

## COMMENTARY

# REAFFIRMING THE CEO EFFECT IS SIGNIFICANT AND MUCH LARGER THAN CHANCE: A COMMENT ON FITZA (2014)

TIMOTHY J. QUIGLEY\* and SCOTT D. GRAFFIN

Management Department, Terry College of Business, University of Georgia, Athens, Georgia, U.S.A.

*A recent study by Fitza argued that the prior estimates of the Chief Executive Officer (CEO) effect are conflated with events outside the CEO's control, are largely the result of random chance, and that the true CEO effect is smaller than has been previously estimated. We suggest that the empirical methodology employed by Fitza to support these claims substantially overstates the "random chance" element of the CEO effect. We replicate Fitza's findings, highlight methodological issues, offer alternative conclusions, and using multilevel modeling (MLM), suggest that his analyses mischaracterize the CEO effect. Copyright © 2016 John Wiley & Sons, Ltd.*

## INTRODUCTION

There is a long history of debate centered on how much impact Chief Executive Officers (CEOs) have over the outcomes of the organizations they lead (Chandler, 1962; Hambrick and Finkelstein, 1987; Hannan and Freeman, 1977; Mackey, 2008; Quigley and Hambrick, 2015; Salancik and Pfeffer, 1977). Starting with Lieberman and O'Connor (1972), a series of studies have sought to quantify the "CEO effect," frequently arguing that prior studies have over- or under-estimated the "true" effect while seeking to demonstrate that a new approach better captures the underlying phenomenon (Fitza, 2014; Hambrick and Quigley, 2014). Like all empirical studies, this body of work relies on association, not necessarily causality. Nevertheless, the plurality of prior work is suggestive of

a causal link with estimates of the CEO effect most frequently near 15 percent (for a review, see Hambrick and Quigley, 2014; Mackey, 2008; also see Hambrick and Quigley, 2014 for an approach that estimates a significantly larger CEO effect). Calling this consensus into question, however, Fitza (2014) argued that prior estimates of the CEO effect are conflated with events outside the CEO's control and largely the result of random chance. Fitza suggests that when this random chance is accounted for, "the effect CEOs have on company performance that can clearly be distinguished from the effect of chance is between 3.9 and 5.0 percent" (2014: 1847).

We believe that the empirical methodology Fitza employed, however, substantially overstates the "random chance" element of the CEO effect. By extension, Fitza's conclusions underestimate the true impact CEOs have on organizational performance. There are two reasons for this. First, Fitza based his conclusions on an incomplete assessment of the quality of the ANOVA models used in his analysis. Specifically, full consideration of model fit statistics, adjusted R-squared, in

Keywords: CEO effect; firm performance; variance decomposition; upper echelons; CEOs

\*Correspondence to: Timothy J. Quigley, Terry College of Business, University of Georgia, 409 Brooks Hall, Athens, GA 30602, U.S.A. E-mail: tquigley@uga.edu

particular, suggests his findings are not properly interpreted. Second, Fitza draws his conclusions using ANOVA, a method that has been shown to be problematic when data are nested (Bliese and Hanges, 2004; Crossland and Hambrick, 2011; Misangyi *et al.*, 2006). Using multilevel modeling (MLM), which is more appropriate in this context given the nested nature of the data (as we explain in detail below), we find substantially different results.

We address these issues by first replicating Fitza's findings using firm performance data from Compustat, which we refer to as "Compustat data," and with data where ROA is randomized (following Fitza's approach) across CEOs and years, which we refer to as "simulated data" (we explain these techniques in detail below). Our results suggest that evaluating the fit of Fitza's models using Compustat and simulated data provide for alternative conclusions that are more consistent with prior results of CEO effects studies. Next, we re-analyze the Compustat and simulated data using MLM. This method has been used in a number of studies seeking to understand the relative impact of factors that affect firm performance (Crossland and Hambrick, 2011; Hough, 2006; Misangyi *et al.*, 2006), and unlike ANOVA, which was the methodology employed in the Fitza (2014) study, it is designed to account for and remain robust to the challenges of data with an inherently nested structure (Klein and Kozlowski, 2000; Raudenbush and Bryk, 2002). Using this method, our results show that the CEO effect using Compustat data is approximately 22 percent, which is comparable to the results of the most recent CEO effects study using MLM (Quigley and Hambrick, 2015). We also use MLM to re-analyze the simulated data and find a CEO effect of 0.1 percent, which, we assert, is to be expected when employing a randomly generated dependent variable. Our study thus contributes to research on the CEO effect by reconciling findings regarding the impact CEOs have on their firms.

## BACKGROUND

Management scholars have long sought to understand the magnitude of impact leaders have on their organizations. Properly estimating the "CEO effect," or the proportion of variance in firm performance that can be attributed to CEOs, is foundational to the field of strategic management. Indeed, multiple domains within strategic management are

strongly influenced by the answer to this question. For instance, properly understanding the CEO effect has important implications for corporate governance research in terms of how CEOs are evaluated, rewarded, and retained or dismissed, and in the policy and regulation of how management are overseen. Further, research streams ranging from strategic change, diversification, exploration and exploitation, upper echelons, executive reputation, and so on could all be informed by a clear and definitive answer to this important question: How much do CEOs really matter?

For more than 40 years, research has employed various variance partitioning methodologies (VPM) to calculate the CEO effect. Conceptually, the CEO effect is estimated after isolating the effects of contextual factors, namely, yearly macro-economic trends ("year effect"), industry trends ("industry effect"), and firm trajectory ("firm effect"). A larger CEO effect arises when many individual CEOs deliver distinctive performance by deviating, positively or negatively, from the expectations driven by contextual factors. Figure 1 demonstrates this with three hypothetical CEOs. To simplify, we include a line at the mid-point of the y-axis representing "the context" or the performance we might expect from any given CEO based on conditions in a given calendar year and industry as well as firm history. All things being equal, this "context" line indicates we would expect an average CEO to deliver a performance of 6. CEO 1 delivers performance largely in line with that expectation. If a sample of CEOs contained numerous individuals that looked like CEO 1, we would observe a small or statistically insignificant CEO effect. In contrast to this, CEOs 2 and 3 delivered performances over their tenures that systematically deviated from expectations—one positively and one negatively. If our sample included many CEOs that looked like these two, we would observe a large CEO effect.

Empirically, the measurement of these effects is typically undertaken using large panel datasets with firm performance, most often return on assets (ROA), as the dependent variable and a series of dummy variables representing calendar years, industries, firms, and CEOs. The analysis generally begins with the entry of calendar year variables and continues with entry of industry, firm, and finally, CEO. At each level, the R-squared is calculated and any incremental gain at a given level is attributed to that factor. Crossland and Hambrick's (2007) results typify the prevailing findings in this stream. Using




Figure 1. How individual CEOs contribute to the CEO effect

ANOVA, their models attributed 3.6 percent of the variance in ROA to the year effects, 11.8 percent to industry effects, 19.1 percent to firm effects, and 13.4 percent to CEO effects, with the remaining 47.9 percent unexplained.

Calling these and related findings into question, Fitza argued that prior estimates of the CEO effect were artificially inflated as the result of random variance because “studies attribute all performance differences that coincide with different CEO tenures to the CEO (after a company’s industry, the year of measurement, and fixed company effects are controlled for)... [and in] doing so, these studies also attribute the effect of chance, of randomness to the CEO” (Fitza, 2014: 1840). His argument was based on a two-step analysis. First, actual firm performance data from Compustat, or the “Compustat data,” and ANOVA, he found a CEO effect in his sample of 17.7 percent, which is consistent with prior work. Next, he replaced the dependent variable with the “simulated data,” where return on assets (ROA) was replaced with a randomly generated number from a normal distribution having the same mean and standard deviation as the archival data from Compustat. He then repeated the analysis using the simulated data as the dependent variable. The underlying nested structure created by all other data (namely, the dummy variables representing year, industry, firm, and CEO) remained unchanged. Since the outcome variable in the simulated data was randomly generated, we would expect the resulting CEO effect to be zero. Surprisingly, however, with simulated data as the dependent variable, the measured CEO effect averaged 13.3 percent across 100 simulated trials (90% confidence interval of 12.8–13.8%).

Drawing on these results, Fitza argued that “13.3 percent represent[s] the statistical artifact by which the CEO effect is inflated by the effect of

randomness” (2014: 1845). He further notes that “any CEO effect that is below 13.3 percent cannot be distinguished from the effect of random chance” (1845) and “given the confidence interval, to be statistically significantly different (with a  $p < 0.05$ ) from randomness, a measured CEO effect needs to be larger than 13.8 percent [the upper end of the 90% confidence interval]” (1845). Subtracting this baseline (13.8%) from the initial result of 17.7 percent yields a CEO effect “that is statistically significant over the effect of randomness” of 3.9 percent. Fitza concluded that “past variance decomposition studies on the so-called CEO effect result in inflated CEO effects” (1849).

## METHODS

### Replication of Fitza and an alternative analysis

We believe that the approach taken by Fitza overstates the “random chance” element of the CEO effect primarily by misinterpreting the ANOVA results and ignoring key diagnostics from models using simulated data. To explore these issues, we first construct a sample with actual data from Compustat and Execucomp and recreate Fitza’s findings using Compustat and simulated data. Using these results, we highlight problems with Fitza’s interpretations. We then apply a different, and we argue, more appropriate, methodological approach before presenting our alternative findings and conclusions.

### Sample

Our sample was created following the process outlined by Fitza (2014). Namely, we began with all CEOs in the Execucomp database from 1993 to 2012. We removed financial institutions (SIC codes beginning with 6) and government and unclassified industries (SIC codes beginning with 9); removed firms having only one CEO over the entire sample; removed firms with less than \$20 million in assets; and removed firm-years where ROA was above the 99th or below the first percentile of the sample. Different from Fitza, however, we also removed CEOs who served just one year as their effects are perfectly predicted by their CEO dummy, and as a result, artificially inflate the CEO effect (below we highlight the impact this has on the Compustat and simulated samples). Like Fitza, our dependent variable, ROA, was calculated after eliminating extraordinary items. Further, we created 100 randomized

ROA variables selected from a normal distribution using the same mean and standard deviation of the Compustat sample (mean of 4.10, standard deviation of 13.72).<sup>1</sup> We then generated dummy variables representing each of the calendar years, industries, firms, and CEOs in the sample.

## RESULTS

### Replicating Fitza using ANOVA

Table 1 provides descriptive statistics for the samples as well as comparisons to samples used by Fitza. The number of firm-years, industries, and firms as well as the mean and standard deviation for ROA are all comparable to Fitza (2014).<sup>2</sup> Detailed results of our replication ANOVA models are provided in Table 2, while the results of our MLM analyses are offered in Table 3. For ease of comparison across models, Table 4 provides a summary of Fitza's reported results, our replication of those findings, and our additional analyses.

Columns 1 and 2 in Table 4 compare Fitza's ANOVA results with our replication using Compustat data. The reported effect sizes are similar. Specifically, as shown in Column 1, Fitza reported a CEO effect of 17.7 percent, and we report a CEO effect of 18.0 percent in Column 2. Columns 3 and 4 compare average results from 100 trials that employed the simulated data for the dependent variable (ROA). Here, our reported CEO effect is somewhat smaller than Fitza's. Specifically, Fitza reported a mean CEO effect across 100 random samples of 13.3 percent, while our mean effect across 100 trials was 11.0 percent—a difference of 2.3 percentage points. As noted earlier, our sample differed from Fitza's in one important

way—we removed CEOs that served just a single year. When we added these CEOs back into the sample, we found a CEO effect in the simulated data of approximately 13.0 percent, nearly identical to the 13.3 percent shown by Fitza. By comparison, in the Compustat data, including single-year CEOs increased the CEO effect to 19.5 percent (suggesting an approximately 2% increase as a result of including single-year CEOs in the sample). In short, while our samples are slightly different, our results are comparable to Fitza's (2014) in both the Compustat and simulated data.<sup>3</sup>

With comparable results to Fitza's using ANOVA, we now turn our attention to the various model specification tests. We first consider model significance. As shown at the bottom of Column 1 in Table 2, the ANOVA using Compustat data is statistically significant ( $p < 0.000$ ), while, as shown in Column 4, the average of the results from 100 trials using simulated data is not statistically significant (mean  $p = 0.483$ ). Though not reported, the same can be said for each set of categorical variables in the model—in all cases each variable is statistically significant in the Compustat data, but is not statistically significant using the simulated data. This lack of statistical significance for each category using the simulated data means the variance explained by them is unlikely to differ from zero.

We next considered the variance explained by each model. Past CEO effects studies typically consider only the incremental R-squared as each level of predictor is entered into the model (this value can alternatively be obtained by dividing the sum-of-squares for a particular level by the total sum-of-squares). However, related variance decomposition studies that use ANOVA to consider how corporate parents, business units, and industry membership impact performance have routinely considered adjusted R-squared as well (e.g., Hough, 2006; Rumelt, 1991). As reported in Column 2 of Table 4, using Compustat data, the R-squared, or proportion of total variance explained by the model is 55.6 percent, while the adjusted R-squared is 46.0 percent, a decline of 9.6 percentage points. Column 4 shows the model R-squared for the simulated data is 18.0 percent,

<sup>1</sup> We repeated our analysis with simulated data using the exact mean and standard deviation reported by Fitza and the results were the same.

<sup>2</sup> We are alert to the fact that our count of CEOs is nearly 1,000 more than reported by Fitza (2014), and that it would be higher still had we included single-year CEOs as Fitza did. After extensive evaluation, we believe Fitza's reported number is likely a typographical error. First, it seems infeasible that we replicated all other aspects of the sample within  $\pm 3$  percent, yet have such a large difference (37%) in CEO count. Second, as noted in the description of the sample formation process, firms with just a single CEO over the entire 20-year sample must be dropped as their effect would be indistinguishable from the firm effect. By extension, this means the CEO count must be, at minimum, double the number of firms. As shown in Table 1, Fitza reports just 1.85 CEOs per firm, which is why we attribute this reported CEO count as a typographical error.

<sup>3</sup> As a robustness check, we repeated all of our reported analyses using a sample that included single-year CEOs. While reported CEO effects were always approximately two percentage points higher, the conclusions drawn remained substantively unchanged.



Table 1. Descriptive statistics—Compustat versus simulated data

	1 Fitza	2 This study
<b>Total firm-years</b>	19,746	20,348
<b>Industries</b>	220	213
<b>Firms</b>	1,425	1,399
<b>CEOs</b>	2,634	3,627
<b>CEO per firm</b>	1.85	2.59
<b>Mean ROA</b>		
Compustat data	3.86 <sup>a</sup>	4.10
Compustat data after dropping at 1/99th percentile		4.63
Simulated data <sup>b</sup>		4.10
Simulated data after dropping at 1/99th percentile <sup>b</sup>		4.09
<b>ROA—standard deviation</b>		
Compustat data	13.26 <sup>a</sup>	13.72
Compustat data after dropping at 1st and 99th percentile		7.64
Simulated data <sup>b</sup>		13.72
Simulated data after dropping at 1/99th percentile <sup>b</sup>		12.83

<sup>a</sup> Mean and SD taken from footnote 9 of Fitza (2014:1845)

<sup>b</sup> Simulated data descriptive statistics reflect average from 100 randomly created samples.

Table 2. ANOVA replication results

Factor	1	2 Compustat data		3	4	5 Simulated data <sup>a</sup>		6
	Sum of squares	Incremental R-squared	Incremental adjusted R-squared	Sum of squares	Incremental R-squared	Incremental adjusted R-squared	Sum of squares	Incremental R-squared
Year	29,649.7	2.5%	2.4%	3,266.8	0.1%	0.0%		
Industry	111,019.6	9.3%	8.4%	34,975.8	1.0%	0.0%		
Firm	306,866.3	25.8%	22.2%	195,026.3	5.8%	0.0%		
CEO	213,293.9	18.0%	13.0%	368,583.5	11.0%	0.0%		
Residual	526,651.9			2,747,956.8				
Total	1,187,481.4	55.6%	46.0%	3,349,396.0	18.0%	0.0%		
Model statistics								
<i>F</i> -statistic	5.750			1.001				
<i>P</i> -value	0.000			0.483				

<sup>a</sup> Average of 100 trials.

while the adjusted R-squared is 0.0 percent. This difference is worth noting because R-Squared values increase, or at least, cannot decrease, with each additional predictor, even if model fit is not improved. Adjusted R-squared, however, corrects for this. Adjusted R-squared only increases when new terms added improve the model's predictive power beyond chance and beyond the cost, in degrees of freedom, of the additional terms (Kennedy, 2003). Thus, when we examine the incremental adjusted R-squared (Table 2) at the CEO-level using the Compustat data, we find a statistically significant impact of 13.0 percent (see Column 3, Table 2). This stands in contrast to

the incremental adjusted R-squared of 0.0 percent using the simulated data (see Column 6, Table 2).

### Alternative approach: multilevel modeling

While ANOVA has been the most frequently used method for assessing the CEO effect (Crossland and Hambrick, 2007; Lieberman and O'Connor, 1972; Mackey, 2008), its use has been called into question. First, with ANOVA, the ordering of variable entry can have a significant impact on the reported effect sizes (Bowman and Helfat, 2001; Brush and Bromiley, 1997; Weiner, 1978). For example, entering the CEO variables into our ANOVA models first

Table 3. Multilevel model results

Factor	1		2		3		4		5		6	
	Variance estimates		Variance estimates		Variance estimates		Variance estimates		Variance explained		Variance explained	
	Unconditional models		Models incorporating year effects		Models incorporating year effects		Models incorporating year effects		Variance explained		Variance explained	
	Compustat	Simulated <sup>a</sup>	Compustat	Simulated <sup>a</sup>	Compustat	Simulated <sup>a</sup>	Compustat	Simulated <sup>a</sup>	Compustat	Simulated <sup>a</sup>	Compustat	Simulated <sup>a</sup>
<b>Year</b>									1.6%		0.1%	
<b>Industry</b>	2.759	0.034	2.867	0.034	2.867	0.034	2.867	0.034	4.6%		0.0%	
<b>Firm</b>	11.681	0.126	13.258	0.129	13.258	0.129	13.258	0.129	21.1%		0.1%	
<b>CEO</b>	14.817	0.177	13.720	0.179	13.720	0.179	13.720	0.179	21.8%		0.1%	
<b>Residual</b>	33.575	187.972	31.975	187.779	31.975	187.779	31.975	187.779	50.9%		99.7%	
<b>Total</b>	62.832	188.308										
<b>Model Statistics</b>												
<b>Chi-square</b>	5,741.453	0.987	6,128.895	0.944	6,128.895	0.944	6,128.895	0.944				
<b>P-value</b>	0.000	0.821	0.000	0.827	0.000	0.827	0.000	0.827				

<sup>a</sup> Average of 100 trials.

Table 4. Results comparison

Factor	1	2	3	4	5	6
	Fitza - Compustat ANOVA	This study - Compustat ANOVA	Fitza - Simulated ANOVA, <sup>ab</sup>	This Study - Simulated ANOVA <sup>b</sup>	This Study - Compustat MLM	This Study - Simulated MLM <sup>b</sup>
<b>Year effect</b>	2.0%	2.5%	0.1%	0.1%	1.6%	0.1%
<b>Industry effect</b>	7.3%	9.3%	1.2%	1.0%	4.6%	0.0%
<b>Firm effect</b>	33.4%	25.8%	7.3%	5.8%	21.1%	0.1%
<b>CEO effect</b>	17.7%	18.0%	13.8/13.3% <sup>a</sup>	11.0%	21.8%	0.1%
<b>Unexplained</b>	39.6%	44.4%	77.6%	82.0%	50.9%	99.7%
<b>Model Statistics</b>						
<b>R-squared/variance explained</b>	60.4%	55.6%	22.4%	18.0%	49.1%	0.3%
<b>Adjusted R-squared (ANOVA)</b>	Not reported	46.0%	Not reported	0.0%	n/a	n/a
<b>Model p-value</b>	Not reported	0.000	Not reported	0.483	0.000	0.827

<sup>a</sup> Fitza reports the upper range of the 90% confidence interval from the 100 random trials for each level; his mean CEO effect as 13.3%. While our discussion is based on the mean CEO effect, using the upper limit has no substantive impact on the results.

<sup>b</sup> Average of 100 trials.

yields a CEO effect of 52.8 percent in the Compustat data and approximately 18.0 percent in the simulated data (though adjusted R-squares were 43.7 and 0.0%, respectively). More important, a key assumption of ANOVA is that various categorical data (and their associated error terms) are independent (Bliese and Hanges, 2004). In our data, however, roughly 14 percent of the variance in ROA is shared across the four levels of predictors. To address this, recent CEO effects studies (Crossland and Hambrick, 2011; Hambrick and Quigley, 2014; Quigley and Hambrick, 2015) and related studies assessing the importance of business segments and corporate parents versus industry effects (Hough, 2006; Misangyi *et al.*, 2006) have adopted MLM.

As a brief summary of the method, calculating CEO effects with MLM entails estimating two models (we used xtmixed in STATA 12.1). The first, referred to as the unconditional model (Hough, 2006; Misangyi *et al.*, 2006; Raudenbush and Bryk, 2002), is a three-level model with yearly CEO performances nested within firms, and firms nested within industries. A second model then incorporates calendar year effects at the lowest level. The variance estimates from the second model are used to calculate the relative effect sizes by dividing the variance component for industry, firm, and CEO by the total from the unconditional model. The remaining difference, caused by the addition of calendar years to the second model,

represents the year effect (for a detailed description of the use of MLM for this purpose, see Hough, 2006; Misangyi *et al.*, 2006).

Details from these models are provided in Table 3, while Columns 5 and 6 of Table 4 summarize these results and provide a comparison to the various ANOVA models reported earlier. Column 5 in Table 4 shows that with MLM the CEO effect was 21.8 percent in the Compustat data (models were statistically significant with  $p < 0.000$ ), which is in line with the most recent studies using MLM (Hambrick and Quigley, 2014; Quigley and Hambrick, 2015). In contrast, when MLM was applied to the simulated data, as shown in Column 6, the mean CEO effect across 100 trials was 0.1 percent and the models were not significant (mean  $p = 0.827$ ). In short, MLM provides results consistent with the expectations we have for each dataset: with Compustat data, we see a substantial CEO effect that is similar to that found in prior studies and somewhat larger than reported with ANOVA models; with a simulated (randomized) dependent variable, where we would not expect to find a CEO effect, we, in fact, find an effect size of essentially zero (0.1%).

## DISCUSSION AND CONCLUSION

We aimed to achieve one objective in our study—to accurately estimate the CEO effect, which refers to the impact CEOs have on organizational outcomes. As we noted in our introduction, this is one of the most fundamental research questions in strategic management scholarship. Indeed, properly estimating the CEO effect has implications for researchers who study corporate governance, the resourced-based view of organizations, dynamic capabilities, mergers and acquisitions as well as numerous other topics. Properly understanding the CEO effect also has implications for practitioners in terms of properly evaluating and compensating executives. Using MLM, we find a CEO effect of 21.8 percent.

Our findings are in contrast to a recent study that estimated a much smaller CEO effect. Specifically, Fitza (2014) asserted that the properly estimated CEO effect, “is between 3.9 and 5.0 percent” (1847). We demonstrated that these findings overlooked the overall poor fit of the ANOVA models used to arrive at these conclusions

and also that using a more appropriate statistical technique, MLM, provided quite different results.

Our study highlights the importance of interpreting adjusted R-squared, which has two implications. First, this shows that Fitza’s interpretation of results using randomized simulated data was not appropriate. Second, it also means that prior CEO effects research using ANOVA overstated the reported CEO effect as well (and not just that for CEOs, but for year, industry, and firm as well). Nevertheless, these prior studies used actual firm data and found effects that were presumably significant as demonstrated by our replication. Beyond that, our study also brings attention to the emergence of MLM as an analytical technique that can more appropriately model the CEO effect in the nested data structures inherent to this phenomenon.

While our findings stand in contrast to Fitza’s, it is important to acknowledge and discuss some of the compelling issues raised by his efforts. Most notably, Fitza argued the CEO effect is inflated by two factors—a “statistical artifact” caused by the statistical method and also random chance driven by the fact that “these studies attribute all performance differences that coincide with different CEO tenures to the CEO” (Fitza, 2014: 1840). While we believe the use of MLM largely rules out inflation caused by statistical artifact, we have not addressed his second claim that the CEO effect can be inflated by events that, through happenstance, exactly correspond to a given CEO’s tenure. As Fitza rightly pointed out, “random events that coincide with the tenure of one CEO and that cannot be attributed to the industry or the general (annual) economic conditions will be attributed to that CEO” (Fitza, 2014: 1842) and would inflate the CEO effect. For example, we can imagine a new CEO benefitting from a new product developed under a former CEO but launched under the tenure of a new CEO. Or, like Mary Barra, who inherited the ignition recall at GM, a CEO might be burdened by lawsuits related to decisions made years ago, nullifying what otherwise may have been distinctly superior performance. Referring back to Figure 1, these types of “(un)lucky” events might make CEO 1 look, instead, like CEO 2 or 3 purely by (bad) luck rather than any meaningful impact by the current CEO. While this is certainly a concern, the exact opposite phenomenon might also exist. That is, Fitza’s arguments seem to presume that there are many potential random factors that artificially create systematic variance that results in distinctively performing CEOs (that is, events

that drive up the CEO effect by turning CEO 1 in Figure 1 into CEO 2 or CEO 3). However, it seems just as likely that some CEOs might, left to their own devices, deliver truly distinctive performance (like CEO 2 or 3), but have their “true” impact muted by these same random occurrences. That is, the same random process that might boost the CEO effect could also reduce it by making CEOs 2 or 3 from Figure 1 look like CEO 1 instead.

In short, we agree that Fitza’s efforts to highlight the impact of stochastic processes are important. But, unless these forces disproportionately enhance rather than inhibit the likelihood of distinct performances over a CEO’s tenure, it seems most likely that, over a large number of CEOs, these random effects would cancel out each other. As this remains a challenging issue for numerous research streams, it remains a ripe area for future theoretical and empirical work.

Fitza also argued that variance decomposition works best when we have many observations per entity. While MLM enhances the robustness of variance decomposition, relatively short CEO tenures dictates this will remain a challenge for studies that employ this methodology. Finally, echoing Lieberson and O’Connor’s (1972) original work in this area, Fitza also reminds us of the potential differences found when incorporating a lag structure into the analysis. Similar to Fitza, our own analysis (not reported) shows a slight increase in the measured CEO effect (using MLM) when performance is lagged one year. The CEO effect declines substantially for two- and three-year lags. Regardless of the results, additional theoretical and empirical efforts are needed to better disentangle if and how much of the initial period of a CEO’s tenure should be credited to former versus current CEO.

Since Lieberson and O’Connor (1972) first studied the CEO effect more than 40 years ago, scholars from various perspectives have fiercely debated the relative importance of leaders. While those from the strategic management and choice perspective have highlighted the importance of CEOs and the impact they can have through planning and decision making (Child, 1972, 1997; Hambrick and Mason, 1984; Miles and Snow, 1978), others have highlighted their limited potential for impact as a result of institutional pressures, resource constraints, and inertia (DiMaggio and Powell, 1983; Hannan and Freeman, 1977; Haveman, 1993). In many ways, this argument has evolved to a common understanding that the potential

for impact is contingent on a variety of factors. For example, Hambrick and Finkelstein (1987) introduced the conception of managerial discretion highlighting that a CEO’s latitude of action is a function of contextual conditions in the external environment and firm as well as the psychological traits of the leader. Empirical research has shown that consideration of managerial discretion helps explain, among other things, the impact of CEOs across national settings (Crossland and Hambrick, 2007, 2011), the magnitude of CEO pay (Finkelstein and Boyd, 1998), and the amount of strategic change undertaken by a firm (Peteraf and Reed, 2007; Quigley and Hambrick, 2012).

While progress has been made, empirical evidence incorrectly concluding that the CEO effect is trivial could stifle consideration of the important ways leaders impact their organization’s outcomes. With our findings here, as well as a preponderance of prior research evidence pointing to a nontrivial (Mackey, 2008; Wasserman, Anand, and Nohria, 2010), growing (Quigley and Hambrick, 2015), and perhaps even more substantial (Hambrick and Quigley, 2014), CEO effect that varies across national settings (Crossland and Hambrick, 2007, 2011), we believe it is time to move on from the debate regarding the *existence* of a CEO effect, and instead, focus on the factors that drive *differences* in the effects CEOs have on their firms. To that end, we agree with Blettner, Chaddad, and Bettis who highlighted that “the CEO performance effect is determined in aggregate by a complex set of interdependencies” (2012: 994). They further point to a number of empirical and theoretical challenges relating to the “fit” among these factors that must be solved in order to garner a more thorough understanding of when and how leaders impact their organizations. With our findings affirming the existence of a sizable CEO effect, we hope greater attention will now focus on these broader, but more complex, issues regarding CEO and firm performance.

## ACKNOWLEDGEMENTS

We sincerely thank Vilmos Misangyi for providing helpful comments on earlier versions of this article and Bob Vandenberg for helping us sharpen our methodological arguments. We would also like acknowledge the gracious assistance of Markus Fitza, who answered questions and provided his SAS code during the final development of our



article. Finally, we greatly appreciate *SMJ* Editor Richard Bettis's efforts shepherding this article through the review process.

## REFERENCES

- Blettner D, Chaddad F, Bettis R. 2012. The CEO performance effect: challenges of empirical modeling and a complex fit perspective. *Strategic Management Journal* **33**(8): 986–999.
- Bliese PD, Hanges PJ. 2004. Being both too liberal and too conservative: the perils of treating grouped data as though they were independent. *Organizational Research Methods* **7**(4): 400–417.
- Bowman EH, Helfat CE. 2001. Does corporate strategy matter? *Strategic Management Journal* **22**(1): 1–23.
- Brush TH, Bromiley P. 1997. What does a small corporate effect mean? A variance components simulation of corporate and business effects. *Strategic Management Journal* **18**(10): 825–835.
- Chandler AD. 1962. *Strategy and Structure: Chapters in the History of the Industrial Enterprise*. MIT Press: Cambridge, MA.
- Child J. 1972. Organizational structure, environment and performance: the role of strategic choice. *Sociology* **6**(1): 1–22.
- Child J. 1997. Strategic choice in the analysis of action, structure, organizations and environment: retrospect and prospect. *Organization Studies* **18**(1): 43.
- Crossland C, Hambrick DC. 2007. How national systems differ in their constraints on corporate executives: a study of CEO effects in three countries. *Strategic Management Journal* **28**(8): 767–789.
- Crossland C, Hambrick DC. 2011. Differences in managerial discretion across countries: how nation-level institutions affect the degree to which CEOs matter. *Strategic Management Journal* **32**(8): 797–819.
- DiMaggio PJ, Powell WW. 1983. The iron cage revisited: institutional isomorphism and collective rationality in organizational fields. *American Sociological Review* **48**(2): 147–160.
- Finkelstein S, Boyd BK. 1998. How much does the CEO matter? The role of managerial discretion in the setting of CEO compensation. *Academy of Management Journal* **41**(2): 179.
- Fitza MA. 2014. The use of variance decomposition in the investigation of CEO effects: how large must the CEO effect be to rule out chance? *Strategic Management Journal* **35**(12): 1839–1852.
- Hambrick DC, Finkelstein S. 1987. Managerial discretion: a bridge between polar views of organizational studies. *Research in Organizational Behavior* **9**: 369–406.
- Hambrick DC, Mason PA. 1984. Upper echelons: the organization as a reflection of its top managers. *Academy of Management Review* **9**(2): 193–206.
- Hambrick DC, Quigley TJ. 2014. Toward more accurate contextualization of the CEO effect on firm performance. *Strategic Management Journal* **35**(4): 473–491.
- Hannan MT, Freeman J. 1977. The population ecology of organizations. *American Journal of Sociology* **82**(5): 929.
- Haveman HA. 1993. Ghosts of managers past: managerial succession and organizational mortality. *Academy of Management Journal* **36**(4): 864–881.
- Hough JR. 2006. Business segment performance redux: a multilevel approach. *Strategic Management Journal* **27**(1): 45–61.
- Kennedy P. 2003. *A Guide to Econometrics*. MIT Press: Cambridge, MA.
- Klein KJ, Kozlowski SWJ. 2000. From micro to meso: critical steps in conceptualizing and conducting multilevel research. *Organizational Research Methods* **3**(3): 211.
- Lieberson S, O'Connor JF. 1972. Leadership and organizational performance: a study of large corporations. *American Sociological Review* **37**(2): 117–130.
- Mackey A. 2008. The effect of CEOs on firm performance. *Strategic Management Journal* **29**(12): 1357–1367.
- Miles RE, Snow CC. 1978. *Organizational Strategy, Structure, and Process*. McGraw-Hill: New York.
- Misangyi VF, Elms H, Greckhamer T, Lepine JA. 2006. A new perspective on a fundamental debate: a multilevel approach to industry, corporate, and business unit effects. *Strategic Management Journal* **27**(6): 571–590.
- Peteraf M, Reed R. 2007. Managerial discretion and internal alignment under regulatory constraints and change. *Strategic Management Journal* **28**(11): 1089–1112.
- Quigley TJ, Hambrick DC. 2012. When the former CEO stays on as board chair: effects on successor discretion, strategic change, and performance. *Strategic Management Journal* **33**(7): 834–859.
- Quigley TJ, Hambrick DC. 2015. Has the “CEO effect” increased in recent decades? A new explanation for the great rise in America's attention to corporate leaders. *Strategic Management Journal* **36**(6): 821–830.
- Raudenbush S, Bryk A. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. Sage Publications: Thousand Oaks, CA.
- Rumelt R. 1991. How much does industry matter? *Strategic Management Journal* **12**: 167–185.
- Salancik GR, Pfeffer J. 1977. Constraints on administrator discretion: the limited influence of mayors on city budgets. *Urban Affairs Review* **12**(4): 475–498.
- Wasserman N, Anand B, Nohria N. 2010. When does leadership matter? In *Handbook of Leadership Theory and Practice*, Nohria N, Khurana R (eds). Harvard Business School Publishing Corporation: Boston, MA; 27–64.
- Weiner N. 1978. Situational and leadership influences on organization performance. In *Annual Meeting of the Academy of Management Proceedings*; 230–234; doi:10.5465/AMBPP.1978.4976895.