Obscured transparency? Compensation benchmarking and the biasing of executive pay

Mathijs de Vaan
Haas School of Business, UC Berkeley, Berkeley, CA, USA, mdevaan@haas.berkeley.edu

Benjamin Elbers, Thomas A. DiPrete
Department of Sociology, Columbia University, New York, NY, USA, tad61@columbia.edu be2239@columbia.edu

The disclosure of compensation peer groups is argued to provide shareholders with valuable information that can be used to scrutinize CEO compensation. However, research suggests that there are substantial incentives for executives and directors to bias the compensation peer group upward such that the CEO can extract additional rent. We leverage the idea that reciprocated peer nominations are unlikely to be biased in order to construct counterfactual peer groups that allow us to measure the bias of disclosed peer groups. Analyses of eleven years of comprehensive data on compensation peer groups demonstrate that the average firm uses an upwardly biased peer group. The size of the bias increases when incentives and opportunities to do so are more pronounced. Specifically, results show that bias is larger when financial targets are not met and when exercising discretion in the selection of peer firms is justifiable. More importantly, upward bias in compensation peer groups is highly predictive of an increase in CEO compensation – suggesting that there is a strong incentive for CEOs to strategically select peers. Finally, while average peer group bias has gone down in recent years, the predictive effect of bias on pay has gone up. These findings call into question the practical impact of recent efforts to introduce greater transparency into the process for determining executive compensation.

Key words: Executive compensation; Social comparisons; Peer Benchmarking; Income Inequality

---

1. Introduction

It is well known that rising income inequality in the U.S. has, in recent decades, been driven by changes in the top of the income distribution and particularly in the pay of "supermanagers" (Piketty 2014) who dominate the top 0.1% of this distribution (Bakija, Cole, and Heim 2010, Keister 2014). Average Chief Executive Officer (CEO) compensation increased from $625,000 to $9,840,000 between 1980 and 2004 – an average growth rate of 12% per year. More recently, while median market capitalization for the constituents of the S&P 1500 grew by 22% between 2007 and 2014, median CEO compensation increased by 39% in that same period.

A potential explanation for this trend draws from the way in which companies benchmark the compensation of their CEOs (DiPrete, Eirich, and Pittinsky 2010, Faulkender and Yang 2013, Albuquerque, De Franco, and Verdi 2013, Pittinsky and DiPrete 2013). Compensation benchmarking involves the assembly of a group of peer firms that serves as a lens through which CEO compensation can be evaluated (Barsalou 1983, Podolny 2001, Smith and Chae 2017). The Securities and Exchange Commission (SEC) made disclosure of these peer groups mandatory in 2007. Although the rationale for using compensation peer groups is that a comparison between the focal CEO and a set of CEOs at *similar* firms should enable boards of directors and shareholders to accurately evaluate compensation, DiPrete et al. (2010) note that the inclusion of well compensated CEOs may upwardly bias the peer group and may allow the focal CEO to command more pay because the peer group makes the pay seem appropriate and defensible. The business media and some academic research (Faulkender and Yang 2010, Pittinsky and DiPrete 2013) suggest that companies indeed have a tendency to select larger firms and better paid CEOs into their compensation peer groups.[1] However, recent scholarship has questioned the validity of these claims (Bizjak et al. 2008, Albuquerque et al. 2013, Cadman and Carter 2013, Kim et al. 2015). For example, Albuquerque et al. (2013) compare reported peer groups to propensity score matched peer groups and conclude that although CEOs in reported peer groups are more generously compensated, the gap is accounted for by variation in CEO talent. Cadman and Carter (2013) also reject the claim that peer groups are upwardly biased, but they argue that previous research on compensation benchmarking draws from flawed samples of potential peers. The implications of these studies are important. They suggest that, if anything, the reporting of compensation peer groups facilitates the efficient allocation of compensation in the CEO labor market.

These conflicting conclusions suggest that compensation benchmarking and its effects on pay remain poorly understood. The aim of this study is to shed light on the process by which peer firms are selected and by scrutinizing its relationship with CEO pay. To that end, we develop a novel approach that allows us to construct counterfactual peer groups based on a process that should be unbiased. Specifically, we leverage the idea that reciprocated peer selections (i.e. when firm i nominates firm j and firm j nominates firm i) are unlikely to be biased. We estimate a model for reciprocated nominations of compensation peers and use the estimates from this model to simulate samples of counterfactual peer groups based on an author-constructed dataset of more than 3,400 firms that reported their compensation peer groups to the SEC between 2006 and 2016.

Using those data and methods, we find substantial evidence that companies benchmark their CEOs to a biased sample of well compensated CEOs. We then examine the conditions under which peer group bias should be most pronounced. We draw from research on social comparisons (Audia et al. 2016, Smith and Chae 2017) and predict that peer group bias should increase when the firm fails to meet its financial targets and when firms can exercise benchmarking discretion as a

[1] See “A Better Way to Compare C.E.O. Pay” and “Cozy relationships and “peer benchmarking” send CEOs’ pay soaring” for examples.

result of ambiguity about the set of firms that are most similar to the focal firm, i.e. its "natural" or "counterfactual" peers.[2] Since CEO compensation is typically coupled to several performance targets including growth in market capitalization and profits (Frydman and Jenter 2010), CEOs of firms that fail to meet its performance targets should have a strong *incentive* to bias its peer group upward. Doing so could allow the CEO to keep receiving generous compensation packages without drawing excessive scrutiny (Bertrand and Mullainathan 2001). Our results are consistent with this prediction. We also demonstrate that the level of bias in a firm's peer group is positively associated with the amount of discretion that firms can exercise in the selection of compensation peers. Specifically, when the compensation of CEOs employed by the natural peers of a firm varies widely, firms have the *opportunity* to cherry pick peers with well compensated CEOs without compromising the integrity of the peer selection process in obvious ways.

Our fourth and perhaps most important result demonstrates that bias in compensation peer groups has important consequences: it is associated with increased compensation for the focal CEO. These results are robust to the inclusion of a variety of predictors of executive compensation and the use of different methods for constructing counterfactual peer groups.

Finally, we explore whether the practice of peer group benchmarking has changed over time. We demonstrate that the greater environmental pressure for market-based – i.e. unbiased – peer groups appears to increase conformity with established norms for how compensation peer groups should be selected: both the variance of bias across firms and mean levels of bias are decreasing over time. However, bias remains large and positive even in the later years of our study period. Furthermore, we find that the predictive effect of bias on compensation is actually *increasing* over time, counter-acting the decrease in overall bias levels almost completely. In sum, these findings strongly suggest that peer groups are designed more to reduce environmental opposition to high levels of executive compensation than to provide a market signal to determine the market-appropriate level of compensation. Ten years after the disclosure requirements of the SEC came into effect, the distorting influence of peer group bias in the labor market for CEOs remains substantial.

2. The principles of compensation benchmarking

2.1. Bias in compensation peer groups

Benchmarking CEO compensation has become common practice. Early scholarly debates addressed the question of whether "internal" or "external" labor markets should be used to determine and evaluate compensation (Roberts 1956, Simon 1957, DiPrete et al. 2010). A key feature in this debate is the idea that a set of executives employed by similar firms could serve as a lens through

[2] We will use the terms "natural peer group" and "counterfactual peer group" interchangeably to refer to impartially selected peer groups.

which the focal executive's compensation can be evaluated. Specifically, comparing the focal CEO with a set of CEOs at *similar* organizations is argued to allow for an accurate evaluation because these CEOs are navigating similar market conditions. Both the potential to increase evaluation accuracy and the increased labor mobility of executives in the last decades of the 20^{th} century, have made benchmarking against "a list of companies comparable to [the focal] organization" a dominant practice for setting and evaluating executive compensation (Cook 1981, p.39).

Arguments in support of the introduction of mandatory peer group disclosure build on the assumption that peer groups would be formed by selecting *similar* firms into the peer group. An important incentive to assemble an impartial compensation peer group is the potential for push-back by agents of corporate governance (Faulkender and Yang 2010). For example, if a proposed compensation peer group includes CEOs who are compensated at much higher levels than the CEOs of companies similar to the focal firm, the compensation committee or shareholders – with support from watchdog organizations – can take action to adjust the peer group and correct the bias (Zuckerman 2012, ISS 2015). As a result, compensation peer groups should provide information to evaluate whether there is a gap between the price (i.e., compensation) and the value (i.e., what the CEO's contribution is worth) of a CEO.

There is, however, a substantial body of research showing that corporate governance varies between firms and sometimes fails to provide adequate tools to address rent seeking behavior (Bebchuk and Weisbach 2010). Moreover, there are clear incentives for stakeholders in the pay process to support upwardly biased peer groups. CEOs themselves have a clear financial incentive to select firms with well compensated CEOs into the peer group, because upwardly biased peer groups make overly generous pay to CEOs seem appropriate and defensible. Moreover, even though boards and compensation committees should push back and correct biases initiated or supported by the CEO (Shleifer and Vishny 1997, Iliev et al. 2015, Dicks 2012), there are various reasons why that may not be the case. For example, board members have incentives to maintain cordial social relations with the CEO (Gulati and Westphal 1999, Khurana 2011, Stern and Westphal 2010). Also, board members are often executives in their own right, and biased peer groups provide an indirect justification for a rise in their own compensation (Westphal and Zajac 1995). Finally, board members are responsible and accountable for hiring and employing a CEO who is above average, and above average CEO pay is arguably a validation that they are doing so (Khurana 2011). Thus, both the CEO and the principals of an organization may therefore actively push for or at least support the selection of upwardly biased peer groups.

2.2. Peer group bias: incentives and opportunities

Research on social comparisons suggests that there should be variation in the tendency to benchmark against similar others (Festinger 1954, Audia et al. 2016, Smith and Chae 2017) because there are differences in *incentives* and *opportunities* to select dissimilar comparison peers. Building on this research, we propose that weak performance on the metrics that affect CEO compensation creates an incentive to bias the peer group upward, while ambiguity in the set of firms most similar to the focal firm creates an opportunity to do so.

First, prior studies have demonstrated that if there are incentives to make a social actor look more favorable, the selection of a set of biased[3] comparison actors may achieve that goal. For example, experimental research has demonstrated that low performance motivates study participants to select a benchmark group of other low performers and that doing so dampens the perception of failure (Audia et al. 2016). Likewise, Elsbach and Kramer (1996) study how members of U.S. business schools respond to disappointing Business Week rankings. Their study suggests that disappointment with the rankings motivated business school members to selectively categorize their school in ways that placed them in more favorable interorganizational comparison groups. These studies imply that comparison groups provide a reference point which may shift the subjective interpretation of the performance of a social actor. In the context of CEO compensation, CEOs (and board members) of firms that fail to meet performance targets should have an especially strong incentive to manipulate the peer group. Assuming that CEOs aim to keep improving their compensation, failure to meet salient performance targets prevents CEOs or board members from citing prior achievements as an argument for an improved compensation package. An upwardly biased peer group may give them an alternative justification for a pay increase without attracting unwanted scrutiny from watchdogs or shareholders (Bertrand and Mullainathan 2001). Note that such behavior would also be consistent with earlier research that demonstrates that CEOs are rewarded for good luck but not penalized for bad luck (Garvey and Milbourn 2006). We therefore predict that CEOs of firms that have failed to impress financially are more likely to be benchmarked against an upwardly biased peer group.

Second, the ability to exercise discretion in the selection of compensation peers should create an opportunity to bias the compensation peer group upward. This argument builds on the idea that each firm varies in the extent to which there are other firms that one could plausibly argue to be similar. For example, Smith and Chae (2017) demonstrate that evaluators of organizations can exercise more discretion in the selection of comparison peers when the organization is atypical. They argue that evaluators of atypical organizations "have additional leeway to consider idiosyncratic attributions of performance when that performance deviates from a middle point." Their argument builds on earlier work that demonstrates that ambiguity about the right reference point allows actors to exploit that ambiguity in order to be evaluated more positively (Audia and Brion 2007).

[3] The bias may be upward or downward depending on which makes the actor look more favorable.

The idea of variation in ambiguity about the right reference point extends to the context of CEO compensation. While all firms self-select their compensation peers, some firms are able to exercise more discretion than others because they lack an obvious and homogeneous set of comparison peers. Variation in the ability to exercise discretion comes from two main sources. One is the number of natural peers. As we will describe later in this paper, the median reported peer group contains 17 firms. If the number of natural peers is larger than the typical size of a named peer group, firms have the opportunity to cherry pick peers with better compensated CEOs without engaging in obviously deviant behavior. A second source of discretion in the selection of peers is variation in compensation within the group of natural peers. Specifically, if the range in pay of CEOs of natural peers is wide, firms can use discretion to cherry pick the better compensated CEOs. Jointly, the number of natural peers and the spread of CEO compensation in this group influence the discretion firms can exercise when selecting their peer group. And since the discretion can be exercised within a group of firms that is easily justifiable, this behavior is unlikely to receive substantial scrutiny.

In sum, we predict that failure to meet performance targets that are tied to CEO compensation provides an incentive while benchmarking discretion provides an opportunity for firms to introduce bias into their compensation peer group. Before we turn to the analyses that test our predictions, we briefly review prior research on compensation peer groups.

2.3. Prior Evidence of Peer Selection

Some empirical evidence suggests that compensation peer groups are biased. An early study by Porac, Wade, and Pollock (1999) finds that, although selected peers are typically active in the same industry as the focal firm, firms select peers beyond industry boundaries when they perform poorly compared to other firms in the industry, when CEOs are highly paid, and when shareholders are powerful and active. More recently, Faulkender and Yang (2010, 2013) developed a method that builds on propensity score matching to construct counterfactual peer groups for each named peer group and then compare the compensation levels of CEOs in the named and counterfactual peer groups. They find that compensation in the named peer group tends to be larger than compensation in the counterfactual peer group. Moreover, in a regression of pay on these compensation differences, Faulkender and Yang (2010, 2013) establish that there is a positive association between peer compensation and compensation of the focal CEO. Using an algorithmic approach to identify counterfactual peer groups, Pittinsky and DiPrete (2013) also find that named peers are better compensated than natural peers.

Other evidence contrarily suggests that peer groups are constructed based on impartial selection methods. A recent study by Cadman and Carter (2013) suggests that previous studies suffered from a misspecification of the pool of potential peers, which, they argue, included firms that are unsuitable for benchmarking. For example, they criticize Faulkender and Yang (2010) because Faulkender and Yang limit their analyses of disclosed and potential peers to firms in the S&P 900. Cadman and Carter (2013) argue that many of the firms selected as peers but not in the S&P 900 are smaller and that the results therefore provide an incomplete picture of peer group selection. They also warn against samples of potential peers that are too broad. They claim that if the set of potential peers for the largest firm in the sample includes some of the smallest firms in the sample, peer selection would appear opportunistically biased towards larger firms. To address these issues, Cadman and Carter construct a sample of potential peers by taking the union of a focal firm's actual peers, the firms listed as peers of the focal firm's peers (i.e. peers of the peer firms), and any firms in the full sample of firms that list the focal firm as a peer. They use these data to predict peer selection with differences in company financials between the focal firm and potential peers, and they find that evidence of bias is weaker in their sample than in other samples that have been used to measure peer group bias. In a second test, Cadman and Carter regress compensation of the focal CEO on differences in company financials between the median named peer and the focal firm. The authors find a weak association between pay and upward social comparison, and they interpret this as evidence that bias in peer selection does not affect pay.

Although we agree with the proposition that how one operationalizes the pool of potential peers influences the estimate of peer group bias, the strategy proposed by Cadman and Carter (2013) does not address the issues they raise. Peer groups are generally biased in observable characteristics, a phenomenon that is well documented. Operationalizing the peer group by relying on *actual* peer choices, as in their strategy, necessarily oversamples larger and well-compensating firms. Because Cadman and Carter's operationalization of potential peers is biased, it is not surprising that Cadman and Carter report lower levels of bias in named peer groups. Note that this problem also plagues other papers cited here that use the method developed by Faulkender and Yang (2010, 2013). This method finds counterfactual peers *that match reported peers* when the right approach is to find counterfactual peers that match the focal firm. An additional issue that makes interpreting the claims of Cadman and Carter less than straightforward is that even when they use their preferred set of counterfactual peers they find that named peer groups are generally biased.

The recent study by Albuquerque et al. (2013) presents another argument to counter the claim that peer groups are biased in ways that allow CEOs to extract rent. Their study argues that the compensation distribution of the reported peer group is indeed often shifted to the right when compared to a natural peer group, but that the reason for this shift is that the CEOs in these named peer groups are more representative of the talent of the focal CEO than the CEOs in the natural peer group. To test this hypothesis, the authors construct counterfactual peer groups using the method introduced by Faulkender and Yang (2010). They then construct proxy measures for CEO talent and CEO self-serving behavior. The talent proxies include the number of media mentions, past abnormal performance of the firm(s) the CEO worked for in the past three years, and the size of the firm(s) the CEO worked for in the past three years, while the self-serving variables include a dummy for busy boards, the percentage of the board hired after the CEO, board size, the GIM corporate governance index, small institutional ownership, and small insider ownership. The authors then propose a decomposition of the pay gap between the median CEOs in the named and counterfactual peer groups into a self-serving component and a talent component. They report that while the self-serving component of peer group bias is positively associated with CEO compensation, the talent component has a much stronger association with CEO compensation, and they therefore conclude that what appears to be peer group bias is mostly an adjustment to take CEO talent into account.[4]

There are issues, however, that the study by Albuquerque et al. (2013) does not address.[5] The first issue is that Albuquerque et al. do not explain why firms with especially talented CEOs do not simply select an unbiased peer group and transparently compensate the CEO at a higher percentile. Given the talent of the CEO, doing so is unlikely to draw scrutiny.[6] A second problem is empirical. If Albuquerque et al.'s explanation for bias as a reward for talent is correct, then talent would be expected to have a strong positive effect on bias. This, however, is not the case in their data. The combined talent measures account for less than 5% of the variation in bias in their data. Even if one accepted that a small component of bias is explainable by CEO talent, it still leaves 95% of variation in bias that their paper does not explain. A final issue is methodological. Their "decomposition" of bias into two components (which together explain very little of their bias measure) obscures the negligible effects of the main predictors and also obscures the statistical association between their large unexplained component of bias and CEO compensation. We estimate a model that allows a direct test of Albuquerque et al.'s proposed mechanism in Appendix D. This model fails to support the "reward-for-talent" hypothesis proposed in their paper.

Also note that our prediction that firms that meet their financial targets are less likely to use upwardly biased compensation peers contradicts the finding by Albuquerque et al. (2013) that there

[4] Two other studies suggest that peer benchmarking is not affected by bias. The results of these studies, however, are difficult to interpret. Bizjak, Lemmon, and Naveen (2008) do not use data on reported peer groups, but put firms into two size bins (i.e. large vs. small firms) per industry and use other firms in the same bin as the focal firm as peers. Kim et al. (2015) do use data on named peers but they report that the software they used to extract peers from SEC filings harvests an average of 5.4 peers per peer group for the period 2006-2010. As we show below, our algorithm extracted an average of 19.6 peers per firm and a median of 16 peers per firm. Moreover, Cadman and Carter (2013) manually collect their data and report a median of 16 peers.

[5] We briefly discuss these issues here and discuss them in greater detail in Appendix D.

[6] Related to this issue are the problems with reliably measuring talent. Existing measures show very low correlations, questioning the validity of the concept (see Appendix D section A).

is a positive association between CEO talent and peer group bias. Assuming that more talented CEOs are better able to generate value for shareholders, the proposition of Albuquerque et al. (2013) would imply a negative association between the performance of the firm and peer group bias.

In the next sections, we introduce a strategy that should allow for the identification of an unbiased set of counterfactual peers and we then compare the outcome of that strategy to alternative measures. We then conduct analyses and provide tests of our central predictions that the size of bias responds to incentives and opportunities to increase it. We further establish that peer group bias is strongly associated with CEO compensation and (as demonstrated in Appendices B and D) that this result is robust to different methods for constructing counterfactual peer groups. Finally, we test for trends in the bias of named peer groups in response to environmental pressure for more transparent and normative compensation practices, and we assess the implications of trends for the importance of peer group bias as a strategic factor in the compensation of CEOs.

3. Data and Sample Selection

The main dataset used in this paper is constructed by the authors and is directly based on information contained in the DEF 14A or DEF 14C forms (also known as a definitive proxy statement) that firms are required to file annually with the Securities and Exchange Commission (SEC). These DEF 14 forms typically include a "Compensation Discussion and Analysis" section in which the reporting firm provides a rationale for the compensation package of its CEO and in which the various components of total executive compensation are described in detail. This is also the section in which peer firms are listed that are used for benchmarking executive compensation. We downloaded these forms for all firms in Morningstar's Executive Insight database, an executive compensation database, for the years 2006 to 2016.[7] While Executive Insight contains more than 7,500 unique firms, only a subset of those firms are publicly traded and file a DEF 14 form.[8]

After downloading the DEF 14 forms for the firms in Executive Insight, we created a computer algorithm that identifies and extracts the peers from these forms. Compared to existing algorithms, our algorithm allowed us to greatly reduce the level of missing information – both missing firms that actually report peers but are treated as if the information is not present, and companies that

[7] DEF 14s are submitted and made publicly available at the end of the fiscal year. Many fiscal years end during the spring – a period typically referred to as proxy season – so our data covers fiscal years 2006 to 2016.

[8] Not every firm selected into a peer group reports using peers itself. It is important to include these non-peer-reporting firms in the analysis in order to avoid error in the calculation of the extent to which peer groups are biased. Other studies that have analyzed the structure of peer groups have relied on the smaller Execucomp compensation data, but this is problematic because it often does not cover the whole peer group. The use of the larger Morningstar database addresses this problem. Because coverage of Executive Insight is a bit spotty in early years (2007,2008), we use Execucomp data where possible to fill in minor gaps.

Table 1: Observations by year

Year	Harvested	Harvested: S&P 1500	Count	Count: S&P 1500	Mean	Median	Peer group changed (%)
2006	1321	603	1035	571	17.15	15.00	
2007	1869	867	1587	838	18.99	16.00	46
2008	1955	952	1748	926	19.84	16.00	60
2009	1993	992	1810	974	19.86	16.00	62
2010	2082	1011	1867	985	20.23	16.00	61
2011	2106	1052	1943	1038	19.62	16.00	62
2012	2085	1081	1945	1066	19.81	16.00	66
2013	2091	1084	1936	1061	19.70	17.00	66
2014	2113	1071	1977	1054	19.85	17.00	63
2015	2020	1041	1914	1022	20.18	17.00	64
2016	1951	1023	1765	988	19.48	17.00	68
All	21586		19527		19.63	16.00	62
Unique firms	4290		3426				

Column 1 ("Harvested") shows the number of firms for which the algorithm identified a peer group and column 2 ("Harvested: S&P 1500") shows the number of firms from column 1 that are also in the S&P 1500. Columns 3 and 4 show the number of firms that remain once cases with missing data are removed. The final 3 columns show the mean and median number of firms in the compensation peer groups and the fraction of all firms that changes its peer group from one year to the next.

are actually named as peers but are overlooked by the computer algorithm. The resulting data set contains the named peers of 4,290 unique firms that reported having used compensation peers in at least one year between 2006 to 2016. This sample contains more firms for a larger number of years than prior studies that have used data on named compensation peers.[9]

Table 1 presents sample statistics by year. The first column shows the number of cases for which the algorithm was able to harvest peer group information, with the second column showing how many of those firms were in the of S&P 1500. In the later years, when peer compensation discussions become increasingly standardized, these numbers indicate that more than two thirds of the firms in the S&P 1500 used peer group benchmarking. The next two columns show the number of cases in the estimation sample (i.e., with cases removed due to missing data and firm identifier mismatches). Comparing these numbers with the first two columns shows that missing data becomes less of a problem over time, and that firms in the S&P 1500 are well covered. All remaining statistics reported in this paper always refer to the estimation sample. The table also shows that the median number of peers in a peer group has grown modestly while the mean remains rather stable between 2006 and 2016. Between 46% and 68% of all disclosing firms update their peer group from year to year by adding or removing peers.

After collecting the data from the DEF 14 forms, we merged each observation with financial data from Compustat. Variables included in Compustat and used in this study include market

[9] For example, Bizjak et al. (2011) collected data on 707 firms reporting in 2007. Faulkender and Yang (2010) analyzed a sample of 657 firms, and Faulkender and Yang (2013) employed a sample of 763 firms that have reported in at least one of four years resulting in a dataset of 2,066 firm-years (the average firm is represented in 2.7 years). Finally, Albuquerque et al. (2013) collected data from 3 years (2007 to 2009) for 1,273 unique firms that select peers and 3,221 firm-year observations, but that number drops to 2,158 because of missing data issues.

capitalization, revenue, return on assets (ROA), total assets, net income, market-to-book ratio, number of employees, and geographic location (state). Apart from compensation data, we also merged the collected data with board member data from Morningstar's Executive Insight and data from Thomson Reuters on institutional ownership. Indicators of institutional ownership are typically used as proxies for the quality of corporate governance in a firm. Finally, we merge our data with Institutional Brokers' Estimate System (IBES) data which provides us with information about which stocks are co-covered by which security analysts (Zuckerman 2004). We use this co-coverage data to capture similarity between firms that is not captured by other variables such as size and industry.

4. Determining the Peer Pay Gap

In line with prior research, we define the Peer Pay Gap (PPG) as the difference in total compensation between the CEOs of the named peers and the CEOs in a set of counterfactual peer firms. We first estimate a model that should represent an impartial selection process. We then use that selection process to generate multiple *counterfactual* peer groups for each *named* peer group. The goal of specifying the impartial selection processes is to generate peer groups that firms would have constructed had there not been any form of bias in the selection process.

Specifying a counterfactual selection process is not straightforward. Compensation consultants and corporate watchdog organizations agree that industry and size similarity are two dimensions on which a focal firm and its peers must match (Reda et al. 2008, ISS 2015). Moreover, a recent report on peer groups issued by Equilar, Inc. listed the top 10 dimensions that firms claim they use in the peer selections process. From most to least often used, the list includes the following dimensions: Industry, Revenue, Market Capitalization, Competition for Talent, Business Model, Direct Competitor, Geographic Location, Assets, Number of Employees, and Profitability. While this information is helpful to determine which dimensions to include in a model, the question remains of how these dimensions should be weighted and what the acceptable boundaries are for labeling two firms as similar.

Prior work mostly relied on propensity score matching (PSM) to construct counterfactual peer groups (Faulkender and Yang 2010, Bizjak et al. 2011, Faulkender and Yang 2013, Albuquerque et al. 2013). These studies create a risk set of potential peers for each selecting firm, leading to N (number of selecting firms) * M (number of firms in the pool of potential peers) observations. The dependent variable is binary and captures whether the dyad was realized (i.e. whether firm i actually selected firm j as a peer). The independent variables are dummy variables that capture similarity for categorical variables and whether the firms are in the 50% to 200% range of each other for continuous variables. To construct the counterfactual peer group, the predicted probabilities from this model are used. Each named peer is matched to an unnamed peer that has the closest absolute distance to the named peer in terms of their propensity scores, and that has not yet been matched to another named peer. This generates a counterfactual peer group of the same size as the named peer group, and, by construction, does not include any of the named peers. The PPG is then measured as the ratio of the median pay of the named and counterfactual peer group.

While this method dominates the research done in this area, it has substantial shortcomings. First, the method is set up to find counterfactual peers that are similar to the named peers, not similar to the focal firm. Obviously, if there is bias in the observed characteristics of the named peers, the counterfactual peer group would suffer from the same bias. Second, the method forces named peers to be matched to firms not named as peers. This ignores the fact that at least some firms, such as airlines, are partially constrained in their peer choice, thus indicating a set of natural peers that should realistically always be included (and whose absence will be noted by the firm's shareholders). In those cases, some or all of the alternative peers represent a much worse fit than the named peers, leading to counterfactual peer groups that are *less* likely to be regarded as a natural peer group by the firm and its shareholders. Third, the dummy variables in the model do not capture the magnitude of the difference in observed covariates, because a potential peer is considered a match if it is within a 50-200% range of the focal firm. The model does not distinguish, for instance, a potential peer that is 10% larger than the focal firm from a potential peer that is twice the size – although, arguably, the first represents a much better match. The model may also find a much smaller counterfactual peer for a much larger named peer and vice versa.[10]

To address these issues, we developed a method that builds on the idea that reciprocated peer nominations – that is firm i chooses firm j and firm j chooses firm i – are unlikely to suffer from a biased selection process. The unit of analysis in this method is a tie between two firms, which can be either reciprocated or not. Since a nomination can only be reciprocated if both firm i and firm j report peer groups, we limit this analysis to firms that report peer groups. The number of possible ties between these firms is the total number of combinations, $\binom{N}{2} = \frac{N!}{2(N-2)!}$, where N is the number of peer-reporting firms. Naturally, only a very small fraction of the total number of possible ties is reciprocated (in our sample, around .2% of about 19.2 million possible ties). We then estimate a logit model predicting reciprocity, using a set of variables including industry, revenue, geography, market capitalization, assets, and number of employees. We capture differences in size by including more fine-grained dummy variables, and we include a graded measurement of industry similarity.

[10] For example, BJs Restaurants Inc., a named peer of Famous Daves of America Inc. with a revenue of $316 million, can be matched to potential peer MGM Resorts International, which has 24 times as much revenue - $7.7 billion. Since BJs Restaurants Inc. was less than half the size of Famous Daves of America Inc. and MGM Resorts International was more than twice as large, the value of the dummy variable for size was the same for the potential peer MBM Resorts International and the named peer BJs Restaurants, Inc.

Since there is not a single existing database that provides information on a firm's business model or direct competitors, we proxy these variables by including a measure of security analyst co-coverage. This is a dyadic version of the firm-level measure introduced by Zuckerman (2004). Using IBES data on the coverage of stocks by security analysts, we calculate the cosine similarity of analyst coverage for each pair of firms.[11] Higher cosine similarity means that the pair of firms is covered by similar analysts. Because these data are not available for all firms in our sample, we include two "missing" indicators: the first equals 1 if one of the firms in the pair attracts coverage from one or zero analysts and the second equals 1 if each of the two firms is covered by one or zero analysts. The reference group includes pairs of firms that are both covered by two or more analysts.

Table 2 shows the coefficient estimates of this model. The model indicates that industry similarity, being covered by similar analysts, and not being headquartered in the same state make firms much more likely to nominate each other as peers. It also shows that firms are more likely to select other firms that are similar in size. Net of the other size criteria, revenues are most strongly predictive and assets are least predictive of choosing a peer who reciprocates the choice. A pair of firms that matches perfectly on every dimension is predicted to have a probability of reciprocation of 99%.

We then use the estimates from this model to establish a probability of reciprocal peer nomination using the universe of Compustat firms. Given that peer groups are typically upwardly biased, we include all Compustat firms for which we have financial data. Restricting the sample only to firms that select peers or are selected as peers may result in a shortage of smaller firms that are necessary to produce balanced peer groups. The intuition is that the model assigns high probabilities to firms that could have been or were reciprocated peer nominations for a focal firm, and which thus should represent unbiased choices.

The distribution of predicted probabilities for each focal firm is often skewed, with a small number of potential peers that are predicted to be very good matches and a long tail of firms with substantially lower predicted probabilities. We argue that each firm is at least partially aware of this distribution: Some peer choices are obvious to the firm and its shareholders, and if those are not included, this will likely draw attention to the firm's peer group. The number of "obvious" peer choices naturally varies among firms, indicating that some firms have more discretion in constructing their peer group than others. To capture the uncertainty in this process, i.e. the number of peer choices that are "obvious," we implement a simulation approach to the construction of counterfactual peers. In each iteration of the simulation, we sample N_{it} firms from this population of potential peers (where N_{it} is the number of peers selected by firm i in year t) using their predicted probabilities from the reciprocity model as weights. Because the weights alone do not ensure that a

[11] i.e., similarity(A, B) $= \sum^n A_i B_i / (\sqrt{\sum^n A_i^2}\sqrt{\sum^n B_i^2})$, where A and B are vectors of length n reflecting the number of unique analysts. A_i is 1 if the security analyst i covers firm A, 0 otherwise (and same for B_i).

Table 2: Prediction Equation for Peers from the Simulation Model

	Dependent variable: Reciprocated tie (1/0)
	Logit
Industry similarity, 1-digit	1.618^{***} (0.026)
Industry similarity, 2-digit	2.914^{***} (0.022)
Industry similarity, 3-digit	3.575^{***} (0.024)
Industry similarity, 4-digit	4.553^{***} (0.019)
Analyst coverage coherence (IBES)	5.286^{***} (0.036)
Missing: Either firm	-0.300^{***} (0.014)
Missing: Both firms	-0.698^{***} (0.028)
Same state	0.619^{***} (0.015)
Revenue within 10-20%	0.216^{***} (0.047)
Revenue within 20-30%	0.571^{***} (0.043)
Revenue within 30-40%	0.950^{***} (0.041)
Revenue within 40-50%	1.286^{***} (0.040)
Revenue within 50-60%	1.553^{***} (0.040)
Revenue within 60-70%	1.724^{***} (0.040)
Revenue within 70-80%	1.779^{***} (0.040)
Revenue within 80-90%	1.833^{***} (0.040)
Revenue within 90-100%	1.879^{***} (0.041)
Market cap within 10-20%	0.357^{***} (0.039)
Market cap within 20-30%	0.668^{***} (0.038)
Market cap within 30-40%	0.867^{***} (0.037)
Market cap within 40-50%	1.004^{***} (0.037)
Market cap within 50-60%	1.114^{***} (0.038)
Market cap within 60-70%	1.227^{***} (0.038)
Market cap within 70-80%	1.220^{***} (0.038)
Market cap within 80-90%	1.262^{***} (0.039)
Market cap within 90-100%	1.246^{***} (0.039)
Employees within 10-20%	0.778^{***} (0.043)
Employees within 20-30%	1.098^{***} (0.042)
Employees within 30-40%	1.324^{***} (0.041)
Employees within 40-50%	1.464^{***} (0.041)
Employees within 50-60%	1.527^{***} (0.041)
Employees within 60-70%	1.634^{***} (0.041)
Employees within 70-80%	1.632^{***} (0.041)
Employees within 80-90%	1.631^{***} (0.042)
Employees within 90-100%	1.614^{***} (0.042)
Assets within 10-20%	0.937^{***} (0.060)
Assets within 20-30%	1.334^{***} (0.058)
Assets within 30-40%	1.527^{***} (0.057)
Assets within 40-50%	1.663^{***} (0.057)
Assets within 50-60%	1.734^{***} (0.057)
Assets within 60-70%	1.810^{***} (0.058)
Assets within 70-80%	1.836^{***} (0.058)
Assets within 80-90%	1.882^{***} (0.058)
Assets within 90-100%	1.919^{***} (0.058)
Constant	-12.515^{***} (0.067)
Observations	19,274,475
Log Likelihood	−137,350.800

Note: *p<0.1; **p<0.05; ***p<0.01, standard errors are in parentheses.

balanced peer group is drawn, we implement the following procedure: For each iteration and firm, one out of four covariates (market cap, assets, revenue, number of employees) is randomly chosen. Based on the value of this covariate for the focal firm, we split the population of potential peers into two sub-populations; one that has covariate values larger than the focal firm, and the other has smaller covariate values than the focal firm. We then draw $N_{it}/2$ peers from each of the two sub-populations. In the limit, this process generates peer groups that are both balanced on the four covariates and represent the most natural peer choices. We repeat this process 500 times for each named peer group, generating 500 counterfactual peer groups.[12]

This strategy corrects for at least two problems with methods used in prior studies. First, we correct for the problem that Cadman and Carter (2013) identified, namely, that a large pool of potential peers – which often includes relatively smaller firms – increases the chance to find positive bias in peer group selections. We correct for this problem by forcing reciprocated peer selections to guide the identification of counterfactual peers. Second, by finding counterfactual peers that are similar to reciprocating peers rather than the full set of actually selected peers, we address the concern that bias on observables mechanically affects the counterfactual peer group. Our method should substantially limit the potential of bias entering the process whereby counterfactual peers are identified. Moreover, the uncertainty in the process of finding counterfactual peers is quantified by our simulation approach.

In addition to the theoretical arguments presented above, we find that the reciprocity method produces counterfactual median peer pay that, in combination with the PPG, does a better job of predicting actual CEO compensation than does the Faulkender and Yang method, which was also used in Albuquerque et al. (2013) (reported in Appendix Table B1). Furthermore, Figure 1 shows a visual comparison of the two methods (plots for other covariates are in Appendix B). The solid line demonstrates that peer groups are biased in observed covariates, by plotting the logged ratio of the market cap of the focal firm and the median market cap in the observed peer group. This distribution is shifted to the right of the balanced reference point of zero. The dotted and dashed lines then compare our simulation method to the method developed by Faulkender and Yang, comparing the median market cap of the counterfactual peer group to the market cap of

[12] Constructing peer groups for very large or very small firms raises the issue of ceiling effects and, to a lesser extent, floor effects. Floor effects are less likely because we include the whole universe of Compustat firms, which include many much smaller firms than are represented in our sample of peer-nominating firms. To ensure that ceiling or floor effects do not drive our results, the very few firms that are so large or so small that they do not have sufficient potential peer groups to choose from are dropped from our analysis. Specifically, we only include firms for which we found more than 250 valid peer groups (out of 500). Using this criterion, only 23 firm-years are dropped from the analysis. Because it is logically impossible to produce *any* balanced peer group for these firms, any peer benchmarking process is inherently inefficient for these cases. In a set of sensitivity analyses (Appendix Table C1), we also trimmed both the largest and smallest 5% and 10% firms on each of our four size criteria. The results from the sensitivity analyses are very similar to those for the entire sample.

Figure 1: Logged balance ratios comparing observed peer groups to counterfactual methods.

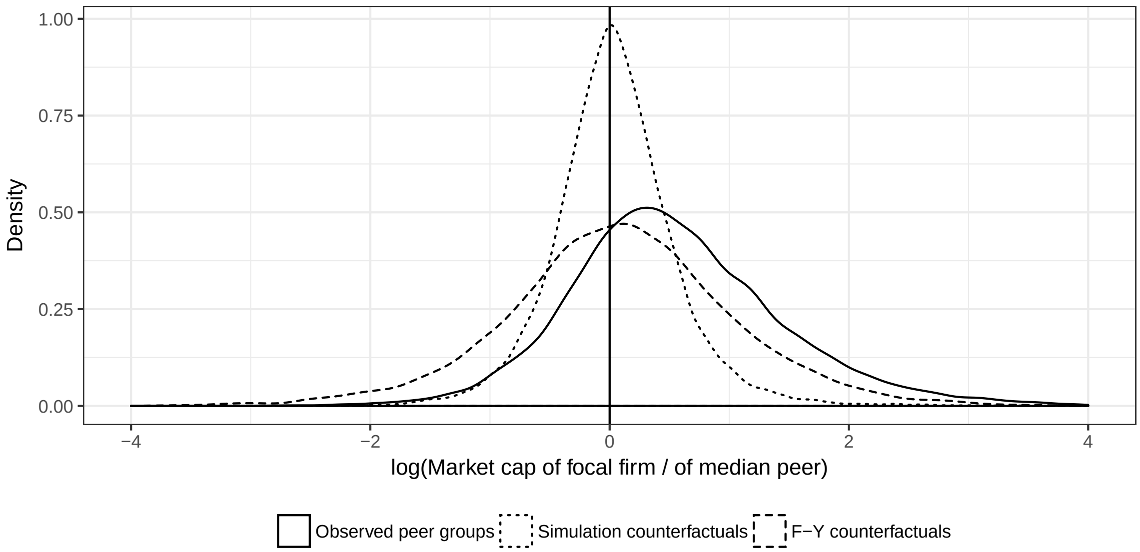

the focal firm. We replicated Faulkender and Yang's method using our own data, to make sure that these comparisons are meaningful. While both methods produce balanced peer groups, the simulation method does so more precisely. Given these two empirical comparisons, we proceed by adopting the results from the reciprocity method in the rest of the paper.

By drawing 500 natural peer groups for each firm, we capture uncertainty about the PPG. Specifically, the simulations generate a distribution of possible compensation benchmarks (i.e., the distribution of median pay of each of the 500 simulated peer groups) and any of the 500 counterfactual peer groups could plausibly be argued to represent a "natural" peer group. The median compensation of any of these peer groups could then be justified as an appropriate benchmark. As a result, firms with a wider range of median compensation in its 500 natural peer groups are expected to be able to exercise more discretion in selecting its peer group.

The range of median compensation in these peer groups is a function of two factors. The first is the overlap or intersections of all 500 peer groups. In other words, if the 500 peer groups look very similar in composition, the rate of overlap is high. The rate of overlap can be captured by the distribution of predicted probabilities of the potential peers of the focal firm. If there is a limited number of potential peers with high predicted probabilities, the firm is constrained in its peer choice by its environment. We call this *peer group constraint* and measure it as the median predicted probability of the 20 potential peers with the highest predicted probabilities.[13] The second factor is the distribution of pay in the set of natural peers. We call this *benchmarking discretion*, which reflects the variation in median pay in the 500 natural groups. For each firm, we take the median

[13] We chose 20 because it is the mean peer group size, but the exact number does not drive any of the results.

pay in each of the 500 peer groups. For this distribution of 500 values, we then calculate the difference between the 75^{th} and the 25^{th} percentile and divide it by median peer compensation (i.e., $\frac{q_{75}-q_{25}}{q_{50}}$). Our prediction is that firms with more discretion in choosing firms with well compensated CEOs within their set of natural peers are more likely to select an upwardly biased peer group.

To effectively account for the simulation uncertainty in our regression models, all models reported below are averages over 500 separate estimations. We adopt the procedure of Little and Rubin (1989, p. 305) in averaging over the outcomes of 500 different regressions. The value of any regression coefficient $\hat{\mu}$ is given by

$$\hat{\mu} = \frac{1}{M}\sum_{l=1}^{M}\mu_l,$$

where M represents the total number of simulations, and $l = 1 \ldots M$. The associated standard error $\hat{SE}$ equals

$$\hat{SE} = (\bar{U} + \frac{M+1}{M}B)^{1/2},$$

which involves average-within simulation variability $\bar{U}$ and between-simulation variability B:

$$\bar{U} = \sum_{l=1}^{M}\frac{SE_l^2}{M} \qquad \text{and} \qquad B = \sum_{l=1}^{M}\frac{(\mu_l - \hat{\mu})^2}{M-1}$$

Whenever we report sample statistics that concern values derived from the simulated peer groups (such as Figure 1), we first take the median over the distribution of values for each firm-year, and then report the appropriate statistic using the averaged distribution.

5. Determinants of Bias

Figure 2 positions the compensation of the focal CEO with respect to the named and counterfactual peer groups. The difference between using the named and the counterfactual peer group is obvious; while CEOs generally appear to be paid below the median named peer (at about the 45^{th} percentile), they are paid above the median (at around the 55^{th} percentile) counterfactual peer.[14]

Next, we assess firm characteristics that are associated with a larger or smaller PPG. Descriptive statistics of the variables are presented in Table 3.[15] Recall that our two main predictions are

[14] If all firms reciprocated their peer choices, the percentile position would be constrained to be at the 50^{th} percentile. The simulation method builds on reciprocated ties, but then counterfactual peers are chosen from the universe of potential peers based on their probability of reciprocating. This frees the simulation from being constrained to select counterfactual peer groups with a mean percentile of 50.

[15] Due to a change in reporting requirements issued by the SEC that applies to proxy statements for fiscal years ending on or after December 20th, 2009, the total compensation variable is not consistently measured before and after 2009. The sources for this variable are Morningstar's ExecutiveInsight and Compustat's Execucomp database,

Figure 2: Percentile in counterfactual peer group at which the CEO is compensated.

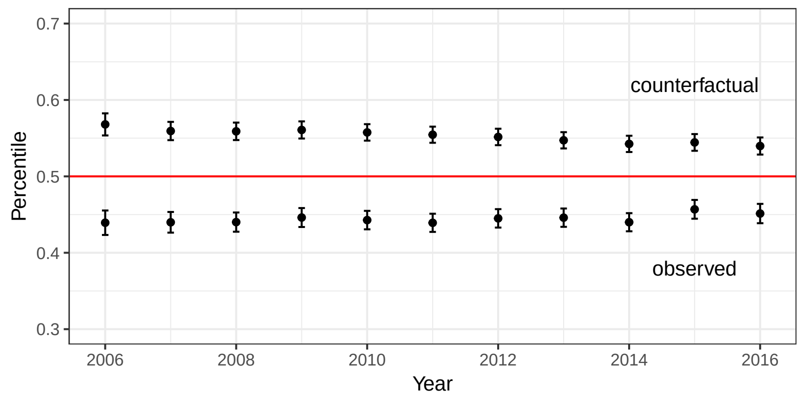

Note: The graph shows the mean of the percentiles in the counterfactual peer groups at which the CEOs in our sample are compensated. The error bars show the 95% confidence intervals around the mean.

that the PPG should be higher when discretion in the selection of peers can be exercised and when the company fails to meet important financial targets to which CEO compensation is tied. To test this second prediction, we include performance metrics at the firm level that are most likely to affect CEO compensation through the pay-for-performance contract. If firms perform well on these metrics, the CEO is unlikely to need a biased peer group to improve pay because recent achievements should justify generous pay. Research on how CEO compensation is tied to performance metrics at the firm level suggests that there is substantial variation in the metrics used to incentivize CEOs (Jensen and Murphy 2010, De Angelis and Grinstein 2015). Most firms use stock and option packages which respond directly to growth in market capitalization. Some firms also use bonuses for accounting performance. We therefore include two variables that should capture performance and that should affect compensation of the CEO, namely market capitalization and return on assets (ROA).[16]

Table 4 shows the results of regressions in which the PPG is the dependent variable. We estimate the PPG as a function of measures of institutional ownership, the mean PPG of other firms for

which in turn record the compensation value that the companies report as "Total compensation" in their annual proxy statement. The reporting change affected the calculation of stock and option awards. It is unclear how and if this reporting change systematically affects our models. To ensure that our results are not driven by the effects of this reporting change, we split our sample into two periods: 2007-2008 and 2010-2016. (We leave out 2009 because companies may be slow to comply.) We then estimate our models on the two samples separately (see Appendix Table C2) and find that the results are similar to the models containing all years. We thus proceed by presenting these models in the paper. We thank an anonymous reviewer for pointing us to this issue.

[16] Since there is an important link between the talent argument proposed by Albuquerque et al. (2013) and our argument about bias and how it relates to companies meeting their financial targets, we also evaluated talent measures, including those proposed by Albuquerque et al. (2013), and found that the correlations between them are generally very low (see Appendix D). We conclude that there is no reliable method to measure CEO talent and that measurement issues may be the reason for our predictions and their findings to be irreconcilable.

Table 3: Variables definitions and their means and standard deviations

Variable	Description and source	Median	Mean	SD
Total compensation	As reported in the company's proxy statement. This item includes the components below (Morningstar ExecutiveInsight/Compustat Execucomp) (in thousands)	3540.12	5439.09	5766.44
	– Salary	675.00	725.77	342.75
	– Bonus	755.44	1273.37	1642.41
	– Equity-based compensation	1899.48	3190.33	3828.60
	– Other compensation	87.11	560.83	1316.00
Peer compensation*	Median compensation in the observed peer group (in thousands)	4108.83	5271.05	4050.10
Counterfactual comp.*	Median compensation in the counterfactual peer group (in thousands)	3226.53	4155.19	3142.29
PPG*	Ratio of median pay in the observed peer group and median pay in the counterfactual peer group	1.23	1.34	0.52
Benchmarking discretion	Spread of median pay in the counterfactual peer across simulations (for the same firm-year). Defined as (q75-q25)/q50, where q refers to quantiles of the distribution (own calculations)	0.23	0.25	0.11
Peer group constraint	Median value of the 20 highest predicted probabilites of reciprocation	0.10	0.12	0.10
Revenue	Annual revenue (Compustat) (in millions)	937.56	4552.69	12061.54
Market cap	Market value at end of year (Compustat) (in millions)	1300.33	6371.19	17621.94
Employees	Number of employees (Compustat) (in thousands)	2.63	13.07	32.83
Assets	Total assets (Compustat) (in millions)	1690.12	8609.16	24075.27
Return on assets	Ratio of net income to assets (own calculation based on Compustat)	0.03	-0.01	0.19
Book value per share	Shareholders' equity divided by outstanding shares (Compustat) (in $)	10.52	13.70	13.45
Institutional ownership	Percentage of stock held by institutional owners (Thomson Reuters)	0.76	0.72	0.21
Board mean PPG	The average PPG in other firms with which the focal firm's board members are interlocked. For each board member in the focal firm, we identified whether there was an interlock with the boards of other firms. For the interlocked members, we measured the PPG at the interlocked firms. For each committee member, we computed the mean of these values as the measure of that individual's attitude, and then we computed the mean for all board members as the focal firm's board attitude. When we can't identify board members, we set the value to the median within the 2-digit industry.	0.05	0.06	0.07

Note: All distributions have been winsorized at the $.5^{th}$ and 99.5^{th} percentiles to remove extreme outliers.

which board compensation committee members are directors, market capitalization and ROA, the log of the benchmarking discretion and peer group constraint measures. We also include year fixed effects and dummy variables for the various compensation consultants. We thus estimate a model in the form:

$$log(\mathrm{PPG}_{it}) = X_{it}\beta + \alpha_t + u_{it},$$

where X denotes a matrix of covariates, β is a vector of coefficients for these covariates, α_t estimates year-specific intercepts, and u denotes the error term. Fixed effects for firms are added

Table 4: Models for the Peer Pay Gap (PPG).

	Dependent variable: PPG					
	log(PPG)	log(PPG)	log(PPG)	log(PPG)	log(PPG)	log(PPG)
	(1)	(2)	(3)	(4)	(5)	(6)
Institutional ownership	-0.051^{**}	-0.044	0.038	0.012	0.036	0.021
	(0.023)	(0.031)	(0.023)	(0.031)	(0.022)	(0.031)
Board mean PPG	0.147^{***}	0.111^{*}	0.312^{***}	0.118^{**}	0.265^{***}	0.105^{*}
	(0.057)	(0.058)	(0.058)	(0.057)	(0.056)	(0.057)
ROA			-0.147^{***}	-0.055^{*}	-0.152^{***}	-0.054^{*}
			(0.024)	(0.030)	(0.023)	(0.030)
log(Market cap)			-0.032^{***}	-0.081^{***}	-0.023^{***}	-0.076^{***}
			(0.003)	(0.008)	(0.003)	(0.008)
log(Benchmarking discretion)					0.119^{***}	0.049^{***}
					(0.010)	(0.009)
log(Peer group constraint)					-0.039^{***}	-0.034^{***}
					(0.006)	(0.009)
Firm fixed effects	No	Yes	No	Yes	No	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Compensation consultant dummies	Yes	Yes	Yes	Yes	Yes	Yes
Number of simulations	500	500	500	500	500	500
R2	0.016	0.017	0.044	0.038	0.074	0.043
Adj. R2	0.015	-0.194	0.042	-0.169	0.073	-0.163
Observations	19,527	19,527	19,527	19,527	19,527	19,527

Note: *p<0.1; **p<0.05; ***p<0.01, standard errors are in parentheses.

in some models. The standard errors reported in Table 4 and all following models are robust and adjusted for correlation within firms.

Model 1 is a pooled OLS regression while model 2 includes firm fixed effects. In these models we find that bias is generally greater when board members are directors at other firms that themselves have higher bias in their named peer groups and that there is no relationship between institutional ownership and the PPG.

Model 3 and 4 include our two measures of firm performance. Both the pooled regression in model 3 and the estimates in the fixed effects model show that both variables are negatively associated with PPG. ROA has a meaningful association with the PPG: a one standard deviation decrease in ROA is associated with an almost 3% increase in the PPG according to the estimates from model 3. In the model that includes firm fixed effects, the coefficient reduces to -0.055 (which implies a 1% PPG increase with each standard deviation decrease in ROA) and is marginally significant. The market capitalization variable is significantly different from zero in each of the two models and the coefficient is in the predicted direction. Model 4 suggests that a 10% drop in market capitalization is associated with an increase of the peer pay gap of almost 1%. These results demonstrate that CEOs who create value for shareholders are benchmarked against peer groups that are less biased and they are consistent with the idea that biased peer groups may serve as an alternative justifications of generous compensation packages.

In models 5 and 6 in Table 4 we examine the relationship between benchmarking discretion and peer group constraint on the one hand and PPG on the other. Recall that our prediction is that greater benchmarking discretion makes firms upwardly bias the peer group and that having higher peer group constraint should reduce the PPG. The model results in Table 4 align with those predictions and demonstrate that discretion and constraint affect the PPG net of other structural characteristics of firms. The effects are substantively meaningful. Imagine a firm with average benchmarking discretion and a PPG of 1.10, i.e. the median compensation of the named peer is 10% higher than average of the counterfactual simulations. If an otherwise comparable firm had a benchmarking discretion that is one standard deviation higher than the mean, its PPG is predicted from Table 4, model 5 to be 1.14 (1.1+0.119*(log(.36)-log(.25))), which is slightly less than a 50% increase in peer group bias. The within-firm changes of discretion are, of course, smaller than the between-firm differences, but the effect of benchmarking discretion on bias remains statistically significant; the coefficient estimates in Table 4 imply that a one standard deviation increase in discretion for the same firm would raise a 1.10 PPG to a 1.12 PPG. The descriptive relationship between benchmarking discretion and the PPG can also be seen in Figure 3. The upper panel of this figure shows firms with little discretion while the bottom panel shows firms with high discretion. This figure demonstrates that the PPG varies systematically across the distribution of benchmarking discretion.

The relationship between the PPG and peer group constraint is also in the predicted direction. Firms with more ambiguity (i.e. lower peer group constraint) select peer groups with more upward bias in pay. The size of the association is substantial: in the model that includes firm fixed effects, a one standard deviation decrease in peer group constraint increases the PPG by about 6%.

The results in Table 4 raise an important question: Is the use of biased peer groups limited to firms that fail to meet their performance targets and have a lot of benchmarking discretion? Or do these conditions merely increase the expected bias beyond an expected positive bias even for the firms that are performing well and have little benchmarking discretion? Our models suggest that even high performing firms in unambiguous structural positions assemble biased peer groups. This is seen by examining the combined impact of a one standard deviation reduction in benchmarking discretion and a 10% increase in market capitalization (and assuming average peer group constraint and ROA). The combined impact of these predictors of bias (from Table 4, model 6) is

Figure 3: The distribution of PPG across Benchmarking Discretion

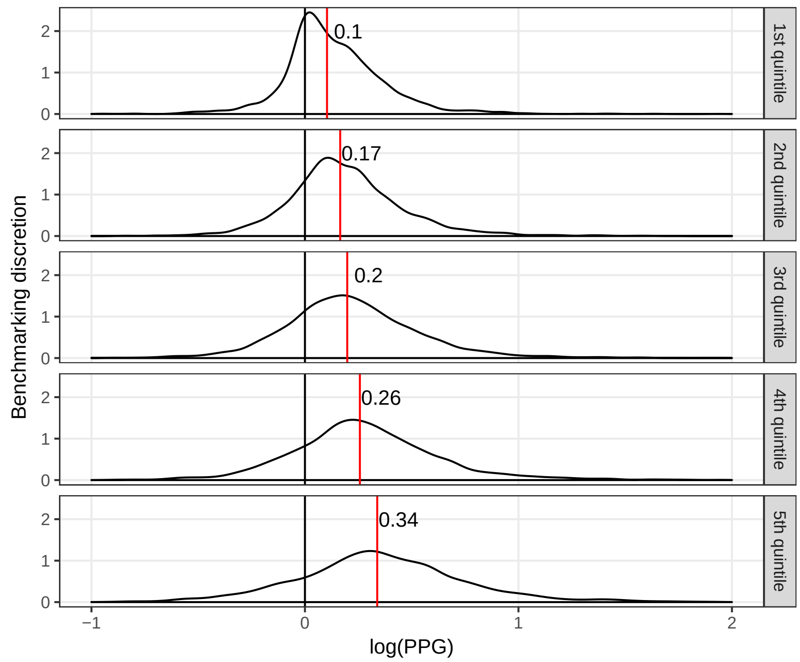

-.036 ((.049*(log(.14)-log(.25))) + (-0.076*(log(7008)-log(6371)))); in other words, the mean bias of about 1.34 from Figure 2 would be reduced to 1.30, which means that the expected named peer group still has an expected 30% higher median pay than does the normative peer group from the simulation model.

In sum, the variation in the PPG across firms is largely unexplained by firm characteristics, but part of the variation is systematic. We find that more discretion and weak firm performance are associated with the tendency to assemble biased peer groups.

6. The Relationship between Bias and CEO Compensation

Although the models presented in Table 4 suggest that the PPG is real, the practical implications of our findings would be limited if the PPG and the actual pay of the CEO were not positively associated. As a first test of the prediction that bias in the benchmarking process is associated with compensation levels of CEOs, we compare the observed compensation that CEOs received and the compensation that the CEO would have received under two alternative scenarios. In the first scenario, which we call the "50^{th} percentile" scenario, we computed what the compensation of the CEO would have been had the CEO been compensated at the 50^{th} percentile of the counterfactual peer group using the simulation method. In the second scenario, the "$n^{th} percentile$" scenario, we compute the percentile in the named peer group at which the CEO was compensated and look up the compensation level that corresponds to that percentile in the counterfactual peer group.

Figure 4: Compensation under two alternative scenarios

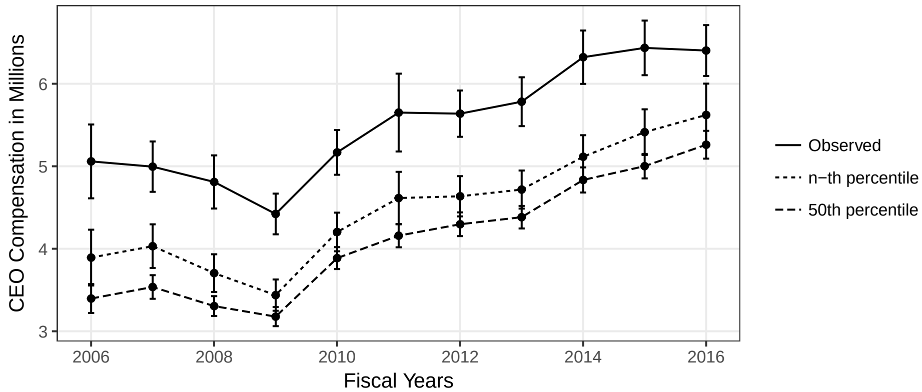

Note: The graph shows the yearly trends of observed CEO compensation and CEO compensation under two alternative scenarios. The error bars show the 95% confidence intervals around the mean.

The results of these analyses are shown in Figure 4. All three scenarios show a similar trend: CEO compensation went down during the first years of the financial crisis, but it increased again starting in 2010. The differences between the three scenarios are pronounced and consistent. The compensation that CEOs actually received is substantially higher than the compensation that CEOs would have received under the two alternative scenarios. The difference averages around \$1 million when compared to compensation under the "$n^{th} percentile$" scenario and about \$1.75 million under the "$50^{th} percentile$" scenario. Although this evidence does not show that bias in benchmarking *caused* compensation to increase, it is certainly consistent with our prediction that bias in the benchmarking process is potentially used to justify high levels of compensation.

Next, we estimate a model for CEO compensation to assess whether compensation is higher when bias in the named peer group increases, net of other factors. We model compensation as a function of the pay of the median peer in the counterfactual peer group along with structural characteristics of the firm, firm performance, and bias in the named peer group. We thus estimate a model of the form:

$$log(\text{Compensation}_{it}) = \gamma_1 log(\text{Counterfactual peer compensation}_{it}) + \gamma_2 log(\text{PPG}_{it}) + X_{it}\beta + \alpha_t + u_{it},$$

where X denotes a matrix of control variables, α_t estimates year-specific intercepts, and u_{it} denotes the error term. Of special interest is the parameter for the PPG, γ_2. Fixed effects for firms are added in some models. Table 5 shows the results of these regressions. Model 1 shows that every percent change in the pay of the median counterfactual peer produces an equal percent change in the expected pay of the focal CEO if the peer group is unbiased. This result gives additional confidence that the reciprocity method is producing a reasonable set of counterfactual peers. Net of

Table 5: Models for Executive Compensation.

	Dependent variable: Compensation					
	log(Comp.)	log(Comp.)	log(Comp.)	log(Comp.)	log(Comp.)	log(Comp.)
	(1)	(2)	(3)	(4)	(5)	(6)
log(Counterfactual comp.)	1.034^{***}	0.473^{***}	0.960^{***}	0.415^{***}	0.614^{***}	0.213^{***}
	(0.009)	(0.023)	(0.011)	(0.024)	(0.020)	(0.026)
log(PPG)	0.536^{***}	0.257^{***}	0.513^{***}	0.230^{***}	0.427^{***}	0.145^{***}
	(0.019)	(0.019)	(0.019)	(0.018)	(0.019)	(0.018)
Book value per share			0.003^{***}	0.007^{***}	-0.001^{*}	0.0004
			(0.001)	(0.001)	(0.001)	(0.001)
Market-to-book ratio			0.027^{***}	0.005	−0.008	−0.004
			(0.006)	(0.003)	(0.005)	(0.003)
Return on assets			0.022	0.239^{***}	-0.219^{***}	0.009
			(0.033)	(0.039)	(0.036)	(0.039)
Institutional ownership			0.349^{***}	0.315^{***}	0.384^{***}	0.189^{***}
			(0.038)	(0.050)	(0.039)	(0.049)
log(Revenue)					−0.002	0.013
					(0.007)	(0.013)
log(Market cap)					0.180^{***}	0.175^{***}
					(0.009)	(0.011)
log(Employees)					0.009	0.066^{***}
					(0.006)	(0.022)
log(Assets)					0.030^{***}	0.084^{***}
					(0.008)	(0.023)
Firm fixed effects	No	Yes	No	Yes	No	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Number of simulations	500	500	500	500	500	500
R2	0.664	0.167	0.67	0.18	0.693	0.215
Adj. R2	0.664	-0.012	0.67	0.004	0.693	0.046
Observations	19,527	19,527	19,527	19,527	19,527	19,527

Note: *p<0.1; **p<0.05; ***p<0.01, standard errors are in parentheses.

the counterfactual pay, CEO compensation also rises with the bias in the named peer group; each 1% increase in the PPG produces a 0.54% increase in the compensation of the CEO. Because the PPG is the ratio of the median pay of the named peers and the median pay of the counterfactual peers, this specification is algebraically equivalent to a model that is linear in the logs of the median counterfactual and median named peer groups. In other words, model 1 implies that, net of the median pay of the named peer group, each 1% increase in the pay of the median counterfactual peer implies a 0.5% increase (1.03-.54) in the focal CEO compensation.[17] In short, firms act as if they are paying attention to the pay in their natural peer group as well as in their named peer group. However, because bias is typically positive, they are paying higher compensation than would be the case if they only paid attention to the pay distribution in the natural peer group.

Model 2 in Table 5 adds firm fixed effects. The coefficients for counterfactual compensation and bias are reduced in size but remain positive. In other words, firms respond to changes over time in the median pay of the natural peer group but they are also responding to the changing amount of bias in the named peer group. Of course, we are not making an assertion about the direction of causality here. These models do not say that firms exogenously receive named peer groups, look at the pay in these groups along with the pay in the natural peer group, and then set the pay of their executive. Rather, we conjecture that the named peer group is determined partly in response to what the compensation committee of the board of directors is willing to pay their CEO, thus creating a justification for his compensation package. If firms were unable to adjust the named peer group for several years once chosen, then the coefficient for the PPG in model 2 (with fixed effects for the firm) would more plausibly reflect a causal effect, i.e., changes in compensation for the peer group would drive changes in focal CEO compensation. However, firms do change their named peer groups on a regular basis, as we have shown in Table 1, and even if they do not change their peer group, they *choose* not to do so. It is therefore difficult to argue that pay changes in the peer groups are exogenous.

Models 3 to 6 include additional predictors for revenue, market cap, employees, assets, financial attributes of the firm, and the percent of the firm's shares held by institutional shareholders. Models 3 and 4 include predictors that should be largely uncorrelated with counterfactual compensation, and models 5 and 6 include the four financial variables that have been used in the counterfactual benchmarking process. These are thus by construction related to counterfactual compensation. The fact that the R^2 of the models increases only slightly when including the four firm financials attests to the fact that our measure of counterfactual peer compensation captures much of the relevant information about the CEO's compensation.

Having a large ownership stake of institutional shareholders is associated with higher executive compensation, net of other factors. The CEOs of larger firms in terms of market capitalization, employees, and assets are paid more, even net of the benchmark of the median counterfactual peer and the median named peer. Importantly, however, CEO pay is tied to the amount of bias in their named peer group even after size, other structural characteristics, firm performance, and

[17] Because of the mathematical equivalence of the log-ratio and log-linear forms of our compensation model, model 1 also implies that, net of the median pay of the counterfactual peer group, each 1% increase in the pay of the median named peer produces a 0.54% increase in the compensation of the focal CEO.

the natural benchmark are taken into account. A change in peer group bias from 10% to 20% for the same firm over time implies an increase of 1.45% in compensation ($72,500 against a $5M compensation baseline), net of any changes in the counterfactual median peer pay, and net of any changes in the structural characteristics of a firm.

Also note that in the full model (model 6), market capitalization is positively associated with compensation while the effect of ROA on pay is positive but not significantly different from zero. Recall that our argument about how weak performance should be associated with more bias in the peer group builds on the idea that CEOs are affected financially by failure to meet performance targets and should therefore find an alternative strategy to justify an improvement in pay. Model 6 suggests that market capitalization is a salient performance metric for CEOs.

In summary, our results support the conclusion that CEOs benefit financially from a positive PPG. Firms that have more bias in their named peer group have CEOs with higher compensation. For the same firm, fluctuations over time in the size of the PPG are associated with changes in the same direction of CEO pay. We do not know whether the PPG produces higher pay through the mechanism of (a) convincing the compensation committee to recommend higher CEO compensation to the full board, (b) convincing board members who are not on the compensation committee that a higher level of compensation is warranted, (c) reducing resistance from stakeholders who otherwise might conclude that the proposed compensation for the firm's CEO is excessive, or some combination of these mechanisms. However, regardless of the specific channels by which the PPG has become associated with higher executive compensation, the results in Table 5 are consistent with the interpretation that CEOs and/or compensation committees establish bias in a peer group in order to legitimize a level of executive compensation that exceeds the level that is predicted by economic determinants.

7. Trends in Bias and their Implications for Compensation

Firms have been under pressure to be more transparent and normative in their procedures for paying their executives. Because the use of biased peer groups is considered to be wrong by compensation consultants and corporate watchdog organizations alike (Reda et al. 2008, ISS 2015), one might expect the PPG to be decreasing over the period of time when the watchdog organizations began considering bias in forming their recommendations on "say on pay" votes. Our data cover the period during which this pressure was being felt, and so we next examine whether there is evidence for a decline in the PPG.

Figure 5, Panel A, shows that the median bias has indeed decreased between 2006 and 2016. Moreover, Panel B shows that the distribution of bias across firms has tightened, both for the entire population in our data (solid lines) and for the subset of firms for which we have data on

Figure 5: The Evolution of Bias Norms

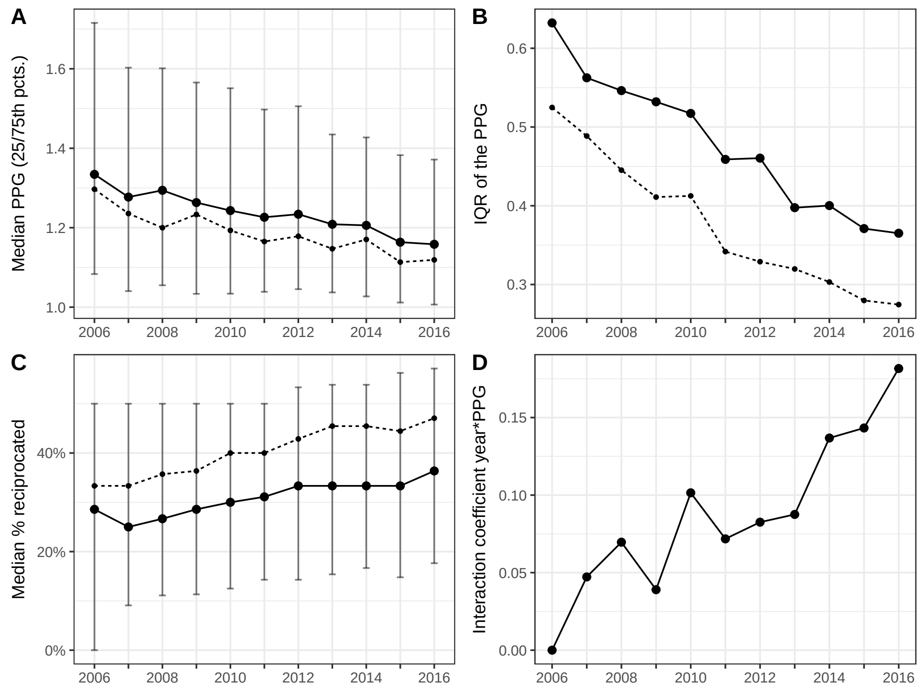

Note: Error bars in Panels A and C report the 25^{th} and 75^{th} percentile. Solid lines report results for all firms. Dashed lines report results for those 327 firms for which we have data for all 11 years (without error bars).

named peers in all years (dashed). These results suggest that the pattern of peer selection of firms is converging towards an ideal type. This is also indicated by a slow increase in reciprocated ties (Panel C), which we take as another piece of evidence in this regard. A reciprocation rate of 100% would completely remove bias from the system of peer nominations, but this, of course, is not a necessary condition for unbiasedness. Similarly, an increase in reciprocated ties does not necessarily suggest a connection to declining bias if the non-reciprocated were chosen specifically to maintain bias. Nonetheless, it is worthy of note that peer group assembly is slowly trending towards an ideal type process.

The patterns in Panels A, B, and C of Figure 5 raise the question of whether the potential impact of biased peer groups on CEO compensation is also diminishing. The answer to this question is negative, because while the average size of the PPG has diminished, the coefficient of bias in the compensation model has been rising over time, and the rise in the coefficient of bias is offsetting the decline in the magnitude of bias. We estimated a model that includes interaction effects between the PPG and dummy variables for each year. The model coefficients are in Table C2 in the appendix, and are plotted in Panel D of Figure 5.

The net effect of this change can be seen in Figure 6, which plots compensation under different scenarios. We restricted this analysis to the firms included in the S&P 1500 to make the numbers

Figure 6: Estimated Impact of Bias and Bias-Year Interaction on Compensation for the S&P 1500

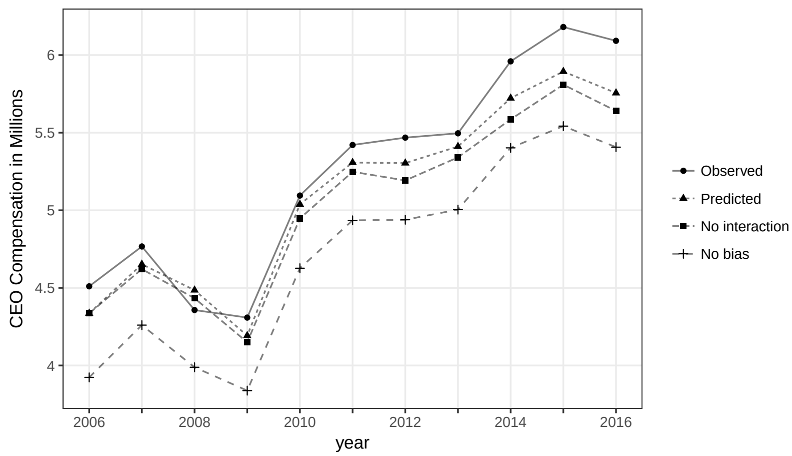

comparable over time. Figure 6 shows that mean CEO compensation fell in the early years of the Great Recession, but then resumed its climb, rising from about $4 million to over $6 million. The solid line shows how the mean CEO compensation has changed over the years covered by our data. The dashed line marked by triangles shows the prediction from our compensation model that includes interaction terms between the PPG and dummy variables for years. The predicted values track the observed values closely, indicating a good model fit. In addition, we show two counterfactual scenarios. The dashed line marked by plus signs shows predicted compensation from our model if bias is set to zero for all firms. It is evident that an important component of CEO compensation is associated with the PPG across all the years. The dashed line marked by squares shows the impact of fixing the coefficient for bias at its 2006 value, when its estimated effect is the smallest. The graph shows that the gap between the predicted line and the "no interaction" scenario gradually increases as the median PPG across the firms declines. This growing gap is the countervailing effect of the increasing coefficient on the PPG.

In summary, bias retains a significant association with compensation throughout the period of time under analysis here and is substantively important despite the downward trend. Moreover, the impact of reductions in the size of the PPG are being offset by increases in the size of the coefficient of the PPG over time.

8. Discussion and Conclusion

In this paper we have scrutinized the assembly of compensation peer groups and its effect on CEO compensation. We leveraged the idea that reciprocated peer nominations are unlikely to result from opportunistic behavior. Specifically, we used the estimates from a model that predicts reciprocated peer nominations to construct counterfactual peer groups.

Using this strategy, we found that while there is a clear tendency towards the selection of similar firms as peers, firms also tend to include in the peer group CEOs who are better compensated than the CEOs of firms that would have been selected had the process been impartial. This produces bias in the named peer group. We also found that benchmarking discretion is generally associated with an increase in bias in the named peer group: Firms that *can* choose a higher benchmarking value to compare themselves to because of ambiguity about the composition of their natural peer group, often do so. Thirdly, we found that the size of the bias is larger when the CEO is unable to cite recent financial achievement as a justification for higher pay. We then found that the size of peer group bias is consequential: greater bias is strongly associated with better compensation for the CEO. Finally, we find evidence that the growing normative pressure on firms to use similar firms as peers may be having an effect on the composition of the named peers because both the mean and the variance in peer group bias has been diminishing in recent years. But this tendency is offset by the increasing effect of peer group bias on CEO compensation.

The evidence presented in this paper challenges the argument that peer group benchmarking enables evaluators to accurately assess whether CEO are compensated appropriately. Instead, we find pervasive bias in peer group benchmarks, and we find that this bias is associated with higher pay for CEOs. The practical impact of our findings is substantial if one considers the arguments put forward in DiPrete et al. (2010). They found that biased compensation practices in a small number of firms will diffuse throughout the corporate comparison network and have substantial downstream effects (i.e. higher compensation across the entire sample). Thus peer group bias has both local effects in the focal firm and global effects through the benchmarking process.

Better market governance has the potential to align pay-setting practices more closely with market mechanisms, and indeed, we find that peer group bias has declined as a presumed consequence of the greater attention to this issue by watchdog agencies. However, the fact that the predictive power of bias on pay has been growing to offset the impact of reduced bias reinforces the message that the process of compensating CEOs is complex and subtle, and strategic behavior is not easily controlled.

Acknowledgments

This research has been supported in part by a grant from the Russell Sage Foundation and by the National Science Foundation under Grant No. SES-1359666.

References

Albuquerque AM, De Franco G, Verdi RS (2013) Peer choice in ceo compensation. *Journal of Financial Economics* 108(1):160–181, ISSN 0304-405X.

Audia P, Brion S, Greve HR (2016) Self-assessment, self-enhancement, and the choice of comparison organizations for evaluating organizational performance. Advances in Strategic Management, Volume on Cognition and Strategy.

Audia PG, Brion S (2007) Reluctant to change: Self-enhancing responses to diverging performance measures. *Organizational Behavior and Human Decision Processes* 102(2):255–269, ISSN 0749-5978.

Bakija J, Cole A, Heim B (2010) Jobs and income growth of top earners and the causes of changing income inequality: Evidence from US tax return data. *Department of Economics Working Papers* .

Barsalou LW (1983) Ad hoc categories. *Memory & Cognition* 11(3):211–227, ISSN 0090-502X, 1532-5946.

Bebchuk LA, Weisbach MS (2010) The state of corporate governance research. *The Review of Financial Studies* 23(3):939–961, ISSN 0893-9454.

Bertrand M, Mullainathan S (2001) Are ceos rewarded for luck? the ones without principals are. *The Quarterly Journal of Economics* 116(3):901–932, ISSN 0033-5533.

Bizjak J, Lemmon M, Nguyen T (2011) Are all ceos above average? an empirical analysis of compensation peer groups and pay design. *Journal of Financial Economics* 100(3):538–555, ISSN 0304-405X.

Bizjak JM, Lemmon ML, Naveen L (2008) Does the use of peer groups contribute to higher pay and less efficient compensation? *Journal of Financial Economics* 90(2):152–168, ISSN 0304-405X.

Cadman B, Carter ME (2013) Compensation peer groups and their relation with ceo pay. *Journal of Management Accounting Research* 26(1):57–82, ISSN 1049-2127.

Cook FW (1981) The compensation director and the board's compensation committee. *Compensation & Benefits Review* 13(2):37.

De Angelis D, Grinstein Y (2015) Performance terms in ceo compensation contracts. *Review of Finance* 19(2):619–651, ISSN 1572-3097.

Dicks DL (2012) Executive compensation and the role for corporate governance regulation. *The Review of Financial Studies* 25(6):1971–2004, ISSN 0893-9454.

DiPrete TA, Eirich GM, Pittinsky M (2010) Compensation benchmarking, leapfrogs, and the surge in executive pay. *American Journal of Sociology* 115(6):1671–1712, ISSN 0002-9602.

Elsbach KD, Kramer RM (1996) Members' responses to organizational identity threats: Encountering and countering the business week rankings. *Administrative Science Quarterly* 41(3):442–476, ISSN 0001-8392.

Faulkender M, Yang J (2010) Inside the black box: The role and composition of compensation peer groups. *Journal of Financial Economics* 96(2):257–270, ISSN 0304-405X.

Faulkender M, Yang J (2013) Is disclosure an effective cleansing mechanism? the dynamics of compensation peer benchmarking. *Review of Financial Studies* 26(3):806–839, ISSN 0893-9454, 1465-7368.

Festinger L (1954) A theory of social comparison processes. *Human Relations* 7:117–140, ISSN 1741-282X(Electronic);0018-7267(Print).

Frydman C, Jenter D (2010) Ceo compensation. *Annual Review of Financial Economics* 2(1):75–102.

Garvey GT, Milbourn TT (2006) Asymmetric benchmarking in compensation: Executives are rewarded for good luck but not penalized for bad. *Journal of Financial Economics* 82(1):197–225, ISSN 0304-405X.

Gulati R, Westphal JD (1999) Cooperative or controlling? the effects of ceo-board relations and the content of interlocks on the formation of joint ventures. *Administrative Science Quarterly* 44(3):473–506, ISSN 0001-8392,.

Iliev P, Lins KV, Miller DP, Roth L (2015) Shareholder voting and corporate governance around the world. *The Review of Financial Studies* 28(8):2167–2202, ISSN 0893-9454.

ISS (2015) 2015 U.S. compensation policies: Frequently asked questions. Technical report, Institutional Shareholder Services, New York, NY.

Jensen MC, Murphy KJ (2010) Ceo incentives-it's not how much you pay, but how*. *Journal of Applied Corporate Finance* 22(1):64–76, ISSN 1745-6622.

Keister LA (2014) The one percent. *Annual Review of Sociology* 40(1):null.

Khurana R (2011) *Searching for a Corporate Savior: The Irrational Quest for Charismatic CEOs* (Princeton University Press), ISBN 1400841097.

Kim JW, Kogut B, Yang JS (2015) Executive compensation, fat cats, and best athletes. *American Sociological Review* 80(2):299–328, ISSN 0003-1224, 1939-8271.

Little RJ, Rubin DB (1989) The analysis of social science data with missing values. *Sociological Methods & Research* 18(2-3):292–326.

Piketty T (2014) *Capital in the Twenty-First Century* (Cambridge Massachusetts: Belknap Press), first edition edition edition, ISBN 9780674430006.

Pittinsky M, DiPrete TA (2013) Peer group ties and executive compensation networks. *Social Science Research* 42(6):1675–1692, ISSN 0049-089X.

Podolny JM (2001) Networks as the pipes and prisms of the market. *American Journal of Sociology* 107(1):33–60.

Porac JF, Wade JB, Pollock TG (1999) Industry categories and the politics of the comparable firm in ceo compensation. *Administrative Science Quarterly* 44(1):112–115.

Reda JF, Reifler S, Thatcher LG (2008) *The Compensation Committee Handbook, Third Edition* (New York: John Wiley and Sons).

Roberts DR (1956) A general theory of executive compensation based on statistically tested propositions. *The Quarterly Journal of Economics* 70(2):270–294.

Shleifer A, Vishny RW (1997) A survey of corporate governance. *The Journal of Finance* 52(2):737–783, ISSN 1540-6261.

Simon HA (1957) The compensation of executives. *Sociometry* 20(1):32–35.

Smith EB, Chae H (2017) The effect of organizational atypicality on reference group selection and performance evaluation. *Organization Science* 28(6):1134–1149, ISSN 1047-7039.

Stern I, Westphal JD (2010) Stealthy footsteps to the boardroom: Executives' backgrounds, sophisticated interpersonal influence behavior, and board appointments. *Administrative Science Quarterly* 55(2):278–319.

Westphal JD, Zajac EJ (1995) Who shall govern? ceo/board power, demographic similarity, and new director selection. *Administrative Science Quarterly* 40(1):60–83, ISSN 0001-8392.

Zuckerman EW (2004) Structural incoherence and stock market activity. *American Sociological Review* 69(3):405–432.

Zuckerman EW (2012) Market efficiency: A sociological perspective. Cetina KK, Preda A, eds., *The Oxford Handbook of the Sociology of Finance*, 223–248 (OUP Oxford), ISBN 978-0-19-164135-0.