PRELIMINARY,
NOT FOR QUOTATION

Do individual directors matter?

Evidence from the S&P 1,500

Markus Senn*

Abstract

We ask whether directors on corporate boards contribute to firm performance as individuals. From the universe of the S&P 1,500 firms since 1996 we track 2,062 directors who serve on multiple boards over extended periods of time. Our initial findings suggest that the presence of these directors is associated with substantial performance shifts (director fixed effects). Closer examination shows that these effects are statistical artifacts and we conclude that directors are largely fungible. Moreover, we contribute to the discussion of the fixed effects method. In particular, we highlight that the selection of the randomization method is pivotal when generating placebo benchmarks.

* Institut für Finanzmanagement, University of Bern, markus.senn@ifm.unibe.ch, +41 31 631 3478.
This project was made possible by the support of the Swiss National Science Foundation. We are grateful for the valuable input from Stefan Aebischer, Morten Bennedsen, Demian Berchtold, Oliver Dichter, Michel Habib, Alexis Kunz, Claudio Loderer, Urs Peyer, Lukas Roth, Berk Sensoy, Dennis Sheehan, René Stulz, Urs Waelchli, Michael Weisbach, and Jonas Zeller. We also received helpful feedback from the participants of the Empirical Corporate Finance course in Zurich, the First Macro Uni Bern Conference in Hasliberg, the PhD seminar at OSU in Columbus, and the IFM research workshop in Novello. All errors are ours.

1. Introduction

Academic research considers director characteristics important determinants of board behavior and performance. Frequently discussed characteristics include directors' expertise (Agrawal and Knoeber 2001; Agrawal and Chadha 2005; Defond, Hann, and Hu 2005; Güner, Malmendier, and Tate 2008), experience (Fich 2005; Fahlenbrach, Low, and Stulz 2010), access to a network of relationships (Goldman, Rocholl, and So 2009), and number of directorships (Ferris, Jagannathan, and Pritchard 2003; Perry and Peyer 2005). Other studies are concerned with directors' effort. Thus, they investigate incentives arising from directors' compensation (Vafeas, 1999; Bryan and Klein, 2004; Yermack, 2004), their ownership stake in the firm (Walkling and Long, 1984; Morck, Shleifer, and Vishny, 1988; Gilson 1990), and the risk of being dismissed (Yermack 2004; Wälchli 2008). Yet another strand of literature focuses on the relationships between directors and firms. These articles investigate the consequences of directors' dependence on the firm (Bhagat and Black), interlocks between the CEO and the directors (Hallock 1997; Loderer and Peyer 2002; Bizjak, Lemmon, and Whitby 2009; Fahlenbrach, Low, and Stulz 2010), or board representation of the firm's founding family (Anderson and Reeb 2004; Villalonga and Amit 2006).

This literature typically investigates corporate boards from an aggregate point of view, such as the presence of a financial expert or the total fraction of voting rights controlled by the board. Generally, boards are, therefore, treated as portfolios of directors' characteristics, incentives, and relationships. A specific example of this approach is papers such as Bhagat and Black (2002), who analyze the relation between the overall fraction of independent directors and firm performance. Using soccer as an analogy, this approach resembles assessing the strength of a team by measuring the fraction of forwards—a description that conveys valuable information about the team but reveals little about the impact of individual players. Soft evidence from corporate practice, however, suggests that the individuals on the board matter. For example, the media often cover the nomination of new directors. Moreover, specialized firms and databases provide help in finding suitable candidates for board positions.

So far, most empirical knowledge regarding the role of individual directors comes from event studies (see Yermack (2006) for an extensive literature review). These articles investigate market reactions to the announcement of director nominations and departures. For example, Rosenstein and Wyatt (1990) document positive announcement returns upon the nomination of outside directors. Other articles find that the market does not, on average, react to these events but that announcement effects depend significantly on the nomination procedure (Shivdasani and Yermack 1999) and on characteristics of the nominated directors (Fich 2005; Goldman, Rocholl, and So 2009; Fahlenbrach, Low, and Stulz 2010; von Meyerinck, Oesch, and Schmid 2012). Nguyen and Nielsen (2010) investigate sudden deaths of directors. They document that these events are followed by declining stock prices. The magnitude of this reaction is related to factors such as the departed director's degree of independence or board structure.

Measuring the market reaction to events such as a director turnover is a fairly simple and useful way to assess what effect these events have on firm value. However, relying only on the market's judgment of directors' impact has drawbacks: First, the market benchmarks its reaction against previous expectations. Since these expectations are typically unknown to the empirical researcher it is difficult to properly interpret the reaction. For example, a positive reaction to the nomination of a new director would typically be interpreted as the market considering this individual a valuable addition to the board. An alternative explanation is that finding a new director was easier and, thus, cheaper than expected by the market. Second, event studies usually investigate very short observation windows. In efficient markets, where new information is quickly and correctly processed and incorporated into the share price, this reduces the likelihood of having confounding events. However, in the case of director nominations, it is debatable whether market participants can immediately assess the full effect an individual will have on its firm. After all, a director's board service usually lasts several years. Third, isolating the announcement return of a director joining the board is particularly difficult because the event is not clearly defined. On the one hand, nominations of directors are often clustered or they are announced together with other information such as retirement of incumbent directors, annual reports, or strategy changes. In these cases, it is unclear what information the stock market reacts to. On the other hand, candidates for directorship are typically nominated by the board previously to the annual stockholder meeting. However, it is only after the election at the annual stockholder meeting that they become directors. While stockholders typically vote in favor of the nominees, there is always a degree of uncertainty as elections may be contested.

Given these drawbacks, we suggest using a different method to cross-check our knowledge of the individual board members' contribution to performance. We do so by estimating director fixed effect regressions. That is, for each director we create an indicator variable equal to one if the director serves on a firm's board in a given year and equal to zero otherwise. We include these indicators in regressions that explain firm performance. This method is restricted to measuring constant parallel shifts in performance (fixed effects) associated with the presence individual directors. In return, it allows the measurement of unobserved individual effects. This technique has been applied before to gauge the influence of individual managers (Bertrand and Schoar 2003) and blockholders (Cronqvist and Fahlenbrach 2009) on firm behavior and performance.

We use data from IRRC and Compustat to construct a sample covering the boards of the S&P 1,500 firms from 1996 to 2009. We restrict our analysis to independent directors serving on two or more boards for at least 3 years and, thus, track 2,062 individuals. Our initial findings suggest that the explanatory power of performance regressions increases significantly if we include director fixed effects. These effects are distributed roughly symmetrically around zero. They are of large magnitude, the difference between a third and first quartile director's effect on ROA, for example, being 6.19 percent. This distribution is comparable to what Bertrand and Schoar (2003) find for managers and Cronqvist and Fahlenbrach (2009) for blockholders. The effects we estimate are driven by a large number of directors rather than being associated with few extraordinary individuals. In the ROA regressions, for example, 41.86 percent of the director effects are individually significant at the 10 percent level and 25.35 percent at the 1 percent level.

Fee, Hadlock, and Pierce (2013) point out that inference in the fixed effects environment could be troubled by the presence of serial correlation or because asymptotic theory may not apply. We evaluate whether the effects we present are economic reality or statistical artifacts by comparing these findings to the results of a fixed effects analysis on random data. Using four different randomization mechanisms we generate 400 placebo datasets on which we estimate fixed effects regressions. The outcome of these regressions serves as a benchmark for our previous results. We expect that the actual data lead to stronger results than placebo data if director fixed effects are economic reality. Should the results from actual data be indistinguishable from the placebo results we are lead to conclude that the director fixed effects are statistical artifacts.

We find that whether the results from real data stand out depends on how we generate the placebo data. Using a placebo benchmark that focusses on the number of possible outcomes, rather than on mimicking the original data precisely, we obtain economically meaningless results. That is, we find significantly stronger director effects on random than on real data. Moreover, these results are very sensitive against relatively minor changes in the randomization method (that is, reassigning each director individually to new firms instead of full boards at a time). These worrisome results mostly disappear if we restrict the randomization process so that it more accurately mimics the original data. In particular, the results from real data are either stronger than the placebo results or they differ insignificantly. The sensitivity against smaller changes in the randomization process is reduced as well. The cost of restricting the reassignment mechanism is a substantial reduction in the number of possible outcomes from randomization. Therefore, it is more likely that the placebo datasets are partly identical with the real data, masking differences between the real and placebo estimates.

Overall, the real data estimates do not consistently differ from placebo estimates. Hence, as a first finding, we conclude that the individual directors of our sample do not on average influence their firms' performance in a way that can be measured with fixed effects. That is, our results show no evidence of directors who unconditionally improve or harm firm performance. Rather, they draw a picture of fungible directors. Hence, according to our findings, there are no super star directors that firms should be chasing after. On the other hand, we find no super loser directors either.

As a second finding, we point out some obstacles that need to be overcome when using the fixed effects method. First, we confirm that benchmarking the results of fixed effects analyses against placebo results is indispensable (Fee, Hadlock, and Pierce 2013). The test statistics typically used are inflated when working with serially correlated dependent variables. Moreover, even if residuals are well behaved, inference may be flawed as the theoretical distribution of the test statistics is unknown. Second, to our knowledge we are the first to highlight the difficulties in using random benchmarking with fixed effects. In particular, the randomization mechanism has to compromise between accurately representing the original data and keeping the number of possible outcomes high. Finding this compromise is difficult if the structure of the indicator variables is complex, such as is the case with director data. Ultimately, the data at hand define what this compromise looks like.

The remainder of this paper is structured as follows: In the next chapter we clarify why we think that individual outside directors might be associated with shifts in firm performance. Chapter 3 explains how the sample is constructed and provides descriptive statistics on the dataset. These statistics focus particularly on the board service of the directors we track and their movements across firms. The following chapter introduces the regression model and the control variables. In the fifth chapter we estimate the individual director fixed effects. We discuss the overall significance of these effects as well as the distribution of the estimated coefficients. While these intermediate results paint a picture of highly influential directors, the following chapter raises methodological issues. Evaluating these issues we conclude that our test statistics are inflated. Chapter 7 derives the empirical distribution of the inflated test statistics through repeated random draws. Knowing this distribution allows a more accurate interpretation of our initial findings. Chapter 8 concludes.

2. Why should outside directors be associated with firm performance?

The general idea that individual outside directors can influence their firm is widely spread in the academic community and among regulators. In particular, when investigating different classes of directors (e.g., foreigners (Masulis, Wang, and Xie 2012), bankers (Byrd and Mizruchi 2005), and women (Adams and Ferreira 2009)) authors ask whether the presence of one or more of these directors is related to firm outcomes. Hence, these authors assume that even a single outside director can influence a firm. Similarly, the Sarbanes-Oxley Act asks for at least one independent financial expert on the board's audit committee.

We further investigate the role of the individual on the board by asking whether individual directors are associated with shifts in firm performance. To see why we expect such a relation, we have to look at how firms select their board members.[1] Our reasoning begins with a frictionless setting: In this ideal world, a firm (1) has enough information to select optimal board candidates, (2) has access to optimal candidates at all times, (3) can quickly and at a low cost adapt board composition to its needs, and (4) does not face agency problems in the selection of directors. Such a firm should always staff its board with individuals whose influence maximizes firm value. This kind of situation, where all directors influence their firms optimally, cannot be distinguished from a situation where the directors have no influence at all using fixed effects analysis. This is the case because fixed effects is a relative concept: It implies that the presence of a given director leads to a shift in performance relative to another board.[2] Such relative differences cannot occur if all directors influence their firms ideally at all times.

In practice, firms' selection of directors is unlikely to be frictionless: (1) The nominating firm cannot perfectly observe the individual's aptitude for the position. (2) The nominating firm has a limited set of candidates (e.g., super star directors are scarce, non-compete agreements hinder the nomination of high profile industry experts, reputation concerns deter promising prospects). (3) Board composition is

[1] A large body of literature investigates how firms compose their boards depending on their needs. For example, Coles, Daniel, and Naveen (2008) provide empirical evidence that board size is adjusted to firm complexity. With respect to individual directors, Agrawal and Knoeber (2001) document that utility firms add politically connected directors to their boards when politics get more important for the industry.

[2] Depending on the model specification, this alternative board can be the board of another firm or the same firm's board at some time the respective director is absent.

sticky. For example, as of 2002, the majority of S&P firms had staggered boards (Bebchuk and Cohen 2005). That is, the directors of these firms are elected to multi-year terms and only a fraction of the board can be replaced each year. (4) Board members are not always nominated for their impact on firm value. Examples include family firms that elect relatives to the board, blockholders who seek board representation, CEOs who nominate interlocking directors or directors 'not rocking the boat' (Hermalin and Weisbach 1998; Adams, Hermalin, and Weisbach 2010). These frictions in the selection of directors cause firms to end up with directors who are suboptimally suited for the board.[3] This suboptimal influence leads to variation in board performance which may take the form of director related shifts in firm performance.

So far, we discussed why the directors' influence could lead to observable shifts in performance. However, similar data patterns might result from certification (Fahlenbrach, Low, and Stulz 2010): Assume a firm knows it heads for a period of above average performance. Such a firm might seek to certify the imminent outperformance to outsiders by adding a reputation sensitive individual to its board. Data patterns resembling directors' influence emerge if the certifying directors have gained their reputation by previous service on boards of outperforming firms or if they successfully certify multiple firms. In these cases, directors would seem to bring about positive shifts in performance although the causality is opposite. This effect can occur independently of whether boards are ideally composed. In particular, certification is in line with the ideal world setting if the certifying directors do not distort board decisions and the benefit of certification outweighs its cost.

Table 1 provides an overview of how possible findings of a fixed effects analysis are related to the causal effects we discuss. We can identify certification as the cause of observed fixed effects if we are confident that board selection is not distorted (i.e., no frictions). If board selection is distorted, fixed effects observed can be induced either by certification or by influence. For this case, Cronqvist and Fahlenbrach (2009) suggest an identification strategy that distinguishes the different causes based on the timing of the effects. Similarly, there are multiple explanations why we may not observe fixed effects. In a frictionless world, observing no effects means either that all individual directors have optimal influence or that they have no influence at all. We further conclude that firms do not engage in certification if we observe no fixed effects in a world without frictions. Dropping the assumption of optimal board composition, the absence of fixed effects lets us conclude that individual directors do not matter in a way consistent with fixed effects. Unless the cost of certification erodes the imminent outperformance, we also conclude that there is no certification being done based on pervious performance on boards. Our quest for fixed effects begins in the next chapter where we discuss sample construction and describe the data.

[3] Extant evidence that firms can end up with CEOs whose personality (i.e., preferences, skills, and experiences) harms performance. For example, Malmendier and Tate (2009) find that winning prestigious business awards changes CEOs' preferences towards spending more time on activities outside the firm. Meanwhile, firm performance deteriorates after the CEO wins the award. Further evidence comes from family firms: Pérez-González (2006) documents that if a CEO is replaced by a relative of his, subsequent firm performance drops. The magnitude of the decrease in performance is related to the successor's education. Hence, as family status dominates the CEO selection we observe CEOs with a lack of ability.

3. Sample and data

Sample construction

We base our analysis on two different data sources. The first source is RiskMetrics, which allows us to track directors across firms and over time. This dataset lists all directors of the S&P 1,500 firms starting in 1996. It further classifies directors as internal or external. In 2007 the data provider introduced a new director identification number and partially discontinued the old one.[4] Hence, we devise our own director identification that incorporates information from the old and the new ID. This identification is verified by hand based on the directors' full name, age, and firm affiliations. The second dataset we use is Compustat from where we draw the firms' annual financial data. From this data we construct the performance measures that will be the dependent variables in our regressions and the time varying firm level control variables.

Starting with the combination of these two datasets, we exclude a number of firms: (1) firms from the financial and utilities sectors, (2) firms that are in the sample for less than three years[5], and (3) firms, that drop out of our analysis because of missing data (the selection and definition of our variables follows Cronqvist and Fahlenbrach (2009), see Appendix A). These restrictions produce a sample of 2,417 firms and 15,139 firm years. For our further analysis we select a subset of all directors classified as independent outsiders in this sample: First, we exclude all mandates (i.e., director firm-combinations) with less than three years overall tenure. We do so because it seems unlikely that directors influence their firms immediately after appointment. Sencond, we include only directors that serve on multiple boards within our dataset. This restriction will help us distinguish director effects from firm fixed effects. These two selection criteria follow those in Bertrand and Schoar (2003) for the case of managers. Ultimately, 2,062 outside directors meet these criteria and are tracked in our analysis.

Distribution of tracked directors across firm years

Table 2 describes how the directors we track are distributed across sample firms and years. Panel A shows the number of total sample boards the directors serve on. The observation that no tracked directors serve on less than two sample boards is a consequence of our selection criteria. Roughly 90 percent of directors serve on either two or three boards. Accordingly, the average director in our sample is tracked in 2.43 firms. No director was tracked across more boards than Ms. Gail R. Wilensky with a total of 9 different firms. Panel B investigates the number of boards a given director is simultaneously tracked on. In total, our sample covers 19,764 active director year-combinations, the majority of these were served on a single board. In less than 2 percent of the director year-combinations the individual held three or more positions simultaneously. In rare cases, directors are tracked in up to 7 firms in one year. On average a director is tracked in 1.58 firms at any given time. Panel C shows how a director's number of simultaneously held positions is distributed over the sample years. The majority of directors (1,749) is absent from the sample in at least one year and only

[4] This change in IDs followed the acquisition of the initial data provider (IRRC) by RiskMetrics in 2005.
[5] This follows from limiting our analysis to directors who are present in their firms for at least 3 years (see below).

21 directors are always present in 2 or more firms. Our sample contains 250 directors who never hold more than one board position at a time. The maximum number of 7 positions in a single year is reached by 3 individuals (Ms. Claudine B. Malone, Ms. Rozanne L. Ridgway, Mr. Willie D. Davis). Panel D provides information about how many different sample years a given director is tracked in. Here, the effect of another selection criterion becomes apparent: directors are tracked in at least 3 years. The median director is tracked in 10 out of 15 sample years,[6] the mean is 9.58 years. As a last characteristic, the duration of the director mandates within our sample is described in Panel E. Again, we see the minimum 3 years required by our selection criteria. On the other hand, we see that some directors stay with their firms for the full sample period of 14 years.

While Table 2 takes the directors' perspective to discuss how their board service is distributed across the sample, Table 3 takes the firms' perspective on this distribution. Panel A provides the number of directors that are simultaneously tracked in a given firm year. A large group of firm years (27.16 percent) have no tracked directors at all. In another roughly 50 percent of the firm years we track between 1 and 3 directors. The highest number of directors we track in a single firm at a time is 11. This distribution results in 2.06 tracked directors in the average firm year. Panel B shows the minimum and maximum number of simultaneously tracked directors for a given firm over time. From a total of 2,417 firms 999 never have a tracked director, 1,074 firms have always one or more tracked directors, and the remaining 344 firms have tracked directors in some and none in other years. The upper end of the distribution is marked by two firms for which we never track less than 8 directors simultaneously. These firms are Ameritech Corp. (1996 to 1998) and TRW Inc. (1996 to 2000).

Describing the sample firms

The disparity between the 40 percent of firms that never have any tracked directors and other firms that always have as many as 8 tracked directors suggests that the individuals we track are not arbitrarily distributed across the sample firms. Table 4 investigates this notion by comparing the financial and board data between firm years without and with tracked directors. For each variable, the values in the first row refer to the full sample. The second and third rows describe the subsamples of firm years without and with tracked directors. The columns show the number of observations, mean, standard deviation, and quartiles of each variable. The last column of the table tests for differences between the two subsamples providing the t-statistic (second row) and the z-value from Wilcoxon non-parametric tests (third row).

By sample definition, all variables that enter our regressions are observed in a total of 15,139 firm years.[7] These observations divide into 4,112 firm years without any tracked directors and 11,027 firm

[6] Our sample includes all directors elected between 1996 and 2009, implying a sample period of 14 years. The reason why the table lists up to 15 firm years is the following: The first directors were elected in early 1996 and, thus, were present starting in firm year 1996. The last directors were elected in late 2009 and thus were present until firm year 2010. However, no firm can be in the sample for more than 14 years.

[7] Note that Board ownership and Board tenure were often missing. To avoid losing too many observations we substituted zero for missing values and included an indicator variable that marks

years with one or more of them. From Panel B of Table 3 we know that 999 firms are always in the subsample without and 1,074 in the subsample with tracked directors. The remaining 344 firms switch between subsamples at least once.[8] Note that for variables we are not using in the regression analyses the number of observations may be lower.

The table first shows the measures of firm performance that are going to be the dependent variables in the fixed effects regressions: Return on assets, Return on equity, Sales to assets, and Tobin's Q. The first two variables measure operating performance and are commonly used in the corporate governance literature, Sales to assets has been proposed as a measure for the efficiency of asset utilization (Yermack 1996; Ang, Cole, and Lin 2000), and Tobin's Q represents the market valuation of the firm's assets. Overall, firm years with one or more tracked directors seem to perform better than firm years without. These differences are statistically significant for both measures of operating performance and for the efficiency measure. Moreover, the difference in operating performance is economically large. For example, the median Return on equity is 2.16 percent higher for firm years with tracked directors, comparing to a full sample inter quartile range (IQR) of 12.91 percent. Firm years with tracked directors are also more efficient. For every dollar of assets they employ they generate 3.53 cents more sales than firm years without tracked directors. However, this difference is small relative to the variation within the subsamples. For Tobin's Q, the mean values of the two subsamples are comparable while the Wilcoxon points out significant differences in the distributions. Indeed, the median, first, and third quartiles are higher for firms with tracked directors. However, the magnitude of the differences is small.

Looking at firm characteristics other than performance, the difference between the firm size in the two subsamples is particularly striking. The median total assets are nearly three times larger in firm years with tracked directors than in firm years without. This difference is also statistically highly significant. Further, firm years with tracked directors have significantly lower relative capital expenditures (Investment) and higher Leverage than the ones in the other subsample. Also, they grow somewhat slower and are about 40 percent older.

In terms of board characteristics, firm years with tracked directors have larger boards (9 as opposed to 8 directors) and a greater fraction of independent directors (71 vs. 57%). Note that, these differences might arise endogenously from our selection criteria. The likelihood that at least one director of a firm gets selected is higher, the more independent directors a firm has on its board. Similarly, based on requiring a minimum tenure for tracking, we might expect to observe higher average board tenure in firm years with tracked directors. Somewhat surprisingly, we find the opposite regarding tenure. The likelihood that a single person serves as both CEO and COB is slightly higher in firm years with tracked directors. Further, directors in these firm years own a smaller fraction of the firm's stock (voting rights). This latter relation is to some extent driven by the difference in firm size. In fact, the average USD amount of stock held by the board is larger in firms with tracked directors.

---

these substitutions in the regressions. This substation biases the descriptive statistics in this table towards zero.

[8] Constructing the subsamples based on whether firms ever have a tracked director leads to comparable results. The differences between the subsamples increase for most variables.

Overall, the relation between firm characteristics and the presence of tracked directors further supports the notion that tracked directors cluster around specific firms while they are rarely found in others. In this sense, the individuals we track do not represent directors at large. Pinpointing exactly what makes these directors special is difficult as there are many different theories about how directors are matched with firms and what the role of a director is. Fama and Jensen (1983), for example, argue that there is a competitive market for directors. Therefore, only effective monitors retain their positions and get to build a reputation that can result in additional mandates. Accordingly, a focus on long-term, multi-firm directors would select highly performing individuals and bias our sample towards overestimating directors' influence. An alternative argument is that different firms have different expectations regarding their boards' work. For example, Coles, Daniel, and Naveen (2008) find that firm specific knowledge is particularly valuable in an R&D-intensive environment. Consequently, we would expect such organizations to rely more heavily on insiders and have fewer tracked directors. In this case, it is not clear whether and how the sample selection biases our findings. Independent of whether our subsample of directors is biased, we think the group is large and prevalent enough to warrant our attention even if our findings do not easyily to generalize. After all, these particular directors' services are in demand by the majority of the S&P 1,500 firms. Further, this demand is primarily particularly high in larger, arguably more sophisticated, and financially powerful firms.

4. The regression model

Table 4 shows that tracked directors cluster around particular firm years. Most notably, these firm years are characterized by high performance. This overall difference in performance might confound our evaluation of individuals' contribution. Hence, this chapter evaluates whether the difference persists in the multivariate panel regression setting we use in our further analysis (Table 5). The chapter also discusses the estimates of the control variables' coefficients and compares them to previous findings.

The dependent variables in the regressions are Return on assets, Return on equity, Sales to assets, and Tobin's Q. The estimates control for year and firm fixed effects and the lagged logarithm of total assets. This specification follows Bertrand and Schoar (2003) and Cronqvist and Fahlenbrach (2009). In addition to the models used in these articles, we control for the lags of Investment (Yermack 1996) and Leverage (Ang, Cole, and Lin 2000; Faleye 2007). To avoid confounding effects, we further include board characteristics that are frequently discussed in the context of firm performance. These characteristics are board size, independence, ownership, and tenure (Yermack 1996; Anderson, Mansi, and Reeb 2004; Faleye 2007). Including the squares of these metrics allows for a possible non-linear component in these relations. The last board characteristic we control for is an indicator variable for firm years where the CEO serves as COB simultaneously.

The sample split from Table 4 is reproduced by an indicator variable that equals one in firm years with tracked directors and zero otherwise. This variable will later be replaced by the set of variables that indicate the presence or absence of each individual tracked director. The insignificant coefficient of this variable across all performance measures shows that the performance differential we observed in

Table 4 disappears in the multivariate setting. Thus, what appears to be outperformance in a univariate setting turns out to be normal performance in our regression context. Hence, having a director who qualifies for tracking is not, on average, associated with firm performance.

Overall, our regressions show that board characteristics are barely related to firm performance. This finding is somewhat surprising, given that these characteristics are much discussed in the extant literature. One possible explanation for this finding is that many board characteristics vary little over time. Consequently, their effect might be picked up by the firm fixed effects. Indeed, most articles investigating how board characteristics affect firm performance do not control for firm fixed effects (Yermack 1996; Anderson and Reeb 2004; Faleye 2007; Coles, Daniel, and Naveen 2008; Anderson, Duru, and Reeb 2009). An exception is, for example, the indicator for firm years where a single person holds the positions of CEO and COB. This variable's coefficient is positive and significant in the Return on assets and the Sales to assets regressions. This finding is in line with studies that argue that unitary leadership is beneficial to firms (Donaldson and Davis 1991; Brickley, Coles, and Jarrell 1997). Board size and its square term are only significant in the Tobin's Q regression. Figure 1 compares our binomial estimate of this relation to previous findings. The three lines show the predicted change in Tobin's Q when board size increases by one director (marginal effects). The solid line represents our estimate, the dashed and dotted lines result from Yermack's (1996) log-linear and Faleye's (2007) linear predictions, respectively. The overlaid histogram shows how often each board size occurs in our data set. The white bars represent 90.23 percent of data. In the range where most of our observations lie the predicted changes of all three studies are close to one another. Notable differences arise with very small or large boards. However, as only few observations occur in these areas, the differences are most likely statistical artifacts.

The financial control variables are significantly related to all measures of firm performance. We find a negative association between performance and firm size (Assets) and Leverage, respectively. Investment is positively related to Sales to assets and Tobin's Q, unrelated to Return on assets, and negatively related to Return on equity. These results are difficult to compare to the fixed effects studies which we base our model upon because neither of them reports individual coefficients for their performance regressions. If we consider other studies for comparison, the coefficients they estimate vary substantially. Even within studies the coefficients can change from significantly positive to significantly negative depending on how the dependent variable is defined (see, for example, the size coefficients in Anderson and Reeb (2003)).

The explanatory power (adjusted R square) of our Return on assets regression is 0.61. Thus, it lies between what Bertrand and Schoar (2003) and Cronqvist and Fahlenbrach (2009) find before including manager or block holder fixed effects, respectively. Note that the explanatory power of these studies varies widely (between 0.57 and 0.72). This difference might be related to the different sample periods they cover. The R square of our Tobin's Q regression is 0.64 which is again lower than the one from Cronqvist and Fahlenbrach (2009).

5. Individual director effects

Overall, the discussion of Table 5 has shown that our regression setting is similar to previous analyses: our basic specifications are common in the literature, there is no inherent performance differential between firm years with and without tracked directors, and the results roughly match our expectations. Table 6 replaces the indicator variable for firm years with tracked directors by the individual director indicators. Given that our sample includes 2,062 directors and we estimate a separate coefficient for each of them, it is impractical to report all coefficients individually. Instead, Table 6 tests for joint significance of groups of coefficients (Panel A). It reports F-values associated with the hypothesis that the respective group of coefficients is jointly zero. Put differently, the null hypothesis is that the respective group of coefficients does not contribute to the model's explanatory power. The groups of coefficients are Individual director effects, Board characteristics, Firm level financial control variables (X_{it}), Firm fixed effects, and Year fixed effects; these groups include the variables introduced in Table 5.

We find that the Individual director effects contribute significantly to the model's explanatory power for all four performance measures. This finding is confirmed by the increase in adjusted R square compared to the previous table. The 0.05 increase in explanatory power in the Return on assets and Tobin's Q regressions are particularly noteworthy as the models without individual director coefficients already explain a large part of the depending variable's variation. This finding compares to the two previous fixed effects studies where the inclusion of blockholder or manager fixed effects increases explanatory power by 1–5 percent (Bertrand and Schoar 2003; Cronqvist and Fahlenbrach 2009). Similarly, the groups of Firm level financial control variables, the Firm fixed effects, and the Year fixed effects are highly significant. However, the group of Board characteristics is barley or not at all significant. This latter finding is in line with the observations we made in Table 5.

Panel B describes the distribution of the group of coefficients and the individual significance of the director effects. First, we note that 1,854 coefficients are estimated in each of the performance regressions. Thus, we estimate 208 coefficients less than the number of directors that qualify for tracking. These indicators were dropped from the regressions because of collinearity. For example, this could be the case if a director remains active for the full sample period on every board she serves. It would be impossible to distinguish this director's fixed effect from firm fixed effects. The table then provides the quartiles of the distribution of the 1,854 coefficient estimates. For comparison, it also reports the median of the regressions' dependent variable (Median (y)). The coefficients are distributed roughly symmetrically and their median is close to zero. This finding is in line with the observation that having a director who qualifies for tracking does not, on average, influence performance (see discussion of Table 5). However, the effects' IQR indicates that there are substantial differences between individual directors. For example, a firm with a third quartile director is predicted to have a ROA that is 6.19 percent higher than a firm with a first quartile director. Given a median ROA of 10.18 percent this difference is economically important. For Tobin's Q we estimate director effects of similar relative magnitude, for ROE the magnitude is even larger while for Sales to assets it is somewhat smaller.

While the estimated coefficients vary greatly, we so far only test for their joint significance. However, this test might reflect few extraordinary individuals with a large impact on their firms whereas the majority of coefficients are insignificant. In this case, talking of director effects in general would be misleading. To assess this notion we provide the fraction of individually significant coefficients (t-tests). At the 0.01 confidence level 16.02–28.21 percent of the 1,854 individual coefficients are significant, depending on how we measure performance. Similarly, 25.19–38.08 percent and 32.36–45.04 percent of coefficients are significant at the 0.05 and 0.1 level, respectively. The observed effect, therefore, is widely spread and not limited to a handful of exceptional directors, granted that the directors in our sample are already exceptional in the sense of the sample criteria.

6. Methodological issues

The F-values in Table 6 are strikingly high compared to what we observe in similar studies. Cronqvist and Fahlenbrach (2009), for example, find F-values of 1.27 and 1.75 when assessing the importance of large shareholders on firm performance. The F-values from Bertrand and Schoar (2003) are somewhat larger but, the highest one being 53.48, they are still small compared to the ones we find. These high test statistics might indicate methodological issues. The first potential problem is serial correlation. Bertrand, Duflo, and Mullainathan (2004) show that test statistics in difference-in-differences regressions can be severely inflated by the presence of serial correlation. The typical model they discuss is:

$$Y_{ist} = A_s + B_t + cX_{ist} + \beta I_{st} + \epsilon_{ist} \ ,$$

where Y_{ist} is the outcome of firm I in group s at time t. A_s and B_t are fixed effects for groups and years, respectively. X_{ist} are firm-level controls, I_{st} is a group-level treatment indicator, and ϵ_{ist} is the error term. This model is reasonably similar to the fixed effects approach we use, which can be written as:

$$Y_{it} = A_i + B_t + cX_{it} + \beta I_{kit} + \epsilon_{it} \ .$$

We note two differences: First, our treatment is not group-specific. Bertrand, Duflo, and Mullainathan (2004) consider the case where treatment is administered, for example, on a state level. That is, all firms incorporated in the same state receive the same treatment simultaneously. In our setting, treatment is the presence of a certain director k in firm i at time t. Thus, firms receive different treatments at different times and firms may even receive multiple treatments simultaneously. Consequently, it makes little sense for us to control for state fixed effects. Instead, we control for fixed effects of the firm (A_i). The second major difference is that, in the typical case of Bertrand, Duflo, and Mullainathan (2004), once treatment has occurred the treatment indicators take the value of one and remain constant from then on. Our director indicators, however, can return to zero if a director leaves a board.

Despite these differences, the models are similar enough for us to believe that the arguments of Bertrand, Duflo, and Mullainathan (2004) remain valid in our setting. Thus, we expect that our test statistics are inflated in the presence of serial correlation. Table 7 evaluates whether the residuals are indeed serially correlated. It does so, by regressing the residual on its first, second, and third lags. The table shows a positive correlation between the residual and its first lag, and negative correlations with the further lags. The only exception to this finding is the Return on equity model where all lags are negatively correlated to the current residual. The explanatory power is substantial with R squares between 0.11 and 0.16. Thus, serial correlation seems to be an issue in our estimates.

Serial correlation can sometimes be eliminated by including the lagged dependent variable on the right hand side of the regression.[9] Table 8 estimates the fixed effects regressions from Table 6 adding the lagged dependent variable. The italic printed rows of Panels A and B compare how this change affects our estimates relative to Table 6. In the Sales to assets and Tobin's Q regressions F-values for the group of director indicators drop. For the other two performance measures, however, the F-values increase. The number of observations (firm years) and the number of effects estimated remain unchanged for the three measures of operating performance. In the Tobin's Q regression they both drop slightly because the lagged market price of equity does not always exist. The IQR of the estimated director coefficient is reduced for most performance measures. Here, the exception is Return on equity where the IQR widens slightly. The fractions of individually significant coefficients do not vary much between the two tables. Altogether, our results are somewhat affected by including the lagged dependent variable on the right hand side of the model. However, we do not find a marked decrease in test statistics as we would expect if the serial correlation problem were solved. Accordingly, Table 9 shows that there is substantial serial correlation left in the residuals by, again, regressing them on their lags. Thus, controlling for the lagged dependent variable does not solve the issue of serial correlation in our case.

Independent of the potential serial correlation problems, the second methodological issue is that F-tests might be invalid in the context of fixed effects (Wooldridge 2002). When testing for the significance of fixed effects, the number of coefficients grows alongside the number of cross sectional observations. Therefore, asymptotic theory does not apply and we can no longer assume that the test statistic is F-distributed. Fee, Hadlock, and Pierce (2013) propose testing this concern by repeating the fixed effects analysis using placebo data. In particular, they suggest randomly recombining the director indicators and firms' financial data. Using this scrambled data we expect not to find a significant effect. If the tests turn out significant, we should conclude that they are invalid.

The first step in placebo testing is generating a randomized dataset. In deciding how exactly this should be done we have to trade off the number of possible random outcomes against the accurate representation of the original data structure in the placebo dataset. That is, we impose restrictions on the random reassignment to ensure that certain features of the original data are maintained. At the same time, adding restrictions reduces the number of possible outcomes. Therefore, some authors decide against employing any restrictions when randomizing. For example, Cronqvist and Fahlenbrach (2009) randomly reassign block holder data and financial data on a firm year basis. This procedure has a very large number of possible outcomes. While the resulting dataset may appropriately represent the authors' ownership data, it would poorly mimic our director data. After all, director

[9] See Keele and Kelly (2006) for a discussion of the exact circumstances under which controlling for the first lag of the dependent variable improves the estimates.

indicators are not randomly distributed across firm years but show a number of distinctive characteristics.

We think that one crucial characteristic of our data is that directors are typically active in a firm during multiple consecutive years. That is, the likelihood that a given director is active in a firm year depends greatly on whether that director was active in that same firm one year before. Thus, the director indicators are strongly serially correlated. To maintain this feature we make sure that randomization does not affect the years during which a director is active. For example, a director serving at Firm A from 1999 to 2004 will, after randomization, appear to have served at some other Firm B but during the same period of time. This restriction also prevents that consecutive years of service on one board are reassigned to a single year across multiple firms. Consequently, the number of firms a director is active in cannot increase during randomization.

Another restriction we consider is keeping whole boards together when reassigning the directors to new firms. We refer to mechanisms using this restriction as 'board level randomization'. If this restriction is in place and we reassign one director of Firm A to Firm B, all other directors of A are reassigned to the same B. Employing this restriction makes sure the distribution of the number of different directors we track in a firm year is the same in the placebo and the original data. Without this restriction we reassign each director to a different firm ('director level randomization'). That is, if we track two directors at Firm A one would be reassigned to Firm B and the other to Firm C.[10] This less restricted procedure allows for a greater number of possible outcomes. However, it abandons the original dataset's distribution of the number of directors tracked in a firm year. Instead, director level randomization implies a uniform probability distribution of mandates across firms. While the distribution of directors across firm years is unknown (it depends on the length and overlap of the mandates), it is unlikely to exhibit similarly clustered tracked directors as the original data. As the accuracy of the approximation and the consequences of deviations are unclear we report results for both types of placebo data.

The placebo data we generate by these two reassignment mechanisms are used in Table 10 to estimate fixed effects regressions. The results based on the two randomization methods are shown side by side and they turn out to be similar. As in Table 6, Panel A conducts F-tests for joint significance of the groups of variables. Most importantly, the group of Individual director effects remains highly significant for all performance measures. This result suggests that individual directors are related to firm performance even though these individuals did not serve on these firms' boards. Such a relation does not make sense economically and, hence, our previous concerns regarding the validity of the tests are substantiated. However, given that the test statistics' distribution is unknown, we should not rely on significance alone. Indeed, a closer comparison between Table 6 and Table 10 shows that the results differ: The F-values and the estimates' explanatory power from placebo data tend to be lower than the ones from real data. The change in IQR of the estimated individual effects

[10] In fact, we select a new random firm without restricting it to be different from either firm A or firm B. Therefore, the placebo data might also report the director at her original firm A or at the same firm B the first director was reassigned to. However, both cases are unlikely and would be purely coincidental.

depends on which randomization method we use. The estimates based on board-level placebo data tend to show a wider IQR while director-level placebo results show a narrower IQR. If real data results consistently stick out from placebo results, this would indicate that director effects are real after all (Bertrand, Duflo, and Mullainathan 2004). Therefore, the next chapter evaluates whether these differences are systematic.

7. Benchmarking against placebo data

Improved comparability and 100 random draws

So far we have discussed the outcome of a single random draw. Hence, we cannot tell whether the differences are typical or whether they occur just randomly in this one draw. We can increase the likelihood that differences are systematic if we analyze multiple placebo datasets. Therefore, in Table 11 we repeat the random reassignment and generate 100 different draws of the placebo data.

Moreover, to properly interpret the differences, we have to make sure the results are actually comparable. That is, placebo data has to accurately represent real data. In the case of Table 6 and Table 10, comparison is complicated because for the placebo results we estimate only about one third of the individual director effects we estimate in the real results.[11] This decline occurs during randomization: We reassign director and firm data in an unbalanced panel dataset while keeping the years unchanged. Thereby, we match director years with inexistent firm years. For example, a director who serves on the board of a firm from 1996 to 1999 might be reassigned to a firm that enters our sample only in 2002. In this case, we lose information because the firm and director activity do not overlap. Chances are that this director will no longer meet our selection criteria (i.e., serving on multiple boards for at least three years) and drop out of the sample.

The exact consequences of the lower number of directors are difficult to anticipate. However, we cannot safely assume that the outcome remains unaffected. For example, including more director indicators should lead to higher standard errors of the coefficients if the director indicators are correlated with one another. Consequently, the fraction of individually significant coefficients is decreases and the dispersion of the coefficients is widens. Indeed, comparing Panel B between Table 6 and Table 10, we find an increase in the fraction of individually significant coefficients. The change in the IQR of the coefficients depends on the choice of the randomization mechanism.

Table 11 improves comparability of the results by generating alternative real data estimates using a smaller number of director indicators (‘Subset’). These estimates deliberately exclude two thirds of the director indictors from the regressions. A random process decides which directors are included and excluded.

[11] Note that the number of effects is the same for each dependent variable. This is the case because we use the same random datasets to analyze all four performance measures. That is, in we use only one random data set for board and one for director level placebo data. Alternatively, we could generate new datasets for each dependent variable. However, there is no clear benefit from having individual dataset for each dependent variable.

The table is divided into four blocks of rows representing the four performance measures studied throughout this article. Each of these blocks is, yet again, made up of four rows: The first row is a summary of the analysis from Table 6. Row 2 represents the 'Subset' analysis, that is, the real estimates with a deliberately reduced number of director indicators. Rows 3 and 4 are based on datasets randomized at board and director-level, respectively. 'N of draws' shows how many repeated draws the respective results are based on. The full sample real data analysis does not involve a random mechanism and, hence, is conducted only once. All other analyses use random selection or reassignment and are repeated on 100 different draws. Next, the table displays the average number of individual director coefficients estimated per draw (N of directors). As we discussed earlier, the number of director effects in the full sample real data analysis is larger than in the placebo analyses. This difference is evened out by the subset analysis: The average difference between the subset and the placebo results is 11 and 1 directors for the two randomization methods, respectively. Note that a perfect match is impossible since the number of directors included varies slightly for each draw. This variation arises because collinearity depends on the random assignment and on which directors are included, respectively.

The table then investigates the F-values associated with the hypothesis that all director effects are zero. Since the estimates were repeated on 100 draws, the table describes the distribution of the corresponding F-values across these draws. The full set real data estimates lead to higher F-values for Sales to assets and Tobin's Q than the subsets. For Return on assets, both estimates result in comparable test statistics. Return on equity is the only performance measure where the F-value from the full sample analysis is low relative to the subset estimate. In particular, for two of our performance measures, the results of the full sample analysis are significant outliers compared to the subset analysis: only 6 and 5 random draws of the subsets produce F-values larger than the full set for Sales to assets and Tobin's Q, respectively. This finding confirms the notion that we cannot compare the full sample real estimates to the placebo estimates.

Comparing the F-values from the subset analysis with the ones from board-level randomized data, the random data tend to produce consistently higher test statistics. All quartiles of the distribution of F-values are higher in placebo data than subset data for all dependent variables. Moreover, the Kolmogorov-Smirnov (KS) test and the Wilcoxon test both mark the difference in distributions as significant. The sign of these differences is surprising. We would expect the statistics from real data to be higher if there are economic director effects. At most, the distributions should be indistinguishable if there are no effects. However, the results suggest that there is a stronger association between individual directors and firm performance in randomly selected firm years than in the firm years they actually serve. Such an association is economically meaningless.

These meaningless results mostly disappear when we look at the results from director level randomized data instead. While the Wilcoxon test for Return on equity still finds the F-values from placebo data to be larger than the ones from the subset analysis, the Kolmogorov Smirnov test indicates that this difference is insignificant. For Tobin's Q the two distributions of F-values are indistinguishable. However, for Return on assets and Sales to assets the subset analysis leads to higher F-values than the placebo analysis. This finding is in line with what we expected if there are economic director fixed effects.

Another statistic that might help distinguishing estimates on real data from ones on randomized data is R square. We expect the real director indicators to have a greater explanatory power than the placebo indicators. Comparing the subset with the board level placebo results we find lower average R squares in the placebo estimates for all performance measures. However, the largest difference is 0.39 percent (in the Sales to assets estimates), thus these differences are economically small. Looking at director level randomization we find even smaller decreases of R square for Return on equity and Sales to assets. For Return on assets and Tobin's Q the R square is greater when using placebo data. Thus, we are hesitant to interpret the changes in explanatory power as evidence for the materiality of the estimated fixed effects.

Cronqvist and Fahlenbrach (2009) distinguish real from placebo results by looking at the distribution of the estimated coefficients. Based on finding a wider dispersion of the real effects compared to the ones from placebo data they conclude that the effects do indeed exist and are economically significant. The last four columns of Table 11 replicate this comparison for our director fixed effects. If we compare the IQR from the full data estimate to the subset estimate, we observe that the coefficients of the model with fewer variables are consistently narrower distributed. This is in line with our expectation for the case that director indicators are correlated. Comparing the subset with the board level placebo data, the Kolmogorov-Smirnov test finds the dispersion of the subset to be significantly smaller for all performance measures. As with the distribution of F-values, this finding suggests that the association of directors and performance is stronger in random firm years than in the ones they were actually present. This finding, too, is difficult to reconcile with economic intuition and, as before, the findings are less problematic under director-level randomization. Though the IQR in the Return on assets and Tobin's Q estimates are wider for placebo than for real data, the differences are smaller than for board level randomization. Moreover, the IQR for Return on equity and Sales to assets are greater in real data than placebo data.

In conclusion, Table 11 shows that when using board-level randomization, we find substantial evidence suggesting that the placebo effects are stronger than real effects. This finding is absurd and indicates a problem with the random data set. When using director-level randomization these absurd results largely disappear. To be convinced of the existence of individual director effects, we would require the different statistics to agree that the real effects are stronger than placebo estimates. Sales to assets is the only performance measure for which this is the case. For all other dependent variables, the placebo results either produce indistinguishable F-values or a wider distribution of the estimated coefficients.

Investigating the absurd findings

The absurd findings when using board level randomization and the fact that many results depend on whether we randomize at the board or the director level are worrisome. There are multiple explanations for why our analysis may be troubled: The first potential problem is that randomizing at the board level may not produce truly random data. That is, the reassignment process may be restrictive to an extent that many boards are matched with their original firms. In this case, parts of our placebo data would be identical with the original data. Counting the number of potential outcomes we note that the first board is reassigned to one out of 2,417 firms, the second firm is reassigned to one out of 2,416 firms, and so on. Ultimately, 2,417! distinctive outcomes are possible.[12] Given this larger number of possible outcomes, it is unlikely that major parts of our placebo data are identical to the original data. Moreover, while insufficient randomness might explain the contradictions between the two randomization methods, it does not explain why placebo effects are significantly stronger than the real effects.

Second, these worrisome results might obtain because there are not enough random draws. Despite 100 repetitions, extraordinary outcomes might still dominate the findings. To get additional insights into this potential problem, Figure 2 looks at what happens to the differences between real and placebo results as the number of random draws grows from 1 to 100. If the outcome is stable over the course of the first 100 draws, we can be confident that our conclusions are robust against adding even more draws. If the outcome varies much over the course of the 100 draws, we might have to add further repetitions until the results turn stable.

Figure 2 plots the key statistics of Table 11 and their development as the number of draws increases. The dark lines compare the results from the subset analysis with the results from board-level randomization. The light lines look at director-level randomization, instead. For the comparison of F-values (left column), the statistics displayed are the p-value of the Kolmogorov-Smirnov test (solid lines, primary vertical axis) and the difference between the median of the real and placebo F-values (dashed lines, secondary vertical axis). The right column of the figure compares the distributions of individual director coefficients. Again, the solid lines represent the p-value from the Kolmogorov-Smirnov test. The dashed lines illustrate the difference in the IQR. For example, the first row of Figure 2 looks at the Return on assets regressions. The left chart compares the distributions of real and placebo F-values. For both randomization mechanisms, the p-value testing the equality of these distributions drops below 0.1 after around 30 draws (solid lines). For board-level placebo data the difference between the distributions remains significant thereafter. For director-level data, the p-value rebounds and reaches a peak of about 0.12 after 77 draws. Subsequently, it drops quickly to almost zero. Thus, concerning p-values, our conclusion is very stable when looking at board-level randomized data but somewhat less stable when looking at director-level placeboes. Quite the opposite is true for the difference between the median F-value. Here, it is the director-level data that produces stable results, whereas the board-level data fluctuates. However, both the sign and the order of magnitude of both differences remain constant after about 20 draws. Comparing the estimated individual director coefficients we find a similar picture. The Kolmogorov-Smirnov test for the equality of distribution of subset and board-level data is significant from the first draw and remains so everafter. For director level data, the test turns significant after 24 draws and is mostly significant thereafter, except for a

[12] For director level randomization there are 2,417! outcomes for each director. Thus, there are 2,417!^2,062 possible outcomes. However, given that a director is only active in 2.43 firms on average, many of these outcomes are identical. If each director served on exactly two boards, we would be left roughly (2,417 x 2,416)^2,062 distinct outcomes.

short rebound to 0.11 after 72 draws. Also, the differences between IQR are very stable over the course of the 100 draws. Thus, we conclude that 100 draws is enough for board-level randomization. For director-level randomization, the situation is less clear. However, as the numbers are relatively stable over the last 25 draws, it seems unlikely that our conclusions would change when we add more random draws.

Looking at the other performance measures, the outcome of the board-level placebo data is always highly stable (dark lines). For the director-level data, the outcome of comparing individual director coefficients is also stable. However, comparing F-values, odds are that our conclusion in the Return on equity regressions would change. Here, the p-value from the Kolmogorov-Smirnov test crisscrosses the threshold of 0.1, changing between significant and insignificant.

Overall, it is unlikely that the number of draws is the source for our worrisome findings. In particular, the results from board-level placeboes are very robust against changes in the number of draws, while we noted earlier that these results are most troubling.

The third issue that might trouble our analysis is the structure of the placebo data. The placebo benchmark is valid only if the structure does not differ substantively between real and random data. To see whether this is the case we identify a number of features that are characteristic for our original data. One such feature is that the director indicators are serially correlated. Our randomization mechanisms maintain this specific feature by keeping the years of service intact when reassigning directors to new firms. Three other features that are characteristic for our data are the number of tracked directors in a given firm year, the number of different boards a director is tracked on, and the duration of the mandates. These features are potentially affected by randomization. Figure 3 analyzes whether this is the case. For each feature it compares the two kinds of placebo datasets (board-level randomization is represented by squares and director-level randomization by triangles) with the full sample real data in the first and the subset real data in the second column (represented by the respective grey bars).

Panel A looks at the first feature, that is, it compares how directors are distributed across firm years in real and placebo data. Both kinds of placebo data sets have more firm years without any tracked directors than the full original dataset. Partly, this is the case because the total number of tracked directors is smaller in the placebo data, which is why we introduced the subset representation of real data. Accordingly, the right hand side of Panel A compares these subsets with the placebo data and shows that the differences decrease considerably. However, the placebo data still show more firms without any tracked directors than the real data. Knowing that the number of directors is adjusted quite accurately (see Table 11), this observation implies that randomization decreases the number of mandates per director or the duration of the mandates. Indeed, Panels B and C demonstrate that the directors are tracked on fewer boards and during fewer years in placebo data than in real data. These differences result from randomly matching director years with inexistent firm years. While we re-enforce our original sample criteria on the placebo data to avoid the most drastic changes in data structure (e.g., directors serving during less than three years or on single boards), substantial distortions remain.

Figure 3 further allows comparing the outcomes of board- and director-level randomization (squares and triangles, respectively). The figure shows substantial differences between the two methods. As we expected, director-level randomization distributes tracked directors more evenly across the sample (Panel A). This distortion shows in the overrepresentation of moderate cases (i.e., one or two directors) and underrepresentation of extreme cases (i.e., zero or more than two directors) relative to the board-level randomization. We argued earlier that board-level randomization should better mimic the original dataset as it carries over its distribution. Surprisingly, the opposite appears to be the case: It turns out that the distortions from breaking up boards coincidentally offset part of the distortions from reassigning directors to inexistent firm years, leading to a closer match in the frequency of firm-years with few directors. The only advantage of board-level randomization is that it more accurately matches the rare cases of firm years with three or more directors. In terms of the other characteristics of our data, the outcomes of board and director level randomization perform nearly identical (see Panels B and C).

Alternative random benchmark

Whether these remaining distortions in the data structure drive our results is easiest to find out by generating an alternative random benchmark that more precisely mimics the original data. Such a placebo is obtained by further restricting the set of firms which a director can be reassigned to. One way of doing so is reassigning directors only to firms that are in the sample during exactly the same years as their original firms.[13] For example, a director who worked for a firm that was present from 1996 to 2002 and absent during the remaining sample years can only be reassigned to other firms that were present from 1996 to 2002 exclusively. This restriction prevents director years from being reassigned to inexistent firm years. Hence, randomization distorts neither the number of firms a director is tracked in nor the duration of these mandates. Also, no tracked directors are lost in the process. The cost of this improvement is that the additional restrictions substantially reduce the total number of outcomes possible. Specifically, given that we identify 358 subsets of firms with identical years active, randomization takes place among 6.75 firms on average.

Figure 4 compares the data structure between the real data and the datasets produced by this alternative, more restricted randomization process. The alternative placebo data mimics the original data structure rather precisely compared to the datasets in Figure 3. In particular, Panels B and C show that the distributions of the number of different boards a director is tracked on and the duration of the mandates are exactly replicated. Panel A shows that board-level randomization exactly replicates the original distribution of the number of tracked directors per firm year. This outcome is in line with our initial expectations. Meanwhile, director-level randomization flattens the distribution of directors across firm years (i.e., fewer extreme and more moderate cases). In Figure 3, the distortions of this flattening

[13] Note, that we could also reassign directors to firms who are present in *at least* the same years. However, at some point we would lose directors. Consider what happens, for example, if we reassign the directors of Firm A to another Firm B that is present in more years than A. B's cannot be reassigned directors to A, as the latter firm is not active in all years that these directors are. While we may initially find a third firm active in at least the same years as B, we will eventually run out of options because the criterion is asymmetrical.

happened to partly offset the ones from reassigning directors to inexistent firm years. Now that the latter distortions are eradicated, this corrective aspect no longer applies.

Table 12 presents the results from estimating fixed effects regressions using these alternative placebo data. The third column of the table shows that the number of effects estimated on the placebo data matches the one from the original data quite accurately. Therefore, contrary to Table 11, no further adjustments are necessary to ensure comparability and the full sample results can be benchmarked directly against the placebo results. This direct comparability simplifies our analysis of the F-values (column 4). Instead of comparing two distributions as in Table 11, we can now compare a single value to a distribution. That is, we count the random draws that produce a test statistic greater than the F-value from our real data estimates. The fraction of random draws meeting this requirement is interpreted analogously to a p-value. The F-value for Sales to assets is the only one that, according to this pseudo p-value, stands out from the placebo results. For all other performance measures, the pseudo p-values are clearly higher than the 0.1 benchmark.

As a second criterion to distinguish the estimates we again consider the explanatory power. The adjusted R square of the placebo estimates is smaller than the one of real estimates for Return on equity and Sales to assets. For the other two measures our model has higher explanatory power on placebo data than on real data.

The last criterion we consider is the distribution of the individual director coefficients. The only significant difference in these distributions concerns the Return on equity estimates. For this performance measure we document a wider IQR of the coefficients estimated on real data. For the other three performance measures, the director coefficients are wider distributed in the placebo estimates. However, according to the Kolmogorov-Smirnov and the Wilcoxon test, these latter differences are insignificant. Hence, this criterion points towards the existence of economically material director effects for the Return on equity but not the other performance measures. This finding is at odds with the results of our previous tests where ROE was the performance measure giving us least reason to think that economic effects exist.

Figure 5 investigates whether Table 12 uses enough random draws to obtain stable results. The charts comparing the F-values of real and random results are simpler than ones in Figure 2. The reason is that we compare a single value to a distribution rather than two distributions. Therefore, the charts in the left column now show only the development of the pseudo p-value as the number of draws increases. The comparison of individual director coefficients (right column) is analogous to Figure 2. The figure shows that the p-values resulting from the distribution of F-values as well as the pseudo p-values form Kolmogorov-Smirnov tests are very stable. Neither of them crisscrosses a typical level of confidence. The differences in IQR of the individual director coefficients are stable as well. In particular, the sign of neither of these differences changes after about 10 random draws.

Overall, Table 12 presents little evidence that individual directors matter. While some tests turn out significant they do not aggregate to form a consistent picture of influential directors because they often don’t agree. That is, the distribution of F-values suggests there the only real effects are found when looking at Sales to assets. Meanwhile, the comparison of coefficients Return on equity is the only performance measure that shifts with the presence of certain directors. Trading off the randomization methods from Table 11 against the ones from Table 12, we note that the two issues bothering us largely disappeared. First, board- and director-level randomization typically agree in Table 12. Second, Table 12 produces barely any results suggesting that the placebo effects are significantly stronger than the real effects. Both improvements result from the additional restriction on the randomization process. That is, the placebo data in Table 12 more accurately mimic the original data. The cost of this better representation is that there are fewer possible outcomes. It follows that there is a higher likelihood that the placebo data are partially identical to the real data. Thus, we use an overly conservative benchmark for our results.

8. Conclusion

The extant literature documents that certain outside directors single handedly influence firm outcome. Examples for directors that matter are women, bankers, and foreigners. We go beyond this classification and ask whether outside directors as individuals are associated with shifts in firm performance. That is, we look for individual outside directors that are persistently associated with extraordinary performance of their firms. This question is answered using a fixed effects method that tracks selected directors across firms and over time.

The findings of this article are twofold. First, the substantive finding of our analysis is that the individual directors in our sample are not, on average, associated with performance shifts in their firms. When we discussed our expectations regarding fixed effects we noted that, in a frictionless world, this finding could mean either that directors don’t matter or that all directors influence their firm ideally. As this setting assumes that optimal directors are available at all times, these results contradict the notion of unique super star directors. Similarly, if we believe that board composition is not always optimal, our results suggest that directors don’t influence their firms in a way consistent with fixed effects. Thus, the directors we track appear to be overall fungible. This finding puts a question mark behind firms running after particularly high profile individuals for their boards.

Note that these findings do not rule out the possibility that individuals matter at all. In particular, the fixed effects method is ill-suited to detect directors’ influence that is conditional on firm situations such as raising capital, replacing CEOs, or conducting major acquisitions. The same is true if the individual director’s effect arises only through interaction with specific other board members, that is, in team dynamics settings. Moreover, we should investigate different samples of directors. On the one hand, our sample excludes directors who serve on a single board only. If these directors are the super stars, we overlook them. On the other hand, literature suggests successful directors are offered multiple mandates. In this case, we would have them in our sample. Nevertheless, our empirical methods might not be powerful enough to detect them if only a small fraction of directors are super stars. That is, our analysis might not detect them because of the noise induced by the remaining directors. This problem can be reduced by looking at particular groups of directors for whom we believe that they are particularly important to their firms.

Second, the article makes a methodological contribution to the discussion of the fixed effects method. We confirm that the reservations about the fixed effects method outlined by Fee, Hadlock, and Pierce (2013) should be taken seriously. Looking at individual directors' contribution to firm performance, we find extraordinarily high F-values when testing for the joint significance of director fixed effects. These test statistics suggest that including indicators for the presence of individual directors on the board substantially increases the explanatory power of our performance regressions. However, we find test statistics of comparable magnitude using placebo data. This result would suggest that individual directors are significantly related to the performance of firms they did not actually work for. Such a relation is economically meaningless, thus, we conclude that the test statistics are inflated and the tests invalid. This finding highlights the importance of benchmarking ones findings against placebo data when using the fixed effects method.

In principle, using repeated draws of placebo datasets allows learning the empirical distribution of the test statistics and, thereby, gaining insights into the economical magnitude of the findings. However, if the data at hand is sufficiently complex it is not obvious how the random datasets should be generated. We employ four different randomization procedures and show that the choice of procedure greatly influences our conclusions. When debating which mechanism is appropriate, we find it instructive to think of the mechanisms as different restrictions of the random reassignment. The more restrictive a mechanism is, the more accurately the placebo data will represent the structure of the original data. For our director data we consider four characteristics: (1) Serial correlation of director indicators. (2) Board size. (3) Busyness of directors. (4) Duration of mandates.

We identify restrictions to the randomization mechanisms that accurately preserve these characteristics. However, adding restrictions reduces the number of possible outcomes of randomization. Thereby, it becomes more likely that the placebo data is partly identical with the original data. It follows that the benchmark statistics are higher than if the data were truly placebo. For example, unrestricted randomization (i.e., recombining financial and director data on a firm year basis) would lead to the highest number of possible outcomes. Indeed, this mechanism is employed by Cronqvist and Fahlenbrach (2009) to generate a benchmark for their blockholder fixed effects. For our director data, such a benchmark would be inadequate as it disregards data structure entirely. To our knowledge, we are the first ones to illustrate this trade off when benchmarking the fixed effects results.

A follow-up study to this article will continue searching for directors that are associated with performance shifts. This search will push in three broad directions: (1) We will formulate and test hypotheses about the individual director effects. Given that the criticism of the fixed effects method only affects inference but leaves the coefficient estimates unbiased, these tests will provide further evidence on whether our estimates are more than pure noise. (2) We will assess the statistical power of combining the fixed effects method with placebo benchmarking. We try establishing a lower bound for this power by replicating Bertrand and Schoar's (2003) original work on manager fixed effects and benchmarking the findings against placebo results. The idea behind this approach is that managers are more closely involved with the firm's operations than directors and, thus, are more likely to have a tangible influence on firm performance. Thus, our method might detect manger effects even if it is not strong enough to detect director effects. (3) We will further limit our analysis to directors who are considered particularly important by the literature. This limitation should reduce the noise from tracking a potentially large number of irrelevant directors. Thus, it increases the likelihood that our analysis detects super stars or super losers if they exist.

Appendix A: Variable definitions

Variable	Definition (Compustat item names in parentheses; all amounts in 2009 USD)
Assets	Natural logarithm of the book value of assets (at). Winsorized at the 0.5% level.
Return on assets	Operating income before depreciation (oibdp) divided by the book value of total assets (at) of the previous year times 100. Winsorized at the 0.5% level.
Return on equity	Net income (ni) divided by the total common equity (ceq) of the previous year (Fahlenbrach and Stulz 2011) times 100. Winsorized at the 0.5% level.
Sales to assets	Sales (sale) devided by the book value of total assets (at) of the previous year times 100.
Tobin's Q	Market value of assets divided by the book value of assets (at), where the market value of assets is the sum of the book value of assets (at) and the market value of equity (mv) less common equity (ceq) and deferred taxes (txdb). Winsorized at the 0.5% level.
Investment	Capital expenditures (capex) divided by the previous year's net property, plant, and equipment (ppent) times 100. Winsorized at the 0.5% level.
Leverage	Total long term debt (dltt) and debt in current liabilities (dlc) divided by the sum of total long term debt (dltt), debt in current liabilies (dlc), and common equity (ceq) times 100. Winsorized at the 0.5% level.
Sales growth	Sales (sale) of current year minus sales of previous year divided by sales of previous year times 100. Winsorized at the 0.5% level.
Age (incorporation)	Number of years since the firm first incorporated. Data courtesy of Loderer, Stulz, and Waelchli (2013).
Board size	Number of directors on a firm's board in a given year.
Board independence	Fraction of independent directors (as classified by IRRC / RiskMetrics) on a firm's board.
Board tenure	The median number of years a board's directors have served at a given time. If the starting date of the mandate is missing for some directors, these directors are ignored. If the starting point is missing for all directors, Board tenure is set to zero. An indicator variable that tracks these cases is included in all regressions. Winsorized at the 0.5% level
CEO is COB	Indicator variable that is 1 in firm years where the chairman of the board (COB) serves as a chief executive (CEO) and is classified as an employee of the firm. Serving as a CEO is not sufficient, as this variable also equals 1 for outside CEOs, that is, for independent directors that serve as CEO in an outside firm.

Variable	Definition (Compustat item names in parentheses; all amounts in 2009 USD)
Board ownership (%)	Total percentage of votingpower controlled by the board members (pcnt_ctrl_votingpower). If this variable is missing for some directors, these directors are assumed to have not voting power. If the variable is missing for all directors, Board ownership is set to zero. An indicator variable that tracks these cases is included in all regressions. The same is true if the total percentage of ownership controlled is greater than 100 percent. Winsorized at the 0.5% level.
Board ownership (USD)	Board ownership (%) multiplied by the market value of equity.

Table 1
Interpretation of the director fixed effects results

		Do frictions distort firms' selection of directors?	
		no	yes
Do we observe fixed effects? (Assuming our tests have the power needed to detect them)	no	*Influence:* - All directors influence their firms optimally at all times, or - individual directors don't matter. *Certification:* - Firms do not use directors who have been successful in the past to certify their own outlook.	*Influence:* - Individual directors don't matter in a way consistent with fixed effects. *Certification:* - Firms do not use directors who have been successful in the past to certify their own outlook, or - cost of certification fully erodes outperformance.
	yes	*Influence:* - All directors influence their firms optimally at all times, or - individual directors don't matter. *Certification:* - Firms use directors who have been successful in the past to certify their own outlook.	*Influence / certification:* - Either directors influence their firms, or firms use directors to certify their own outlook, or both happen simultaneously. - Use timing to identify the causal effect at play (Cronqvist and Fahlenbrach 2009).

Table 2
Distribution of tracked directors: Directors' perspective – This table describes how board service is distributed across the sample from a directors' perspective. It does so by providing the frequencies and means of five measures for the directors' activities. Panel A counts the number of different boards the directors are tracked on. Panel B reports the number of simultaneous positions a director is tracked on for every year this director is active. Panel C shows the distribution of the number of simultaneous mandates over a directors' sample-live by reporting the minimum and maximum number of simultaneously tracked positions the director holds. Panel D counts the number of different sample years a director is tracked in. Panel E describes the duration of the tracked mandates, that is, the number of years a director is tracked on a given board.

Panel A: Number of tracked board positions per director

	Number of directors	Percent	Cumulative		
1	0	0.00	0.00		
2	1,460	70.81	70.81		
3	412	19.98	90.79		
4	135	6.55	97.33		
5	31	1.50	98.84		
6	15	0.73	99.56		
7	5	0.24	99.81		
8	3	0.15	99.95		
9	1	0.05	100.00		
Total	2,062	100.00		Mean	2.43

Panel B: Number of simultaneously held board positions per director year

	Director years	Percent	Cumulative		
1	10,576	53.51	53.51		
2	7,442	37.65	91.17		
3	1,390	7.03	98.20		
4	283	1.43	99.63		
5	51	0.26	99.89		
6	16	0.08	99.97		
7	6	0.03	100.00		
Total	19,764	100.00		Mean	1.58

Panel C: Distribution over time of simultaneously held positions per directors

	Maximum							
Minimum	1	2	3	4	5	6	7	Total
0	243	1,212	238	46	8	1	1	1,749
1	7	178	70	28	5	3	1	292
2	0	3	7	7	3	0	1	21
Total	250	1,393	315	81	16	4	3	2,062

Table 2 (continued)

Panel D: Number of years a director appears in the sample

	Directors	Percent	Cumulative		
1	0	0.00	0.00		
2	0	0.00	0.00		
3	32	1.55	1.55		
4	62	3.01	4.56		
5	131	6.35	10.91		
6	177	8.58	19.50		
7	221	10.72	30.21		
8	192	9.31	39.52		
9	198	9.60	49.13		
10	198	9.60	58.73		
11	173	8.39	67.12		
12	191	9.26	76.38		
13	174	8.44	84.82		
14	294	14.26	99.08		
15	19	0.92	100.00		
Total	2,062	100.00		Mean	9.58

Panel E: Duration of board mandates

	Frequency	Percent	Cumulative		
1	0	0.00	0.00		
2	0	0.00	0.00		
3	1,060	21.16	21.16		
4	779	15.55	36.71		
5	705	14.07	50.79		
6	581	11.60	62.39		
7	442	8.82	71.21		
8	348	6.95	78.16		
9	309	6.17	84.33		
10	209	4.17	88.50		
11	184	3.67	92.17		
12	153	3.05	95.23		
13	107	2.14	97.36		
14	132	2.64	100.00		
Total	5,009	100		Mean	6.22

Table 3
Distribution of tracked directors: Firms' perspective – The table describes the distribution of directors in the sample from the firms' perspective. Panel A reports the frequency and average of the number of directors tracked in a given firm year. Panel B shows the distribution of the number of tracked directors over a firm's sample-live. It does so by reporting the minimum and maximum number of simultaneously tracked directors in a firm over its sample years.

Panel A: Number of directors tracked per firm year

	Frequency	Percent	Cumulative		
0	4,112	27.16	27.16		
1	3,399	22.45	49.61		
2	2,564	16.94	66.55		
3	1,853	12.24	78.79		
4	1,237	8.17	86.96		
5	836	5.52	92.48		
6	541	3.57	96.06		
7	275	1.82	97.87		
8	186	1.23	99.10		
9	88	0.58	99.68		
10	38	0.25	99.93		
11	10	0.07	100.00		
Total	15,139	100.00		Mean	2.06

Panel B: Distribution over time of the number directors tracked per firm

Minimum tracked directors	Maximum tracked directors												
	0	1	2	3	4	5	6	7	8	9	10	11	Total
0	999	184	85	45	18	6	2	3	1	0	0	0	1,343
1	0	212	142	84	35	13	7	3	2	2	0	0	500
2	0	0	98	61	57	32	21	6	6	3	0	1	285
3	0	0	0	39	37	24	20	15	4	1	3	0	143
4	0	0	0	0	16	18	21	12	7	4	2	2	82
5	0	0	0	0	0	9	10	11	2	8	5	1	46
6	0	0	0	0	0	0	3	3	4	0	1	0	11
7	0	0	0	0	0	0	0	1	1	2	1	0	5
8	0	0	0	0	0	0	0	0	1	1	0	0	2
Total	999	396	325	229	163	102	84	54	28	21	12	4	2,417

Table 4
Descriptive statistics: Firm years with and without tracked directors – The table provides distributional statistics (i.e., number of observations, mean, standard deviation, first quartile, median, and third quartile) for all variables used in the regression analysis. Three additional variables are added that improve our understanding of the data (Firm age, Sales growth, and Board ownership measured in USD). For each variable separate statistics are shown for three different data samples. The first row includes the full sample of firm years. The second and third rows represent subsamples without any and with one or more tracked directors. The two subsamples' distributions are tested for differences in the last column. The first test statistic refers to a t-test and the second one to a Wilcoxon test (z-value). ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively.

	N	Mean	Std. dev.	1st quartile	Median	3rd quartile	Test stat.
Return on assets	15,139	10.53	11.19	5.21	10.18	15.97	
	4,112	9.38	13.14	3.61	9.46	16.10	-7.73***
	11,027	10.95	10.33	5.69	10.37	15.93	-6.36***
Return on equity	15,139	6.63	42.71	2.76	9.73	15.67	
	4,112	2.67	47.69	0.62	8.09	13.64	-6.98***
	11,027	8.11	40.60	3.60	10.25	16.45	-13.11***
Sales to assets	15,139	121.25	78.59	68.26	103.52	153.20	
	4,112	117.09	77.71	61.78	100.83	152.45	-3.98***
	11,027	122.80	78.86	70.20	104.36	153.56	-4.52***
Tobin's Q	15,139	1.98	1.33	1.18	1.56	2.25	
	4,112	1.96	1.37	1.14	1.53	2.21	-0.92
	11,027	1.99	1.32	1.20	1.57	2.27	-3.67***
Assets	15,139	6,559	25,160	554.94	1,331	3,870	
	4,112	1,734	5,791	344.57	658.44	1,410	-14.51***
	11,027	8,358	29,064	746.29	1,852	5,328	-40.91***
Investment	15,139	31.41	33.42	13.88	21.97	36.01	
	4,112	40.19	45.24	14.73	26.20	46.49	19.98***
	11,027	28.14	27.04	13.64	20.94	32.99	14.59***
Leverage	15,139	32.00	26.25	8.24	31.09	47.70	
	4,112	27.54	27.06	1.34	23.20	45.40	-12.82***
	11,027	33.66	25.74	12.90	33.21	48.42	-14.84***
Sales growth	12,380	6.06	22.84	-3.92	4.28	13.40	
	2,783	6.96	24.64	-3.77	4.94	15.84	2.36**
	9,597	5.80	22.29	-3.97	4.12	12.90	3.26***
Age (incorporation)	5,455	42.41	30.73	18.00	33.00	63.00	
	1,512	34.79	28.62	14.00	26.00	45.00	-11.47***
	3,943	45.33	31.01	19.00	37.00	70.00	-12.39***
Boardsize	15,139	8.99	2.43	7.00	9.00	10.00	
	4,112	7.82	2.06	6.00	8.00	9.00	-37.69***
	11,027	9.42	2.41	8.00	9.00	11.00	-37.7***
Board independence	15,139	0.66	0.18	0.56	0.69	0.80	
	4,112	0.56	0.20	0.43	0.57	0.71	-42.21***
	11,027	0.70	0.16	0.60	0.71	0.82	-37.13***

Table 4 (continued)

	N	Mean	Std. dev.	1^{st} quartile	Median	3^{rd} quartile	Test stat.
Board tenure	15,139	7.15	4.35	4.00	6.00	9.00	
	4,112	7.59	5.48	4.00	6.00	10.50	7.51***
	11,027	6.99	3.82	4.00	6.00	9.00	1.35
CEO is COB	15,139	0.61	0.49	0.00	1.00	1.00	
	4,112	0.58	0.49	0.00	1.00	1.00	-4.65***
	11,027	0.62	0.48	0.00	1.00	1.00	-4.64***
Board ownership (%)	15,139	9.39	16.49	0.00	2.30	10.10	
	4,112	13.43	19.24	1.00	5.00	18.00	18.62***
	11,027	7.88	15.07	0.00	1.90	7.60	22.36***
Board ownership (USD)	13,193	411.12	1,605	0.00	34.62	170.14	
	3,375	339.21	1,252	9.76	48.67	193.44	-3.02***
	9,818	435.84	1,709	0.00	29.26	162.46	12.14***

Table 5
The regression model – Firm performance is measured by four different variables. All measures are separately regressed on board characteristics, firm-level time varying control variables, an indicator for firm years with one or more tracked directors, and firm and year fixed effects. `Squared' refers to the squared term of the variable above. `Missing indicator' is a variable that indicates firm years where the variable above was missing and replaced by zero. Variable definitions are in Appendix A. Residuals are clustered by firms. ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively.

	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q (x 100)
Board size	0.11	1.37	0.81	-8.67**
squared	0.00	-0.05	0.01	0.35*
Fraction of independent directors	3.74	1.97	-5.15	21.32
squared	-3.14	2.41	6.71	-21.74
Board tenure	0.15	0.50	0.10	0.53
squared	-0.01	-0.02	-0.01	-0.01
missing indicator	-2.35	-18.94**	-14.84*	0.11
CEO is COB	0.70**	1.40	3.47***	3.73
Board ownership	-0.05*	-0.01	-0.15	0.19
squared	0.00	0.00	0.00*	-0.00
missing indicator	-0.55	1.65	-9.35*	9.05
Assets	-4.41***	-8.35***	-37.88***	-68.59***
Investment	0.01	-0.07***	0.05***	0.22***
Leverage	-0.04***	0.03	-0.20***	-0.35***
Firm year tracked indicator	0.12	-0.59	2.84	6.47
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Adjusted R^2	0.61	0.24	0.88	0.67
N	15,139	15,139	15,139	15,139

Table 6
Director fixed effects – The regressions estimated follow the specifications introduced in Table 4. The specification differs in that the indicators for firm years with tracked directors are replaced by the full set of director indicators. This change increases the number of regression variables to a level that makes reporting individual coefficients impractical. Instead, Panel A groups coefficients and reports F-values from testing their joint significance. ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively. Panel B provides additional information on the estimated individual director coefficients. 'N (directors)' is the number of director effects included in the regression. The distribution of the estimates is described by its quartiles. To facilitate comparison, 'Median (y)' provides the median of the dependent variable. 'Individually significant coefficients' refers to the fraction of coefficients significant at the usual confidence levels.

Panel A: Fixed effects regressions

	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q
Individual director effects	69,789***	4,309***	4,600,000***	770,000***
Board characteristics	1.81**	0.70	2.90***	1.33
X_{it}	50.74***	10.07***	153.25***	53.43***
Firm fixed effects	140,000***	310,000***	130,000,000***	6,400,000***
Year fixed effects	13.12***	6.71***	18.69***	32.48***
Adjusted R^2	0.66	0.26	0.91	0.71
Increase in adjusted R^2	0.05	0.02	0.03	0.05
N (firm years)	15,139	15,139	15,139	15,139

Panel B: Director fixed effects coefficients

	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q
N (directors)	1,854	1,854	1,854	1,854
Distribution of coefficients				
1st quartile	-3.12	-10.52	-12.82	-0.32
Median	0.02	0.14	0.38	0.00
3rd quartile	3.07	12.87	14.05	0.35
Median (y)	10.18	9.73	103.53	1.56
Individually significant coefficients				
α = 0.01	25.35%	16.02%	28.21%	27.24%
α = 0.05	35.38%	25.19%	38.08%	36.03%
α = 0.1	41.86%	32.36%	45.04%	43.15%

Table 7
Serial correlation of residuals – The sample is the residuals from the regressions in Table 5. Serial correlation is evaluated by regressing each residual on its one, two, and three year lags. The table reports the regressions coefficients, the adjusted R square, the p-value of the test of overall significance of the model, and the number of observations. ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively.

	Return on assets			Return on equity			Sales to assets			Tobin's Q		
1st lag	0.19***	0.29***	0.30***	-0.21***	-0.19***	-0.21***	0.15***	0.23***	0.24***	0.12***	0.21***	0.22***
2nd lag		-0.23***	-0.16***		-0.28***	-0.29***		-0.17***	-0.11***		-0.13***	-0.08***
3rd lag			-0.17***			-0.27***			-0.17***			-0.16***
R^2 adjusted	0.04	0.11	0.16	0.04	0.08	0.11	0.02	0.07	0.11	0.02	0.06	0.11
Prob > F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	12,381	10,123	8,264	12,381	10,123	8,264	12,381	10,123	8,264	12,381	10,123	8,264

Table 8
Director fixed effects, controlling for the lag of the dependent variable – The performance regressions are estimated and presented analogously to Table 5. The specification is altered by including the 1 year lag of the dependent variable as an additional explanatory variable. The table shows how this change affects the F-value associated with the director effects, the Adjusted R square, the number of effects estimated, the IQR of the estimated coefficients, and the fraction of effects that is individually significant at the 1 percent confidence level (italic print).

Panel A: Fixed effects regressions				
	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q
Individual director effects	900,000***	6,792***	43,550***	49,682***
increase from Table 5	*830,211*	*2,483*	*-4,556,449*	*-720,318*
Board characteristics	0.98	0.76	2.75***	1.24
X_{it}	57.77***	9.90***	160.05***	34.12***
Firm fixed effects	380,000***	670,000***	250,000***	350,000***
Year fixed effects	11.78***	6.43***	18.71***	34.89***
Adjusted R^2	0.71	0.26	0.92	0.73
increase from Table 5	*0.05*	*0.00*	*0.01*	*0.02*
N (firm years)	15,139	15,139	15,139	14,905

Panel B: Director fixed effects coefficients				
	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q
N (directors)	1,854	1,854	1,854	1,846
increase from Table 5	*0.00*	*0.00*	*0.00*	*-8.00*
Median (y)	10.18	9.73	103.53	1.56
Dist. of coef. (IQR)	5.05	24.30	22.69	0.57
increase from Table 5	*-1.14*	*0.91*	*-4.18*	*-0.10*
Ind. sign. coef. (α=0.01)	25.62%	16.02%	27.72%	25.19%
increase from Table 5	*0.27%*	*0.00%*	*-0.49%*	*-2.05%*

Table 9
Serial correlation of residuals, controlling for the lag of the dependent variable – The sample is the residuals from the regressions in Table 7. Serial correlation is evaluated by regressing each residual on its one, two, and three year lags. The table reports the regressions coefficients, the adjusted R square, the p-value of the test of overall significance of the model, and the number of observations. ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively.

	Return on assets			Return on equity			Sales to assets			Tobin's Q		
1st lag	-0.08***	-0.03***	-0.01	-0.16***	-0.12***	-0.14***	-0.05***	-0.01	0.01	-0.06**	0.01	0.04
2nd lag		-0.18***	-0.15***		-0.29***	-0.28***		-0.11***	-0.09***		-0.09***	-0.06***
3rd lag			-0.17***			-0.27***			-0.16***			-0.16***
R^2 adjusted	0.01	0.04	0.06	0.02	0.07	0.11	0.02	0.07	0.11	0.00	0.01	0.05
Prob > F	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
N	12,381	10,123	8,264	12,380	10,122	8,263	12,381	10,123	8,264	12,187	9,966	8,142

Table 10

Real versus placebo director fixed effects – The table repeats the analyses from Table 5 on two different sets of placebo data. The datasets referred to as board and director level data and are based on different randomization mechanisms. Board-level placebo data is generated by reassigning the complete set of director indicators from one firm to another. That is, directors who serve on the same board in real data will still serve on the same board after randomization. The mechanism that generates director-level placebo data reassigns each director individually to a random firm. Both sets of placebo data have in common that the years in which directors are active are not affected by randomization. Thus, existing boards are torn apart. The table presents the results by showing F-values associated with the significance of groups of effects in Panel A. ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively. The individual director effects estimated are discussed in Panel B by providing the number of effects estimated, the IQR of the coefficients, and the fraction of individually significant coefficients at the 1 percent confidence level. The changes resulting from using placebo instead of real data are highlighted by adding a comparison of the results with the ones from Table 5 (italic print).

Panel A: Fixed effects regressions

	Board-level				Director-level			
	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q
Individual director effects	33,178***	27,380***	110,000***	64,404***	19,903***	660,000***	6,238***	5,122***
increase from Table 5	*-36,612*	*23,071*	*-4,490,000*	*-705,596*	*-49,886*	*655,691*	*-4,593,761*	-764,878
Board characteristics	1.94**	0.97	3.00***	1.13	1.68*	0.90	2.89***	1.01
X_{it}	37.08***	9.91***	137.05***	70.67***	38.78***	10.11***	118.25***	74.55***
Firm fixed effects	1,500,000***	18,000,000***	6,800,000***	3,900,000***	1,600,000***	3,700,000***	73,749***	46,870***
Year fixed effects	18.03***	7.24***	21.37***	44.16***	16.10***	7.68***	20.09***	40.71***
Adjusted R^2	0.63	0.24	0.89	0.69	0.63	0.25	0.89	0.69
increase from Table 5	*-0.02*	*-0.02*	*-0.02*	*-0.02*	*-0.03*	*-0.01*	*-0.02*	*-0.02*
N (firm years)	*15,139*	15,139	15,139	15,139	15,139	15,139	15,139	15,139

Panel B: Director fixed effects coefficients

	Board-level				Director-level			
	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q	Return on assets (%)	Return on equity (%)	Sales to assets (%)	Tobin's Q
N (directors)	651	651	651	651	647	647	647	647
increase from Table 5	*-1,203*	*-1,203*	*-1,203*	*-1,203*	*-1,207*	*-1,207*	*-1,207*	*-1,207*
Dist, of coef, (IQR)	7.51	17.33	28.67	0.82	6.19	15.10	21.48	0.62
increase from Table 5	*1.32*	*-6.07*	*1.80*	*0.15*	*0.00*	*-8.29*	*-5.39*	*-0.05*
Ind. sign. coef. (α=0.01)	36.25%	28.88%	36.71%	36.56%	35.86%	26.89%	32.61%	32.30%
increase from Table 5	*10.90%*	*12.86%*	*8.50%*	*9.32%*	*10.51%*	*10.87%*	*4.40%*	*5.06%*

Table 11

Benchmarking of the real director fixed effects against repeated draws of the placebo effects – The table evaluates whether the results from real data differ from placebo results in repeated draws. ‘Normal’ refers to the estimates from Table 5. ‘Subset’ estimates the effects on real data including only a random subset of the director indicators (33 percent). ‘Board’ and ’Director’ refer to placebo data analyses based on datasets generated using the respective mechanism introduced in Table 9. The table shows the number of random draws for each of the datasets and the average number of directors included in the estimates of each draw. It describes how the F-values associated with joint significance of the director indicators are distributed across the draws. ‘KS test’ and ‘Wilcoxon’ represent Kolmogorov-Smirnov and Wilcoxon tests for differences in the distribution of F-values between the respective placebo results and the real results represent by ‘Subset’. Further, the table presents the adjusted R square averaged over draws. The last four columns describe the distribution of the estimated individual director fixed effects coefficients. Besides the IQR and the Median of the coefficients, it provides Kolmogorov-Smirnov (KS test) and Wilcoxon test statistics for differences in the distribution. ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively.

	N of draws	N of directors	F-values					Adjusted R^2 in %	Director fixed effects coefficients			
			1st quartile	Median	3rd quartile	KS test	Wilcoxon		IQR	Median	KS test	Wilcoxon
Return on assets												
Normal	1	1,854		69,789				65.56	6.19	0.02		
Subset	100	619	9,154	21,251	73,005			62.67	5.98	0.04		
Board	100	608	18,703	62,336	261,731	0.25***	-3.56***	62.62	7.22	0.02	0.05***	-1.36
Director	100	620	3,070	11,719	52,271	0.24***	2.92***	62.97	6.05	0.03	0.01**	-1.88*
Return on equity												
Normal	1	1,854		4,309				25.82	23.39	0.14		
Subset	100	619	7,732	29,815	133,966			24.80	17.35	0.10		
Board	100	608	27,279	83,618	298,112	0.32***	-3.84***	24.49	18.57	0.14	0.02***	0.98
Director	100	620	14,491	49,601	339,100	0.16	-2.10**	24.59	14.46	0.20	0.04***	0.02
Sales to assets												
Normal	1	1,854		4,600,000				90.77	26.87	0.38		
Subset	100	619	5,973	22,330	210,429			89.21	25.51	0.77		
Board	100	608	25,245	101,534	395,071	0.31***	-3.89***	88.83	27.82	0.35	0.03***	3.24***
Director	100	620	2,944	10,178	35,483	0.20**	2.67***	88.99	23.71	0.22	0.03***	4.60***
Tobin’s Q												
Normal	1	1,854		770,000				71.25	0.66	0.00		
Subset	100	619	5,312	15,788	70,214			68.80	0.64	0.01		
Board	100	608	14,939	43,187	192,038	0.29***	-3.95***	68.56	0.73	0.00	0.04***	2.27**
Director	100	620	3,866	13,933	91,579	0.10	0.48	68.85	0.67	-0.01	0.02***	3.93***

Table 12
Benchmarking of the real director fixed effects against the placebo effects resulting from restricted reassignment – The table introduces an alternative randomization mechanism. This mechanism restricts the reassignment of each director or board to firms that are active in the exact same years as the original firm. The mechanism can be applied to the board or director level (see description of Table 9). ‘Board’ and ’Director’ refer to placebo data analyses based on 100 draws of these alternative mechanisms. ‘Normal’ refers to the estimates from Table 5. The table shows the number of random draws for each of the datasets and the average number of directors included in the estimates of each draw. It provides the F-value testing for the joint significant of director effects for the real data analysis. For the placebo results, the fraction of draws that produce a higher F-value than the real estimate is presented (pseudo p). Further, the table presents the adjusted R square averaged over draws. The last four columns describe the distribution of the individual director fixed effects coefficients. Besides the IQR and the Median of the coefficients, it provides Kolmogorov-Smirnov (KS test) and Wilcoxon test statistics for differences in the distribution between real and placebo estimates. ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively.

					Director fixed effects coefficients			
	N of draws	N of directors	F-value / pseudo p	Adjusted R^2	IQR	Median	ks test	Wilcoxon
Return on assets								
Normal	1	1,854	69,789	65.56	6.19	0.02		
Board	100	1,613	0.46	66.62	6.85	0.01	0.03	0.11
Director	100	1,867	0.58	67.02	6.64	0.05	0.02	-0.04
Return on equity								
Normal	1	1,854	4,309	25.82	23.39	0.14		
Board	100	1,613	0.98	24.92	20.18	0.12	0.04***	1.17
Director	100	1,867	0.99	24.83	18.85	0.15	0.05***	0.93
Sales to assets								
Normal	1	1,854	4,600,000	90.77	26.87	0.38		
Board	100	1,613	0.03	90.54	27.32	0.49	0.01	0.04
Director	100	1,867	0.07	90.67	26.16	0.48	0.01	0.14
Tobin’s Q								
Normal	1	1,854	770,000	71.25	0.66	0.00		
Board	100	1,613	0.17	72.03	0.70	0.00	0.02	0.18
Director	100	1,867	0.17	72.48	0.67	0.01	0.02	0.13

Figure 1
Board size and Tobin's Q – The three lines in this figure illustrate the marginal effects implied in three independent estimates of the relation between Board size (i.e., number of directors) and Tobin's Q: the solid line represents our own polynomial estimate from Table 4; the dashed line results from Yermack's (1996) estimate based on a logarithmic relation (Tobin's Q = –0.337 x ln(Board size)); and the dotted line is from Faleye (2007) and results from a linear relation (Tobin's Q = –0.015 x Board size). The bars show the relative frequency of board sizes in our data. Note that the axis corresponding to these bars is not shown. Instead, the white and gray bars roughly indicate the central 90% of observations and 5% upper and lower extremes, respectively.

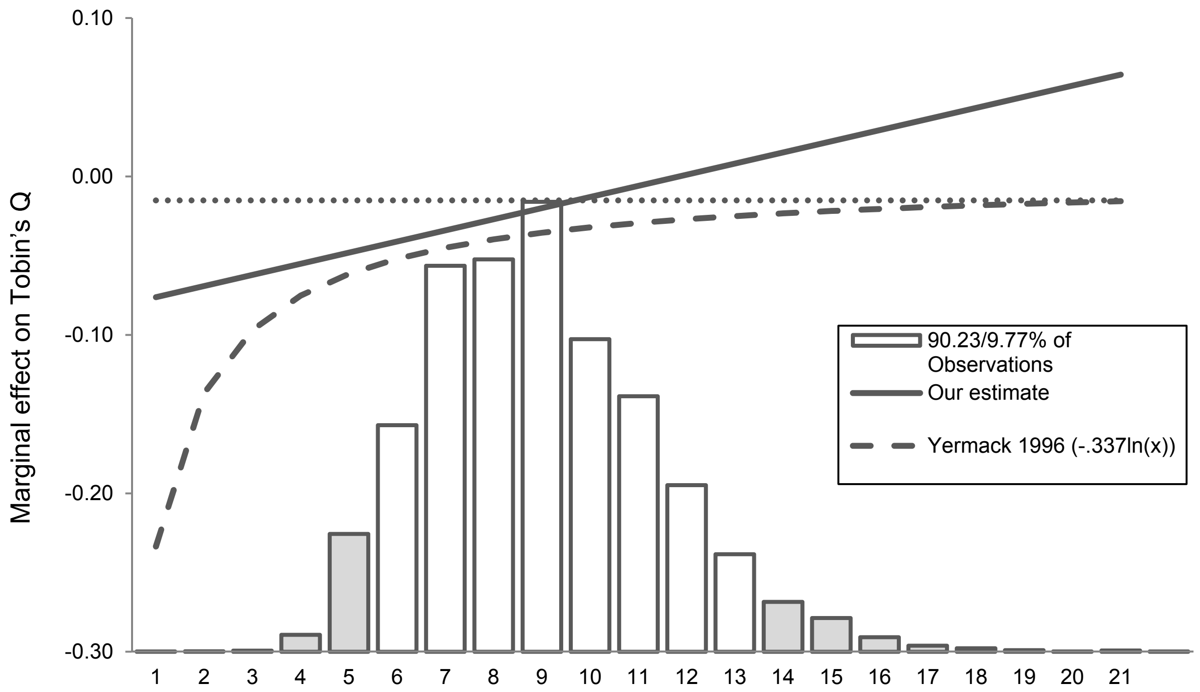

Figure 2
Differences between real and placebo results with growing number of random draws – Table 10 presents a snapshot of how the results based on real (subset) and placebo data compare after 100 draws. This figure shows how these differences evolve as the number of draws grows from 1 to 100. The left column compares distributions of F-values and the right column compares distributions of individual director coefficients. The solid lines are p-values from the Kolmogorov-Smirnov test for equality of distributions. These lines use the primary vertical axis. The dashed lines and the secondary vertical axis represent the differences between real and placebo results in the median F-value and the IQR of the estimated coefficients, respectively. Positive differences indicate that real data results are stronger than placebo results while negative values indicate the opposite. The dark lines compare subset and board-level placebo results, whereas the light lines compare subset and director-level placebo results.

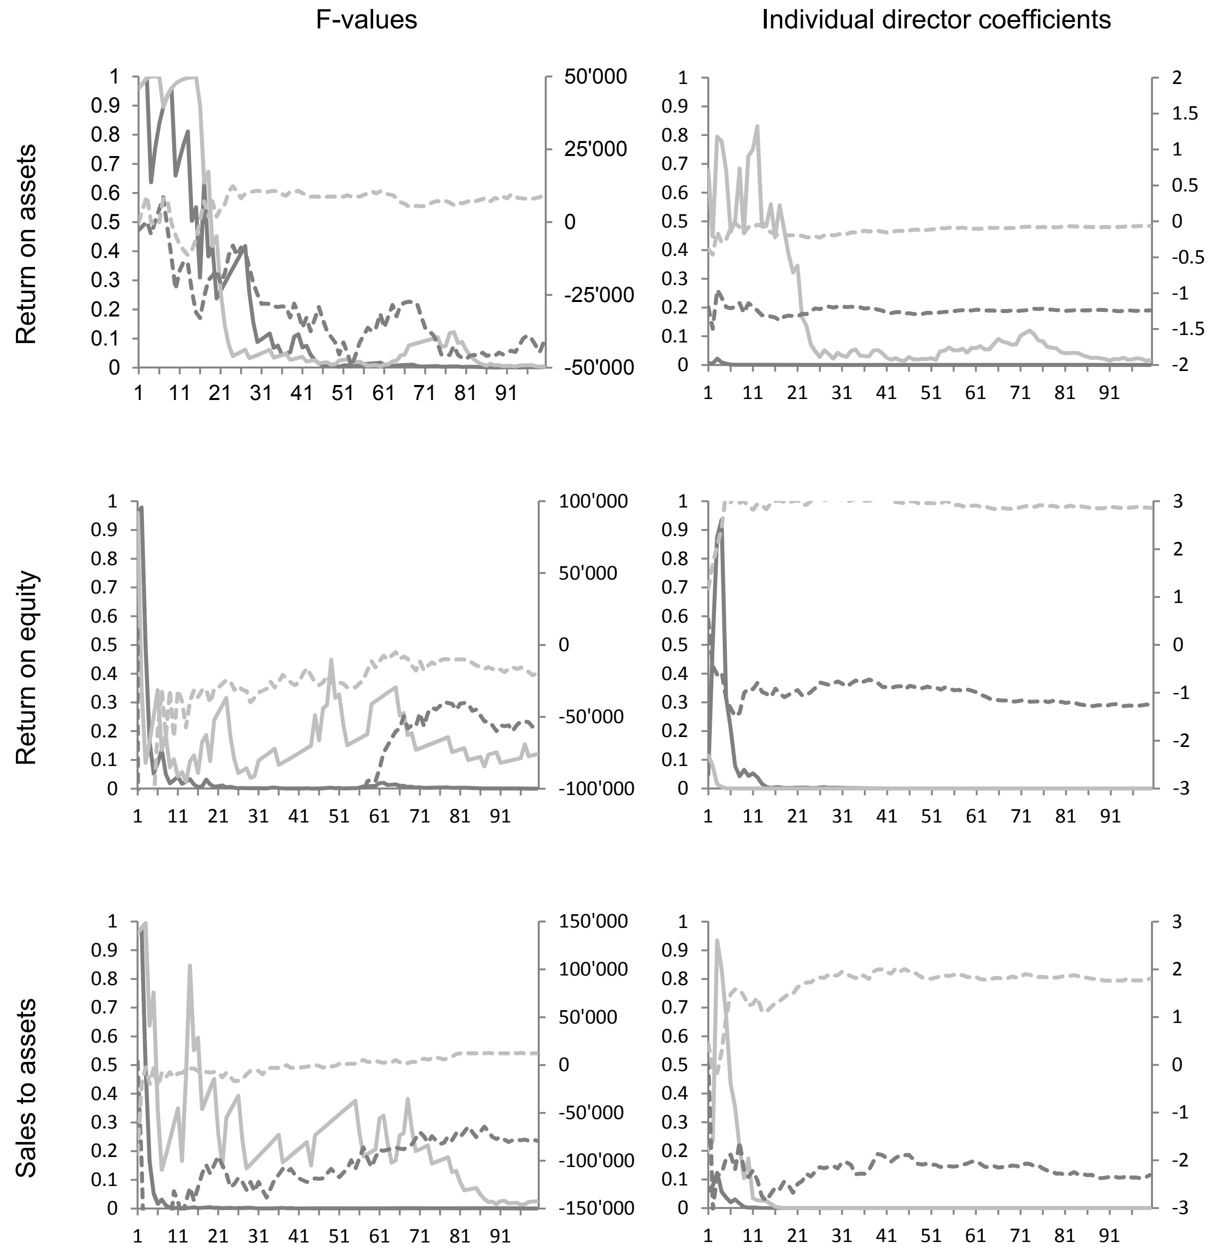

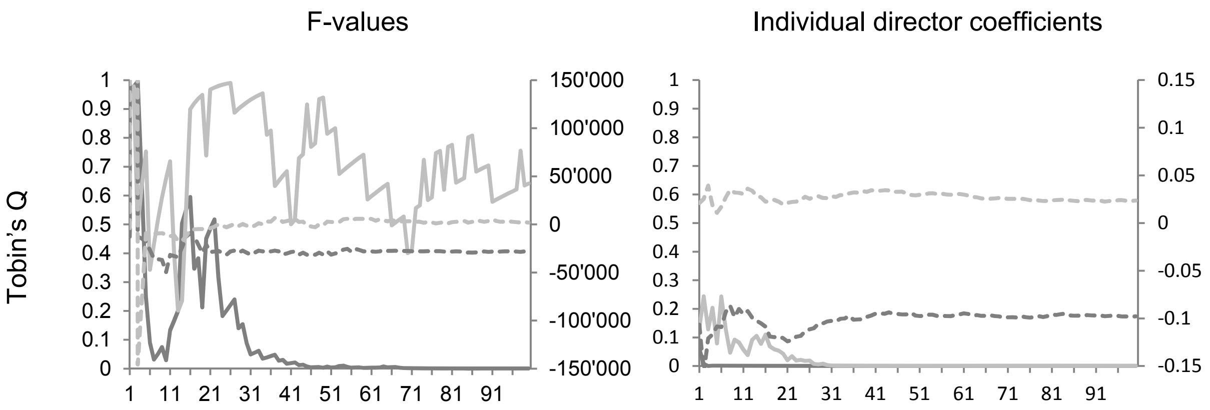
F-values
Individual director coefficients
Tobin's Q
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
150'000
100'000
50'000
0
-50'000
-100'000
-150'000
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
1 11 21 31 41 51 61 71 81 91

Figure 3
Characteristics of real and placebo data – The figure compares the real director data to the placebo data used for the analysis in Table 10. The placebo data is generated by reassigning whole boards (squares) or individual directors (triangles) to new firms. The randomization mechanism moves the boards or the directors en bloc, that is, without changing the years active. No further restrictions are employed with regard to the firms considered for reassignment. The bars in the first column represent the full sample real data. In the second column, the bars represent real data when taking random subsets of directors. Three characteristics of the data are considered: the number of tracked directors per firm year (Panel A), the number of board position a director is tracked in (Panel B), and the duration of mandates, that is, director-firm combinations (Panel C). Relative frequencies of the values are shown.

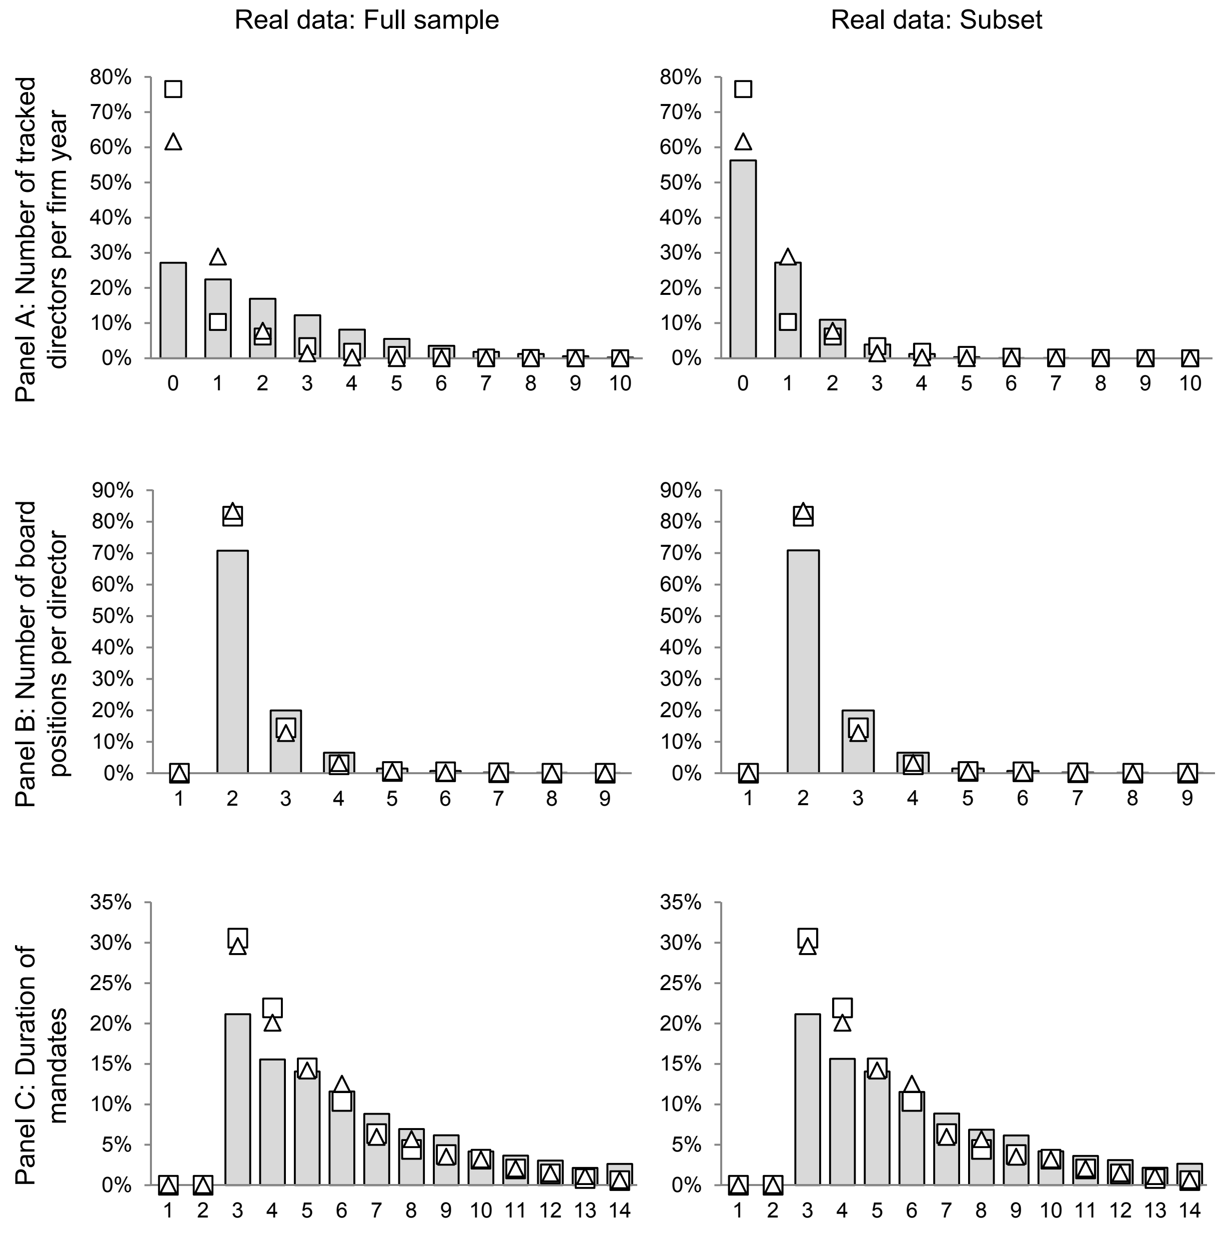

Figure 4
Characteristics of real and placebo data (restricted reassignment) – This figure compares the real director data to the placebo data used for the analysis in Table 11. The placebo data is generated by reassigning boards (squares) or directors (triangles) to new firms. The randomization mechanism moves the board or directors en bloc, that is, without changing the years active. Only firms that are active in exactly the same years as the original firms are considered for reassignment. The real data is the full sample (bars). Three characteristics of the data are considered: the number of tracked directors per firm year (Panel A), the number of board position a director is tracked in (Panel B), and the duration of mandates, that is, director-firm combination (Panel C). Relative frequencies of the values are shown.

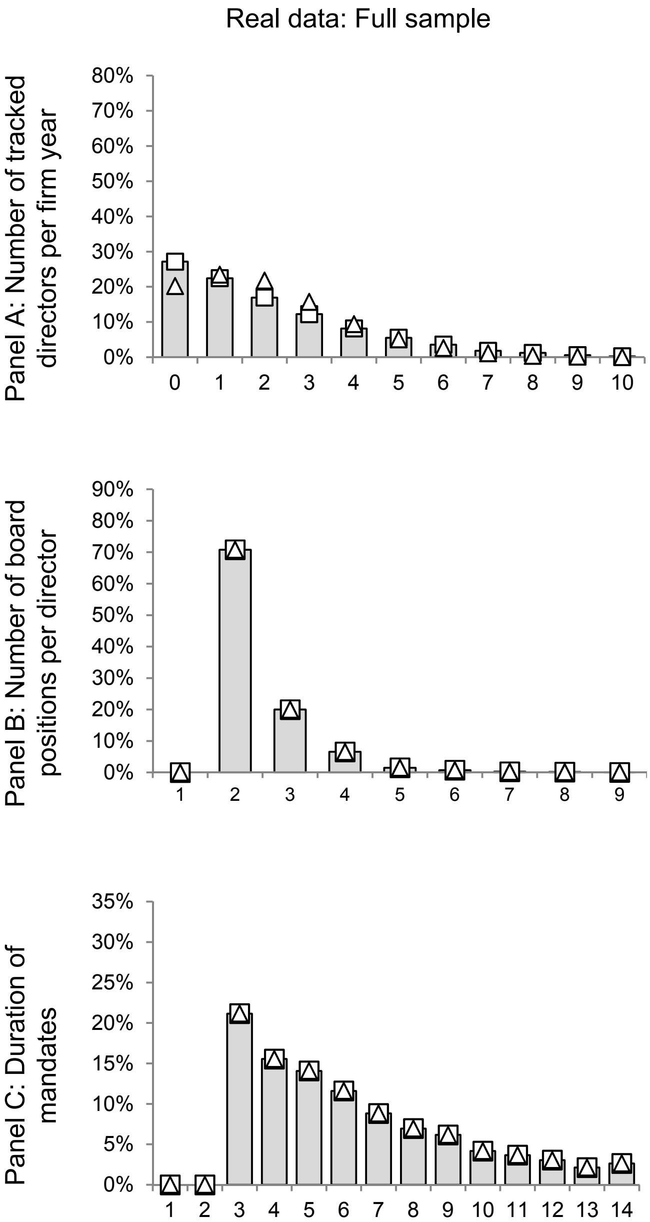

Figure 5
Differences between real and placebo results with growing number of random draws (restricted reassignment) – Table 11 presents a snapshot of how the results based on real and placebo data compare after 100 draws. This figure shows how these differences evolve as the number of draws grows from 1 to 100. The left column compares distributions of F-values and presents the respective pseudo p-values, that is, the number of random draws which lead to a higher F-value than the real data. The right column compares distributions of individual director coefficients. The solid lines are p-values from the Kolmogorov-Smirnov test for equality of distributions. These lines use the primary vertical axis. The dashed lines and the secondary vertical axis represent the differences between real and placebo results in the IQR of the estimated coefficients. Positive differences indicate that real data results are greater than placebo results while negative values indicate the opposite. The dark lines compare subset and board-level placebo results, whereas the light lines compare subset and director-level placebo results.

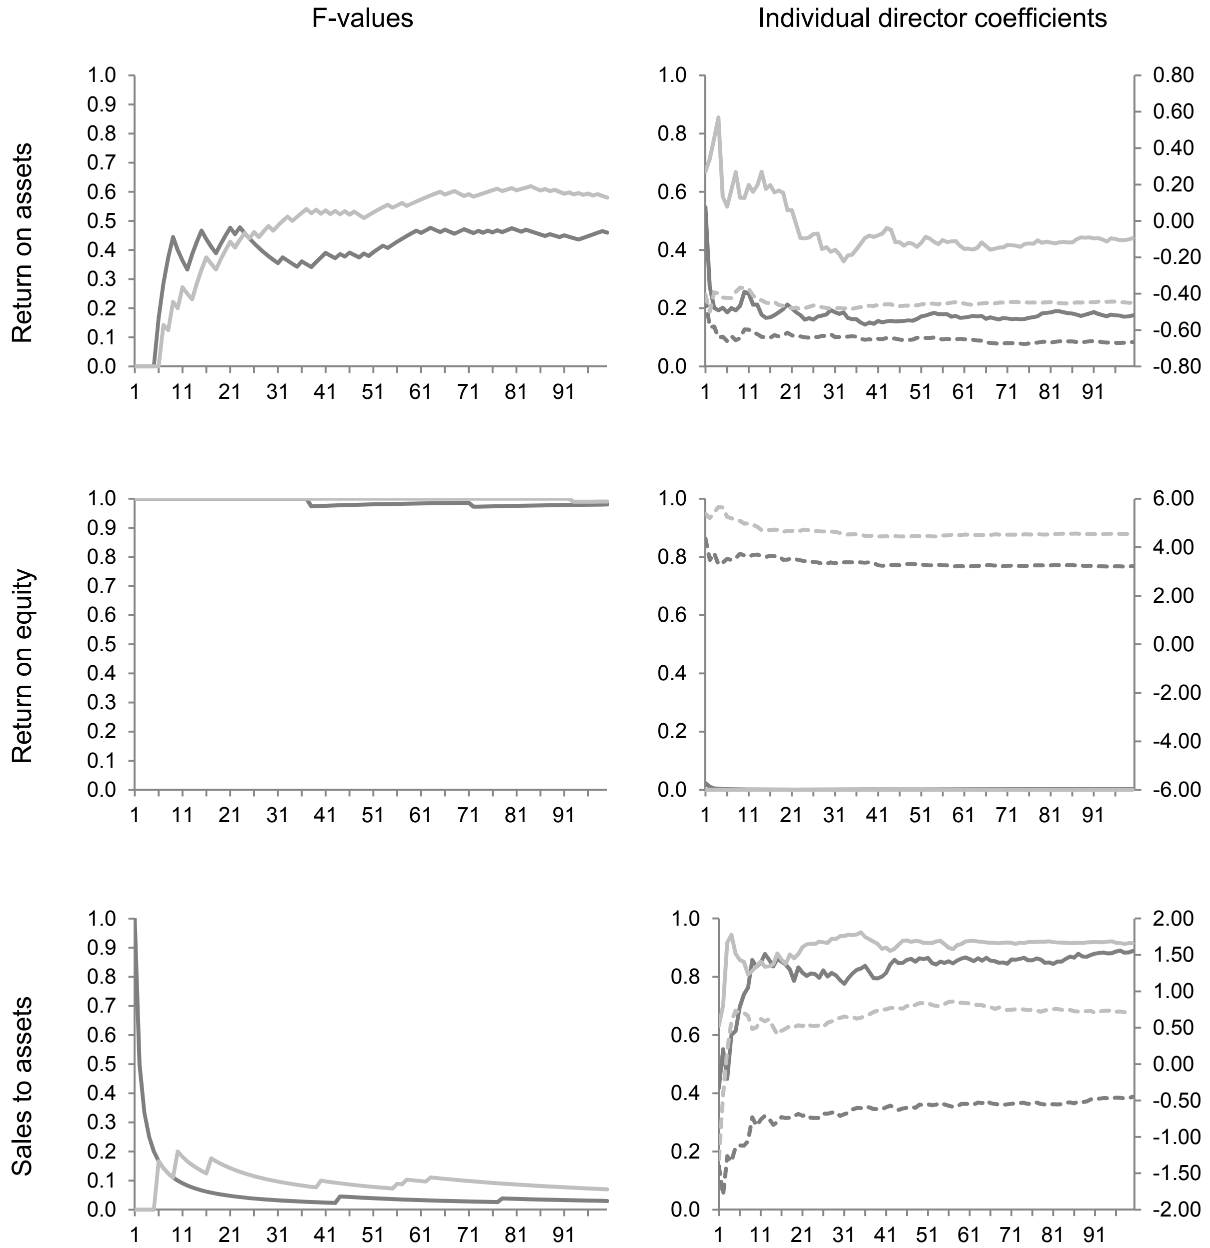

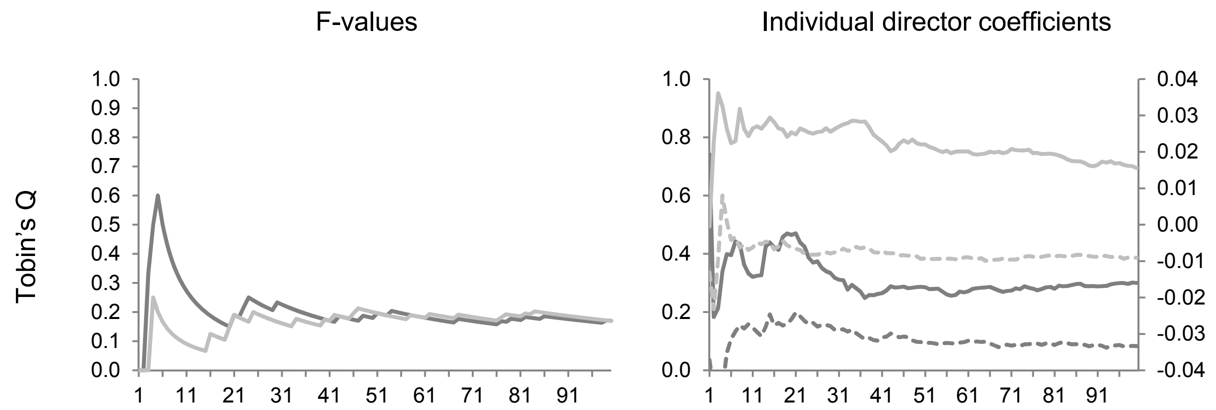
F-values
Individual director coefficients
Tobin's Q
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
1 11 21 31 41 51 61 71 81 91
0.04
0.03
0.02
0.01
0.00
-0.01
-0.02
-0.03
-0.04

References

Adams, R. B. and D. Ferreira, 2009, Women in the boardroom and their impact on governance and performance, *Journal of Financial Economics* 94, 291-309.

Adams, R. B., B. E. Hermalin, and M. S. Weisbach, 2010, The role of boards of directors in corporate governance: A conceptual framework and survey, *Journal of Economic Literature* 48, 58-107.

Agrawal, A. and S. Chadha, 2005, Corporate governance and accounting scandals, *Journal of Law and Economics* 48, 371-406.

Agrawal, A. and Charles R. Knoeber, 2001, Do some outside directors play a political role?, *Journal of Law and Economics* 44, 179-198.

Anderson, R. C., A. Duru, and D. M. Reeb, 2009, Founders, heirs, and corporate opacity in the united states, *Journal of Financial Economics* 92, 205-222.

Anderson, R. C., S. A. Mansi, and D. M. Reeb, 2004, Board characteristics, accounting report integrity, and the cost of debt, *Journal of Accounting and Economics* 37, 315-342.

Anderson, R. C. and D. M. Reeb, 2003, Founding-family ownership and firm performance: Evidence from the s&p 500, *The Journal of Finance* 58, 1301-1327.

Anderson, R. C. and D. M. Reeb, 2004, Board composition: Balancing family influence in s&p 500 firms, *Administrative Science Quarterly* 49, 209-237.

Ang, J. S., R. A. Cole, and J. W. Lin, 2000, Agency costs and ownership structure, *The Journal of Finance* 55, 81-106.

Bebchuk, L. A. and A. Cohen, 2005, The costs of entrenched boards, *Journal of Financial Economics* 78, 409-433.

Bertrand, M., E. Duflo, and S. Mullainathan, 2004, How much should we trust differences-in-differences estimates?, *Quarterly Journal of Economics* 119, 249-275.

Bertrand, M. and A. Schoar, 2003, Managing with style: The effect of managers on firm policies, *Quarterly Journal of Economics* 118, 1169-1208.

Bhagat, S. and B. Black, 1999, The uncertain relationship between board composition and firm performance, *Business Lawyer* 54, 921.

Bhagat, S. and B. Black, 2002, The non-correlation between board independence and long-term firm performance, *Journal of Corporation Law* 27, 43.

Bizjak, J., M. Lemmon, and R. Whitby, 2009, Option backdating and board interlocks, *The Review of Financial Studies* 22, 4821-4847.

Brickley, J. A., J. L. Coles, and G. Jarrell, 1997, Leadership structure: Separating the ceo and chairman of the board, *Journal of Corporate Finance* 3, 189-220.

Bryan, S. H. and A. Klein, 2004, Non-management director options, board characteristics, and future firm investments and performance, Law and Economics Research Paper Series, New York University.

Byrd, D. T. and M. S. Mizruchi, 2005, Bankers on the board and the debt ratio of firms, *Journal of Corporate Finance* 11, 129-173.

Coles, J. L., N. D. Daniel, and L. Naveen, 2008, Boards: Does one size fit all, *Journal of Financial Economics* 87, 329-356.

Cronqvist, H. and R. Fahlenbrach, 2009, Large shareholders and corporate policies, *Review of Financial Studies* 22, 3941-3976.

Defond, M. L., R. N. Hann, and X. Hu, 2005, Does the market value financial expertise on audit committees of boards of directors?, *Journal of Accounting Research* 43, 153-193.

Donaldson, L. and J. H. Davis, 1991, Stewardship theory or agency theory: Ceo governance and shareholder returns, *Australian Journal of Management* 16, 49-64.

Fahlenbrach, R., A. Low, and R. M. Stulz, 2010, Why do firms appoint ceos as outside directors?, *Journal of Financial Economics* 97, 12-32.

Fahlenbrach, R. and R. M. Stulz, 2011, Bank ceo incentives and the credit crisis, *Journal of Financial Economics* 99, 11-26.

Faleye, O., 2007, Classified boards, firm value, and managerial entrenchment, *Journal of Financial Economics* 83, 501-529.

Fama, E. F. and M. C. Jensen, 1983, Separation of ownership and control, *Journal of Law & Economics* 26, 301-326.

Fee, C. E., C. J. Hadlock, and J. R. Pierce, 2013, Managers with and without style: Evidence using exogenous variation, *Review of Financial Studies* 26, 567-601.

Ferris, S. P., M. Jagannathan, and A. C. Pritchard, 2003, Too busy to mind the business? Monitoring by directors with multiple board appointments, *The Journal of Finance* 58, 1087-1112.

Fich, E. M., 2005, Are some outside directors better than others? Evidence from director appointments by fortune 1000 firms, *Journal of Business* 78, 1943-1971.

Gilson, S. C., 1990, Bankruptcy, boards, banks, and blockholders: Evidence on changes in corporate ownership and control when firms default, *Journal of Financial Economics* 27, 355-387.

Goldman, E., J. Rocholl, and J. So, 2009, Do politically connected boards affect firm value?, *The Review of Financial Studies* 22, 2331-2360.

Güner, A. B., U. Malmendier, and G. Tate, 2008, Financial expertise of directors, *Journal of Financial Economics* 88, 323-354.

Hallock, K. F., 1997, Reciprocally interlocking boards of directors and executive compensation, *Journal of Financial and Quantitative Analysis* 32, 331-344.

Hermalin, B. E. and M. S. Weisbach, 1998, Endogenously chosen boards of directors and their monitoring of the ceo, *American Economic Review* 88, 96-118.

Keele, L. and N. J. Kelly, 2006, Dynamic models for dynamic theories: The ins and outs of lagged dependent variables, *Political Analysis* 14, 186-205.

Loderer, C. and U. Peyer, 2002, Board overlap, seat accumulation and share prices, *European Financial Management* 8, 165.

Loderer, C., R. M. Stulz, and U. Waelchli, 2013, Limited managerial attention and corporate ageing, Working Paper.

Malmendier, U. and G. Tate, 2009, Superstar ceos, *Quarterly Journal of Economics* 124, 1593-1638.

Masulis, R. W., C. Wang, and F. Xie, 2012, Globalizing the boardroom—the effects of foreign directors on corporate governance and firm performance, *Journal of Accounting and Economics* 53, 527-554.

Morck, R., A. Shleifer, and R. W. Vishny, 1988, Management ownership and market valuation: An empirical analysis, *Journal of Financial Economics* 20, 293-315.

Nguyen, B. D. and K. M. Nielsen, 2010, The value of independent directors: Evidence from sudden deaths, *Journal of Financial Economics* 98, 550-567.

Pérez-González, F., 2006, Inherited control and firm performance, *American Economic Review* 96, 1559-1588.

Perry, T. and U. Peyer, 2005, Board seat accumulation by executives: A shareholder's perspective, *The Journal of Finance* 60, 2083-2123.

Rosenstein, S. and J. G. Wyatt, 1990, Outside directors, board independence, and shareholder wealth, *Journal of Financial Economics* 26, 175-191.

Shivdasani, A. and D. Yermack, 1999, Ceo involvement in the selection of new board members: An empirical analysis, *The Journal of Finance* 54, 1829.

Vafeas, N., 1999, Determinants of the adoption of director incentive plans, *Journal of Accounting, Auditing & Finance* 14, 453-474.

Villalonga, B. and R. Amit, 2006, How do family ownership, control and management affect firm value?, *Journal of Financial Economics* 80, 385-417.

von Meyerinck, F., D. Oesch, and M. Schmid, 2012, The value of director industry experience.

Wälchli, U., 2008, The causes and consequences of ceo, cob, and board turnover, Working Paper, University of Bern.

Walkling, R. A. and M. S. Long, 1984, Agency theory, managerial welfare, and takeover bid resistance, *The RAND Journal of Economics* 15, 54-68.

Wooldridge, J. M., 2002, *Econometric analysis of cross section and panel data*, MIT Press, Cambridge, Massachusetts.

Yermack, D., 1996, Higher market valuation of companies with a small board of directors, *Journal of Financial Economics* 40, 185-211.

Yermack, D., 2004, Remuneration, retention, and reputation incentives for outside directors, *The Journal of Finance* 59, 2281-2308.

Yermack, D., 2006, Board members and company value, *Financial Markets and Portfolio Management* 20, 33-47.