Selecting Directors Using Machine Learning

ISIL EREL

LÉA STERN

CHENHAO TAN

MICHAEL S. WEISBACH*

December 13, 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. Our results suggest that 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. 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 about the way that potential nominees will 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. These algorithms 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. They are particularly useful in applications when there is no clear guidance for how to build expectations. Because there is no one size fits all "good" governance, many covariates potentially matter and interact in non-linear ways when predicting director performance, which is the reason we wish to rely on a rigorous data-driven model selection that produces more accurate predictions.[3]

---

[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. 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 and Cai, Nguyen, and Walkling (2017), who document that more complex firms are more likely to appoint directors who are connected to the CEO or the existing board.

[3] See Athey and Imbens (2017) and Mullainathan and Spiess (2017).

We construct a large database of publicly traded U.S. firms and independent directors appointed between 2000 and 2014. We employ 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. Voting based measures are individualized, market-based measures of performance that capture investors' preferences. Therefore, we task the algorithm with predicting the average excess voting support relative to the slate of directors up for reelection over the first three years of director tenure.

We employ several machine learning algorithms to predict the performance of a potential director at a 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. 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 have worse performance out of sample than those it predicted would do well.

We ensure that the out of sample predictive accuracy of the algorithm is not dependent on one particular measure of director performance by considering two alternative outcome measures in addition to excess votes. The first is a dichotomous variable that captures low absolute shareholder support – i.e., strong dissent against a director. The second is based on the idea that a director who leaves shortly after being appointed likely reflects a poor director-firm match (Ferreira et al., 2017, Bates et al., 2016). We therefore use a dichotomous variable for turnover within two years of appointment as an alternative performance measure.[4] In unreported tests, we also use total instead of excess votes to measure director performance and draw the same conclusions. Finally, we consider the model's ability to predict announcement returns of director appointments. We find that the algorithm's predictions of shareholder votes to reelect directors are also strongly related to announcement returns around director appointments.

We observe director performance for directors who were 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). 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. Although we do not observe the outcome (the *label*) of those potential candidates at the focal firm (this is the essence of the selective labels problem), the design of our candidate

[4] Our results are unchanged if we use turnover within three years instead of two.

pools allows us to observe what we refer to as their *“quasi-label”*: their performance on the board they effectively joined. We use the distribution of quasi-labels to impute the distribution of labels we would have observed for passed-over directors at the focal firm, and assess whether the algorithm’s predictions are informative about how a director will eventually perform relative to alternative candidates.

We find that directors the algorithm predicted would perform poorly (well) do perform poorly (well) when compared to potential available alternatives as well. The average director in the bottom decile of predicted performance ranks at the 23rd percentile in the distribution of quasi-labels. In contrast, those in the top decile rank at the 80th percentile. 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 the dependent variable. However, in recent years, there have been rapid developments in the growing strand of the machine learning literature referred to as “*Explainable AI” (XAI) or “Interpretable ML”*, which focuses on improving model interpretability (e.g., Lundberg and Lee, 2017 and Ribeiro, Singh and Guestrin, 2016 as well as Vilone and Lungo, 2020, for a review). We employ methods from this literature 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 tend to have worse governance structures. This pattern is consistent with the view that firms choosing predictably bad directors is a manifestation of underlying agency conflicts within a boardroom.

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. Corporate finance applications are developing.[5] 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 paper is the first, to our knowledge, to apply supervised machine learning to improve our understanding of director selection and, hence, corporate governance. We emphasize strongly, however, that algorithms hold promise to *complement* (rather than substitute) human judgement.

2. Constructing a Sample on which Algorithms Can Select Directors

2.1. Measuring Director Quality

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

[5] 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).

board that year.[6] Our results are qualitatively unchanged if we task the algorithms with predicting the total 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 affect 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 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 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,

---

[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).

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 ownership (> 26%) 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 recommendations 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.

Shareholders do on occasion oppose newly nominated directors even though shareholder support in uncontested elections is typically very high. Therefore, we also train the algorithm to predict strong dissent, measured by an indicator variable equal to one if a director receives low (less than 90%) support, as an alternative measure of director performance using shareholder votes.

Finally, director turnover is another measure of director-firm match quality used in the literature (see e.g., Ferreira, Ginglinger, Laguna, and Skalli, 2017). Therefore, we also use whether a new director leaves within two (or three) 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. Internet Appendix IA1 provides detailed variable definitions.

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 in 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 reelection 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. As previously documented in the literature on uncontested director elections, the overall level of shareholder support is typically very high. Given that 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. Therefore, 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.

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

---

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

(serving on three or more boards) experience low shareholder support more frequently than non-busy directors.

Although some variables appear to affect director performance, 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 busy female 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 one reason why machine learning algorithms are successful at prediction. They are designed for 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 and allow for rigorous, data-driven model selection (Athey, 2017).

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 (see Internet Appendix IA1) 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).* 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.[11]

[11] The algorithms rely on a regularizer that balances out in-sample fit and out-of-sample overfitting.

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 these out-of-sample predictions to those from an AIC-selected OLS model (see Appendix C). 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, although 70-30% or 80-20% splits are typically used in practice.[12] We use an 80-20% split but our results do not depend on the way in which we split the data into training and test periods.[13] Because of the possibility that the Sarbanes-Oxley Act (SOX) affected the process by which independent directors are chosen, we have also used a training sample starting in 2003 (keeping the test sample between 2012-2014). We report results using this training sample in Figure IA2 in the Internet Appendix, which are very similar to those discussed below.

3.2. Predictions of Director Performance

A way to evaluate the quality of a model predicting performance is to evaluate the way in which actual performance increases related to 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* and *Lasso* perform best at predicting the subsequent success of directors using excess votes.[14] Directors predicted to be in the bottom percentile as predicted by

---

[12] See Hastie et al. (2009) for a discussion of methodological issues involved in choosing training and testing sets.
[13] Our results are similar when we include appointments in 2011-2014 in the test set (see, for example, Internet Appendix Figure IA1)
[14] *XGBoost* can be easily utilized on a standard laptop using the *Xgboost* package for Python, available at https://pypi.org/project/xgboost/

XGBoost and *Lasso* have an average *observed* excess shareholder support of -3.1% and -2.6%, respectively, whereas the average observed excess support is 1.2% and 1.8% 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 predicted support 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.

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 C, 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.[15] For each model, we present specifications with *excess votes* as the dependent variable (with and without year and industry fixed effects) as well as with *total votes* (including firm-year fixed effects, see Gormley and Matsa, 2014). To compare the out-of-sample predictable power of these specifications, we calculate the Akaike Information Criterion (AIC) for each model. 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 (5) of the Appendix C Table, which has the lowest AIC error.[16]

[15] Cai, Garner and Walkling (2009) examine the determinants of shareholder votes at annual director elections using OLS. The R^2 of these models drops significantly when explaining votes for new directors only and when removing variables not available to board members at the time of hiring (e.g. ISS recommendation).

[16] In an earlier version of this paper, we ensure that our results are not driven by poorly performing firms. Excluding firms that had negative abnormal returns in the year prior to the nomination give similar results to those reported here. See Erel et al. (2018).

4. Comparing Directors who were Appointed to Potential Alternative Choices

Our results so far document that directors identified by our algorithm as likely to have low (high) future shareholder support, are in fact more likely to have low (high) support in subsequent elections. The key issue raised by this finding is that when firms hire directors that turn out to be poor, it is possible, at least some of the time, to predict this poor performance before the choice was made. In other words, firms regularly hire *predictably poor* directors.

One possible explanation for this practice is that there are simply no alternatives available to the firms, and although it is predictable that these directors will do poorly, others whom the firm could have hired would have been even worse. This possibility illustrates a form of selection bias referred to in the machine learning literature as the "selective labels problem" (Kleinberg et al. 2017). This problem refers to the fact that only the firms' choices for directors that actually occurred are observed, so it is impossible to know how potential alternative directors that they did not select would have performed. This possibility limits our ability to assess how algorithmic decision aids could improve on boards' decisions. Put differently, while showing that the algorithm predicts well out of sample is important, it is not sufficient to assess whether algorithmic predictions could actually improve nominating decisions.

Our setting allows us, however, to evaluate the importance of this issue. We consider a sample of individuals who would have taken the directorship with high probability. We classify these alternative candidates as those directors who, within one year of the appointment, joined the board of a smaller neighboring firm.[17] 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.[18]

[17] A neighboring firm is defined as a firm whose headquarters is within 100 miles of the focal firm's headquarters.
[18] We redo our analyses relaxing the assumption that candidates join the board of a smaller firm and find that the results remain robust. We report the results using this larger pool of candidates in the Internet Appendix Figure IA3. Our results also hold when we restrict pools to candidates who joined the board of a company in the same industry as the focal firm.

While we do not observe the performance of these potential candidates at the focal firm, we do observe their performance on the board they did join. 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"*. Quasi-labels allow us to impute the performance distribution for alternative candidates on the focal board. Importantly, this distribution is independent from the performance predictions $\hat{y}$. We estimate the average (median) rank of the focal director's performance in the distribution of quasi-labels for each decile of predicted performance $\hat{y}$, and examine whether it is correlated with $\hat{y}$. This procedure allows us to evaluate whether algorithmic predictions could help improve boards' nominating decisions even in the presence of selective labels.

Consider directors with low $\hat{y}$. Intuitively, if their observed performance were to lie in the right tail of quasi-labels, this would imply that they ended up doing well relative to available alternatives, despite the fact that our algorithm had predicted they would do poorly. The focal board could have selected on unobservables, and the high rank in the distribution of quasi-labels would suggest that the unobservables were used as *signal*. On the other hand, if their observed performance were to lie in the left tail of quasi-labels, this would imply that our algorithm identified *ex ante* that these directors would perform poorly, and relative to alternatives, they 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.

We evaluate this procedure in two ways. In the first, we use the algorithms to generate performance predictions for potential candidates on the focal boards and use these predictions to narrow down the candidate pools. For focal directors predicted to do poorly (well), we narrow down the candidate pool to those the algorithm predicted were most (least) promising. 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 poorly performed noticeably worse than available promising candidates, while nominated directors predicted to do well performed better than available unpromising alternative candidates. *XGBoost* and *Lasso* again appear to be the preferred algorithms because they can best discriminate the directors who will do well from those who will not. In contrast, the predictions from the OLS model are uninformative about subsequent performance relative to available alternative candidates.

As a second test of this procedure, we compare the performance of focal directors to the performance of *all* potential candidates and present this comparison in Figure 2. The mean and median percentile in the distribution of quasi-labels increase across deciles of *predicted* performance, using the predictions from the *XGBoost* algorithm. This relation suggests that as the algorithm's prediction of a director's performance increases, her effective rank when compared to plausible candidates is higher. The difference between decile 1 and 10 is statistically significant at the 1% level for both the mean and the median. The results using *Lasso* are similar and presented in Appendix A.[19]

The selective labels problem is a central challenge to evaluating predictive models and their ability to improve on human decisions (Kleinberg et al. 2017). In the context of the application we are considering, it simply states that one cannot evaluate the quality of a firm's choices without knowing about possible alternative choices. The quasi-labels approach provides such a potential alternative. We show that firms that hire predictably bad directors could have hired other directors who would have done a better job.[20]

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) argue 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.

[19] In the rest of the paper, we will report results using *XGBoost* in the main text and the ones using *Lasso* in Appendix A to save space. *XGBoost and Lasso* perform similarly well in most cases; however, *XGBoost* provides better predictions using alternative definitions of director performance.

[20] We present a formalization of the quasi-labels approach to evaluate algorithmic predictions in Appendix B.

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.[21] 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), who tend to rely on ISS recommendations to a smaller extent.

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.[22]

Table 5 reports the mean CARs using a (-1; +1) window around announcements. 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, and also with larger event windows, with similar results.[23]

5.2. 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

---

[21] 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.

[22] We collect announcement dates from *BoardEx*, *CapitalIQ* and *Lexis-Nexis*.

[23] An alternative approach to looking at the stock price reaction to the announcement of the director's appointment is to examine whether it is possible to evaluate whether the predicted quality of a director is associated with changes in the firm's profitability. An earlier draft of this paper performed this analysis and concludes that excess votes are closely related to profitability, so that if the model is trained on excess shareholder votes, it also predicts profitability and vice versa. See Erel et al. (2018).

votes in their reelection (within the next three years) from shareholders, and use it as an alternative measure 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 3, 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.

5.3. Predicting Director Turnover

Director turnover shortly after being appointed is likely to be associated with a poor fit between the director and the firm, suggesting that the decision to hire him was not a good one. For this reason, we train the *XGBoost* algorithm to predict whether a director will leave within two years following her appointment. We choose the two-year span to ensure that the turnover happens before reelection; however, the results are similar if we use turnover over a three-year period.

Figure 4 documents the way in which the *XGBoost* algorithm's predictions compare with 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 average fraction of directors who departed within two years in the test sample and the algorithm's predictions, as is documented 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, the machine learning algorithm 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. With the goal of understanding what leads models to make specific predictions, the machine learning literature has focused on developing methods to improve model interpretability and to make black-box models more transparent (e.g., Lundberg and Lee, 2017 and Ribeiro, Singh and Guestrin, 2016).

In this section, we employ one such method, SHAP (Shapley Additive exPlanations), to understand which variables are most important in our model's assessment of the quality of a potential director. SHAP values were introduced by Lundberg and Lee (2017) as a game theoretic approach to improve the interpretability of a complex model's output. A feature x's SHAP value estimates the extent to which it contributes to pushing the model's output away from the base value (the unconditional expectation). It is defined as the change in the expected model output, averaged across all possible orderings of the other (non- x) features. Each feature's incremental contribution to the model's output is estimated by considering the output of all possible combinations of features. This analysis provides a way to identify the variables that are especially important when training the algorithm.[24]

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 examine the importance of characteristics that our algorithm uses to predict a given director's performance. We use SHAP values to quantify the marginal contribution of each feature to predicting director performance. A features' SHAP values can be computed for each observation, i.e. each individual prediction, to improve the model's transparency (*local interpretability*). This approach is helpful as it shows which variables were instrumental in generating a specific prediction. Additionally, SHAP values for each feature can be averaged across

[24] For this analysis, we use the SHAP library for Python available at https://github.com/slundberg/shap.

observations. This analysis produces a ranking of variables to understand which contribute the most to the model's output (average across observations) - i.e. improve the *global interpretability* of the model.

6.1.1 Global interpretability

Figure 5 presents the ten features that contribute the most to *XGBoost's* predictions of director performance, measured by SHAP values averaged across observations. Panel A uses excess votes as the performance measure, with the variables presented in red 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 increases the predicted performance of the incoming directors while the total number of boards she sits on currently and sat on in the past decrease her predicted performance.

In Panel B, we calculate SHAP values when the model predicts whether a director receives less than 90% of the votes. This panel indicates that the attributes with the largest SHAP values when predicting excess votes largely overlap with those with the largest SHAP values when the model predicts dissent.

In Panel C, we present SHAP values for models with director turnover as the performance measure. Using this measure, mostly individual characteristics of these directors, appear to explain turnover. Having a classified board, which gives directors longer term, has a negative effect on board turnover. In addition, three other important variables predicting turnover within two years of joining a board are being the chairman, having an entrepreneurial background and having graduated from an Ivy League university.

We emphasize that SHAP values do not establish causality. Instead, they quantify features' contribution to the model based on correlations. However, SHAP values do provide a helpful basis for expanding our qualitative understanding of how machine learning models generate predictions.

6.1.2 Local interpretability

It is also possible to compute features' SHAP values for individual observations. In Figure 6, we report the features 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) increase (decrease) the model's outcome, i.e. the individual prediction $\hat{y}$. The value of the attribute (x) 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). For each observation, the difference between the model's output and the base value (mean prediction) is the sum of all features' SHAP values.

For example, in Panel A, where the model predicts excess votes, the most important variable affecting the model's prediction is that the 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 model's performance prediction upwards. The standard deviation of incumbent directors' time on the board is six years (i.e. two years above the sample average). This increased 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 of these attributes increased $\hat{y}$. Interestingly, some attributes sometimes push $\hat{y}$ up and sometimes down (e.g. classified board). This pattern 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, or a value of an attribute that is "too high" for one firm or one director, could be irrelevant or even have a positive effect for another. This supports our rationale for using machine learning in the context of predicting director quality.

6.2. Overvalued Director Characteristics

In this section we consider the characteristics of *predictably bad directors*, i.e. directors who the model predicted that they would do poorly but were chosen anyway, and subsequently did poorly. These characteristics are likely to be 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.[25] The patterns of discrepancies between these two groups reflect the types of directors that tend to be overvalued in the nomination process.

In Table 6, we report characteristics of these predictably bad directors. 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.

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 tend to be chosen as directors in other firms. 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?

The pattern that some firms regularly choose directors who can be predicted to do a poor job suggests that corporate governance is problematic at these firms. Our model provides a way to address this issue directly, by comparing a firm's susceptibility to hiring predictably bad directors with other measures of corporate governance. To measure the effect of corporate governance on director choices, Table 7 provides estimates of probit equations that estimate the likelihood that a firm hires a predictably bad director.

[25] 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.

As a measure of governance quality, we first use Bebchuk et al.'s (2009) Entrenchment Index (E-index). The index is based on six governance attributes -staggered boards, limits to shareholder bylaw amendments, poison pills, golden parachutes, supermajority requirements for mergers and for charter amendments- and is constructed so that it increases as the firm-level governance gets worse. Note that this E-index is constructed by evaluating the relative importance of the Investor Responsibility Research Center (IRRC)'s 24 provisions, which Gompers, Ishii and Metrick (2003) had included in their G-index.

The estimates in Table 7 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. Moreover, confirming this interpretation of the results, we find 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).

A second measure of governance quality is 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, so co-option likely reflects a board that does not act in the interest of shareholders. 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 2002, requiring exchange-listed firms to have a majority of independent directors. Consistent with the agency view, the estimates in Table 7 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 7 suggest that agency conflicts distort nominating decisions as firms that nominate directors who were predictable poor choices have worse governance structures.

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 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 particular, firms that select predictably bad directors appear to have lower quality corporate governance.

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 use alternative measures of director performance: the firm's abnormal returns at the time of the announcement of a director's appointment, high dissent, and turnover shortly after the appointment.

Our algorithm can predict the distribution of outcomes using these different 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) that machine learning is a valuable approach for prediction problems in the social sciences.

An issue we need to address before we can conclude that algorithms can help us understand and potentially improve the director nomination process is that we observe the predictive accuracy of our algorithm only for directors who were nominated. We design a procedure to address this issues that exploits the fraction of votes plausible candidates received at the company whose board they joined as an indication of their performance. Our results suggest 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 identify features that are overrated in the director nomination process. Comparing predictably bad directors to promising candidates suggested by the algorithm, it appears that predictably bad directors 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 well-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]

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 (and more diverse) 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 addition, if firms progressively start reshuffling their boards based on algorithmic decision aids, behaviors and decisions will start to change. This would likely have general equilibrium effects which would need to be evaluated.

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

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

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 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.

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

Berle A. and G. Means (1932) The Modern Corporation and Private Property. *Commerce Clearing House, New York.*

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.

Erel, I., L. Stern, C. Tan and M.S. Weisbach (2018) Selecting Directors Using Machine Learning, NBER Working Paper 24435.

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.

Gompers, P., J. Ishii, and A. Metrick (2003) Corporate Governance and Equity Prices. Quarterly Journal of Economics, 118, 107-155.

Gormley, T. and D. Matsa (2014) Common Errors: How to (and Not to) Control for Unobserved Heterogeneity, *The Review of Financial Studies* 27: 617–661.

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.*
Vilone, G. and L. Lungo (2020) Explainable Artificial Intelligence: a Systematic Review, working paper, arXiv:2006.00093
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.

Appendix A: *LASSO* RESULTS

This Appendix presents figures analogous to those in the paper using the *Lasso* algorithm rather than the *XGBoost* algorithm.

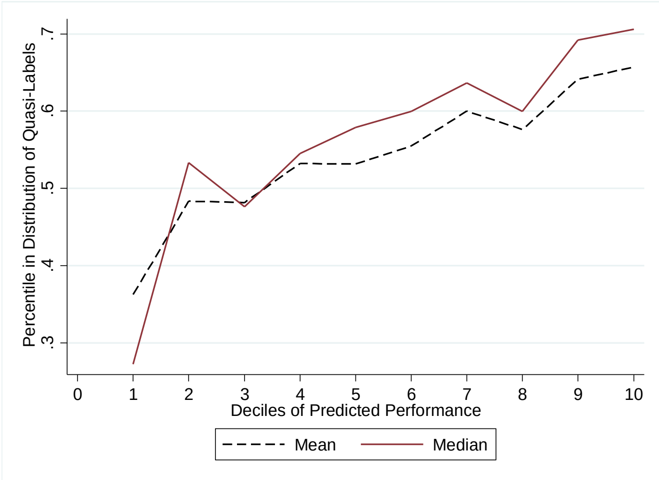

FIGURE A1: MEAN AND MEDIAN RANK IN QUASI-LABEL DISTRIBUTION ACROSS DECILES OF *Lasso*-PREDICTED PERFORMANCE

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.

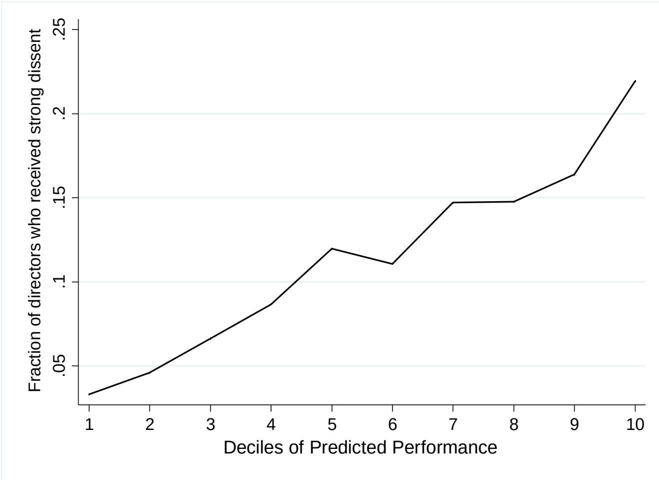

FIGURE A2: MEAN OBSERVED *DISSENT* VS. LASSO-PREDICTED *DISSENT*

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 *Lasso*-predicted dissent in the 2012-14 test set. Directors in decile one (ten) were predicted by *Lasso* as least (most) likely to face strong decent. Unconditional mean dissent is about 11%.

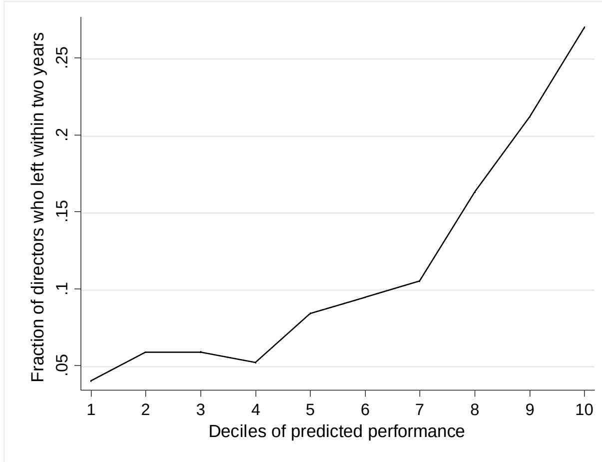

FIGURE A3: MEAN OBSERVED *DIRECTOR TURNOVER* VS. LASSO-PREDICTED *TURNOVER*

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 *lasso*-predicted turnover in the 2012-14 test set. Unconditional mean turnover is about 11%.

Characteristics	*Coefficient*
Total number unlisted boards sat on	-0.0181
Incumbent directors standard deviation time on board	0.0061
Compensation committee	-0.0051
Fraction owned by institutions	0.0045
Audit committee	0.0045
Busy	-0.0041
Fraction incumbent directors with law background	0.0040
Incumbent directors standard dev. time in company	0.0026
Generation X (born 1965-1980)	-0.0024
CEO has filed certification documents as required under section 302 of the Sarbanes-Oxley Act of 2002	0.0024

TABLE A1: LIST OF TOP TEN LASSO-SELECTED FEATURES (EXCESS VOTES)

This table reports the ten features with the largest absolute coefficients as estimated by Lasso in predictions of excess votes. Note however that these coefficients cannot be interpreted in the same way as OLS coefficients. See Mullainathan and Spiess (2017) for limitations in interpreting Lasso results.

Characteristics	*Coefficient*
Fraction incumbent directors with experience as Partner	0.1116
Succession Factor	0.0701
Net debt issues	-0.0687
Other Chair	-0.0675
Newly retained earnings	-0.0655
Compensation chair	-0.0639
Governance chair	-0.0608
Nomination chair	-0.0591
President	-0.0577
Audit chair	-0.0540

TABLE A2: LIST OF TOP TEN LASSO-SELECTED FEATURES (TURNOVER)

This table reports the ten features with the largest absolute coefficients as estimated by Lasso in predictions of turnover within two years. Note however that these coefficients cannot be interpreted in the same way as OLS coefficients. See Mullainathan and Spiess (2017) for limitations in interpreting Lasso results.

Characteristics	*Coefficient*
Fraction incumbent directors with experience in University	-0.1450
Fraction owned by top five institutional investors	0.1034
Auditor Moss Adams	-0.0920
CEO has not filed Certification Documents as required under section 302 of SOX	0.0896
Incumbent directors average number of directorships in same industry	0.0846
Auditor Grant Thornton	0.0777
Incumbent directors average number of jobs in a quoted company	0.0741
Newly retained earnings	-0.0721
Incumbent directors average number of jobs in same industry	0.0714
Chairman	0.0659

TABLE A3: LIST OF TOP TEN LASSO-SELECTED FEATURES (DISSENT)

This table reports the ten features with the largest absolute coefficients as estimated by Lasso in predictions of dissent (shareholder support below 90%). Note however that these coefficients cannot be interpreted in the same way as OLS coefficients. See Mullainathan and Spiess (2017) for limitations in interpreting Lasso results.

Appendix B: A FRAMEWORK TO EVALUATE ALGORITHMIC PREDICTIONS

We develop a framework using the template laid out in 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, respectively. $\mathcal{Z}$ 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 i in candidate pool $\mathcal{D}$, the board's payoff is characterized as:

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

$g_j(x)$ 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.

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

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

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, $E[y|\ \alpha]$. 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. Therefore, it could be the case that even for directors (x, y) who were predictably poor directors, there were no other available director (x', y') such that $\pi_j(x', y') > \pi_j(x, y)$. In other words, there could be supply considerations which make it difficult to evaluate whether algorithmic predictions can help boards improve nominating decisions.

We overcome this challenge by constructing a plausible pool $\mathcal{D}$ of candidate directors *-i* for each focal director (x_i, y_i). Our goal is to compare their performance to the performance of the focal director. We do not observe the outcome $E[y|\alpha]$ for alternative candidates. We do observe, however, their performance on the board of the (smaller) neighboring firm they joined around the same time, i.e. their *quasi-label* y_{-i}. Let $P(y_{-i} \leq y_i)$ be the director (x_i, y_i)'s rank in the distribution of y_{-i}. We sort all (x, y) in our test set into deciles d_n according to their algorithm-predicted performance $\widehat{y_i}$ and estimate the mean $P(y_{-i} \leq y_i)$ for directors in each decile: $\frac{1}{|d_n|}\sum_{(x,y)\in d_n} P(y_{-i} \leq y_i)\ \ \forall\, n = (1, \ldots, 10)$. To test whether the algorithm's predictions of director performance are useful even when we consider alternative candidates, we test the null that directors' rank when compared to alternatives is the same for directors predicted to do well as for those predicted to do poorly.

$$H_0:\ \ \frac{1}{|d_1|}\sum_{(x_i,y_i)\in d_1} P(y_{-i} \leq y_i) = \frac{1}{|d_{10}|}\sum_{(x_i,y_i)\in d_{10}} P(y_{-i} \leq y_i)$$

The alternative is that directors predicted to do poorly do worse than those predicted to do well when compared to alternative candidates.

$$H_1:\ \ \frac{1}{|d_1|}\sum_{(x_i,y_i)\in d_1} P(y_{-i} \leq y_i) < \frac{1}{|d_{10}|}\sum_{(x_i,y_i)\in d_{10}} P(y_{-i} \leq y_i)$$

A rejection of the null would suggest that the algorithm's predictions were informative about future performance, even when considering plausible candidates.

It is important to note that quasi-labels are not perfect substitutes for missing labels. We need two conditions to use quasi-labels to impute the distribution of missing labels. The first one relates to how quasi-labels are collected. Quasi-labels must be independent from the predictions: $\boldsymbol{y_{-i}} \perp \widehat{y_i}$. In our case, the set of firms that hired directors in the candidate pools is independent from the algorithms' predictions of director performance. They are firms that hired a director around the same time and whose headquarters is nearby, regardless of whether the algorithm predicted the focal director would do well or whether she would do poorly. Some of these hires by nearby firms were likely good choices and some likely were not. But importantly, the way candidate directors (and therefore $\boldsymbol{y_{-i}}$) are considered is independent of $\widehat{y_i}$. The second one is the assumption that the difference between the quasi-label and the missing label is not systematically negatively correlated with predicted director performance, i.e. $(E[y|\alpha] - y_{-i}) \perp \widehat{y_i}$.[28]

With these conditions, we 1) use the distribution of quasi-labels to impute the distribution of labels we would have observed for passed-over directors at the focal firm and 2) evaluate whether algorithmic predictions can help improve board decisions by assessing whether they are informative about how a director will eventually perform relative to alternative candidates.

[28] 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 driven, at least in part, by this positive (negative) difference between the true labels and the quasi-labels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Dependent variable	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes
Compensation chair	-0.007	-0.007	0.001	-0.001	-0.001	-0.005	-0.001	-0.001	-0.006	-0.001	-0.001	-0.003	0.001	0.001	-0.002
	(-1.307)	(-1.264)	(0.086)	(-0.475)	(-0.607)	(-1.349)	(-0.338)	(-0.377)	(-0.780)	(-0.153)	(-0.271)	(-0.399)	(0.344)	(0.279)	(-0.535)
Audit chair	-0.006	-0.006	-0.005	0.005***	0.005***	0.009***	0.001	0.001	0.005	0.003	0.003	0.004	0.005***	0.005***	0.009**
	(-1.089)	(-1.245)	(-0.480)	(2.779)	(2.742)	(2.731)	(0.376)	(0.318)	(0.681)	(0.811)	(0.793)	(0.571)	(2.716)	(2.865)	(2.422)
Gender	-0.002	-0.002	-0.000	-0.001	-0.001	-0.001	0.002	0.002	-0.001	0.002	0.002	-0.000	0.000	0.000	-0.004
	(-0.843)	(-0.847)	(-0.007)	(-0.631)	(-0.562)	(-0.216)	(0.968)	(1.014)	(-0.204)	(0.978)	(1.163)	(-0.120)	(0.007)	(0.019)	(-1.525)
Number current other boards	-0.002**	-0.002**	-0.003				-0.004***	-0.004***	-0.003***	-0.003***	-0.003***	-0.003***			
	(-2.328)	(-2.362)	(-1.382)				(-6.440)	(-6.763)	(-3.719)	(-4.829)	(-5.315)	(-3.145)			
Director age > 65	-0.003	-0.002	0.001				0.001	0.001	-0.001	-0.000	0.001	0.001	-0.003**	-0.003**	-0.001
	(-0.778)	(-0.525)	(0.159)				(0.503)	(0.608)	(-0.284)	(-0.031)	(0.251)	(0.113)	(-2.063)	(-2.155)	(-0.197)
Fraction owned by director	-9.945	-15.763	-5.054												
	(-0.308)	(-0.496)	(-0.099)												
Governance chair	0.005	0.005	0.011	0.002	0.002	-0.003	0.006	0.003	0.003	0.002	-0.001	-0.008	0.003	0.003	-0.001
	(0.392)	(0.416)	(0.424)	(1.009)	(0.824)	(-0.830)	(0.618)	(0.344)	(0.159)	(0.242)	(-0.115)	(-0.410)	(1.329)	(1.270)	(-0.327)
E-index	0.000	-0.000	-0.006				0.001	0.001*	0.004	0.000	0.000	0.005*			
	(0.062)	(-0.055)	(-1.407)				(1.492)	(1.672)	(1.434)	(0.510)	(0.793)	(1.683)			
Governance chair * E-index	-0.003	-0.003	-0.002				-0.002	-0.001	-0.004	-0.001	-0.000	-0.001			
	(-0.635)	(-0.750)	(-0.208)				(-0.696)	(-0.378)	(-0.790)	(-0.420)	(-0.017)	(-0.084)			
Classified board	0.002	0.000	-0.004							0.002	0.001	-0.007	0.001	0.001	-0.003
	(0.719)	(0.132)	(-0.572)							(1.119)	(0.681)	(-1.635)	(1.409)	(0.819)	(-1.065)
ln(number institutional owners)	0.003*	0.005**	0.006							0.002	0.001	0.007			
	(1.722)	(2.533)	(0.720)							(0.693)	(0.543)	(1.040)			
Industry adj. EBITDA	-0.026**	-0.008	0.010							-0.016*	-0.006	0.052*	0.002	0.003	-0.017
	(-2.097)	(-0.991)	(0.221)							(-1.831)	(-0.970)	(1.886)	(0.683)	(1.068)	(-1.290)
Industry adj. 12-months returns	0.002	0.001	-0.002							0.000	-0.000	0.001	0.001	0.001	-0.002
	(0.667)	(0.496)	(-0.269)							(0.154)	(-0.053)	(0.245)	(0.837)	(0.910)	(-1.257)
Nomination chair				0.002	0.001	-0.009	-0.002	-0.004	-0.040*	-0.005	-0.008	-0.041**	-0.007	-0.008	-0.017*
				(0.353)	(0.216)	(-1.002)	(-0.219)	(-0.373)	(-1.816)	(-0.530)	(-0.813)	(-2.019)	(-1.454)	(-1.583)	(-1.770)
Industry experience				-0.000	-0.000	-0.002	-0.003	-0.003	-0.004						
				(-0.317)	(-0.291)	(-0.679)	(-1.097)	(-1.004)	(-0.803)						
Background finance				0.002*	0.002	0.003	0.005**	0.004*	-0.000						
				(1.725)	(1.386)	(1.175)	(2.102)	(1.781)	(-0.051)						
Background law				-0.004**	-0.004**	-0.009**	-0.010***	-0.009***	-0.014**						
				(-2.062)	(-2.112)	(-2.401)	(-3.054)	(-2.872)	(-2.446)						
MBA				0.001	0.001	-0.000	0.000	-0.000	0.001						
				(0.810)	(0.976)	(-0.067)	(0.037)	(-0.005)	(0.343)						
Ivyplus				0.000	0.000	0.001	-0.000	-0.000	-0.003						
				(0.044)	(0.056)	(0.574)	(-0.044)	(-0.095)	(-0.724)						
Director age				-0.000	-0.000	-0.000									
				(-0.647)	(-0.499)	(-0.854)									

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Dependent variable	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes	Excess votes	Excess votes	Total votes
Number of qualifications				0.000	0.000	0.002***	-0.000	-0.000	0.001	0.001	0.001	0.002	0.000	0.000	0.001*
				(0.324)	(0.041)	(2.816)	(-0.440)	(-0.605)	(0.852)	(0.899)	(0.836)	(1.192)	(0.249)	(0.219)	(1.837)
ln(assets)				0.001**	0.001***	-0.011**	0.001	0.001	-0.004	0.001	0.001	0.001	0.001***	0.001***	-0.013**
				(2.545)	(2.649)	(-2.476)	(0.866)	(1.195)	(-0.328)	(0.986)	(1.146)	(0.077)	(2.839)	(2.660)	(-2.479)
Leverage				-0.005**	-0.004**	0.025**	-0.001	-0.002	0.020	0.001	-0.003	0.047	-0.006**	-0.005**	0.003
				(-2.264)	(-2.024)	(2.024)	(-0.172)	(-0.294)	(0.633)	(0.221)	(-0.544)	(1.511)	(-2.342)	(-2.351)	(0.214)
MB				-0.000	-0.000	0.000	-0.000	-0.000	-0.000	-0.000	-0.000	0.000	-0.000	-0.000	0.000
				(-0.132)	(-0.336)	(1.062)	(-0.303)	(-0.384)	(-0.317)	(-0.139)	(-0.208)	(0.031)	(-0.361)	(-0.500)	(0.536)
Largest 5 institutional owners %				0.005	0.009**	-0.049***	0.017*	0.022**	-0.075**	0.011	0.016	-0.110**	0.006	0.011***	-0.029
				(1.383)	(2.490)	(-2.819)	(1.747)	(2.449)	(-2.201)	(0.787)	(1.167)	(-2.378)	(1.497)	(2.791)	(-1.513)
ROA				0.002	0.002	-0.001	0.011	0.009	-0.023						
				(1.022)	(0.716)	(-0.220)	(1.420)	(1.298)	(-1.445)						
Product market fluidity				-0.000	-0.000**	0.000									
				(-0.614)	(-2.030)	(0.563)									
12-months returns				0.000	-0.000	-0.002	-0.000	-0.000	0.005						
				(0.040)	(-0.021)	(-1.381)	(-0.135)	(-0.099)	(1.442)						
Dividend payer				0.001	0.001	0.005	-0.001	-0.001	0.014						
				(0.848)	(0.504)	(0.828)	(-0.375)	(-0.523)	(1.194)						
Board size				-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.001	-0.000	-0.000	0.000
				(-1.081)	(-1.560)	(-0.489)	(-0.130)	(-0.391)	(-0.286)	(-0.227)	(-0.533)	(-0.417)	(-0.138)	(-0.866)	(0.129)
Fraction female on board				-0.000	0.000	0.003	0.000	0.000	-0.002						
				(-0.135)	(0.315)	(1.375)	(0.300)	(0.228)	(-0.652)						
Fraction independent directors				-0.006*	-0.006	0.010	-0.008	-0.007	0.055	-0.013	-0.012	0.029	-0.005	-0.004	0.013
				(-1.701)	(-1.512)	(0.549)	(-0.955)	(-0.853)	(1.393)	(-1.501)	(-1.415)	(0.695)	(-1.175)	(-1.029)	(0.608)
Average director age				-0.000	0.000	0.001	-0.000	0.000	-0.002						
				(-0.204)	(0.759)	(1.110)	(-0.331)	(0.340)	(-1.625)						
Chairman CEO duality							0.002	0.001	0.011	0.001	0.000	-0.004	0.001	0.001	-0.004
							(1.185)	(0.787)	(1.362)	(0.821)	(0.211)	(-0.458)	(0.864)	(0.575)	(-0.690)
Fraction co-opted directors										-0.004	-0.004	-0.002			
										(-1.463)	(-1.588)	(-0.174)			
Busy													-0.006***	-0.006***	-0.008***
													(-5.552)	(-5.573)	(-3.656)
Constant	0.039**	0.011***	0.975***	-0.001	0.005	0.987***	0.010	0.005	1.068***	0.034*	0.013	0.935***	0.012	0.005	1.044***
	(2.184)	(2.846)	(59.837)	(-0.102)	(0.778)	(19.192)	(0.389)	(0.329)	(8.806)	(1.748)	(1.603)	(8.982)	(1.077)	(1.256)	(24.446)
Observations	1,345	1,345	1,255	10,601	10,601	11,092	3,136	3,136	3,064	3,040	3,040	2,823	8,773	8,773	8,634
R-squared	0.051	0.015	0.632	0.012	0.005	0.611	0.037	0.022	0.668	0.034	0.015	0.662	0.015	0.007	0.602
Calendar year FE	yes	no	no	yes	no	no	yes	no	no	yes	no	no	yes	no	no
Industry FE	yes	no	no	yes	no	no	yes	no	no	yes	no	no	yes	no	no
Firm-Year FE	no	no	yes	no	no	yes	no	no	yes	no	no	yes	no	no	yes
AIC	-4940	-4993	-4665	-36989	**-37023**	-36264	-10956	-11015	-11229	-10778	-10826	-10344	-31045	-31092	-28344

Appendix C: *OLS* MODELS

This table reports coefficients from various OLS models of excess votes and total 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.

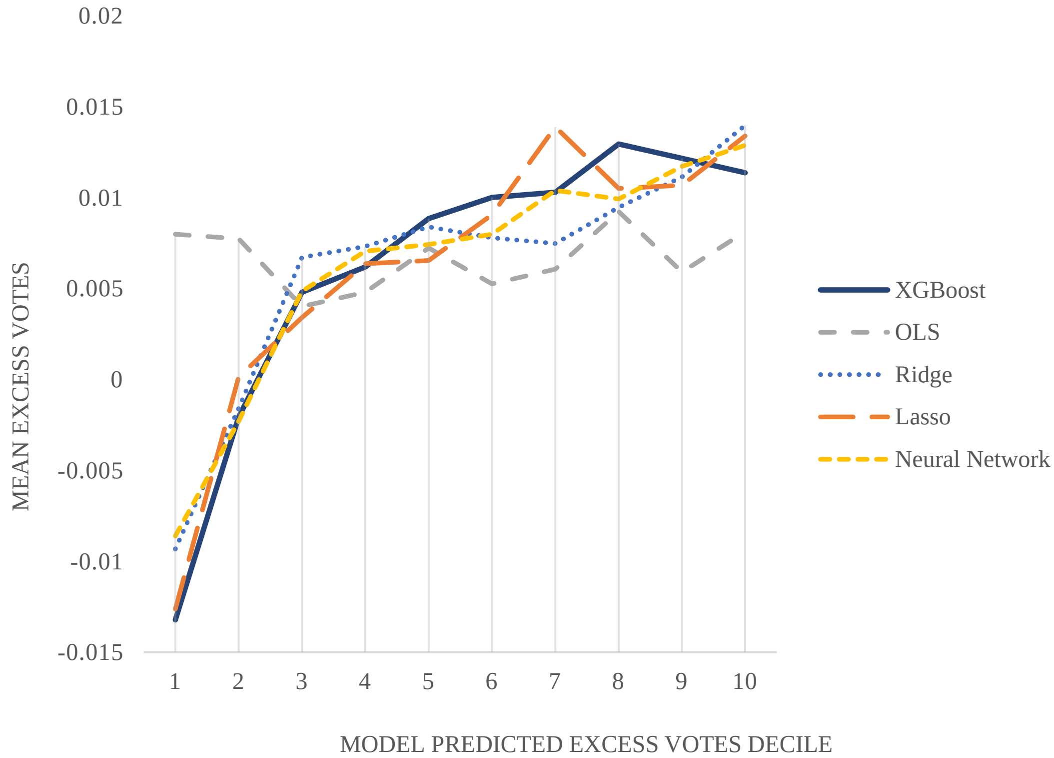

FIGURE 1: MEAN OBSERVED *EXCESS VOTES* VS. PREDICTED *EXCESS VOTES*

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.

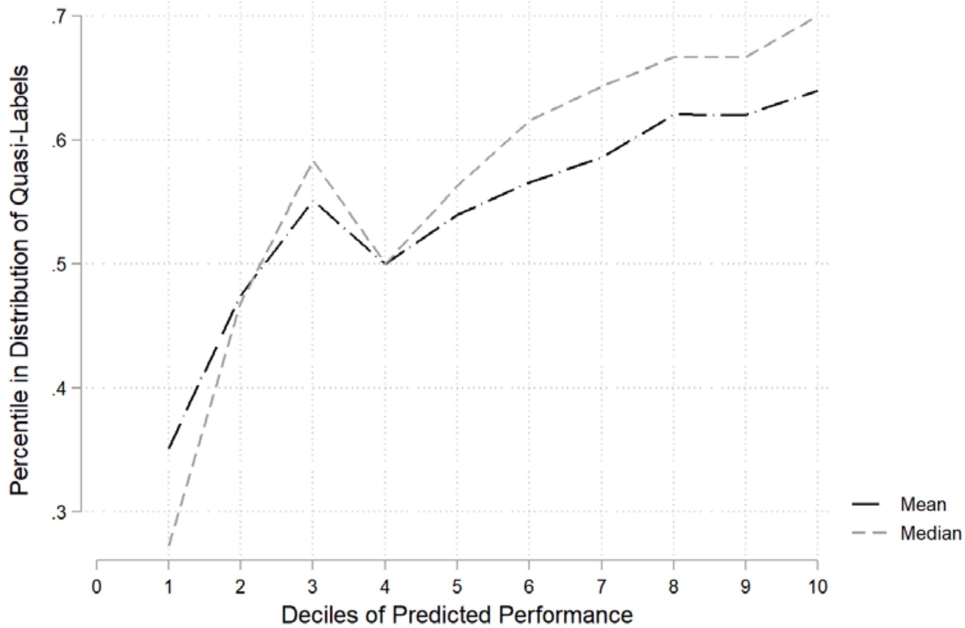

FIGURE 2: MEAN AND MEDIAN RANK IN QUASI-LABEL DISTRIBUTION ACROSS DECILES OF PREDICTED PERFORMANCE

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

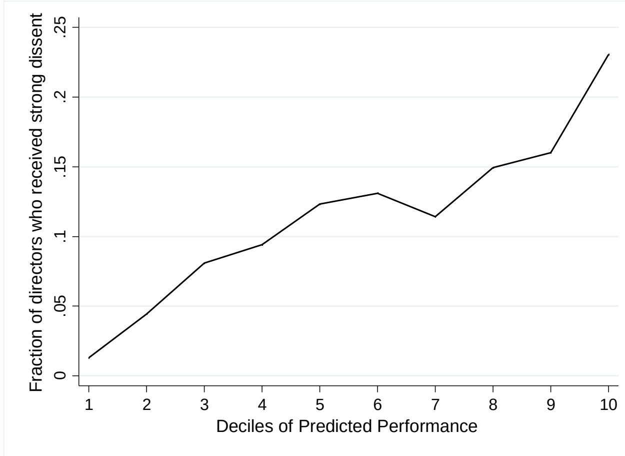

FIGURE 3: MEAN OBSERVED *DISSENT* VS. PREDICTED *DISSENT*

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%.

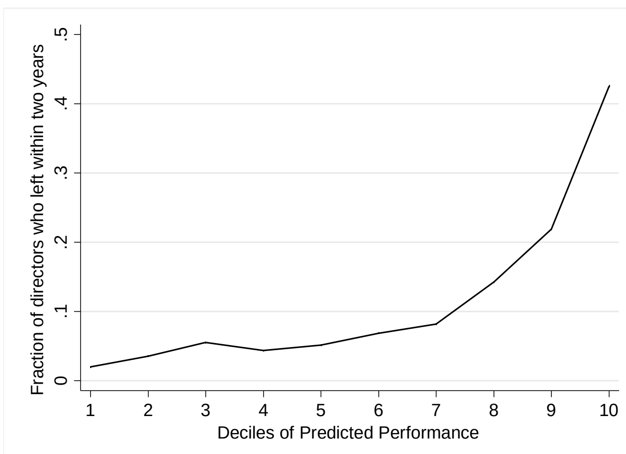

FIGURE 4: MEAN OBSERVED *DIRECTOR TURNOVER* VS. PREDICTED *TURNOVER*

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

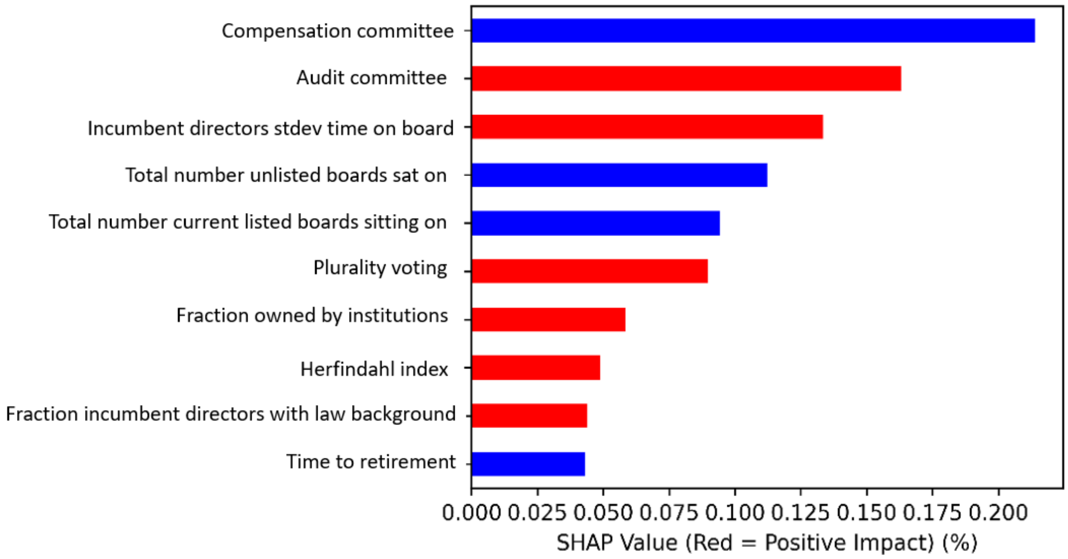

Panel B: (Larger than 10%) Dissent as a Measure of (Worse) Director Performance

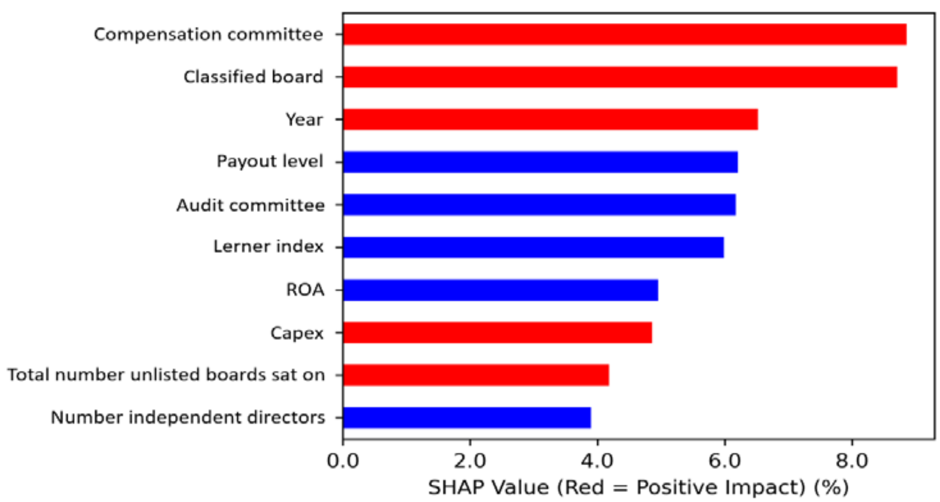

Panel C: Director Turnover as a Measure of (Worse) Director Performance

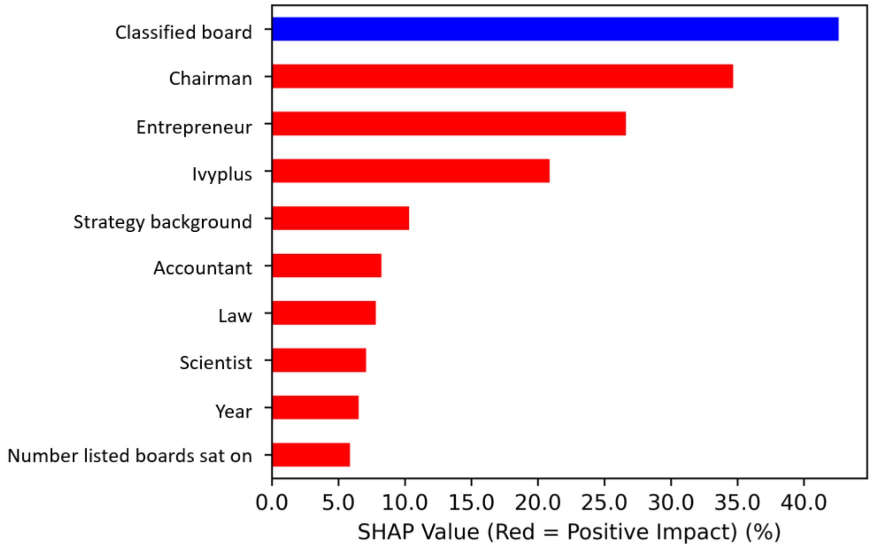

FIGURE 5: VARIABLE IMPORTANCE PLOT

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

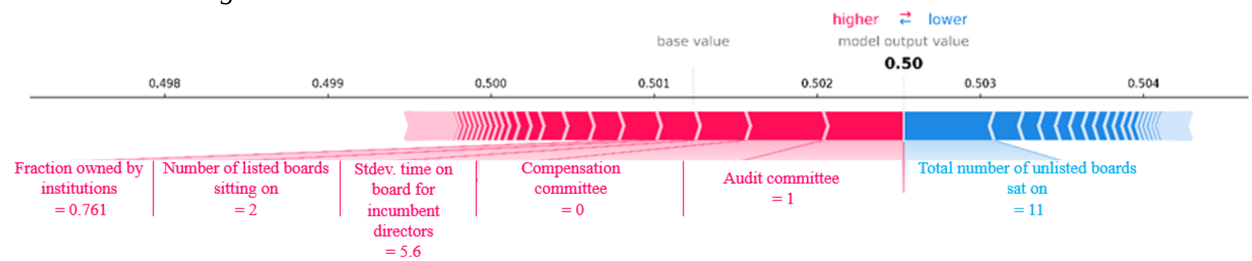

Panel B: Predicting Dissent

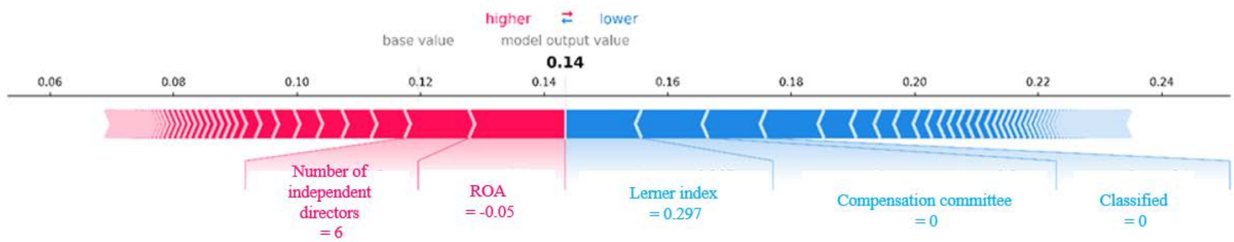

Panel C: Predicting Turnover

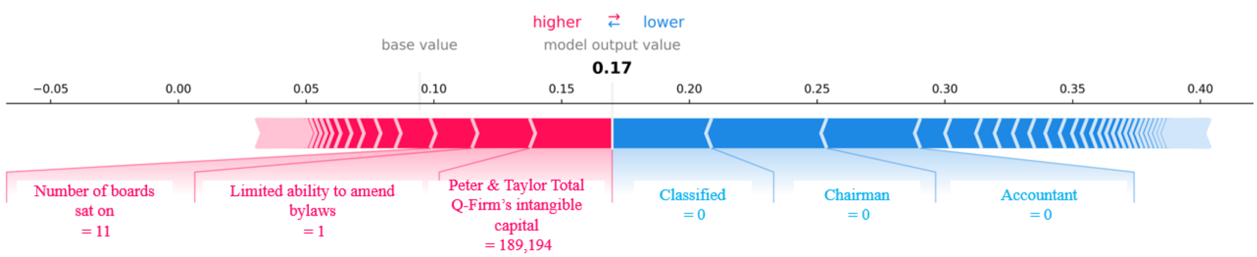

FIGURE 6: INDIVIDUAL SHAP VALUES FOR LOCAL INTERPRETABILITY

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	Mean Total Votes	Mean Excess Votes	Mean Dissent	Mean 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.

	Focal 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 6: 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 7: 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 include the following additional sections:

1. Section IA1 provides detailed variable definitions and data sources.

2. Section IA2 provides details for Ram and CPU requirements.

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

4. Figure IA2 presents XGBoost prediction results using a training period of 2003-2010 (post-SOX only) rather than 2000-2010.

5. Figure IA3 presents the mean and median rank in Quasi-Label Distribution across Deciles of *XGBoost*-Predicted Performance, where we relax the restriction that candidates joined a *smaller* firm.

IA1. DATA DEFINITIONS

IA1.1. 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

IA1.2. 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

IA1.3 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' website)
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 2002
CFOSO1	Equals to one if the CFO is exempt from filing Certification Documents as required under section 302 of the Sarbanes-Oxley Act of 2002
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)

IA1.4. 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)

IA2. RAM AND CPU REQUIREMENTS

XGBoost is readily available using the *XGBoost* package for Python available at https://pypi.org/project/xgboost/.

The models were run on a machine with the following specifications:
Memory: 256GB
CPU: Intel Xeon E7 (48 cores)
System: Ubuntu 18.04

The computing time to train the model for focal directors are:
Lasso: 536 ms ± 30.8 ms per loop (mean ± std. dev. of 7 runs, 1 loop each)
Xgboost: 18.9 s ± 284 ms per loop (mean ± std. dev. of 7 runs, 1 loop each)

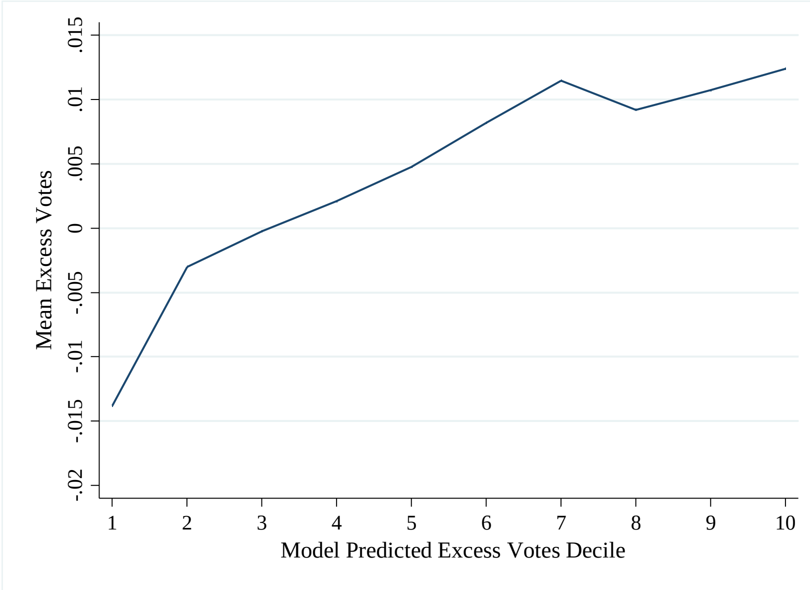

FIGURE IA1: MEAN OBSERVED *EXCESS VOTES* VS. *XGBOOST*-PREDICTED *EXCESS VOTES WHEN TEST SET INCLUDES APPOINTMENTS BETWEEN 2011-2014.*

This figure shows the average observed level of excess shareholder support for directors in the test set across the ten deciles of predicted performance for the *XGBoost* model when the test set includes appointments between 2011 and 2014.

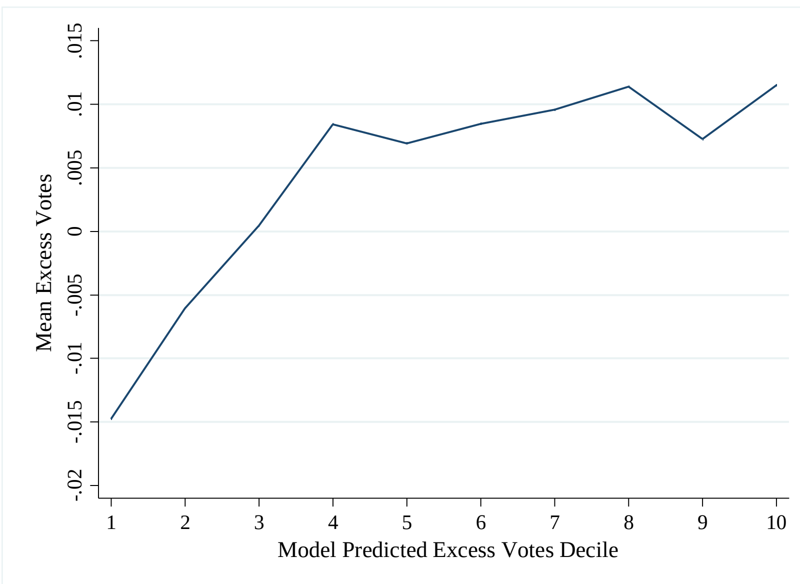

FIGURE IA2: MEAN OBSERVED *EXCESS VOTES* VS. *XGBOOST*-PREDICTED *EXCESS VOTES WHEN TRAINING SET INCLUDES APPOINTMENTS BETWEEN 2003-2010 (POST-SOX ONLY)*

This figure shows the average observed level of excess shareholder support for directors in the test set across the ten deciles of predicted performance for the *XGBoost* model when the training set includes appointments between 2003 and 2010.

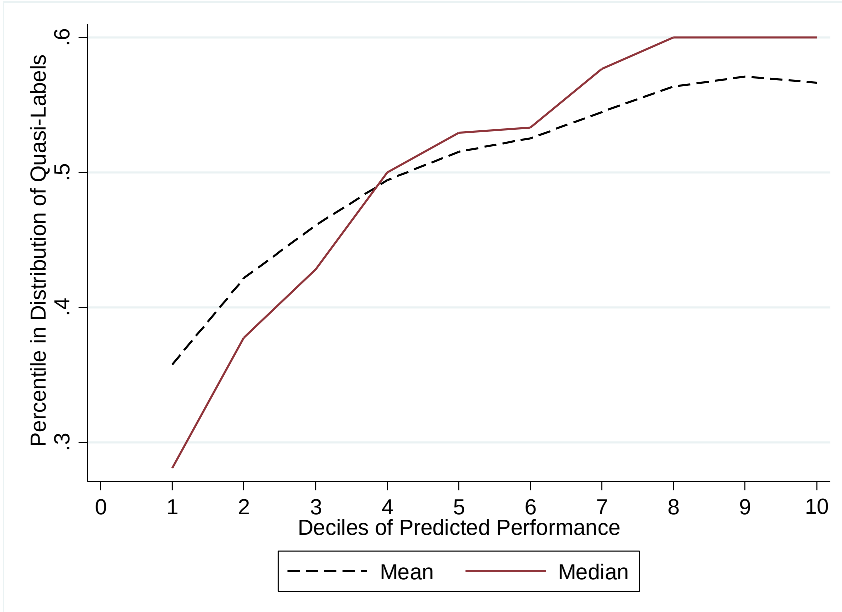

FIGURE IA3: MEAN AND MEDIAN RANK IN QUASI-LABEL DISTRIBUTION ACROSS DECILES OF *XGBOOST*-PREDICTED PERFORMANCE (FIRM SIZE ASSUMPTION RELAXED)

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.