

Do individual directors matter?

Inaugural dissertation submitted by Markus Senn in fulfillment of the requirements for the degree of Doctor rerum oeconomicarum at the Faculty of Business, Economics and Social Sciences of the University of Bern.

Submitted by

Markus Senn

from Geltwil, AG

2014

Original document saved on the web server of the University Library of Bern



This work is licensed under a Creative Commons Attribution-Non-Commercial-No derivative works 2.5 Switzerland licence. To see the licence go to <http://creativecommons.org/licenses/by-nc-nd/2.5/ch/> or write to Creative Commons, 171 Second Street, Suite 300, San Francisco, California 94105, USA.

Copyright Notice

This document is licensed under the Creative Commons Attribution-Non-Commercial-No derivative works 2.5 Switzerland.

<http://creativecommons.org/licenses/by-nc-nd/2.5/ch/>

You are free:



to copy, distribute, display, and perform the work

Under the following conditions:



Attribution. You must give the original author credit.



Non-Commercial. You may not use this work for commercial purposes.



No derivative works. You may not alter, transform, or build upon this work.

For any reuse or distribution, you must take clear to others the license terms of this work.

Any of these conditions can be waived if you get permission from the copyright holder.

Nothing in this license impairs or restricts the author's moral rights according to Swiss law.

The detailed license agreement can be found at:

<http://creativecommons.org/licenses/by-nc-nd/2.5/ch/legalcode.de>

The faculty accepted this work as dissertation on December 18, 2014 at the request of the two advisors Prof. Dr. Claudio Loderer and Prof. Dr. Dennis Sheehan, without wishing to take a position on the view presented therein.

Acknowledgments

« Nous sommes comme des nains juchés sur des épaules de géants (les Anciens), de telle sorte que nous puissions voir plus de choses et de plus éloignées que n'en voyaient ces derniers. Et cela, non point parce que notre vue serait puissante ou notre taille avantageuse, mais parce que nous sommes portés et exhaussés par la haute stature des géants »

Bertrand de Chartres

If my dissertation sees a little further, this success is also owed to my professional and personal environment. First, my gratitude goes to my advisor Claudio Loderer. 'Capo' inspired me in many ways. He sparked my interest in corporate finance and my passion for academic research. Whenever I was getting lost, he reliably brought me back on track.

In my years as a doctoral student, I had the privilege to meet, work, and debate with many great people at the IFM and elsewhere. My research has benefitted from uncountable discussions with all of you. Your input is invaluable and remains truly appreciated. Special thanks go to Dennis Sheehan for acting as my second advisor; Urs Wälchli for offering endless advice without expecting anything in return; Demian Berchtold for helping me to keep perspective when things threatened to get out of hand; and Karin Dolder for uncountable good deeds that made my life easy. I thank the Swiss National Science Foundation for funding my project.

Regardless of the nurturing professional setting, ups and downs are inherent in research. Therefore, I am deeply grateful for the unconditional support and affection that I experienced through my family and friends. Thank you. This dissertation is your achievement, as much as it is mine.

Content

Paper 1:

Pilot study of Swiss firms **1**

- 1. Introduction..... 2
- 2. Literature review..... 4
 - 2.1 Organizational characteristics of boards 4
 - 2.2 Characteristics of directors 5
 - 2.3 Individuals and the corporation..... 6
- 3. Data..... 7
 - 3.1 Sample construction 7
 - 3.2 Data description..... 8
- 4. The board and firm performance..... 11
 - 4.1 Organizational board characteristics and the model of performance..... 11
 - 4.2 Individual directors and performance 11
- 5. Influence or matching? 12
- 6. Director styles..... 14
- 7. Director characteristics..... 15
- 8. Conclusions..... 15
- Appendix A: Variable definitions..... 18
- Tables..... 20

Paper 2:**Evidence from the S&P 1,500 31**

1.	Introduction.....	32
2.	Why should outside directors be associated with firm performance?	35
3.	Overview of the research approach	37
3.1	Testing strategy	37
3.2	Generating placebo data	38
4.	Sample and data	42
4.1	Sample construction	42
4.2	Distribution of tracked directors across firm-years	43
4.3	The sample firms	44
4.4	The regression setting	46
5.	Individual director effects	48
5.1	Estimating director fixed effects.....	48
5.2	Issues with the testing strategy	49
6.	Benchmarking against the empirical distribution.....	51
6.1	Improving comparability and obtaining the empirical distribution.....	51
6.2	Investigating the counterintuitive findings.....	53
6.3	Benchmarking against different placebo outcomes.....	56
7.	Conclusions.....	58
	Appendix A: Variable definitions.....	61
	Appendix B: Number of possible random observations	63
	Tables	64
	Figures.....	77

Paper 3:**Robustness, statistical power, and important directors 87**

1.	Introduction.....	88
2.	Data and methodology	91
3.	Additional tests of the director fixed effects	93
3.1	Robustness to model specification	93
3.2	Hypotheses on the estimated effects	94
4.	Statistical power of individual fixed effect regressions with placebo benchmarking	97
4.1	CEO fixed effects.....	98
4.2	A simulation approach to evaluate the statistical power of our test	102
5.	Focus on influential directors	106
5.1	Director characteristics data	106
5.2	Special directors	107
5.3	Special situations.....	109
5.4	Multivariate analysis	111
6.	Conclusions.....	114
	Appendix A: Variable definitions.....	116
	Tables.....	118
	Figures.....	133

References 134**Selbständigkeitserklärung 142**

Paper 1:

Pilot study of Swiss firms

Abstract

Using a fixed effects approach, we investigate whether the presence of specific individuals on Swiss firms' boards affects firm performance and the policy choices they make. We find evidence for a substantial impact of these directors' presence on their firms. Moreover, the director effects are correlated across policies and performance measures but uncorrelated to the directors' background. We find these results interesting but conclude that they should to be substantiated on a dataset that is larger and better understood by researchers. Also, further tests are required to rule out methodological concerns.

1. Introduction

The extant literature typically investigates corporate boards from an aggregate point of view, such as the number of directors or committees or the presence of a dual leadership structure. Generally, boards are therefore treated as portfolios of personal characteristics that should reflect monitoring skills, incentives, expertise, experience, access to valuable connections, and reputation.¹ 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. Overall, this literature produces very mixed evidence and, thus, raises questions about the actual function of boards.

Using soccer as an analogy, the portfolio 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. However, we believe individual directors' contribution is an important issue for at least three reasons: (1) When nominating, compensating, or replacing directors, firms need to know whether a director has an individual effect or whether all that matters is his contribution to the board as a group. (2) Current legal practice allows holding individual directors liable for their actions. A recent example that gained substantial public attention in Switzerland is the lawsuit against board members of SAir Group. This liability is appropriate only if directors can individually affect firm performance. (3) There is ample evidence suggesting that individuals can influence corporations. Studies evaluate, for example, the impact of CEOs' idiosyncrasies on firms. Therefore, we ask whether board members make individual contributions to board decisions and overall firm performance. This question is related to the category of recent papers mentioned by Fahlenbrach, Low, and Stulz (2010), who “emphasize the importance of going beyond broad board characteristics and analyze specific types of directors.”² These authors contribute to this new perspective by examining the relevance of outside CEO directors.

The empirical method we use to assess the importance of individual directors is director fixed effect regressions. That is, for each director we create a dummy variable equal to one, if the director serves on a firm's board in a given year, and equal to zero otherwise. We then include these dummy variables in regressions that explain corporate decisions and firm performance. The fixed effects method restricts the impact of individual directors to constant parallel shifts. 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.

Using a sample of listed Swiss firms we find highly significant and economically large individual director effects. For example, Return on assets is 3.91 percent higher or 2.01 percent lower for first or third quartile directors, respectively. Moreover, these effects are not limited to few directors but rather

¹ See, among many others, Weisbach (1988), Kaplan and Reishus (1990), Booth and Deli (1999), Güner, Malmendier, and Tate (2008), and Kaplan, Klebanov, and Sorensen (forthcoming).

² Fahlenbrach, Low, and Stulz (2010), p. 14.

widely distributed: 27.2 percent of the director effects on Return on assets are individually significant at the 10 percent level. These effects might result from the directors' skills and preferences influencing firm performance. However, similar effects would arise from assortative matching between firms and directors. If firms with high performance succeed in recruiting individuals with a positive track record, whereas poorly performing firms can recruit only less successful individuals, this matching mechanism is in line with the fixed effects we document. However, the effects would arise endogenously. That is, firm performance would cause the presence of certain directors, rather than the directors influencing firm performance. We test whether the effects we find result from endogenous matching but find no evidence for this hypothesis. Therefore, it seems that these effects are the outcome of heterogeneity in director skills, education, and personality. Consistent with that hypothesis, individual directors also affect specific firm decisions such as number of acquisitions or level of cash holdings. Moreover, the marginal impact of individual directors on firm decisions correlates with their marginal impact on firm performance. For example, directors associated with high performance tend to be associated with high Investment outlays. These preliminary results indicate that directors are individually important. We also investigate whether these individual effects are related to personal characteristics of the directors such as education and expertise. However, we find barely any evidence that would support such a relation.

To the best of our knowledge we are the first to apply the fixed effects method to directors in a panel setting. Doing so contributes to the academic discussion about board governance. At the most fundamental level, we evaluate whether boards matter as a governance institution. By bringing the fixed effects approach to individual directors, we observe the effects of individual directors' behavior. This gives valuable insight into the inner workings of the board and complements the extant literature that focuses on committees and meeting frequencies. The article also has implications for firms, directors, and legal practice. Should firms evaluate their directors individually or as a team? The answer to this question affects the incentives directors receive from their firms. After all, optimal board compensation would seem to depend on how much an individual can contribute to the outcome of the group process.

In the remainder of the article we first provide an overview of the relevant literature. Chapter three describes our data sources and how we construct our sample. As Swiss data are not well known in the literature we discuss the characteristics of our sample in depth. First, we illustrate how directors typically move from one firm to the other. Then, we provide summary statistics for all the variables we use in our analyses and evaluate the consequences of the criteria by which we select directors for tracking. Chapter four starts by introducing the regression model and evaluating the effect of the selection criteria in the multivariate setting. It then estimates director fixed effects on performance. In chapter 5 we lay out two different causal interpretations for our findings and provide evidence that the effects result from influence rather than matching. Chapter 6 documents what Bertrand and Schoar (2003) would refer to as director styles, namely the correlation of individual director effects across different policy choices. Chapter 7 seeks to relate these styles to director characteristics. The last chapter concludes and discusses limitations associated with the use of Swiss data.

2. Literature review

Before we start our investigation we set the stage by reviewing the literature that is relevant to our problem. In a broad sense, this would include literature on the team problem (Holmstrom 1982), multi-agent situations (Mookherjee 1984; Itoh 1991), upper echelons (Hambrick and Mason 1984), and group processes (Shaw 1971). However, these strands of literature are large and diverse so that reviewing them extensively would exceed the scope of this article. Therefore, our review focuses on the finance literature.

From a finance perspective, the board of directors is a crucial element in the governance of modern corporations as it links ownership and control, represented by shareholders and management, respectively (Fama and Jensen 1983). However, boards represent many, and at times conflicting, interests. In fact, introducing a board of directors does not solve the agency problem but adds another layer to it. After all, directors have their own personal interests which need not always be in line with shareholder interests. In some situations directors' interests are more similar to management's interests than to shareholders' interests. For example, takeover offers are usually considered beneficial to the target shareholders. The target's directors, however, are often removed from office if a bid is successful (Walkling and Long 1984; Harford 2003). It has been suggested that proper organization of the board could mitigate such conflicts of interest.

2.1 Organizational characteristics of boards

One of the most frequently discussed characteristics of boards is how their processes are organized. Anecdotal and survey evidence describe a substantial change in board processes between the late 1960s and the 1990s. Over these three decades, boards allegedly have become more vigorous and professional in pursuing shareholder interests.³ These claims are difficult to assess because publicly available information on the inner workings of boards is scarce. The evidence on board committees provided by Klein (1998) and Hayes, Mehran, and Schaefer (2004) suggests that board processes are adapted to firms' needs and that they matter for firm performance. Vafeas (1999a) investigates board meeting frequencies and shows that board activity increases subsequent to declines in share price. This reaction improves operating performance of the firm.

Board size is another characteristic that has been examined. On the one hand, larger boards are more likely to provide wide expertise. On the other hand, according to Lipton and Lorsch (1992) and Jensen (1993), both the coordination costs and the likelihood of free-riding behavior increase with the number of directors. In line with the latter argument, Kini, Kracaw, and Mian (1995) find that boards are downsized following disciplinary takeovers. Several articles investigate the effect of board size and find an inverse relation between board size and firm value (Yermack 1996; Eisenberg, Sundgren, and Wells 1998; Faleye 2007). Evidence provided by Coles, Daniel, and Naveen (2008) suggests a more complex relation between board size and Tobin's Q.

Whether one or two persons occupy the positions of CEO and COB is another readily observable and much debated feature of board governance. A one-person solution benefits from clear-cut leadership

³ See Mace (1971), Lorsch and MacIver (1989), and MacAvoy and Millstein (1999).

and low information costs. This view is supported by the empirical results of Donaldson and Davis (1991) and Brickley, Coles, and Jarrell (1997), among others. The flipside of this solution is that it abandons the separation of decision management and decision control. Arguably, such a solution is associated with less rigorous board monitoring (Jensen 1993). The implied negative effect is documented in several empirical articles (Rechner and Dalton 1991; Pi and Timme 1993; Goyal and Park 2002). Palmon and Wald (2002) suggest that the net effect resulting from trading off the two aspects depends on firm size. Other articles find no relation at all between firm performance and leadership duality (Daily and Dalton 1993; Baliga, Moyer, and Rao 1996).

One way firms seek to overcome the possible conflict of interest between directors' personal incentives and the interests of the shareholders is by setting the proper financial incentives for their directors by an appropriate level of directors' ownership stake in the firm. Morck, Shleifer, and Vishny (1988) find a non-monotonic overall relation between board ownership and Tobin's Q. Yet there is evidence that board ownership can help overcome conflicts of interest (Walkling and Long 1984). Moreover, according to Gilson (1990), directors' ownership tends to increase following debt-restructurings, which could be bondholders' attempt to boost board monitoring. The second source of financial incentives for directors is their compensation. The empirical results of Perry (2000) and Adams and Ferreira (2008) suggest that proper incentive compensation can lead to the desired board behavior. The downside of incentive compensation is that it can lead to collusion between the board of directors and management.

Another board characteristic that is often examined is entrenchment. This term describes a number of board features that make dismissal of underperforming directors difficult. Bebchuk and Cohen (2005) discuss staggered board elections as one method of entrenchment and provide evidence that this mechanism reduces firm value. Other forms of entrenchment include anti-takeover charter provisions such as poison pills or classified boards (Gompers, Ishii, and Metrick 2003).

2.2 Characteristics of directors

Empirical analyses of the board's organizational structure often treat individual board members as homogenous. However, directors may differ with respect to skills and incentives. Recent empirical studies, therefore, analyze how certain types of directors are related to board actions and performance. One way to distinguish types of directors is to look at whether they are employed by the firm on whose board they serve or whether they are outsiders. This distinction is a frequent measure for directors' independence from the CEO, a feature that is considered crucial for effective monitoring. Contrary to the recommendations of many corporate governance codes, evidence about the impact of the fraction of outside directors on firm value or performance is inconclusive. Nevertheless, board independence appears to be an important governance characteristic as it seems to affect specific board decisions (see Bhagat and Black (1999), Hermalin and Weisbach (2003), and Bebchuk and Weisbach (2010) for literature surveys).

Interlocks are a specific form of director dependence that is ignored by the distinction between insiders and outsiders. This term refers to a situation where two executives serve simultaneously on each

other's board. Empirical evidence suggests that interlocks are fairly common and that they may lead to undesirable outcomes.⁴

Following Fama and Jensen (1983), a considerable body of literature has looked at how the expertise and experience of directors impacts firm policy and performance. The list includes directors who are CEOs of other firms (Fich 2005; Fahlenbrach, Low, and Stulz 2010), accounting experts (Agrawal and Chadha 2005; Defond, Hann, and Hu 2005), external bankers (Jensen and Meckling 1976; Diamond 1984; Booth and Deli 1999; Kroszner and Strahan 2001; Güner, Malmendier, and Tate 2008), venture capitalists (Baker and Gompers 2003), foreigners (Masulis, Wang, and Xie 2012), and politically connected directors (Agrawal and Knoeber 2001; Goldman, Rocholl, and So 2009). Because of the many roles a board has to fulfill, it is not surprising that the evidence is mixed. The presence of external bankers, for example, could increase the board's monitoring efficiency. At the same time, however, it could also fuel the conflict of interest between shareholders and creditors.

Another controversially discussed topic is seat accumulation of directors. Holding multiple mandates simultaneously reduces the time a director can dedicate to any one firm. Empirical evidence suggests that a large number of simultaneous mandates is problematic. On the other hand, firms benefit from directors with several mandates, as that would mean greater experience and a larger network of connections. Moreover, multiple mandates may be a characteristic of successful directors.⁵ Arguably, reputable directors are likely to be offered additional mandates (Fama and Jensen 1983). According to Fahlenbrach, Low, and Stulz (2010), reputational damage may have repercussions on a director's primary employment. Most arguments concerning director reputation refer only to the public perception of directors as monitors. However, this perception may differ from that of insiders. Adams, Hermalin, and Weisbach (2010) argue that CEOs favor directors with an intact public reputation but who are privately known for "not rocking the boat."

A common feature of most empirical studies mentioned above is that they treat boards as portfolios of director characteristics. For example, articles evaluating the role of independence do so by measuring the fraction of outside directors on the board. While this perspective tells us something about the importance of board composition, it says little about individual directors' contribution to firm performance. However, we know some things about the contribution of individuals to their firms.

2.3 Individuals and the corporation

Most of our knowledge about the role of individual directors results from event studies. These articles look at how the market reacts to the announcement that a new director is appointed to the board. This method was first used to gauge directors' contribution to firm value by Rosenstein and Wyatt (1990). This early study is also the only one that explicitly finds a significant overall announcement effect. The studies that follow either don't test for an average announcement effect or find it to be insignificant.

⁴ See Hallock (1997), Loderer and Peyer (2002), Bizjak, Lemmon, and Whitby (2009), and Fahlenbrach, Low, and Stulz (2010).

⁵ See Mace (1971), Bacon and Brown (1975), Fama and Jensen (1983), Kaplan and Reishus (1990), Rosenstein and Wyatt (1994), Booth and Deli (1996), Carpenter and Westphal (2001), Loderer and Peyer (2002), Ferris, Jagannathan, and Pritchard (2003), Perry and Peyer (2005), Fich and Shivdasani (2006), and Adams, Hermalin, and Weisbach (2010).

However, these studies document significant differences in the market reactions. For example, announcement returns are lower when the CEO is involved in the director nomination process (Shivdasani and Yermack 1999). Further, the announcement return is higher if the appointee is an outside CEO (Fich 2005; Fahlenbrach, Low, and Stulz 2010), is politically connected (Goldman, Rocholl, and So 2009), and has industry experience (von Meyerinck, Oesch, and Schmid 2012). Another type of event is investigated by Nguyen and Nielsen (2010). They look at unexpected deaths of board members and find significantly negative abnormal returns. Moreover, they identify independence as the crucial determinant of a director's value to the firm.

The fact that so much of what we know about the role of individual directors is based on a single econometric method (i.e., event studies) is somewhat disconcerting. For example, it means that our knowledge depends entirely on the market's ability to quickly assess the impact of a director. This reliance on a single method is not common in related strands of literature: Event studies are also used to assess the contribution of management to the firm (Bruce Johnson, Magee, Nagarajan, and Newman 1985; Bennedsen, Nielsen, Perez-Gonzalez, and Wolfenzon 2007), however, different methods substantiate these findings. Articles look, for example, at award winning CEOs (Malmendier and Tate 2009) or at hospitalization of CEOs (Bennedsen, Perez-Gonzalez, and Wolfenzon 2011). Others compare operating performance before and after CEO turnover (Denis and Denis 1995; Huson, Malatesta, and Parrino 2004). Related to this approach is the managerial fixed effects method used by Bertrand and Schoar (2003), Frank and Goyal (2007), and Graham, Li, and Qiu (2012). To our knowledge, the latter method has not yet been used to estimate individual director effects in a panel setting.

Overall, this literature review shows that there is substantial academic interest in the board as a mechanism to mitigate agency conflicts in public firms. While recent articles investigate the role of several director characteristics in the board processes, our knowledge on the role of the individual board member is still limited. In particular, it concentrates on results obtained using the case study approach. We propose an alternative method: director fixed effects. Implementing this approach, we control for several of the previously discussed board characteristics. Chapter 3 introduces the data we use in our analyses.

3. Data

3.1 Sample construction

Estimating director fixed effects implies tracking individual directors across firms and over time. Thus, we require data on board composition that allow this sort of tracking. In particular, we have to be able to uniquely identify directors throughout the sample. The data from Waelchli (2008) meet this requirement. Wälchli hand-collected corporate governance data from Swiss firms' annual reports. The data set provides annual information on board composition since 1995 and until 2009. This information is organized such that each individual (i.e., board members and CEOs) has a unique and constant personal identifier. These identifiers also link to individuals' background information such as age,

gender, education, and nationality. Further, the dataset provides ownership information for the sample firms. The dataset covers the firms listed on SIX Swiss Exchange (formerly SWX) excluding investment firms, firms where annual reports were no longer available in 2005, and firms that were listed for less than three years during the sample period.

We match this corporate governance data with several other sources to complete our dataset. Most importantly, we obtain firms' financial data from Compustat global. Further, we use mergers and acquisitions (M&A) data from SDC Platinum. Following Bertrand and Schoar (2003), we exclude firms from the financial and the utilities industries. This procedure yields a panel of 202 firms over 15 years and a total of 2,033 firm year observations.

3.2 Data description

A total of 2,042 individuals serve on the boards in our sample. We restrict our fixed effects analysis to a subset of these directors using the following criteria⁶: First, mandates lasting less than three years are excluded because it seems unlikely that a director can influence the firm immediately after appointment. After all, von Meyerinck, Oesch, and Schmid (2012) show that it may take some time for independent directors to familiarize themselves with the firm and sometimes even its industry. As a second criterion, we consider only directors that have mandates in multiple firms. Excluding single-firm directors helps distinguishing director effects from firm fixed effects. To see how these sample criteria work, consider for example Mr. Hans Ziegler. The initial dataset lists Mr. Ziegler on the boards of Charles Vögele, Elma Electronic, Global Natural Resources, OC Oerlikon, Schlatter, and Swisslog. His service for Charles Vögele and OC Oerlikon started in 2008. Thus, these mandates appear in our sample for two years only and do not qualify for tracking. The four remaining mandates last three or more years and qualify for the tracked subsample.

Table 1-1 further illustrates this selection process. Panel A provides the number of years each director spends on the board of a given firm and Panel B the number of board positions each director holds. Both panels distinguish between the full dataset (columns 2–4) and the tracked subsample (columns 5–7). Out of the 2,042 directors in our initial sample 236 qualify for tracking. 1,700 directors were dropped because they serve on only one board. The remaining directors were excluded because less than two of their mandates last at least three years. In terms of director-firm combinations (henceforth referred to as mandates), 618 out of 2,616 qualify for the tracked subsample. 742 mandates are excluded because they last less than three years. Another 1,256 mandates are not considered for further analysis because they are the only mandates of a given director lasting three or more years. This selection process increases the average duration of a mandate in the tracked subsample by about 1.5 years compared to the full dataset. The average number of board positions per director rises from 1.28 to 2.62. Note, that these numbers result from within our dataset: Revisiting the example of Mr. Ziegler, this means that his additional involvements with Globus and Interdiscount are not represented in Table 1 because these firms were active for less than three years in the sample period. Similarly, his term with Usego-Trimerco is missing because it does not coincide with the sample period. Further, several positions in firms not listed at SIX Swiss Exchange do not enter Table 1-1.

⁶ Originally, these criteria were used from Bertrand and Schoar (2003) to select managers for their fixed effects analysis.

Evidently, we focus on a group of rather long-term serial directors when applying the above selection criteria. As we have seen in the discussion of the literature, there is reason to believe that these directors are special. In particular, successful directors should be awarded multiple mandates (Fama and Jensen 1983) while underperforming individuals are expected to be dismissed from the board (Gilson 1990; Harford 2003). Thus, the directors we track may be more likely to influence their firms in a positive way than average directors. In this sense, our results do not necessarily apply to directors in general but should be understood as statements about this special class of directors. This notion is supported by a comparison of director backgrounds. Table 1-2 documents a significantly increased probability of finding university graduates, individuals holding a doctoral degree, financial experts, men, and Swiss citizens among tracked directors. Further, tracked directors are on average nearly five years older than their peers. Legal experts, on the other hand, are similarly distributed across the two subsamples.

Table 1-3 illustrates the dynamics in the tracked subsample. First, it reports how many different years the 236 directors appear in (Panel A). The fact that no director is tracked in less than three years results from the sample selection criteria. We see that the board service of most directors is widely spread across the sample period and concentrations of multiple mandates on a short sub-period are the exception. Specifically, 78.81 percent of the directors are tracked in eight or more different years and the average director is tracked in 10.33 years. These numbers result in a total of 2,439 tracked director-years (Panel B). In 44.69 percent of these years, the director serves on a single board only. In the remaining 55.31 percent of years the director serves on multiple boards. On average, a director is tracked on 1.74 boards at a time with a maximum of five simultaneously tracked mandates. Panel C shows 19 directors that are never tracked in more than one firm at a time and 20 directors that are never tracked in fewer than two firms at a time. Hence, relatively few directors have strictly consecutive or perfectly overlapping mandates—the majority of mandates overlap partially.

The consequences of the selection criteria and the board dynamics were so far discussed from the directors' perspective. Table 1-4 changes to the firms' vantage point. Panel A shows how many directors were tracked in a given firm year. The lower end of the distribution is marked by 466 firm years without any tracked directors. On the other end, three firm years have a maximum of 11 simultaneously tracked directors (SAirGroup in the years 1995 to 1997). On average, 2.08 directors were tracked in each firm year. According to Panel B, 53 of the totally 202 firms never have any tracked directors. Another 29 firms have some years without tracked directors and 120 firms have at least one tracked director at all times. Two firms, Georg Fischer and SAirGroup, never have fewer than 5 tracked directors.

We have already shown that directors who qualify for tracking typically differ from their untracked colleagues. Table 1-5 asks whether firms that employ these directors are special, too. The table introduces the variables we are going to use in the subsequent analyses and provides mean and median values (see Appendix A for variable definitions). Statistics are shown for the full sample and for the subsamples of firm years with and without tracked directors separately. t-tests and Wilcoxon-tests compare the variables across the subsamples. Our measure for firm size is total assets. The

median firm in our analysis has CHF 505.86 million assets. Hence, on an international scale, many of our firms are rather small. The mean size is considerably larger (CHF 3,688.35 million), indicating that the distribution is skewed. Mean and median size are significantly greater in firm years with tracked directors. Three different performance measures are used: Return on assets (ROA), Return on equity (ROE), and Tobin's Q. All of these measures are higher for firm years with tracked directors in terms of mean and median values. The differences in ROE and Tobin's Q, however, are only significant according to Wilcoxon tests but not t-tests. Looking at firm policies, firm years without tracked directors seem to rely more on internal growth (i.e., through Investment and R&D) whereas firm years with tracked directors grow more externally (i.e., through Acquisitions). In financial policies we find fewer significant differences. Concerning board characteristics we find that firm years with tracked directors have larger boards. This finding might, to some extent, result endogenously from our selection criteria: The probability of having one or more directors that qualify for tracking increases in the number of directors on the board. Given that our second criterion for tracking is a minimum tenure of three years, it is somewhat surprising that mean board tenure is higher for firms without tracked directors. This observation might indicate that directors, who do not qualify for tracking because they serve on a single board, stay with their firm longer. About 80 percent of the directors in our sample are independent. While differences in independence between the two subsamples are statistically significant they are economically small: The mean values are 0.07 percent apart, corresponding to less than half a director at the typical board size. Further, Table 1-5 shows that blockholders are fairly common in our sample firms. The ownership data is based on the mandatory disclosure of block ownership. Thus, we have no information on the distribution of voting rights that are not part of a block. The threshold above which ownership stakes are considered blocks was 5 percent before 2007 and 3 percent since. To maintain consistency in our data we ignore all blocks smaller than 5 percent. The holder of the largest block controls on average more than a third of the total votes. This fraction is even larger in the subsample of firms without tracked directors and the difference is significant. Distinguishing different types of blockholders, we see that ownership by the founding family is prevalent. Family involvement is particularly pronounced in firms without tracked directors. The board members and the government, on the other hand, do not usually control major stakes in the sample firms.

These differences between firms with and without tracked directors suggest that our investigation not only focuses on a rather special group of directors but also that the firms who employ these directors are not representative for all listed firms. However, some of the difference might be accounted for by control variables. Therefore, we test whether the distinction of firm years with and without tracked directors remains significant in the multivariate setting of the regression model we use for all performance analyses.

4. The board and firm performance

4.1 Organizational board characteristics and the model of performance

Table 1-6 introduces the model that relates firm performance to board characteristics. In the next section, this model will be added the individual director indicators. Here, the model instead includes an indicator for firm years with at least one tracked director. This indicator represents the sample split from the previous tables in a multivariate setting. Performance is measured by Return on assets, Sales to assets, and Tobin's Q. Board variables include Board size, Fraction of independent directors, Board tenure, and an indicator for unitary leadership (CEO is COB). Further, we control for firm size (measured as the natural logarithm of total assets), Investment, Leverage, R&D, and the ownership variables discussed in Table 1-5. The regression model also includes firm and year fixed effects. Our model specification and variable definitions are based on Cronqvist and Fahlenbrach (2009). The indicator for tracked firm years, the board related variables, ownership, and the respective squared terms measures we added off our own.

We find that most governance variables we control for in the regressions produce insignificant coefficients. Presumably, these characteristics vary little over time and, thus, their influence is to some extent captured by firm fixed effects. One notable exception is board size. We find coefficients that imply hump-shaped relations between board size and Return on assets and Tobin's Q. The turning points of these two relations are at 9.77 and 9.02 directors, respectively. About 90 percent of boards are smaller than this. While, this might indicate that the negative square term represents a flattening rather than an inversion of the relations, the majority of residuals is negative for board sizes larger than 9. Thus, the quadratic shape we estimate typically overestimates performance for these firms and the decrease in performance is even greater than predicted. This negative slope is in line with prior findings for US firms (Yermack 1996; Faleye 2007). The only voting rights that are significantly related to performance are the ones held by the government. Government controlled firms seem to have lower operating performance. The coefficient for the variable identifying firm years with tracked directors is not significant in the Return on equity and Tobin's Q regressions. It is significant only at the 5 percent level in the Return on assets regressions. Thus, the differences in performance between the two subsamples we observed in Table 1-4 largely disappear in the multivariate setting.

4.2 Individual directors and performance

Table 1-7 follows the specifications of Table 1-6 and asks whether individual directors are associated with firm performance. Instead of one indicator variable for firm years with tracked directors, the regressions now include the full set of individual director indicators. As this change substantially increases the number of explanatory variables, the table does not present individual coefficients. Rather, Panel A of the table tests for the overall significance of different sets of coefficients and reports the associated F-values. The F-values representing the set of director indicators are highly significant. Thus, allowing for individual director effects significantly improves the fit of our model for each performance measure. The contribution of the individual director indicators is also reflected by the increase in R square which amounts to 3 to 5 percent. As for the remaining variables, voting rights and time varying firm level controls (X_{it}) are significant for all performance measures. Similarly, firm and

year fixed effects are significant. In contrast, the group of board characteristics is not jointly significant in any regression.

Panel B of Table 1-7 examines the estimated director effects. 146 to 149 coefficients are estimated in each regression. Thus, roughly 90 indicator variables are missing compared to the 236 directors meeting the selection criteria. These variables are omitted either because of collinearity or because of missing data. The table reports the quartiles of the distribution of the estimated director fixed effects. By comparing the inter quartile range (IQR) to the median of the dependent variable of the regressions (Median (y)), we find the director effects to be economically large. In the Return on assets regression, for example, the IQR of the individual director coefficients equals 5.81 percentage points with a median Return on assets of 6.83 percent. The effects on the other performance measures are of similar magnitude.

Overall, the size of the F-values and the magnitude of the effects are larger than we would expect. For example, Cronqvist and Fahlenbrach (2009) find no F-values greater than 3.57, whereas the smallest F-value in Table 7 is 19'519. Also, Fahlenbrach, Low, and Stulz (2010) find barely any difference in operating performance of firms that appoint CEO or non CEO outside directors. This difference might result from us using a larger sample. Benchmarking our findings against placebo results similarly to Cronqvist and Fahlenbrach (2009) might provide additional insights into the relative magnitude of the effects we measure. However, doing so is beyond the scope of this pilot study.

The F-tests in Panel A document that the group of director indicators is overall significant. However, they give no indication as to whether the effects represent a widespread phenomenon or few extraordinary directors. Panel B therefore also investigates the individual significance of the fixed effects using t-tests. We report the fraction of significant coefficients relative to the total number of coefficients. For example, in the Return on assets regression, 8.05 percent of the coefficients are individually significant with 99 percent confidence. This fraction is much higher than what one would expect just randomly. Comparable fractions of significant coefficients are found for the remaining performance measures. Thus the effect we document is spread widely among the tracked directors. Overall, this evidence suggests that firm performance is substantially related to the presence of certain individuals on the board of directors.

5. Influence or matching?

The relation between firm performance and the presence of individuals on the board as documented in Table 1-7 can result from more than one causal effect. The first interpretation is that individual directors affect their firms. An alternative interpretation is that firms and directors are matched endogenously. Poorly performing firms and firms with poor governance, for example, might look for

directors who don't cause a stir and these firms might indeed end up attracting them.⁷ If so, the causality would not go from directors to poor performance, but rather the other way.

We follow Cronqvist and Fahlenbrach (2009) and disentangle the two interpretations based on the different implications they have on the timeline: Directors can only exert influence *after* they join the board. For a matching to take place, however, a nominee would have to observe the firm's characteristics *before* joining its board. Therefore, we would expect to find similar performance and policy effects during the years prior to the appointment if the matching hypothesis holds.

To estimate the effects prior to a director's appointment we start by randomly dividing the director's mandates into two groups. The first group of mandates is used to estimate the ordinary director effects. We use the second group of mandates to generate indicator variables that take the value of 1 in the 2 years before the director joined the board and 0 otherwise. We refer to these as the 2-years-before effects. We estimate the two effects on separate subsamples to avoid mechanically inducing a negative correlation between director and 2-year-before effects. To see why this would be the case, remember that a mandate lasts 6.68 years on average (see Table 2). Further, the average firm is in the sample in 10.06 years. Hence, only 3.38 years are left where the director is not present. If we observe a positive director effect that means that the performance of the director's firms is above average when the director is present and, consequently, below average when the director is absent. As the director is absent in only few years it is likely that the 2-years-before effect would mechanically reflect this below average performance. The opposite is true if we observe a negative director effect. Thus, the director and 2-years-before effects would be negatively correlated.⁸ The downside of using separate subsamples is that we use only half of the directors' mandates to estimate their fixed effects. As the average director is only present in 2.62 firms, this revives the concern of distinguishing firm and director effects (see our discussion of the sample criteria). We cannot maintain the initial sample criteria because only a handful of directors pass after their mandates are split into two groups.

Panel A of Table 1-8 shows the results of the regressions estimating the director effects and the 2-years-before effects. The director indicators keep contributing significantly to the model's explanatory power for each performance measure. The R squares are lower than in Table 1-7, where we included the full set of director indicators, but still higher than the ones in Table 1-6, where only a single indicator variable for firm years with tracked directors was included. The 2-years-before effects are overall significant. This latter finding is important for the matching hypothesis as this hypothesis suggests that the firm is special prior to the director's appointment. Further, the matching hypothesis predicts that the 2-years-before effects are related to the actual director effects. However, Panel B shows that a director's probability of having positive effect on Return on assets and Return on equity is lower for directors that have a positive 2-years-before effect. This finding is the opposite of what the matching hypothesis predicts and might indicate that underperforming firms appoint directors they

⁷ See Mace (1971), DeAngelo and DeAngelo (1989), Lorsch and MacIver (1989), and Adams, Hermalin, and Weisbach (2010).

⁸ This explanation is somewhat simplified to intuitively convey the underlying the problem. The actual situation is complicated by firm and year fixed effects and the fact that all directors are observed in multiple firms.

expect to turn the trend or that directors nominated by highly performing firms are unable to maintain this performance level. The probability of a positive director effect on Tobin's Q does not significantly depend on the 2-years-before effect. Thus, we do not find any evidence in line with the matching hypothesis. Rather, our findings seem to document directors' influence on the performance of their firm.

6. Director styles

Presumably, directors affect firm performance because they contribute to shaping firm decisions and policies. Therefore, we investigate whether director fixed effects explain firms' growth and financial strategy. Growth strategy is measured by Investment, R&D outlays, the total number of acquisitions, and the number of diversifying acquisitions. Our measures for financial strategy are Leverage, Cash holdings, and Dividend payouts. Table 1-9 reports the estimates of a regression specification similar to that in Table 1-7, except for the different firm controls, which we take from Cronqvist and Fahlenbrach (2009). The results, too, are similar to those in Table 1-7. As before, director coefficients are jointly highly significant (Panel A). However, ranging from 1.86 in the diversifying acquisitions regression to 6.10 in the Leverage regression, they are markedly lower than the ones we observed for performance. Also, the fractions of individually significant coefficients are mostly smaller than in the performance regressions. In particular, the ones for the number of acquisitions and the number of diversifying acquisitions are barely higher than we would expect just randomly (3.18% and 1.91% at the 1% level, respectively). Other regressions, such as Leverage and R&D, still have a substantial number of individually significant coefficients. The relative magnitude of the estimated coefficients is larger and comparable to what we found in the performance regressions (Panel B). For example, median Leverage is 31.15 percent. We expect this to be 7.93 percent lower or 7.21 percent higher if a first or third quartile director serves on the firm's board. Thus, a firm's directors seem to be an important determinant of its policy decisions as well as its performance.

Bertrand and Schoar (2003) document that the effects a manager has on different policies of her firm are correlated. They interpret this finding as evidence for different management styles. In a similar vein, Table 1-10 uses the individual directors' coefficients on performance (Table 7) and policies (Table 9) and correlates them across performance measures and policies. We report Spearman rank correlation coefficients to control for potential outliers. Many pairs of director coefficients are significantly correlated. Directors associated with good operating performance also tend to be associated with a high number of acquisitions, low Leverage, and small Cash holdings. The ones who are associated with high Tobin's Q are also associated with high Return on assets. Meanwhile, there is no association between Tobin's Q and any of the policy preferences. Directors with a preference for high Investment are also associated with a high number of Acquisitions, high R&D expenditures, and low Leverage.

7. Director characteristics

So far, we show that the fixed effects method can ex post determine a director's association with performance and management styles. When appointing directors, firms often cannot observe candidates on multiple boards over several years. Rather, they need to know ex ante what policies a given director stands for and what impact will result from his board membership. Therefore, it is important to know if there are readily observable factors that are related to an individual's performance as a board member. Extant studies suggest that, for example, experience in top management could be such a factor. Table 1-11 relates a set of similarly easily observable characteristics of the directors to the director effects estimated in Table 1-7. In particular, we consider the director's level of education, field of expertise, provenance, and birth cohort. The only significant relation we find is that non-Swiss directors have above average effects on operating firm performance. However, Swiss citizenship does not seem to relate to a director's effect on Tobin's Q or on his investment or financial decisions. Whether a director has graduated from a university, holds a doctoral degree, is a financial or legal expert, or is born before 1944 (the median year of birth of the tracked directors) does not relate to his effect on the firm.

It is not clear how we should interpret the absence of a measurable relation between director characteristics and the directors' effect on the firm. One explanation is that our measures of the characteristics are inadequate to capture the differences between the directors. This might be the case because the skills, experience, and preferences of directors are too complex to be associated with easily observable variables such as the ones we consider. Also, directors undergo a selection process that should rule out, for example, lack of expertise more effectively, than we can measure it in our dataset. An alternative explanation is that our tests lack the statistical power to find associations between personal characteristics of a director and his impact on the firm. This concern arises because our test is based on only about 150 observations of director effects. Moreover, there are only few observations for some of the director characteristics. For example, only 15 of the tracked directors are not Swiss citizens and only 28 are legal experts. The statistical power is further reduced because each of the effects is subject to measurement error. After all, the standard errors of the effects are typically substantial (e.g., on average 3.00 percent in the Return on assets regression).

8. Conclusions

The board of directors has the crucial role of linking the owners and the management in corporations. Consequently, it has been subject to a major debate in the literature on corporate governance and agency problems. The focus of this discussion is largely on institutional characteristics of the board and on board composition. Meanwhile, surprisingly little is known about the role and importance of individual board members.

This article shows that if we track individual directors over time and across firms, we find significant and large effects on firm performance. We investigate whether these effects are related to the firm's

performance prior to the director's nomination and find a negative or no relation at all. Thus, the individual effects seem to result from directors' influence rather than from a mechanism matching directors with firms based on previous performance. We further show that individual directors are associated with characteristic policy choices as well. These policy choices are related to each other and to the effect on performance. Following Bertrand and Schoar (2003), these sets of related effects could be considered director styles. These styles, however, seem to be barely related to easily observable personal characteristics of directors.

To our knowledge we are the first to apply the fixed effects method to directors in a panel setting. Moreover, the present article differs from the existing literature with respect to the firms investigated. While many existing studies focus on large U.S. corporations we examine Swiss firms. Looking at different kinds of firms may lead to new insights in corporate governance (Zingales 2000). However, the unusual sample has two important drawbacks. First, Switzerland is a small country and, consequently, we have a small sample although we include as many data points as possible. In our case, we observe 202 firms and track about 150 directors. These numbers are small compared to sample sizes in recently published articles. More importantly, sample size affects the statistical power of our analysis. This issue may be particularly important when asking whether director effects are associated with personal characteristics. Second Swiss firms may differ fundamentally from the firms usually investigated. While these differences might allow observing different solutions to governance problems, they also complicate the comparison of our findings to previous results. Some of the most striking differences are: (1) Although the average market capitalization of CHF 5,317 million corresponds to what we expect to find for S&P 500 firms, many of the firms in our sample are rather small. In fact, the median capitalization of CHF 354 million ranks among the smaller of the S&P 600 Smallcap firms. (2) The firms in our sample have large blockholders. The average and median ownership stake of the largest blockholder are 36 and 30 percent, respectively. For U.S. firms, Holderness (2009) provides figures of 26 and 17 percent. (3) Switzerland's score on the anti-director rights index of La Porta, Lopez-de-Silanes, Shleifer, and Vishny (2002) is 2 out of 6. Thus, Swiss directors sit rather firmly in their saddle once they are elected. These differences may lead to unique governance outcomes that are difficult to compare internationally.

We therefore suggest expanding this study to U.S. firms. This would substantially increase sample size and, along with that, the statistical power of our analyses. Consequently, we hope to get additional insights into how the director effects differ between types of directors and depending on the directors' background. Further, it would be interesting to see whether the individual directors maintain their importance if the market for directors is more competitive. After all, we would expect directors with a negative effect on performance to be removed from boards rather quickly.

In addition to these considerations with regard to the data, three issues need to be addressed: First, our findings are stronger and of greater magnitude than we would expect based on previous literature. Second, Wooldridge (2002) cautions that conventional test statistics may be invalid when dealing with fixed effects. Third, Bertrand, Duflo, and Mullainathan (2004) show that serially correlated residuals can lead to inflated test statistics in difference-in-difference estimates. As the estimation techniques these authors discuss are related to the fixed effects method we apply, this latter concern affects our

analyses similarly. Thus, further testing is required to rule out the possibility that the effects we present are statistical artifacts. One possible way to do so is placebo benchmarking. That is, we plan on comparing our findings to the outcomes of analogous analyses on random data.

Appendix A: Variable definitions

Variable	Definition; Compustat item names in parentheses
Assets	Natural logarithm of the book value of assets (at). Winsorized at the 2.5% level.
Return on assets	Operating income before depreciation (oibdp) less depreciation (dp) divided by the book value of total assets of the previous year (at). Winsorized at the 2.5% level.
Return on equity	Income before extraordinary items (ib) plus total income taxes (txt) divided by the total common equity (ceq) of the previous year. Winsorized at the 2.5% level.
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 2.5% level.
Investment	Capital expenditures (capex) divided by the previous year's net property, plant, and equipment (ppent). Winsorized at the 2.5% level.
R&D	Research and development expenditures (xrd) divided by the previous year's book value of assets (at). Winsorized at the 2.5% level.
Acquisitions	Number of acquisitions a firm completed during a given year according to SDC.
Diversifying acquisitions	Number of acquisitions a firm completed outside its industry during a given year according to SDC. 10-industry classification according to Fama and French is used.
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 liabilities (dlc), and common equity (ceq). Winsorized at the 2.5% level.
Cash holdings	Cash and short term investments (che) divided by last year's net property, plant, and equipment (ppent). Winsorized at the 2.5% level.
Dividends	Cash dividends (dv) divided by operating income before depreciation (oibdp). Winsorized at the 2.5% level.
Board size	Number of directors on a firm's board in a given year.
Board independence	Fraction of independent directors on a firm's board. A board member counts as independent if he or she has not been an executive of this firm in the past 3 years and no significant business relations to this firm other than the board membership.
Board tenure	The median number of years a firm's directors have served at a given time.
CEO = COB	Indicator variable that is 1 if the chairman of the board (COB) also serves as chief executive (CEO) and 0 otherwise.
Votes largest block	Fraction of voting rights controlled by the firm's largest blockholder. All blocks larger than 5% are considered.

Variable	Definition; Compustat item names in parentheses
Votes board	Total fraction of voting rights controlled by the board of directors. All blocks larger than 5% are considered.
Votes government	Total fraction of voting rights controlled by the government. All blocks larger than 5% are considered.
Votes family all	Total fraction of voting rights controlled by the firm's founders or members of the founding family. All blocks larger than 5% are considered.

Tables

Table 1-1

Description of the sample directors' board service – The full dataset covers all board members of listed Swiss firms 1995 to 2009, excluding the financial and utilities sectors. The “tracked subsample” results from excluding directors who serve in only one sample firm and mandates lasting less than 3 years. Panel A counts the number of different boards each director serves on. Panel B lists the mandates according to their duration.

<i>Panel A: Number of board positions by director</i>						
	Full dataset			Tracked subsample		
	Number of directors	Percent	Cumulative	Number of directors	Percent	Cumulative
1	1'700	83.3	83.3	0	0.0	0.0
2	210	10.28	93.54	147	62.29	62.29
3	70	3.43	96.96	51	21.61	83.90
4	33	1.62	98.58	21	8.90	92.80
5	23	1.13	99.71	15	6.36	99.15
6	3	0.15	99.85	2	0.85	100
7	3	0.15	100	0	0	100
Total	2'042	100		236	100	
	Mean: 1.28			Mean: 2.62		

<i>Panel B: Duration of mandates (director-firm combinations)</i>						
	Full dataset			Tracked subsample		
	Mandates	Percent	Cumulative	Mandates	Percent	Cumulative
1	431	16.48	16.48	0	0	0
2	311	11.89	28.36	0	0	0.00
3	324	12.39	40.75	102	16.50	16.50
4	294	11.24	51.99	87	14.08	30.58
5	241	9.21	61.20	64	10.36	40.94
6	202	7.72	68.92	83	13.43	54.37
7	176	6.73	75.65	66	10.68	65.05
8	138	5.28	80.93	43	6.96	72.01
9	132	5.05	85.97	44	7.12	79.13
10	117	4.47	90.44	38	6.15	85.28
11	58	2.22	92.66	20	3.24	88.51
12	48	1.83	94.50	23	3.72	92.23
13	34	1.30	95.80	8	1.29	93.53
14	48	1.83	97.63	17	2.75	96.28
15	62	2.37	100	23	3.72	100
Total	2'616	100		618	100	
	Mean: 5.19			Mean: 6.68		

Table 1-2

Comparison of tracked and untracked directors – The table presents contingency tables for a number of personal characteristics of the directors in our sample. It compares these characteristics across the tracked and untracked subsamples. Chi-squared tests evaluate the significance of the association in these tables. For the year born, mean values are reported for the full sample and the subsamples. A t-test (top) and a Wilcoxon test (bottom) are provided. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

<i>University graduate</i>				<i>Doctoral degree</i>			
	No	Yes	Total		No	Yes	Total
Total	502	2'595	3'097	Total	2'054	1'032	3'086
	16.21%	83.79%	100%		66.56%	33.44%	100%
Tracked	15	140	155	Tracked	85	70	155
	9.68%	90.32%	100%		54.84%	45.16%	100%
Untracked	487	2'455	2'942	Untracked	1'969	962	2'931
	16.55%	83.45%	100%		67.18%	32.82%	100%
Chi ² : 5.13**				Chi ² : 10.07***			
<i>Financial expertise</i>				<i>Legal expertise</i>			
	No	Yes	Total		No	Yes	Total
Total	1'615	1'576	3'191	Total	2'485	603	3'088
	50.61%	49.39%	100%		80.47%	19.53%	100%
Tracked	59	93	152	Tracked	122	29	151
	38.82%	61.18%	100%		80.79%	19.21%	100%
Untracked	1'556	1'483	3'039	Untracked	2'363	574	2'937
	51.20%	48.80%	100%		80.46%	19.54%	100%
Chi ² : 8.88***				Chi ² : 0.01			
<i>Female</i>				<i>Swiss citizen</i>			
	No	Yes	Total		No	Yes	Total
Total	3'367	196	3'563	Total	986	2'453	3'439
	94.50%	5.50%	100%		28.67%	71.33%	100%
Tracked	157	0	157	Tracked	16	141	157
	100%	0%	100%		10.19%	89.81%	100%
Untracked	3'210	196	3'406	Untracked	970	2'312	3'282
	94.25%	5.75%	100%		29.56%	70.44%	100%
Chi ² : 9.56***				Chi ² : 27.47***			
<i>Year born</i>							
	Mean	Test					
Total	1948.29						
Tracked	1943.88	5.22***					
Untracked	1948.53	5.61***					

Table 1-3

Distribution of tracked directors in the sample: Director perspective – The table represents all directors and mandates that meet the selection criteria. Panel A shows how many different sample years each director is tracked in. Panel B counts the number of simultaneously tracked mandates in each year a given director is present. Panel C lists the minimum and maximum number of simultaneously tracked mandates for each director.

<i>Panel A: Number of years in the sample</i>			
	Directors	Percent	Cumulative
1	0	0	0
2	0	0	0
3	3	1.27	1.27
4	7	2.97	4.24
5	7	2.97	7.20
6	15	6.36	13.56
7	18	7.63	21.19
8	16	6.78	27.97
9	30	12.71	40.68
10	23	9.75	50.42
11	25	10.59	61.02
12	25	10.59	71.61
13	15	6.36	77.97
14	27	11.44	89.41
15	25	10.59	100
Total	236	100	
	Mean: 10.33		

<i>Panel B: Number of board positions held simultaneously in a given director year</i>			
	Director years	Percent	Cumulative
1	1'090	44.69	44.69
2	1'007	41.29	85.98
3	245	10.05	96.02
4	87	3.57	99.59
5	10	0.41	100
Total	2'439	100	
	Mean: 1.74		

<i>Panel C: Minimum and maximum number of positions tracked simultaneously by director</i>					
Maximum	Minimum				Total
	1	2	3	4	
1	19	0	0	0	19
2	137	11	0	0	148
3	41	4	0	0	45
4	14	2	1	1	18
5	5	1	0	0	6
Total	216	18	1	1	236

Table 1-4

Distribution of tracked directors in the sample: Firm perspective – Panel A of the table counts the number of tracked directors for each firm year in the sample. Panel B provides the minimum and maximum number of tracked directors for each firm.

<i>Panel A: Number of tracked directors by firm year</i>			
	Firm years	Percent	Cummulative
0	466	22.92	22.92
1	436	21.45	44.37
2	322	15.84	60.21
3	379	18.64	78.85
4	258	12.69	91.54
5	103	5.07	96.61
6	38	1.87	98.48
7	20	0.98	99.46
8	4	0.20	99.66
9	0	0.00	99.66
10	4	0.20	99.85
11	3	0.15	100
Total	2'033	100	
	Mean: 2.08		

<i>Panel B: Minimum and maximum number of tracked directors by firm</i>							
Maximum	Minimum						Total
	0	1	2	3	4	5	
0	53	0	0	0	0	0	53
1	16	15	0	0	0	0	31
2	5	13	6	0	0	0	24
3	7	13	9	3	0	0	32
4	1	11	13	3	3	0	31
5	0	3	4	7	2	0	16
6	0	2	1	2	2	0	7
7	0	0	1	1	1	1	4
8	0	1	0	1	0	0	2
9	0	0	0	0	0	0	0
10	0	0	0	1	0	0	1
11	0	0	0	0	0	1	1
Total	82	58	34	18	8	2	202

Table 1-5

Descriptive statistics – The table provides mean (first row) and median (second row) values of the variables used in the analysis. Statistics are provided separately for the full dataset, the firm years without tracked directors (“untracked”), and the firm years with one or more tracked directors (“tracked”). The last column tests for differences between the two subsamples, showing t-values (first row) and z-values resulting from Wilcoxon tests (second row). For the variable indicating firm years where the same person is CEO and COB no median values are shown and the test statistic is a proportion test. Appendix A provides detailed variable definitions. The values shown in this table deviate from these definitions in two ways: first, variables are summarized before winsorizing, and second, Assets is measured in million CHF instead of logarithms. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	N	All	Untracked	Tracked	test
<i>Firm size</i>					
Assets (mCHF)	2'033	3'688.35	1'391.52	5'106.29	-5.95***
		505.86	249.16	849.08	-17.63***
<i>Firm performance</i>					
Return on assets (%)	1'960	6.50	5.08	7.35	-2.94***
		6.84	6.45	7.13	-2.60***
Return on equity (%)	1'971	11.80	8.32	13.90	-1.35
		13.91	12.04	15.16	-3.27***
Tobin's Q	1'956	1.67	1.63	1.70	-1.00
		1.25	1.22	1.27	-3.08***
<i>Firm policies</i>					
Investment (%)	1'929	23.89	32.20	18.93	2.48**
		15.17	15.44	15.01	1.40
R&D (%)	862	6.69	9.28	5.63	5.67***
		4.11	5.94	3.74	4.62***
Acquisitions	2'033	0.54	0.31	0.68	-7.43***
		0.00	0.00	0.00	-7.33***
Diversifying acquisitions	2'033	0.20	0.10	0.26	-6.65***
		0.00	0.00	0.00	-6.56***
Leverage (%)	2'029	30.90	28.45	32.41	-1.76*
		30.56	24.85	33.81	-6.41***
Cash holdings (%)	1'967	206.07	367.59	109.29	5.56***
		40.68	39.35	40.87	-0.15
Dividends (%)	1'709	13.12	13.74	12.79	0.17
		10.41	9.88	10.84	-0.79
<i>Board characteristics</i>					
Board size	2'033	6.67	5.96	7.11	-10.73***
		6.00	5.00	7.00	-12.18***
Board independence	2'033	0.78	0.74	0.81	-9.61***
		0.83	0.80	0.83	-9.45***
Board tenure	1'957	7.18	7.55	6.95	2.66***
		6.09	6.09	6.08	1.27
CEO = COB (bin)	2'033	0.24	0.26	0.23	1.53
<i>Voting rights</i>					
Votes largest block (%)	2'019	35.96	40.29	33.33	5.89***
		30.00	36.80	26.00	5.63***
Votes board (%)	2'019	1.01	0.98	1.02	-0.12
		0.00	0.00	0.00	-0.10
Votes government (%)	2'019	1.28	0.34	1.86	-4.32***
		0.00	0.00	0.00	-4.01***
Votes family all (%)	2'019	24.60	28.58	22.17	4.79***
		5.60	20.44	0.00	5.12***

Table 1-6

The regression model – The table shows regressions of performance on an indicator for firm years with one or more tracked directors. The regressions also control for the variables described in Table 1-5. Additionally, the regressions include firm and year fixed effects. Standard errors are estimated using the Huber / White sandwich method. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Return on assets (%)	Return on equity (%)	Tobin's Q * 100
Firm year tracked	1.39**	2.28	-3.68
Board size	1.05**	0.92	5.84*
squared	-0.05**	-0.04	-0.32*
Board independence	2.06	10.31	-15.72
squared	-0.57	-9.93	8.54
Board tenure	0.05	0.42	-0.11
squared	-0.00	-0.02	-0.02
CEO is COB	-1.39**	-0.50	-7.29
Votes board	0.11	0.23	-0.23
squared	-0.00	-0.00	0.01
Votes government	-0.33**	-0.85**	0.32
squared	0.00	0.00	0.00
Votes largest blockholder	-0.00	-0.05	0.39
squared	0.00	0.00	-0.00
Votes family	-0.09	-0.21	-0.75
squared	0.00	0.00	0.01
Ownership data missing	-0.68	-6.37	-17.01
Assets	-2.18***	-8.87***	-36.74***
Investment	-0.01	0.00	0.24*
Leverage	-0.03**	0.07	-0.25**
R&D	0.21	-0.05	0.99
missing	0.17	-0.32	3.83
Year FE	Yes	yes	yes
Firm FE	Yes	yes	yes
Adjusted R ²	0.60	0.43	0.73
N	1'764	1'775	1'746

Table 1-7

Director effects on firm performance – The table estimates individual director effects in the regression setting described in Table 1-6. Panel A groups coefficients and reports F-values associated with the hypothesis that all coefficients in the respective group are jointly zero. The groups are Individual director effects, Board characteristics, Voting rights, other firm level time varying control variables (X_{it}), and firm and year fixed effects. Panel B further describes the estimated individual director effects. It provides the overall number and the quartiles of the estimated effects. The median of the dependent variable (Median (y)) is given for comparison. Further, it lists the fractions of coefficients that are individually significant at standard levels. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

<i>Panel A: Fixed effects regressions</i>			
	Return on assets (%)	Return on equity (%)	Tobin's Q * 100
Individual director effects	510'000***	19'519***	74'674***
Board characteristics	1.38	1.12	0.87
Voting rights	4.14***	3.66***	3.01***
X_{it}	2.20*	5.24***	4.92***
Firm fixed effects	75'000'000***	1'600'000***	26'000'000***
Year fixed effects	6.20***	6.50***	8.32***
Adjusted R ²	0.63	0.48	0.76
N (firm years)	1'764	1'775	1'746
<i>Panel B: Director fixed effects coefficients</i>			
	Return on assets (%)	Sales to assets (%)	Tobin's Q * 100
N (directors)	149	149	146
Median (y)	6.83	13.76	123.24
Distribution of coefficients			
1 st quartile	-1.95	-4.82	-27.02
Median	0.11	2.85	-4.85
3 rd quartile	3.86	12.10	31.27
Individually significant coefficients			
$\alpha=0.01$	7.38%	7.38%	17.81%
$\alpha=0.05$	18.12%	16.78%	28.77%
$\alpha=0.1$	29.53%	28.86%	38.36%

Table 1-8

Influence vs. matching – Table 8 randomly assigns each director's board positions to two subsamples. One subsample is then used to estimate director effects as before. The indicators of the second subsample are modified so they equal 1 in the two years before the respective director actually joins the firm. Both types of indicators are included in the regression model known from previous tables. Panel A shows F-tests for joint significance of different groups of coefficients. The coefficients resulting from the modified indicators are referred to as 2-years-before effects. Panel B presents contingency tables for the sign of the standard and the modified effects of the directors. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

<i>Panel A: Fixed effects regressions</i>									
	Return on assets (%)		Return on equity (%)		Tobin's Q				
Individual director effects	1'425.98***		2'262.41***		87'483***				
2-years-before effects	5'104***		3'475***		15'094***				
Board characteristics	1.55		0.67		0.89				
Voting rights	3.38***		4.65***		0.82				
X _{it}	1.27		2.21*		3.48***				
Firm fixed effects	8'109.79***		7'875.23***		19'000'000***				
Year fixed effects	6.77***		6.97***		7.97***				
Adjusted R ²	0.61		0.45		0.75				
N (firm years)	1'764		1'775		1'746				

<i>Panel B: Association of 2-years-before and individual director effects</i>									
2-years-before	Return on assets		2-years-before	Return on equity		2-years-before	Tobin's Q		
	< 0	> 0		< 0	> 0		< 0	> 0	
< 0	11	22	< 0	5	23	< 0	13	17	
	33.33%	66.67%		17.86%	82.14%		43.33%	56.67%	
> 0	20	17	> 0	23	19	> 0	20	20	
	54.05%	45.95%		54.76%	45.24%		50.00%	50.00%	
Total	31	39	Total	28	42	Total	33	37	
	44.29%	55.71%		40.00%	60.00%		47.14%	52.86%	
	Chi ² : 3.04*			Chi ² : 9.53***			Chi ² : 0.31		

Table 1-9

Director effects on firm policies – This table estimates individual director effects on measures of firm policies. Depending on the firm policy under consideration, different control variables were included in the panel regressions: all regression models control for the lagged log of assets. The Investment regression also controls for lagged Tobin's Q. Additionally, all other regressions control for Return on assets. The Diversifying acquisitions model further includes an indicator for whether the firm made any acquisitions in this year, while the R&D, Leverage, Cash, and Dividends regressions control for lagged cash flow. Panel A provides F-values for the joint significance of different groups of variables. Panel B describes the distribution of the estimated director effects, provides the median of the dependent variable (median (y)) for comparison, and describes the individual significance of the director effects. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

<i>Panel A: Fixed effects regressions</i>							
	Investment	Acquisitions	Diversifying acquisitions	R&D	Leverage	Cash	Dividends
Individual director effects	2.94***	1.87***	1.86***	5.35***	6.10***	1.93***	3.05***
Board characteristics	3.15***	2.04**	1.30	0.55	1.36	2.10**	3.17***
Voting rights	1.27	0.33	0.85	2.92***	2.45**	2.04**	4.63***
X _{it}	16.49***	6.82***	98.16***	15.62***	30.04***	13.47***	1.05
Firm fixed effects	2'600.21***	99.72***	77.28***	300.84***	261.73***	1'680.71***	23.86***
Year fixed effects	2.59***	2.67***	1.43	2.45***	5.71***	2.29***	5.46***
Adjusted R ²	0.50	0.36	0.42	0.88	0.70	0.64	0.43
N (firm years)	1'754	1'889	1'889	786	1'799	1'798	1'522
<i>Panel B: Director fixed effects coefficients</i>							
	Investment	Acquisitions	Diversifying acquisitions	R&D	Leverage	Cash	Dividends
N (directors)	147	157	157	63	152	152	126
Median (y)	15.03	0.00	0.00	4.09	31.13	41.46	10.95
Distribution of coefficients							
1 st quartile	-5.05	-0.44	-0.19	-1.00	-7.93	-57.72	-5.02
Median	-0.88	0.03	-0.01	0.29	0.06	5.33	-1.55
3 rd quartile	4.63	0.34	0.18	1.61	7.21	73.46	4.19
Individually significant coefficients							
α=0.01	8.16%	3.82%	2.55%	11.11%	12.50%	3.95%	7.14%
α=0.05	15.65%	11.46%	9.55%	22.22%	23.03%	13.16%	18.25%
α=0.1	24.49%	17.20%	17.83%	36.51%	30.92%	28.29%	26.98%

Table 1-10

Director styles – This table relates the estimated director effects on performance and firm policies to one another. The table shows pairwise Spearman rank correlations. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Return on assets	Return on equity	Tobin's Q	Investment	Leverage
Return on equity	0.74***				
Tobin's Q	0.34***	0.07			
Investment	0.10	0.27***	-0.05		
Acquisitions	0.21***	0.25***	-0.03	0.19**	0.00
Diversifying acquisitions	0.04	0.01	-0.00	0.09	0.03
R&D	-0.02	-0.04	-0.00	0.34***	0.03
Leverage	-0.17**	-0.34***	0.05	-0.19**	
Cash holdings	-0.28***	-0.14*	0.01	0.12	-0.16**
Dividends	0.01	-0.02	0.07	0.08	-0.12

Table 1-11

Director effects and director characteristics – The estimated director effects on firm performance and policies are related to the directors' background. The table distinguishes directors that graduated from a university, hold a doctoral degree, have financial or legal expertise, are Swiss citizens, and were born earlier than the median director. Mean values are reported for the full sample and the subsamples separately. t-tests and Wilcoxon tests evaluate the significance of the observed differences. ***, **, and * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Return on assets	Return on equity	Tobin's Q	Investment	Leverage
<i>All directors</i>					
N	149	149	146	147	152
Mean	0.62	2.91	-0.02	-0.28	-0.89
<i>University graduate</i>					
no (mean)	-0.10	-0.94	-0.12	-0.52	-0.17
yes (mean)	0.71	3.23	-0.01	-0.10	-0.95
t-value	-0.61	-1.03	-0.81	-0.18	0.24
z-value	-0.86	-1.33	-0.81	0.13	0.52
<i>Doctoral degree</i>					
no (mean)	0.70	2.59	-0.03	0.36	0.13
yes (mean)	0.53	3.06	-0.02	-0.73	-2.11
t-value	0.21	-0.19	-0.11	0.81	1.12
z-value	0.00	-0.64	-0.43	0.24	1.31
<i>Financial expertise</i>					
no (mean)	0.98	2.84	-0.04	-1.27	-0.68
yes (mean)	0.55	2.68	0.01	0.35	-0.22
t-value	0.52	0.06	-0.65	-1.15	-0.22
z-value	0.74	0.39	-0.19	-1.20	-0.17
<i>Legal expertise</i>					
no (mean)	0.87	3.68	-0.01	0.12	-1.35
yes (mean)	0.10	0.97	-0.08	-1.01	1.74
t-value	0.72	0.83	0.65	0.66	-1.20
z-value	0.28	0.96	1.06	0.53	-0.94
<i>Swiss citizen</i>					
no (mean)	3.92	11.68	-0.09	-0.94	-1.43
yes (mean)	0.25	1.92	-0.01	-0.21	-0.83
t-value	2.86***	2.45**	-0.54	-0.33	-0.17
z-value	2.92***	2.60***	-0.24	-0.10	0.02
<i>Year born</i>					
After median (mean)	0.46	1.98	-0.04	-0.35	0.10
Before median (mean)	0.82	4.08	0.00	0.06	-2.47
t-value	-0.45	-0.85	-0.43	-0.31	1.27
z-value	0.34	-0.57	0.15	-0.33	1.12

Paper 2:

Evidence from the S&P 1,500

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.

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; Waelchli 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 1999), board interlocks (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. Casual observations from corporate practice, however, suggest that the individuals on the board matter. For example, specialized firms and databases provide help in finding suitable candidates for board positions. Thus, this article evaluates whether some individual directors are associated with positive or negative effects on their firms' performance.

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. Yet another example is that the market may update its assessment of the a firm that succeeds in recruiting reputable board members (see the certification argument of Fahlenbrach, Low, and Stulz (2010)). Second, event studies usually investigate very short observation windows. Efficient markets process new information quickly and correctly. Thus, the share price reflects this information almost immediately. In this setting, a short observation window allows observing the effect of an announcement, without subsequent events confounding the outcome. 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, it is even unclear when the relevant information is disclosed. Candidates for directorship are typically nominated by the board previously to the annual stockholder meeting. However, it is only after the formal 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 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 first and third 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 made difficult by serial correlation or because asymptotic theory may not apply. We evaluate whether the effects we present are economic reality or statistical artifacts by comparing them to the results of a fixed effects analysis on random data. Using four different randomization procedures we generate 400 placebo datasets with which we estimate fixed effects regressions. The outcome of these regressions serves as a benchmark for our previous results (treatment). We expect that the actual treatment data lead to stronger results than placebo data if director fixed effects are economic reality. Should the results from treatment data be indistinguishable from the placebo results we would have to conclude that the director fixed effects are statistical artifacts.

We find that the outcome of this analysis depends on how we generate the placebo data. Using a placebo benchmark that focuses on allowing a large number of possible observations, rather than on replicating the characteristics of the directors of interest in the sample (later referred to as *Simple* randomization), we obtain economically meaningless results. That is, we find significantly stronger director effects with placebo than with actual treatment data. Moreover, these results are very sensitive to relatively minor changes in the randomization method, such as reassigning each director individually to new firms instead of full boards at a time. These contradictory results mostly disappear if we restrict the randomization process so that it more accurately preserves the characteristics of the treatment data (referred to as *Lossless* randomization). In that case, the results from treatment data are either stronger than the placebo results or they differ insignificantly. The sensitivity to small changes in the randomization process is reduced as well. The cost of restricting the reassignment mechanism to preserve data characteristics of the treatment sample is a substantial reduction in the number of possible observations from randomization. In particular, *Lossless* randomization, risks ending up with placebo data that overlap with the treatment data, masking differences between the treatment and placebo estimates.

Overall, the treatment 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; at least not 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, the findings draw a picture of fungible directors. That is, there are no super star directors that firms should be chasing after. Rather, different individuals seem to do similarly well in terms of firm performance. 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 the error terms of the associated regressions are serially correlated. 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 with director data.

The remainder of this paper is structured as follows: In the next section we clarify why we think that individual outside directors might be associated with firm performance. In Section 3 we provide an overview of our approach to measure director effects on firm. Section 4 explains how the sample is constructed and provides descriptive statistics. These statistics focus particularly on the board service of the directors we track and their movements across firms. The section ends with an introduction of the regressions setting. In Section 5 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, methodological issues arise that question those results. We test for these issues and conclude that our test statistics are inflated. Section 6 generates the empirical distribution of the fixed effects and the associated test statistics through repeated random draws. Knowing this distribution allows a more accurate interpretation of our initial findings. In particular, the director effects become insignificant. Section 7 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 the hiring of individual directors could lead to changes in firm performance. To see how that could happen, let's assume first a frictionless world. In this ideal world, a firm (1) has all the information needed 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. Moreover, technology is constant. Firms in such a world would always staff their board with individuals whose influence maximizes firm value. The departure and the hiring of new directors per se would not induce a change in performance. This scenario cannot be distinguished from a world in which directors have no influence at all. Using fixed effects to measure performance contribution of individual directors, we would not measure an effect because board composition would always be at its optimum.

In practice, the world 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 industry experts, and reputation concerns may deter promising candidates). (3) Board composition cannot be adjusted costlessly to exogenous changes, such as changes in technology. 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 members of the founding family to the board, blockholders who seek board representation, and CEOs who nominate interlocking directors or directors 'not rocking the boat' (Hermalin and Weisbach 1998; Adams, Hermalin, and Weisbach 2010). These frictions in the appointment of directors can cause firms to end up with directors who are suboptimally suited for firm performance.¹ This suboptimal influence can lead to variation in board performance which may take the form of director fixed effects (relative shifts in performance).

So far, we discussed why the directors' influence could lead to observable changes in operating performance. However, similar data patterns might result if high-profile individuals are appointed to the board to certify the quality of the firm to the outside world (Fahlenbrach, Low, and Stulz 2010): Assume a firm knows it is headed for a period of above average operating performance. Such a firm might seek to certify the imminent outperformance to outsiders (e.g., the stock market) 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. In these cases, individual directors would seem to bring about positive shifts in performance although the causality is opposite. That is, the presence of a given director is determined by the firm's performance outlook, rather than the director's influence is driving firm performance.

When we discuss 'director effects on firm performance', we do not imply a causal interpretation of the association. Yet, we adhere to this expression for the sake of simplicity. In the next section we give an overview of how we intend to measure director effects on firm performance.

¹ While relatively little is known about suboptimal directors, extant evidence suggests 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. Overview of the research approach

3.1 Testing strategy

This section provides an overview of how we proceed to answer the question whether individual directors matter. At the basis of our analyses lies a panel regression model of the form

$$y_{it} = \delta_{it} + \lambda_{it} + \mathbf{X}_{it}\boldsymbol{\beta} + \mathbf{Z}_{it}\boldsymbol{\Gamma} + \varepsilon_{it},$$

where y_{it} is the outcome of firm i at time t . The outcomes we consider are *Return on assets*, *Return on equity*, *Sales to assets* and *Tobin's Q*. δ_{it} and λ_{it} are fixed effects for firms and years, respectively. Bold print indicates that the respective symbol refers to a vector. \mathbf{X}_{it} is a vector of firm-level, time-varying control variables. \mathbf{Z}_{it} is a vector of K treatment indicators, with K being the total number of tracked directors. That is, the model includes a separate indicator variable for each director. This variable equals one if the respective director k serves on the board of firm i at time t , and zero otherwise. ε_{it} is the error term. We are primarily interested in whether including the individual director indicators (\mathbf{Z}_{it}) improves the explanatory power of our model. Therefore, we test of the joint linear hypothesis

$$H_0: \boldsymbol{\Gamma} = \mathbf{0},$$

where $\boldsymbol{\Gamma}$ is the $(K \times 1)$ vector of individual director coefficients. If individual directors matter, we expect to reject this hypothesis. We will primarily rely on a Wald test but consider other test statistics, too.

This approach follows Bertrand and Schoar (2003) for the case of individual managers. However, two caveats have been mentioned with regard to that test design: First, serially correlated error terms may lead to inflated test statistics (Bertrand, Duflo, and Mullainathan 2004) and second, the fixed effects framework may violate assumptions that are crucial in deriving the Wald test's F-distribution property (Wooldridge 2002).² Thus, we evaluate the validity of this test by replicating our analysis on data with placebo treatments (Fee, Hadlock, and Pierce 2013). That is, we estimate effects for indicator variables that are not actually associated with any treatment. If the Wald test is valid in our setting, we expect to find only insignificant outcomes from these placebo treatments. However, our tests indicate significant effects for the placebo indicators. Hence, the test statistics are inflated and conclusions regarding significance based on the F-distribution are invalid.

As the F-distribution leads to incorrect conclusions about the significance of our test statistics, our next step is to obtain the test statistic's empirical distribution and use that to determine the statistical significance of the director effects we measure. That is, we conduct repeated analyses using placebo treatments and keep track of the outcomes. From the resulting distribution of placebo outcomes we can infer the conditional probability of obtaining the treatment outcome if no actual treatment effect exists. This probability can be interpreted analogously to a p-value and we therefore refer to it as empirical p-value. Ultimately, this empirical p-value is our main measure to evaluate the significance of individual director effects. Other statistics we use to compare treatment and placebo outcomes are the R square, the fraction of individually significant director effects, and the distribution of the director effects. The latter criterion is borrowed from Cronqvist and Fahlenbrach (2009). They compare inter

² Section 5.2 elaborates in more detail on these two caveats.

quartile range (IQR) of the treatment effects to the IQR of placebo effects to evaluate whether the blockholder fixed effects are economically significant. They test for the equality of the distributions using a Kolmogorov Smirnov test.

3.2 Generating placebo data

To implement the testing strategy outlined in the previous section, we have to generate placebo data. This data is supposed to consist of financial and director data that closely resembles the treatment counterparts. We follow the extant literature in generating placebo data through randomization of the treatment data (Cronqvist and Fahlenbrach 2009; Fee, Hadlock, and Pierce 2013). That is, we extract the director indicators from our dataset. Then, we use different procedures to reassign individual directors or entire boards to random firms. The randomized data pretends that directors serve for firms they never actually served for. Thus, this data represents a placebo, rather than an actual treatment, and does not contribute to explaining variation in firm performance.

In choosing a specific procedure to randomize the director data, we have to decide what characteristics of the treatment data should be preserved in the placebo data.³ To understand the differences between random reassignment procedures, it helps to think of the $(1 \times K)$ vector of director data \mathbf{Z}_t as individual data points $z_{kit} \in \{0, 1\}$ located in the three dimensional space spanning across directors (k), firms (i), and years (t). Each of these data points is an indicator that equals 1 if the director-firm-year combination referred to by its coordinates represents active service (i.e., director k served on the board of firm i in year t) and 0 otherwise (see Panel A of Figure 2-1).

Randomization corresponds to relocating the data points by drawing random coordinates (without replacement) while leaving their value (i.e., 0 or 1) intact. If no restrictions are imposed on this relocation, each point in this space has an independent and identical probability of being assigned a value of one. That is, each director-firm-year combination has an independent and identical probability of representing active board service.⁴ Typically, this probability distribution inadequately represents the treatment data at hand. Therefore, restrictions need to be imposed to preserve characteristic properties of the treatment data. The more restrictions there are, the more data characteristics are preserved. However, additional restrictions also reduce the variety of random observations that may result from the process.

Depending on the restrictions imposed, the resulting randomization procedures preserve different characteristics of the treatment data. For example, we may decide to maintain the treatment data's director coordinate (k), while randomly relocating indicators across firms (i) and time (t).⁵ By

³ This question arises with all approaches to generating placebo data. For example, an alternative approach to obtain placebo data would be to simulate data from scratch. Typically, doing so would involve formulating models that characterize financial and director data, respectively. While the models could be parameterized using estimates from treatment data, we would need to determine what characteristics of the financial and director data our models should represent and how the error terms behave.

⁴ The probability equals the total number of active director-firm-years ($\sum \mathbf{Z}_{it}$) divided by the total number of points in the space: number of directors (K) \times number of firms (I) \times number of sample years (T). In our case, this probability equals $19,764 / (2,062 \times 2,417 \times 15) = 0.0264$ percent.

⁵ In addition to entirely preserving one or more dimensions, movements along the three dimensions can be restricted in complex ways. For example, the indicators may be allowed to shift over time (t),

maintaining the director coordinate, we ensure that active board service is not reassigned from one director to another. Thus, this procedure would preserve the total number of years any director has served on boards. However, by randomly relocating board service across firms and years, the procedure ignores, among other things, the fact that board service is correlated over time. That is, the resulting data does not represent the relation between a director serving on a given board today and the probability that the director serves on the same board in the next year. Therefore, this procedure would be ill-suited for our case. Our analyses use two different randomization procedures.⁶

(1) Individual randomization. Randomization only takes place along the firm dimension (i). Along this dimension, randomization is not restricted in any way. In particular, the set of target firms is not limited to those firms that actually have tracked directors. The time and director coordinates, t and k , are preserved from treatment data. Further, all years of active board service in one firm are reassigned to the same random firm. That is, reassignment acts on entire spells of a director on the board of a firm (i.e., director-firm combinations), rather than individual years of board service. This procedure corresponds to reassigning entire stacks of data points from one firm to another (see Panel B of Figure 2-1).

The datasets obtained will resemble treatment data in several aspects: the number of positions a director holds, the correlation of board service within a firm over time, and the directors' busyness are preserved. Despite the restrictions applied in this process, some deviations from treatment data still remain. In particular, the probability of being assigned active board service is independent and identically distributed across all firms. However, as we show in Section 4, treatment data suggests that tracked directors cluster around some firms but stay away from others. Thus, the probability appears not to be independent and identically distributed across firms in treatment data.

(2) Common board randomization. In addition to the restrictions imposed by Individual randomization, we require that all directors of a given firm are jointly reassigned to the same random firm. That is, entire boards are relocated across the firm dimension (i), rather than individual directors. Panel C of Figure 2-1 illustrates this process. It corresponds to relocating arrays of data points that are stacked further along the director dimension (k). This additional restriction ensures that board service tracked in placebo data is clustered analogously to treatment data. Yet, placebo board service may cluster around different firms than the treatment.

When applying these randomization procedures to our data, additional complications arise from dealing with unbalanced panel data (i.e., not all firms are present throughout the sample): Active board service may be relocated to a firm-year for which we have no financial data. In particular, financial data may be missing because the respective firm does not exist yet or has ceased existing in that particular

while preserving dependence across time. For illustration, assume we want to reassign board service that lasted from 1998 to 2002. Under the complex restriction, this board service could get relocated to any years, as long there remain 4 consecutive years of service on a given board. The restriction we use is more strict. That is, this service could only be relocated the years 1998 to 2002 of a given board.

⁶ We use the same draw of placebo director data to analyze all four performance measures. Alternatively, we could generate new placebo data for each dependent variable. However, we see no clear benefit from having individual placebo data for each dependent variable, whereas computation time would increase notably.

year. For example, board service of a director of *The Coca Cola Company* in the 1990s may be randomly reassigned to *Google Inc.* As the latter firm enters our sample only in 2004, the financial data to match this director's board service is missing after the reassignment. Thus, information on this board service does not enter our analysis.⁷ In consequence, the directors of our placebo datasets will appear to be active in fewer firm-years than the treatment directors. We try to avoid differences between the characteristics of treatment and placebo data by re-applying the original criteria for tracked directors (see Section 4.1) to the placebo datasets. Indeed, numerous directors that are part of the treatment data do not pass these sample criteria on placebo data. Thus, our analyses of placebo data cover overall fewer directors than the treatment's. Moreover, according to the placebo data, the directors hold fewer board positions with shorter duration.

Ex ante, it is unclear whether these distortions in the data characteristics are material, in the sense that they affect the outcome of our analyses. Therefore, we generate a different set of placebo data with an additional restriction in place. So far, data points were relocated arbitrarily to any firm in the entire sample. Now, we restrict relocation along the firm dimension (i). In particular, we only allow data points to be reassigned within groups of firms that are active in the exact same years. For example, board service of a firm that exists exactly from 1996 to 2002 can only be relocated to another firm that exists exactly from 1996 to 2002. This restriction prevents board service from being assigned to sample firms that do not exist yet or have ceased existing in a particular year. Thus, this randomization process does not cause a loss of information on board service, and ends up with placebo sample sizes equal to the treatment sample sizes. Therefore, we refer to datasets that were generated in this way as 'Lossless'. Henceforth, the randomization procedures that do not apply this additional restriction are referred to as 'Simple'.

In sum, we use four types of placebo datasets in our analyses. They differ in that reassignment can occur at the level of the *Individual* or the *Common board*. Further, both procedures can be used in their *Simple* form or with the additional *Lossless* restriction in place. *Individual randomization* reassigns each spell of a director with a firm (i.e., board position or mandate) to another randomly selected firm. Hence, generating one set of placebo data involves as many random draws as there are director-firm spells (5,009). This procedure allows for multiple board positions to be reassigned to the same firm, thus, it is a random assignment with replacement. Throughout this article, we sometimes refer to this procedure as reassigning 'individual directors'. We do so as it seems natural that all spells of a director are reassigned if the director serves on multiple boards. In the *Simple* form of this procedure, we ignore whether the firm is active (i.e., our sample provides financial data on that firm-year) during the years the randomly assigned director holds the position. Thus, if a director is reassigned to a firm that is not active in the years of the director's board service, the information on this director is dropped and we move on to reassign the next director.

⁷ Analogously, director years with missing data is sometimes assigned to existing firm-years. However, our data only lists directors' board service and not the absence thereof. Hence, missing director data is by default interpreted as indicating the absence of the director in a given firm-year and, thus, set to zero in preparing the data for the regression analysis.

The *Lossless* procedure avoids dropping any information on board service. One obvious approach to do so would start by reassigning a board position to any random firm. Then, this approach would check whether the randomly drawn firm is active in all years that the spell lasts and, if this is not the case, draw a new reassignment for the same position. This reassignment and check of the outcome would be iterated until the outcome no longer entails a loss of information. However, we do not follow this obvious approach as it has two important drawbacks: First, it is relatively complex in its implementation and slow in its execution because it may require several iterations until an appropriate outcome is achieved. Thus, given that we reassign 5,009 board positions in 200 draws of placebo data, it potentially takes much longer to generate the placebo data than a procedure without iterations. Second, this approach would generate distorted placebo data. In particular, firms that are active in more sample years would have a disproportionate number of placebo tracked directors compared to firms that are active in fewer sample years. The design of our *Lossless* procedure avoids these two drawbacks. The procedure avoids iterations by identifying all potential outcomes that maintain all information on board service before randomly reassigning the directors. Further, it prevents the distortions described above by restricting directors to be reassigned only to firms that are active in the same number of years as the original firms. In sum, recalling that we do not allow directors' activity to shift over time (see above), the set of firms a director can be reassigned to is limited to the set of firms that are active in the exact same years as the director's original firm.

The *Common board randomization* procedure reassigns the spells of all directors of a firm to the same random firm. Thus, the population of this reassignment is groups of director positions. Generating one set of placebo data now involves as many random draws as there are firms in our sample (2,417). Again, we often refer to groups of directors or simply boards for simplicity, implying that service on multiple different boards is reassigned separately. This procedure no longer allows replacement. That is, once a board has been randomly assigned a specific firm, no other boards can be reassigned to this firm. As with *Individual randomization*, this procedure can be applied in the *Simple* or the *Lossless* form. The difference between these two methods is the same as in the case of *Individual randomization* explained above.

Given that *Common board randomization* preserves more characteristics of the treatment data than *Individual randomization*, and *Lossless* preserves more than *Simple*, it may not be obvious why we use all four procedures instead of only *Lossless Common board*. We do so because adding more restrictions reduces the number of possible random matches (Appendix B estimates the numbers of possible random observations for the four randomization procedures). In particular, for about 25 percent of the sample firms, the *Lossless* procedure entails a substantial risk that placebo data overlap with the actual treatments (see Section 6.3 for a detailed discussion on this issue). These overlaps imply that the individual director effects estimated with placebo data contain treatment effects because they are based on partly the same firms. Thus, the placebo director effects could be contaminated and, if individual directors really matter, end up setting a benchmark under the H_0 that is too high. This inflated benchmark means that the statistical power of our tests (i.e., the probability that they reject H_0 if it is false) is reduced. Therefore, we should only use *Lossless* randomization if *Simple* randomization fails to generate placebo data that adequately represent the treatment data.

Comparing our randomization procedures to what previous studies do is difficult as the authors typically provide few details about their randomization procedures. Cronqvist and Fahlenbrach (2009) say that they “[reassign] each blockholder in [their] dataset to a random firm-year observation” (p. 3963). Notably, the authors mention firm-years in the context of randomization, rather than just firms. This wording might indicate that they preserve only the blockholder dimension, while randomly reassigning across firms (i) and time (t). Such a procedure would be ill-suited for director data. By reassigning across time, the procedure would ignore that a director is more likely to be present on a board if she is present on the same board in the previous year. However, such a procedure might be reasonable to randomize blockholder data if equity investment is relatively mobile. However, later on the same page, the authors mention that “blockholdings [are] randomly distributed across firms.” Thus, they might reassign only across firms, while preserving the time dimension. In this case, their procedure would most likely resemble the *Individual randomization* procedure we use. Fee, Hadlock, and Pierce (2013) only say that they “randomly assign each CEO-to-CEO mover to a different hiring firm” (p. 593). This statement indicates that the procedure is closely related to our *Common board randomization* procedures. *Individual randomization* does not seem sensible for CEOs, as it could generate outcomes where multiple CEOs serve for the same firm simultaneously.

Unfortunately, neither of the two articles provides descriptive statistics of the placebo data, mentions a loss of information, or discusses any additional restrictions. Thus, it is unclear whether and how they address the additional difficulties posed by unbalanced panel data (e.g., our *Loss/less* modification).

In the next section, we discuss our sample and the data which we will use to implement the methods outlined above.

4. Sample and data

4.1 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. Hence, we design our own director identification procedure 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 (y_{it}) and the time-varying firm-level control variables (\mathbf{X}_{it}).

Starting with the combination of these two datasets, we exclude two types of firms: (1) firms from the financial and utilities sectors, (2) firms that are in the sample for less than three years.⁸ Further, a

⁸ This follows from limiting our analysis to directors who are present in their firms for at least 3 years (see below).

number of firms-years, drop out of our regression analysis because they have missing data on the control variables (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 of overall tenure. We do so because it seems unlikely that directors influence their firms immediately after their appointment. Second, 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.

4.2 Distribution of tracked directors across firm-years

Table 2-1 describes how the directors we track are distributed across sample firms and years. Panel A shows the number of different sample boards a director serves 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 either simultaneously or over time. 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 simultaneously serves on. In total, our sample covers 19,764 active director-year combinations; the majority of these refer to 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 any 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 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 at the same time 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,⁹ 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 2-1 takes the directors' perspective to discuss how board service is distributed across the sample, Table 2-2 takes the firms' perspective on this distribution. Panel A describes the distribution of the number of directors we track in each 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

⁹ 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 fiscal year 1996. The last directors were reelected in late 2009 and thus were present until fiscal year 2010. However, no single firm can be in the sample for more than 14 years.

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).

4.3 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 might indicate that not all firms are equally likely to recruit one or more tracked directors. Table 2-3 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, respectively. The columns show the number of observations, and the 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.¹⁰ These observations divide into 4,112 firm-years without any tracked directors and 11,027 firm-years with one or more tracked directors. From Panel B of Table 2-2 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. We have also constructed subsamples based on whether firms ever have a tracked director. The outcome is overall comparable but the difference between the subsamples increases for most variables. The table also displays variables that do not enter our regressions (e.g. *Sales growth*, *Age (incorporation)*). For these variables the number of observations may be lower than the 15,139 firm-years that constitute our regression sample.

The table first shows the measures of firm performance: *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, than in the full sample. Firm-years with tracked directors are also more efficient. They generate 3.53 cents more sales for every dollar of assets they employ than firm-years without tracked directors. However, this difference is small relative to the variation within the

¹⁰ *Board ownership* and *Board tenure* are often missing. We avoid losing numerous observations by substituting zero for missing values and include an indicator variable that marks these substitutions in the regressions. This substitution biases the descriptive statistics in this table towards zero.

subsamples. For *Tobin's Q*, the mean values of the two subsamples are comparable while the Wilcoxon test 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.

The extant literature provides several findings that might help explaining these differences in performance. First, the differences could result from certification: According to Fahlenbrach, Low, and Stulz (2010), high-quality firms certify their outlook by adding particularly reputable individuals to their boards. Given our selection criteria, the directors we track may qualify for this purpose. Second, Field, Lowry, and Mkrtychyan (2013) document that busy directors (i.e., directors who hold many positions simultaneously) have a positive impact on young firms (post IPO). Many of the firms and directors they investigate are probably part of our sample. Further, many of the directors we track would qualify as busy. Thus, the differences we observe are in line with their findings. Third, the differences may relate to the lifecycle of the firms. For the case of CEOs, there is evidence that firms at different stages of their life cycles hire different types (Custódio and Metzger 2014). The firms' preferences with regard to directors may change similarly as they mature. Indeed, Table 2-3 shows that the firms who hire tracked directors are significantly larger (in terms of total assets) and older; they grow slower, and invest less. Moreover, recruiting high-profile individuals, such as the directors we track, may be only possible for large firms (Gabaix and Landier 2008). At the same time, only reputable directors with a proven track record in smaller firms may have access to positions on these firms' boards. If the difference is indeed related to the firms' lifecycle, it would likely disappear once we control for other firm characteristics. However, if either of the first two explanations is true, we would expect to also find the difference in performance in a multivariate setting. In Section 4.4 we introduce such a setting.

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%). These differences might arise endogenously from our sample selection criteria. The likelihood that at least one director of a firm makes our sample 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. It is unclear what this finding means for our analysis. As we analyze only non-executive directors, we might exclude the most important board member if the COB serves as CEO as well. However, this concentration of power might also indicate that monitoring by independent directors is particularly important. Further, the board of firm-years with tracked directors owns a smaller fraction of the firm's stock in terms of voting rights. This 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.

Overall, the relation between firm characteristics and the presence of tracked directors further supports the notion that tracked directors cluster around specific firms but 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 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 favor highly performing individuals and bias our sample towards estimating overly positive director effects. An alternative argument is that firms have different expectations regarding their board's 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 easily generalize. After all, these particular directors' service is in demand by the majority of the S&P 1,500 firms. Further, this demand is particularly high in larger, arguably more sophisticated, and financially powerful firms.

4.4 The regression setting

Table 2-3 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 individual directors' contribution. Hence, this section evaluates whether the performance difference persists in the multivariate panel regression setting we use in our further analysis (Table 2-4; see Section 3 for a formal introduction of the regression model). The section also discusses the estimates of the control variables' coefficients and compares them to previous findings.

The dependent variables in our 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). We add the lags of *Investment* (Yermack 1996) and *Leverage* (Ang, Cole, and Lin 2000; Faleye 2007) to the control variables used in these articles. 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 the relations. The last board characteristic we control for is an indicator variable for firm-years where the CEO serves simultaneously as COB.

The sample split from Table 2-3 is reproduced with an indicator variable that equals one in firm-years with tracked directors and zero otherwise (*Firm-year tracked indicator*). This variable will later be replaced by the set of variables that indicate the presence or absence of each individual tracked director (Z_{it}). The insignificant coefficient of this variable across all performance measures shows that the performance difference observed in Table 2-3 disappears in a 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 superior 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 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). The indicator for firm-years where a single person holds the positions of CEO and COB is an exception in that its 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 2-2 compares our polynomial 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 effect). The solid line represents our own 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 large boards. However, these differences may be artifacts occurring because only few observations lie in these areas. Indeed, in our estimate all 210 firm-years with board sizes greater than 14 have negative residuals. Thus, the quadratic shape of our regression function leads to a poor prediction of the few extremely large outcomes. In particular, the positive coefficient on the square term of *Board size* should probably not be interpreted as reversing the negative primary effect. Rather, the primary effect seems to weaken for large boards. This interpretation also seems in line with intuition, as it suggests that increasing board size from 6 to 7 directors has a larger impact than going from 13 to 14. The predictions for board sizes smaller than 4 are not interpreted since the data cover no such observations.

The financial control variables are significantly related to all measures of firm performance. We find negative associations 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 of Bertrand and Schoar (2003) and Cronqvist and Fahlenbrach (2009) because neither of them reports individual coefficients for their performance regressions. If we include studies that estimate similar regressions without fixed effects in our comparison, the estimated coefficients vary substantially. Even within studies the coefficients can change from significantly positive to significantly negative depending on how the dependent variable is defined (e.g., 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 and block holder fixed effects, respectively. The explanatory power of these two studies is significantly different (between 0.57 and 0.72). This difference might be due to the different

time 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

5.1 Estimating director fixed effects

Overall, the discussion of Table 2-4 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 those of the extant literature. Table 2-5 replaces the indicator variable for firm-years with tracked directors by indicators associated with the individual directors. 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 2-5 tests for joint significance of groups of coefficients (Panel A). It reports test statistic associated with the hypothesis that the respective group of coefficients is jointly zero (Wald test). Put differently, the null hypothesis is that the respective group of coefficients does not contribute to the model's explanatory power. Under certain assumptions, these Wald statistics are F-distributed. This distribution is the basis for the asterisks in the table, exhibiting significance of the respective test. The groups of coefficients are *Individual director effects*, *Board characteristics*, *Firm level financial control variables*, *Firm fixed effects*, and *Year fixed effects*; these groups include the variables introduced in Table 2-4.

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 may seem small at first. However, given that the models without individual director coefficients already explain a large part of the depending variable's variation, they are noteworthy after all. This finding compares to the outcome of previous fixed effects studies. Including manager and blockholder fixed effects, respectively, 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 barely or not at all significant. This latter finding is in line with the observations made in Table 2-4.

Panel B describes the distribution of the group of coefficients and the individual significance of the director effects. 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 directors' indicators are omitted from the regressions because they are linearly dependent on other directors' indicators and firm fixed effects. For example, this could occur if a director remains active for the full sample period on every board she serves. In that case, it would be impossible to distinguish this director's fixed effect from firm fixed effects. Thus, we would omit the director's indicator. The table 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 roughly in line with the observation that having a director who qualifies for tracking does not, on average, influence firm performance (see discussion of Table 2-4). However, the effects' interquartile range (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 a 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 their joint significance. Panel B of Table 2-5 evaluates how many of the estimated effects are individually significant. It does so by providing 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 the performance measure. 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. This indicates that the observed effect is widely spread among the directors in our sample.

5.2 Issues with the testing strategy

The Wald statistics in Table 2-5 are strikingly high compared to what we observe in similar studies. Cronqvist and Fahlenbrach (2009), for example, find test statistics of 1.27 and 1.75 when assessing the importance of large shareholders on firm performance. The statistics from Bertrand and Schoar (2003) are somewhat larger but they are still small compared to the ones we find (the highest is 53.48). This difference might partly be due to the greater number of individuals tracked in our analysis. However, these high test statistics might also indicate issues with the testing strategy we use so far. 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. Consequently, the null hypothesis is often rejected although it is true. 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 (see Section 3). We note two differences: First, Bertrand, Duflo, and Mullainathan (2004) consider the case where treatment is administered within groups (e.g., at state level), whereas our treatment is not group-specific. Consequently, we control for fixed effects of the firm (δ_i), rather than state fixed effects. Moreover, in their model all firms of a given state receive the same treatment at the same time. In contrast, our treatment being the presence of a director k , our sample firms can in principle receive up to K different treatments simultaneously or over time. Indeed, many sample firms receive multiple treatments simultaneously. The second major difference is that, in the typical case described by Bertrand, Duflo,

and Mullainathan (2004), the treatment indicators remain constant once treatment has occurred. Our director indicators differ in that they return to zero once a director leaves a board.

Despite these two 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 2-6 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 could 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. Table 2-7 estimates the fixed effects regressions from Table 2-5 adding the lagged dependent variables. The rows of Panels A and B set in italic print compare how this change affects our estimates relative to Table 2-5. In the *Sales to assets* and *Tobin's Q* regressions, the Wald statistics for the group of director indicators drop. For the other two performance measures, however, the statistics 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 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 2-8 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. For this case, Bertrand, Duflo, and Mullainathan (2004) suggest using a block-bootstrapping method to generate the empirical distribution of the test statistics. This technique is closely related to the placebo benchmarking we apply in Section 6 of this article.

Independent of the potential serial correlation problems, the second methodological issue is that Wald test statistics may not follow an F-distribution in the context of fixed effects: 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 (Wooldridge 2002). Fee, Hadlock, and Pierce (2013) propose testing this concern by repeating the fixed effects analysis with placebo data. In particular, they suggest generating placebo treatments by randomly recombining the director indicators and firms' financial data. Using this scrambled data we expect not to find a significant effect of the placebo treatment, that is, of the randomized presence of individual directors. If including the placebo treatment significantly improves the fit of the regression model, we should conclude that the tests used are invalid.

The first step in placebo testing is generating a randomized dataset. Section 3 of this article discusses in detail what we mean by ‘randomization’. In particular, it presents the two different procedures we use to generate placebo data (i.e., *Individual randomization* and *Common board randomization*) and what we mean by *Simple* as opposed to *Lossless* randomization. Placebo data generated by the *Simple* versions of these two reassignment procedures are in Table 2-9 to estimate fixed effects regressions. The results based on the two alternative randomization methods are shown side by side. It turns out that both procedures lead to similar outcomes. As in Table 2-5, Panel A conducts Wald 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 randomization ensures that the directors do 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.

Given our discussion of potential issues with the testing strategy, this significant outcome also means that the Wald statistics are not F-distributed. Thus, we should not rely on significance alone. Indeed, a closer comparison between Table 2-5 and Table 2-9 shows that the results differ: The Wald statistics and the estimates’ explanatory power from placebo data tend to be lower than the ones from treatment data. If outcomes consistently differ between treatment and placebo analyses, this would indicate that director effects are real, after all (Bertrand, Duflo, and Mullainathan 2004). Therefore, the next section evaluates whether these differences are systematic. It does so by repeating the placebo analysis on a larger number of random draws. From this repeated analysis, it generates the empirical distribution of the test statistics and uses this distribution to evaluate significance.

6. Benchmarking against the empirical distribution

6.1 Improving comparability and obtaining the empirical distribution

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 draw. We learn more if we analyze multiple placebo datasets. Therefore, we now repeat the random reassignment and analyze 100 different draws of the placebo data.

Table 2-10 provides four rows of information for each dependent variable. The respective first row (*Treatment (full)*) summarizes the outcome from the treatment analysis in Table 2-5. The third and fourth row each show the outcome of 100 repeated analyses of placebo data generated by the *Common board* and *Individual randomization* procedure in their *Simple* form, respectively. As we note in Section 3, the *Simple* form of the randomization procedures leads to a loss of information as active director service may be randomly reassigned to firm-years with missing financial data. In Table 2-10, this loss shows in the *Number of directors* for whom we estimate individual effects. The number drops from 1,854 in treatment to about 600 in placebo data. This drop occurs because some of the board service in our sample is reassigned to inexistent firm-years. As this service now refers to missing financial data, the information on it drops out of the sample. Directors affected by this loss may no

longer meet the sample criteria (serving on 2 or more boards for at least 3 years). For example, the director's mandates may be cut to less than three years or be reassigned to missing firm-years entirely. Due to this difference in the number of coefficients tested, we cannot simply compare the test statistic of treatment and placebo data estimates. The number of effects tested can be aligned either by reducing it in the treatment data analysis or by raising it in the placebo data analyses. The first approach is pursued in the second row of each performance measure in Table 2-10 (*Subset*). The second approach is represented by the *Lossless* modification to the randomization procedures. The outcome of this analysis presented in Table 2-11 and discussed in Section 6.3. *Subset* deliberately reduces the number of director effects in the treatment by randomly excluding 67 percent of director indicators from the analysis. In particular, each tracked director is assigned a random number and the indicators of all directors whose number is lower than the 67th percentile are excluded. To avoid analyzing one specific random observation, 100 random draws of *Subset* are generated.

Table 2-10 shows that the average *Number of directors* included across these 100 draws is 619. Thus, it is relatively close to the 608 and 620 director effects estimated on the different placebo datasets. A perfect match is impossible because the exact number of directors included in *Subset* analysis varies slightly between draws. This variation arises because some of the director indicators are linearly dependent of one another and of firm fixed effects. This linear dependence is resolved if one of the linearly dependent indicators is coincidentally chosen for exclusion by the random *Subset* procedure. If none of the linearly dependent indicators is excluded by *Subset*, one of them needs to be omitted in the regression estimation. Thus, the exact number of indicators in the regressions depends on the realization of the random *Subset* procedure.

The table then investigates the test statistics associated with the hypothesis that all director effects are zero (Wald test). Since the estimates were repeated on 100 draws, the table describes the distribution of the test statistic across these draws. For two of our performance measures (*Sales to assets* and *Tobin's Q*), the results of the full sample treatment analysis are significant outliers compared to the *Subset* analysis: only 6 and 5 *Subset* draws produce Wald statistics larger than the full set for *Sales to assets* and *Tobin's Q*, respectively. This finding is in line with the notion that we should not compare Wald tests with different degrees of freedom, such as one from the full sample treatment estimates and the placebo estimates. Thus, for the remaining discussion of Table 2-10 we refer to the *Subset* analysis when comparing treatment to placebo outcomes.

Comparing the Wald statistics from the treatment effects with the ones from *Common board* placebo effects, the placebo effects tend to produce consistently higher outcomes. All quartiles of the distribution of Wald statistics are higher in placebo data for all measures of performance. Moreover, the Kolmogorov Smirnov test (KS test) and the Wilcoxon test both reject the null hypothesis of no difference. The sign of these significant differences is surprising. We would expect the statistics from treatment data to be higher if directors have an impact on performance. However, the results suggest that there is a stronger association between individual directors and firm performance in randomly selected firm-years. This finding makes no sense.

These illogical results mostly disappear when we look at the results generated with the *Individual randomization* procedure instead. For *Return on assets* and *Sales to assets* the treatment now leads to higher Wald statistics than the placebo. This finding is in line with our expectations if there are true director fixed effects. For *Tobin's Q* the two distributions of test statistics are indistinguishable. The only performance measure, for which the Wald statistics from placebo data tend to be larger than the ones from the treatment, is *Return on equity*. However, the Wilcoxon and Kolmogorov Smirnov tests disagree with regard to the significance of this difference.

Another statistic to compare treatment and placebo estimates is the R square. We expect the director indicators in the treatment sample to have greater explanatory power than the placebo indicators. Comparing the treatment and the *Common board* 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 *Individual 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 true director fixed effects.

Cronqvist and Fahlenbrach (2009) distinguish treatment from placebo results by looking at the distribution of the estimated fixed effects. Based on finding a wider dispersion of the treatment effects compared to placebo effects, they conclude that the effects do indeed exist and are economically significant. The last four columns of Table 2-10 replicate this comparison for our director fixed effects. Comparing the treatment with the placebo data from *Common board randomization*, the Kolmogorov Smirnov test finds a significantly smaller dispersion on the treatment for all performance measures. As with the distribution of Wald statistics, this finding suggests that the randomized association of directors and performance is stronger than the association with performance in firm-years they were actually present. This finding, too, is difficult to reconcile with economic intuition. As with previous criteria, the findings of *Individual randomization* are less problematic. Here, the counterintuitive outcomes result only with two of the performance measures (*Return on assets* and *Tobin's Q*).

In conclusion, Table 2-10 shows that when using *Common board randomization*, we find substantial evidence suggesting that the placebo effects are stronger than treatment effects. This finding is illogical and indicates a problem with the placebo data. It is even more surprising as the *Common board randomization* procedure is supposed to better represent the treatment data than *Individual randomization* (see Section 3). Therefore, the next section investigates this counterintuitive finding by looking more closely at the placebo data.

6.2 Investigating the counterintuitive findings

The illogical findings when using *Common board randomization* and the fact that many results depend on which randomize procedure we apply are worrisome. There are multiple explanations for why our analysis may be problematic: 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 treatment data. Appendix B estimates the number of possible outcomes to be $e^{16,418}$.

Given this immense number of possible random observations, it is very unlikely that significant parts of our placebo data are identical to the treatment 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 treatment 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 2-3 looks at what happens to the differences between treatment 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 these 100 draws, we might have to add further repetitions until the results turn stable.

Figure 2-3 plots the key statistics of Table 2-10 and their development as the number of draws increases. The dark lines compare the results from the *Subset* analysis with the results from *Common board randomization*. The light lines look at *Individual randomization*, instead. For the comparison of Wald statistics (left column), the figure displays the p-value of the Kolmogorov Smirnov test (solid lines, primary vertical axis) and the difference between the median of the treatment and placebo test statistics (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, Panel A of Figure 2-3 looks at the *Return on assets* regressions. The left chart compares the distributions of treatment and placebo Wald statistics. For both randomization mechanisms, the p-value testing the equality of these distributions drops below 0.1 after around 30 draws (solid lines). For *Common board* randomized placebo data the difference between the distributions remains significant thereafter (dark line). For data randomized using the *Individual* procedure, the p-value rebounds and reaches a peak of about 0.12 after 77 draws (light line). Subsequently, it drops quickly to almost zero. Thus, concerning p-values, our conclusion is very stable when looking at *Common board* placebo data but somewhat less stable when looking at *Individual* placebos. Quite the opposite is true for the difference between the medians of the Wald statistics (dashed lines). Here, it is the *Individual randomization* placebo that produces stable results, while the *Common board* placebo outcomes fluctuate more. However, the sign and the order of magnitude of both differences remain steady after about 20 draws. Comparing the estimated individual director coefficients we find a similar picture (right chart). The Kolmogorov Smirnov test for the equality of the distribution of treatment and *Common board* placebo outcomes is significant from the first draw and remains so ever after. For director level data, the test turns significant after 24 draws and is mostly significant thereafter, except for a 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 *Common board randomization*. For *Individual randomization*, it is less clear whether we have enough random draws. However, as the numbers are relatively stable over the last 25 draws, it seems unlikely that our conclusions would change if we add more random draws.

Looking at the other performance measures, the outcome of the *Common board* placebo data is stable (dark lines). Since this randomization procedure is the source of the results we are worried about, it seems unlikely that adding more random draws would change these results. Thus, having not enough random draws does not seem to explain our counterintuitive findings. For the data obtained from *Individual randomization*, the outcome of comparing individual director coefficients is also stable. However, comparing Wald statistics of *Return on equity* regressions, the p-value from the Kolmogorov Smirnov test crisscrosses the threshold of 0.1. That is, it changes between significant and insignificant. While this indicates that our conclusion might change if we add more random draws, this change would not resolve the counterintuitive findings.

The third issue that might cause the counterintuitive findings is the structure of the placebo data. As we discuss in Section 3, the benchmark is valid only if the placebo data accurately represent the treatment data in all material aspects. 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 board service is highly serially correlated. That is, whether a director serves on a board in a given year significantly affects the probability that she serves on the same board in the following year. As discussed in Section 3, our randomization mechanisms are specifically designed to maintain this feature. 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 in, and the duration of the mandates. These features are potentially affected by randomization. Figure 2-4 analyzes whether this is the case. For each feature it compares the two kinds of placebo datasets (squares represent *Common board* and triangles *Individual randomization*) to the full sample treatment data in the first and the *Subset* treatment 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 treatment and placebo data. Both kinds of placebo data sets have a greater fraction of firm-years without any tracked directors than the full treatment dataset. Partly, this is the case because the total number of tracked directors is smaller in the placebo data. To counter this effect, we introduced *Subset* to roughly align the overall number of tracked directors in treatment and placebo data. Accordingly, when comparing *Subset* with placebo data the differences decrease considerably (second column of Panel A). However, the placebo data still show more firms without any tracked directors than the treatment data. Knowing that the number of directors is adjusted quite accurately (see Table 2-10), this observation implies that randomization decreases the number of mandates per director or the duration of these mandates. Indeed, Panels B and C demonstrate that the directors are tracked on fewer boards and during fewer years in the placebo in the treatment data. These differences result from randomly reassigning board service to inexistent firm-years. Thus, while we try to avoid the most drastic changes in data structure (e.g., directors serving during less than three years or on single boards) by requiring the placebo treatments to pass our original sample criteria, substantial distortions remain.

Figure 2-4 also allows comparing the outcomes of *Common board* and *Individual randomization*. The figure shows substantial differences between the two methods. As we expected, *Individual randomization* distributes tracked directors more evenly across the sample than *Common board*

randomization, which preserves the distribution of the treatment data (Panel A). This distortion results 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 *Common board randomization*. We argued earlier that *Common board randomization* preserve characteristics of the treatment data more accurately than *Individual randomization*. Surprisingly, the opposite appears to be the case. The distortions from breaking up boards turn out to offset part of the distortions from reassigning directors to inexistent firm-years. By chance, these counteracting effects lead to a closer match in the frequency of firm-years with few directors. The only advantage of *Common board 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 identically (see Panels B and C).

In sum, both *Individual* and *Common board randomization* appear to distort the characteristics of the treatment data significantly when applied in their *Simple* form. These distortions are caused by the loss of information when board service is reassigned to sample firms that do not exist yet or have ceased existing in a particular year. In the next section, we prevent such reassignments from occurring using the *Lossless* procedure introduced in Section 3.

6.3 Benchmarking against different placebo outcomes

We repeat the analyses from the previous section, using the *Lossless* modification of the randomization procedures. That is, we restrict the set of firms which data points can be reassigned to. In particular, reassignment takes place only among groups of firms that are active during exactly the same years. For example, board service in 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 exactly from 1996 to 2002.

Figure 2-5 illustrates the outcome of building groups according to this criterion. It describes the fraction of firms that are reassigned within groups of a given size. The median firm's board service is reassigned within a group of 35 firms. The largest group within which board service is reassigned contains 230 members. Overall, about $e^{5,663}$ and $e^{5,663 \times 2,062}$ random observations are possible for *Common board* and *Individual randomization*, respectively (see Appendix B). While these numbers are much smaller than their respective counterparts from *Simple* randomization, they are still immense. However, the size of these numbers is misleading with regard to the risk of generating overlaps between treatment and placebo data. The numbers are mostly driven by the reassignment possibilities in the larger groups of firms. Yet, about 8 percent of the sample firms are in atomic groups. That is, these firms' board service cannot be reassigned to any other firm under *Lossless* randomization. Thus, for these firms the treatment and *Lossless* placebo data necessarily overlap. Another 15 percent of the sample faces a substantial risk of overlaps as they are in groups of 2 to 7 firms. Hence, overlaps of treatment and placebo are real in this part of our sample. As these overlaps allow the hypothesized treatment effect to contaminate the *Lossless* placebo benchmark, the benchmark is overly conservative. Thus, the probability that we fail to reject the null hypothesis (H_0) although it is false is higher than when using an entirely random placebo benchmark. In other words, the tests relying on the *Lossless* procedure have a relatively low statistical power.

While the *Lossless* restriction creates the issue of overlaps, it also prevents the loss of information we suffer when active board service is reassigned to firm-years with missing financial data. Therefore, the placebo data more accurately preserve the characteristics of the treatment data. Figure 2-6 compares the data characteristics between the treatment data and the placebo data produced by the *Lossless* randomization procedures. The *Lossless* placebo data represent the treatment data more accurately than the *Simple* placebo datasets in Figure 2-4. 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 preserved. Panel A shows that *Common board randomization* also exactly preserves the original distribution of the number of tracked directors per firm-year. Meanwhile, *Individual randomization* flattens the distribution of directors across firm-years (i.e., fewer extreme and more moderate cases). This outcome is in line with the expectations we formulate in Section 3.

Table 2-11 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 ones from the original data quite accurately. Somewhat surprisingly, even more director effects are estimated for the *Individual* placebo than the treatment. This difference results because *Individual randomization* distributes directors more evenly across the sample relative to the treatment data. A more even distribution has fewer director indicators that are linearly dependent of each other and of firm fixed effects. Thus, fewer indicators need to be excluded when estimating the regressions (see our discussion of why the number of effects in *Subset* varies between draws in Section 6.1). Nevertheless, the degrees of freedom in the treatment and placebo tests match relatively precisely. Hence, the full sample results can be directly benchmarked against the placebo results. This direct comparability simplifies our analysis of the Wald statistics (column 4). Instead of comparing two distributions, as in Table 2-10, we can now compare a single value to a distribution. That is, we can simply count the number of random draws that produce a greater Wald statistic than the treatment and interpret the fraction of draws that meet this requirement analogously to a p-value. According to this empirical p-value, the Wald test for *Sales to assets* is the only significant. For all other performance measures the empirical 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 treatment estimates for *Return on equity* and *Sales to assets*. For the other two measures, our model has higher explanatory power with placebo data than with treatment data. However, the differences between R square treatment and placebo estimates are minute.

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 with treatment data. This finding suggests that director effects on *Return on equity* are significant, which is at odds with the results from comparing Wald statistics, where ROE was the performance measure with the least significant outcome. For the other three performance measures, the director coefficients are more widely distributed in the placebo estimates. However, according to the Kolmogorov Smirnov and the Wilcoxon test, these latter differences are insignificant.

Overall, Table 2-11 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 disagree with each other. That is, the distribution of Wald statistics suggests significant treatment effects only for *Sales to assets*. The R squares criterion shows barely any difference between treatment and placebo effects. Finally, *Return on equity* is the only performance measure where treatment effects are significantly more widely distributed than placebo effects.

Figure 2-7 investigates whether Table 2-11 covers enough random draws to obtain stable results. The charts in the left column show the change in empirical p-value of the Wald statistic as the number of draws increases. Thus, the charts are simpler than the ones in Figure 2-3. The reason is that we compare a single value to a distribution rather than two distributions. The comparison of individual director coefficients (second column) is analogous to Figure 2-3. The figure shows that the empirical p-values resulting from the distribution of Wald statistics as well as the p-values from Kolmogorov Smirnov tests are very stable. Neither of them crisscrosses a typical level of confidence. The differences in IQR of the estimated effects are stable as well. In particular, the sign of neither of these differences changes after about 10 random draws. Thus, we expect the findings of Table 2-11 to be robust to adding more placebo draws.

Comparing the *Simple* randomization methods used in Table 2-10 to the *Lossless* ones in Table 2-11, we note that the two issues we worried about earlier in the paper have largely disappeared. First, *Common board* and *Individual randomization* typically agree in Table 2-11. Second, Table 2-11 provides no evidence that the placebo effects are significantly stronger than the treatment effects. Both improvements result from additionally including the *Lossless* restriction in the randomization process. That is, the placebo data in Table 2-11 more accurately mimic the original data. The cost of this better representation is that the placebo data generated by the *Lossless* procedure partly overlap with the treatment data. Ultimately, this reduces the statistical power of our tests. We considered the *Simple* randomization procedures as examples of more powerful methods. Unfortunately, they fail to accurately preserve important characteristics of the treatment data. Therefore, they lead to illogical results and cannot be applied to directors.

7. Conclusions

The extant literature documents that certain outside directors influence firm performance. 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 changes in firm performance. That is, we look for individual outside directors that have a measurable systematic effect (positive or negative) on the 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 presence of individual directors is not, on average, associated with firm performance. Hence, directors do not

appear to be super stars or losers. This finding questions the benefits from trying to go after high-profile candidates for board positions. It seems that staffing the board with different individuals would lead to similar firm performance. That is, the directors we track appear to be overall fungible.

An alternative interpretation of our findings could be that directors do actually matter but that firms always staff their board with individuals whose influence maximizes firm value. In this case, the departure and the hiring of new directors per se would not induce a change in performance. However, this interpretation relies on a frictionless market for directors. The extant literature provides evidence for a number of practical frictions in this market. They include staggered boards, non-compete agreements between candidates and their previous employers, scarcity of talented directors, and agency problems inside the nominating firms. Thus, we doubt that firms manage to optimally staff their boards at all times.

Nevertheless, our findings do not entirely rule out the possibility that some directors matter. In particular, we should investigate different samples of directors. It might be argued that the super star directors serve on single boards only. In this case, our sample excludes would exclude them. However, this view is at odds with the literature suggesting that successful directors are offered multiple mandates. If the latter notion is true and the directors actually accept these mandates, we would have them in our sample. However, our empirical method might not be powerful enough to detect them if only a small fraction of the directors we track are super stars. This problem can be reduced by looking at particular groups of directors with special characteristics such as reputation or educational background.

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 Wald statistics associated with 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 do not actually work for. This outcome does not make sense, which implies that the test statistics are inflated and the tests invalid in our setting.

In principle, using repeated draws of placebo datasets allows learning the empirical distribution of the test statistics and, thereby, gaining insights into the economic significance of the estimated fixed effects. However, 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 procedure is appropriate, we find it instructive to think of the procedures as applying different restrictions to the random reassignment. The more restrictive the randomization process, the more accurately the placebo data will represent the structure of the original data. For our director data we consider four characteristics: (1) the correlation of board service within a firm over time, (2) the board size, (3) the busyness of directors, and (4) the duration of mandates.

We identify restrictions which ensure that the randomization mechanisms accurately preserve these characteristics. However, adding restrictions reduces the possible number observations from randomization. For example, unrestricted randomization (i.e., relocating individual observations of board service across time, firms, and directors) would lead to the highest number of possible outcomes. However, placebo data generated this way would disregard all four characteristics of the treatment data mentioned above. As restrictions are added, the possible number of random observations decreases and it becomes more likely that the placebo data overlap partly with the treatment data. In our case, randomization needs to be severely restricted for it to preserve all four characteristics mentioned above. As a result, more than 20 percent of our placebo data face a significant risk of overlapping with treatment data. It follows that the benchmark statistics are higher than if the data were truly placebo. Thus, we compare our treatment effects to an excessively conservative benchmark, which decreases the statistical power of our tests. To our knowledge, we are the first ones to illustrate this trade-off between accuracy and power when fixed effects are benchmarked against placebo 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 individual director effects (e.g., that the effects possess out of sample predictive power). 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) As Fee, Hadlock, and Pierce (2013) are somewhat unclear about what placebo data they use, we will replicate Bertrand and Schoar's (2003) original work on manager fixed effects and benchmark the treatment effect against our placebos. Doing so will provide additional insights into the workings of our testing strategy. In particular, as managers are more closely involved with the firms' operations than directors, they are more likely to have a tangible influence on firm performance. Thus, our method might detect manager effects even if it is not strong enough to detect director effects. Moreover, managers typically work for a single firm at a time. Consequently, there are fewer data characteristics to be preserved in randomization. In this setting, the less restricted randomization methods might turn out to be useful. (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 will detect 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) divided 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 (mkvalt) less common equity (ceq) and deferred taxes (txdb). Winsorized at the 0.5% level.
Investment	Capital expenditures (capx) 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 liabilities (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 date 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 from RiskMetrics / IRRC). If this variable is missing for some directors, these directors are assumed to have no 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.

Appendix B: Number of possible random observations

This appendix provides the numbers of possible observations for different randomization procedures. While these computations refer to all random observations possible, it could be argued that the number of truly distinct observations possible is more relevant. As our dataset is sparsely populated (i.e., many data points have a value of 0), the numbers of distinct observations will indeed be significantly smaller. However, the possible number of distinct outcomes is difficult to compute as it depends on the values of each indicator variable for board service (Z_{it}). We expect that the ratio between the number of all possible observations and the number of unique possible observations is particularly high for *Individual randomization*.

- *Simple Common board randomization* relocates the entire board of a firm at once. As there are 2,417 firms in our sample, the first board can be relocated to 2,417 different firms, the second to 2,416 firms and so on. In total, there are 2,417! possible random observations. According to Stirling's formula this factorial approximates to $e^{16,418}$ (Cormen, Leiserson, Rivest, and Stein 2001).
- *Simple Individual randomization* follows the same process as *Simple Common board randomization* but repeats it for each of the 2,062 director-firm spells independently. Thus, there are $e^{16,426 \times 2,062}$ possible random observations.
- *Lossless Common board randomization* is restricted to randomly relocate within groups of firms that are active in the exact same years. Thus, the total number of possible outcomes is

$$\prod_{g=1}^G N_g!$$

where g refers to a group of firms that are active in the exact same firm-years, G is the total number of these groups in the sample, and N_g is the number of firms in group g . In our sample, we identify 358 groups consisting of 6.75 firms on average. Given our sample's distribution of firms across these groups, the number of possible observations approximates $e^{5,663}$.

- *Lossless Individual randomization*, again repeats this reassignment for each director. The total number of possible outcomes therefore is $e^{5,663 \times 2,062}$.

Tables

Table 2-1

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.

<i>Panel A: Number of tracked board positions per director</i>					
	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

<i>Panel B: Number of simultaneously held board positions per director-year</i>					
	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

<i>Panel C: Distribution over time of simultaneously held positions per directors</i>								
	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

<i>Panel D: Number of years a director appears in the sample</i>					
	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

<i>Panel E: Duration of board mandates</i>					
	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 2-2

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-life. It does so by reporting the minimum and maximum number of simultaneously tracked directors in a firm over its sample years.

<i>Panel A: Number of directors tracked per firm-year</i>			
	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

<i>Panel B: Distribution over time of the number directors tracked per firm</i>													
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 2-3

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 0.01, 0.05, and 0.1-level, respectively.

	N	Mean	Std. dev.	1 st quartile	Median	3 rd 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***

	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 2-4

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 0.01, 0.05, and 0.1-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 2-5

Director fixed effects – The regressions estimated follow the specifications introduced in Table 2-4. 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 where reporting individual coefficients is impractical. Instead, Panel A groups coefficients and reports Wald statistics from testing their joint significance. Under standard assumptions, these statistics follow an F-distribution. ***, **, * 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 (returns are in percent). For 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.

<i>Panel A: Fixed effects regressions</i>				
	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
Financial control variables	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
<i>Panel B: Director fixed effects coefficients</i>				
	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				
1 st quartile	-3.12	-10.52	-12.82	-0.32
Median	0.02	0.14	0.38	0.00
3 rd 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 2-6

Serial correlation of residuals – The sample is the residuals from the regressions in Table 2-5. 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 associated with the overall significance of the model, and the number of observations. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Return on assets			Return on equity			Sales to assets			Tobin's Q		
1 st 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***
2 nd lag		-0.23***	-0.16***		-0.28***	-0.29***		-0.17***	-0.11***		-0.13***	-0.08***
3 rd 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 2-7

Director fixed effects, controlling for the lag of the dependent variable – The performance regressions are estimated and presented analogously to Table 2-5. 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 Wald statistics 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 (in italic print).

<i>Panel A: Fixed effects regressions</i>				
	Return on assets	Return on equity	Sales to assets	Tobin's Q
Individual director effects	900,000***	6,792***	43,550***	49,682***
<i>increase from Table 5</i>	<i>830,211</i>	<i>2,483</i>	<i>-4,556,449</i>	<i>-720,318</i>
Board characteristics	0.98	0.76	2.75***	1.24
Financial control variables	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
<i>increase from Table 5</i>	<i>0.05</i>	<i>0.00</i>	<i>0.01</i>	<i>0.02</i>
N (firm-years)	15,139	15,139	15,139	14,905
<i>Panel B: Director fixed effects coefficients</i>				
	Return on assets	Return on equity	Sales to assets	Tobin's Q
N (directors)	1,854	1,854	1,854	1,846
<i>increase from Table 5</i>	<i>0.00</i>	<i>0.00</i>	<i>0.00</i>	<i>-8.00</i>
Median (y)	10.18	9.73	103.53	1.56
Dist. of coef. (IQR)	5.05	24.30	22.69	0.57
<i>increase from Table 5</i>	<i>-1.14</i>	<i>0.91</i>	<i>-4.18</i>	<i>-0.10</i>
Ind. sign. coef. ($\alpha=0.01$)	25.62%	16.02%	27.72%	25.19%
<i>increase from Table 5</i>	<i>0.27%</i>	<i>0.00%</i>	<i>-0.49%</i>	<i>-2.05%</i>

Table 2-8

Serial correlation of residuals, controlling for the lag of the dependent variable – The sample is the residuals from the regressions in Table 2-7. 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 0.01, 0.05, and 0.1-level, respectively.

	Return on assets			Return on equity			Sales to assets			Tobin's Q		
1 st 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
2 nd lag		-0.18***	-0.15***		-0.29***	-0.28***		-0.11***	-0.09***		-0.09***	-0.06***
3 rd 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 2-9

Real versus placebo director fixed effects – The table repeats the analyses from Table 2-5 on two different sets of placebo data. They result from the *Individual randomization* and *Common board randomization* procedures described in Section 3 of this article. Both procedures are applied in their *simple* form. The table presents the results by showing Wald statistics 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 (returns are in percent), and the fraction of individually significant coefficients at the 1 percent confidence level. The changes resulting from using placebo instead of treatment data are highlighted by adding a comparison of the results with the ones from Table 5 (in italic print).

<i>Panel A: Fixed effects regressions</i>								
	Common board randomization				Individual randomization			
	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***
<i>increase from Table 5</i>	<i>-36,612</i>	<i>23,071</i>	<i>-4,490,000</i>	<i>-705,596</i>	<i>-49,886</i>	<i>655,691</i>	<i>-4,593,761</i>	<i>-764,878</i>
Board characteristics	1.94**	0.97	3.00***	1.13	1.68*	0.90	2.89***	1.01
Financial control variables	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
<i>increase from Table 5</i>	<i>-0.02</i>	<i>-0.02</i>	<i>-0.02</i>	<i>-0.02</i>	<i>-0.03</i>	<i>-0.01</i>	<i>-0.02</i>	<i>-0.02</i>
N (firm-years)	15,139	15,139	15,139	15,139	15,139	15,139	15,139	15,139

<i>Panel B: Director fixed effects coefficients</i>								
	Common board randomization				Individual randomization			
	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
<i>increase from Table 5</i>	<i>-1,203</i>	<i>-1,203</i>	<i>-1,203</i>	<i>-1,203</i>	<i>-1,207</i>	<i>-1,207</i>	<i>-1,207</i>	<i>-1,207</i>
Dist. of coef. (IQR)	7.51	17.33	28.67	0.82	6.19	15.10	21.48	0.62
<i>increase from Table 5</i>	<i>1.32</i>	<i>-6.07</i>	<i>1.80</i>	<i>0.15</i>	<i>0.00</i>	<i>-8.29</i>	<i>-5.39</i>	<i>-0.05</i>
Ind. sign. coef. ($\alpha=0.01$)	36.25%	28.88%	36.71%	36.56%	35.86%	26.89%	32.61%	32.30%
<i>increase from Table 5</i>	<i>10.90%</i>	<i>12.86%</i>	<i>8.50%</i>	<i>9.32%</i>	<i>10.51%</i>	<i>10.87%</i>	<i>4.40%</i>	<i>5.06%</i>

Table 2-10

Benchmarking of the treatment effects against repeated draws of the placebo effects – The table evaluates whether the results from treatment data differ from placebo results in repeated draws. *Treatment (full)* refers to the estimates from Table 2-5. *Treatment (Subset)* estimates the effects on treatment data including only a random subset of the director indicators (33 percent). *Common board* and *Individual* refer to analyses on placebo data generated by the respective procedure introduced in Section 3 (both in their *Simple* form). 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 Wald statistics 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 Wald statistics between the respective placebo results and the treatment 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. 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 0.01, 0.05, and 0.1-level, respectively.

	Number of draws	Number of directors	Wald statistics				Adjusted R ² in %	Director fixed effects coefficients				
			1 st quartile	Median	3 rd quartile	KS test		Wilcoxon	IQR	Median	KS test	Wilcoxon
Return on assets												
Treatment (full)	1	1,854		69,789				65.56	6.19	0.02		
Treatment (Subset)	100	619	9,154	21,251	73,005			62.67	5.98	0.04		
Common board	100	608	18,703	62,336	261,731	0.25***	-3.56***	62.62	7.22	0.02	0.05***	-1.36
Individual	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												
Treatment (full)	1	1,854		4,309				25.82	23.39	0.14		
Treatment (Subset)	100	619	7,732	29,815	133,966			24.80	17.35	0.10		
Common board	100	608	27,279	83,618	298,112	0.32***	-3.84***	24.49	18.57	0.14	0.02***	0.98
Individual	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												
Treatment (full)	1	1,854		4,600,000				90.77	26.87	0.38		
Treatment (Subset)	100	619	5,973	22,330	210,429			89.21	25.51	0.77		
Common board	100	608	25,245	101,534	395,071	0.31***	-3.89***	88.83	27.82	0.35	0.03***	3.24***
Individual	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												
Treatment (full)	1	1,854		770,000				71.25	0.66	0.00		
Treatment (Subset)	100	619	5,312	15,788	70,214			68.80	0.64	0.01		
Common board	100	608	14,939	43,187	192,038	0.29***	-3.95***	68.56	0.73	0.00	0.04***	2.27**
Individual	100	620	3,866	13,933	91,579	0.10	0.48	68.85	0.67	-0.01	0.02***	3.93***

Table 2-11

Benchmarking of treatment director effects against the outcome *Lossless* placebo datasets – The analysis of Table 2-10 is repeated on placebo datasets resulting from procedures with the *Lossless* restriction. *Common board* and *Individual* refer to placebo data analyses based on 100 draws of these alternative procedures. *Normal* refers to the estimates from Table 2-5. 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 Wald statistic testing for the joint significance of director effects for the treatment analysis. For the placebo outcomes, it provides the fraction of draws with a higher Wald statistic than the treatment estimate (*empirical p-value*). 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 treatment and placebo estimates. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

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


Figures

Figure 2-1


Illustration of the randomization procedures. This figure illustrates director data as individual data points $z_{kit} \in \{0,1\}$ in the three dimensional space spanned by directors (k), firms (i), and time (t). Each of these data points is an indicator that equals 1 if the director-firm-year combination referred to by its coordinates represents active service (i.e., director k served on the board of firm i in year t) and 0 otherwise. Panel A illustrates this space with one data point z_{kit} (red box).

Randomization of director data consists of relocating each data point to random coordinates, while leaving its value (i.e., 0 or 1) unaffected. This relocation process is restricted in the randomization procedures applied throughout this article. First, data points are only relocated along the firm dimension (i). Second, the relocation jointly applies to larger arrays of data points. The arrays relocated in *Individual randomization* are obtained by stacking data points over time (Panel B). *Common board randomization* additionally stacks arrays across directors (Panel C).

Panel A: Individual value (z_{kit})



Panel B: Array of data points as jointly relocated in Individual randomization



Panel C: Array of data points as jointly relocated in Common board randomization




Figure 2-2

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 2-4; the dashed line results from Yermack's (1996) estimate based on a logarithmic relation ($Tobin's\ Q = -0.337 \times \ln(Board\ size)$); and the dotted line is from Faleye (2007) and results from a linear relation ($Tobin's\ Q = -0.015 \times Board\ size$). The bars show the relative frequency of board sizes in our data. The axis corresponding to these bars is suppressed. Instead, the white and gray bars roughly indicate the central 90% of observations and 5% upper and lower extremes, respectively.


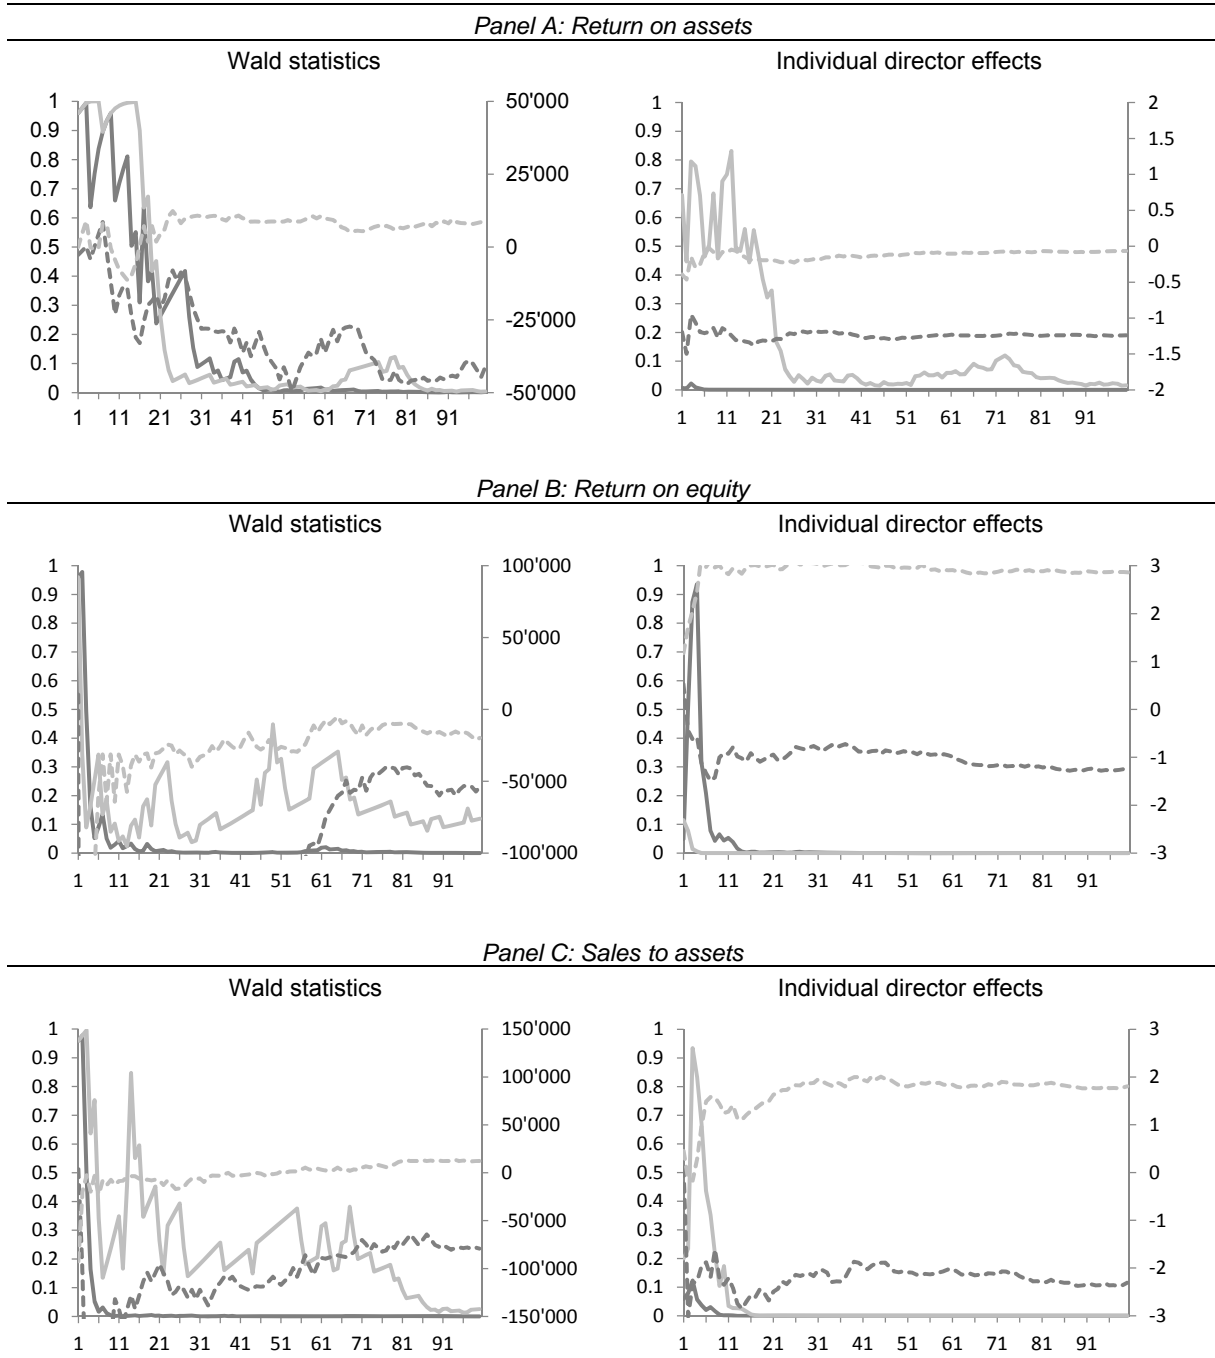


Figure 2-3

Differences between treatment and placebo results with growing number of random draws – This figure shows how the *Subset* treatment and placebo outcomes relate as the number of draws increases from 1 to 100. The left column compares distributions of Wald statistics and the right column compares distributions of individual director effects. 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 in the median Wald statistic and the IQR of estimated coefficients, respectively. Positive differences indicate that treatment effects are stronger than placebo effects. The dark lines compare *Subset* and *Common board* placebo results, whereas the light lines compare *Subset* and *Individual* placebo results.



Panel D: Tobin's Q

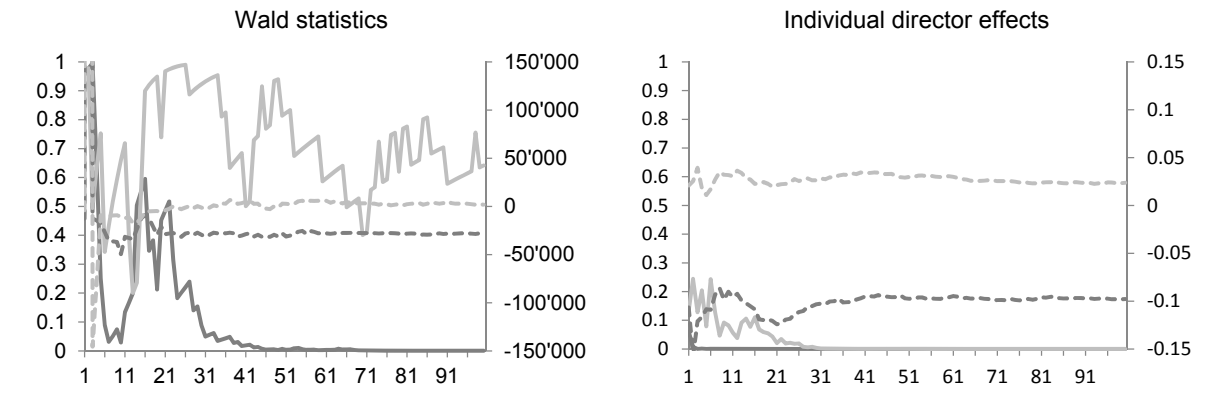




Figure 2-4

Characteristics of treatment and placebo data – The figure compares the treatment data to the placebo data used for the analysis in Table 2-10. The placebo data is generated using the *Common boards* (squares) and the *Individual* procedures (triangles) in their *Simple* form. The bars in the first column represent the full sample of treatment data. In the second column, the bars represent *Subset* treatment data, that is, when including randomly selected 33 percent of tracked directors. Three characteristics of the data are considered: the number of tracked directors per firm-year (Panel A), the number of board positions a director is tracked in (Panel B), and the duration of mandates, that is, director-firm combinations (Panel C). The diagrams show relative frequencies.

Panel A: Number of tracked directors per firm-year



Panel B: Number of board positions per director



Panel C: Duration of mandates




Figure 2-5

Groups of firms for *Lossless* randomization. The *Lossless randomization procedure* restricts reassignment of board service to take place within specific groups of firms. These groups represent unique combinations of years the sample firms are active in. That is, all firms inside a group are active in the exact same years. In total, our sample firms form 358 unique combinations of years active in. Thus, reassignment takes place within 358 different groups.

This figure illustrates the fraction of sample firms that are reassigned within groups of a given size. Group size is measured as the number of firms inside a group. For example, the most common combination of years active and, thus, the largest group contains 230 firms (about 10 percent of the sample). Conversely, 202 firms (about 8 percent of the sample) are inside groups of size one. That is, there are 202 unique combinations of years active, which apply to only 1 firm each. Inside these atomic groups, board service cannot be reassigned to any firm but the original one. Thus, the placebo necessarily overlaps with the treatment.






Figure 2-6

Characteristics of treatment and placebo data (*Lossless* reassignment) – This figure compares the full sample of treatment data to the placebo data used for the analysis in Table 2-11. The placebo data is generated by *Common board* (squares) and *Individual randomization* procedures (triangles) including the *Lossless* modification. The treatment data is represented by the grey bars. Three characteristics of the data are considered: the number of tracked directors per firm-year (Panel A), the number of board positions a director is tracked in (Panel B), and the duration of mandates, that is, director-firm combinations (Panel C). The diagrams show relative frequencies.

Panel A: Number of tracked directors per firm-year



Panel B: Number of board positions per director



Panel C: Duration of mandates


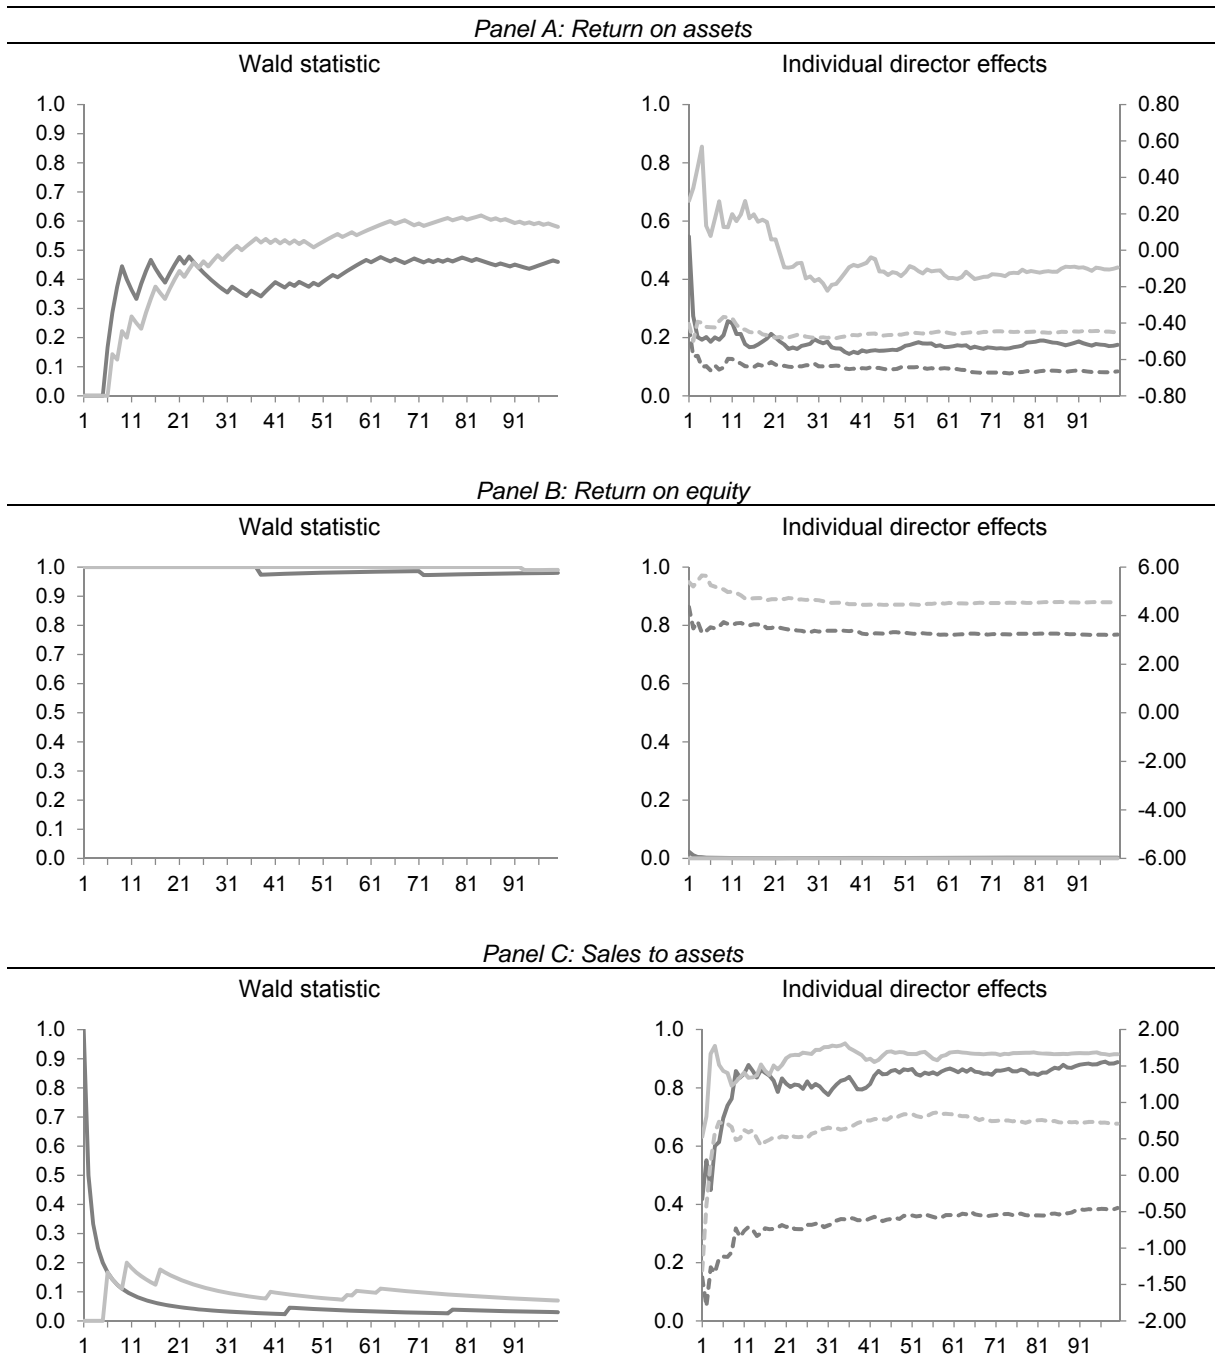



Figure 2-7

Differences between treatment and placebo outcomes with a growing number of random draws (*Lossless*) – Table 2-11 presents a snapshot of how the outcome of treatment and placebo data analyses compare after 100 draws. This figure shows how the treatment and placebo outcomes evolve relative to each other as the number of draws grows from 1 to 100. The left column compares distributions of Wald statistics and presents the respective *empirical p-values*, that is, the number of random draws which lead to a higher Wald statistics than the treatment data. The right column compares distributions of individual director effects. 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 the IQR treatment and placebo effects. Positive differences indicate that treatment data results are more widely distributed than placebo results. The dark lines compare full sample treatment and *Common board* placebo results, whereas the light lines compare full sample treatment and *Individual* placebo results.



Panel D: Tobin's Q



Paper 3:

Robustness, statistical power, and important directors

Abstract

This article provides a better understanding of Senn's (2014) findings: The outcome that individual directors are unrelated to firm performance proves robust against different estimation models and testing strategies. By looking at CEOs, the statistical power of the placebo benchmarking test is evaluated. We find that only the stronger tests are able to detect CEO fixed effects. However, these tests are not suitable to analyze directors. The suitable tests would detect director effects if the inter quartile range of the true effects amounted to 3 percentage points ROA. As Senn (2014) finds no such effects for outside directors in general, we focus on groups of particularly important directors (e.g., COBs, non-busy directors, successful directors). Overall, our evidence suggests that the members of these groups are not individually associated with firm performance either. Thus, we confirm that individual directors are largely fungible. If the individual has an effect on performance, it is of small magnitude.

1. Introduction

Senn (2014) uses fixed effects analysis to investigate the importance of individual independent directors. Initially, he finds effects that appear to be highly significant and have a magnitude that is comparable to what Bertrand and Schoar (2003) find for executives and Cronqvist and Fahlenbrach (2009) for blockholders. These findings turn out to be misleading, as the effects' test statistics are indistinguishable from placebo benchmarks. In sum, the findings suggest that independent directors do not have a measureable individual impact on firm performance and are therefore fungible.

This result is surprising because it conflicts with the image of boards suggested by the literature. Three different strands of corporate governance literature suggest that a single director can affect a firm. The first strand of literature suggests that directors bear personal responsibility for the firm outcome. Gilson (1990) documents that directors typically lose their position if their firm faces financial distress. Moreover, these directors also lose positions on other boards over the following years. This latter finding is in line with Fama and Jensen's (1983) argument that directors build a reputation based on the performance of the firms they previously served for. Overall, an outside director's personal wealth is estimated to change by 11 cents per USD 1,000 change in shareholder wealth (Yermack 2004). Implicitly, this suggests that individual directors can affect the outcome of firm activities.

Second, event study analyses suggest that individual directors matter. They show that the stock market considers some directors more valuable to their firms than others. In particular, Fich (2005) documents that firms can expect significantly higher abnormal stock returns when announcing that an outside CEO will join their board than when appointing a non-CEO director. This difference is seen to originate from the CEO directors' "superior managerial talent and unique expertise". Other articles investigate announcement returns of a wider range of different directors but acknowledge that "One could argue that independent directors are valuable for shareholders [...] for their abilities, expertise, and skills" (Nguyen and Nielsen (2010), p. 559).

The third strand of literature documents the effect of different classes of directors on firm outcomes. For example, Masulis, Wang, and Xie (2012) find that foreigners on the board of U.S. public companies have a positive impact on acquisition returns but a negative impact on operating firm performance. Similarly, Giannetti, Liao, and Yu (forthcoming) find directors with foreign experience to have a positive effect on market valuation, productivity, and profitability of Chinese firms. Field, Lowry, and Mkrтчyan (2013) show that busy directors have a positive impact on book-to-market ratios of IPO firms, which the authors attribute to these directors' superior experience and the associated valuable advice. Earlier studies found the impact of busy directors to be negative, which allegedly results from time constraints and inefficient monitoring (Fich and Shivdasani 2006). While some classes discussed are in line with individual director effects (e.g., financial experts (Güner, Malmendier, and Tate 2008) or foreign directors in U.S. firms (Masulis, Wang, and Xie 2012)), others are not (e.g., a director may be an industry expert in one firm but not in others (von Meyerinck, Oesch, and Schmid 2012)). Nevertheless, this literature indicates that directors can single-handedly affect firm outcomes.

Outside the academic literature, lawmakers have adopted the notion that single directors can make a difference. For example, the Sarbanes-Oxley Act of 2002 regulates board composition and asks for at least one financial expert on the board's audit committee. Meanwhile, only few academic articles find that directors do not affect firm performance (e.g., Fahlenbrach, Low, and Stulz (2010) for the case of outside CEOs).

These previous findings suggest that, given the situation a given firm is in, the addition of a particular director has an impact. The analysis of Senn (2014) helps interpreting these findings. In particular, there seem to be many identical individuals who can provide the special skills a firm requires from its directors. This interpretation contrasts the notion that there are individuals with special skills that would contribute in principle to any firm. In sum, there appear to be no special individuals but there are special situations.

In the present article, we propose a number of additional tests to better understand this outcome. First, we ask whether the outcome depends on the regression model or is robust to different specifications. In particular, the specification used so far relies on firm-fixed effects to control for firm and industry related differences in performance. However, it is unclear whether firm and industry effects remain constant over a sample period of 16 years. We therefore repeat the analysis using an alternative specification to evaluate whether, for example, time varying industry effects confound our analysis. In this specification, the performance measures are centered around the annual industry averages instead of controlling for firm and year fixed effects. Thus, we no longer rely exclusively on within-firm variation to estimate director effects but make use of annual within-industry variation. The outcome of this robustness test is comparable to the outcome of the original analysis in all material aspects. That is, directors still appear to be fungible.

Second, we note that the individual director coefficients estimated in the fixed effects regressions are unbiased, although inference is questionable. Based on this notion we evaluate whether the estimated effects are real by testing additional hypotheses about the coefficients. We test two hypotheses: (1) director effects persist out of sample; and (2) director effects correlate across performance measures. Overall, these tests agree with the previous finding that the director effects estimated on treatment data are indistinguishable from effects estimated on placebo data. Thus, they do not provide any consistent evidence that individual directors are associated with firm performance.

The finding that the presence of a specific individual on the board is not associated with firm performance can be interpreted in two different ways: Either individual directors do not matter for firm performance (i.e., they are fungible) or the effect of individual directors is too small for our tests to detect it. We seek to learn more about the two explanations by asking what magnitude of director effects on firm performance our tests would detect. To answer this question we replicate Bertrand and Schoar's (2003) analysis of CEOs. Arguably, CEOs are more closely tied to firm operations than independent directors. Thus, if any individuals have an impact on firm performance, we expect the CEO effects to be stronger than the director effects. If we find an effect, its magnitude would be an upper bound for the strength of individual director effects. Using the ExecuComp sample we estimate CEO and top-management fixed effects. Our initial results suggest individual executives have a

significant impact on firm performance. This finding is similar to what the original authors and Fee, Hadlock, and Pierce (2013) document. Fee, Hadlock, and Pierce (2013) proceed to showing that the effects are indistinguishable from a placebo benchmark. We, however, find that the Wald tests obtained from treatment data are significant under the distribution of placebo outcomes if we use Senn's (2014) *Simple* randomization procedure (i.e., if we assign our tracked directors randomly to other firms without ensuring that the two firms are active during the same years). That is, we document a significant relation between the presence of certain CEOs and firm performance. This result differs from previous ones. The reason may be that we use a different sample, a different estimation model, or a different process to generate placebo data.

In sum, the results regarding CEOs suggest that the fixed effects method with placebo benchmarking is capable to detect the association of an individual's presence with firm performance. However, the statistical power of the method depends on how the placebo is obtained. If we use a placebo based on *Simple* randomization, the method is powerful enough to detect the effect of CEOs. The more restricted randomization (referred to as *Lossless*) is too weak to do so.

When investigating directors, *Simple* randomization fails to preserve important patterns of director data. These patterns are distorted when director data is reassigned to a firm that does not exist yet or has ceased existing in that particular year. Thus, we have to rely on the less powerful randomization method (*Lossless*). We use simulated data to conduct an analysis of the statistical power this method has. This analysis shows that the method works reliably if the effects are strong enough. For example, the method detects director effects with 0.71 probability at the 0.01 confidence level if the interquartile range of this effect on *Tobin's Q* is 0.36 (in line with previous findings, our simulations assume a mean director effect of zero). However, this magnitude seems to exceed the magnitude of the effect we would expect for typical outside directors. For a rough benchmark, we can look at Fahlenbrach, Low, and Stulz (2010), who compare announcement returns of firms appointing outside CEOs to the board with appointments of other outside directors. These announcement returns differ by 1 percent. The IQR of 0.36 that our tests would require corresponds to 18 percent of mean *Tobin's Q*. Thus, the method might overlook the effects we are after.

To counter this limited statistical power, we focus our analysis on cases where we expect a particularly strong association between the presence of directors and firm performance. Based on previous literature, we believe that these situations are characterized either by the economic environment of the firm, the organization of the board, or personal characteristics of the director. Examples include COBs, individuals with top executive experience, directors whose appointment triggers a significant market reaction, and foreign, highly paid, non-busy, and reputable directors. We find that hardly any of the resulting subgroups of directors and positions yield significant results. That is, even those independent directors that the literature considers particularly important are not measurably associated with firm performance. In a multivariate analysis, the director fixed effects do not appear to be significantly related to personal characteristics either.

Overall, we confirm Senn's (2014) results. Using alternative tests and specifications does not materially affect the outcome. Thus, we conclude that the methods used are reliable. Regarding the

statistical power of the tests we confirm the expectation that it depends on how the placebo benchmark is obtained. Randomization processes that focus on a large number of possible outcomes (e.g., *Simple* randomization) have a greater statistical power than the more restricted ones (e.g., *Lossless* randomization). These relatively powerful methods are strong enough to detect CEO fixed effects. Unfortunately, they lead to absurd outcomes when applied to director data, namely that placebo director data have greater explanatory power than real treatment data. The less powerful methods are unable to detect CEO fixed effects or effects of particularly important directors. A power analysis provides an upper bound for effect sizes that would likely lead to this null result. This upper bound states that we would almost certainly find director effects on *Return on assets* and *Tobin's Q* if the standard deviation of these effects were greater than 20 to 25 percent of the respective performance measure's standard deviation in the pooled sample.

The remainder of this article is structured as follows: Section 2 reviews the data and methods used by Senn (2014) to obtain the individual director effects. Section 3 conducts a robustness analysis and additional tests to corroborate the finding that the director fixed effects are insignificant. In Section 4 we estimate CEO fixed effects to evaluate the statistical power of the placebo benchmarking procedure. It turns out that the methods available to study individual directors are relatively weak. Exactly how powerful they are is evaluated using simulations. As the outcome suggests that the effects of typical outside directors would not be detected, Section 5 focusses on directors that are believed to be particularly important. Section 6 concludes.

2. Data and methodology

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. 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 years¹, 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. Second, we include only directors that serve on multiple boards within our dataset. This restriction will help us distinguish director effects from

¹ This criterion follows from restricting our analysis to directors who are present in their firms for at least 3 years (see below).

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.

In our basic analysis, we follow Senn (2014) in using these data to estimate director fixed effects in a model of the form:

$$y_{it} = \delta_{it} + \lambda_{it} + \mathbf{X}_{it}\boldsymbol{\beta} + \mathbf{Z}_{it}\boldsymbol{\Gamma} + \varepsilon_{it},$$

where y_{it} is the outcome of firm i at time t . The outcomes we consider are *Return on assets*, *Return on equity*, *Sales to assets* and *Tobin's Q*. δ_{it} and λ_{it} are fixed effects for firms and years, respectively. Bold print indicates that the respective symbols refer to vectors rather than scalars. \mathbf{X}_{it} is a vector of firm-level, time-varying control variables. \mathbf{Z}_{it} is a vector of K treatment indicators, with K being the total number of tracked directors. That is, the model includes a separate indicator variable for each director. This variable equals one if the respective director k serves on the board of firm i at time t , and zero otherwise. ε_{it} is the error term.

The outcome of these estimates is compared to placebo results to evaluate significance. Placebo data are generated by randomly reassigning directors to sample firms. Senn (2014) describes various possible ways to perform this randomization. In particular, *Individual* and *Common board randomization* are distinguished, both of which are implemented in a *Simple* and a *Lossless* form. As the differences between *Individual* and *Common board randomization* turn out to be relatively minor, the present article focuses on *Common board randomization*, unless it is explicitly stated. However, the distinction between *Simple* and *Lossless* forms of randomization is maintained throughout the article as it may significantly affect the outcome.

Recall that *Simple* randomization reassigns board service to firms without further restricting the set of possible outcomes. This procedure ignores the fact that we use unbalanced panel data. In consequence, boards may be reassigned to inexistent firm-years. For example, The Coca-Cola Company's board may be assigned to Google Inc. In this case, as the latter firm enters our sample only in 2004, information on the former firm's board would be lost for much of the sample period. Thus, while this randomization procedure offers a large number of possible outcomes, it leads to distortions in potentially important data characteristics (e.g., the number of board positions per director, the duration of tracked positions, or the number of tracked directors per firm-year).

Lossless randomization addresses the issues that arise from using an unbalanced panel. That is, it reassigns firm and board data within subsets of firms that are in the sample in the exact same years. Restricting reassignment in this manner, we succeed in maintaining the original data characteristics. However, this advantage comes at the cost of having fewer possible outcomes. In particular, some of the subsets contain only a handful of firms. Thus, there is a significant risk that these firms' treatment data overlaps with placebo data. This risk is material for roughly 20 to 25 percent of the sample firms.

Using these randomized datasets as a placebo benchmark for the outcome of director fixed effects regressions, Senn (2014) finds that directors have no detectable impact on performance and are therefore fungible. The next section provides additional tests of this conclusion. First, we test the

robustness of the finding to different specifications of the estimation model. Then, we investigate hypotheses regarding the properties of the estimated effects.

3. Additional tests of the director fixed effects

3.1 Robustness to model specification

Using the above described firm-fixed effects specification helps separating the effect of individual directors from firm specific performance differences. Thus, among other things, the firm fixed effects are supposed to control for differences between industries. They do so under the assumption that firms remain in the same industry and that industry differences are constant over time. As our sample spans a period of 14 years, it is questionable whether these assumptions hold. After all, firms may venture into new industries and industry specific characteristics may vary over time. In this case, our model for normal performance needs to be adjusted. To evaluate whether our estimates are driven by a model misspecification, we repeat the analysis using industry-year standardized performance as the dependent variable and no longer control for firm and year fixed effects. The analysis shown in this article uses the Fama and French 48 industry classification. Comparable results obtain using the 17 industry classification (not shown).

Table 3-1 describes the outcome of fixed effect regressions using this alternative model. These outcomes are compared to placebo benchmarks obtained by estimating the same model on randomized datasets. The placebo datasets used in Table 3-1 result from *Lossless* randomization. The first columns of the table show that we have a single treatment dataset and 100 draws of placebo data (*N of draws*). Further, the number of tracked directors is identical in both types of data (*N of directors*). Next, the table analyzes the Wald statistics associated with the overall significance of director indicators: for treatment estimates, the actual test statistic is shown (upper row); for estimates on simulated data, the table provides the fraction of random draws that lead to Wald statistics higher than the treatment estimate. This fraction can be interpreted analogously to a p-value. We note that this p-value is 0.03 for *Return on equity* (ROE) and 0.05 for *Tobin's Q*, implying that the director indicators contribute significantly to the explanatory power of our model. For the other two performance measures, the empirical p-value suggests that the treatment estimates do not differ significantly from the placebo estimates.

Cronqvist and Fahlenbrach (2009) evaluate the economic significance of blockholder fixed effects by comparing the distribution of the individual treatment coefficients estimated to that of placebo effects. They find that the treatment effects are significantly more widely dispersed than placebo effects and conclude that the effects are sizeable. Table 3-1 shows that the interquartile range (*IQR*) of the treatment effects on *ROE* is 0.63 standard deviations compared to 0.61 for the placebo effects. The difference in distributions is significant based on Kolmogorov-Smirnov and Wilcoxon rank-sum tests. According to Cronqvist and Fahlenbrach (2009) this result would suggest that the effects on *ROE* are real. However, given that the difference between the *IQR* of treatment and placebo effects amounts to only 0.02 standard deviations, we conclude that their magnitude is trivial. Using *Tobin's Q* as the

dependent variable, we find that the IQR of the treatment effects is even significantly smaller compared to placebo outcomes. This result is meaningless as random data should not lead to significantly higher test statistics than treatment data. Thus, we are reluctant to interpret the findings of Table 3-1 as evidence that individual directors affect performance.

Table 3-2 repeats the analysis using the *Simple* randomization procedure to generate placebo data. This procedure has the benefit that the number of possible outcomes is greater. This benefit is achieved by possibly allowing distortions in the data characteristics to occur. The most obvious distortion is an overall smaller number of tracked directors. In Table 3-2 we avoid possible effects of this difference by reducing the number of directors in the treatment analysis to correspond to that in the placebo analysis. That is, we randomly select a subset of directors. To allow for a general interpretation of the results, this random selection is repeated 100 times. Thus, Table 3-2 compares two distributions of Wald statistics resulting from 100 draws each, instead of comparing a single treatment test statistic to a distribution of 100 placebo outcomes as in Table 3-1.

The table shows for all performance variables that each quartile of the treatment Wald statistics is smaller than the respective quartile of the placebo statistics. Moreover, the difference in distribution is significant according to Kolmogorov Smirnov (K/S test) and Wilcoxon tests. The table also compares the distribution of the estimated fixed effects between treatment and placebo data. Contrary to what Cronqvist and Fahlenbrach (2009) predict, we find that the effects from placebo data are more widely dispersed than the ones based on treatment data. Moreover, this finding is significant and holds for all performance variables. This outcome implies that the economic magnitude of the placebo effects is significantly larger than that of treatment effects. As no economic effect whatsoever underlies placebo data (they are random indicators), this finding is economically meaningless.

Overall, the results from using the model specification with annual industry standardization instead of firm fixed effects largely agree with the findings of our original analysis. While the effects on *ROE* might be significant when compared to the *Lossless* benchmark, their economic magnitude is too small to be of interest. All other performance variables and the use of an alternative placebo benchmark suggest that individual directors have no measurable impact on firm performance.

Large parts of the analysis so far rely on Wald statistics to test the overall significance of the estimated effects. One exception is when we follow Cronqvist and Fahlenbrach (2009) and test for the overall significance of fixed effects based on the estimated coefficients. They note that treatment effects should be more widely distributed than placebo effects if they indeed represent a real economic association. In the next section, we conduct further tests that are based on the estimated director effects. In particular we formulate and test additional hypotheses about real treatment effects.

3.2 Hypotheses on the estimated effects

When comparing Wald test statistics to the F-distribution, we overestimate the joint significance of fixed effects in the presence of serial correlation and because asymptotic arguments may not apply. Despite these issues, coefficient estimates are unbiased. Thus, we can test whether the estimated

effects are noise by testing hypotheses about individual director effects. These hypotheses go beyond the original question of whether the individual directors are associated with firm performance.

Out of sample predictive power

The first additional hypothesis we test is that director effects have out of sample predictive power. If directors have an effect on their firms because of certain constant characteristics, this effect should be valid out of sample. That is, a director should have a similar effect over time and across different firms (Fee, Hadlock, and Pierce 2013).

We assess the out of sample predictive power over time by estimating two separate effects for each director: one effect for the early and a separate effect for late half of the sample period (i.e., before and after 2002). Table 3-3 tests whether our estimates for these two effects are associated. We measure association using Spearman rank correlations as well as comparing the coefficients' signs. The directors' early and late effects appear to be systematically related. The Spearman correlations (0.26 – 0.34) are sizable and statistically significant (p-value). The sign test shows that the early and late effects of 60 to 63 percent of tracked directors have the same sign. This is significantly more than the 50 percent we would expect just randomly. That is, a director with a positive effect in the early sample-half is likely to have a positive effect in the late sample-half.

Serial correlation could also affect this test for out-of-sample predictive power. Thus, we ask whether placebo outcomes show comparable out-of-sample predictive power. That is, we test whether the early and late sample-half coefficients obtained from placebo data are correlated. These placebo associations serve as a benchmark for our treatment outcomes. The table reports the fractions of placebo outcomes with an association that is stronger than the one estimated on treatment data (empirical p-values). Most of these fractions range from 0.27 (Spearman correlation for *ROE* using *Lossless* randomization) to 0.50 (Sign test for *Tobin's Q* using *Lossless* randomization). As they can be interpreted analogously to p-values, these findings indicate that the association estimated on treatment data is similar to that observed with placebo data. Thus, the association we find between in-sample and out-of-sample estimates does not seem to be related to the presence of specific individuals.

So far, we test whether the director effects we estimate have out of sample predictive power over time. That is, we evaluate whether the effects estimated on an early sample are related to the effects estimated on a later sample. However, if an individual has the skill to improve firm performance unconditionally, the observed effect should also be similar across firms. Table 3-4 assesses whether this is the case for the effects we estimate. Thus, we again split the sample in two subsamples. For both subsamples we estimate separate effects for each director. The analysis differs from Table 3-3 in the way the sample is split. Whereas in Table 3-3 we split the sample horizontally into pre and post 2002 subsamples, we now split it vertically. That is, the various spells of each director with different firms are randomly assigned to the two subsamples. For example, assume a director served on the board of firms A, B, and C. A possible outcome of the random split is that we estimate an overall effect in firms A and B and a separate effect for the service in firm C. This split also affects our initial sample criteria. For the original estimates, we required tracked directors to have served for three or more

years in at least two firms. The analysis in Table 3-4 also tracks only directors who meet these criteria. However, after splitting the sample, the effects we estimate for each subsample do not necessarily meet the criteria when considered separately. That is, the two subsample-effects may each refer to a single firm only. If we were to require the sample criteria for each subsample-effect separately, this requirement could only be met by directors who served on 4 or more boards in total. In our sample, only a handful of directors would qualify.

Analogously to Table 3-3, Panel A of Table 3-4 investigates whether the separate subsample-effects are systematically related. We would expect this to be the case if individual directors have the same impact on the performance of all firms they serve. However, the resulting Spearman correlation coefficients are close to zero and mostly insignificant. The only significant (at the 10 percent level) correlation between the two coefficients is found when investigating *Sales to assets*. The table also provides the fraction of directors whose coefficients have the same sign in both subsamples. This fraction is very close to 50 percent for all dependent variables. Consequently, the relation between the sign of the two subsamples' coefficients is insignificant. Given that the associations are insignificant, there is no need to benchmark this analysis against placebo outcomes.

Panel B of the table repeats the analysis of Panel A using 100 different random subsample assignments. This repetition prevents us from interpreting an odd outcome of randomization. The table shows that the findings of Panel A are rather typical. For *Sales to assets* we still find statistically significant correlation between the two subsamples' estimates. However, the correlation is economically negligible (the mean correlation across 100 draws is 0.05). The only change compared to the single draw in Panel A occurs when investigating *ROE*. While the correlation between fixed effects on *ROE* was insignificant in the single draw, it turns significant in slightly more draws than we would expect just randomly. However, opposite to what out of sample predictive power would suggest, the coefficient is negative. Moreover, its magnitude is economically negligible. The fraction of directors for whom the coefficients of the two subsamples have the same sign is not significantly greater than 50 percent for *Sales to assets* and *Tobin's Q*. For *ROA* we observe slightly more observations that are significant at the 10 percent level than we would expect. However, this finding does not hold for the 5 and 1 percent levels. For *ROE* we find more significant results than expected at the 10 and 5 percent levels. However, the fraction of directors whose effects have the same sign is 49 percent (mean across 100 draws). The difference from 50 percent has the wrong sign and is economically negligible.

Overall, we find little evidence that the individual director coefficients persist over time. Similarly, there is barely any association between the effect a director has on one firm and her effect on another firm.

Correlation across performance measures

The second additional hypothesis we test is that the estimated effect of a director on one performance measure is related to the effect of this director on other performance measures. Bertrand and Schoar (2003) and Cronqvist and Fahlenbrach (2009) find correlations among the effects on different dependent variables and interpret these correlations as evidence for the 'style' of a CEO or blockholder, respectively. Fee, Hadlock, and Pierce (2013) would be willing to accept these correlations as evidence for the effects being real but find no such correlations in their data.

Panel A of Table 3-5 reports the coefficients from regressing the estimated director effects on one another. Separate regressions are estimated for each pair of dependent variables. For example, in the top left cell the effects on *ROA* are regressed on the effects on *ROE*. The three coefficients result from least squares regression with Huber-White standard errors (top), least squares regressions weighted by the inverse of the effects' standard error (middle), and median regressions (bottom). The regression coefficients are positive and highly significant for all pairs of dependent variables and all estimation methods. For example, the table suggests that a director with a 1 percentage point higher effect on *ROE* has on average a 0.07 to 0.12 percentage points higher effect on *ROA*. Following the previous literature, this finding suggests that the estimated effects are real.

In Panel B, we perform an analogous analysis on placebo datasets. For both the *Simple* and the *Lossless* randomization procedures we find highly significant positive correlations between the effects on different performance measures. These significant correlations occur for all combinations of the performance measures and independent of the estimation method used. Thus, the correlation between individual director effects on different performance measures can be observed also with randomized data. This finding suggests that the correlations in Panel A do not, in fact, indicate a real treatment effect related to the presence of directors. Rather, the findings may result from a component of performance that is shared by all four measures considered and that is not accounted for by the control variables. This residual correlation might then affect the director indicators or their placebo, respectively. While it remains unclear what exactly the correlation between the individual effects on different performance measures represents, it does not distinguish treatment from placebo director effects. Thus, our findings disagree with previous studies that suggest interpreting this form of correlation as evidence for real treatment effects.

4. Statistical power of individual fixed effect regressions with placebo benchmarking

The analyses in the previous section imply that there is no discernible association between the presence of a given individual director and firm outcome. This result may obtain because the hypothesized effect does not exist. In this case, individual directors are fungible and neither super stars nor losers occur. Alternatively, individual directors may actually be associated with firm outcome but this effect is too weak for our estimation methods to detect. This alternative explanation is particularly important because we assess the significance of our estimates by benchmarking them against placebo distributions.

The present section investigates which of these two competing explanations is more likely to hold. First, it evaluates whether the fixed effects method with placebo benchmarking is, in general, capable to detect any effects of an individual on the firm. To answer this question, we investigate CEOs. Arguably, the CEO is the most influential individual in many firms. In particular, we expect her to have a stronger impact than an outside board member. Thus, it should be easier for our method to detect CEO effects than director effects. Indeed, we find that some randomization methods allow detecting

CEO fixed effects. Thus, the second part of this section focuses on the randomization method most suited to analyze directors (i.e., *Lossless Common board randomization*). Using simulated data we estimate the minimum magnitude of director effects that our method would detect.

4.1 CEO fixed effects

We argue that a CEO would find it easier to have an effect on her firm than an outside director. Presumably, this is the case because the CEO is better informed than outside board members who largely depend on information provided by the management (Coles, Daniel, and Naveen 2008). Further, she has extensive personal control over the firm's resources (Shleifer and Vishny 1997), whereas directors are limited to acting as members of boards of about 10. Ultimately, the CEO is positioned relatively close to the firm's operations compared to directors who meet only about 7 times a year (Adams and Ferreira 2008).

Methodology and data

We estimate CEO fixed effects using similar models as in the directors fixed effects analysis. That is, the regression model controls for the natural logarithm of total assets, the level of investments, and leverage (see Appendix A for variable definitions). In addition, it includes year and firm fixed effects. As we study CEOs instead of board members, we drop the board characteristics from our model of normal performance. Our model is the one Bertrand and Schoar (2003) use for their *Return on assets* regression, except that we also include the lagged *Investment* and *Leverage* variables as in Senn (2014). In an unreported robustness check, the analysis is repeated using the exact Bertrand and Schoar (2003) specification (i.e., without the additional control variables). The outcome of this analysis does not differ materially from the results reported in the present section.

The data needed for these estimations are largely identical with the data of our previous analyses. While we no longer require RiskMetrics data on directors, we now use ExecuComp data to track CEOs. Thus, our sample period is 1992 to 2012. As in our analysis of directors, we employ the tracking criteria from Bertrand and Schoar (2003) (i.e., we focus on individuals who are CEOs of multiple firms for three or more years). 53 individuals meet these criteria. Moreover, to facilitate comparability with what previous studies on CEO fixed effects do (i.e., Bertrand and Schoar (2003) and Fee, Hadlock, and Pierce (2013)), we exclude firms that never have a tracked CEO from our sample. The resulting sample consists of 1,524 firm-years. By excluding firms that never employ a tracked individual, our analysis of CEOs differs from the one of directors. In an unreported robustness check, we evaluate whether our findings regarding CEOs are sensitive to the decision to include or exclude these firms (not shown). We find that this decision does not qualitatively affect the outcomes reported in the present section.

Table 3-6 describes how the service of the tracked CEOs is distributed across the regression sample. Panel A documents the number of positions each CEO is tracked in. By sample selection, none of the executives is tracked in a single firm only. Meanwhile, only one is tracked in three firms and none in more than that. Typically, these tracked positions are sequential, with only 29 out of 512 CEO years of service being simultaneous. These cases are mostly related to M&A activity. For example, we may observe a merger with Compustat continuing to consider the two firms as separate entities.

Alternatively, we may look at a spin off, during which one individual is the CEO of both firms. The results we present are robust to the exclusion of CEOs that hold multiple positions simultaneously (not shown). Panel C lists the number of years the CEOs are active in. The average CEO appears in nearly 10 sample years; that is, about half of the sample period. Related to this, Panel D describes the distribution of the duration of the mandates. However, the actual service of a CEO may begin before the firm enters the sample or continue after it leaves. Thus, the measure of duration is biased downwards in comparison to the actual service for the firm. The average sample CEO remains for little more than 5 years with his firm while mandates lasting less than three years are excluded by sample selection criteria. Although firms that never have a tracked CEO are excluded, there are about 65 percent of firm-years without tracked individuals (Panel E).

Results

The outcome of the regression estimates is described in Table 3-7. This table does not provide actual regression coefficients because the number of coefficients is too large to be displayed. Instead, the table provides Wald statistics associated with the overall significance of groups of coefficients. These groups are *Individual CEO effects*, time-varying firm-level control variables (X_{it}), *Firm fixed effects*, and *Year fixed effects*. We note that the group of individual CEO coefficients is highly statistically significant for all dependent variables based on the Wald tests.

Whether we can trust these Wald tests is evaluated in Table 3-8. This table compares the treatment data CEO fixed effects to outcomes based on placebo data. The placebo data used in Table 3-8 are obtained using the *Lossless* randomization procedure.² As discussed in Section 2, this procedure differs from the *Simple* procedure in that it more accurately mimics characteristics of the original data. That is, it maintains the distributions of the number and the duration of CEO positions, and the number of tracked CEOs per firm-year. This additional accuracy comes at the cost of a lower number of possible random matches.

The table shows that the Wald statistic of treatment data analysis is insignificant under the distribution of placebo test statistics. For all four performance measures, the treatment Wald statistic lies between the first and third quartile of the placebo statistics. According to these results, the treatment outcome does not differ significantly from placebo outcomes. The table then analyzes the actual CEO coefficients and compares their distribution between treatment and placebo data. For most performance measures, the treatment effects are more widely distributed than the placebo effects. *Sales to assets* is the exception to this. The Kolmogorov Smirnov (K/S) and Wilcoxon tests show that the differences between distributions are insignificant. Thus, when using the placebo that focuses on maintaining data structure, we conclude that CEOs are not significantly related to firm outcomes.

The design of the randomization procedure used to obtain these results (*Lossless*) focuses on preserving potentially important characteristics of the treatment data. In particular, these restrictions preserve the fact that the identity of the current CEO is largely determined by the identity of last year's

² The procedure we use is analogous to *Common board randomization*. *Individual randomization* is not applicable to CEOs. This procedure would allow for multiple simultaneous COEs in the same firm, which is in most cases unrealistic.

CEO. Also, they ensure that firms typically have only a single CEO at a time. Further, they preserve the duration of CEO spells with a firm and the number of different firms a CEO serves for. Preserving these characteristics requires restrictions on the randomization process that, in turn, cause the number of possible random matches to be lower than for unrestricted randomization. Most importantly, this *Lossless* procedure causes parts of the randomized dataset to not actually be random but identical to the original treatment dataset. Hence, when using this kind of placebo data, our benchmarks may be excessively high.

Table 3-9 repeats the analysis using a placebo benchmark obtained from a less restricted randomization process (*Simple*). In particular, the randomization also allows for CEOs to be reassigned to firms that were not active in exactly the same years as the CEO himself. In consequence, we lose observations of CEO service. This loss of information is problematic as it leads to different degrees of freedom in the real and placebo Wald statistics. Avoiding the loss of information is possible but comes at the cost of limiting the number of possible random observations. We evaluate this option later in this section (*Lossless* randomization). For the moment, we ensure that the treatment and placebo statistics are comparable by adjusting the number of CEOs we track in the treatment data. In particular, we randomly select 33 percent of the sample of tracked CEOs. This decrease in the number of tracked CEOs ensures that the treatment and placebo Wald statistics have similar degrees of freedom. This random selection is repeated 100 times. As the table shows, this procedure is fairly successful in generating comparable numbers of tracked CEOs in treatment and placebo data (*N of CEOs*). We then compare the distributions of treatment and placebo Wald statistics. All quartiles of the treatment statistics turn out to be greater than the respective placebo quartiles. This observation is true independent of the performance measure considered. Moreover, this difference in distributions is significant (Kolmogorov Smirnov test) and the treatment Wald statistics typically rank higher (Wilcoxon test). In sum, the presence of individual CEOs is significantly related to firm performance. At this point, we explicitly avoid a causal interpretation of the association between the presence of individual CEOs and firm performance, as we are only interested in the statistical power of our method. For a discussion of causality we refer, for example, to Fee, Hadlock, and Pierce (2013).

The last 4 columns of Table 3-9 evaluate the magnitude of the CEO related performance shifts. First, the dispersion of the estimated effects is analyzed. The IQR shows the difference in firm performance between firms with a first versus a third quartile executive. Analyzing *ROA*, this change is the same in treatment and placebo data. For *Sales to assets*, placebo effects predict an even greater difference than treatment effects. However, the treatment effects on *ROE* and *Tobin's Q* seem to be significantly wider distributed than placebo effects. For *ROE*, the placebo analysis suggests that we find a difference of 25.76 percentage points between the first and third quartile CEO (IQR) when using random data. The IQR estimated for treatment data is 2.86 percentage points greater. This compares to a median *ROE* of 9.73. For *Tobin's Q*, the difference between treatment and placebo IQR is 0.15, compared to a sample median of 1.56. Thus, at least for these two performance variables, we conclude that individual CEO effects are significant and economically material.

Conclusions

Our findings suggest significant CEO effects on firm performance. Thereby, they differ from what Fee, Hadlock, and Pierce (2013) document. There are at least three explanations for why we don't obtain the same results:

- (1) *Different randomization method.* Our findings change depending on the randomization process used (compare Table 3-8 to Table 3-9). Unfortunately, Fee, Hadlock, and Pierce (2013) provide little information about how they generate placebo data. They state that they "assign each CEO-to-CEO mover to a different hiring firm than the one he or she actually joins" (p. 593). Thus, randomization seems to be restricted only in the sense that multiple years of CEO service for one firm must not be distributed across multiple firms. The authors do not mention further restricting this process. These observations suggest that the authors' procedure is similar to our *Simple* randomization. However, *Simple* randomization loses CEO observations, whereas the authors do not mention losing any CEOs (they do not provide descriptive statistics on placebo estimates either). Therefore, it is unclear whether our randomization mechanism is comparable to theirs.
- (2) *Different sample.* As mentioned before, we use the sample of CEOs that are identified in the ExecuComp database. This sample covers the current and former S&P 1,500 firms. Fee, Hadlock, and Pierce (2013) use the larger sample of all Compustat firms. This allows them to track 131 CEOs as compared to our 53. Moreover, their sample period (1990 to 2007) begins and ends earlier than ours (1992 to 2012).
- (3) *Different specification.* Fee, Hadlock, and Pierce (2013) leave room for interpretation regarding the specifications they use. In particular it is unclear how their specifications control for firm fixed effects and year fixed effects. We apply the original specification of Bertrand and Schoar (2003). Unreported analyses show that CEO effects largely disappear if we exclude firm and year fixed effects entirely.

Our significant findings obtain only when benchmarking the treatment outcome against the outcomes of placebo data generated by the *Simple* procedure. The treatment outcome is indistinguishable from the *Lossless* placebo outcomes. This finding is arguably related to differences in statistical power of the two methods. *Simple* randomization is designed with a focus on high statistical power. This power comes at the cost of allowing distortions of data characteristics in the placebo relative to the treatment. For the case of directors, these distortions are severe enough to create illogical outcomes. Thus, we cannot apply the *Simple* randomization method to directors. As the outcomes regarding CEOs are largely in line with our expectations, it appears that CEO data is less sensitive to distortions created by *Simple* randomization. This lower sensitivity appears to be related to a number of differences between the two datasets. Two of the main differences between CEO and director data are, first, that directors serve on widely varying number of boards, whereas few CEOs are observed in more than 2 firms. And, second, that directors can and often do serve on multiple boards simultaneously. Conversely, we typically track CEOs only in a single firm at a time.

Lossless randomization is designed specifically to maintain the characteristics of director data. However, preserving these characteristics requires severe restrictions in the randomization process. Ultimately, these restrictions cause the treatment and placebo data to overlap in a significant fraction

of the sample. That is, the treatment effects contaminate the placebo benchmark, making it excessively conservative. Thus, *Lossless* randomization is expected to be less powerful than *Simple* randomization, which might explain why it does not detect individual CEO effects.

As the *Lossless* randomization method is too weak to detect CEO fixed effects, it is still unclear what size of director effects it would detect. The next section answers this question by analyzing the statistical power of the method using a simulation approach.

4.2 A simulation approach to evaluate the statistical power of our test

The statistical power of a test is the probability that it correctly rejects a false null hypothesis. That is, the greater the statistical power of a test, the less likely it is to falsely accept the null hypothesis. The power of a test depends on a number of factors that vary with the testing situation. Generally speaking, it is a function of the sample size, the size of the treatment effects, and the confidence level pursued.

When working with financial data, the sample is typically given exogenously. Thus, we are interested in describing the likelihood to reject the null hypothesis (H_0) as a function of the effect size and the confidence level. Based on this relation, we can investigate whether our test is likely to detect director fixed effects of a reasonable size. What effect size is considered reasonable can be inferred from the extant literature. For example, Masulis, Wang, and Xie (2012) find that *ROA* of firms with foreign independent directors is about 1 percentage point lower than that of firms without FIDs. Fich and Shivdasani (2006) document that *ROA* is about 0.24 percentage points lower in firms with busy boards. Conversely, given our null result, the power analysis provides an upper bound for the size of the hypothetical individual director effects.

We find the relation between the size of the treatment effects and the rejection rate by applying our test repeatedly to different datasets. Simulations are used to generate suitable datasets.

Simulation method

Each run of our simulation starts by randomizing the vector of K director indicators (\mathbf{Z}_{it}). This randomization uses the *Lossless* procedure. The resulting random director indicators are referred to as $\tilde{\mathbf{Z}}_{it}$. As with placebo data, these indicators are not linked to any treatment and therefore have no explanatory power with respect to actual firm performance (y_{it}).

Next, we generate a $(K \times 1)$ vector of random treatment effects ($\tilde{\mathbf{\Gamma}}$). That is, we assign one individual fixed effect to each director (K). These effects are independent random numbers drawn from a normal distribution and they remain fixed throughout the sample. The mean of this distribution is set to zero. This parameter reflects the finding that the average tracked director has no effect on firm performance and that the estimated effects are roughly symmetrically distributed around zero. Hence, the standard deviation of the distribution (i.e., the dispersion of the effect across directors) alone controls the size of the simulated treatment effect. Thus, our simulated treatment implies that some directors have a positive impact on performance while others have a negative impact. However, this negative impact should not be interpreted as the director actually harming the firm. Rather, given that our estimates control for firm fixed effects, it is a relative statement. That is, it compares the firm performance with

this director on the board to the average performance with his predecessor and successor. Thus, if there is variation in the impact of directors, some directors will have a negative effect relative to the others. We define effect size (e) as a fraction of the standard deviation of firm performance in the pooled sample ($\hat{\sigma}_y$). Thus, the simulated treatment effects are:

$$\tilde{\Gamma} \sim N_K(0, \sigma^2),$$

$$\text{where } \sigma = e\hat{\sigma}_y.$$

These treatment effects are applied to all firm-years, in which the respective director is present according to the randomized director indicators. That is, manipulated firm performance is obtained by adding the simulated effects to actual firm performance:

$$\hat{y}_{it} = y_{it} + \tilde{Z}_{it}\tilde{\Gamma}.$$

The manipulated performance constitutes the data with known treatment effects that we need to gauge the statistical power of our tests. That is, we estimate our director fixed effects model on these simulated data:

$$\hat{y}_{it} = \delta_{it} + \lambda_{it} + X_{it}\beta + \tilde{Z}_{it}\Gamma + \epsilon_{it}.$$

The outcome of this estimate is then benchmarked against the outcomes of 100 placebo estimates. Analogously to the analysis of actual treatment data, the placebo estimates are obtained by randomizing the simulated treatment data. From placebo benchmarking, we obtain the significance of the (simulated) treatment outcome with respect to the placebo outcomes (empirical p-value).

Obtaining the empirical p-value completes one run of the simulation. To learn about the power of our tests, this simulation is run repeatedly and with different sizes of the simulated effect (e). In doing so, we keep track of the resulting empirical p-values. Ultimately, we compute the rejection rates of the null hypothesis for each effect size at standard confidence levels.

Results

We evaluate the statistical power of our tests for all four performance measures. Effects of 7 different sizes, ranging from 0 to 0.3, are simulated and tested for each performance measure. Given our definition of effect size, this means that we gradually increase the standard deviation of the simulated director effects from 0 to 30 percent of the pooled standard deviation of firm performance. For each performance measure and effect size, we simulate 100 datasets and test at standard confidence levels ($\alpha = 0.01, 0.05, \text{ and } 0.1$).

Figure 3-1 shows the outcome of this analysis. Each row refers to a different performance measure, whereas columns refer to different confidence levels. The horizontal axes show e , the size of the simulated director effects as a fraction of $\hat{\sigma}_y$. The vertical axes show the rejection rate of H_0 . That is, the fraction of repeated tests that reject the null hypothesis. The dashed lines mark the respective confidence levels on the vertical axis. The points connected by solid lines represent the rejection rates we observe in our simulation for each effect size at the given confidence level. Thus, for example, the

first figure in the top row refers to *Return on assets* estimates, testing at the 99 percent confidence level. The squares connected by solid lines show rejection rates. For example, given an effect of size 0.15 (vertical axis), the null hypothesis is rejected in 30 percent of the simulated trials.

Before discussing the power of our test, a first important outcome of the simulation analysis is that our test does not overreject the null hypothesis. That is, if there is no treatment effect, our test does not indicate a significant effect. In particular, in the absence of a treatment, we expect to find the fraction of rejections to equal the confidence level (alpha or type I error). Figure 3-1 documents that for *Sales to assets* and *Tobin's Q* the outcomes almost exactly meet our expectation for all confidence levels. That is, the rejection rate (points / solid line) coincides with the confidence level (dashed line) for effect size 0. For *Return on equity*, the rejection rates are initially somewhat lower than expected, whereas they are somewhat higher for *Return on assets*. We believe that these deviations are purely random and would disappear if more manipulations were analyzed. This conclusion is based on three observations: First, the deviations are relatively small. Second, the deviations have opposite signs for *Return on assets* and *Return on equity*. Third, for *Return on assets*, rejection rates decrease as effect size increases from no effect to a small effect (confidence levels $\alpha = 0.05$ and 0.1). This decrease is difficult to reconcile with an overrejecting test.

With regard to the statistical power of test, we note that the rejection rates barely rise above α for small treatment effects. That is, the tests fail to reject the false null hypothesis for treatment effects of this size. As the simulated effects increase, the observed rejection rates increase as well. That is, our tests are more likely to detect larger effects than small ones. In the analysis of *Sales to assets* the rejection rate approaches 1 for $e = 0.15$ at the 0.01 confidence level. That is, our tests nearly always detect effects on *Sales to assets* if the variation between director effects is 15 percent of the standard deviation of performance. For the other performance measures, only larger effects lead to rejection rates of this magnitude. Thus, our tests are most powerful in the analysis of *Sales to assets*. Tests on *Return on assets* and *Tobin's Q* are reliable for effect sizes of 0.2 to 0.25. Tests on *Return on equity* fail to reliably reject the null at the 0.01 confidence level even for the largest simulated effect ($e = 0.3$).

Conclusions

The results of the power simulation displayed in Figure 3-1 indicate that the fixed effects method with placebo benchmarking based on *Lossless Common board randomization* is too weak to detect the hypothesized director effects. For example, the method has a 0.79 probability to detect effects of size 0.2 on *Return on assets* at the 0.01 confidence level. This effect size translates into a difference between the first and third quartile director's effect (IQR) on *Return on assets* of 3 percentage points, given the pooled sample standard deviation of *ROA* of 11.19 percentage points and the normal distribution of effects across directors. Smaller effects have an even lower probability to be detected. Conversely, an IQR of 3.77 (effect size 0.25) is detected in 99 of 100 trials. This required effect size is somewhat larger than the director effects we would expect to find, based on the extant literature. Masulis, Wang, and Xie (2012) find that *ROA* of firms with foreign independent directors is about 1 percentage point lower than that of firms without FIDs. Fich and Shivdasani (2006) document that *ROA* is about 0.24 percentage points lower in firms with busy boards. For director effects on *Tobin's Q*, analogous considerations show that an IQR of 0.36 would be detected with probability 0.71.

Yermack (1996) finds that Tobin's Q falls by about 0.23 if board size doubles and by about 0.13 if board size rises 50%. Moreover, Fahlenbrach, Low, and Stulz (2010) compare announcement returns of firms appointing outside CEOs to the board with appointments of other outside directors. These announcement returns differ by 1 percent. Thus, the required effect size seems unrealistic.

The outcome of our analysis may be improved by adjusting our test in three broad directions:

(1) *Adjust benchmark.* As discussed above, the statistical power is expected to be relatively low when using placebo data generated by the *Lossless Common board randomization* procedure. This expected low power is due to the overlaps the procedure creates for numerous firms between the treatment and the placebo. Hence, using a different randomization procedure might lead to a more powerful test. An ideal procedure would preserve enough data characteristics to produce sensible results, while avoiding overlaps between treatment and placebo data.

An obvious candidate for a procedure that does not create overlaps consists of using *Lossless Common board randomization* on a sample that excludes firms with excessive risk of overlaps. Avoiding overlaps suggests that statistical power should increase. However, the sample would shrink by nearly 25 percent, suggesting lower statistical power. While a sample size of about 11,000 firm years may still seem to be large, however, we also estimate several thousand coefficients. Thus, it is difficult to assess the overall effect on statistical power.

Approaches that are more fundamentally different arise if the restrictions on reassignment along the time (t) and director (k) dimensions are loosened. That is, a spell of board service with a firm would be allowed to shift to different years and to another director, respectively. Both approaches would increase the number of possible random observations and, therefore, reduce the risk of overlaps. Thus, the methods are expected to have higher statistical power. However, loosening the restrictions on randomization also exposes data characteristics to possible distortions. The type of distortions depends on the specific design of the procedure. Moreover, the distortions' effect on the outcome of our test is difficult to predict: Consequences can range from being negligible to producing entirely irrelevant results. An example for the latter outcome is generating placebo data using *Simple* randomization in the case of directors.

(2) *Increase sample size.* The statistical power of a method generally increases with sample size. While there is no simple way to increase the cross-sectional size of our sample, we might extend the sample period. Our current dataset covers 14 years, ending in 2009. Thus, updating the data would add three years of observations, corresponding to an increase in size of 21 percent. This increase would improve the power of our tests. Given that our sample is already reasonably large, it is uncertain whether updating the sample alone would sufficiently increase power to detecting individual director effects.

(3) *Investigate stronger treatments.* The required difference in the effect on *Return on assets* of 3 percent between the first and third quartile director seems to be excessively high. Yet, the extant literature shows that it is not too far off. Thus, we may find individual director effects if we focus on subsamples of particularly important directors.

The next section follows this third suggestion by focusing on directors who are considered particularly important in the extant literature.

5. Focus on influential directors

The extant literature has found numerous characteristics of boards and directors that affect the outcome of board processes. In this section, we use these insights to identify directors who are expected to be particularly influential. As RiskMetrics provides limited information on the background of the individuals covered, we use matched BoardEx data. This allows us to separately consider individuals whose personal characteristics suggest that they are highly important (*Special directors*) and directors who are active in an environment that emphasizes board importance (*Special situations*). Lastly, we conduct a multivariate analysis.

5.1 Director characteristics data

We gather personal data on the directors we track and test whether this information is related to the directors' fixed effects. Detailed information on the individuals is obtained from BoardEx data. As RiskMetrics and BoardEx data provide no common identifier, we match the datasets by hand relying on the directors' name, age, and affiliations.

Table 3-10 describes what we know about the directors we track. The median director was born in 1942. That is, she was 54 years old at the beginning of the sample period and 67 at its end. *Selective undergraduate education* is an indicator variable that equals 1 if the director graduated from a school that was ranked 'Very Competitive' or higher by Barron's 1981 Profiles of American Colleges and 0 otherwise (Pérez-González 2006). 70 percent of the sample directors attended a selective undergraduate education. *Network contacts* is the number of links BoardEx lists for the director. A link is a connection to another individual in the database that was established through a common affiliation. The sample directors have between 1 and 3,611 network contacts and the distribution of this variable is somewhat skewed. *Mandates (lifetime)* is the number of board and management positions an individual has held over the course of her career. As reputable individuals supposedly are offered more positions, this variable serves as a proxy for reputation (Fama and Jensen 1983; Gilson 1990; Yermack 2004). *Mandates (simultaneously, mean)* is the average number of positions a director has held in a given year. While the literature considers being offered many mandates as something good, the opinion is divided when the positions are simultaneous (Fich and Shivdasani 2006; Field, Lowry, and Mkrtchyan 2013). The median sample director holds 10 positions over the course of his career and 3.30 positions simultaneously (average across all sample years the director is active). *Has been CEO* is an indicator variable for individuals whom BoardEx lists as CEO of a firm over the course of their careers (Fich 2005). *Loss before* and *Loss during* are indicators for directors who more than once joined a firm in a year after it posted negative earnings and directors who more than once experienced net loss during their tenure. Thus, the variables identify directors who served during financial distress or were appointed after financial distress. Arguably, it is in these situations that the board is particularly important (Weisbach 1988). While 33 percent of directors experience financial distress in more than one of their firms, only few directors are more than once appointed immediately after distress. *ASCAR (join)* and *ASCAR (leave)* are the average standardized cumulative abnormal returns in a window of one day before to one day after it was announced the individual would join or leave the board. The announcement dates are provided by BoardEx and refer to the public announcement of the director's

election. A director's compensation is evaluated both relative to other outside directors of the same board (*Total compensation relative to board*) and to all outside directors of the firm's industry (*Total compensation relative to sector*). The differences from the peer group's median compensation are expressed as a fraction of the median compensation. These numbers are averaged over the director's service. Thus, the sample directors receive on average 9 percent more than their board's median compensation and they receive 82 percent more than their industry's median compensation. This latter figure seems to be in line with the earlier observation that the tracked directors are special and are found predominantly in large firms. This information on the directors' backgrounds is used in the next section to form subgroups of directors that are expected to be particularly important.

5.2 Special directors

Table 3-11 presents the findings of fixed effects regressions explaining *Return on assets* and *Tobin's Q* for subgroups of directors that are particularly important according to the extant literature. While we only present the findings for two performance measures, analogous analyses have been conducted for the other two dependent variables without materially different outcomes. In Panel A, the subgroups are built based on the personal background of the directors. For example, according to Adams and Ferreira (2009), having one or more women on the board affects firm outcome. Female directors have fewer attendance problems than their male colleagues and are associated with better monitoring. To evaluate whether the presence of individual women is correlated to firm performance, the table presents the results of fixed effects regressions tracking only female directors (*Female*). The upper row for each performance measure displays the outcome of the treatment estimates. The lower row is the benchmark based on 100 draws of placebo data generated by *Lossless* randomization. To gain additional statistical power, we randomly reassign *Individual* directors rather than the *Common board*. The differences between the two procedures are discussed by Senn (2014). In particular, it is shown that both procedures lead to similar outcomes. *N directors* provides the number of director effects estimated for treatment data and the average number of effects estimated for placebo data. These numbers are in general close to each other. The small difference cannot be avoided as the actual number of directors included depends on the outcome of the random assignment. As director indicators may be linearly dependent on other indicators and of firm fixed effects, some of them need to be omitted from the regressions. When reassigning directors individually, the number of omitted indicators depends on the outcome of the random reassignment. For treatment data, the next column of the table presents the Wald statistic associated with the hypothesis that all director indicators are jointly zero. For the placebo benchmark this column presents the fraction of draws with a Wald statistic greater than the treatment statistic above. This fraction is interpreted analogously to a p-value. For the case of female directors these empirical p-values are 0.34 for *ROA* and 0.81 for *Tobin's Q*. Hence, we do not reject the hypothesis that the treatment Wald statistic results from the same distribution as the placebo statistics. That is, there is no significant individual effect for female directors. The table then describes the distribution of the individual director coefficients. The median coefficient for women on the board is close to zero for both performance measures, and treatment and placebo results alike. The distribution of placebo effects is somewhat narrower (*IQR*) than the treatment effects for both performance measures. However, the Kolmogorov-Smirnov (*K/S*) test indicates that the distributions do not differ significantly from each other.

Waelchli and Zeller (2013) document that the age of a firm's COB is associated with firm performance. In particular, old COBs seem to harm the firm. Therefore, the table investigates subgroups of directors based on their age. *Young* directors are younger (median age across all board positions) than the sample median whereas *Old* directors are older. The sample median age is 64 years. Thus, the directors of our sample are somewhat older than the ones Waelchli and Zeller (2013) investigate (in their sample, 64 years marks the third quartile). As the effect allegedly results from decreasing cognitive abilities, it should be more pronounced in older individuals. The benchmarking of Wald statistics suggests that neither young nor old directors are associated with *Tobin's Q*. However, the empirical p-value of 0.04 for young directors and *ROA* suggests weak significance. The distributions of the estimated effects do not differ.

Masulis, Wang, and Xie (2012) investigate the effect of having foreign independent directors (FIDs). They find that firms with FIDs make better cross border acquisitions in the FID's home market. However, these firms' overall *ROA* is below average. Our fixed effects regressions show no association between the presence of individual FIDs and firm performance.

Next, we consider whether the individual has been one of the 400 wealthiest Americans (Forbes 400). These individuals are highly successful in terms of personal finances. We hypothesize that the factors leading to this success may make these individuals successful in a business setting as well. Table 3-11 shows that the effect of these directors is insignificant (Wald statistic). Thus, wealthy individuals do not seem to be associated with firm performance.

Panel B of Table 3-11 investigates subgroups of directors based on their education. First, we replicate the criterion of Pérez-González (2006): We obtain the name of the institution where the individual attended his or her undergraduate education from the BoardEx database. We classify the institutions based on their selectivity according to Barron's 1981 Profiles of American Colleges.³ The analysis of the Wald statistics and the distribution of estimated coefficients do not indicate that the outcome of treatment data estimates differs from what we expect just randomly. Hence, we investigate other aspects of education. *Ivy-League graduate* are individuals who have obtained a degree from an Ivy-League school. As these schools are considered particularly prestigious elite schools, holding an Ivy-League degree is an alternative measure of ability and access to a valuable network. Indeed, the analysis of Wald statistics indicates significance at the 5 percent level in the *ROA* regressions. However, the distributions of the estimated treatment and placebo effects are barely distinguishable. Moreover, the median of the treatment effects of Ivy-League graduates is -0.35 percent *ROA* whereas the placebo effects' median is 0.14. If anything, the Ivy-League graduates seem to underperform other directors. Similarly, we find little evidence that these directors are associated with *Tobin's Q*. While *Selective undergrad education* and *Ivy-league graduate* focus on selectiveness of education, *Doctoral degree* focuses on its level. However, the presence of individuals with a doctoral degree is not significantly associated with firm performance. The same is true for individuals with a *Law degree*.

³ We are interested in the rating of the institutions at the time our sample individuals attended them. Given that the median individual was born in 1942, she would have attended college in the early 1960s. The oldest edition of Barron's Profiles of American Colleges we could obtain was 1981.

Panel C focuses on the directors' work experience. Fich (2005) finds that CEOs are particularly valuable directors because of their experience and their exceptional managerial abilities. Our analysis does not support an association between the presence of individuals that have experience as a CEO and firm performance. Thus, our findings are in line with the results from Fahlenbrach, Low, and Stulz (2010) who find that *ROA* is unaffected by the appointment of an outside CEO to the board. Agrawal and Knoeber (2001) find that some firms demand directors who play a political role. Similarly, Goldman, Rocholl, and So (2009) show that adding politically connected directors can increase firm value. However, focusing our fixed effects analysis on individuals having served as *Government employee* does not yield results that significantly differ from the placebo results.

Panel D distinguishes directors based on their success in the market. First, we consider directors who have held a large number of mandates. According to Yermack (2004) only successful directors are offered additional positions. Thus, we examine the total number of positions the directors have held over their lifetime as reported by BoardEx for the year 2009. To avoid confounding effects with the directors' age, we regress the number of positions on director age and select all directors with a positive residual. While the analysis of Wald statistics indicates that the presence of these directors may indeed be associated with firm performance, the effects estimated with placebo data are more widely distributed. Thus, the association is economically immaterial.

The last two criteria in Table 3-11 are based on the market reaction when the director's appointment or retirement is announced. A number of articles considers abnormal announcement returns a measure for the impact a director will have on her firm (Rosenstein and Wyatt 1990; Shivdasani and Yermack 1999; Fich 2005; Nguyen and Nielsen 2010). We use announcement dates from BoardEx and select all directors who have a significant average announcement return from the day before the announcement until three days thereafter. When these directors join or leave a board, the market reacts significantly on average. This reaction can either be positive or negative. Hence, we would expect to include both 'good' and 'bad' directors. However, in terms of director fixed effects, this market appreciation does not show up in the performance data. The Wald statistics for both performance measures and groups of directors are indistinguishable from placebo benchmarks.

5.3 Special situations

While Table 3-11 focuses on directors who we expect to be particularly influential because of their background, Table 3-12 investigates circumstances under which the individuals have an increased likelihood of influencing the firm. The first set of situations (Panel A) is based on the individual's role within the board of directors. The criterion we consider is whether the individual is the chairman (*COB*). The *COB* typically chairs the board meetings and sets the board's agenda. Therefore, it seems likely that she has greater influence on board decisions compared to other external directors. Next, we focus on directors who are members of the audit, compensation, or nomination committees. The academic literature (Klein 1998) as well as lawmakers' decisions (Sarbanes Oxley Act) suggest that committees largely drive the outcome of board operations. For similar reasons we consider directors that are highly compensated relative to the other members of the board (data on director compensation are from BoardEx). Arguably, these individuals are paid more than their colleagues because they provide more or higher value services to their firms. Small boards are another situation

where the individual director is important. It is easier for an individual director to impose her personal preferences on a small board than on a larger board. Furthermore, it has been argued that directors holding multiple mandates simultaneously may lack the time to really care about any of those positions (Fich and Shivdasani 2006). The authors define busy directors as individuals holding more than 3 positions at a time. We use BoardEx data to get a wider sample of positions held than the S&P 1,500 firms of the RiskMetrics data. In this sample, we compute the average number of positions held over time, and exclude directors with numbers of 3 or greater. Thus, these directors are less busy and therefore more productive. By doing so, we follow Fich and Shivdasani (2006) who consider individuals being 'busy' if they hold 3 or more positions simultaneously. Lastly, we look at directors who are older than the median board member. We hypothesize that older, more experienced board members find it easier to assert themselves than their younger colleagues.

Given these eight criteria to build subgroups of directors and the two performance measures for each group, we conduct a total of 16 independent tests. Of these tests, only two treatment Wald statistics are significant at the 5 percent level compared to the distribution of placebo statistics (*Nomination committee* and *Small board* with *Tobin's Q* as the performance measure). In both cases the distribution of the estimated coefficients does not differ significantly between treatment and placebo effects. Thus, although the Wald statistics suggest that treatment and placebo results differ significantly at the 5 percent level, the effects are of trivial magnitude according to the test proposed by Cronqvist and Fahlenbrach (2009).

While Panel A considers individual directors who are in a special situation because of how the board is organized, Panel B focuses on directors who serve for firms in situations where the board is particularly important. One such situation is young firms. Arguably, young firms depend more heavily on the individual experience and expertise of directors. This hypothesis is in line with Field, Lowry, and Mkrtchyan (2013) who find that IPO firms have a greater need for advice from their directors. The second situation where directors are particularly important is when the firm is in a crisis. A substantial literature investigates how boards react to poor performance (e.g., by replacing the CEO). Thus, we build a subgroup of individuals whose firms post a net loss during their tenure. The last situation is directors who are appointed in the year after a net loss is posted. Arguably, directors who are appointed shortly after a crisis should be selected with a focus on turning the firm around. All of these subgroups lead to insignificant Wald statistics, given the distribution of placebo statistics. The distribution of individual effects differs significantly for directors serving on firms that report losses. However, the difference is such that placebo effects are more widely dispersed than treatment effects; the opposite of what we would expect if the treatment effects were real.

Overall, we investigate 25 subgroups of directors whom we expect to be particularly influential. For each of these subgroups we repeat the test for two different performance measures. 5 of these 50 tests are significant at the 10 percent level when benchmarked against placebo outcomes (Wald statistics). This number is what we would expect just randomly. Moreover, the significant outcomes do not provide a coherent image of important directors: none of the subgroups seems to matter for both performance measures: 3 groups seem to matter for *ROA* and 2 for *Tobin's Q*. Further, testing according to Cronqvist and Fahlenbrach (2009), the distribution of the treatment effects is significantly

more dispersed in only 1 of these 5 cases. Moreover, in 3 out of 5 cases the placebo estimates have a greater dispersion than the treatment effects. Therefore, we conclude that not even the presence of directors that were selected with regard to their expected importance is associated with abnormal firm performance.

One explanation for the absence of effects for these supposedly important individuals is that we only analyze one director characteristic at a time. The presence of confounding effects may make it necessary to control for multiple characteristics simultaneously. However, subset analysis is an inappropriate tool to do so. Given that we investigate a total of 25 subsets, a bivariate subset analysis would consist of 300 unique pairs of criteria. While some of these combinations would be pointless, such as young and old directors, the number of subsets is impractical. Moreover, this number grows rapidly, as we consider triplets or even larger combinations of criteria. Thus, the next section relates the estimated director effects to director characteristics in a multivariate regression setting.

5.4 Multivariate analysis

The full set of estimated director effects is regressed on the director characteristics described in Table 3-10.⁴ The regressions are non-parametric median regressions. We show the outcome of these non-parametric estimates because we are concerned that the dependent variables (i.e., director effects) are not well behaved. These concerns arise because the director effects are not independent and because their estimation may be affected by serial correlation. However, we arrive at the same conclusions estimating OLS regressions with heteroscedasticity robust standard errors (not shown).

Table 3-13 presents the outcomes of these estimates. In the case of the relation between the director's characteristics and their effect on *ROA* (column (1)), the results suggest that directors of larger firms (*Firm size*) and directors whose appointment triggers a higher abnormal return (*ASCAR (join)*) are typically associated with higher operating performance. Both coefficients are significant at the 5 percent level. Similarly, reputable directors (*Mandates (life time)*) are associated with higher and busy directors (*Mandates (simultaneously)*) with lower operating performance. These latter two effects are in line with what we would expect. However, they are significant only at the 10 percent level. When investigating the director effects on *Tobin's Q* (column (3)), we find that *Firm size* is again positive and significant at the 5 percent level. Being an *Ivy League graduate* is significantly negative. Also significantly negative, although at a lower confidence level, is being present while the firm reports negative earnings (*Loss during*). Busyness (*Mandates (simultaneously)*) is positive at the 10 percent level in the *Tobin's Q* regression. This latter finding is somewhat at odds with the outcome of the *ROA* regression, where it is significantly negative.

These outcomes suggest that there is a statistically significant relation between directors' fixed effects and their personal characteristics. However, evaluating significance in quantile regressions is not straightforward. In particular, the reported significance levels rely on the assumption that the observations of the dependent variable are independent and identically distributed (iid). In our case,

⁴ Correlating director effects with personal characteristics could be viewed as a third additional hypothesis regarding the director effects in the sense of Section 3.2 of this paper. Thus, a significant relation would be sufficient if not necessary condition for the effects being more than noise.

the independence assumption is violated as director fixed effects are estimated jointly in a regression model. Thus, Table 3-13 provides empirical p-values regarding the estimates' overall significance based on placebo benchmarking. We obtain this benchmark by estimating the analogous regressions using placebo director fixed effects as dependent variables. From these estimates we obtain a distribution of 100 pseudo R squares which we compare to the treatment pseudo R square. The reported empirical p-value represents the probability of obtaining a pseudo R square of the magnitude observed using treatment data, given the distribution of placebo outcomes. For *ROA* this value is 0.05, suggesting overall significance, whereas the 0.13 of the *Tobin's Q* estimate is insignificant. Thus, this analysis suggests that director effects on *ROA* are significantly related to director characteristics whereas the effects on *Tobin's Q* are not.

To better understand this finding, we group the director characteristics into three types. The characteristics of the first type relate to the individual (e.g., education). Characteristics of the second type describe the environment an individual is active in. These characteristics are derived entirely from firm-level data (i.e., Board size, Firm size, Firm age, and Loss during / before). The third type of characteristic depends on the interaction of the individual with the firm (e.g., compensation, and relative age). That is, these characteristics are based on data at the director-firm-year level.

In the regressions shown in Table 3-13, the coefficients of the characteristics of type 2 and type 3 are potentially affected by spurious correlation. This concern arises because firm characteristics are on both sides of the regression equation: The director effect on performance (dependent variable) is the firms' abnormal performance averaged across the years the respective director is active in. The director characteristics (independent variables) are based on firm characteristics (type 2) in the years the director is active and characteristics of the director-firm interaction (type 3), respectively. Thus, if there is a direct relation between these firm characteristics and abnormal firm performance, this relation might affect the estimates in Table 3-13. In particular, *Loss during*, the indicator for directors who are on the board of a firm that posts negative earnings, might be mechanically related to the firms abnormal performance. Thus, it would also be related to our estimate of director effects on performance. This relation would be particularly strong if the negative earnings flag persistent underperformance, rather than a transitory event. This scenario seems likely in our data. After all, the residuals of our performance regressions are serially correlated (Senn 2014). If there is a mechanical correlation between the dependent and independent variables in the table, a significant correlation does not necessarily reflect a relation between the director's impact and director characteristics. Thus, this significance cannot be interpreted as evidence that director effects are real.

We provide evidence that spurious correlation presumably contributes to the relation between director effects on *Return on assets* and director characteristics of type 2. For this purpose, we regress the placebo director effects on *ROA* on director characteristics measured after randomization (not shown). That is, the firm characteristics underlying the director characteristics are obtained from the firms the director is randomly reassigned to. As the director effect is a placebo, any correlation between this effect and the firm characteristics would be considered spurious in the context of Table 3-13. Indeed, we find a correlation that is significant at the 10 percent level in 98 of 100 placebo datasets (79 at the

1 percent level). By analogy, we expect the treatment estimates in column (1) to be similarly affected by spurious correlation.

As the director effects on *Tobin's Q* are overall insignificantly correlated with director characteristics, conducting an analogous analysis for column (3) is unnecessary. In addition, the overall insignificance indicates that the *Tobin's Q* estimates are less affected by spurious correlation than *ROA* estimates. This observation, combined with the earlier finding that serial correlation is weaker for *Tobin's Q*, supports the notion that serial correlation exacerbates the effect of spurious correlation in Table 3-13.

The relation between director effects and director characteristics of type 3 may be similarly affected by spurious correlation. After all, these characteristics also include a firm specific component. However, evaluating whether this spurious correlation is likely to affect the outcome in Table 3-13 is infeasible. This evaluation would imply estimating the director characteristics after randomization. However, data on the interaction between the directors and the firms they are randomly reassigned to is unavailable as the interaction is not observed. For example, we would need to know the compensation a director would have received at a firm she never actually served for.

If the estimates of Table 3-13 are affected by spurious correlation, the placebo benchmark used to evaluate the overall significance of these regressions is invalid. The benchmark does not reflect the spurious correlation because director characteristics are measured before randomization (i.e., at the directors' original firms). In consequence, it is too low. For characteristics of type 2, a valid benchmark could be obtained by measuring director characteristics after randomization. However, as type 3 characteristics cannot be measured after randomization, it is unclear how an overall valid benchmark could be obtained.

Columns (2) and (4) of Table 3-13 evaluate the importance of this issue. They show regression estimates that exclude director characteristics of type 2. The empirical p-value associated with the overall significance of the *Return on assets* regression increases to 0.10 compared to 0.05 in column (1). Thus, the significant correlation between director effects on *ROA* and director characteristics in column (1) is strongly driven by type 2 characteristics and therefore is likely to contain a spurious component. In the *Tobin's Q* regression, the empirical p-value increases from 0.13 to 0.55. That is, the estimate turns even less significant. Moreover, the benchmarks used may still be biased downward, as some characteristics that might be affected by spurious correlation remain in the regression (type 3). Indeed, the empirical p-values drop further if these characteristics are excluded (not shown). In sum, the estimates turn overall insignificant when characteristics that are prone to spurious correlation are excluded. This outcome indicates that the director effects are unrelated to director characteristics. Similar results obtain for the remaining two performance measures. These analyses are not shown for brevity.

We conclude that the initially observed correlations between the estimated director effects and the individuals' personal characteristics result from spurious correlation. Thus, they should not be interpreted as evidence for a relation between director effects and director characteristics. When excluding the characteristics that are prone to spurious correlation, the analysis provides no evidence

that the estimated director effects represent an actual treatment effect. Rather, it is in line with the notion that directors are fungible.

6. Conclusions

Senn (2014) illustrates the methodological difficulties when using the fixed effects approach to evaluate the association of individual stakeholders with firm outcomes. The paper concludes that the method should only be used with placebo benchmarking. Doing so, the analysis shows firm performance is not significantly related to the presence of individual directors. In this paper we want to better understand this finding and how it relates to the extant literature. We do so by testing its robustness to an alternative regression specification. Next, we evaluate whether significant effects are missed by simply focusing on the comparison of overall Wald statistics between treatment and placebo results. For example, Cronqvist and Fahlenbrach (2009) distinguish treatment from placebo results by comparing the dispersion of the distribution of the effects. We also test whether the effects have out of sample predictive power and whether the effects correlate across different performance measures (Bertrand and Schoar (2003) refer to this kind of correlation as 'style'). The outcomes of these additional analyses support the conclusion that individual outside directors have little impact on performance and are therefore fungible.

Since these findings seem to be at odds with articles suggesting that individual directors can influence firm outcomes, we investigate the statistical power of our tests. That is, what magnitude of effect could be detected using our data and methods? This question is difficult to answer, as the power of the tests depends on how we generate the placebo datasets used as a benchmark for statistical significance.

We look at CEO fixed effects to gain additional insights into this issue. As the effect of CEOs is expected to be stronger than that of outside directors, CEOs are an ideal testing ground for our methods. We find that individual CEOs are significantly related to firm performance. The difference between a first and third quartile CEO effect on *ROE* is 2.86 percentage points greater using treatment instead of placebo data. The respective difference on *Tobin's Q* is 0.15. These numbers give us an idea of the size of effect directors would need to have for us to detect them. Further, these numbers result from placebo benchmarks that focus on statistical power. The findings disappear when using a more restrictive benchmark that better preserves certain characteristics of the treatment data. Hence, the association of individual CEOs and performance is strong enough to be picked up by some variants of the method but not by others.

Unfortunately, the randomization procedure that detects CEO effects is unavailable to analyze director data. That is, if it is applied to directors, it generates placebo data that do not accurately represent the characteristics of treatment data. These distortions lead to placebo effects of a significantly greater magnitude than the treatment effect – a meaningless result. The randomization that is available leads to overall less powerful tests. A power analysis using simulated data finds that director effects would need to be large for this method to detect them. For example, the standard deviation of the effect on

ROA and *Tobin's Q* across directors would need to be 20 to 25 percent of the standard deviation of the respective performance measure.

Given the limited statistical power of the available method, we focus our search for individual director effects on specific subgroups of board members. In particular, we consult the extant literature to select directors who are expected to be particularly influential. We obtain as many significant results as we would expect just randomly due to multiple testing (5 out of 50 independent tests are significant at the 10 percent level). Moreover, we note that for none of the subgroups the outcome is significant for both ROA and *Tobin's Q*. With regard to Cronqvist and Fahlenbrach's (2009) criterion of economic significance, we find a significant difference between the distribution of treatment and placebo estimated individual director coefficients only in 1 of the 5 cases for which the Wald statistic suggests a significant relation. Thus, we conclude that not even those directors that the literature suggests to be particularly important are significantly associated with firm performance. In a multivariate regression setting, we find no significant relation between individual director effects on performance and director characteristics either.

Overall, the results of Senn (2014) turn out to be robust. Despite extensive testing, we find hardly any evidence that the presence of individual directors is associated with firm performance. Thus, we maintain the conclusion that directors are fungible. That is, there is no evidence that individual outside directors consistently improve or harm the financial performance of the firms they work for. In that sense, there are no super stars or losers. With regard to the fixed effects method, our analysis of CEOs shows that placebo benchmarking is sensitive to how we obtain the benchmark. This finding reinforces the notion that placebo benchmarking is a double edged sword: On the one hand, it allows us to evaluate the significance of fixed effects in panel regressions despite the issues with Wald tests (Fee, Hadlock, and Pierce 2013). On the other hand, when choosing a benchmark, we need to trade off accurate representation of the data structure against statistical power. The absence of an objective measure for the adequacy of the benchmark used leaves considerable discretion to the researcher.

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) divided 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 (mkvalt) less common equity (ceq) and deferred taxes (txdb). Winsorized at the 0.5% level.
Investment	Capital expenditures (capx) 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 liabilities (dlc), and common equity (ceq) times 100. Winsorized at the 0.5% level.
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 date 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 from RiskMetrics / IRRC). If this variable is missing for some directors, these directors are assumed to have no 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.

Tables

Table 3-1

The table estimates individual director effects on firm performance. The performance measures are centered around industry mean performance and standardized annually. Industries are defined according to the Fama and French 48 industries classification. The estimates control for *Size*, *Investment*, *Leverage*, *Board size*, *Board independence*, *Board tenure*, *Board ownership* and *CEO is COB* (variable definitions according to Appendix A). Squared terms of all board related variables are included. Rows labeled *Treatment* provide estimation outcomes based on the actual treatment data. *Placebo* refers to outcomes obtained from data randomized by the *Lossless* procedure. *N of draws* lists the number of random draws. *N of directors* is the number of director indicators included in the regressions, averaged across the different draws. For the treatment data, *Wald test / empirical p-value* provides the Wald statistic associated with the joint significance of the individual director indicators. For placebo, it reports the fraction of draws leading to Wald statistics greater than the treatment outcome. The distribution of estimated director fixed effects is described by its inter quartile range (*IQR*) and *Median*. Kolmogorov Smirnov (*KS test*) and *Wilcoxon* test compare the distributions of treatment and placebo effects. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	N of draws	N of directors	Wald test / empirical p-value	Director fixed effects coefficients			
				IQR	Median	KS test (d)	Wilcoxon (z)
Return on assets							
Treatment	1	2,058	1,000,000	0.78	0.00		
Placebo	100	2,058	0.19	0.78	0.00	0.02	1.07
Return on equity							
Treatment	1	2,058	6,100,000	0.63	0.04		
Placebo	100	2,058	0.03	0.61	0.02	0.04**	2.64***
Sales to assets							
Treatment	1	2,058	47,711	0.81	-0.38		
Placebo	100	2,058	0.50	0.80	-0.02	0.03	2.21**
Tobin's Q							
Treatment	1	2,058	2,200,000	0.76	-0.01		
Placebo	100	2,058	0.05	0.79	-0.04	0.05***	3.53***

Table 3-2

The table estimates individual director effects on firm performance. The estimated model specification is identical to Table 3-1. The differences between the tables are: Rows labeled *Treatment* provide estimation outcomes based on the actual treatment data, where only a random subset of 33 percent of directors is tracked in the regressions. *Placebo* refers to outcomes based on data randomized using the *Simple* procedure. The table describes the distribution of the Wald statistics associated with the overall significance of director indicators over 100 random draws. The distributions of treatment and placebo effects are compared using Kolmogorov Smirnov (*KS test*) and *Wilcoxon* tests. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	N of draws	N of directors	1 st quartile	Median	Wald test			Director fixed effects coefficients			
					3 rd quartile	KS test (d)	Wilcoxon (z)	IQR	Median	KS test (d)	Wilcoxon (z)
Return on assets											
Treatment	100	681	971.48	1'642	2'645			0.66	0.00		
Placebo	100	649	10'222	29'458	152'411	0.83***	-11.56***	0.88	-0.01	0.07***	4.56***
Return on equity											
Treatment	100	681	855.97	2'113	5'125			0.45	0.03		
Placebo	100	649	19'866	73'225	235'432	0.72***	-10.82***	0.62	0.01	0.08***	10.76***
Sales to assets											
Treatment	100	681	1'264	2'646	5'410			0.77	-0.03		
Placebo	100	649	11'753	55'226	229'583	0.70***	-9.99***	0.90	-0.05	0.05***	7.25***
Tobin's Q											
Treatment	100	681	1'290	2'038	4'070			0.70	-0.03		
Placebo	100	649	14'722	47'148	204'065	0.82***	-11.35***	0.87	-0.07	0.08***	14.60***

Table 3-3

This table evaluates the out of sample predictive power of director fixed effects. It does so by estimating 2 separate effects per director: one for the period up to 2002, the other for the rest of the sample. The table investigates whether these two sets of coefficients are related. For this purpose, it reports *Spearman correlations* and cross-tables of the two coefficients (*Sign*). *p-value* tests the significance of the association against 0. The empirical p-values test the significance of the association against the outcome of a placebo analysis. That is, they are the fractions of placebo draws that lead to a stronger association than the treatment data. *Simple* and *Lossless* refer to the respective randomization procedure used to generate the placebo benchmark.

	Return on assets		Return on equity		Sales to assets		Tobin's Q	
Spearman correlation	0.30		0.34		0.34		0.26	
p-value	0.00		0.00		0.00		0.00	
empirical p-value (<i>Simple</i>)	0.31		0.36		0.41		0.41	
empirical p-value (<i>Lossless</i>)	0.36		0.27		0.50		0.38	
Sign	After 2002		After 2002		After 2002		After 2002	
	< 0	> 0	< 0	> 0	< 0	> 0	< 0	> 0
Before or in 2002	< 0	> 0	< 0	> 0	< 0	> 0	< 0	> 0
	0.28	0.23	0.31	0.21	0.25	0.20	0.29	0.22
	0.17	0.32	0.16	0.32	0.21	0.35	0.17	0.32
Fraction with same sign	0.60		0.63		0.60		0.61	
p-value	0.00		0.00		0.00		0.00	
empirical p-value (<i>Simple</i>)	0.33		0.35		0.36		0.46	
empirical p-value (<i>Lossless</i>)	0.33		0.31		0.47		0.50	

Table 3-4

Similarly to Table 3-3, this table tests the out of sample predictive power of director fixed effects. While the previous table makes a cross sectional sample cut in the year 2002, Table 3-4 divides the directors' positions into two random longitudinal subsamples. *Spearman correlation* and cross-tables of the two coefficients (*Sign*) evaluate whether the two subsamples' coefficients are associated. *p-value* tests the significance of the association against 0. Panel A considers a single random sample split, Panel B repeats the formation of subsamples 100 times and reports the fraction of outcomes that are significant at different confidence levels. As the association is typically insignificant, no placebo benchmarking is required.

	Return on assets	Return on equity	Sales to assets	Tobin's Q				
<i>Panel A: Single draw</i>								
Spearman correlation	-0.01	-0.02	0.06	0.01				
p-value	0.66	0.51	0.08	0.80				
Sign	Subgroup 2		Subgroup 2		Subgroup 2		Subgroup 2	
	< 0 > 0		< 0 > 0		< 0 > 0		< 0 > 0	
Subgroup 1	< 0	> 0	< 0	> 0	< 0	> 0	< 0	> 0
	0.25	0.23	0.24	0.25	0.26	0.24	0.23	0.25
	> 0	0.27	> 0	0.27	> 0	0.26	> 0	0.25
	0.26	0.24	0.26	0.24	0.26	0.24	0.25	0.27
Fraction with same sign	0.51	0.48	0.50	0.50				
p-value	0.63	0.28	0.86	0.85				
<i>Panel B: 100 draws</i>								
Spearman correlation								
Rho (mean)	0.00	-0.02	0.05	0.00				
Fraction of significant outcomes								
$\alpha=0.10$	0.00	0.11	0.41	0.00				
$\alpha=0.05$	0.00	0.07	0.19	0.00				
$\alpha=0.01$	0.00	0.02	0.04	0.00				
Sign								
Fraction with same sign (mean)	0.51	0.49	0.51	0.50				
Fraction of significant outcomes								
$\alpha=0.10$	0.13	0.12	0.02	0.02				
$\alpha=0.05$	0.03	0.07	0.00	0.00				
$\alpha=0.01$	0.00	0.00	0.00	0.00				

Table 3-5

Panel A uses the director fixed effects coefficients estimated in Senn (2014). Effects on the performance measure in the column-head are regressed on the effects on the measure in the row-head. Individual regressions are estimated for each pair of dependent variables. The three values per pair are based on least squares regressions with Huber-White standard errors (top), least squares weighted by the inverse of the individual coefficients' standard errors (middle), and median regressions (bottom). Panel B repeats the analysis using coefficients obtained from randomized data. The data was generated using board-level randomization that focuses on the number of outcomes possible (*Simple*) or on maintaining data characteristics (*Lossless*). ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

<i>Panel A: Original data</i>							
	Return on assets	Return on equity	Sales to assets				
Return on equity	0.07*** 0.12*** 0.07***						
Sales to assets	0.13*** 0.15*** 0.14***	0.16*** 0.14*** 0.15***					
Tobin's Q	4.22*** 4.71*** 4.72***	4.73*** 6.68*** 5.69***	9.60*** 11.57*** 10.09***				
<i>Panel B: Placebo data</i>							
	<i>Simple randomization</i>			<i>Lossless randomization</i>			
	Return on assets	Return on equity	Sales to assets	Return on equity	Return on assets	Return on equity	Sales to assets
Return on equity	0.12*** 0.19*** 0.14***			Return on equity	0.10*** 0.17*** 0.11***		
Sales to assets	0.14*** 0.15*** 0.16***	0.16*** 0.21*** 0.17***		Sales to assets	0.14*** 0.16*** 0.16***	0.17*** 0.20*** 0.18***	
Tobin's Q	4.22*** 4.61*** 4.94***	4.33*** 5.69*** 5.07***	9.41*** 10.56*** 10.28***	Tobin's Q	4.26*** 4.84*** 4.75***	5.34*** 6.95*** 5.63***	8.92*** 10.37*** 9.87***

Table 3-6

The table describes the service of the CEOs we track in our analysis. Panel A counts the number of different firms the CEOs are tracked in. Panel B reports the number of firms a CEO is simultaneously tracked in for every year this executive is active. Panel C counts the number of different sample years a CEO is tracked in. Panel D describes the duration of the tracked mandates, that is, the number of years a CEO is tracked in a given firm. Panel E provides the number of tracked CEOs in each sample firm-year.

<i>Panel A: Number of tracked positions by CEO</i>			
	Number of CEOs	Percent	Cumulative
2	52	98.11	98.11
3	1	1.89	100
Total	53	100	Mean 2.02

<i>Panel B: Number of positions held simultaneously in a given CEO year</i>			
	CEO years	Percent	Cumulative
1	483	94.34	97.39
2	29	5.66	100
Total	512	100	Mean 1.06

<i>Panel C: Number of years in the sample</i>			
	CEOs	Percent	Cumulative
5	1	1.89	1.89
6	5	9.43	11.32
7	6	11.32	22.64
8	11	20.75	43.40
9	5	9.43	52.83
10	5	9.43	62.26
11	9	16.98	79.25
12	3	5.66	84.91
13	4	7.55	92.45
14	2	3.77	96.23
17	1	1.89	98.11
20	1	1.89	100
Total	53	100	Mean 9.66

<i>Panel D: Duration of mandates (CEO-firm combinations)</i>			
	Mandates	Percent	Cumulative
3	34	31.78	31.78
4	21	19.63	51.41
5	19	17.76	69.17
6	12	11.21	80.38
7	9	8.41	88.79
8	3	2.80	91.59
9	1	0.93	92.52
10	4	3.74	96.26
11	1	0.93	97.19
13	2	1.87	99.06
15	1	0.93	100
Total	107	100	Mean 5.05

<i>Panel E: Number of tracked CEOs by firm-year</i>			
	Firm-years	Percent	Cummulative
0	983	64.50	64.50
1	541	35.50	100
Total	1'524	100	Mean 0.35

Table 3-7

CEO fixed effects are estimated using specifications that follow the estimates of Table 3-5 in Senn (2014). There are 2 major differences in the specification: (1) Firms that never have a tracked CEO are not included in the regression sample. (2) When estimating individual director effects, Senn (2014) controls for board characteristics. Doing so is unnecessary when investigating CEO effects. These 2 changes improve comparability with the estimates of Bertrand and Schoar (2003) and Fee, Hadlock, and Pierce (2013). The table groups coefficients and reports Wald statistics from testing joint significance of these groups. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Return on assets	Return on equity	Sales to assets	Tobin's Q
Individual CEO effects	620.52***	159.08***	655.43***	584.51***
X_{it}	2.97**	5.15***	7.42***	8.12***
Firm fixed effects	120'000***	9'577***	120'000***	150'000***
Year fixed effects	3.42***	2.88***	2.37***	3.92***
Adjusted R ²	0.54	0.11	0.77	0.62
N (firm-years)	1'524	1'524	1'524	1'524

Table 3-8

The outcome of CEO fixed effects regressions (*Treatment*) is benchmarked against placebo estimates. *Placebo* reports results from datasets generated by *Lossless* randomization. *N of draws* indicates the number of repeated random draws used in the placebo benchmark. *N of CEOs* provides the number of CEOs tracked in the regression. *Wald test* describes the test statistic associated with the joint significance of CEO fixed effects. For placebo data, the quartiles of the distribution of Wald statistics across the repeated random draws are provided. The Wald statistic of treatment data is listed under *Median*. *Empirical p-value* is the fraction of placebo draws with Wald statistics greater than the treatment statistic. *CEO fixed effects coefficients* describes the distribution of the estimated individual CEO coefficients (i.e., their *IQR* and *Median*). *K/S test* and *Wilcoxon* provide the test statistics of Kolmogorov Smirnov and Wilcoxon rank-sum tests comparing the distributions of treatment and placebo effects. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Number of draws	Number of CEOs	Wald test				CEO fixed effects coefficients			
			1 st quartile	Median	3 rd quartile	empirical p-value	IQR	Median	KS test (d)	Wilcoxon (z)
Return on assets										
Treatment	1	52		278.58			7.37	-1.02		
Placebo	100	52	213.28	408.79	930.24	0.65	6.58	-0.90	0.08	0.36
Return on equity										
Treatment	1	52		289.93			30.95	-0.53		
Placebo	100	52	77.51	170.69	434.87	0.35	28.25	-3.49	0.12	1.18
Sales to assets										
Treatment	1	52		847.20			24.75	1.26		
Placebo	100	52	251.40	533.36	1'004	0.29	25.80	-2.52	0.12	0.68
Tobin's Q										
Treatment	1	52		1'474			0.93	-0.09		
Placebo	100	52	948.54	1'512	2'461	0.53	0.86	-0.03	0.06	0.09

Table 3-9

The outcome of CEO fixed effects regressions is benchmarked against placebo estimates. The placebo benchmark (*Placebo*) obtains from datasets that are randomized in a way that focuses on a high number of possible random outcomes (*Simple*). Treatment estimates are adjusted for the lower number of CEO indicators included in the placebo regressions by selecting a random subset (33 percent) of CEOs to be tracked (*Subset*). *N of draws* indicates the number of repeated random draws used to select the subset of tracked CEOs and to generate the placebo benchmark, respectively. *N of CEOs* provides the number of CEOs tracked in the regression. *Wald test* describes the distribution of Wald statistic associated with the joint significance of CEO fixed effects across the repeated random draws. *KS test* and *Wilcoxon* provide the test statistics of Kolmogorov Smirnov and Wilcoxon rank-sum tests comparing the distributions of treatment and placebo Wald statistics. *CEO fixed effects coefficients* describes the distribution of the estimated individual CEO coefficients (i.e., their *IQR* and *Median*). *KS test* and *Wilcoxon* provide the test statistics of Kolmogorov Smirnov and Wilcoxon rank-sum tests comparing the distributions of treatment and placebo effects. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Number of draws	Number of CEOs	Wald test					CEO fixed effects coefficients			
			1 st quartile	Median	3 rd quartile	KS test (d)	Wilcoxon (z)	IQR	Median	KS test (d)	Wilcoxon (z)
Return on assets											
Treatment	100	17.67	21.10	29.15	53.88			6.58	-1.08		
Placebo	100	18.15	7.34	10.98	19.91	0.58***	7.76***	6.58	-0.25	0.09***	-3.16***
Return on equity											
Treatment	100	17.67	6.69	9.15	21.72			28.62	0.61		
Placebo	100	18.15	4.05	7.43	15.69	0.31***	3.08***	25.76	0.67	0.06***	-0.71
Sales to assets											
Treatment	100	17.67	11.95	25.98	46.14			17.15	-0.24		
Placebo	100	18.15	5.81	10.63	25.79	0.32***	4.41***	22.53	0.59	0.08***	-0.67
Tobin's Q											
Treatment	100	17.67	39.76	104.10	171.94			0.87	-0.12		
Placebo	100	18.15	5.17	8.68	16.19	0.74***	10.17***	0.72	-0.07	0.07***	0.28

Table 3-10

The table provides summary statistics on a set of variables describing the background of tracked directors and the environment they are active in. Most of these variables are provided directly by RiskMetrics or BoardEx. *Selective undergraduate education* is a binary variable indicating directors holding an undergraduate degree that was ranked as 'Very Competitive' or higher by Barron's 1981 Profiles of American Colleges. *Graduated with distinction* is a binary variable indicating directors graduating with a distinction such as summa cum laude. *Network contacts* is the number of contacts listed for a director's network by BoardEx. *Mandates (life time)* is the total number of positions a director has held over her career according to BoardEx. *Mandates (simultaneously, mean)* is the number of positions a director holds simultaneously, averaged across the years this director is tracked in our dataset. *Tenure* is the average number of years the director is tracked with a firm inside our sample. *Loss before* and *Loss during* are binary variables indicating directors nominated immediately after the firm posted a loss and directors who experienced a loss during their tracked mandates, respectively. *ASCAR (join / leave)* is the average standardized cumulative abnormal returns of the firm's stock upon announcement of the director's nomination and departure, respectively. *Age relative to board* is the director's age in years relative to the median age of the board members, averaged over the firm-years this director is tracked in our sample. *Board size* is the average number of members of the boards the director in question is tracked on. *Total compensation relative to board* is the director's compensation relative to the board's median compensation measured in percent and averaged over the director's service. *Total compensation relative to sector* is the director's compensation relative to the sector median in percent and averaged over the director's service. *Firm size* is the average of the natural logarithm of total assets of the firms the director serves. *Firm age* is the average number of years since incorporation of the firm the director serves.

	N	Mean	Std. dev.	Median	IQR	Min.	Max.
Year of birth	1'798	1941	7.32	1942	11.00	1929	1955
Female	1'798	0.14					
Non-US citizen	1'798	0.23					
Selective undergraduate education	1'626	0.70					
Graduated with distinction	1'706	0.06					
Holds doctorate	1'706	0.26					
Holds law degree	1'706	0.14					
Ivy League graduate	1'706	0.37					
Forbes 400 richest Americans	1'798	0.01					
Network contacts	1'798	1'290	979	1'015	1'229	1	3'611
Mandates (life time)	1'798	11.23	5.81	10.00	7.00	4.00	25.00
Mandates (simultaneously, mean)	1'797	3.69	1.68	3.30	2.24	1.42	7.69
Has been CEO	1'798	0.64					
Has been chairman	1'798	0.65					
Has worked for government	1'798	0.05					
Tenure	1'798	6.16	1.91	6.00	3.00	3.50	10.00
Loss before	1'798	0.01					
Loss during	1'798	0.33					
ASCAR (join)	562	-0.02	0.79	0.00	0.99	-1.55	1.49
ASCAR (leave)	1'034	0.09	0.75	0.05	0.88	-1.33	1.72
Board size	1'798	10.15	1.63	10.13	2.42	7.33	13.21
Age relative to board	1'798	0.61	5.27	1.09	7.35	-10.18	9.67
Total compensation relative to board	1'777	0.09	0.24	0.00	0.09	-0.10	0.93
Total compensation relative to sector	1'777	0.82	1.17	0.40	1.03	-0.29	4.31
Firm size	1'798	8.19	1.14	8.11	1.72	6.32	10.28
Firm age	1'798	32.85	16.71	30.35	26.08	9.58	66.37

Table 3-11

Analogously to Table 2-11 of Senn (2014), director fixed effects estimates (respective upper row) are compared to placebo results (respective lower row). Placebo datasets are generated using the *Lossless Individual randomization* procedure. The table focuses on different groups of directors. These groups are formed based on the directors' characteristics (see Table 3-10). *Female* are directors that are listed as female by BoardEx. *Young* and *Old* directors are individuals whose age is smaller or greater than the sample median. *Non-US citizens* are individuals whose nationality is listed as anything else than 'American' by BoardEx. *Wealthy* directors are individuals who have appeared on the Forbes list of the 400 richest Americans. A director has a *Selective undergraduate education* if she holds an undergraduate degree (according to BoardEx) from a school that is ranked 'Very Competitive' or higher by Barron's 1981 Profiles of American Colleges. *Ivy-League graduate* are directors holding a degree from Ivy-League member school. *Law degree* and *Doctoral degree* are directors holding a respective degree. *CEOs* are directors who have held a position as a CEO at some point prior to their directorship (according to BoardEx). A director qualifies as a *Government employee* if BoardEx reports her as having held a position working for a government. *High number of mandates* includes directors who have held relatively many director positions given their age. That is, the total number of mandates held is regressed on age. Directors with positive residuals are selected for the subgroup. *Significant AR when joining / leaving* include directors whose announcements to join or retire from boards trigger significant abnormal stock returns on average. The dependent variables analyzed are Return on assets (ROA) and Tobin's Q. *N of directors* provides the number of directors tracked in the estimates. *Wald test / empirical p-value* provides the Wald statistic associated with the overall significance of director indicators for treatment estimates. For placebo estimates, the column gives the fraction placebo outcomes that are greater than the treatment outcome. *Director fixed effects coefficients* describes the distribution of individual director effects (*IQR* and *Median*). *KS test* provides the Kolmogorov Smirnov test statistic comparing the distributions of treatment and placebo outcomes. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

		N of directors	Wald test / empirical p-value	Director fixed effects coefficients		
				IQR	Median	K/S test (d)
<i>Panel A: Personal background</i>						
Female	ROA	259	8,488	5.89	-0.03	
		259.28	0.34	6.35	-0.07	0.03
	Tobin's Q	259	1,067	0.61	-0.02	
		259.28	0.81	0.63	0.01	0.04
Young	ROA	854	1,654,992	6.00	0.23	
		859.22	0.04	6.41	0.09	0.04
	Tobin's Q	854	212,345	0.62	0.03	
		859.22	0.29	0.65	0.03	0.02
Old	ROA	895	149,707	6.40	-0.22	
		901.2	0.41	6.27	0.00	0.04
	Tobin's Q	895	80,425	0.68	-0.02	
		901.2	0.45	0.65	-0.01	0.02
Non-US citizen	ROA	53	220.06	5.02	0.48	
		52.98	0.76	6.07	0.61	0.09
	Tobin's Q	53	492.03	0.72	-0.01	
		52.98	0.41	0.62	0.04	0.07
Wealthy	ROA	18	195.50	9.86	0.79	
		18	0.84	6.69	0.80	0.10
	Tobin's Q	18	271.79	0.89	-0.15	
		18	0.49	0.64	0.02	0.20
<i>Panel B: Education</i>						
Selective under-graduate education	ROA	1,137	2,743,863	6.25	0.04	
		1,144	0.40	6.42	0.02	0.02
	Tobin's Q	1,137	271,078	0.65	0.00	
		1,144	0.80	0.65	0.01	0.02

		N of directors	Wald test / empirical p-value	Director fixed effects coefficients		
				IQR	Median	K/S test (d)
Ivy-League graduate	ROA	635	376,487	6.56	-0.35	
		637.06	0.05	6.28	0.14	0.05*
	Tobin's Q	635	4,906	0.65	-0.03	
		637.06	0.64	0.65	0.01	0.05*
Law degree	ROA	241	1,609	5.66	0.79	
		240.95	0.78	5.83	-0.01	0.10**
	Tobin's Q	241	1,512	0.64	0.06	
		240.95	0.64	0.61	0.01	0.07
Doctoral degree	ROA	444	1,950	5.86	0.17	
		443.53	0.76	6.07	-0.05	0.04
	Tobin's Q	444	6,592	0.65	0.03	
		443.53	0.37	0.64	0.01	0.04
<i>Panel C: Experience</i>						
CEOs	ROA	1,130	30,341	5.70	0.39	
		1,136	0.95	6.42	0.11	0.04
	Tobin's Q	1,130	13,144	0.65	0.00	
		1,136	0.99	0.65	0.01	0.02
Government employee	ROA	139	3,780	6.13	-0.48	
		138.96	0.68	6.07	-0.17	0.08
	Tobin's Q	139	58,063	0.68	-0.03	
		138.96	0.15	0.62	-0.01	0.05
<i>Panel D: Market evaluation of director</i>						
High number of mandates	ROA	704	2,446,705	6.20	0.30	
		706.78	0.02	6.39	0.01	0.03
	Tobin's Q	704	9,257	0.66	0.00	
		706.78	0.38	0.64	0.00	0.03
Significant AR when joining	ROA	44	26.91	4.96	0.20	
		44	0.92	5.45	-0.32	0.14
	Tobin's Q	44	181.65	0.63	0.03	
		44	0.37	0.59	0.04	0.16
Significant AR when leaving	ROA	69	135.16	3.39	0.50	
		68.99	0.85	5.41	0.21	0.16*
	Tobin's Q	69	70.47	0.57	-0.03	
		68.99	0.97	0.60	0.00	0.07

Table 3-12

Analogously to Table 2-11 of Senn (2014), director fixed effects estimates (respective upper row) are compared to placebo results (respective lower row). The placebo datasets are generated using the *Lossless Individual randomization* procedure. The table focuses on different groups of directors. These groups are formed based on the circumstances of the directors' board activity. In the *COB* subgroup we track only the chairmen of the boards. *Audit*, *Compensation*, and *Nomination committee* track the members of committees. *Highly compensated directors* receive a higher compensation than the majority of non-executive members of their board in most of the years served. *Small board* includes directors who on average serve on boards smaller than the median board of the sample. Directors classified as *Non busy* hold on average less than 3 positions simultaneously according to BoardEx. *Relatively old* directors are part of the older half of the board in the majority of years they serve. In *Young firm*, directors are only tracked in firms that are younger than the median firm (i.e., 21.25 years since incorporation). *Loss during / before mandate* tracks mandates during which a loss occurred and mandates that started in the year after the firm posted a loss. The dependent variables analyzed are Return on assets (*ROA*) and *Tobin's Q*. *N of directors* provides the number of directors tracked in the estimates. *Wald test / empirical p-value* provides the Wald statistic associated with the overall significance of director indicators for treatment data estimates. For placebo estimates, the column gives the fraction placebo outcomes that are greater than the treatment outcome. *Director fixed effects coefficients* describes the distribution of individual director effects (*IQR* and *Median*). *KS test* provides the Kolmogorov Smirnov test statistic comparing the distributions of treatment and placebo outcomes. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

		N of directors	Wald test / empirical p-value	Director fixed effects coefficients		
				IQR	Median	KS test (d)
<i>Panel A: Role on board</i>						
COB	ROA	873	293,100	5.78	0.21	
		875.29	0.20	6.29	-0.01	0.03
	Tobin's Q	873	71,344	0.63	0.01	
		875.29	0.45	0.63	0.00	0.03
Audit committee	ROA	701	95,344	4.80	0.19	
		700.76	0.38	5.42	0.10	0.03
	Tobin's Q	701	194,780	0.56	0.01	
		700.76	0.29	0.56	0.02	0.03
Compensation committee	ROA	701	160,262	5.03	-0.01	
		701.74	0.37	5.55	0.12	0.05
	Tobin's Q	701	162,236	0.53	0.00	
		701.74	0.38	0.57	0.00	0.02
Nomination committee	ROA	601	830,127	5.27	-0.05	
		600.99	0.15	5.21	-0.07	0.03
	Tobin's Q	601	12,191,448	0.49	0.00	
		600.99	0.04	0.54	0.00	0.03
Highly compensated	ROA	202	927	7.93	0.49	
		201.93	0.99	7.07	-0.23	0.08
	Tobin's Q	202	9,046	0.76	0.07	
		201.93	0.24	0.71	-0.03	0.08
Small board	ROA	635	1,629	8.04	-0.11	
		641.85	0.93	6.98	0.02	0.05*
	Tobin's Q	635	1,033,824	0.79	-0.03	
		641.85	0.05	0.69	0.01	0.04
Non busy	ROA	1'406	655'759	6.44	-0.06	
		1'416	0.14	6.46	0.04	0.02
	Tobin's Q	1'406	32'006	0.60	-0.03	
		1'416	0.40	0.66	0.01	0.05***

		N of directors	Wald test / <i>empirical</i> p-value	Director fixed effects coefficients		
				IQR	Median	KS test (d)
Relatively old	ROA	824	448,680	6.15	-0.22	
		829.17	0.22	6.47	0.07	0.05*
	Tobin's Q	824	114,578	0.66	-0.03	
		829.17	0.23	0.66	-0.01	0.03
<i>Panel B: Situation of the firm</i>						
Young firm	ROA	512	16'763	7.66	-0.03	
		521.57	0.26	7.20	0.11	0.04
	Tobin's Q	512	5'345	0.82	0.01	
		521.57	0.51	0.72	0.01	0.05
Loss during mandate	ROA	589	15'017	6.59	-1.13	
		599.02	0.45	6.75	-0.15	0.11***
	Tobin's Q	589	1'869	0.59	-0.05	
		599.02	0.80	0.67	-0.01	0.05*
Loss before mandate	ROA	18	22.06	5.89	-0.30	
		18	0.68	6.43	-0.06	0.15
	Tobin's Q	18	343.24	0.63	-0.01	
		18	0.14	0.62	0.05	0.14

Table 3-13

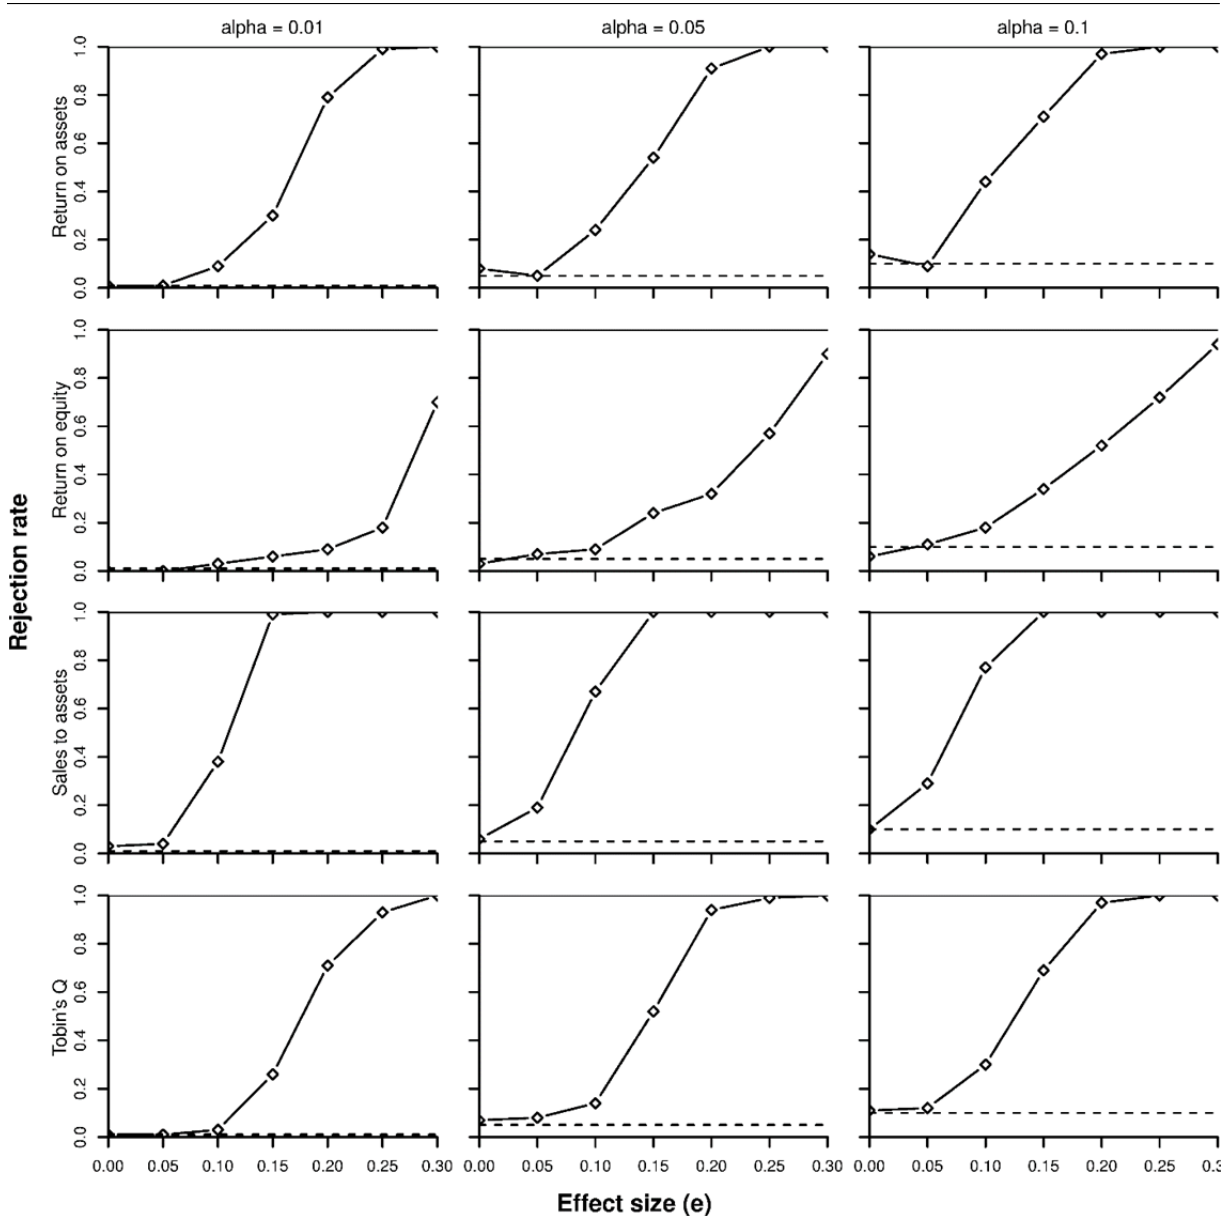
The table presents nonparametric median regressions of the director fixed effects estimated in Table 2-5 of Senn (2014) on the director characteristics described in Table 3-10. The *empirical p-value* assesses the overall significance of these regressions. It is obtained from performing analogous regressions using the director effects estimated from placebo data (*Lossless*) as the dependent variable. The empirical p-value is the fraction of placebo draws with greater pseudo R square than the treatment regression. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.1-level, respectively.

	Return on assets		Tobin's Q	
	(1)	(2)	(3)	(4)
Year of birth	0.02	0.01	0.00	0.00
Female	-0.46	-0.41	-0.05	-0.05
Non-US citizen	0.14	-0.10	0.03	-0.01
Selective undergrad education	0.08	0.03	-0.02	-0.02
Graduated with distinction	-0.78	-0.44	-0.06	-0.09
Holds doctorate	0.42	0.36	0.04	0.02
Holds law degree	0.40	0.56	0.03	0.05
Ivy League graduate	-0.33	-0.11	-0.09**	-0.08**
Forbes 400 richest Americans	-0.49	-0.33	-0.04	-0.05
Network contacts	0.00	0.00	0.00	0.00
Mandates (life time)	0.06*	0.05	-0.01	-0.01
Mandates (simultaneously)	-0.22*	-0.25**	0.03**	0.03**
Has been CEO	0.16	0.30	-0.03	-0.02
Has been chairman	0.17	0.31	0.02	0.02
Has worked for government	-0.61	-0.51	-0.05	-0.05
Tenure	0.02	-0.04	0.00	0.00
Loss before	1.64		-0.18	
Loss during	-0.29		-0.06*	
ASCAR (join)	0.66**	0.61**	-0.01	-0.01
ASCAR (leave)	0.16	0.25	0.03	0.03
Board size	-0.14		-0.02	
Age relative to board	-0.07	-0.07	0.00	0.00
Total compensation relative to board	0.19	-0.11	0.11*	0.08
Total compensation relative to sector	-0.10	0.12	-0.02	0.00
Firm size	0.62**		0.06**	
Firm age	-0.01		0.00	
Constant	-0.20	-0.38	0.05	0.04
R-squared	0.02	0.01	0.02	0.01
N	1,610	1,610	1,610	1,610
empirical p-value	0.05	0.10	0.13	0.55

Figures

Figure 3-1

The figure analyzes the statistical power of director fixed effects regressions with *Lossless Common board* benchmarking. Displayed are the results from the simulation analysis described in Section 4.2. The horizontal axis shows different sizes of the director effects used in the simulation (e). The vertical axis shows the rejection rate of H_0 . Each point represents the observed rejection rate at confidence level α for effect size e . The dashed lines mark the expected rejection rates if the null hypothesis is true, given confidence level α (i.e., α or type I error). Rows represent different dependent variables, whereas columns represent the typical confidence levels.



References

Adams, R. B. and D. Ferreira, 2008, Do directors perform for pay?, *Journal of Accounting and Economics* 46, 154-171.

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 C. 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.

Bacon, J. and J. K. Brown, 1975, *Corporate directorship practices: Role, selection and legal status of the board*, a Joint Research Report from the Conference Board and the American Society of Corporate Secretaries, Inc., New York, NY.

Baker, M. and P. A. Gompers, 2003, The determinants of board structure at the initial public offering, *Journal of Law and Economics* 46, 569-598.

- Baliga, B. R., R. C. Moyer, and R. S. Rao, 1996, CEO duality and firm performance: What's the fuss?, *Strategic Management Journal* 17, 41-53.
- Bebchuk, L. A. and A. Cohen, 2005, The costs of entrenched boards, *Journal of Financial Economics* 78, 409-433.
- Bebchuk, L. A. and M. S. Weisbach, 2010, The state of corporate governance research, *The Review of Financial Studies* 23, 939-961.
- Bennedsen, M., K. M. Nielsen, F. Perez-Gonzalez, and D. Wolfenzon, 2007, Inside the family firm: The role of families in succession decisions and performance, *Quarterly Journal of Economics* 122, 647-691.
- Bennedsen, M., F. Perez-Gonzalez, and D. Wolfenzon, 2011, Estimating the value of the boss: Evidence from CEO hospitalization events, Working Paper.
- 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.
- Booth, J. R. and D. N. Deli, 1996, Factors affecting the number of outside directorships held by CEOs, *Journal of Financial Economics* 40, 81-104.
- Booth, J. R. and D. N. Deli, 1999, On executives of financial institutions as outside directors, *Journal of Corporate Finance* 5, 227-250.
- 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.
- Bruce Johnson, W., R. P. Magee, N. J. Nagarajan, and H. A. Newman, 1985, An analysis of the stock price reaction to sudden executive deaths: Implications for the managerial labor market, *Journal of Accounting and Economics* 7, 151-174.
- 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.
- Carpenter, M., A. and J. D. Westphal, 2001, The strategic context of external network ties: Examining the impact of director appointments on board involvement in strategic decision making, *The Academy of Management Journal* 44, 639-660.
- Coles, J. L., N. D. Daniel, and L. Naveen, 2008, Boards: Does one size fit all, *Journal of Financial Economics* 87, 329-356.
- Cormen, T. H., C. E. Leiserson, R. L. Rivest, and C. Stein, 2001, *Introduction to algorithms*, MIT Press.
- Cronqvist, H. and R. Fahlenbrach, 2009, Large shareholders and corporate policies, *Review of Financial Studies* 22, 3941-3976.
- Custódio, C. and D. Metzger, 2014, Financial expert CEOs: CEO's work experience and firm's financial policies, *Journal of Financial Economics* 114, 125-154.
- Daily, C. M. and D. R. Dalton, 1993, Board of directors leadership and structure: Control and performance implications, *Entrepreneurship: Theory & Practice* 17, 65-81.
- DeAngelo, H. and L. DeAngelo, 1989, Proxy contests and the governance of publicly held corporations, *Journal of Financial Economics* 23, 29-59.
- 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.
- Denis, D. J. and D. K. Denis, 1995, Performance changes following top management dismissals, *The Journal of Finance* 50, 1029-1057.
- Diamond, D. W., 1984, Financial intermediation and delegated monitoring, *Review of Economic Studies* 51, 393-414.
- Donaldson, L. and J. H. Davis, 1991, Stewardship theory or agency theory: CEO governance and shareholder returns, *Australian Journal of Management* 16, 49-64.
- Eisenberg, T., S. Sundgren, and M. T. Wells, 1998, Larger board size and decreasing firm value in small firms, *Journal of Financial Economics* 48, 35-54.
- 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 1,000 firms, *Journal of Business* 78, 1943-1971.
- Fich, E. M. and A. Shivdasani, 2006, Are busy boards effective monitors?, *The Journal of Finance* 61, 689-724.
- Field, L., M. Lowry, and A. Mkrtchyan, 2013, Are busy boards detrimental?, *Journal of Financial Economics* 109, 63-82.
- Frank, M. Z. and V. K. Goyal, 2007, Corporate leverage: How much do managers really matter?
- Gabaix, X. and A. Landier, 2008, Why has CEO pay increased so much?, *Quarterly Journal of Economics* 123, 49-100.
- Giannetti, M., G. Liao, and X. Yu, forthcoming, The brain gain of corporate boards: Evidence from china, *Journal of Finance*.
- 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.
- Gompers, P., J. Ishii, and A. Metrick, 2003, Corporate governance and equity prices, *Quarterly Journal of Economics* 118, 107.
- Goyal, V. K. and C. W. Park, 2002, Board leadership structure and CEO turnover, *Journal of Corporate Finance* 8, 49-66.
- Graham, J. R., S. Li, and J. Qiu, 2012, Managerial attributes and executive compensation, *Review of Financial Studies* 25, 144-186.
- 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.

- Hambrick, D. C. and P. A. Mason, 1984, Upper echelons: The organization as a reflection of its top managers, *The Academy of Management Review* 9, 193-206.
- Harford, J., 2003, Takeover bids and target directors' incentives: The impact of a bid on directors' wealth and board seats, *Journal of Financial Economics* 69, 51-83.
- Hayes, R., H. Mehran, and S. Schaefer, 2004, Board committee structures, ownership, and firm performance, Working Paper.
- Hermalin, B. E. and M. S. Weisbach, 1998, Endogenously chosen boards of directors and their monitoring of the CEO, *American Economic Review* 88, 96-118.
- Hermalin, B. E. and M. S. Weisbach, 2003, Boards of directors as an endogenously determined institution: A survey of the economic literature, *Economic Policy Review* 9, 7-26.
- Holderness, C., 2009, The myth of diffuse ownership in the united states, *The Review of Financial Studies* 22, 1377-1408.
- Holmstrom, B., 1982, Moral hazard in teams, *The Bell Journal of Economics* 13, 324-340.
- Huson, M. R., P. H. Malatesta, and R. Parrino, 2004, Managerial succession and firm performance, *Journal of Financial Economics* 74, 237-275.
- Itoh, H., 1991, Incentives to help in multi-agent situations, *Econometrica* 59, 611-636.
- Jensen, M. C., 1993, The modern industrial revolution, exit, and the failure of internal control systems, *The Journal of Finance* 48.
- Jensen, M. C. and W. H. Meckling, 1976, Theory of the firm: Managerial behavior, agency costs and ownership structure, *Journal of Financial Economics* 3, 305-360.
- Kaplan, S. N., M. M. Klebanov, and M. Sorensen, forthcoming, Which CEO characteristics matter, *The Journal of Finance*.
- Kaplan, S. N. and D. Reishus, 1990, Outside directorships and corporate performance, *Journal of Financial Economics* 27, 389-410.
- Kini, O., W. Kracaw, and S. Mian, 1995, Corporate takeovers, firm performance, and board composition, *Journal of Corporate Finance* 1, 383-412.
- Klein, A., 1998, Firm performance and board committee structure, *Journal of Law and Economics* 41, 275-304.
- Kroszner, R. S. and P. E. Strahan, 2001, Bankers on boards: Monitoring, conflicts of interest, and lender liability, *Journal of Financial Economics* 62, 415-452.
- La Porta, R., F. Lopez-de-Silanes, A. Shleifer, and R. Vishny, 2002, Investor protection and corporate valuation, *The Journal of Finance* 57, 1147.

- Lipton, M. and J. W. Lorsch, 1992, A modest proposal for improved corporate governance, *Business Lawyer* 48, 59-77.
- 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.
- Lorsch, J. W. and E. MacIver, 1989, *Pawns or potentates: The reality of America's corporate boards*, Harvard Business School Press, Boston.
- MacAvoy, P. W. and I. M. Millstein, 1999, The active board of directors and its effect on the performance of the large publicly traded corporation, *Journal of Applied Corporate Finance* 11, 8-20.
- Mace, M., 1971, *Directors, myth, and reality*, Harvard Business School Press, Boston.
- 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.
- Mookherjee, D., 1984, Optimal incentive schemes with many agents, *Review of Economic Studies* 51, 433-446.
- 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.
- Palmon, O. and J. K. Wald, 2002, Are two heads better than one? The impact of changes in management structure on performance by firm size, *Journal of Corporate Finance* 8, 213-226.
- Pérez-González, F., 2006, Inherited control and firm performance, *American Economic Review* 96, 1559-1588.
- Perry, R. T., 2000, Incentive compensation for outside directors and CEO turnover, *Working Paper*, Arizona State University.
- Perry, T. and U. Peyer, 2005, Board seat accumulation by executives: A shareholder's perspective, *The Journal of Finance* 60, 2083-2123.
- Pi, L. and S. G. Timme, 1993, Corporate control and bank efficiency, *Journal of Banking & Finance* 17, 515-530.
- Rechner, P. L. and D. R. Dalton, 1991, CEO duality and organizational performance: A longitudinal analysis, *Strategic Management Journal* 12, 155-160.

- Rosenstein, S. and J. G. Wyatt, 1990, Outside directors, board independence, and shareholder wealth, *Journal of Financial Economics* 26, 175-191.
- Rosenstein, S. and J. G. Wyatt, 1994, Shareholder wealth effects when an officer of one corporation joins the board of directors of another, *Managerial and Decision Economics* 15, 317-327.
- Senn, M., 2014, Do individual directors matter? - Evidence from the S&P 1,500, Working Paper.
- Shaw, M. E., 1971, *Group dynamics: The psychology of small group behavior*, McGraw-Hill.
- Shivdasani, A. and D. Yermack, 1999, CEO involvement in the selection of new board members: An empirical analysis, *The Journal of Finance* 54, 1829.
- Shleifer, A. and R. W. Vishny, 1997, A survey of corporate governance, *The Journal of Finance* 52, 737-783.
- Vafeas, N., 1999a, Board meeting frequency and firm performance, *Journal of Financial Economics* 53, 113-142.
- Vafeas, N., 1999b, 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.
- Waelchli, U., 2008, The causes and consequences of CEO, COB, and board turnover, Working Paper, University of Bern.
- Waelchli, U. and J. Zeller, 2013, Old captains at the helm: Chairman age and firm performance, *Journal of Banking & Finance* 37, 1612-1628.
- Walkling, R. A. and M. S. Long, 1984, Agency theory, managerial welfare, and takeover bid resistance, *The RAND Journal of Economics* 15, 54-68.
- Weisbach, M. S., 1988, Outside directors and CEO turnover, *Journal of Financial Economics* 20, 431-460.
- 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.

Zingales, L., 2000, In search of new foundations, *The Journal of Finance* 55, 1623.

Selbständigkeitserklärung

Ich erkläre hiermit, dass ich diese Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen benutzt habe. Alle Koautorenschaften sowie alle Stellen, die wörtlich oder sinngemäss aus Quellen entnommen wurden, habe ich als solche gekennzeichnet. Mir ist bekannt, dass andernfalls der Senat gemäss Artikel 36 Absatz 1 Buchstabe o des Gesetzes vom 5. September 1996 über die Universität zum Entzug des aufgrund dieser Arbeit verliehenen Titels berechtigt ist.

Markus Senn

19. November 2014