	0.010	0.010	0.013	0.012	0.011	0.0000
Algorithm trained on **total votes**	Average observed profitability	-0.003	-0.032	-0.031	-0.018	0.024	0.029	0.058	0.075	0.086	0.100	0.0000
	Average observed shareholder support	0.920	0.937	0.946	0.948	0.950	0.957	0.957	0.966	0.972	0.977	0.0000

TABLE 6: COMPARING SHAREHOLDER SUPPORT MODELS WITH PROFITABILITY MODELS

This table reports the actual performance for each decile of *XGBoost*-predicted performance. *XGBoost* is trained to predict 1) firm profitability three years after the director has been appointed (EBITDA/Total Assets), 2) total votes, and 3) excess votes. The results are for our test set only (out-of-sample performance for directors appointed between 2012-2014).

	Hired directors with predicted and observed low shareholder support	Promising candidates for this board position	
	Mean	Mean	Difference *p-value*
Male	0.88	0.83	0.000
Number of qualifications	2.07	2.23	0.000
Ivy League	0.15	0.15	0.960
MBA	0.52	0.49	0.079
Network size	1714	1261	0.000
Total number of listed boards sat on	6.4	2.6	0.000
Total number of unlisted boards sat on	10.6	2.6	0.000
Total current number of boards sitting on	3.1	1.6	0.000
Number previous jobs same industry	0.07	0.10	0.004
Number previous directorships same industry	0.21	0.10	0.000
Busy	0.56	0.15	0.000
Director age	54.4	56.8	0.000
Background academic	0.043	0.014	0.000
Background finance	0.223	0.178	0.000
International work experience	0.142	0.052	0.000

TABLE 7: OVERVALUED DIRECTOR CHARACTERISTICS

This table reports the mean of director features for directors in our test set (out of sample predictions) whom our *XGBoost* algorithm predicted would be in the bottom decile of performance and indeed ended up in the bottom decile of actual performance (i.e. predictably low-quality directors) and compares it to the mean for potential candidates the board could have nominated instead, whom our *XGBoost* algorithm predicted would be in the top decile.

	(1)	(2)	(3)	(4)	(5)	(6)
E-index	0.239***	0.245*	0.265**	0.296**	0.452***	0.552***
	(2.620)	(1.946)	(2.070)	(2.193)	(2.695)	(2.682)
% co-opted directors		1.238***	1.102**	1.107**	1.102**	1.196**
		(2.767)	(2.427)	(2.446)	(2.122)	(2.099)
% board independent			-3.627***	-3.908***	-3.106*	-4.044**
			(-2.908)	(-3.007)	(-1.903)	(-1.991)
% board busy				0.616	0.361	0.365
				(0.889)	(0.397)	(0.348)
Average board tenure					-0.057	-0.067
					(-1.322)	(-1.251)
Number of institutional owners					-0.429	0.853
					(-1.284)	(1.216)
Board size					-0.071	-0.014
					(-0.939)	(-0.166)
Firm age					-0.005	-0.005
					(-0.421)	(-0.311)
ln(assets)						-0.727**
						(-2.085)
MB						-0.009
						(-0.446)
Leverage						1.319
						(1.237)
ROA						-0.420
						(-0.320)
Excess 12 month-returns						-0.849
						(-1.643)
Constant	-3.413***	-4.188***	-1.194	-1.218	1.056	-0.748
	(-9.863)	(-6.994)	(-1.067)	(-1.084)	(0.508)	(-0.271)
Observations	3,596	1,775	1,773	1,773	1,713	1,693
Pseudo R-squared	0.0309	0.109	0.174	0.180	0.266	0.353

TABLE 8: WHO HIRES PREDICTABLY BAD DIRECTORS

This table reports results from a Probit regression where the dependent variable is a dummy variable equal to one for directors flagged by *XGBoost* as most likely to face strong decent (decile 10 of dissent predictions) and ended up facing strong dissent, on various firm-level governance measures and other board-, and firm-level controls.

Internet Appendix (IA) for

"Selecting Directors Using Machine Learning"

ISIL EREL, LÉA H. STERN, CHENHAO TAN, AND MICHAEL S. WEISBACH

This Internet Appendix reports the following additional sections:

1. Section IA1 describes Machine Learning algorithms -*lasso*, *ridge*, *neural networks* and *gradient boosting trees (XGBoost)*- used in the paper to predict director performance.

2. Sections IA2 presents the framework we developed in the spirit of Kleinberg et al. (2017) to understand the issues faced when assessing the prediction accuracy of our algorithms.

3. Section IA3 provides detailed variable definitions and data sources.

4. Figures IA1-4 present various graphs on the distribution of votes and excess votes.

5. Figure IA1-5 presents XGBoost prediction results using a test period of 2011-2014 rather than 2012-2014.

6. Figure IA1-6 presents mean and median rank in quasi-label distribution across deciles of Lasso-predicted performance, rather than XGBoost-predicted one.

7. Table IA1 reports coefficients from an OLS regression of excess votes on various director, firm, and board characteristics.

8. Table IA2 provides univariate comparisons of directors predicted to be in the bottom vs Top deciles of performance by XGBoost ML algorithm.

9. Table IA3 reports multivariate comparisons: estimates from an OLS model of directors predicted performance with firm, board, and director characteristics.

IA1. ALGORITHMS USED TO PREDICT PERFORMANCE

We employ algorithms designed to make an ex ante prediction of directors' level of relative shareholder support, averaged over the first three years of their tenure. The algorithms use a set of observable director, board, and firm features that are available to the nominating committee at the time of the nominating decision. There are a number of well-known machine learning algorithms that can be used for our prediction exercise. We use four of these algorithms to predict director performance, and give a brief summary of each in this section.

IA1.1. *Lasso* and *Ridge*

OLS regressions tend to generate poor out-of-sample predictions as they are designed to minimize the in-sample residual sum of squares. This observation is known as the bias-variance tradeoff in the machine learning literature: if an algorithm fits in-sample data too well (low bias), it has high variance and thus does not perform as well on out-of-sample data. *Lasso* and *ridge* are both linear models that use a regularization term to achieve a balance between bias and variance. They do so by minimizing a loss function that includes in-sample fit and a penalty term that favors simple models, thereby reducing variance. Prediction accuracy is thus improved by setting some coefficients to zero and shrinking others. To achieve this goal, Lasso and ridge combine the minimization of the sum of the squared errors with the norm of parameters. The lasso estimator solves the problem:

$$\min_{\beta} \sum_{j=1}^{k} (y_i - x_i\beta)^2 \; + \; \lambda \cdot \|\beta\|_1$$

where $\|\beta\|_1$ is the ℓ_1-norm (least absolute deviation). The penalty weight (λ) on the sum of the absolute values of coefficients is set using the default parameter in scikit-learn[28]. *Ridge* is similar to *lasso* except that the bound on the parameter estimates is the ℓ_2-norm (least squares), therefore shrinking estimates smoothly towards zero, as opposed to setting some estimates to zero as Lasso does.[29]

[28] http://scikit-learn.org/stable/

[29] For a detailed discussion of sparse estimators, we refer interested readers to Hastie, Tibshirani and Wainwright (2015).

IA1.2. *Gradient Boosting Trees*

Gradient Boosting Trees are similar to random forest algorithms. A *random forest* algorithm is an ensemble method that combines multiple decision trees. Intuitively, a single decision tree presents a flow chart where a data point can follow the flow starting from the root to a leaf node associated with its final prediction. The selection of attributes at each node in decision trees is inspired by information theory to maximize information gain. In the *random forest* algorithm, multiple trees are estimated by using a random subset of covariates in each tree. Among those, the covariate that provides the best binary split based on information gain is used to split the data into two partitions and functions as the root of the tree. The algorithm repeats this process until it reaches the bottom of the tree, where each "leaf" or terminal node is comprised of similar observations. Then, a new data point can start at the top of each tree and follow the splits at each node all the way to a leaf node. The prediction for this new data point is the average outcome of observations in the leaf it ends up in. The random forest algorithm takes an average of the predictions from all the decision trees.

Similar to the random forest algorithm, the *gradient boosting trees* algorithm is an ensemble method that combines multiple trees. The key difference lies in that the final prediction is a linear sum of all trees and the goal of each tree is to minimize the residual error of previous trees. A decision tree is the basic building block of random forests. A decision tree defines a tree-shape flow graph to support decisions. An instance is classified by starting from the root of the tree, testing the feature specified by the node, moving down the branch corresponding to the feature value in the given instance. The ***XGBoost*** algorithm provides an efficient implementation of this algorithm that is scalable in all scenarios (Chen and Guestrin, 2016).

A key difference between decision tree learning and Ridge and Lasso regression lies in the fact that there is no explicit objective function that a decision tree optimizes. Instead, the learning process is a greedy recursive algorithm that finds the best feature to split the current data based on a criterion. In our paper, we use a decision tree regressor where the criterion aims to minimize the mean squared error in each branch. Refer to Mitchell (1997) for more details on decision tree learning.

IA1.3. *Neural Networks*

A neural network is structured in layers of neurons connected by synapses. The first layer comprises the input neurons and the final layer represents the output. Layers of neurons between the first and final layers are hidden layers. The figure below depicts the structure of a basic neural network with two hidden layers. Neurons x_i are input neurons connected to the next layer of neurons by synapses which carry weights w^1. Each synapse carries its own weight. An activation function (usually a sigmoid to allow for non-linear patterns) is embedded in each neuron in the hidden layers to evaluate its inputs. The set of weights carried by the synapses that reach a neuron are fed into its activation function, which will determine whether or not that neuron is activated. If activated, it triggers the next layer of neurons with the value it was assigned, with weight w^2 (again with each synapse carrying its own weight). Similar to the neurons in the hidden layers, the output neuron judges its input via an activation function and decides from which neurons to accept the triggered values. The output is a weighted sum of the activated neurons in the last hidden layer. Training a network involves modifying the weights on the synapses to minimize a cost function (e.g. the sum of squared errors).

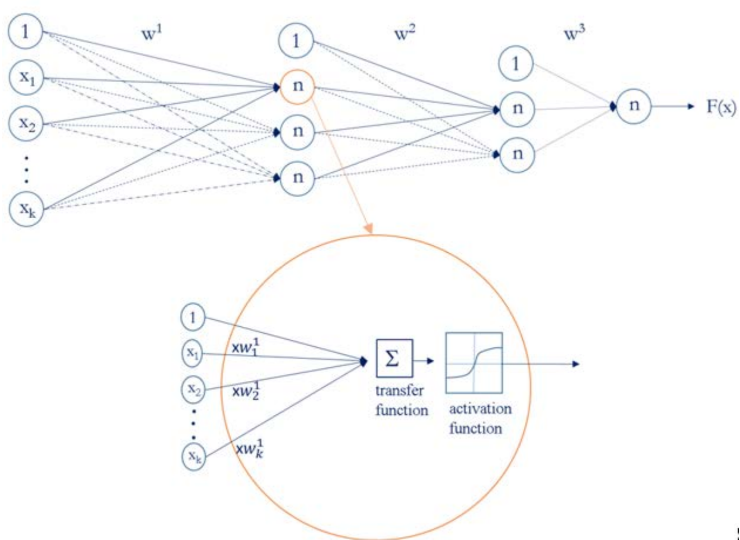

IA2. A FRAMEWORK TO ASSESS ALGORITHMS' PREDICTIONS

We develop a framework in the spirit of Kleinberg et al. (2017) to understand the issues faced when assessing the prediction accuracy of our algorithms. Suppose that the true data generating process is given by $\mathcal{Y} = \mathcal{F}(\mathcal{W}, \mathcal{Z})$, where $\mathcal{W}$ and $\mathcal{Y}$ are operationalized by W, our vector of inputs and Y, our outcome variable (i.e., director performance). $\mathcal{Z}$ represents a set of features that affect director performance and that are observable by the decision maker (board/CEO) but not by the algorithm. An example of such a feature would be idiosyncratic knowledge of the firm or its industry that would make a potential director more valuable.

In addition, there are features $\mathcal{B}$ that do *not* affect director performance and are unobservable to the algorithm, but could nonetheless affect boards' nominating decisions. Examples of such features could be a candidate's political views, or the neighborhood where she grew up. The board's preferences for certain features in $\mathcal{B}$ could be conscious or could represent an implicit bias of which they are unaware of. The key point is that these attributes can influence boards' decisions even though they are *not* correlated with performance.

$\mathcal{F}$ is operationalized by a functional form $\mathscr{f}$. For the purpose of predictive modeling, we are interested in finding a function that closely matches the function $\mathscr{f}$ in out-of-sample data. Compared to classic causal hypothesis testing, we do not make strong assumptions about the structure of $\mathcal{F}$ and thus do not focus on examining the estimated parameters and claim that these parameters match $\mathscr{f}$. In other words, our supervised machine learning algorithm seeks to learn a functional form that maps features W into predictions $\hat{\mathscr{f}}(W)$ that generalize well on out-of-sample data (Shmueli, 2010).

A director is characterized by $\vec{x}$, composed of three vectors of features as well as of outcome y:

$$\vec{x} = \begin{bmatrix} W \\ \mathrm{Z} \\ \mathrm{B} \end{bmatrix}$$

Note that x may include not only director characteristics but also firm and board level characteristics so that both the board and the algorithm try to assess a director's future performance for a specific board position.

For the purpose of the model and similar to Kleinberg et al. (2017), we shrink the dimension of $\vec{x}$ to a vector with three unidimensional characteristics *w*, *z* and *b*. In addition, we assume that the sum of *w* and *z* is distributed between 0 and 1 and that their sum equals *y* on average:

$$E[Y = y|W = w, Z = z] = E[y|w, z] = w + z$$

Each board *j* has a payoff function π_j that is a function of the director's performance as well as of the director's characteristics as defined by $\vec{x}$.

For each director (x, y) in the candidate pool $\mathcal{D}$ of size *k*, the board's payoff is characterized as:

$$\pi_j(x, y) = \underbrace{u_j y}_{\substack{benefits\ from \\ director's\ performance}} + \underbrace{v_j g_j(x)}_{\substack{benefits\ from\ hiring \\ director\ with\ characteristics\ x}}$$

$g_j(x)$ is a board specific function that maps directors' characteristics into a score. We can think of $g_j(x)$ as a measure of the utility the board derives from nominating a director with specific characteristics; for example, they could derive private benefits from nominating someone from their own network. The variables u_j and v_j represent weights that board *j* puts on director performance and on the benefits it derives from nominating a director with certain features, respectively.

We assume that board *j* chooses a nominating rule h_j such that it maximizes its expected payoff.

$$h_j \in \{0,1\}^k \ and \ \|h_j\|_0 = 1$$

$$\Pi_j(h_j) = \sum_{i \in \mathcal{D}} h_{j,i} E[\pi_j(x_i, y_i)]$$

The nominating rule h_j depends on $k_j(x)$, the board's *assessment* of future performance for a director with characteristics x. For a given $g_j(x)$, the board chooses the director with the highest $k_j(x)$. We do not observe boards' relative weights on director performance, u_j, and their own preferences for directors with particular characteristics, v_j. In a world of perfect corporate governance, boards are only concerned with their mandate (i.e. representing shareholders' interests) and $v_j = 0$.

We set $v_j = 0$ not because we believe in a world of perfect governance but because our question is: can an algorithm identify a director x'' with better performance than director x' nominated by board *j*, whom the board will like at least equally well? In other words, conditional on $g_j(x'') \geq g_j(x')$, can an algorithm recommend a nominating rule α that produces a higher payoff than the baseline: the outcome of board *j*'s actual nominating decision?

The difference in the expected payoffs between the two nominating rules α_j and h_j is:

$$\Pi_j(\alpha_j) - \Pi_j(h_j) = \sum_{i \in \mathcal{D}} \alpha_{j,i} E\left[\pi_j(x_i, y_i)\right] - \sum_{i \in \mathcal{D}} h_{j,i} E\left[\pi_j(x_i, y_i)\right]$$

$$= \underbrace{E[y|\ \alpha]}_{missing\ label} - \underbrace{E[y|\ h]}_{observed\ label}$$

We do not observe the performance of directors who would be nominated under the alternative nominating rule produced by the algorithm. As discussed in Kleinberg et al. (2017), missing labels are often dealt with in the machine learning literature by various imputation procedures. However, this approach would assume that if a director shares the same set of observable feature values, *w*, as the nominated director, their performance would be identical. This is the equivalent of assuming that unobservables, *z*, play no role in nominating decisions. For a given *w*, the imputation error would therefore be:

$$E[y|\alpha, w] - E[y|h, w] = E[w + z|\alpha, w] - E[w + z|h, w]$$

$$= E[w|\alpha, w] - E[w|h, w] + E[z|\alpha, w] - E[z|h, w]$$

$$= E[z|\alpha, w] - E[z|h, w]$$

This imputation error points up the *selective labels problem*. In our setting, it refers to the possibility that directors who were nominated, although they might share the same exact observable features as other directors not nominated, might differ in terms of unobservables. These unobservables could lead to different average outcomes for nominated vs. not nominated, even if both are identical on the basis of observable characteristics.

We exploit the design of our pool of candidate directors for each board seat in order to compare the performance of our algorithm to board decisions. Although we do not have labels for nominees generated by the algorithm's nominating rule, $E[y|\alpha]$, we observe their *quasi-label*: their performance on the smaller neighboring board they joined around the same time.

We are interested in evaluating the quality of boards' nominating decisions. Our approach is to contrast those decisions to an alternative nominating rule that our algorithm would have chosen. For example, using the notation introduced in this section, if the algorithm predicted a director with characteristics x' would perform very poorly and there were fifty other candidates the algorithm predicted would do better, there are effectively fifty alternative nominating rules α that would yield a higher payoff in terms of benefits derived from director performance. To allow boards to use unobservables to make their nominating decisions, we add the assumption that among those fifty alternative nominees, there exists at least one director with characteristics x'' such that $g_j(x'') \geq g_j(x')$. When we analyze the quasi-labels of those potential candidates, we explore whether they indeed do better on average than director x' when x' was predicted to do poorly, and worse when x' was predicted to do well.

IA3. DATA DEFINITIONS

IA3.1. Individual Director Features

Source: BoardEx except if stated otherwise
(as of when the director joins the board)

Variable	**Definition**
Age	Director age
Audit chair	Equals one if director is chair of the audit committee
Audit member	Equals one if director is a member of the audit committee
Avgtimeothco	The average time that a director sits on the board of quoted companies
Bkgd academic	Equals one if job history includes in title one of the following: "professor" "academic" "lecturer" "teacher" "instructor" "faculty" "fellow" "dean" "teaching"
Bkgd finance	Equals one if job history includes in title one of the following: "underwriter" "investment" "broker" "banker" "banking" "economist" "finance" "treasure" "audit" "cfo" "financial" "controller" "accounting" "accountant" "actuary" "floor trader" "equity" "general partner" "market maker" "hedge fund"
Bkgd hr	Equals one if job history includes in title one of the following: "hr " "recruitment" "human resource"
Bkgd law	Equals one if job history includes in title one of the following: "lawyer" "legal" "attorney" "judge" "judicial"
Bkgd manager	Equals one if job history includes in title one of the following: "manager" "vp" "president" "director" "administrator" "administrative" "executive" "coo" "chief operating" "operation" "secretary" "founder" "clerk" "division md" "employee" "associate" "head of division"
Bkgd marketing	Equals one if job history includes in title one of the following: "marketing" "publisher" "mktg" "sales" "brand manager" "regional manager" "communication" "merchandising" "comms" "distribution" "media"
Bkgd military	Equals one if job history includes in title one of the following: "captain" "soldier" "lieutenant" "admiral" "military" "commanding" "commander" "commandant" "infantry" "veteran" "sergeant" "army"
Bkgd politician	Equals one if job history includes in title one of the following: "politician" "senator" "political" "deputy" "governor"
Bkgd science	Equals one if job history includes in title one of the following: "researcher" "medical" "doctor" "scientist" "physician" "engineer" "biologist" "geologist" "physicist" "metallurgist" "science" "scientific" "pharmacist"
Bkgd technology	Equals one if job history includes in title one of the following: "technology" "software" "programmer" " it " "chief information officer" "database" "system administrator" "developer"
Bonus	Annual bonus payments (in thousands)
Busy	Equals one if directors sits on three or more boards

Chairman	Equals one if director is chairman of the board
Compensation chair	Equals one if director is chair of the compensation committee
Compensation committee	Equals one if director is a member of the compensation committee
Experience CEO	Equals one if director has experience as CEO of a publicly traded company
Experience CFO	Equals one if director has experience as CFO of a publicly traded company
Experience Chairman	Equals one if director has experience as Chairman of a publicly traded company
Experience exec VP	Equals one if director has experience as executive VP of a publicly traded company
Experience President	Equals one if director has experience as President of a publicly traded company
Experience entrepreneur	Equals one if director has experience as an entrepreneur
Experience government & policy	Equals one if director has government and policy experience
Experience risk management	Equals one if director has risk management experience
Experience strategic planning	Equals one if director has strategic planning experience
Experience sustainability	Equals one if director has experience in sustainability
Experience org type 1	Equals one if director has experience working in Armed Forces
Experience org type 2	Equals one if director has experience working in Charities
Experience org type 3	Equals one if director has experience working in Clubs
Experience org type 4	Equals one if director has experience working in Government
Experience org type 5	Equals one if director has experience working in Medical
Experience org type 6	Equals one if director has experience working in a Partnership
Experience org type 7	Equals one if director has experience working in the private sector
Experience org type 8	Equals one if director has experience working in a quoted company
Experience org type 9	Equals one if director has experience working in Sports
Experience org type 10	Equals one if director has experience working in Universities
Foreign	Equals one if director's nationality is not American
GenBBB	Equals one if director was born between 1946 and 1964
GenDepBB	Equals one if director was born in or before 1926
Gender	Equals one if director is male
GenMature	Equals one if director was born between 1927 and 1945
GenX	Equals one if director was born between 1965 and 1980
GenY	Equals one if director was born in 1981 or after
Governance chair	Equals one if director is chair of the governance committee
Governance member	Equals one if director is a member of the governance committee
HistInternational	Equals one if job history includes a position outside the United States
HistInternational_Africa	Equals one if job history includes a position in Africa
HistInternational Asia	Equals one if job history includes a position in Asia

HistInternational Canada	Equals one if job history includes a position in Canada
HistInternational Caribbean	Equals one if job history includes a position in the Caribbean
HistInternational Europe	Equals one if job history includes a position in Europe
HistInternational Middle East	Equals one if job history includes a position in the Middle East
HistInternational South America	Equals one if job history includes a position in South America
Ivy league	Equals one if director went to an Ivy League college
Job accountant	Equals one if director is an accountant
Lead_independent	Equals one if director is lead independent director
MBA	Equals one if director holds an MBA degree
Mean past voting outcome	Average shareholder support during the first three years of tenure for previous board positions (*Source: ISS Voting Analytics*)
Mean_support_3yrs	Average shareholder support over the first three years of tenure. Source: ISS Voting Analytics
Network size	Network size of director (number of overlaps through employment, other activities, and education)
Nomination chair	Equals one if director is chair of the nomination committee
Nomination member	Equals one if director is a member of the nomination committee
Number connections	Number of established connections to incumbent board members prior to joining the board
Number qualifications	Number of qualifications at undergraduate level and above
Nb current seats diff ind	Number of current board seats in different FF48 industry
Nb current seats same ind	Number of current board seats in same FF48 industry
Nb prev seats diff ind	Number of previous board seats in different FF48 industry
Nb prev seats same ind	Number of previous board seats in same FF48 industry
Nb prev jobs industry	Number of previous jobs in same FF48 industry
Time prev jobs industry	Time spent on jobs in same FF48 industry
Nb prev jobs different industry	Number of previous jobs in different FF48 industry
Time prev jobs different industry	Time spent on jobs in different FF48 industry
Other chair	Equals one if director is chair of a committee other than compensation, audit, governance or nomination
Other member	Equals one if director is a member of a committee other than compensation, audit, governance or nomination
Other compensation	Value of annual *ad hoc* cash payments (in thousands)
Perf to total compensation	Performance to total - Ratio of Value of LTIPs Held to Total Compensation
Salary	Base annual pay in cash (in thousands)
Timeretirement	Time to retirement (assumed to be 70 years old)
Time previous seats	Time spent on previous board seats
Time prev seats diff ind	Time spent on previous board seats in a different industry
Time prev seats same ind	Time spent on previous board seats in same industry
Tot Current Nb Listed Boards sitting on	The number of Boards of publicly listed companies that an individual serves on

Tot Current Nb Other Boards sitting on	The number of Boards for organizations other than publicly listed or private companies that an individual serves on
Tot Current Nb Unlisted Boards sitting on	The number of Boards of private companies that an individual serves on
Tot Nb Listed Boards sat on	The number of Boards of publicly listed companies that an individual has served on
Tot Nb Other Boards sat on	The number of Boards for organizations other than publicly listed or private companies that an individual has served on
Tot Nb unlisted Boards sat on	The number of Boards of private companies that an individual has served on
Total Compensation	Salary + Bonus
Total director compensation	Salary plus Bonus plus Other Compensation plus Employers Defined Retirement/Pension Contribution
Total equity linked wealth	Valuation of total wealth at the end of the period for the individual based on the closing stock price of the last annual report
Value of shares held	Value of shares held at the end of the reporting period for the individual based on the closing stock price of the annual report

IA3.2. Board-level features

Source: BoardEx except if stated otherwise
(as of when the director joins the board)

Variable	**Definition**
Attrition rate	Number of Directors that have left a role as a Fraction of average number of Directors for the preceding reporting period
Average age	Average age of directors on the board
Average age less than 50	Equals one if average age of directors is less than 50
Average age more than 67	Equals one if average age of directors is more than 67
Average busy	Fraction of directors currently sitting on three or more boards
Average foreign	Fraction of directors with nationality other than American
Average independent	Fraction of non-executive directors on the board
Average Ivy League	Fraction of directors who went to an Ivy League college
Average MBA	Fraction of directors holding an MBA
Average nb qualifications	Average number of qualifications at undergraduate level and above of directors on the board
Average nb qualifications lt1	Fraction of directors whose number of qualifications at undergraduate level and above is less than one
Average nb qualifications mt3	Fraction of directors whose number of qualifications at undergraduate level and above is more than three
Average network size	Average network size of directors on the board (number of overlaps through employment, other activities, and education)

Average tenure	Average board tenure of directors on the board
Average time in company	Average time in company for executive and non-executive directors on the board
Average timebrd lt3	Fraction of directors whose number of years on the board is less than three
Average timebrd mt12	Fraction of directors whose number of years on the board is more than twelve
Avg tot current nb listed boards	The average number of boards of publicly listed companies directors currently serve on
Avg tot nb listed boards sat on	The average number of boards of publicly listed companies directors have served on
Avg totcurrnolstdbrd_less_1	Fraction of directors whose number of other listed boards sitting on is less than one
Avg totcurrnolstdbrd_more_3	Fraction of directors whose number of other listed boards sitting on is more than three
Avg totnolstdbrd_less_1	Fraction of directors whose number of other listed boards sat on is less than one
Avg totnolstdbrd_more_5	Fraction of directors whose number of other listed boards sat on is more than five
Avg Bkgd academic	Fraction of directors with an academic background (job history)
Avg Bkgd CEO	Fraction of directors with a CEO background (job history)
Avg Bkgd finance	Fraction of directors with a finance background (job history)
Avg Bkgd hr	Fraction of directors with a human resources background (job history)
Avg Bkgd law	Fraction of directors with a law background (job history)
Avg Bkgd manager	Fraction of directors with a manager background (job history)
Avg Bkgd marketing	Fraction of directors with a marketing background (job history)
Avg Bkgd military	Fraction of directors with a military background (job history)
Avg Bkgd politician	Fraction of directors with a political background (job history)
Avg Bkgd science	Fraction of directors with a scientific background (job history)
Avg Bkgd technology	Fraction of directors with a technology background (job history)
Avg Experience CEO	Fraction of directors with experience as CEO of a publicly traded company
Avg Experience CFO	Fraction of directors with experience as CFO of a publicly traded company
Avg Experience Chairman	Fraction of directors with experience as Chairman of a publicly traded company
Avg Experience exec VP	Fraction of directors with experience as executive VP of a publicly traded company
Avg Experience President	Fraction of directors with experience as President of a publicly traded company
Avg GenBBB	Fraction of directors born between 1946 and 1964
Avg GenDepBB	Fraction of directors born in or before 1926
Avg Mature	Fraction of directors born between 1927 and 1945
Avg GenX	Fraction of directors born between 1965 and 1980
Avg GenY	Fraction of directors born in 1981 or after
Avg HistInternational_Africa	Fraction of directors with experience in Africa

Avg HistInternational Asia	Fraction of directors with experience in Asia
Avg HistInternational Canada	Fraction of directors with experience in Canada
Avg HistInternational Caribbean	Fraction of directors with experience in the Caribbean
Avg HistInternational Europe	Fraction of directors with experience in Europe
Avg HistInternational Middle East	Fraction of directors with experience in the Middle East
Avg HistInternational S.America	Fraction of directors with experience in South America
Avg Experience org type 1	Fraction of directors with experience in Armed Forces
Avg Experience org type 2	Fraction of directors with experience in Charities
Avg Experience org type 3	Fraction of directors with experience in Clubs
Avg Experience org type 4	Fraction of directors with experience in Government
Avg Experience org type 5	Fraction of directors with experience in Medical
Avg Experience org type 6	Fraction of directors with experience in a Partnership
Avg Experience org type 7	Fraction of directors with experience in the private sector
Avg Experience org type 8	Fraction of directors with experience in a quoted company
Avg Experience org type 9	Fraction of directors with experience in Sports
Avg Experience org type 10	Fraction of directors with experience in Universities
Avg networksize lt 250	Fraction of directors whose network size is less than 250
Avg networksize mt 3000	Fraction of directors whose network size is more than 3000
Board Pay Slice - salary	Tot indep comp/ CEO salary
Board Pay Slice - total	Tot indep comp/ CEO total compensation
Board size	Number of directors on the board
BOSS	Equals one if the CEO is also the chairman of the board and the President
CEO bonus	CEO's bonus
CEO salary	CEO's salary
CEO total compensation	CEO total compensation (salary plus bonus)
Chairman duality	Equals one if the CEO is chairman of the board
Classified	Equals one if board is classified
Count Female	Number of women on the board
Entrenched	Equals one if CEO is chairman and has been in company more than five years
Fracdirafter	Coopted Directors as Fraction of Total Board (Data from Lalitha Naveen's website)
Fracdirafterindep	Coopted Independent Directors as Fraction of Total Board (Data from Lalitha Naveen's website)

Twfracdirafter	Tenure Weighted Coopted Directors as Fraction of Total Board (Data from Lalitha Naveen's website)
Twfracdirafterindep	Tenure-Weighted Coopted Independent Directors as Fraction (Data from Lalitha Naveen's website)
Gender ratio	The Fraction of male directors
High female dummy	Equals one if board has three or more female directors
No female dummy	Equals one if board has zero female director
Nationality Mix	Fraction of Directors from different countries
Nb independent	Number of independent directors
Nb international experience	Number of directors with international experience
Stdev age	Standard deviation of directors' age
Stdev current listed board	Standard deviation of the number of listed boards each director currently serves on
Stdev listed board sat on	Standard deviation of the number of quoted boards sat on for all directors on the board
Stdev number qualifications	Standard deviation of the number of qualifications at undergraduate level and above for all directors on the board
Stdev Time in Company	Standard deviation of time in the company for all directors on the board
Stdev Time on Board	Standard deviation of time on board for all directors on the board
Succession Factor	Measurement of the Clustering of Directors around retirement age
Tot indep comp	Sum of all independent directors' total compensation
Tot indep comp scaled	Sum of all independent directors' total compensation divided by the number of independent directors

IA3.3 Firm level features

Source: Compustat /CRSP except if stated otherwise

(as of when the director joins the board)

Variable	Definition
Current assets	Current assets - Total
Asset growth	Past year total asset growth
Asset tangibility	Property plant and equipment over total assets
Acquisitions	Acquisitions
Auditor	Dichotomous variable for each auditing firm
Auop1-4	Dummy for auditor opinion
Averagewordsperparagraph	WRDS SEC Analytics Suite -Average number of words per paragraph 10K
BCW	Equals one if firm was on the Fortune-Best Company to work for list within 10 years preceding the nomination (from Alex Edmans' w
Blank check	Equals one if firm has a blank check provision (from ISS RiskMetrics)
CAPX	Capital expenditures
CEOSO1	Equals to one if the CEO is exempt from filing Certification Documents as required under section 302 of the Sarbanes-Oxley Act of 2
CFOSO1	Equals to one if the CFO is exempt from filing Certification Documents as required under section 302 of the Sarbanes-Oxley Act of 2

CEOSO2	Equals to one if the CEO has not filed Certification Documents as required under section 302 of the Sarbanes-Oxley Act of 2002
CFOSO2	Equals to one if the CFO has not filed Certification Documents as required under section 302 of the Sarbanes-Oxley Act of 2002
CEOSO3	Equals to one if the CEO has filed Certification Documents as required under section 302 of the Sarbanes-Oxley Act of 2002
CFOSO3	Equals to one if the CFO has filed Certification Documents as required under section 302 of the Sarbanes-Oxley Act of 2002
Cash	Cash
Cash flow	(ebitda-txt-xint)/at
Confidential Vote	Equals one for confidential vote (from ISS RiskMetrics)
Cumulative vote	Equals one for cumulative vote (from ISS RiskMetrics)
Div	Dividends
Dividend payer	Equals one if the total amount of dividends to ordinary equity > 0
Div ratio	Dividends/ebitda
Div stock repurchase	Dividends plus stock repurchase over total assets
Dual Class	Equals one for dual class stock (from ISS Riskmetrics)
LT debt	Long term debt - Total
ST debt	Short term debt
Depreciation	Depreciation and amortization -
Dividends	Total amount of dividends to ordinary equity
EBIT	Earnings Before Interest and Taxes
EBITDA	Earnings Before Interest
Fair price	Equals one if fair price provision (from ISS Riskmetrics)
Book debt	Fama French (2000) book debt
Book equity	Fama French (2000) book equity
Finterms_negative	Loughran-McDonald Negative word proportion (from Wrds SEC Analytics Suite)
Finterms_positive	Loughran-McDonald Positive word proportion (from Wrds SEC Analytics Suite)
Finterms_litigious	Loughran-McDonald litigious word proportion (from Wrds SEC Analytics Suite)
Finterms_uncertainty	Loughran-McDonald uncertainty word proportion (from Wrds SEC Analytics Suite)
Firm age	Time since IPO or first occurrence on CRSP
Firm age quartile	Quartile for firm age
Firm size quartile	Quartile for firm size (total assets)
Fsize	Size of annual report file (from Wrds SEC Analytics Suite)
Golden parachute	Equals one if firm has a golden parachute provision (from ISS RiskMetrics)
Gunnin_fox_index	Gunning Fog Readability Index (from Wrds SEC Analytics Suite)
Harvardiv_negative	Harvard General Inquirer negative word count (from Wrds SEC Analytics Suite)
HerfindahlIndex	Industry sales Herfindahl index
Incorp Delaware	Equals one if incorporated in Delaware
IPO year	Year of the IPO

K_int	Peter & Taylor Total Q-Firm's intangible capital estimated replacement cost (from Wrds)
K_int_know	Peter & Taylor Total Q-Firm's knowledge capital replacement cost (from Wrds)
K_int_offbs	Peter & Taylor Total Q-Portion of K_int that doesn't appear on firm's balance sheet (from Wrds)
K_int_org	Peter & Taylor Total Q-Firm's intangible capital estimated replacement cost (from Wrds)
Lerner Index	Industry median ebitda/revenues
limit_abil_amend_bylw	Limited ability to amend corporate bylaws (from ISS RiskMetrics)
limit_abil_amend_charter	Limited ability to amend charter (from ISS RiskMetrics)
limit_abil_written_consent	Limited ability to act by written consent (from ISS RiskMetrics)
Leverage	Total long term debt / total assets
Ln(nb insti blocks)	Logarithm of one plus the number of institutional blockholders.
Ln(nb insti owners)	Logarithm of one plus the number of institutional investors.
Majority vote standard	Equals one if requires a director to receive support from a majority of the shares cast to be elected. (from ISS RiskMetrics)
MB	(common shares outstanding * stock price)/ ordinary equity
Minority interest	Minority interest
Mkt value equity	Market value of equity (price times shares outstanding)
Net debt issue	Net debt issued (Baker and Wurgler, 2002)
Net equity issue	Net equity issued (Baker and Wurgler, 2002)
NumestYr norm	Average Annual Number of Analysts (From EPS estimates from IBES) divided by total assets
Plurality vote	Equals one if a director need only receive one vote to be elected. (from ISS RiskMetrics)
Product Mkt fluidity	Product market fluidity. Hoberg and Phillips
Profitability	EBITDA/total assets
Poison pill;	Poison pill (from ISS Riskmetrics)
Resignation req	Resignation requirement for failed election (from ISS RiskMetrics)
Q_tot	Peter & Taylor Total Q-Total q (from Wrds)
Block ownership %	Fraction owned by blockholders.
Institutional ownership %	Fraction owned by institutional investors.
Largest inst. shr. %	Fraction owned by largest institutional investor.
Largest 10 inst. shr. %	Fraction owned by top ten institutional investors.
Largest 5 shr. %	Fraction owned by top five institutional investors.
RD	Research and development
12-month return	Cumulative stock return in the twelve months leading up to the appointment.
3-month return	Cumulative stock return in the three months leading up to the appointment.
6-month return	Cumulative stock return in the six months leading up to the appointment.
Excess returns 12	Cumulative stock return in the twelve months leading up to the appointment net of market return
Excess returns 3	Cumulative stock return in the three months leading up to the appointment net of market return
Excess returns 6	Cumulative stock return in the six months leading up to the appointment net of market return

RIX	RIX Readability index (from Wrds SEC Analytics Suite)
ROA	Net income / total assets
Sales	Net sales - Total
SG&A	SG&A over total assets
SIlpctyr	Average Annual Short Interest as a % of Shares Outstanding
SI quintile	Quintiles of short interest
Total assets	total assets -
Working cap over assets	Working capital over total assets
Extraordinary items	extraordinary items
R&D	R&D expenses
E-Index	Index of six governance attributes: staggered boards, limits to shareholder bylaw amendments, poison pills, golden parachutes, and supermajority requirements for mergers and charter amendments (from ISS)

IA3.4. Industry and market level features

Source: Compustat /CRSP except if stated otherwise
(as of when the director joins the board)

Variable	Definition
Industry ROA	Return on assets of firms with same 3-digit SIC code
Market3	Cumulative returns on the S&P500 in the three months leading up to the appointment
Market6	Cumulative returns on the S&P500 in the six months leading up to the appointment
Market12	Cumulative returns on the S&P500 in the twelve months leading up to the appointment
ExcessReturns3	Cumulative stock return in the three months leading up to the appointment minus cumulative returns on the S&P500 in the three months leading up to the appointment
ExcessReturns6	Cumulative stock return in the six months leading up to the appointment minus cumulative returns on the S&P500 in the six months leading up to the appointment
ExcessReturns12	Cumulative stock return in the twelve months leading up to the appointment minus cumulative returns on the S&P500 in the twelve months leading up to the appointment
Industry leverage	Industry leverage

Industry cash flow vol	Industry cash flow volatility
Tnic3*	3-digit, text-based industry classifications from Hoberg and Phillips (2010, 2016)

VOTES DISTRIBUTION

Shareholder Support: Fraction of Votes "For"

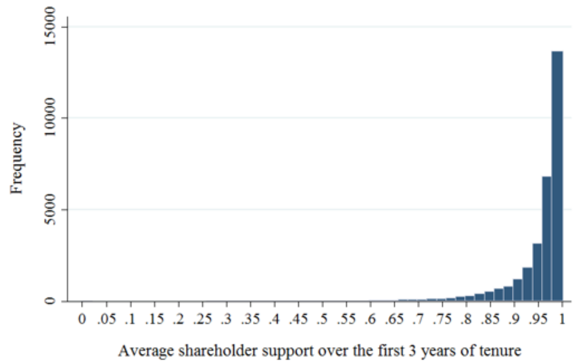

Figure IA1. This figure shows the distribution of average shareholder support, defined as the fraction of votes in favor of a given director over all votes cast for the director's reelection within three years of her tenure. The data is from ISS Voting Analytics.

Distribution of Poor Outcomes: Fraction of Votes "For" Below 95%

Figure IA2. This figure shows the distribution of average shareholder support for values under its mean value of 95%. Shareholder support is defined as the fraction of votes in favor of a given director over all votes cast for the director's reelection within three years of her tenure. The data is from ISS Voting Analytics.

Excess Votes: Fraction of Votes “For” Minus the Slate’s Average

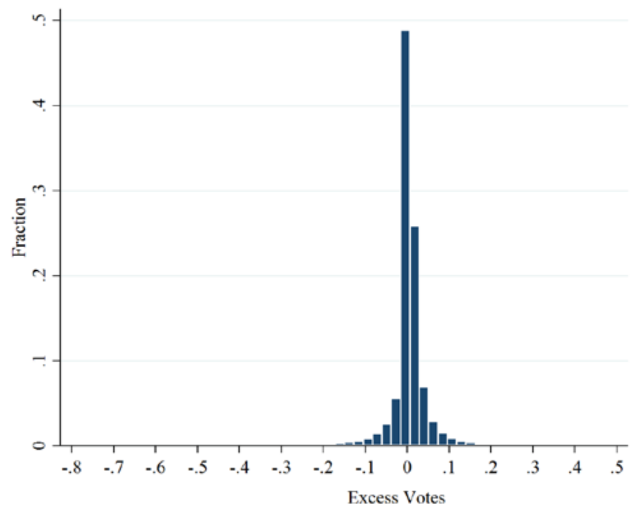

Figure IA3. This figure shows the distribution of *excess votes* for our sample. To compute *Excess votes*, we compute the fraction of votes in favor of a given director over all votes cast for the director. Next, we subtract the average of that variable for the slate of directors up for reelection that year on the focal board. Finally, we take the average of this relative vote measure over the first three years of the new director’s tenure. The data is from ISS Voting Analytics.

Distribution of Poor Outcomes: Excess Votes below -5%

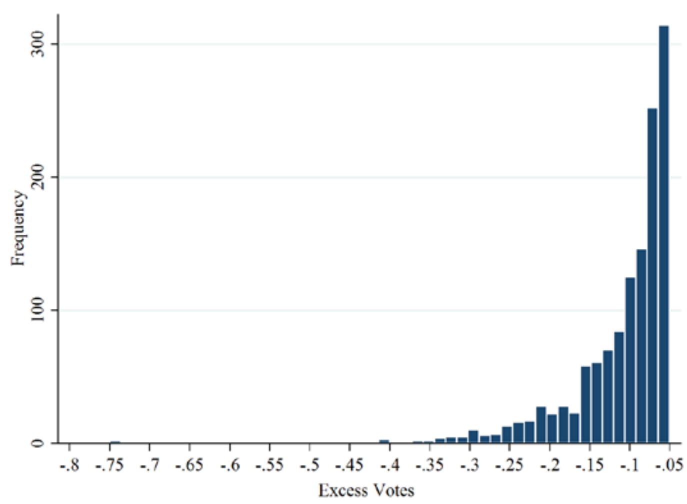

Figure IA4. This figure shows the distribution of *Excess votes* below -5%.

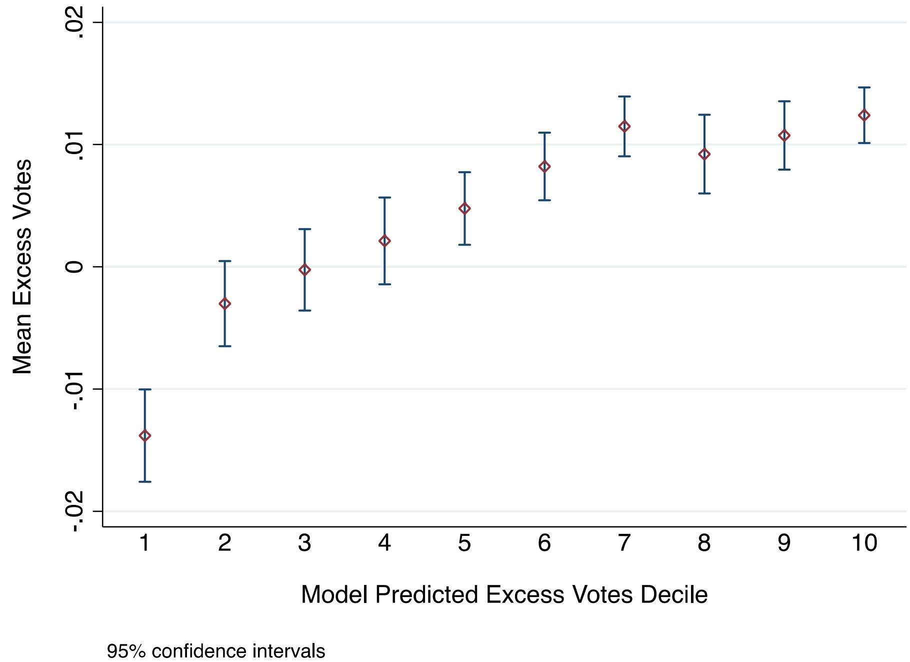

FIGURE IA5: MEAN OBSERVED *EXCESS VOTES* VS. *XGBOOST*-PREDICTED *EXCESS VOTES WHEN TEST SET INCLUDES APPOINTMENTS BETWEEN 2011-2014.*

This figure shows the average observed level of excess shareholder support for directors across the ten deciles of predicted performance for the *XGBoost* model when the test set includes appointments between 2011 and 2014. To compute *excess votes*, we first compute the fraction of votes in favor of a given director over all votes cast for the director. Next, we subtract the average of that variable for the slate of directors up for reelection that year on the focal board. Finally, we take the average of this relative vote measure over the first three years of the new director's tenure.

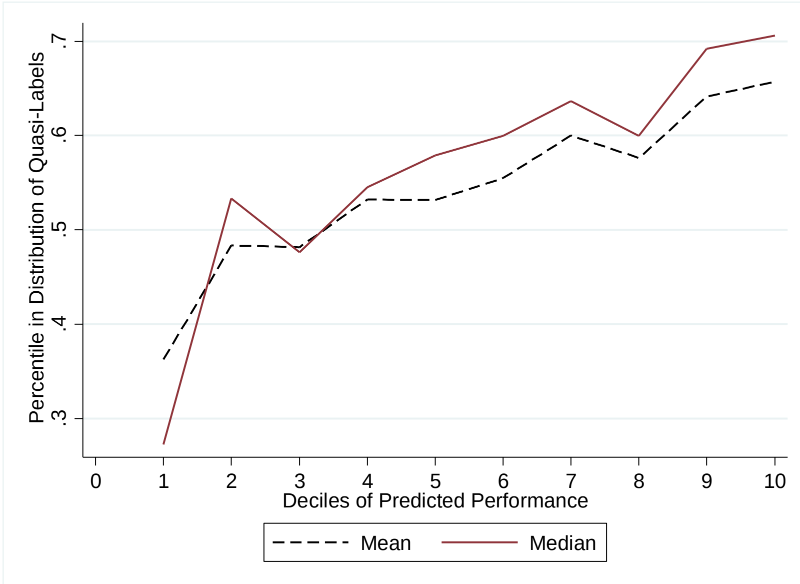

FIGURE IA6: MEAN AND MEDIAN RANK IN QUASI-LABEL DISTRIBUTION ACROSS DECILES OF *Lasso*-PREDICTED PERFORMANCE

This figure shows the mean and median rank in the distribution of quasi-labels for directors in each of the ten deciles of ***Lasso***-predicted performance (*Excess votes*). The observed performance of nominated directors in our test set is compared to the quasi-labels of *all* potential candidates in their respective candidate pool: Each new board appointment in the test set is associated with a candidate pool, comprised of directors who, within one year of the appointment, joined the board of a smaller neighboring firm.

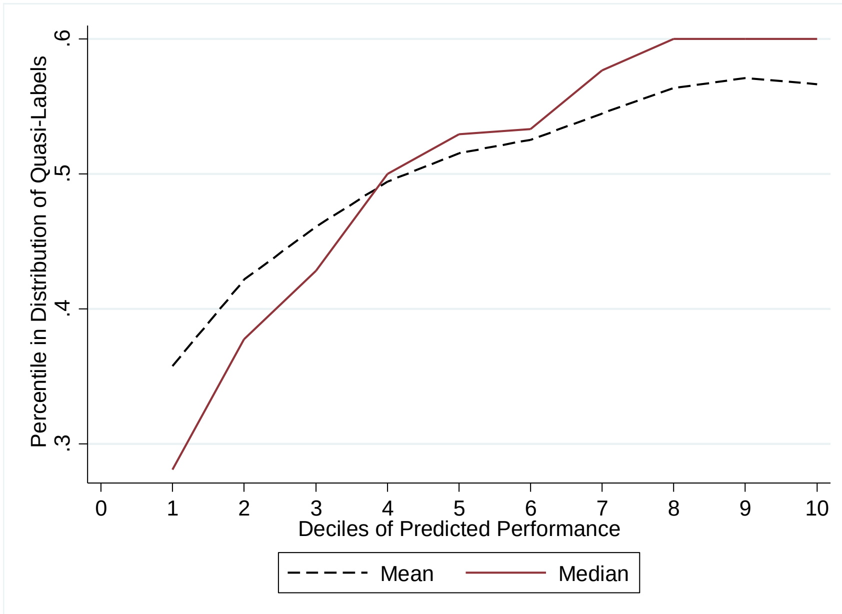

FIGURE IA7: MEAN AND MEDIAN RANK IN QUASI-LABEL DISTRIBUTION ACROSS DECILES OF *XGBOOST*-PREDICTED PERFORMANCE (FIRM SIZE ASSUMPTION RELAXED)

This figure shows the mean and median rank in the distribution of quasi-labels for directors in each of the ten deciles of *XGBoost*-predicted performance (*Excess votes*). The observed performance of nominated directors in our test set is compared to the quasi-labels of *all* potential candidates in their respective candidate pool: Each new board appointment in the test set is associated with a candidate pool, comprised of directors who, within one year of the appointment, joined the board of a neighboring firm. In this graph, we **relax the restriction** that candidates joined a ***smaller*** firm.

Dependent variable: excess votes	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Compensation chair	-0.007	-0.007	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	0.001	0.001
	(-1.307)	(-1.264)	(-0.475)	(-0.607)	(-0.338)	(-0.377)	(-0.153)	(-0.271)	(0.344)	(0.279)
Audit chair	-0.006	-0.006	0.005***	0.005***	0.001	0.001	0.003	0.003	0.005***	0.005***
	(-1.089)	(-1.245)	(2.779)	(2.742)	(0.376)	(0.318)	(0.811)	(0.793)	(2.716)	(2.865)
Gender	-0.002	-0.002	-0.001	-0.001	0.002	0.002	0.002	0.002	0.000	0.000
	(-0.843)	(-0.847)	(-0.631)	(-0.562)	(0.968)	(1.014)	(0.978)	(1.163)	(0.007)	(0.019)
Number current other boards	-0.002**	-0.002**			-0.004***	-0.004***	-0.003***	-0.003***		
	(-2.328)	(-2.362)			(-6.440)	(-6.763)	(-4.829)	(-5.315)		
Director age > 65	-0.003	-0.002			0.001	0.001	-0.000	0.001	-0.003**	-0.003**
	(-0.778)	(-0.525)			(0.503)	(0.608)	(-0.031)	(0.251)	(-2.063)	(-2.155)
Fraction owned by director	-9.945	-15.763								
	(-0.308)	(-0.496)								
Governance chair	0.005	0.005	0.002	0.002	0.006	0.003	0.002	-0.001	0.003	0.003
	(0.392)	(0.416)	(1.009)	(0.824)	(0.618)	(0.344)	(0.242)	(-0.115)	(1.329)	(1.270)
E-index	0.000	-0.000			0.001	0.001*	0.000	0.000		
	(0.062)	(-0.055)			(1.492)	(1.672)	(0.510)	(0.793)		
Governance chair * E-index	-0.003	-0.003			-0.002	-0.001	-0.001	-0.000		
	(-0.635)	(-0.750)			(-0.696)	(-0.378)	(-0.420)	(-0.017)		
Classified board	0.002	0.000					0.002	0.001	0.001	0.001
	(0.719)	(0.132)					(1.119)	(0.681)	(1.409)	(0.819)
ln(number institutional owners)	0.003*	0.005**					0.002	0.001		
	(1.722)	(2.533)					(0.693)	(0.543)		
Industry adj. EBITDA	-0.026**	-0.008					-0.016*	-0.006	0.002	0.003
	(-2.097)	(-0.991)					(-1.831)	(-0.970)	(0.683)	(1.068)
Industry adj. 12-months returns	0.002	0.001					0.000	-0.000	0.001	0.001
	(0.667)	(0.496)					(0.154)	(-0.053)	(0.837)	(0.910)
Nomination chair			0.002	0.001	-0.002	-0.004	-0.005	-0.008	-0.007	-0.008
			(0.353)	(0.216)	(-0.219)	(-0.373)	(-0.530)	(-0.813)	(-1.454)	(-1.583)
Industry experience			-0.000	-0.000	-0.003	-0.003				
			(-0.317)	(-0.291)	(-1.097)	(-1.004)				
Background finance			0.002*	0.002	0.005**	0.004*				
			(1.725)	(1.386)	(2.102)	(1.781)				
Background law			-0.004**	-0.004**	-0.010***	-0.009***				
			(-2.062)	(-2.112)	(-3.054)	(-2.872)				
MBA			0.001	0.001	0.000	-0.000				
			(0.810)	(0.976)	(0.037)	(-0.005)				
Ivyplus			0.000	0.000	-0.000	-0.000				
			(0.044)	(0.056)	(-0.044)	(-0.095)				
Director age			-0.000	-0.000						
			(-0.647)	(-0.499)						
Number of qualifications			0.000	0.000	-0.000	-0.000	0.001	0.001	0.000	0.000
			(0.324)	(0.041)	(-0.440)	(-0.605)	(0.899)	(0.836)	(0.249)	(0.219)
ln(assets)			0.001**	0.001***	0.001	0.001	0.001	0.001	0.001***	0.001***
			(2.545)	(2.649)	(0.866)	(1.195)	(0.986)	(1.146)	(2.839)	(2.660)
Leverage			-0.005**	-0.004**	-0.001	-0.002	0.001	-0.003	-0.006**	-0.005**
			(-2.264)	(-2.024)	(-0.172)	(-0.294)	(0.221)	(-0.544)	(-2.342)	(-2.351)
MB			-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
			(-0.132)	(-0.336)	(-0.303)	(-0.384)	(-0.139)	(-0.208)	(-0.361)	(-0.500)
Largest 5 institutional owners %			0.005	0.009**	0.017*	0.022**	0.011	0.016	0.006	0.011***
			(1.383)	(2.490)	(1.747)	(2.449)	(0.787)	(1.167)	(1.497)	(2.791)
ROA			0.002	0.002	0.011	0.009				
			(1.022)	(0.716)	(1.420)	(1.298)				
Product market fluidity			-0.000	-0.000**						
			(-0.614)	(-2.030)						
12-months returns			0.000	-0.000	-0.000	-0.000				
			(0.040)	(-0.021)	(-0.135)	(-0.099)				
Dividend payer			0.001	0.001	-0.001	-0.001				
			(0.848)	(0.504)	(-0.375)	(-0.523)				
Board size			-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
			(-1.081)	(-1.560)	(-0.130)	(-0.391)	(-0.227)	(-0.533)	(-0.138)	(-0.866)
Fraction female on board			-0.000	0.000	0.000	0.000				
			(-0.135)	(0.315)	(0.300)	(0.228)				
Fraction independent directors			-0.006*	-0.006	-0.008	-0.007	-0.013	-0.012	-0.005	-0.004
			(-1.701)	(-1.512)	(-0.955)	(-0.853)	(-1.501)	(-1.415)	(-1.175)	(-1.029)
Average director age			-0.000	0.000	-0.000	0.000				
			(-0.204)	(0.759)	(-0.331)	(0.340)				
Chairman CEO duality					0.002	0.001	0.001	0.000	0.001	0.001
					(1.185)	(0.787)	(0.821)	(0.211)	(0.864)	(0.575)
Fraction co-opted directors							-0.004	-0.004		
							(-1.463)	(-1.588)		
Busy									-0.006***	-0.006***
									(-5.552)	(-5.573)
Constant	0.039**	0.011***	-0.001	0.005	0.010	0.005	0.034*	0.013	0.012	0.005
	(2.184)	(2.846)	(-0.102)	(0.778)	(0.389)	(0.329)	(1.748)	(1.603)	(1.077)	(1.256)
Observations	1,345	1,345	10,601	10,601	3,136	3,136	3,040	3,040	8,773	8,773
R-squared	0.051	0.015	0.012	0.005	0.037	0.022	0.034	0.015	0.015	0.007
Calendar year FE	yes	no	yes	no	yes	no	yes	no	yes	no
Industry FE	yes	no	yes	no	yes	no	yes	no	yes	no
AIC	-4940	-4993	-36989	**-37023**	-10956	-11015	-10778	-10826	-31045	-31092

Table IA1. OLS MODEL

This table reports coefficients from various OLS models of excess votes on various director, firm, and board characteristics. Excess vote is defined as the average observed level of shareholder support over the first three years of a new director's tenure, minus the average vote for all directors in the same slate. The regression sample contains director appointments between 2000-2011. The Akaike Information Criterion (AIC), reported in the last row, is each estimation's out-of-sample prediction error. It allows us to compare the relative quality of the OLS models presented. In the paper, we use the Model (4) as it gives us the lowest AIC error.

	Mean		
	Bottom decile of predicted performance	**Top** decile of predicted performance	Difference p-value
Director level			
Age	56.3	57.0	0.083
Audit committee	0.236	0.818	0.000
Audit committee chair	0.039	0.077	0.001
Background academic	0.060	0.049	0.330
Background finance	0.190	0.122	0.000
Background lawyer	0.026	0.017	0.233
Background manager	0.335	0.318	0.471
Background marketing	0.084	0.026	0.000
Background military	0.010	0.006	0.405
Background politician	0.029	0.011	0.008
Background science	0.040	0.011	0.000
Background technology	0.021	0.007	0.021
Busy	0.520	0.120	0.000
Chairman	0.098	0.001	0.000
Compensation committee	0.624	0.059	0.000
Compensation committee chair	0.175	0.024	0.000
Foreign	0.156	0.088	0.005
Governance chair	0.045	0.011	0.000
Governance committee	0.168	0.122	0.008
International work experience	0.109	0.037	0.000
Male	0.897	0.746	0.000
Network size	1540	1327	0.000
Nomination chair	0.004	0.001	0.318
Nomination committee	0.023	0.011	0.057
Number of qualifications	2.208	2.282	0.180
Total current number of boards sitting on	2.848	1.545	0.000
Total number of listed boards sat on	5.814	2.289	0.000
Ivy league	0.217	0.109	0.000
MBA	0.466	0.410	0.064
Nb previous jobs same FF48 industry	0.105	0.037	0.000
Nb previous directorships same FF48 industry	0.342	0.037	0.000
Board level			
Gender ratio	0.105	0.153	0.000
Nationality mix	0.128	0.084	0.000
Board attrition	0.102	0.054	0.000
Average tenure of incumbent directors	3.443	9.731	0.000
Average tot. nb of boards incumbent directors sit on	1.672	1.809	0.000
Board size	8.5	10.2	0.000
CEO SOX certified	0.539	0.995	0.000
Chairman is CEO	0.357	0.496	0.001
Chairman is CEO with tenure ≥ 5	0.600	0.983	0.000
Indep. directors compensation over CEO tot. compensation	0.912	1.172	0.280
Mean past voting shareholder support	-0.012	0.011	0.000
Number of female directors	1.007	1.611	0.000
Incumbent directors with finance background	0.117	0.221	0.000
Busy incumbent directors	0.173	0.210	0.000
Average age of incumbent directors	57.5	63.0	0.000
Average network size of incumbent directors	1239	1347	0.007

Firm level			
Dividend payer	0.298	0.630	0.000
Excess returns 12 months leading up to appointment	0.028	-0.018	0.126
Firm age	10	30	0.000
Hoberg-Phillips product market fluidity	7.446	6.237	0.000
Institutional ownership %	0.586	0.711	0.000
Largest 10 institutional shareholders %	0.367	0.421	0.000
Largest 5 institutional shareholders %	0.275	0.303	0.001
Largest institutional shareholder %	0.106	0.102	0.492
Leverage	0.266	0.191	0.000
Log (number of institutional blockholders)	1.010	1.250	0.000
Log (number of institutional owners)	4.971	5.279	0.000
Ownership by blockholders %	0.193	0.226	0.002
ROE	-0.110	0.194	0.353
Stock returns prior 12 months	0.158	0.116	0.188
Total assets	17600	30435	0.087
Number of analysts	8.4	12.0	0.000
Short interest (%)	0.036	0.053	0.000
Peter & Taylor Total Q	4.291	0.990	0.000

TABLE IA2: TOP VS. BOTTOM DECILE OF PREDICTED PERFORMANCE

This table reports the mean of firm and director level features for directors in the bottom decile of predicted excess votes and compares it to the mean for directors in the top decile of predicted excess votes. These results are for directors in our test set. Because we do not need the actual vote outcomes for this exercise but only the predictions, this test set covers appointments up to 2016. The algorithm used to predict performance is *XGBoost*.

Dependent variable: predicted performance	(1)	(2)	(3)	(4)
Busy	-0.006***	-0.005***	-0.005***	-0.005***
	(-24.332)	(-13.183)	(-12.230)	(-12.087)
Male	-0.001***	-0.001	-0.001	-0.001*
	(-4.623)	(-1.398)	(-1.603)	(-1.688)
Age		-0.000**	-0.000**	-0.000**
		(-2.001)	(-2.079)	(-2.242)
MBA		0.000	0.000	0.000
		(1.074)	(1.108)	(1.150)
Ivy league		-0.001**	-0.001*	-0.001*
		(-2.555)	(-1.869)	(-1.864)
Background lawyer			-0.002	-0.001
			(-1.521)	(-1.394)
Background academic			0.000	0.000
			(0.180)	(0.134)
Background finance			-0.001	-0.001
			(-1.344)	(-1.568)
Network size			-0.000***	-0.000***
			(-2.838)	(-2.691)
Ln (Assets)	0.001***	0.000	0.000	0.000
	(9.054)	(0.276)	(0.016)	(0.160)
ROA	0.001***	0.000	0.000	0.000
	(4.280)	(0.024)	(0.156)	(0.124)
Board size	-0.000***	-0.000**	-0.000*	-0.000*
	(-3.303)	(-2.027)	(-1.751)	(-1.843)
Average nb independent directors	0.009***	0.006***	0.006***	0.005**
	(20.814)	(3.055)	(2.906)	(2.529)
Chairman duality		0.001***	0.001***	0.001***
		(3.399)	(3.481)	(3.319)
Excess returns 12 months leading up to appointment		0.000	0.000	0.000
		(1.049)	(1.058)	(1.156)
Number of female directors			0.000	0.000
			(0.172)	(0.554)
Average tenure of incumbent directors		0.000***	0.000***	0.000***
		(6.441)	(6.099)	(4.907)
Log (number of institutional owners)		0.001**	0.001**	0.001*
		(2.155)	(2.431)	(1.798)
Compensation committee chair				-0.002*
				(-1.803)
Audit committee chair				0.002**
				(2.524)
Governance committee chair				-0.002
				(-1.457)
Nomination committee chair				-0.002
				(-0.599)
Firm Age				0.000***
				(2.925)
Constant	-0.002***	0.003	0.003	0.005*
	(-3.762)	(1.053)	(1.128)	(1.867)
Observations	7,738	1,893	1,883	1,883
R-squared	0.153	0.131	0.136	0.146

t-statistics in parentheses

*** p<0.01, ** p<0.05, * p<0.1

TABLE IA3: THE DETERMINANTS OF PREDICTIONS: OLS REGRESSIONS

This table reports the results from OLS regression models of the predicted excess votes in our test set on some firm level and director level features. Because we do not need the actual vote outcomes for this exercise but only the predictions, this test set covers appointments up to 2016. The algorithm used to generate the predictions is *XGBoost*.

about ECGI

The European Corporate Governance Institute has been established to improve *corporate governance through fostering independent scientific research and related activities.*

The ECGI will produce and disseminate high quality research while remaining close to the concerns and interests of corporate, financial and public policy makers. It will draw on the expertise of scholars from numerous countries and bring together a critical mass of expertise and interest to bear on this important subject.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.

Electronic Access to the Working Paper Series

The full set of ECGI working papers can be accessed through the Institute's Web-site (www.ecgi.global/content/working-papers) or SSRN:

Finance Paper Series	http://www.ssrn.com/link/ECGI-Fin.html
Law Paper Series	http://www.ssrn.com/link/ECGI-Law.html