Research Article

Bargaining with the CEO: The Case for "Negotiate First, Choose Second"

Russell Korobkin and Michael Dorff

Whether chief executive officers (CEOs) and other senior executives are too highly compensated is one of the most publicized and divisive issues in corporate governance. In this article, we address this question not by asking whether executives are paid more than the value they create, but by asking whether firms could pay executives less money without reducing quality – thus retaining more money for shareholders–by using a better negotiation strategy. The focus of our attention is a particular feature of the way in which the compensation of CEOs and other high-level employees is often determined, although rarely discussed: the firm first decides which candidate it prefers and only then negotiates the amount of compensation with the desired candidate. We hypothesize that this approach to negotiation, which we call "choose first, negotiate second," is inferior to its alternative, which we call "negotiate first, choose second." We explain the theoretical basis for this hypothesis and then present the results of an experiment designed to test it. We conclude by suggesting a number of possible explanations for firms' failure to take advantage of what we consider to be a superior negotiating strategy.

Russell Korobkin is Richard C. Maxwell Professor of Law at the University of California, Los Angeles School of Law. His e-mail address is korobkin@law.ucla.edu.

Michael Dorff is the Michael and Jessica Downer Endowed Professor and director of the Technology Law and Entrepreneurship Program at Southwestern Law School. His e-mail address is mdorff@swlaw.edu.

10.1111/nejo.12242

Key words: negotiation, executive compensation, CEO pay, salary negotiations, employment negotiations, bargaining power, bargaining strategy, corporate governance.

Introduction

Executive compensation is one of the most publicized and divisive issues in corporate governance. Chief executive officers (CEOs) of publicly traded companies, who receive the most attention in the debate, have seen their pay skyrocket over the last several decades, not only in nominal terms, but in real terms, relative to lower level employees, and relative to corporate profits (Murphy 1999; Murphy and Zabojnik 2004; Bebchuk and Grinstein 2005; Frydman and Jenter 2010; Dorff 2014). Some commentators have defended the current pay rates as justified by the value that high-quality executives can bring to a company's bottom line (Anabtawi 2005; Gabaix and Landier 2008; Edmans and Gabaix 2009; Cao and Wang 2013). Many others contend that CEO performance does not justify either the high absolute amount of current compensation packages or the enormous increases in the value of those packages over time (Barris 1992; Crystal 1992; Bogus 1993; Bok 1993; Elson 1993; Yablon 1999; Bebchuk and Fried 2004; Dorff 2014).

We examine the issue of the appropriateness of executive compensation from a different perspective by asking this question: regardless of whether CEOs (or other executives) create more marginal value for their companies than the cost of their compensation packages, could those companies pay their CEOs less money without reducing CEO quality, thus retaining more money for shareholders, for compensating lower-level employees, or for investment? The focus of our attention is a particular feature of how CEO compensation, as well as the compensation of other high-level employees, is often determined that has been rarely discussed: the firm first decides which candidate it prefers and only then negotiates the amount of compensation with the desired candidate (Wackerle 2001).

This bargaining process, which we call "choose first, negotiate second," or "C1N2," stands in sharp contrast to the way firms hire for most lower-paid positions. For more typical employment categories, the firm is more likely to advertise a wage or salary when recruiting for the position, and both the firm and the job applicants share the understanding that the candidate will work for that amount if hired. We call this approach, "negotiate first, choose second," or "N1C2." These bargaining approaches are actually more like points on opposite ends of a spectrum than they are truly dichotomous, and we acknowledge that many cases fall in the intermediate range of the spectrum. Firms might, for example, advertise a position with a compensation range, suggesting that

there will be some, but limited, post-hiring negotiation, perhaps based on the candidate's current salary or level of experience.

Our hypothesis is that firms could hire CEOs (and other executives) at lower compensation levels than they currently pay without reducing the quality of their executive talent by incorporating elements of the N1C2 approach into their senior level hiring process, rather than waiting to address compensation until they have identified their preferred candidate. We describe the theoretical basis of our hypothesis, provide the results of an experiment we conducted to test the hypothesis, and consider a number of possible explanations for the failure of corporate boards to take advantage of the most theoretically advantageous negotiating strategy for hiring CEOs.

The Conventional Argument about CEO Compensation

The compensation of CEOs has drastically increased in the United States in recent decades. From the 1940s through the mid-1970s, CEO pay at large U.S. public corporations was generally stable, at about $1.4 million per year in 2015 dollars (Frydman and Jenter 2010: 85). In 2014, the median pay at similar companies was $13.6 million in nominal dollars (e.g., then current dollars, unadjusted for inflation, inclusive of performance shares, bonuses, stock options, etc.) (Lublin 2015), and the highest-paid CEO, David Zaslav of Discovery Communications, earned $156 million (Gelles 2015).

Whether ever-more-highly compensated CEOs are overpaid is subject to considerable dispute. Proponents of "rational choice theory" in law and in economics argue that, if the market provides CEOs with extremely rich compensation packages, offering such packages *must* be in the interests of the firms that employ CEOs (Chang, Dasgupta, and Hilary 2010; Cao and Wang 2013; Falato, Li, and Milbourn 2014). Implicit in this claim is the conjecture that the amount by which the marginal revenue product of a highly-paid CEO exceeds his or her cost (the "net marginal revenue product") is greater than the net marginal revenue product of alternative candidates for the CEO position (Thomas 2004). This conjecture is correct, by definition, if the firms that hire CEOs are perfect optimizers of shareholder wealth. The hiring of CEOs and other senior executives is a context in which many observers believe actual practice closely approximates optimization (Dorff 2014).

Critics of this conjecture contend that CEOs are often paid more than their marginal product as a result of imperfectly rational behavior by corporate directors who have difficulty predicting a CEO's future marginal revenue product (Dorff 2014), self-interested behavior on the part of directors who wish to curry favor with CEOs (Elson 1993; Bebchuk and Fried 2004), or tax and accounting rules that encourage

boards to use inefficient pay structures, such as stock options (Murphy 2002; Murphy 2013; Dorff 2014).

The "Bargaining Process" Hypothesis

Our purpose is to provide a method for evaluating the appropriateness of executive compensation that takes no side in the debate over whether firms pay CEOs more than their marginal product. We contend that firms could pay CEOs less than they do without reducing shareholding wealth, *even if it is true* that CEOs provide marginal firm value that is at least equivalent to their marginal compensation. We believe that firms could reduce executive pay by implementing a simple type of process change in the way they negotiate compensation: rather than using the dominant "choose first, negotiate second" (C1N2) process, firms should employ a "negotiate first, choose second" (N1C2) process. Specifically, at some point in the selection process used to winnow the pool of candidates from a large group to just one, the firm should require each candidate to reach agreement on a compensation package that the candidate will receive if selected. We call our conjecture that this approach will lead to lower CEO salaries without reducing CEO quality the "bargaining process" hypothesis.

The literature on executive compensation is vast and often contentious – indeed, one of us has written a book that canvases this literature (Dorff 2014) – but we have not identified any other scholars who have advanced in print the argument that we make here. Why is this? We think the answer is that both sides of the standard debate over CEO compensation make an implicit assumption about the market for executive talent that we find implausible.

The literature generally assumes that firms are price takers (i.e., they are able to contract, if at all, only at a price set by the market over which they have no control) in a thick, undifferentiated market, while simultaneously assuming that CEO-caliber talent is in short supply (Rosen 1981; Rosen 1982; Gabaix and Landier 2008; Terviö 2008). It follows from this assumption that a firm must bid its reservation price (that is, the maximum it is willing to pay) for its first-choice candidate because, if it bids less, the candidate will be snapped up by another CEO-starved company fishing in the same limited talent pool. Adherents of the rational choice conjecture thus seem to assume that CEO candidates wield near-monopoly power over firms, which allows them to capture almost all of their surplus value. In this world, the only relevant question concerning whether firms "overpay" CEOs is whether firms correctly estimate the marginal productivity of CEO candidates.

We believe, however, that CEO recruitment often has elements of bilateral monopoly. Candidates have differential qualities and abilities

(although firms may believe these are larger than they actually are) (Khurana 2002), so firms usually lack a perfect substitute for their top choice. But it is also the case that firms with an open CEO position will often be uniquely desirable to candidates. Every CEO position offers different opportunities, challenges, working conditions, geographic location, etc., and CEO openings are limited in number, which means a candidate is likely to be under consideration for only one such position at a given moment in time. As a result, while a particular CEO candidate might be a firm's uniquely best option, a particular CEO position might also be a candidate's uniquely best option.

In this situation, any compensation level that falls between the parties' reservation prices will be Pareto-efficient, and there is no economic justification for the resulting salary to be set at the firm's reservation price as opposed to any other point within the bargaining zone. Thus, simply determining that the CEO's marginal product exceeds his or her marginal compensation is not enough to confirm that the firm is not paying its CEO more than is in the shareholders' best interests. It is also necessary to ask whether, as a consequence of the bargaining process that it chooses to employ, the firm is paying its CEO more than it needs to in order to obtain the level of quality represented by that candidate. If we resolve this inquiry in the affirmative, it would be, in our view, proper to assert the normative conclusion that the firm has "overpaid" its CEO, even if the net marginal revenue product is positive and larger than would be expected if the firm hired an alternate executive.

We believe, for both rational and behavioral reasons, that firms using the C1N2 bargaining process will overpay their CEOs relative to what they would be able to pay if they were to employ an N1C2 approach.

Game Theory and the Concept of Bargaining Power

In negotiation, a party's leverage, or "power" (Korobkin 2014), depends on beliefs about beliefs. Specifically, bargaining power depends on a party's perception of the subjective value the counterparty places on its best alternative to reaching a deal, often referred to as its "BATNA" (best alternative to a negotiated agreement) (Fisher, Ury, and Patton 1991). This is because a negotiator with a desirable BATNA will have a low reservation price, defined as the point at which she would be indifferent between reaching a deal and pursuing an alternative. When negotiating with a CEO candidate, a firm's BATNA will usually be to hire a different candidate for the position. The firm maximizes its negotiating power if the candidate believes that the firm has a low reservation price: i.e., that it will accept an impasse rather than reach an agreement that requires it to pay more than a relatively low salary amount. The candidate will

believe the firm has a low reservation price, in turn, if the candidate believes the firm places a relatively high subjective value on hiring an alternative candidate.

By identifying the most desirable candidate first and negotiating the compensation package with that candidate second, the firm sacrifices bargaining power in two different ways. First, the firm reveals to the candidate that, holding compensation constant, the firm believes that the candidate in question would be more desirable than any other candidate. This implies that the firm has a higher reservation price than the candidate is likely to infer on the part of the firm if it is unclear which of several candidates the firm would prefer.

Second, the act of offering the position to one candidate and embarking on compensation negotiations can itself create reputational costs should the firm later turn to its second-choice candidate. In other words, employing a C1N2 process reduces the value of the firm's BATNA, thus increasing its reservation price in its negotiation with the candidate. If the preferred candidate is publicly identified as such before compensation is negotiated, or if the identity of the preferred candidate becomes public as a result of an information "leak," several negative consequences can ensue. Public perceptions that the firm was unable to attract its top candidate might be seen by investors as a negative signal concerning the firm's future prospects. Dignitary harm (i.e., the suffering of an indignity) caused to the second-choice candidate could make it harder (and more expensive) for the firm to then hire that candidate. The time lag between the selection of the first-choice candidate and the failure of subsequent negotiations can make the firm more impatient to hire the second-choice CEO to avoid further inferences that the firm is having trouble attracting the CEO or that the ultimate hire was not highly desired, which might force the firm to offer a more generous compensation package to the second-choice candidate than it otherwise would. For any or all of these reasons, a candidate could reasonably infer that, once the firm anoints him or her as its top candidate, the firm's reservation price will increase, thus weakening the firm's bargaining position.

As suggested above, the higher the candidate believes the firm's reservation price to be, the greater the candidate's bargaining power, and the higher the salary the candidate will likely be able to negotiate. For example, assume that the candidate's reservation price is $2 million per year and that the firm's reservation price for hiring that candidate is $10 million dollars per year. We would expect that, on average, the firm will be able to negotiate a lower salary if the candidate believes the firm's reservation price is only $5 million than if the candidate believes the firm's reservation price is $10 million.

Commitment and Consistency

Social scientists have long understood that human beings have a strong desire to act consistently with their past commitments. Failure to do so creates cognitive dissonance, which is uncomfortable (Festinger 1957).

By offering a candidate the CEO position, the firm's directors make a clear commitment to the prediction that the firm will prosper under that candidate's leadership. Acting consistently with this commitment requires treating the candidate as if she were not only adequate, but outstanding. Offering a candidate a relatively low salary in comparison to similarly situated CEOs, which might seem like a good negotiating strategy if the candidate appears to have a relatively low reservation price, could seem inconsistent with the firm's determination that the candidate is of high quality. Confidence in the ability of the candidate would more likely suggest the appropriateness of a high salary compared to peers because a high-quality candidate should, in theory, have many other excellent career opportunities and, therefore, have a high reservation price. In addition, the firm's evaluation of the candidate's high quality is salient, but the fact that similarly situated CEOs were also judged by their boards of directors to be of higher quality than their competitors for the position is unlikely to be salient. This type of thinking can create a spiraling effect, where each CEO hired is considered "above average" and thus offered an above-average salary compared to her peers, and salaries are continually ratcheted up (Hayes and Schaefer 2009; Bizjak, Lemmon, and Nguyen 2011; Elson and Ferrere 2013).

The psychological commitment to consistency can also cause the firm to escalate its compensation offers to avoid the risk of an impasse in negotiations, which would force it to seek another candidate. Once a firm's directors identify their top candidate for the CEO position, consistency suggests that they succeed in hiring that candidate. A potential consequence is that the directors adopt a mindset of "doing whatever it takes" to reach agreement with that candidate, which can result in greater concessions than the directors would have made had they not felt committed to a candidate.

We think that if firms negotiate salaries prior to choosing CEOs, they will suffer from less cognitive pressure to offer those candidates salaries that are higher than average compared to the CEOs of other companies. At the time of the negotiation, of course, the candidate would not be the company's CEO, only a candidate for the position, and the firm will not have a cognitive stake in that candidate's conditional compensation reflecting above-average quality.

Fairness Norms

When a bargaining zone exists between the candidate's reservation price and the firm's reservation price, the specific outcome will often have as much to do with what the parties believe is a fair salary as with the parties' relative bargaining power (Korobkin 2014). The concept of fairness, however, lacks a precise metric, and much depends on which of many comparisons seem to the negotiators to be most relevant in a particular context. We believe that subtle contextual clues are more likely to suggest the appropriateness of a higher salary after a candidate is chosen than before.

After the candidate is offered the CEO position, we hypothesize that parties are more likely to see the salary of existing CEOs of companies with similar features (such as firms in the same or similar industries and with similar revenues or profits) as an appropriate reference point, and that it will be more difficult for the parties to agree that a salary that deviates substantially from those comparison points is fair. Such peer group comparisons are a mainstay of executive compensation negotiations (Elson and Ferrere 2013). Before a candidate is offered the position, however, we think that other reference points – such as the candidate's current salary or the salaries of the candidate's current peers – can also have the patina of legitimacy. This greater range of plausible reference points could make somewhat lower salaries appear to be fair in this circumstance.

The Experiment

The Negotiation Simulation

To test our hypothesis that a C1N2 bargaining process will cause firms to overpay their CEOs, we conducted a high-context negotiation simulation, entitled "Hiring a CEO," in which subjects played either the role of an executive under consideration for a CEO position or the role of a corporate official responsible for hiring a new CEO.

Subjects were randomly assigned to one of four roles:

- The chair of the board of directors of Bartleby Manufacturing., Inc. ("the director"), a corporation that manufactures parts for agricultural machinery and aircraft;
- Quinn Morris, vice president and chief operating officer of Parts Manufacturing, Inc.;
- Sidney Murphy, vice president and chief operating officer of Agricultural Assemblers, Inc.; or
- Casey Morgan, vice president and chief operating officer of Amazing Aircraft, Inc.

All subjects were informed that Bartleby recently announced the retirement of its current CEO, Jamie Miller, and that the director was responsible for hiring Miller's replacement. After a series of initial interviews, the director had narrowed the field to three finalists for the position – Morris, Murphy, and Morgan – and was tasked with hiring one of the three at an agreed-upon salary.

Each participant was randomly assigned to play the role of the director or one of the candidates. Candidate subjects and director subjects were prepped in separate rooms. An experimenter read background information and simulation instructions aloud, while subjects followed along with individual sets of written instructions. After being prepped and given an opportunity to ask questions, candidates and directors were brought together for face-to-face meetings (except in one experimental manipulation, explained below).

Subjects playing the role of the director received biographical information for each candidate that described that candidate's background and qualifications for the job. All three candidates boasted many years of experience as chief operating officer (COO) of a company in a related industry, along with bachelors' and MBA degrees from different but similarly elite institutions (i.e., Harvard University, Stanford University, and the University of Chicago). The biographies were designed with the goal of making all three candidates appear equally qualified for the Bartleby job, but for the sake of realism each biography was different. Directors learned that all candidates currently earned an annual salary of approximately $3 million in their COO positions. They also received a short, confidential report, purportedly prepared by a consulting firm retained by the company, that carefully compared the candidates and concluded that all three were equally well qualified for the CEO position: "None of the candidates would be meaningfully better than the others."

Each candidate received the biographical information about himself/herself and was told that she or he was competing against two other finalist candidates for the CEO position. The candidates did not receive any specific information about the other two finalists, but all were told they should assume that the others had similarly impressive qualifications. All candidates learned that they earned a salary of $3 million per year in their current job, but none had any information about the current salaries of their competitors. All were told they were very interested in the Bartleby CEO position because of the personal challenge that running a large company would present, its likely potential for higher income than they currently enjoyed, and because (for various reasons) it was unlikely, although not impossible, that they would become CEO of their current company.

All director and candidate subjects received a "CEO Salary Table" that listed Bartleby and eight other companies, described their industries (all were similar to Bartleby's) and provided the recent annual sales, annual profits, profit margin, and market capitalization of each company, the length of tenure of each company's CEO, and each CEO's current annual compensation. Bartleby's outgoing CEO, Jamie Miller, was reported to be in his/her eighth year of service at a current annual salary of $13.86 million. The statistics for each company varied, providing material that subjects could use to support fairness claims for various salary levels based on the principle of horizontal equity, but the data did not point to an obvious reference salary for the new Bartleby CEO. All subjects, candidates and directors alike, were informed that the table contained publicly available information to which all subjects would have access.

All subjects were informed that, for a CEO to be hired, the director, and one of the three candidates would have to agree to an amount of annual compensation. Although CEO compensation is often divided between salary, bonuses, stock options, etc., subjects were told that, for simplicity, parties in the simulation would need to agree to a single dollar amount of compensation, to be referred to as "salary," which they should assume would be divided between fixed pay and other forms of compensation after the simulation was completed.

Participants

Two-hundred thirty-four first-year law students – 146 from the University of California at Los Angeles and eighty-eight from Southwestern Law School in Los Angeles – were recruited to participate in the simulation for cash compensation. The compensation arrangement, explained at the outset of the simulation, had two important features. First, to encourage all subjects to put forward a level of effort reflective of what might be expected in the real world, payments were incentive-compatible: subjects who achieved better negotiated outcomes received more money. Second, the correlation between outcomes and cash payments was greater for candidate subjects than for director subjects, reflecting the fact that CEOs will receive more direct cash benefit from negotiating higher salaries than members of the board of directors are likely to receive by negotiating lower CEO salaries for the benefit of shareholders.

All candidate subjects earned a $7 flat fee. Candidates who were hired to be the new CEO of Bartleby earned an additional $1 in cash for every $1 million dollars in annual salary that they were able to negotiate. Candidates who were not hired as CEO kept their current COO positions and earned an additional $1 in cash for every $1 million of their current salary. (Recall that each candidate's "current salary"

was $3 million, and thus unsuccessful candidates earned $7 + $3 for a total of $10.) For clarity, candidates were told that if they wished to maximize their cash earnings from the experiment, they should not accept the CEO position at a salary lower than the salary they currently earned. Candidates knew that they shared the same payoff structure with the other candidates, but they did not know their competitors' current salaries.

Director subjects were told that they earned a (hypothetical) salary of $1.2 million as chair of the board of directors and would receive $1 in real money for each $100,000 of that imaginary salary for their participation in the simulation ($12). In addition, examiners told them that Bartleby had given them an annual salary pool of $25 million, some or all of which – but no more – could be used to pay the new CEO's annual salary. Because they were instructed to keep executive salaries as low as possible consistent with hiring top-quality talent, directors were told they would earn an additional $0.25 in real money for every million dollars that remained in the salary pool after subtracting out the amount they agreed to pay the new CEO.

These instructions and limitations made it possible for candidate subjects to earn no less than $10 (assuming they did not agree to accept the CEO position for less than their current $3 million salary) and no more than $32 (although they did not know the upper boundary prior to the simulation). Director subjects could earn no less than $12 and no more than $18.25 in real money.

To ensure that all subjects understood the incentive system, before the simulation began we required that each complete a worksheet demonstrating an understanding of how the cash compensation would be calculated. In an attempt to avoid anchoring effects (Kahneman 1992; Kahneman and Tversky 1995), we asked the subjects to calculate two payments, one based on a low CEO salary of $3 million and one based on a high CEO salary of $22 million. The vast majority of subjects provided the correct answers to both questions on the first try. Those who made errors were asked to try again. All subjects provided the correct answers by this point and were permitted to continue with the simulation.

Experimental Conditions

We conducted three variations of the "Hiring the CEO" simulation. In all three, subjects were randomly assigned to play either the role of a director or a candidate and then were randomly matched to one another.

Control Condition: Choose First

The control version modeled the C1N2 process that is typical in the hiring of CEOs and other senior executives. In this version, "Choose First," after the director subjects finished reading the complete set of instructions, we gave them ten minutes to select their preferred candidate from among the three finalists based on the candidate biographies and the report from the consulting firm. After directors chose a CEO, the experimenters informed the candidates of the decision.[1] Directors and their chosen candidates then had fifteen minutes to prepare to negotiate with one another, and then another fifteen minutes to attempt to reach a salary agreement in a face-to-face, free-form negotiation session.

All subjects were told that if they failed to reach agreement, the candidate would keep his or her current job and Bartleby would have to reopen its CEO search, but neither the director nor the candidate would be a part of that process. In the event of an impasse, candidates would be paid $10 in real money ($7 show up fee plus $3 for their $3 million current salary) and directors would be paid $12 ($1 for every $100,000 of their $1.2 million salary and no bonus). Subjects did not know the payoff structure for subjects playing the other role in case of impasse. By the end of the time period, each director/candidate pair of subjects recorded their salary agreement or indicated that they failed to reach an agreement.

First Experimental Condition: Multiple Negotiation

In our first experimental condition, "Multiple Negotiation," director subjects were instructed to employ a N1C2 process by negotiating potential salaries with all three candidate finalists coincident with making the final hiring decision. All subjects were given fifteen minutes to prepare for negotiations, and directors then had a total of thirty minutes to negotiate individually and serially with each of their three candidates. Directors could meet with each candidate as many times as they chose, for as long or as short a time period as they chose, in whatever order they chose, so long as they met at least once with each candidate. Directors could not meet with more than one candidate at a time, and candidates could not communicate with each other.

Within the thirty-minute negotiation period, directors could enter into an agreement to hire one of the three candidates at an agreed-upon annual salary. As in Choose First, subjects were told that if no agreement were reached, all candidates would keep their current jobs and Bartleby would have to reopen its CEO search. At the end of the negotiation period, directors submitted a form that indicated the lowest salary each of the three candidates had agreed that they would be willing to accept if hired, and identified which candidate, if any, they had agreed to hire.

We predicted that, as a result of the processes employed in this experimental version, candidates would indicate a willingness to accept lower salaries and directors would thus hire a CEO at a lower salary than in Choose First.

In Choose First, the candidate selected as the CEO designate had no way of knowing whether the director preferred him or her strongly or only weakly, but she or he did know that, at worst, the director would not have a higher reservation price for a different candidate. The CEO designate also knew that, in the event of impasse, the director would have to reopen the selection process at some cost to Bartleby.

In Multiple Negotiation, the director could manipulate candidates' perceptions of their relative rankings and use a lower bid from one candidate as leverage against the other two. None of the candidates knew whether they were the director's first choice for the CEO position, holding salary equal, and the director could switch "favorites" at any point without suffering any costs. This meant that the director could tell any of the candidates that Bartleby had no preference between the finalists and would choose the low bidder, that Bartleby preferred the candidate but another candidate had submitted a sufficiently low bid that the cost/quality combination favored that other candidate, or even that that candidate was the least favored candidate but close enough in perceived quality that a sufficiently low bid could switch the cost/quality combination in his or her favor.

Although any of these claims might be false, no candidate could know that. A candidate who knew the firm strongly preferred him or her over the alternatives could safely make and maintain a salary demand close to his or her perceived marginal product, knowing that the director was unlikely to choose an alternative candidate unless the difference in salaries outweighed the company's strong selection preference. The director's bargaining power comes from the assumption – that we believe a reasonable candidate would make – that, at least in some cases, the firm's preference for one candidate among the finalists would be weak.

We also predicted that psychological biases could skew salaries higher in Choose First. Commitment to hiring the candidate who had been selected as the most desirable, a heightened desire to avoid "losing" the designated CEO, and the potentially greater salience of the salaries of other CEOs as anchor points could all encourage the directors to find higher salaries more appropriate than in Multiple Negotiation.

Second Experimental Condition: Solicited Salary

Our second experimental condition, "Solicited Salary," used a different type of N1C2 bargaining process. We told the three candidates that they

were each one of three finalists for the CEO position and that, as the final step in the hiring process, Bartleby required that they each indicate in writing the minimum salary that they would be willing to accept as CEO if selected for the position. By submitting their minimum salary requirements, the candidates would indicate their agreement to accept the CEO position at that salary. The director would then decide which finalist to hire, if any, based on the combination of the firm's evaluation of each candidate's qualifications and the salary cost of hiring each candidate.

After briefly greeting the director to whom they had been assigned but not speaking about the CEO position,[2] candidate subjects were allowed ten minutes to determine their minimum salary requirement and record it. The experimenters then delivered each candidate's salary requirement to his or her respective director. Finally, director subjects selected one of the three candidates for the CEO position, at the salary indicated by the chosen candidate.

In Solicited Salary, any candidate-provided minimum salary bid higher than $3 million would create the risk that the candidate would lose out on the job that she could have secured had the bid been lower. This risk would have to be balanced against the potential gain that would be obtained if the candidate were hired at the higher salary, but the presence of that risk should create pressure to moderate the salary demand. As in Multiple Negotiation, we hypothesized that this bargaining process would exert greater downward pressure on salaries than Choose First, because the candidates would be uncertain about which candidate the director favored, uncertain about how much the firm preferred its top candidate to the others, and uncertain about whether the firm would hire a less-preferred candidate (quality-wise) if that candidate turned out to be most desirable when salary demands are taken into account.

From a rational choice perspective, we would expect identical results in the two experimental conditions. In both, candidates faced the same analytical problem: they had to balance the desire to demand a higher salary in order to increase the payoff they would receive if chosen to be CEO with the incentive to demand less in order to increase the likelihood of being selected. In Multiple Negotiation, directors had opportunities to warn the candidates that they would not be selected if they did not indicate a willingness to accept a relatively low salary that were unavailable in Solicited Salary. Any such warning, however, would be strictly "cheap talk": a rational director would issue the warning whether or not a reduction in salary was actually necessary for a candidate to be selected and, knowing this, a rational candidate would

make the same minimum demand regardless of whether the director claimed a lower demand was necessary in order for the candidate to be selected.

Results

Multiple Negotiation

Our primary interest is in the comparison of the salaries agreed upon by subjects in the control and experimental conditions. In Choose First, all twenty-nine directors reached an agreement with their chosen candidates in the allotted time period. The agreed upon salaries ranged from a low of $4.3 million to a high of $14 million, with an average of $8.5 million.

In Multiple Negotiation, directors negotiated simultaneously with all three finalist candidates with the goal of reaching agreement with one. All twenty-three directors reached a salary agreement with one candidate during the time period. The average agreement was for a salary of $6.56 million, or nearly $2 million *less* than the average salary negotiated in Choose First. The difference between the averages of the two conditions is highly significant ($p < .001$),[3] suggesting that simultaneously negotiating with three finalists before making a choice enabled the directors to save their (hypothetical) companies a substantial amount of money (Table One).

Twenty of the directors in Multiple Negotiation reached agreement with the candidate who indicated to the director his/her willingness to accept the lowest salary. We believe that this result is normative in the simulation in light of the consultant's report provided to the directors predicting that the company would achieve the same level of profitability under each candidate. It is also the outcome that maximizes the actual monetary compensation of the director subjects.

Table One
Salary Agreements (Choose First versus Multiple Negotiation)

	Choose First	Multiple Negotiation
Number of agreements	29	23
Percentage of teams that reached agreement	100	100
Low salary	$4.3 MM	$4 MM
High salary	$14 MM	$9 MM
Average salary	$8.5 MM	$6.56 MM

Three directors reached agreements with one candidate even though a different candidate had indicated a willingness to accept a lower salary. In all three cases, the director subjects responded to a follow-up question as to why they chose not to hire the candidate willing to work for the lowest salary by explaining that he or she believed that the candidate hired was significantly superior in quality to the low bidder. These three subjects were willing to sacrifice a portion of their "real money" payoff in the experiment to vindicate their ethical role to act in the best interest of the fictional firm, Bartleby Manufacturing, Inc. Had these three directors hired the candidate willing to accept the lowest salary among the three, the average salary would have been $6.24 million, or $2.39 million less than the average salary in Choose First.

The experimental literature on negotiation suggests that outcomes of single-issue, bilateral negotiations can depend significantly on which party makes the first offer (Galinsky and Mussweiler 2001). Although the bargaining context and availability of information can affect whether making the first offer is an advantage or a disadvantage, we predicted that, in our simulation, the party making the first offer or demand would gain an advantage by setting an anchor for the subsequent "negotiation dance" – the exchange of several offers and counteroffers of increasing convergence that often precedes an agreement, especially in single-issue negotiations (Raiffa 1982). Consequently, we believed salaries would be lower when the director makes the first offer than when the candidate makes the first demand. The presence of such an effect would not present a problem for interpreting our results if directors were equally likely to make the first offer in both conditions, but we feared that the difference in negotiation processes might differentially affect the likelihood that directors would make the first offer across the two conditions. This could cause what appears to be a difference in outcomes resulting from the difference in structure between C1N2 and N1C2 to actually result from the differential likelihood of directors making the first offer under the different structures.

In light of this concern, we examined the difference between the two conditions controlling for the identity of the party that made the first offer. In Choose First, we feared that the custom of employers making a first salary offer to a selected candidate might cause most or all of the directors to make the first offer. To avoid this, we advised half of the directors to make the first offer in the negotiation (and their paired candidates to refrain from making a salary demand until after the director had made an offer), and we advised the other half of the directors to refrain from making a salary offer until after their candidate had made a demand (and their paired candidates to make the first demand). Hypothesizing that there would not be a strong default assumption

concerning who should make the first offer in the experimental condition, we provided no instruction in that condition as to which party should make the first offer.

In Choose First, fifteen directors reported that they made the first offer, and fourteen reported that their candidate made the first demand (all twenty-nine reports were confirmed by the paired candidate). In Multiple Negotiation, eight directors reported that they had made the first offer to the candidate that they eventually hired, while fifteen reported that the eventually chosen candidate had made the first demand (Table Two).[4]

In our regression model that includes both the experimental condition and the identity of the party who made the first offer, the identity of the first offeror is significant ($p = .03$) – that is, making the first offer led to more advantageous outcomes, providing support for the existing literature on the first-offeror advantage. But, more importantly, our main effect (i.e., whether the subjects participated in Choose First or Multiple Negotiation) remains highly significant independent of which party made the first offer ($p < .01$).

Our examination of which party made the first offer yielded an interesting secondary finding. In Choose First, when directors made the first offer, the final agreement averaged $7.62 million, compared to $9.45 million when the candidate made the first demand. The average difference of $1.83 million is significant ($p = .03$). In Multiple Negotiation, final salaries were also lower on average when directors made the first offer, consistent with our prediction, but the difference was much smaller ($430,000) and not statistically significant. Thus, making the first offer provided a substantial advantage in Choose First, but far less advantage, if any at all, in Multiple Negotiation.

This difference in the value of making the first offer might be explained with the following intuition: the power of setting the initial anchor in negotiations matters more in a bilateral monopoly setting in which, after the director chooses the candidate, the two have roughly equal bargaining power. In contrast, in the N1C2 setting of Multiple Negotiation, the director can pit the three candidates against each other,

Table Two
First Offer (Choose First versus Multiple Negotiation)

Who Made the First Offer?	Director	Candidate
Choose First	15 (51.7%)	14 (48.3%)
Multiple Negotiation	8 (34.8%)	15 (65.2%)

giving the director a greater ability to negotiate down the salary even of a candidate who sets a high initial anchor point. This power advantage reduces the leverage of making the first offer.

Solicited Salary

In Solicited Salary, all directors successfully hired one of the three candidates (as was also the case in the other two conditions). The salaries of the twenty-one new CEOs averaged $7.75 million (ranging from a low of $4 million to a high of $13 million), nearly $1 million less than the Choose First average, but the difference between these means falls short of statistical significance (Table Three).

Four out of the twenty-one directors chose to hire a candidate who did not submit the lowest salary requirement, notwithstanding that this reduced their real-money compensation from the experiment. (All four of these directors answered a follow-up question by explaining that they believed their chosen candidate was the best qualified and that value justified a slightly higher salary.) Had these four directors chosen the lowest bidders among the three candidates, as they could have, the average salary would have fallen to $7.32 million. The difference between the lowest salaries bid by the Solicited Salary candidates and the salaries negotiated by Choose First candidates is significant ($p = .04$). This suggests that the N1C2 mechanism used in the Solicited Salary manipulation made it possible for directors to hire an equally desirable candidate at a lower salary than in Choose First. Clearly, however, this implementation of an N1C2 hiring process did not have the effect of reducing salaries as much as the process tested in Multiple Negotiation.

What explains the difference between the Multiple Negotiation and Solicited Salary results? One hypothesis is that Solicited Salary candidates make surprisingly high salary demands as a consequence of overoptimism about their quality relative to that of their competitors for the

Table Three
Salary Agreements (Choose First versus Solicited Salary)

	Choose First	Solicited Salary
No. of agreements	29	21
Percent Reached agreement	100	100
Low salary	$4.3 M	$4 M
High salary	$14 M	$13 M
Average salary	$8.5 M	$7.75 M

CEO position, a well-established form of "self-serving bias" (Williams 2014).

A rational candidate will submit a high minimum salary requirement if he or she knows that the director heavily favors him or her. This is because, if the director believes that Bartleby will earn tens of millions of dollars more in annual profit under the leadership of a particular candidate, it will be rational for the director to choose that candidate even if he or she demands several million dollars more than the others, and the higher demand allows the candidate to capture more of the available cooperative surplus than will a lower demand. On the other hand, a rational candidate will submit a bid that is much lower if he or she believes that the director is indifferent between the candidates, favors one of the other two candidates, or only slightly favors the candidate in question. In this situation, a candidate who does not bid at or below the level of the other candidates is unlikely to be hired, in which case he or she will not, of course, enjoy any cooperative surplus.

In the "Hiring a CEO" simulation, the candidates have no reliable information concerning whether they will be the preferred candidate or whether any preference the director might have is strong or weak. The candidates know that they are well-qualified for the CEO position, but they have no specific knowledge of their competitors' qualifications, and the instructions advise them to assume that the other finalists are also well-qualified. We believe that this accurately reflects the relative information candidates will usually have in the real world.

After the candidates in Solicited Salary provided their minimum salary requirements, but before the directors made their selection, we asked all candidates to estimate the percentage likelihood that they would be hired if they and their two competitors all had submitted exactly the same minimum required annual salary. Given the candidates' complete lack of comparative information, we believe that the normative answer for candidates in any of the three roles is 33.3 percent. But a large social science literature documents that, on average, individuals tend to overestimate their skills relative to their peer group (Alicke and Goverun 2005) at least in areas in which they possess some skill in an absolute sense (Moore 2007), and that they tend to be more optimistic about the likelihood of experiencing positive events in the future than facts warrant (Sharot 2011).

Our Solicited Salary candidate subjects estimated their likelihood of being chosen as CEO if they and their two competitors submitted identical minimum salary demands at an average of 61 percent, obviously significantly different from the normative response of 33.3 percent ($p < .001$), suggesting the presence of a strong self-serving bias concerning their job qualifications. Of the sixty-three subjects playing

the role of any of the three candidates in Solicited Salary, only nine provided the normative response of 33.3 percent or 33 percent, and only five others provided "pessimistic" estimates of less than 33 percent. The level of bias varied little, and non significantly, across roles (e.g., it was independent of whether the subject was playing the part of Quinn Morris, Sidney Murphy, or Casey Morgan).

This evidence suggests that overoptimism could have caused Solicited Salary candidates to make higher salary demands than they would have if they had held more reasonable beliefs about their likelihood of being the preferred candidate. The Multiple Negotiation process, in contrast, could have enabled directors to undermine candidates' initial confidence that they were the most preferred candidate. But our results also suggest that the Multiple Negotiation bargaining process did not disabuse candidates of self-serving bias.

We also asked Multiple Negotiation candidates, "How likely do you believe you would have been selected for the position if all three candidates made exactly the same salary demand?" The Multiple Negotiation subjects, however, responded to the question after they had learned whether or not they had been hired. The twenty-three candidates hired in Multiple Negotiation responded, on average, that there was a 66 percent likelihood they would have been hired if all three candidates had demanded the same salary. Perhaps this "optimistic" evaluation was justified in light of their knowledge that they had actually been chosen from among the three finalists to be CEO. But the forty-six candidates who were not hired (and knew this fact!) in Multiple Negotiation provided the statistically indistinguishable average response that there was a 63 percent likelihood that they would have been hired if all three candidates had demanded the same salary (Table Four).

Although it struck us as rather shocking that failure to obtain the job would not reduce Multiple Negotiation candidates' optimistic opinions about their worth, it is consistent with other research findings. One study found that drivers who had just been involved in an auto accident had no less inflated views of their driving ability than drivers who had not had an accident (Preston and Harris 1965), and another found that bridge players' predictions about their likely success in future competitions did not change after they lost matches (Simons 2013).

If the Multiple Negotiation candidates were unpersuaded through the bargaining process that the firm did not consider them the most qualified candidates, why did they accept lower salaries, on average, than the salaries the Solicited Salary candidates were willing to accept? We have two hypotheses. First, directors in the Multiple Negotiation process might have failed to persuade candidates that they were less desirable than their competitors for the CEO position, but might have

Table Four
Perceived Likelihood of Being Hired with Equal Salary Demands

	Average (Percent)
Perceived likelihood of being hired: Solicited Salary	61
Perceived likelihood of being hired: Multiple Negotiation (candidate hired)	66
Perceived likelihood of being hired: Multiple Negotiation (not hired)	63

nonetheless persuaded them that the other candidates were close in quality, or at least acceptable to the firm.

Second, directors might have persuaded Multiple Negotiation candidates that the other candidates were willing to take on the CEO position at a substantially lower salary. Either consideration, or both, could have caused Multiple Negotiation candidates to moderate their initial expectations and accept a lower salary than they would have demanded had they been assigned to the Solicited Salary condition.

Morale Effects of N1C2

Implementing an N1C2 process might well be penny wise and pound foolish if the savings on compensation would come at the cost of irritating, insulting, or embittering the new CEO. Boards of directors might not wish to risk souring candidates on the company, even modestly, as those candidates then might work just a little bit less hard on behalf of the company if hired. And perhaps an N1C2 negotiating process would provoke a negative reaction on the part of candidates, either because the process is not customary, or because it will tend to result in lower compensation.

We tried to test this conjecture in the experiment and were unable to find any data suggesting negative morale consequences of either N1C2 approach. We asked every candidate subject who was hired for the CEO position in all three conditions to rate, on a scale of 1 to 7, their level of enthusiasm for the agreement that was reached, their perception of whether they were treated fairly by the firm in the negotiation process, and their level of motivation to lead Bartleby Manufacturing, Inc. as its new CEO. We found no evidence at all, much less statistically significant evidence, that participants subjected to N1C2 processes were insulted, offended, or otherwise miffed at their treatment in the hiring process.

Table Five
Morale Effects (Choose First versus Multiple Negotiation)

Experimental Condition	Enthusiasm for Agreement	Fairness	Motivation to Lead
Choose First	5.2	5.5	5.7
Multiple Negotiation	5.4	5.4	6.3

Choose First subjects, on average, rated their "level of enthusiasm for the agreement [they] reached" at 5.2, with a score of 1 representing "least enthusiastic" and 7 "most enthusiastic." Multiple Negotiation subjects, who engaged in salary negotiations before being chosen, provided an average response of 5.4.

Choose First subjects, on average, rated the agreement's fairness at 5.5, with 1 representing "least fair" and 7 "most fair." Multiple Negotiation subjects gave an average response of 5.4. Neither of these differences is significant.

Choose First subjects, on average, rated their motivation to lead Bartleby at 5.7, with 1 being "least motivated" and 7 being "most motivated." Multiple Negotiation subjects provided an average rating of 6.3. This difference is significant ($p = .01$) but it runs counter to the predicted direction: the CEOs who might have been offended by the process expressed *greater* motivation to lead (Table Five).

Solicited Salary participants were subject to what might be considered by some to be a coercive hiring process, having been required to state their minimum acceptable annual salary before being offered the CEO position, with the understanding that this would be the amount they would be paid if hired. But, at least for our subjects, this group expressed the most positive feelings about the hiring process and the firm, although their responses were not significantly different than those of

Table Six
Morale Effects (Choose First versus Solicited Salary)

Experimental Condition	Enthusiasm for Agreement	Fairness	Motivation to Lead
Choose First	5.2	5.5	5.7
Solicited Salary	6.0	6.0	6.4

the subjects in the other two conditions. Solicited Salary subjects rated (on average) their level of enthusiasm for the agreement at 6.0, the fairness of the agreement at 6.0, and their motivation at 6.4 (Table Six).

In sum, our simulation detected no hint that employees subjected to N1C2 hiring processes would feel more negatively about the process or the firm than subjects who engaged in typical C1N2 process.

Why didn't our experimental condition candidates report any hint of upset or low morale compared with our Choose First candidates? We recognize that at least part of the answer might be that our student subjects did not have any particular expectations of how they "should" be treated in this type of salary negotiations because they are not actually senior executives. But we think something more is at work. We suspect that experimental condition candidates were more likely to view their current salaries as plausible reference points, along with the salaries of CEOs of other companies, whereas candidates (and directors too) negotiating a salary after already having been anointed the new CEO would be more likely to focus their attention *only* on other CEO salaries, because once they see themselves as the CEO their previous salaries (in this case their salaries as a COO) no longer have salience.

We have some, admittedly indirect, evidence to support this hypothesis. After completing the simulation, we asked all successful candidates to "set aside the role you were asked to play in the simulation" and results, and specify the salary amount that they believed "would have been the objectively most fair amount for Bartleby to pay the new CEO." Choose First candidates on average reported that a salary of $8.79 million would be "most fair," significantly higher than Multiple Negotiation candidates who were hired ($p = .04$) ($7.65 million) and higher (although not significantly so) than Solicited Salary candidates who were hired ($8.13 million). Candidates' differential views of what salary would be fair for the CEO position based on the negotiation process used likely explain why candidates in the experimental conditions reported the same high level of morale as did candidates in Choose First, even though the former were paid less (hypothetically, of course, although these differences were reflected in the subjects' real-money payments).

When salary is bargained – as in the Choose First and Multiple Negotiation conditions – both the candidates' and the directors' perceptions of what salary would be "most fair" are likely to affect the final outcome. Our results reflect this prediction. In both of these conditions, the successful candidates' opinions as to the "most fair" salary are significantly correlated with the negotiated outcomes ($p < .001$ for Choose First; $p = .03$ for Multiple Negotiation). That is, a lower candidate

perception of what salary would be fair predicts a lower actual salary at the end of bargaining.

Like the successful candidates, the directors in Multiple Negotiation reported lower "most fair" salaries ($7.76 million) than did the directors in Choose First ($8.26 million), although this difference in means is not significant. The directors' opinions of the "most fair" salary are significantly ($p < .001$ in Choose First) or marginally significantly ($p = .06$ in Multiple Negotiation) correlated with the negotiated salary outcomes.[5] Although correlational results do not prove causation, this pattern leads us to hypothesize that the N1C2 negotiation processes used in the experimental conditions led the parties to determine that a relatively lower salary amount would be fair, which in turn enabled directors to pay candidates lower salaries without a corresponding loss of morale on the part of the successful candidates.

Why Do Firms Use C1N2?

If firms could, in fact, pay their CEOs and other executives lower salaries by employing an N1C2 process, why do they instead routinely use C1N2? In this section, we consider several possible explanations, which we consider in order of least (in our opinion) to most plausible.

Director Self-Interest

Perhaps firms know that they are overpaying their CEOs and continue to do so because this benefits the directors who hire the CEOs, even though doing so hurts shareholders. Many scholars believe high salaries are common because corporate directors wish to curry favor with CEOs, who, according to the "managerial power" theory of corporate governance (Elson 1993; Bebchuk and Fried 2004), can use their influence to affect directors' reappointment to board seats and to bestow other favors. This "managerial power" theory assumes that at least some CEOs can influence the director nomination process and thereby exercise retributive power over directors who resist their desires on issues – such as CEO compensation – that matter most to them. Managerial power theorists recognize that any such power would be grounded in group dynamics, not law; state corporation statutes vest the power to run the company in the hands of the board of directors, not the executives (Elson 1993; Bebchuk and Fried 2004). Without taking a stand on this theory's validity, we recognize that the same logic – that directors are something less than faithful agents for the firm's owners – could explain why firms choose to use a hiring process that benefits CEOs at the expense of shareholders.

Although this might be a partial explanation, it seems unlikely that director malfeasance can fully account for a practice that appears to be

as universal as C1N2. Surely not *all* directors act out of self-interest and contrary to their fiduciary obligations! In addition, the desire to curry favor with the future CEO cannot explain why firms also often use the same procedure when hiring more junior executives. To take one example, we know from personal experience that law schools usually employ a C1N2 process when they make lateral faculty hires, carefully deciding who they believe is the best candidate for the opening and issuing an "offer" before attempting to agree on compensation. But the administrators who employ this process rarely if ever stand to personally benefit financially from lining the pockets of the chosen faculty members.

Wide Disparities in Evaluations of Candidates

The N1C2 approaches provide firms with greater bargaining power than C1N2 because, in the former, the firm has a more credible threat to hire a different candidate if the first-choice candidate demands a salary that is too high. This advantage will not accrue if firms strongly prefer one candidate to another and candidates know this, even if candidates do not know *which* candidate is the preferred one.

In interviews with public company directors, one of us has found that many assert that it is nearly always the case that boards of directors have an extremely strong preference for one candidate over all others, such that financial sacrifices that a second-choice candidate might offer to make would not outweigh the first-choice candidate's perceived marginal value. The academic literature includes similar claims (Dorff 2014).

If this is both true and common knowledge among candidates, finalists for a position might reason that they would be best off negotiating in exactly the same way that they would under a C1N2 process. If it turns out that the candidate is the firm's strong first choice, demanding the same high salary that she would if in a C1N2 process would not reduce her chances of being hired and would likely result in her securing greater compensation. If it turns out that another candidate was the company's strong first choice, by definition, the candidate will not be hired regardless of her bargaining strategy, so what does she have to lose? If candidates would behave the same way under a C1N2 or an N1C2 approach, firms would have no reason to suffer whatever marginal transaction costs or social costs that would be associated with the N1C2 approach.

The empirical assumptions necessary for this theory to explain the prevalence of C1N2 are quite restrictive and seem unlikely to be satisfied in the real world, however. Although firms will often have a strong preference for one finalist candidate over the others, it seems quite likely that, in at least some cases, boards of directors will assess the two (or more) candidates as being close to each other in value. Because

candidates would usually not know whether they were participating in a process in which one candidate was a strong favorite or one in which the competition was judged as close, candidates would have a strategic incentive to moderate their negotiating demands, at least somewhat, in a N1C2 process as compared to a C1N2 process.

Transaction Costs: Absolute and Relative

The C1N2 approach has the advantage of limiting the number of salary negotiations to one for each CEO (or other employee) chosen. The search process required to hire a CEO or other senior executive is nearly always time-consuming and expensive. This is even more likely to be the case when negotiations must cover a range of compensation issues rather than a single salary figure. And, of course, after the CEO's initial term expires, the firm will find itself in a C1N2 process when negotiating a renewal agreement, unable to use uncertainty over the identity of its first-choice candidate to secure bargaining leverage. That is, the advantages of N1C2 would extend only to the initial term of employment.

But the fact that N1C2 would likely entail higher transaction costs – which is perhaps an explanation for avoiding that approach in the case of lower-level employees – does not readily explain the almost complete absence of this approach in negotiations with CEOs and other senior executives. When millions of dollars are at stake, the higher cost of negotiating with several candidates rather than one seems clearly to be justified.

A more plausible, related hypothesis for why firms eschew N1C2 is that boards might cognitively undervalue the potential absolute gains because the potential savings appears small relative to the company's total expenditures or relative to the perceived marginal value of hiring the best CEO. If choosing the "right" CEO might mean billions in profits, saving a few million in the transaction could seem trivial. Defenders of high levels of executive compensation frequently make just this argument (Murphy 1995; Loewenstein 2000).

In a famous experiment that illustrates the general phenomenon, Amos Tversky and Daniel Kahneman asked subjects to imagine they were about to buy a calculator and a jacket. In one version of the experiment, the calculator priced at $125, while the jacket cost $15; in the other version, the calculator cost $15 and the jacket $125. A hypothetical sales clerk informed the subjects that the calculator was five dollars cheaper at a store twenty minutes away. The experimenters then asked the subjects whether they would be willing to drive twenty minutes to save the five dollars. When the calculator's original price was $125, subjects were much less likely to choose to drive to save the money than when it originally cost $15 (Tversky and Kahneman 1981).

As Richard Thaler has explained, people often evaluate their choices by comparing them to a "reference level that is determined by the context within which the decision arises" (Thaler 1999.: 186). This could hardly be viewed as a best practice in the context of executive compensation. Corporate boards should attempt to secure the services of the best CEO at the lowest price possible, just as they would bargain for the lowest price when buying raw materials or other factors of production. The magnitude of the benefits to be gained may be relevant to the board's reservation price, but it should not affect the directors' eagerness to pay less than that if they can.

Procedural Fairness Norms

Another reason firms might not insist on N1C2 hiring processes is that they fear candidates will resist committing to compensation arrangements before being selected for a position. The firm could insist that candidates engage in a N1C2 process as a condition of being considered for the position, of course, but this could harm the important relationship between firms and their potential executives.

Two possible, slightly different concerns could be in play here. First, candidates put off by a N1C2 process might refuse to participate in the process at all, causing the firm to lose out on the chance of hiring a top candidate. Second, candidates' resentment might, in turn, increase the risk of moral hazard: a new CEO who feels mistreated in the hiring process might not perform as well for the firm, ultimately costing the firm far more in the long run than it might hope to save in the short run. Both concerns are likely to be heightened by the fact that C1N2 is the customary process for senior executive hires.

This explanation is plausible on its face and might rightly give firms pause about instituting N1C2 hiring processes for CEOs and other senior executives. If the hiring process leaves successful candidates believing they have been treated unfairly, sapping their motivation or dedication, the upfront salary savings could be dwarfed by the long-term cost to firm profits.

As described above, we could find no evidence that the two N1C2 hiring processes that we tested experimentally created any such negative feelings in our subjects. But, of course, this is hardly definitive proof. Our experiment ends at the conclusion of the hiring process. Our subjects never proceeded to actually perform the job in question, even hypothetically, so we do not know what feelings they might have had or attitudes they might have displayed while working. And, of course, we recognize the external validity concern: our subjects were students who are not, in fact, senior executives, and they might not share the expectations and sensibilities of actual executives. It is possible that

real executives, imbued with corporate norms and cultures, would react more negatively to N1C2 than student subjects, who lack that experience and, perhaps, knowledge of corporate hiring norms.

Substantive Compensation Norms

A related possibility is that substantive norms concerning what constitutes fair compensation for CEOs, and for other senior executives, might dominate strategic bargaining considerations in executive hiring, such that firms would not actually take advantage of whatever bargaining power advantage they theoretically could generate by choosing an N1C2 hiring process.

Assume that Candidate A is excited about the challenge of becoming the CEO of a firm and would accept the position for a salary of $3 million per year, but industry professionals believe that $7 million is the "fair" salary – perhaps because the prior CEO has earned, or CEOs of similarly sized companies in the industry currently earn, that amount. The board of directors might reasonably decide that the potential savings it could achieve by exercising greater bargaining power would be outweighed by the negative (although difficult to measure) consequences of having an unhappy CEO who felt underpaid.

A policy of paying the CEO based on external comparisons rather than negotiating vigorously over pay might, consistent with the efficiency wage hypothesis (Shapiro and Stiglitz 1984), benefit companies by ensuring corporate leadership is maximally productive. Firms with this belief might rationally choose to implement a C1N2 strategy because of its lower transaction costs.

This theory, like the prior one, depends on an unproven empirical assumption. In particular, the theory demands that firms embrace substantive fairness norms that require them to pay a new CEO based on a particular conception of the "market price" for executive talent, regardless of where the negotiating leverage lies.[6] Further work is necessary to determine whether such norms are widespread, or even exist at all.

Conclusion

In this article, we first articulated a theory, which we call the "bargaining process hypothesis," that "negotiate first, choose second" (N1C2) bargaining approaches would benefit firms and their shareholders compared to the typical "choose first, negotiate second" (C1N2) approach by enabling firms to hire the same executives at lower salaries. We then described the results of a high-context negotiating simulation between experimental subjects playing the roles of corporate directors and CEO candidates that are consistent with that theory.

This leaves us with a puzzle. When hiring senior executives, examples of firms employing N1C2 approaches are virtually non-existent. Why do directors voluntarily give up bargaining leverage, make more desirable salary reference points less salient, and risk being hamstrung by commitment bias by choosing CEOs and other senior executives before they negotiate the candidate's compensation? We have suggested and critiqued five explanations for why widespread practice runs directly contrary to our theory. Some are more plausible than others, but none are fully convincing. We hope this article will stimulate further research into these questions to determine which explanation is correct, and whether N1C2 is a normatively desirable approach to bargaining with the CEO.

NOTES

We thank Sam Pierce for excellent research assistance, Ben Nyblade, director of the UCLA School of Law Empirical Research Group, for conducting statistical analyses, Armound Ghoorchian for help with revisions, and workshop participants at the UCLA School of Law, Vanderbilt Law School, the National Corporate Law Scholars Conference, the Los Angeles Corporate Law Workshop, the Conference on Empirical Legal Studies (CELS), the UCLA-Columbia Experimental Methods in Legal Scholarship Workshop, the International Conference on Contracts (K-Con), and the Alternative Dispute Resolution Works-in-Progress Conference for helpful comments and suggestions.

1. In the control condition, candidate subjects were not assigned roles until director subjects had selected their preferred candidate (thus, all the candidates in that condition were "chosen" by the director).

2. We implemented a somewhat artificial greeting so that subjects in Solicited Salary, like subjects in the other two conditions, would know the face of the person playing the opposite role in the simulation.

3. For this and the other primary findings, differences are also significant when the analysis is limited to the UCLA subjects. Differences for Southwestern subjects, of which there were fewer, are directionally as predicted but not significant. We report all data pooled, and, unless otherwise indicated in the text, p values are calculated based on a t-test of independent samples of the groups being compared. T-tests are one-tailed, given that we have a clear directional hypothesis. We report results with p values at or below the conventional threshold of .05 as statistically significant, and results with p values between .05 and .1 as marginally significant.

4. In two cases, director and candidate provided inconsistent reports after the fact as to which side made the first offer or demand, but the subsequent results were unaffected by whether we credited the memory of the director or the candidate.

5. When the candidates' and the directors' "most fair" salary amounts are considered in the same regression model, the candidates' most fair amounts remain highly significant (Choose First, $p < .001$; Multiple Negotiation, $p = .05$), but the directors' most fair amounts become only marginally significant (Choose First, $p = .09$; Multiple Negotiation, $p = .07$).

6. Notice that this social understanding of "market price" is different than the economic understanding of a "market price," which assumes perfect competition in which all parties are price takers. In salary negotiations for a CEO, it is unusual that other companies would be offering a candidate an identical CEO position at the same salary at that moment. Rather, the understanding of the market price is based on what similar companies happen to be paying their CEOs who are in place.

REFERENCES

Alicke, M. D., and O. Govorun. 2005. The better-than-average effect. In *The self in social judgment*, edited by M. D. Alicke, D. A. Dunning, and J. I. Krueger, 85–106. New York: Psychology Press.

Anabtawi, I. 2005. Explaining pay without performance: The tournament alternative. *Emory Law Journal* 54(4): 1557–1602.

Barris, L. J. 1992. The overcompensation problem: A collective approach to controlling CEO pay. *Indiana Law Journal* 68(1): 59–100.

Bebchuk, L., and J. Fried. 2004. *Pay without performance: The unfulfilled promise of executive compensation*. Cambridge, MA: Harvard University Press.

———, and Y. Grinstein, 2005. The growth of executive pay. *Oxford Review of Economic Policy* 21(1): 283–303.

Bizjak, J., M. Lemmon, and T. Nguyen. 2011. Are all CEOs above average? An empirical analysis of compensation peer groups and pay design. *Journal of Financial Economics* 100(3): 538–555.

Bogus, C. T. 1993. Excessive executive compensation and the failure of corporate democracy. *Buffalo Law Review* 41(1): 1–83.

Bok, D. 1993. *The cost of talent: How executives and professionals are paid and how it affects America*. New York: The Free Press.

Cao, M., and R. Wang, 2013. Optimal CEO compensation with search: Theory and empirical evidence. *Journal of Finance* 68(5): 2001–2058.

Chang, Y. Y., S. Dasgupta, and G. Hilary. 2010. CEO ability, pay, and firm performance. *Management Science* 56(10): 1633–1652.

Crystal, G. 1992. *In search of excess: The overcompensation of American executives*. New York: W. W. Norton.

Dorff, M. B. 2014. *Indispensable and other myths: Why the CEO pay experiment failed and how to fix It*. Berkeley: University of California Press.

Edmans, A., and X. Gabaix. 2009. Is CEO pay really inefficient? A survey of new optimal contracting theories. *European Financial Management* 15(3): 486–496.

Elson, C. 1993. Executive overcompensation – A board-based solution. *Boston College Law Review* 34(5): 937–996.

Elson, C. M. and C. K Ferrere. 2013. Executive superstars, peer groups and overcompensation: cause, effect and solution. *Journal of Corporation Law* 38(3): 487–531.

Falato, A., D. Li, and T. Milbourn. 2014. To each according to his ability? The returns to CEO talent. Available at ssrn.com/abstract=1699384.

Festinger, L. 1957. *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.

Fisher, R., W. Ury, and B. Patton. 1991. *Getting to yes: Negotiating agreement without giving in*, 2nd ed. New York: Penguin Books.

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

Gabaix, X., and A. Landier. 2008. Why has CEO pay increased so much? *Quarterly Journal of Economics* 123(1): 49–100.

Galinsky, A. D., and T. Mussweiler 2001. First offers as anchors: The role of perspective-taking and negotiator focus. *Journal of Personality and Social Psychology* 81(4): 657–669.

Gelles, D. 2015. For the highest-paid CEOs, the party goes on. *New York Times* (May 16). Available at www.nytimes.com/2015/05/17/ business/for-the-highest-paid-ceos-the-party-goes-on.html.

Hayes, R. M., and S. Schaefer. 2009. CEO pay and the Lake Wobegon effect. *Journal of Financial Economics* 94(2): 280–290.

Kahneman, D. 1992. Reference points, anchors, norms, and mixed feelings. *Organizational Behavior and Human Decision Processes* 51(2): 296–312.

———, and A. Tversky. 1995. Conflict resolution: a cognitive perspective. In *Barriers to conflict resolution*, edited by, K. Arrow, R. H. Mnookin, L. Ross, A. Tversky and R. B. Wilson, 44–61. New York: W. W. Norton.

Khurana, R. 2002. *Searching for a corporate savior: The irrational quest for charismatic CEOs*. Princeton, NJ: Princeton University Press.

Korobkin, R. 2014. *Negotiation theory and strategy*, 3rd ed. New York: Wolters Kluwer.

Loewenstein, M. J. 2000. The conundrum of executive compensation. *Wake Forest Law Review* 35(1): 1–30.
Lublin, J. S. 2015. How much the best-performing and worst-performing CEOs got paid. *Wall Street Journal* (June 25). Available at www.wsj.com/articles/how-much-the-best-and-worst-ceos-got-paid-1435104565.
Moore, D. A. 2007. Not so above average after all: When people believe they are worse than average and its implications for theories of bias in social comparison. *Organizational Behavior and Human Decision Processes* 102(1): 42–58.
Murphy, K. J. 1995. Politics, economics, and executive compensation. *University of Cincinnati Law Review* 63(2): 713–748.
———. 1999. Executive compensation. In *Handbook of labor economics, vol 2, edited by* O. Ashenfelter and D. Card, 2485–2563. Amsterdam: Elsevier.
———. 2002. Explaining executive compensation: Managerial power versus the perceived cost of stock options. *University of Chicago Law Review* 69(3): 847–869.
———. 2013. Executive compensation: Where we are and how we got there. In *Handbook of the economics of finance, vol 2A, edited by* G. M. Constantinides, M. Harris and R. Stulz, 211–356. Amsterdam: Elsevier.
Murphy, K. J., and J. Zabojnik. 2004. CEO pay and appointments: A market-based explanation for recent trends. *American Economic Review* 94(2): 192–196.
Preston, C. E., and S. Harris. 1965. Psychology of drivers in traffic accidents. *Journal of Applied Psychology* 49(4): 284–288.
Raiffa, H. 1982. *The art and science of negotiation*. Cambridge, MA: Harvard University Press.
Rosen, S. 1981. The economics of superstars. *American Economic Review* 71(5): 845–858.
———. 1982. Authority, control, and the distribution of earnings. *Bell Journal of Economics* 13(2): 311–323.
Shapiro, C., and J. E. Stiglitz. 1984. Equilibrium unemployment as a worker discipline device. *American Economic Review* 74(3): 433–444.
Sharot, T. 2011. The optimism bias. *Current Biology* 21(23): 941–945.
Simons, D. J. 2013. Unskilled and optimistic: Overconfident predictions despite calibrated knowledge of relative skill. *Psychonomic Bulletin and Review* 20(3): 601–607.
Terviö, M. 2008. The difference that CEOs make: An assignment model approach. *American Economic Review* 98(3): 642–668.
Thaler, R. H. 1999. Mental accounting matters. *Journal of Behavioral Decision Making* 12(3): 183–206.
Thomas, R. S. 2004. Explaining the international CEO pay gap: Board capture or market driven? *Vanderbilt Law Review* 57(4): 1171–1268.
Tversky, A. and D. Kahneman. 1981. The framing of decisions and the rationality of choice. *Science* 211(4481): 453–458.
Wackerle, F. W. 2001. *The right CEO: Straight talk about making tough CEO selection decisions*. San Francisco: Jossey-Bass.
Williams, S. H. (2014) Probability errors: Overoptimism, ambiguity aversion, and the certainty effect. In *The Oxford handbook of behavioral economics and the law*, edited by, E. Zamir and D. Teichman, 335–353. Oxford: Oxford University Press.
Yablon, C. M. 1999. Bonus questions – Executive compensation in the era of pay for performance. *Notre Dame Law Review* 75(1): 271–308.