Are Energy Executives Rewarded for Luck?

Lucas W. Davis[a] *and Catherine Hausman*[b]

ABSTRACT

In this paper, we examine executive compensation data from 78 major U.S. oil and gas companies over a 24-year period. Perhaps in no other industry are the fortunes of so many executives so dependent on a single global commodity price. We find that a 10% increase in oil prices is associated with a 2% increase in executive compensation. This oil price effect holds for both CEOs and non-CEOs and separately for several different individual components of compensation, including bonuses. We find that the oil price effect is larger in companies with more insiders on the board, and asymmetric, with executive compensation rising with increasing oil prices more than it falls with decreasing oil prices. We then discuss potential mechanisms drawn from the broader existing literature on executive compensation.

Keywords: Pay-for-Luck, Executive compensation, Principal-Agent problem, Rent extraction, Performance pay

https://doi.org/10.5547/01956574.41.6.ldav

1. INTRODUCTION

From January 2014 to January 2016, oil prices fell from nearly $100 per barrel to just over $30 per barrel. In those same two years, the CEOs of 30 large U.S. oil and gas exploration companies lost an average of over half a million dollars each in annual compensation. Perhaps in no other industry are the fortunes of so many executives so dependent on a single global commodity price.

In this paper, we analyze executive compensation data from 78 U.S. oil and gas companies over a 24-year period. We document a strong correlation between crude oil prices, company value, and executive compensation. In our preferred instrumental variables specification, a 10% increase in company value driven by oil prices leads to a 2% increase in executive compensation.

We then perform a series of additional analyses to better understand the mechanisms. First, we show that this oil price effect is robust to including time-varying controls for capital expenditures and labor. Second, we show that the oil price effect holds for both CEOs and non-CEOs. Third, we show that the oil price effect is widespread across the different individual components of executive compensation, including not only total compensation, but also bonuses and long-term cash incentives. Fourth, we show that the oil price effect is smaller at better-governed companies. Results are similar with two measures of firm governance, both related to the presence of more insiders on the board. Finally, we show that the oil price effect is asymmetric, with executive compensation increasing more with rising oil prices than it decreases with falling oil prices.

We then discuss potential interpretations, drawing from the existing literature on executive compensation. Much of the broader literature on executive compensation is aimed at reconciling the

a Haas School of Business, University of California, Berkeley, Berkeley, CA 94720. E-mail: lwdavis@berkeley.edu.

b Corresponding author. Ford School of Public Policy, University of Michigan, Ann Arbor, MI 48109. E-mail: chausman@umich.edu.

"rent extraction" and "shareholder value" views.[1] Under the rent extraction view, executives have co-opted the pay-setting process, and are increasing compensation as much as possible. In contrast, under the shareholder value model, pay is set within a competitive executive market, structured in such a way that executives are properly incentivized to exert effort on behalf of the firm.

An influential analysis by Bertrand and Mullainathan (2001) (hereafter B&M) interprets regression results similar to ours as evidence of rent extraction. B&M test whether CEOs are rewarded for shocks to firm performance beyond their control. Their paper uses several different measures of such shocks, but some of the most compelling evidence comes from a case study of the oil industry. They find that CEO compensation responds just as much to changes driven by oil prices as it does to generic changes in company value, an effect which they term "pay-for-luck." As first pointed out by Holmstrom (1979), companies that are maximizing shareholder value should "filter out" oil prices and other forms of observable luck. Executives are risk averse, so the optimal pay-for-performance contract focuses purely on factors under control of the executive. To include luck only makes the contract riskier without providing better incentives. Indeed, it is difficult to reconcile this oil price effect with the predictions of this standard simple contracting model.

Still, there are ways to reconcile the oil price effect with shareholder value maximization. For example, under a simple model of profit maximization, firms buy more inputs when output prices go up. One could imagine that when oil prices are high, additional executive effort is needed, and so compensation rises to induce that effort. This could explain our main oil price effect, but would not in its simplest form explain our additional results relating to governance and asymmetry. As in much of the rest of the literature, we are unable to sharply distinguish shareholder value and rent extraction interpretations.[2] Part of the challenge, as explained by Murphy (2013), is that these two views are not mutually exclusive, with both forces impacting compensation to varying degrees across firms and over time.

In addition to providing an advantageous case study, the U.S. energy industry is of significant intrinsic interest. The United States is the world's largest producer of oil and natural gas. The annual value of U.S. oil and natural gas production exceeds $200 billion, and the firms in our sample had a total market value of almost half a trillion dollars in 2016. Reflecting the size of this industry, the dollar value at stake in executive pay is substantial: total compensation of all oil and gas executives in the latter part of our sample is almost $1 billion per year. That said, using oil price variation means that we must identify our main effects of interest using a single national time series.

The paper proceeds as follows. Section 2 provides background on the related literature and on the oil and gas industry. Section 3 describes our data. Sections 4 and 5 present empirical results. Section 6 discusses the results in the context of the existing competing theories about executive compensation, and Section 7 concludes.

2. BACKGROUND

2.1 Related Literature

There is an immense existing literature on executive compensation. However, other than B&M, we are not aware of any other study focused on the effect of oil prices on executive compen-

1. Surveys of this literature include Abowd and Kaplan (1999), Murphy (1999), Bebchuk and Fried (2003), Bertrand (2009), Edmans and Gabaix (2009), Murphy (2013), Edmans and Gabaix (2016), and Edmans et al. (2017).

2. See Himmelberg and Hubbard (2000); Oyer (2004); Bolton et al. (2006); Axelson and Baliga (2009); Gopalan et al. (2010); Hoffmann and Pfeil (2010); Noe and Rebello (2012); Danthine and Donaldson (2015); Chaigneau and Sahuguet (2018).

sation. While there is a substantial literature since B&M on the question of whether executives are "paid-for-luck," these studies have instead focused on relative performance, i.e., how an executive's own company's performance compares to that of other companies in the same industry, and how this "relative performance" affects the executive's compensation. In these studies, "luck" is measured using within-industry average performance. Just as a contract could filter out the effect of oil prices, so could a contract filter out within-industry average performance. Filtering out the exogenous industry-wide ebbs and flows decreases the variance of compensation, without reducing incentives for executives to take actions to benefit the firm.

The literature testing for relative performance evaluation has found mixed evidence (Antle and Smith, 1986; Aggarwal and Samwick, 1999b; Gibbons and Murphy, 1990; Garen, 1994; Garvey and Milbourn, 2003). Typically these studies take the form of testing whether executive compensation is tied to *absolute* firm performance, which depends in part on industry-wide lucky breaks, or is tied to *relative* firm performance, which filters out observable industry-wide shocks. However, conclusions in that literature depend in large part on how the researcher defines the peer group, and for some peer comparison groups there is evidence of relative performance evaluation (Gong et al., 2011; Lewellen, 2015).[3]

A significant advantage of using relative performance is that this measure is available for all industries, facilitating larger-scale analyses and cross-industry comparisons. Oil prices have certain advantages too, however. Oil prices are both exogenous and highly volatile, driven by worldwide shocks.[4] In contrast, with relative performance studies there is ambiguity about how the "peer group" is defined, which introduces measurement error and potential endogeneity, since compensation boards have some flexibility in these choices. Moreover, in industries that are not perfectly competitive, executives may be able to directly influence competitor market value (Aggarwal and Samwick, 1999a).

One of our additional analyses is motivated by several studies in the relative performance literature that test for asymmetry. Garvey and Milbourn (2006) document that executives are rewarded more for good luck than they are punished for bad, which the authors argue is consistent with rent extraction. This argument is bolstered by their finding that both "pay-for-luck" and the asymmetry are stronger at firms with worse governance.[5] In contrast, Bizjak et al. (2008) argue that asymmetry in "pay-for-luck" could be the result of compensation boards using benchmarking to set wages at market reservation levels, motivated by their finding that the asymmetry appears for firms paying their CEOs below the peer group median. They also argue more generally that other observed empirical facts are consistent with executive compensation being set in a competitive environment rather than as a result of rent-seeking. Campbell and Thompson (2015) also argue that

3. Also related is the literature on benchmarking, which investigates how pay *levels* depend on peer comparisons, rather than on how pay varies with peer performance; see, e.g., Albuquerque et al. (2013); Cremers and Grinstein (2014). Other related work has focused on executive hiring and firing decisions: Jenter and Kanaan (2015) documents that industry movements (i.e. bad luck) are predictive of CEO dismissals, and Kaplan and Minton (2012) documents that both firm-specific movements and overall stock market movements predict CEO turnover.

4. Outside of the executive compensation literature, these features of oil prices have been widely noted. For example, Kline (2008) writes that oil prices "provide ample exogenous variation" and "are well measured, volatile, and difficult to forecast" (p. 3). He also notes other advantages of studying the oil industry, including the importance of oil shocks to the economy as a whole. Kline (2008) uses variation in crude oil prices to study labor market dynamics. Using data from the *Current Employment Survey*, he finds that employment and wages in the U.S. oil and gas field services industry increase with crude oil prices.

5. Asymmetry and governance are further explored in Harford and Li (2007), which examines compensation following acquisitions; and in Bebchuk et al. (2010), which investigates governance and director compensation. Bebchuk and Fried (2003) summarizes additional related work on governance.

retention concerns are better able to explain asymmetric "pay-for-luck" than are explanations relating to rent extraction. Finally, Bell and Van Reenen (2016) examine both asymmetry and the impacts of firm governance, finding evidence of "pay-for-luck" in UK firms.

Another paper in the relative performance evaluation literature is Cremers and Grinstein (2014). They ask whether retention concerns can explain observed "pay-for-luck," arguing that compensation practices depend on whether the pool of executives comes from within the industry or from outside industries. They write that, "in an industry with many outsider CEOs and where the overall supply of CEOs will be relatively inelastic, boards may be forced to raise their CEOs' compensation if there is a positive industry-wide shock. . . . However, in industries with few outsider CEOs, such a competitive labor market argument would be less compelling because CEOs and top executives are beholden to the firm" (p. 947). Interestingly, Cremers and Grinstein (2014) find that the oil and gas industry is one of the sectors with few outside hires. According to their logic, observed "pay-for-luck" in a sector like oil and gas would be difficult to explain by retention concerns in a competitive labor market. We view our paper as complementary to this existing literature on asymmetry and retention concerns, in an industry context in which luck is particularly important and easy to measure.

2.2 U.S. Oil and Gas Sector

Our results provide a window into executive compensation in a dynamic, multi-billion dollar sector. As mentioned in the introduction, the U.S. oil and gas exploration companies in our sample had a total market value of almost half a trillion dollars in 2016. The sector is composed of both large, old firms like ConocoPhillips (which began extracting oil over one hundred years ago), and newer firms such as Anadarko Petroleum Corp (established in 1959) and Chesapeake Energy Corp (founded in 1989). While those three companies are quite large—with over \$10 billion in market value each in our sample—dozens of smaller companies are also publicly traded.

The oil and gas sector has changed dramatically since the period examined by B&M. Most importantly, the rise of hydraulic fracturing and associated innovations has substantially increased total U.S. oil and gas production. Hydraulic fracturing has been called "the biggest energy innovation since the start of the new century" (Yergin, 2011) and has had a large positive impact on the U.S. economy (Hausman and Kellogg, 2015). U.S. oil and natural gas production has increased more than 50 percent since 2008, making the United States the world's largest petroleum and natural gas producer (Energy Information Administration, 2018). Along with this growth in production, there have been dozens of new entrants into the oil and gas industry. While the new entrants tend to be smaller firms, some of the entrants have rapidly become major producers. Concho Resources, for instance, was founded in the mid-2000s but by 2016 was among the top ten publicly-traded U.S. oil and gas firms by market value.

Reflecting the size of the industry, executive compensation is substantial, with average executive compensation over the last decade at \$4 million per year, and average CEO compensation over the last decade at \$8 million. For CEOs, this is more than three times the average CEO compensation in the B&M sample, equal to \$2.4 million in 2016 dollars. Total compensation for the five highest-paid executives at the top 30 companies in our data totaled over \$650 million in 2016. This high level of compensation also reflects rising compensation of executives in the U.S. in all industries. From 1992 to 2016, CEO compensation at the 100 largest firms in the U.S. nearly tripled in real terms, and CEO compensation at large oil and gas exploration firms more than tripled.

Several top oil and gas executives have been publicly criticized for their pay. Ray Irani of Occidental Petroleum Corp was forced out after investors criticized his pay package.[6] In 2013, the shareholders of Apache rejected proposed executive compensation plans in a non-binding say-on-pay vote. Perhaps most vivid is the professional story of the late Aubrey McClendon, cofounder of one of the largest companies in our sample, Chesapeake Energy. Through McClendon's leadership, Chesapeake rose to become a leading producer of natural gas, and McClendon was in the late 2000s one of the highest-paid CEOs in the United States. The company was, however, also plagued by questions about governance and conflict-of-interest, and McClendon was eventually forced out of the company (Gold, 2014).

Our analysis is complementary to descriptive studies of executive compensation in the oil and gas industry. We have reviewed, for example, industry reports on how oil and gas firms structure their pay packages.[7] These reports draw on proxy statements, in which firms are required to provide summaries of their executive pay-setting practices, and the summary reports provide a valuable complement to our regression-based analysis. One recent report (Alvarez and Marsal, 2018) notes a couple of useful facts, which together highlight the value of empirically estimating the responsiveness of pay to market value versus to oil prices.

First, firms use a wide variety of performance metrics when setting pay. Common metrics include total production, health/safety/environmental metrics, the value of reserves, and total shareholder return (both relative and absolute). Some of these measures are correlated with oil prices while others are not. Second, most oil and gas companies use some discretion, rather than solely formulaic plans, when determining annual pay. This use of discretion is relevant because it suggests that, even if relative performance evaluation is named as a strategy in proxy statements, it is possible that the use of discretion undoes some of the relative comparison.[8] Indeed, a blog post by a compensation consultant describes just such a mechanism following the oil price crash at the end of 2015.[9]

These facts indicate that examining realized pay structures can provide a useful summary of pay structures that vary (perhaps endogenously) across firms and across time. Moreover, while the mention of relative total shareholder return as one metric points to the possibility that we might observe the filtering of industry-wide luck from executive pay, the use of multiple other performance metrics, and the possibility that the choice of metrics is itself endogenous to compensation committee goals, suggests that we may not empirically observe much filtering. Indeed, in related work, Bell and Van Reenen (2016) argues that a potential explanation for pay-for-luck is "that CEO

6. Krauss, Clifford. 3 May 2013. "Occidental Chairman Agrees to Step Down Ahead of Schedule." *New York Times*. https://www.nytimes.com/2013/05/04/business/occidental-chairman-irani-agrees-to-leave-company.html

7. Pearl Meyer's "2016 Oil and Gas Market Review: CEO Pay and Practice Trends," available at https://www.pearlmeyer.com/2016-oil-and-gas-market-review-ceo-pay-and-practice-trends.pdf; Alvarez and Marsal's 2018 "Oil and Gas Exploration and Production (E&P) Incentive Compensation Report: Analysis of Compensation Arrangements Among the Largest U.S. E&P Companies," available at https://www.alvarezandmarsal.com/sites/default/files/article/pdf/62276_03274_tax_oil-gas_ep_incentive_comp_report_v15_pages.pdf; and similar Alvarez and Marsal reports from earlier years.

8. This is related to the findings of Wade et al. (1997) regarding how compensation committee justify compensation decisions.

9. The author writes "across a sample of 32 publicly-traded large E&P companies, we found that 50% applied negative discretion or subjective assessment (e.g., through an 'individual performance' component) to reduce annual incentive payouts from the formulaic outcome. . . . This all reinforces the importance of compensation committees maintaining some degree of subjectivity or discretion in determining bonus payouts, especially during volatile commodity cycles." Source: Szabo, Jon. "Post #41: Effective Use of Discretion in Annual Incentives." Available at http://www.meridiancp.com/wp-content/uploads/Effective-Use-of-Discretion-in-Annual-Incentives.pdf

remuneration plans are sufficiently complex that shareholders have difficulty effectively monitoring the contracts." The less able shareholders and boards of directors are able to monitor contracts, the more likely we are to see rent extraction.

3. DATA

Our primary dataset describes executive compensation and firm performance for 78 U.S. oil and gas companies over the period 1992–2016. We assemble these data from Compustat, which collects data from firms' annual proxy statements. The entire Compustat dataset covers 3,664 publicly-traded firms, including all of the S&P 1500. The panel is unbalanced because of mergers, acquisitions, entry, and exit. For our 78 companies we have examined new entrants into the sample and exits from the sample. New entrants tended to coincide with initial public offerings, spin-offs, and S&P 1500 listings; exits tended to coincide with acquisitions by other firms.

We focus on oil and gas exploration and production companies, of which there are 78 in the Compustat data. Names of each company, along with mean compensation and mean market value, are provided in the Appendix (Table A2). These firms, defined by NAICS code 211111, are engaged in "the exploration, development and/or the production of petroleum or natural gas." Oil and gas exploration is delineated separately from firms engaged primarily in other oil and gas activities, such as support activities for oil and gas (drilling on a contract basis) and petroleum refining.[10] Whereas crude oil prices have a clear and direct impact on production companies, the relationship is less clear for these other types of oil and gas companies. Refinery mark-ups, for example, can either increase or decrease with crude oil prices (Borenstein and Kellogg, 2014; Muehlegger and Sweeney, 2019).

In principle, one could perform similar analyses on other types of energy companies, such as natural gas transmission and distribution firms, or electric utilities.[11] We instead focus on oil and gas extraction; whereas oil prices are driven by worldwide supply and demand shocks, natural gas prices can be substantially impacted by U.S. supply conditions, as described in detail in Hausman and Kellogg (2015). Also, these other energy firms in our data are a heterogeneous mix of gas-only, electric-only, and gas- and electric utilities; and of competitive producers and price-regulated utilities. These different types of firms are affected quite differently by energy price shocks, limiting both power and generalizability.

Conceptually, it would also be interesting to look at other competitive commodity markets. We considered coal producers, for example. While coal is traded worldwide and thus has a fairly exogenous price, coal prices are much less variable than oil prices, limiting statistical power. Moreover, while we have 78 oil and gas companies, there are only eight coal mining companies in our data. Another possibility would be companies engaged in gold or silver mining, but there are only nine and two companies, respectively, in these sectors. In addition, we considered agricultural commodities but there are only four companies in our data engaged in corn, poultry, or cotton production; and other agricultural companies (e.g. Monsanto) tend to be broader conglomerates with many different lines of business. In short, none of these other commodity markets have nearly as many large, publicly-traded companies as does oil and gas.

10. For a complete list of NAICS codes with descriptions see https://www.census.gov/cgi-bin/sssd/naics/naicsrch. Note that the "supermajors" Chevron and ExxonMobil, as well as other vertically-integrated companies like Valero and Western Refining, are in NAICS Code 324110 "refining," and thus are excluded from this analysis.

11. Related work on utilities includes Joskow et al. (1993) and Joskow et al. (1996), which examine political pressure and executive pay at regulated utilities.

For each firm, we observe financial measures on an annual basis, including total market value, total book value, net income, total assets, the return on equity, the return on assets, the number of employees, and capital expenditures.[12] Each firm also reports executive compensation for its five highest ranked employees.[13]

Altogether, we observe 934 individual executives at oil and gas companies. As mentioned above, total compensation of all oil and gas executives in the latter part of our sample is almost $1 billion per year, and it peaked at $1.1 billion in 2008.[14] Overall, we have over 4,600 executive-year observations, compared to 827 executive-year observations in B&M. For their analysis of the oil industry, B&M use a dataset on CEO compensation from the 51 largest U.S. oil companies between 1977 and 1994. Our data include almost twice as many companies, with not only the very largest U.S. companies but also many smaller, though still publicly-traded, companies. In addition, whereas B&M examine CEOs only, we look at all top executives. This larger sample size yields a significant improvement in statistical precision, allowing us to perform richer analyses, and to include time-varying controls for capital and labor in most specifications.

For each executive, in addition to total compensation we observe its components: salary, stocks and stock options, bonuses, long-term incentives, and other compensation (such as benefits and perquisites). The value of stocks and options is the grant-date fair value of stock/options awarded. The reporting format for some of the individual components of compensation changes in 2006, as we describe in the Appendix. We also observe whether the executive is the CEO.

We merge in crude oil prices from the Energy Information Administration. We use the West Texas Intermediate (WTI) price; this is the standard benchmark price for U.S. crude and it closely follows other international crude prices.[15,16] We focus primarily on December oil prices, since market value is measured at year-end; however we also examine annual average prices.[17]

As we discuss later in more detail, we have two measures of firm governance. The first is the proportion of executives that are not on the board of directors; we construct this variable from Compustat data. Second, from Institutional Shareholder Services (ISS) we have the proportion of board members that are not insiders (e.g., employees or family members of employees).

To control for macroeconomic conditions, we use real GDP and the unemployment rate from FRED. Finally, we deflate all prices using the CPI—All Urban Consumers: All Items Less Energy, from FRED.

We show summary statistics in the Appendix (Table A3). Average annual compensation in our sample is just over $3 million, with about half coming from stocks and stock options. Market value has missing observations, making up around 7 percent of the sample; these missing values are concentrated in firms not listed on a major S&P index and in firm-year observations just prior to

12. Market value is essentially the year-end stock price times the number of shares. Book value is the stockholders' equity from the balance sheet, with adjustments for deferred taxes, investment tax credits, and preferred stock. The return on equity is the income to common equity ratio, multiplied by 100. The return on assets is the income to assets ratio, multiplied by 100.

13. Some firms report compensation for more than five executives, but we limit the sample to the top five in each firm-year.

14. The average per year for 2007 to 2016 is $900 million, reflecting mean compensation of $4 million annually for around 225 executives.

15. We also collect data on the Brent crude oil price, which diverged somewhat from the WTI price over our time period. Our main results are very similar using the Brent price.

16. For supplementary analysis, we also collect natural gas prices at Henry Hub, a major pipeline hub in Louisiana and the official delivery location for most U.S. natural gas futures contracts.

17. A firm could of course have a fiscal year ending in a month other than December; however, in our sample 97 percent of all fiscal years end in December.

a firm's IPO. Across several variables, mean and median values differ because of skew, so we log transform most variables in the specifications that follow. The difference between mean and median values is especially pronounced for the return on assets and return on equity variables; the negative mean values are driven by very negative values during recession and low oil price years.

Before presenting the details of our empirical analysis, we present descriptive evidence on the positive correlation between executive compensation, oil prices, and market value. The time-series plots in Figure 1 show that executive compensation and market value at oil and gas companies closely follow the pattern for crude oil prices during the period 1992–2016.[18] The left-hand panel shows the national average compensation for energy executives (thick black line) versus executives in all other industries (thin grey line).[19] It is immediately apparent that compensation at oil and gas companies is highly correlated with the crude oil price (thick grey line). Asymmetry also appears in this figure: compensation tracks *upward* movements of oil prices more closely than *downward* movements. Firm market value is even more closely tied to oil prices, as shown in the right-hand panel. Also apparent in the right-hand panel is the highly variable nature of market value for oil and gas firms. Figure 1 motivates the regression analyses that follow, showing that this is an industry where both executive compensation and market value are closely tied to oil prices.

Figure 1: Executive Compensation, Market Value, and Oil Prices

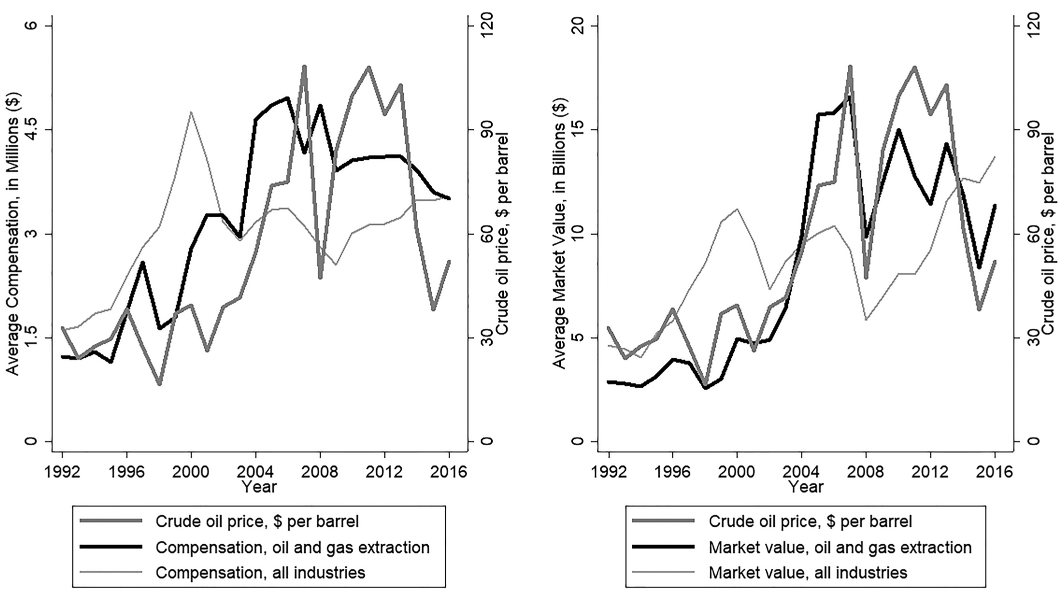

Note: This figure plots the average executive compensation (left panel) and the average year-end market value (right panel) for firms in Compustat. In each panel, the thick black line is for 78 oil and gas firms and the thin grey line for all other firms in Compustat. Executive compensation is averaged across the top five executives at each firm. Crude oil prices (West Texas Intermediate light sweet crude) are plotted in grey. All values are in 2016 dollars, using the CPI-All Urban Less Energy.

18. Bertrand and Mullainathan (2001) document a similar correlation during the period 1977–1994.

19. The data are for an unbalanced panel, as described below; the figure looks very similar if firm fixed effects are removed to adjust for compositional changes.

4. MAIN RESULTS

4.1 Pay for Performance

We begin by measuring pay for performance for U.S. energy executives. That is, we measure the extent to which executive compensation varies with the value of the firm. We estimate the following:

$$ln(C_{i,t,p}) = \alpha + \beta_1 \cdot ln(V_{i,t}) + X_{i,t,p}\Theta + \varepsilon_{i,t,p}, \tag{1}$$

where $C_{i,t,p}$ is compensation at firm i in year t for executive p, V is market value, and X is a vector of controls.[20] Time-series controls (to account for macroeconomic conditions) include GDP growth, unemployment, and a linear trend. Firm fixed effects aid with precision and account for compositional changes. A CEO dummy also aids with precision. Standard errors are two-way clustered by firm (to allow for correlation across individuals and across years within a firm) and by year (to allow for correlation across firms within a year).

Results, shown in Column 1 of Table 1 indeed show a positive coefficient. The coefficient is 0.29, statistically significant at the one percent level, indicating that for every 10 percent increase in market value, executive compensation rises by almost 3 percent.[21] This specification is analogous to Column 3 of Table 1 in B&M. The magnitude we estimate is somewhat smaller than what is estimated by B&M; their coefficient is 0.38.

Table 1: Does Executive Compensation Increase with Oil Prices?

	(1) OLS	(2) IV	(3) OLS
Log market value	0.29*** (0.04)	0.19*** (0.05)	
Log crude oil price			0.19*** (0.06)
First-stage F-statistic		90.59	
Observations	4,673	4,673	4,673
Within R^2	0.52	0.51	0.46

Note: This table reports estimates and standard errors from three separate regressions. The dependent variable in all regressions is log total annual compensation. All regressions include company effects, macroeconomic variables (national GDP growth rate and unemployment rate) and a linear trend, as well as an indicator for whether the executive is the CEO. In Column 2 we instrument for log market value with log crude oil prices. Compensation, market value, and oil prices are normalized to 2016 dollars. Standard errors are two-way clustered by firm and by year. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

The coefficient of 0.29 is economically significant. From 2000 to 2010, the average market value of firms in our sample rose by 84 log points, or $10 billion in real terms. At the same

20. When we take the log of compensation, we add $1 to any zero values; this affects only one observation. The right-hand side variable, market value, has no zero or negative values. In the Appendix we show that results are similar for alternative measures of firm performance, including book value, the return on equity, and the return on assets.

21. Aggarwal and Samwick (1999b) interact market value with a measure of the variance in firm returns, thus allowing heterogeneity in pay-for-performance across the variance of returns. When we estimate an augmented regression with this feature (Appendix Table A4), results are largely unchanged, with a similar slope estimate on the value of the firm and a small and statistically insignificant slope on the interaction term.

time, executive compensation rose 52 log points, or \$1.5 million per executive. The coefficient on market value indicates that compensation would be expected to rise over 20 log points, accounting for around half of the total increase over this time period. Moreover, recall that total compensation includes stocks and options, which we value using the grant-date value. The *exercised* value, i.e. realized pay, will be even more closely tied to the firm's stock price and therefore market value.

Regressions of this form are typically interpreted as measuring pay-for-performance. We next turn to a series of specifications that isolate the extent to which this market value effect is driven by changes in global oil prices, as opposed to generic changes to market value.

4.2 Oil Price Effect

Next, we examine the extent to which executive compensation increases due to changes in market value driven by oil prices. We follow B&M and use an instrumental variables regression of the following form:

$$ln(C_{i,t,p}) = \alpha + \beta_2 \cdot \widehat{ln(V_{i,t})} + X_{i,t,p}\Theta + \varepsilon_{i,t,p}, \tag{2}$$

where $\widehat{ln(V)}$ is the predicted log market value, from a regression of log market value on log oil price:

$$ln(V_{i,t}) = \alpha + \delta \cdot ln(O_t) + X_{i,t}\Theta + \varepsilon_{i,t}. \tag{3}$$

Here O is the December crude oil price.[22] This IV specification allows us to isolate the degree to which changes in market value are driven by global oil prices, over which executives have no control. We can thus compare the impacts of generic changes in market value on compensation to the impacts of oil price-induced changes in market value.

In this IV specification, our identifying variation now comes from an annual time series, whereas in the previous specification, we leveraged panel variation in firm-level market value. As such, the macroeconomic controls and the linear time trend that we include are now particularly important. Below (Section 4.4), we show the sensitivity of our results to alternative sets of controls. In the Appendix, we also show placebo regressions using other industries, such as manufacturing.

The first stage results are shown in the Appendix (Table A5). The coefficient on the oil price has the expected sign and is statistically different from zero at the one percent level. The coefficient of 0.99 indicates that for every 10 percent change in oil prices, market value at the average oil and gas company increases by 9.9 percent. Oil prices are a tremendous driver of value for these firms.[23]

The IV estimates are shown in Column 2 of Table 1.[24] When we instrument for market value using oil prices, we estimate a coefficient on market value of 0.19, statistically significant at the one percent level. For every 10 percent increase in market value driven by an increase in oil prices, executive compensation rises by 1.9 percent. The coefficient is statistically different from

22. For simplicity, we use the same α, $X\Theta$, and ε notation in all equations; obviously the actual estimated values vary across equations. In general, most of our specifications are run at the individual executive level, but the p subscript is dropped from equation 3 as all variables in that equation are at the firm level.

23. The first-stage coefficient of 0.99 also suggests that oil and gas companies are not significantly hedging their exposure to oil price risk. Oil and gas executives, however, are potentially a different story. For executives, not only their income, but also their wealth and human capital, are closely tied to oil prices. Executives are typically forbidden from short-selling company stock or otherwise hedging, as this would prevent performance incentives from having their intended effect (Jin, 2002; Hall and Murphy, 2002). However, in this case given the widespread availability of futures contracts for crude oil it seems reasonable to believe that at least some executives have managed to partially hedge their exposure.

24. The table also reports the first-stage instrument F-statistic, specifically the Kleibergen-Paap (2006) statistic that accounts for the clustered standard errors. The value of 90.59 indicates that we do not have a weak instruments problem.

zero at the 1% level, but not statistically different from 0.29 (the coefficient in Column 1), indicating that we cannot reject that compensation responds just as much to changes in market value due to oil prices as it does to generic changes in market value.

Our estimate is consistent with B&M, but much more precisely estimated. In particular, our standard error is less than one-third the magnitude (0.05 compared to 0.17) of their analogous estimate (Column 4 of Table 1 of their paper). This increased statistical power reflects our larger sample size and somewhat longer time frame. We also have about 20 percent more variation in oil prices as measured by the standard deviation of the log real price. We also have more precision than B&M in specifications using accounting rates of return as the measure of firm value (Table A6). Our point estimate is smaller than B&M (0.19 compared to 0.35). However, when using annual average oil prices, rather than December prices, we estimate a coefficient of 0.31, more comparable to the B&M estimate (see Appendix, Table A6).

Finally, we show in Column 3 of Table 1 the estimates from an alternative regression, in which we regress compensation directly on oil prices:

$$ln(C_{i,t,p}) = \alpha + \psi \cdot ln(O_t) + X_{i,t,p}\Theta + \varepsilon_{i,t,p}. \qquad (4)$$

The coefficient on the oil price is 0.19, statistically significant at the one percent level. This implies that for a 10 percent rise in oil prices, executive compensation rises by 1.9 percent. This is an economically significant amount. The mean absolute year-on-year change in oil prices over this period was 30 log points. Thus in a typical year, oil and gas executives saw their compensation vary by 6 log points because of oil price changes, or $180,000 at the mean pay level. Moreover, from 2013 to 2015, oil prices fell by 100 log points (60 percent)—implying a drop in annual compensation of over $500,000 per executive.

4.3 Controlling for Labor and Capital

B&M interpret their IV estimate as evidence of rent extraction. But under some conditions this pattern could be explained with a simple model of firm profit maximization. When oil prices go up, firms want to produce more—so firms will choose to buy more of all inputs (capital, labor, and executives), in order to produce and sell more oil. More executives could mean hiring more individuals for executive roles or asking executives to work more hours—either of which would increase executive compensation.

In this subsection we test whether results change when we control for labor and capital. A firm will maximize profit by choosing inputs to equate the marginal rate of technical substitution with the input price ratio. That is, the firm equates the marginal product per dollar for each input. When the production technology is, for example, Cobb-Douglas or CES, then the marginal rate of technical substitution does not depend on output prices, so whether the price of oil is $20, $100, or $200 per barrel, the firm will still choose the same ratio of inputs, and controlling for labor and capital is equivalent to controlling for the overall level of demand for inputs.[25]

25. These controls could be viewed as a proxy for firm size. Studies have long shown that large firms pay their executives more than small firms, suggesting that the value of executive skill is increasing in firm size (see, e.g. Murphy, 1999). Tervio (2008), for example, concludes that high levels of top executive compensation are mainly due to firm size rather than the scarce ability of executives, and Gabaix and Landier (2008) find that the sixfold increase in CEO pay between 1980 and 2003 can be attributed to the sixfold increase in market capitalization of large U.S. companies during the same period. Thus a possible interpretation of the observed oil price effect could be that firm size increases with oil prices, increasing the marginal product of top executives.

We thus estimate a series of alternative specifications for our baseline IV regression in which we control for labor and capital. Specifically, we include log employee counts and the log of capital expenditures as control variables. We also show a version with quadratic functions of these variables. Results are shown in Table 2. Column 1 of this table re-creates our main specification from before, that is, Column 2 of Table 1. Column 2 limits the sample to the firm-year observations for which employee and capital expenditure data are available; this drops 2.6% of observations, but results are unchanged. Columns 3 and 4 add the labor and capital controls one-at-a-time. We see that each of these cuts the coefficient on log market value by about 20 to 30 percent, with the capital control having a bigger impact on the market value coefficient. When we include both capital and labor controls, the coefficient on log market value drops to 0.13. A quadratic version (Column 6) drops the market value coefficient a bit further.

Table 2: Controlling for Capital and Labor

	(1) IV	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
Log market value	0.19***	0.19***	0.15**	0.14**	0.13**	0.11**
	(0.05)	(0.06)	(0.05)	(0.05)	(0.05)	(0.05)
Log employees			0.17***		0.10***	0.10***
			(0.04)		(0.03)	(0.03)
Log capital expenditures				0.16***	0.12***	0.06
				(0.05)	(0.04)	(0.07)
Log employees, squared						0.00
						(0.01)
Log capital expenditures, squared						0.01
						(0.01)
First-stage F-statistic	90.59	96.12	89.21	72.49	76.13	84.49
Observations	4,673	4,553	4,623	4,603	4,553	4,553
Within R^2	0.51	0.51	0.52	0.51	0.52	0.52

Note: This table reports estimates and standard errors from six separate IV regressions, identical to Column 2 of Table 1, but with additional labor and capital control variables. Thus, the dependent variable in all regressions is log total annual compensation. Column 1 re-creates the result from Table 1. Column 2 limits the sample to the firm-year observations for which employee and capital expenditure data are available. Columns 3–6 include the additional controls. Compensation, market value, oil prices, and capital expenditures are normalized to 2016 dollars. Standard errors are two-way clustered by firm and by year. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Part of this oil price effect thus appears to be explained by firms choosing to buy more of *all* inputs when oil prices rise. However, a positive coefficient on instrumented log market value remains. Even controlling for other inputs, oil and gas executives see their compensation increase by 1.3 percent for every 10 percent rise in oil price-induced changes to market value. Thus the pattern seems not to be immediately explained by a simple model of firm profit maximization with a Cobb-Douglas or CES production function.

Next we perform a series of additional analyses aimed at better understanding the underlying mechanisms. Throughout, we proceed with the more conservative specification that controls for capital and labor.

4.4 Robustness to Alternative Specifications

In the Appendix, we show robustness of these main results to a variety of additional alternative specifications. First, because the identifying variation is time-series variation, we consider specifications of the regression of compensation on oil prices (Column 3 of Table 1) that use various alternative comparison groups. In these we can examine the *differential* impact of oil prices on compensation in oil and gas companies relative to other industries. These comparisons are designed to assess whether our macroeconomic controls are sufficient. We use manufacturing as a comparison group, then services. We also estimate a specification that includes *all* firms in Compustat, with a separate coefficient on the oil price for each industry. We find an impact of oil prices on compensation for oil and gas extraction firms and not for the comparison groups (Table A7).

We also report estimates (Tables A8 and A9) from a wide variety of alternative specifications for our IV specification. To be conservative, we focus on the specification that controls for capital and labor. We first show that the results do not rely on the macroeconomic controls that we use. Next we show that results with person fixed effects are very similar to the main results. This specification is reassuring because it suggests the results are not driven by compositional changes, such as endogenous entry and exit of executives.[26] Results are similar if we control for the national median level of executive compensation for all other industries. This is reassuring that general trends in executive compensation are not driving the result. Finally, the point estimate drops to only 0.03 when we include a quadratic trend, although it is not statistically different (5 percent level) from the original point estimate of 0.13. Overall, the results are somewhat sensitive to the time series controls that we include, which makes sense given that we are identifying off of annual time-series variation in oil prices.

Results are very similar using various subsets of the data: limiting the sample to firms on the S&P 1500; limiting the sample to a balanced panel; defining the oil and gas sector with SIC sector definitions; weighting by a time-invariant measure of firm size (assets); or using all executives reported in Compustat rather than just the top five.

We show that results are also very similar for alternative variable definitions: an alternative total compensation measure reported in Compustat; using Brent oil prices rather than WTI oil prices; or including the log of natural gas prices as an additional instrument. Thus, overall, the oil price effect is robust across alternative specifications.

Finally, in Table A10, we separate the sample into CEOs and non-CEOs, finding an oil price effect in both subsamples. While the highest paid executive at most of the companies in our sample is the CEO, the other top-five highest paid executives hold a wide range of positions include Chief Financial Officer (CFO), Chief Operating Officer (COO), and General Counsel. That we find essentially identical results for CEOs and non-CEOs is notable. Individuals in these other roles bring skills to the organization that are quite different from the CEO. If one thought that the observed oil price effect was reflecting increased demand for certain skills and ability bundles that are particularly valued when oil prices are high (Eisfeldt and Kuhnen, 2013), then it would not necessarily follow that the very different skills and backgrounds for CEOs and non-CEOs would respond equivalently. Still, as we discuss in Section 6, it is possible that during periods with high oil prices, the marginal product of all executive types rises.

5. ADDITIONAL RESULTS

5.1 Components of Pay

We next test for the oil price effect across the different components of executive compensation. We break total compensation down into five categories: salaries, stocks and stock options, bo-

26. Interestingly, the executives our sample tend to remain in the oil and gas business their entire career. In particular, for this subset of executives 93 percent of all person-year observations are for oil and gas companies. This is notable because it implies that the oil price effect is not driven by movement of executives between this and other sectors, for example, with highly productive executives moving into the industry during high oil price periods.

nuses, long-term incentive programs, and other pay (such as benefits). These separate components are interesting because of secular trends in compensation practices. Nationwide, the use of stocks and stock options expanded dramatically through the 1990s, then plateaued.[27] The growth of stocks and options continued for longer in the energy sector than in other sectors, and as a result, energy relies more heavily on these forms of compensation than other sectors. See Appendix Figure A2. At the end of our sample period, oil and gas executives received an average of over seventy percent of compensation from components other than salary.

We regress each component on market value, instrumented with the oil price. We use the same set of controls as in the primary results shown in Column 4 of Table 2, including capital and labor. We also add a dummy to account for changes in the way Compustat reports these variables after 2006 (details in the Appendix). All five components of pay have some observations equal to zero. We have dropped these from our log regressions, so these results are conditional on a non-zero value. In the Appendix (Table A11), we instead use an inverse hyperbolic sine transformation, thus retaining observations for which the value of that component of compensation is equal to zero.[28]

Table 3: For Which Components of Pay Is There an Oil Price Effect?

	(1) IV Salary	(2) IV Stocks and options	(3) IV Bonuses	(4) IV Other incentives	(5) IV Other pay
Log market value	–0.05*	0.08	0.59***	0.56***	0.21**
	(0.02)	(0.09)	(0.17)	(0.15)	(0.10)
First-stage F-statistic	84.37	101.75	56.65	175.26	84.31
Observations	4,546	3,934	3,083	1,729	4,443
Within R^2	0.34	0.48	0.27	0.38	0.14

Note: This table reports results from five separate IV regressions. The regressions are identical to Column 4 of Table 2, but with alternative dependent variables: each of five components of executive pay (logged). Standard errors are two-way clustered by firm and by year. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

Results are shown in Table 3. Salary (Column 1) does not reflect an oil price effect. If anything, salary appears to decrease with increases in market value driven by oil prices, but the coefficient is close to zero and only marginally statistically significant. Stocks and options (Column 2) have a positive point estimate. When crude oil prices cause market value to rise by 10 percent, the value of stocks and options granted to executives rises by 0.8 percent. That this coefficient is somewhat small and not statistically significant is perhaps surprising. However, we measure the value of stocks and options at the time they are granted, i.e. ex-ante not ex-post, so we have removed most of the mechanical positive correlation between realized pay and oil prices.

Moreover, in the Appendix (Table A12), we show that this appears to be at least in part because the value of stocks and options is tied more closely to *annual average* oil prices than to the December price that we use in our main specifications. This is intuitive if, for instance, these are granted earlier in the year, since their value is computed at the time they are granted. In our sample,

27. One of the reasons nationwide for the rise in the use of stocks and options was legislation in 1993 that limited tax deductions to the first $1 million of salary (Rose and Wolfram, 2002). In the Appendix, we show with histograms that this appears to be somewhat binding for CEO pay, although not for other executive pay. We also show in the Appendix (Table A10) that the oil price effect appears to be approximately equally prevalent with CEOs and other executives, so this tax rule does not appear to be a primary driver of this pattern.

28. The sample size for these regressions should thus match that of Column 5, Table 2. We actually have one more observation in the inverse hyperbolic sine regressions, because there is one executive with zero reported overall compensation.

several years saw enough within-year variability in oil prices for the December and annual average prices to differ by at least 25 percent.[29]

Columns 3–5 examine bonuses, other incentives, and other pay. Together, these make up around 30 percent of pay (see Table A3). When crude oil prices increase market value by 10 percent, bonuses (Column 3) rise by 5.9 percent and other incentives by 5.6 percent, both statistically significant at the one percent level. In the Appendix, we show that the oil price effect for bonuses is still evident when zeroes are included; in fact, the point estimate is substantially larger. The point estimate for other incentives drops, but it is not statistically different from the point estimate shown in Table 3. This evidence on bonuses and long-term incentives is consistent with the metrics that firms report using when setting pay: one recent industry report noted that common metrics for annual bonus payments are production and the value of reserves, both of which are positively correlated with oil prices; and that a common metric for long-term incentive payments is total shareholder return (Alvarez and Marsal, 2018).

For none of these latter components (Columns 3–5) do we reject the null that compensation responds just as much to changes driven by oil prices as it does to generic changes in company value. That is, we fail to reject that the IV estimates are identical to point estimates from analogous OLS regressions, shown in the Appendix (Table A13). And overall, the results in Table 3 indicate that the oil price effect is not simply a mechanical result stemming from stock and option awards.

5.2 Governance

We next test whether the oil price effect varies with firm governance. Our analysis of governance follows B&M, who argue that companies with strong independent leadership from the board should be better able to prevent executives from co-opting the pay process, suggesting that "we should expect more pay for luck in the poorly governed firms" (p. 918). Whereas B&M test governance using relative industry performance as a measure of luck, our analysis is the first to test the effect of governance using oil prices. We leverage two indicators of firm governance.

First, from the Compustat data, we observe whether each of the firm's top-five executives also sits on the board of directors. We construct a variable equal to the proportion of the top five executives that do not sit on the board. Higher levels of this measure thus indicate better governed companies with more independence on the board. The mean in our sample is 0.64. That is, on average about two-thirds of the top-five executives do not sit on the board. Table A3 includes descriptive statistics for all our governance measures.

Second, from ISS data, we observe whether each of the firm's board members has any sort of insider status, for instance because the member is an employee or the close relative of an employee. From this, we construct a variable equal to the proportion of the board members that are not insiders. Again, a higher value indicates a higher level of governance. The mean in our sample is 0.75.[30] Thus on average about three-quarters of board members are not insiders.

For each of the measures, we focus on a time-invariant measure, equal to the mean value observed across all years for each firm. We show results in the Appendix (Table A14) that leverage variation across time within a firm, although we note it is less clear that this variation is exogenous.

We estimate a series of regressions in which we augment our standard IV specification with an interaction of market value with each governance measure. The coefficient on this interac-

29. Another possible explanation for the small coefficient for stocks and options is heterogeneity in the way firms compute the grant-date fair value of stocks and options, detailed in Coles et al. (2014).

30. Missing values arise both because the ISS data do not cover the years 1992-1996, and because they do not cover the smallest firms in our sample.

tion thus tells us how the oil price effect differs for better-governed firms. We display 2SLS results, matching the specification in Column 4 of Table 2.[31] We have two first-stage equations: we instrument for both the log market value and for the interaction of market value with the governance indicator. Our two instruments are log oil prices and the interaction of log oil prices with the governance measure. We show results both for total compensation and separately for bonuses and cash incentives. We analyze this second category separately because this is the component of pay over which boards tend to have the most discretion, and because our analysis earlier showed a significant oil price effect in this category.

Results are shown in Table 4.[32] Overall, the results suggests a smaller oil price effect at better-governed companies. In particular, the interaction of market value and governance has a negative sign in all four specifications, indicating a smaller oil price effect with more independents on the board. Governance seems to particularly matter for bonuses and cash incentives, the components of pay that tend to be most discretionary. The first-stage instrument F-statistics range from 33 to 59, indicating that we do not have a weak instruments problem, though the interaction term is statistically significant in only one of the four specifications. Finally, we also show using OLS (see Appendix Table A15), that executive compensation in well-governed firms is less tied overall to market value, consistent with these firms using more nuanced forms of pay-for-performance.

Table 4: Does Good Governance Reduce the Oil Price Effect?

	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS
	Panel A: All Compensation		Panel B: Bonuses and Cash Incentives	
Log market value	0.17	0.44*	0.96***	1.37***
	(0.20)	(0.23)	(0.22)	(0.36)
Log m.v. X Portion of execs not on board	–0.07		–0.55	
	(0.27)		(0.35)	
Log m.v. X Portion of board non-insiders		–0.46		–1.09**
		(0.28)		(0.44)
First-stage F-statistic	33.25	38.06	43.33	58.95
Observations	4,553	3,406	4,117	3,133
Within R^2	0.52	0.56	0.43	0.50

Note: This table reports estimates and standard errors from four separate 2SLS regressions. The dependent variable in Columns 1 and 2 is log total annual compensation; and log bonuses and non-stock incentives in Columns 3 and 4. In all columns, we have two first stage equations: we instrument for log market value and the governance interaction with both log crude oil prices, and with log oil price interacted with the governance measure. All regressions include the same controls as in Column 4 of Table 2, including capital and labor. Columns 3 and 4 additionally include a dummy for data reporting changes beginning in 2006. The sample size varies across columns because of differential data availability across the different governance measures. Compensation, market value, and oil prices are normalized to 2016 dollars. Standard errors are two-way clustered by firm and by year. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

The magnitudes are suggestive of an economically significant effect across all columns. For instance, removing one of the top five executives from the board removes 8 percent of the oil price effect in Column 1 and 11 percent of the oil price effect in Column 3.[33] Similarly, in Columns

31. We also show OLS results in Appendix Table A15.

32. The sample size drops in Columns 2–4 because of missing data for the non-insiders variable and because of some zero values for bonuses and cash incentives. In the Appendix, we show estimates of the overall oil price effect for the sub-sample that has complete governance variable coverage and non-zero bonus and cash incentives values.

33. Removing one of the top five executives decreases the "portion not on the board" by 20 percentage points, and $-0.07 \cdot 0.20 = -0.014$, or 8 percent of the log market value effect on total compensation of 0.17 if all executives are on the board.

2 and 4, replacing one insider from a nine-member board removes 11 percent of the oil price effect on total compensation and 9 percent of the effect for bonuses and cash incentives.[34] Results with time-varying measures are presented in Appendix Table 14. We again estimate negative coefficients in Columns 2–4, albeit with weak first stages.

Still, these results must be interpreted cautiously. We do not have exogenous variation in governance quality. One might be concerned, therefore, that these governance results are driven by other differences in firms that are correlated with the board composition measures that we have used here. For instance, board composition might be correlated with firm size. To explore this possibility, we estimate alternative specifications (see Appendix Table A16) that control for firm size. Specifically, we include interactions of the market value and oil price variables with a time-invariant measure of firm size, average log employees over our sample period. The coefficients on the governance variables are robust to this additional control, suggesting that heterogeneity along this dimension is not driving the governance results. Nonetheless, without exogenous variation in governance quality it is difficult to make strong causal statements on the basis of these regressions.

5.3 Asymmetry

We next explore whether the oil price effect is symmetric with regard to increases and decreases of the firm's market value. We estimate augmented versions of our standard specifications, separating the coefficients on market value and oil prices into two coefficients each. In particular, we construct an indicator variable equal to one if market value is higher (in real terms) than the previous year. We then interact the log of market value with both this indicator variable and with one minus this indicator variable. The regression thus tells us whether the sensitivity of executive pay to firm performance is the same when the firm's value is rising versus falling.

Results are shown in Table 5. We estimate an economically large and statistically significant difference across the two coefficients. In Column 1, the sensitivity of executive compensation to market value is around one third higher when the firm's value is rising (in real terms). Similarly, we construct an indicator variable for whether oil prices are rising or falling, then include interactions of this indicator variable with the log oil price. Results are shown in Column 2 of Table 5.[35] We again find an economically and statistically significant difference (p-value< 0.01). In particular, the point estimate for increasing oil prices is *three times* larger than the point estimate for decreasing oil prices.

Finally, we estimate a 2SLS version, in which log market value is interacted with the oil price rising/falling dummy. Thus we have two first stage equations, in which our instruments are log crude oil prices if oil prices are rising, and log crude oil prices if oil prices are falling.[36] Results, in Column 3, again indicate an asymmetric oil price effect. In this specification the difference between the two coefficients is not statistically significant (p-value$=0.07$), but the point estimate for increases is again larger than the point estimate for decreases.

34. Removing one insider decreases the "portion non-insiders" by 11.1 percentage points, and $-0.46 \cdot 0.111 = -0.05$, or 11 percent of the log market value effect on total compensation of 0.44 if all board members are insiders.

35. Note that the number of observations drops because we do not observe market value in 1991, and thus we are unable to construct the "rising/falling" dummy for market value in 1992. In the Appendix, we show estimates of the overall oil price effect for the sub-sample that begins in 1993.

36. Ideally, we would estimate a 2SLS version, with the indicator variable for *market value* rising or falling, using three first-stage equations (i.e., instrumenting with the oil price rising indicator variable). However, we find that we lose first stage power when we have three endogenous variables. Thus we interact the market value variable with the oil price rising/falling indicator variable directly.

Table 5: Is the Oil Price Effect Asymmetric?

	(1) OLS	(2) OLS	(3) 2SLS
Log market value, if rising	0.28*** (0.04)		
Log market value, if falling	0.21*** (0.04)		
Log oil price, if rising		0.20*** (0.07)	
Log oil price, if falling		0.06 (0.06)	
Log market value, if oil price rising			0.29*** (0.10)
Log market value, if oil price falling			0.18* (0.10)
p-value, rising versus falling	0.004	0.006	0.067
First-stage F-statistic			17.92
Observations	4,357	4,357	4,357
Within R^2	0.53	0.51	0.52

Note: This table reports estimates and standard errors from three separate regressions. The dependent variable in all regressions is log total annual compensation. We also report the p-value for a test of equality of the “if rising” versus “if falling” coefficients. All regressions include the same controls as in Column 4 of Table 2. Column 1 also includes (not shown) a dummy variable to indicate whether market value is rising; Columns 2 and 3 include a dummy variable to indicate whether oil prices are rising. Column 3 has two first stage equations: “log market value, if oil price rising” and “log market value, if oil price falling” are instrumented with “log oil price, if oil price rising” and “log oil price, if oil price falling.” Compensation, market value, and oil prices are normalized to 2016 dollars. Standard errors are two-way clustered by firm and by year. *** Statistically significant at the 1% level; ** 5% level; * 10% level.

In the Appendix we also test whether this asymmetry differs for executives with compensation levels above and below the median. Oil prices matter for both groups, and asymmetry is present for both groups, but the overall sensitivity to oil prices is stronger for executives paid above their peer group median.[37]

6. DISCUSSION

6.1 Measuring the Degree of Filtering

One of the reasons the oil price effect is interesting is that it is inconsistent with a simple contracting model along the lines of Holmstrom and Milgrom (1987). Executives are risk averse, so the optimal contract would filter out oil prices and other observable measures of luck to make contracts less risky (Holmstrom and Milgrom, 1987). The observed oil price effect (e.g. Table 2), indicates that companies are not completely filtering out oil prices. However, are companies doing partial filtering? To what *extent* are companies filtering out oil prices?

37. This evidence of asymmetry both above and below the median differs from Bizjak et al. (2008), which finds that “pay-for-luck” is driven by CEOs with pay *below* the peer group median. Their context is different as their study uses industry returns as a measure of luck, rather than oil prices, but their discussion is nonetheless quite relevant. In particular, they argue that their evidence is consistent with “competitive benchmarking being used for retention motives” (p. 164), rather than rent extraction: “under [rent extraction], the asymmetry in pay to luck should be most prevalent in the group of highly paid CEOs because these are the ones who have likely captured the pay process” (p. 165).

In this section we perform a back-of-the-envelope calculation aimed at measuring the degree of filtering. We are not arguing that this type of contracting model is the only way to think about executive compensation. Indeed, later in this section we continue our discussion of alternative models. But, we do view this back-of-the-envelope as a useful exercise for evaluating the degree to which these data are consistent or not consistent with one particularly prominent model.

Holmstrom and Milgrom (1987) derives the optimal compensation contract under assumptions about the executive's risk aversion and the process controlled by the executive. This model, and the literature using it (such as B&M and Aggarwal and Samwick, 1999b), abstract away from the potential impact of luck on executive marginal product. Using the notation in Aggarwal and Samwick (1999b), the compensation contract is $w = \alpha_0 + \alpha_1 \pi + \alpha_2 \theta$, i.e. compensation w is a linear function of firm performance π and observable luck θ. The optimal weight on observable luck θ is negative; i.e., in the optimal contract, luck is filtered out.

Aggarwal and Samwick (1999b) show that in the optimal contract the weights on firm performance and observable luck have the ratio $\frac{\alpha_2^*}{\alpha_1^*} = -\beta$, where β is the coefficient from a regression of firm performance on luck. That is, the more firm performance is driven by observable luck, the greater the relative (negative) weight that should be put on luck in the optimal contract. It is worth noting that their model assumes that the executive is risk averse, but that this ratio does not depend on the exact *level* of risk aversion.

We can examine the extent to which firms filter by calculating the ratio $\frac{\alpha_2}{\alpha_1}$ at the theoretical optimum versus in practice. The optimal ratio is simply 0.99, the first-stage coefficient in Table A5—the contribution of luck to firm performance.[38] To calculate the *actual* ratio used by the typical firm, we estimate a regression with log compensation as the dependent variable and both log market value and log oil prices as explanatory variables. Full results are given in Table A18. When we include the same controls as in Table 1, we estimate a coefficient on log market value of 0.31 and a coefficient on log oil price of –0.12.[39] The ratio of the coefficients is $\frac{\alpha_2}{\alpha_1} = -0.12 / 0.31 = -0.4$.

Thus in practice compensation practices are such that firms put a negative weight on oil prices that is only 0.4 times the weight put on firm performance. In contrast, the optimal contract would instead put a negative weight on oil prices that is essentially identical to the weight on firm performance. That is, for this back-of-the-envelope calculation, which abstracts away from executive marginal product effects, we see only 40 percent of the optimal level of filtering being done in practice. Not only do we not find complete filtering, but it appears that about 60 percent of observable luck is not filtered. U.S. oil and gas firms compensate their executives in a way that falls well short of this theoretical optimum.

6.2 Rent Extraction vs Shareholder Value

We have presented a number of results that some papers in the previous literature would have characterized as consistent with rent extraction on the part of executives. But how could rent extraction occur in equilibrium? In particular, if oil and gas executives are overpaid when oil prices

38. The specification in Table A5 includes additional controls, but we show in the Appendix (Table A19) that the 0.99 estimate is not very sensitive to the controls.

39. In the Appendix, we consider two sets of robustness checks. First, results are robust to including additional measures of firm performance; see Table A20. Next, we consider a related specification used in Garvey and Milbourn (2003). Table A21 shows that their regression yields results that are equivalent to a non-linear combination of the parameters that we estimate.

are high, why are outside hires not able to arbitrage this away? One possible explanation is that executives in oil and gas have industry-specific human capital. As documented by Cremers and Grinstein (2014), there is relatively little movement of executives between oil and gas and other sectors. This lack of movement between sectors is evident in our sample as well, with the vast majority of executives in our dataset spending their entire careers in oil and gas (see footnote 26). Thus, it seems plausible that executives are able to exert market power.

Still, it is difficult to make strong conclusions about the mechanisms underlying the observed oil price effect. As we mentioned in the introduction, much of the broader literature on executive compensation is aimed at attempting to reconcile the "rent extraction" and "shareholder value" views. Despite a voluminous literature in this area, it remains very much an open question as to which view best explains the empirical evidence on executive compensation. Part of the challenge is that the two views are not mutually exclusive, with both forces impacting compensation to varying degrees across firms and over time. In addition, there are many alternative approaches for modeling "shareholder value," making it hard to make definitive statements about what is (or what is not) consistent with the best interests of shareholders.

In particular, just because the oil price effect is inconsistent with standard contracting models, need not imply rent extraction. There are alternative approaches for modeling shareholder value in which one would expect executive compensation to increase with oil prices. For example, as we mentioned earlier, under a simple model of profit maximization, firms buy more inputs when output prices go up. This would imply increased demand for executives, and thus increased compensation, during periods with high oil prices. This induced input demand indeed appears to be part of the explanation. As we showed earlier, the oil price effect gets smaller after controlling for capital and labor. This attenuation is consistent with higher executive compensation reflecting, at least in part, higher demand for all inputs.

There are also alternative approaches for modeling shareholder value which focus on retention concerns. Oyer (2004), for example, considers a model in which employees' outside opportunities are correlated with firm performance. Under such a model, executive compensation could increase during high oil prices, not to incentivize effort, but in order to prevent executives from moving to other firms. As we discussed earlier, retention concerns are potentially somewhat less compelling in the market for oil and gas simply because it is a relatively insular market with little movement of executives between this and other sectors, but this paper along with subsequent work in this area nonetheless raises another potential mechanism through which compensation would be correlated with market conditions (see, e.g. Lustig et al., 2011; Cremers and Grinstein, 2014; Campbell and Thompson, 2015).[40]

We find that the oil price effect is pervasive across both CEOs and non-CEOs, and across several individual components of executive compensation. These results are inconsistent with standard contracting models, but could be rationalized by increased input demand or retention concerns. During periods with high oil prices, firms need more effort from non-CEOs, just like they need more effort from the CEO, and outside options matter both for CEOs and non-CEOs. Moreover, there is no clear reason why increased compensation would take the form (or not take the form) of any of the individual components of compensation. Bonuses and cash incentives tend to be the most dis-

40. Yet another potential mechanism comes from DeMarzo et al. (2012) which describes a dynamic agency model in which financing constraints become less binding after positive exogenous shocks: "The intuition is that the marginal cost of compensating the agent is lower when profitability is high because relaxing the agency problem is more valuable when profitability is high. This result may help to explain the empirical importance of absolute, rather than relative, performance measures for executive compensation."

cretionary, so they would seem to be particularly susceptible to rent extraction but these components may also simply be the best way to scale up compensation during oil price shocks. Thus, while these additional results are interesting, they fail to sharply distinguish between rent extraction and shareholder value.

Neither increased input demand nor retention concerns can explain, in their simplest form, the asymmetry results. Under a shareholder value model in which executives are paid their marginal product, we would expect the pattern to be approximately symmetric. Demand for executives would increase with rising oil prices, but then decrease proportionally with falling oil prices. Instead, we find a clear asymmetry. The OLS results in Column 1 of Table 5 are particularly striking, because they indicate that the impact of market value is not explained by a contract in which executive pay is a simple linear function of market value.[41] At first glance, this asymmetric pattern would appear to be more consistent with rent extraction, with executives better able to extract value from the firm when oil prices are rising.

However, there are again alternative approaches for modeling shareholder value that could explain this asymmetric pattern. For example, there could be something inherently asymmetric about capacity investments. For whatever reason, if executive effort is more important for capacity expansions than capacity reductions, that could generate the asymmetric pattern observed here. It could be, simply, that executives are most valuable during the "growth" phase for oil and gas companies, so that compensation increases rapidly during periods of rising oil prices. Thus while the asymmetry initially appears to point toward rent extraction, it is relatively easy to augment the shareholder value model with additional features to match the observed pattern. This underscores the challenge in the broader literature with formulating a test that would allow one to sharply distinguish between rent extraction and shareholder value.

7. CONCLUSION

In the years since B&M, executive compensation has changed in a number of important ways. There has been a dramatic increase in the use of stock options followed by a partial reversal of this trend. There was a major recession, accompanied by an increase in public scrutiny of executive compensation, and regulatory rules regarding compensation disclosure and shareholder involvement have tightened. Finally, the last decade has seen a transformation of the oil and gas industry in the United States, with the rise of hydraulic fracturing and the entry of dozens of new companies. In these changing times, it is striking that the fortunes of oil and gas executives still hinge so dramatically on a single global commodity price.

We find strong evidence that U.S. oil and gas executive compensation is closely tied to oil prices. In our preferred specification, a 10% increase in firm value driven by oil prices is associated with a 2% increase in executive compensation. The bulk of this effect—a 1.3% increase—remains even after controlling for capital and labor. Across specifications, we cannot rule out that executive

41. This pattern could arise through the use of bonuses that are bounded below at zero. This is still consistent with the rent extraction interpretation, however, since the use of such bonuses is a choice by compensation committees. Asymmetry could also arise if bonus criteria change with oil prices. Morse et al. (2011) finds that powerful U.S. CEOs induce board-of-directors to shift the weight on performance measures toward better performing measures, thereby "rigging" incentive pay. Anecdotal evidence suggests this indeed happens in the oil industry. A recent article describes how bonus criteria at Comstock Resources changed over time, from quantitative metrics in years with high oil prices to qualitative metrics (e.g. "leadership development" and "execution of strategic plan") in years with falling oil prices. Source: Denning, Liam. "A Tiny Gas Firm's Big Lesson on Bosses' Pay." 4 June 2018. *Bloomberg Opinion*, https://www.bloomberg.com/view/articles/2018-06-04/comstock-resources-a-lesson-on-oil-gas-executive-pay.

compensation responds just as much to changes in firm value driven by oil prices as it does to generic changes in firm value. The related literature has termed this effect "pay-for-luck" but we more agnostically refer to as the "oil price effect."

Moreover, our analysis supports the following additional conclusions. First, we document the oil price effect for both CEOs and non-CEOs. Second, we show that the oil price effect is pervasive across different components of pay, but is particularly pronounced for bonuses and other incentives. Third, we find a smaller oil price effect for better-governed firms. Results are similar across measures of governance, and robust to controlling for firm size. Fourth, we show that the oil price effect is asymmetric, with executive compensation rising more during oil price increases than it falls during oil price decreases. These additional analyses all follow the broader executive compensation literature, but in all four cases we are the first to perform such tests in the context of oil prices.

ACKNOWLEDGMENTS

We are grateful to Severin Borenstein, Thom Covert, Todd Gerarden, Matthew Kahn, Ryan Kellogg, Daniel Raimi, Martin Schmalz, and seminar participants at Cornell, UC Berkeley, UC Davis, and Wharton for helpful comments. We have not received any financial compensation for this project, nor do we have any financial relationships that relate to this research.

REFERENCES

Abowd, John M. and David S. Kaplan (1999). "Executive Compensation: Six Questions That Need Answering." *Journal of Economic Perspectives* 13(4): 145–168. https://doi.org/10.1257/jep.13.4.145.

Aggarwal, Rajesh K. and Andrew A. Samwick (1999a). "Executive Compensation, Strategic Competition, and Relative Performance Evaluation: Theory and Evidence." *Journal of Finance* 54(6): 1999–2043. https://doi.org/10.1111/0022-1082.00180.

Aggarwal, Rajesh K. and Andrew A. Samwick (1999b). "The Other Side of the Trade-off: The Impact of Risk on Executive Compensation." *Journal of Political Economy* 107(1): 65–105. https://doi.org/10.1086/250051.

Albuquerque, Ana M., Gus De Franco, and Rodrigo S. Verdi (2013). "Peer Choice in CEO Compensation." *Journal of Financial Economics* 108: 160–181. https://doi.org/10.1016/j.jfineco.2012.10.002.

Alvarez and Marsal (2018). *Oil and Gas Exploration and Production (E&P) Incentive Compensation Report: Analysis of Compensation Arrangements Among the Largest U.S. E&P Companies*.

Antle, Rick and Abbie Smith (1986). "An Empirical Investigation of the Relative Performance Evaluation of Corporate Executives." *Journal of Accounting Research* 24(1): 1–39. https://doi.org/10.2307/2490802.

Axelson, Ulf and Baliga, Sandeep (2009). "Liquidity and Manipulation of Executive Compensation Schemes." *Review of Financial Studies*, 22(10): 3907–3939. https://doi.org/10.1093/rfs/hhn095.

Bebchuk, Lucian A., Yaniv Grinstein, and Urs Peyer (2010). "Lucky CEOs and Lucky Directors." *Journal of Finance* 65(6): 2363–2401. https://doi.org/10.1111/j.1540-6261.2010.01618.x.

Bebchuk, Lucian Arye and Jesse M. Fried (2003). "Executive Compensation as an Agency Problem." *Journal of Economic Perspectives* 17(3): 71–92. https://doi.org/10.1257/089533003769204362.

Bell, Brian and John Van Reenen (2016). "CEO Pay and the rise of Relative Performance Contracts: A Question of Governance?" *NBER Working Paper No. 22407*. https://doi.org/10.3386/w22407.

Bertrand, Marianne (2009). "CEOs." *Annual Review of Economics*, 121–149. https://doi.org/10.1146/annurev.economics.050708.143301.

Bertrand, Marianne and Sendhil Mullainathan (2001). "Are CEOs Rewarded for Luck? The Ones without Principals Are." *Quarterly Journal of Economics* 116(3): 901–932. https://doi.org/10.1162/00335530152466269.

Bizjak, John M., Michael L. Lemmon, and Lalitha Naveen (2008). "Does the Use of Peer Groups Contribute to Higher Pay and Less Efficient Compensation?" *Journal of Financial Economics* 90(2): 152–168. https://doi.org/10.1016/j.jfineco.2007.08.007.

Bolton, Patrick, Jose Scheinkman, and Wei Xiong (2006). "Executive Compensation and Short-Termist Behaviour in Speculative Markets." *Review of Economic Studies* 73, 577–610. https://doi.org/10.1111/j.1467-937X.2006.00388.x.

Borenstein, Severin and Ryan Kellogg (2014). "The Incidence of an Oil Glut: Who Benefits from Cheap Crude Oil in the Midwest?" *The Energy Journal* 35(1): 15–33. https://doi.org/10.5547/01956574.35.1.2.

Campbell, T. Colin and Mary Elizabeth Thompson (2015). "Why Are CEOs Paid for Good Luck? An Empirical Comparison of Explanations for Pay-for-Luck Asymmetry." *Journal of Corporate Finance* 35, 247–264. https://doi.org/10.1016/j.jcorpfin.2015.09.006.

Chaigneau, Pierre and Nicolas Sahuguet (2018). "The Effect of Monitoring on CEO Compensation in a Matching Equilibrium." *Journal of Financial and Quantitative Analysis* 53(3): 1297–1339. https://doi.org/10.1017/S0022109017001065.

Coles, Jeffrey L., Naveen D. Daniel, and Lalitha Naveen (2014). "Co-opted Boards." *Review of Financial Studies*, 27(6): 1751–1796. https://doi.org/10.1093/rfs/hhu011.

Cremers, K.J. Martijn and Yaniv Grinstein (2014). "Does the Market for CEO Talent Explain Controversial CEO Pay Practices?" *Review of Finance* 18, 921–960. https://doi.org/10.1093/rof/rft024.

Danthine, Jean-Pierre and John B. Donaldson (2015). "Executive Compensation: A General Equilibrium Perspective." *Review of Economic Dynamics* 18(2): 269–286. https://doi.org/10.1016/j.red.2014.04.003.

DeMarzo, Peter M., Michael J. Fishman, Zhiguo He, and Neng Wang (2012). "Dynamic Agency and the q Theory of Investment." *Journal of Finance* 67(6): 2295–2340. https://doi.org/10.1111/j.1540-6261.2012.01787.x.

Edmans, Alex and Xavier Gabaix (2009). "Is CEO Pay Really Inefficient? A Survey of New Optimal Contracting Theories." *European Financial Management* 15(3): 486–496. https://doi.org/10.1111/j.1468-036X.2009.00500.x.

Edmans, Alex and Xavier Gabaix (2016). "Executive Compensation: A Modern Primer." *Journal of Economic Literature* 54(4): 1232–87. https://doi.org/10.1257/jel.20161153.

Edmans, Alex, Xavier Gabaix, and Dirk Jenter (2017). "Executive Compensation: A Survey of Theory and Evidence." *Handbook of the Economics of Corporate Governance*, 1. https://doi.org/10.1016/bs.hecg.2017.11.010.

Eisfeldt, Andrea L. and Camelia M. Kuhnen (2013). "CEO Turnover in a Competitive Assignment Framework." *Journal of Financial Economics* 109(2): 351–372. https://doi.org/10.1016/j.jfineco.2013.02.020.

Energy Information Administration (2018). U.S. Remains the World's Top Producer of Petroleum and Natural Gas Hydrocarbons. *Today in Energy*, May.

Gabaix, Xavier and Augustin Landier (2008). "Why Has CEO Pay Increased So Much?" *Quarterly Journal of Economics* 123(1): 49–100. https://doi.org/10.1162/qjec.2008.123.1.49.

Garen, John E. (1994). "Executive Compensation and Principal-Agent Theory." *Journal of Political Economy* 102(6): 1175–1199. https://doi.org/10.1086/261967.

Garvey, Gerald T. and Todd T. Milbourn (2003). "Incentive Compensation When Executives Can Hedge the Market: Evidence of Relative Performance Evaluation in the Cross Section." *Journal of Finance* 58(4): 1557–1581. https://doi.org/10.1111/1540-6261.00577.

Garvey, Gerald T. and Todd T. Milbourn (2006.) "Asymmetric Benchmarking in Compensation: Executives are Rewarded for Good Luck but Not Penalized for Bad." *Journal of Financial Economics* 82: 197–225. https://doi.org/10.1016/j.jfineco.2004.01.006.

Gibbons, Robert and Kevin J. Murphy (1990). "Relative Performance Evaluation for Chief Executive Officers." *Industrial and Labor Relations Review* 43: 30–S–51–S. https://doi.org/10.1177/001979399004300303.

Gold, Russell (2014). *The Boom: How Fracking Ignited the American Energy Revolution and Changed the World*. Simon and Schuster.

Gong, Guojin, Laura Yue Li, and Jae Yong Shin (2011). "Relative Performance Evaluation and Related Peer Groups in Executive Compensation Contracts." *Accounting Review*, 86(3): 1007–1043. https://doi.org/10.2308/accr.00000042.

Gopalan, Radhakrishnan, Todd Milbourn, and Fenghua Song (2010). "Strategic Flexibility and the Optimality of Pay for Sector Performance." *Review of Financial Studies* 23(5): 2060–2098. https://doi.org/10.1093/rfs/hhp118.

Hall, Brian J. and Kevin J. Murphy (2002). "Stock Options for Undiversified Executives." *Journal of Accounting and Economics* 33(1): 3–42. https://doi.org/10.1016/S0165-4101(01)00050-7.

Harford, Jarrad and Kai Li (2007). "Decoupling CEO Wealth and Firm Performance: The Case of Acquiring CEOs." *Journal of Finance* 62(2): 917–949. https://doi.org/10.1111/j.1540-6261.2007.01227.x.

Hausman, Catherine and Ryan Kellogg (2015). "Welfare and Distributional Implications of Shale Gas." *Brookings Papers on Economic Activity* Spring: 71–125. https://www.brookings.edu/wp-content/uploads/2016/07/HausmanKellogg_FullText.pdf.

Himmelberg, Charles P. and R. Glenn Hubbard (2000). "Incentive Pay and the Market for CEOs: An Analysis of Pay-for-Performance Sensitivity." *Mimeo*. https://doi.org/10.2139/ssrn.236089.

Hoffmann, Florian and Sebastian Pfeil (2010). "Reward for Luck in a Dynamic Agency Model." *Review of Financial Studies* 23(9): 3329–3345. https://doi.org/10.1093/rfs/hhq062.

Holmstrom, Bengt (1979). "Moral Hazard and Observability." *Bell Journal of Economics* 10(1): 74–91. https://doi.org/10.2307/3003320.

Holmstrom, Bengt and Paul Milgrom (1987). "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55(2): 303–328. https://doi.org/10.2307/1913238.

Jenter, Dirk and Fadi Kanaan (2015). "CEO Turnover and Relative Performance Evaluation." *Journal of Finance* 70(5): 2155–2183. https://doi.org/10.1111/jofi.12282.

Jin, Li (2002). "CEO Compensation, Diversification, and Incentives." *Journal of Financial Economics* 66(1): 29–63. https://doi.org/10.1016/S0304-405X(02)00150-2.

Joskow, Paul, Nancy Rose, and Andrea Shepard (1993). "Regulatory Constraints on CEO Compensation." *Brookings Papers on Economic Activity. Microeconomics* 1993(1): 1–72. https://doi.org/10.2307/2534710.

Joskow, Paul L., Nancy L. Rose, and Catherine D. Wolfram (1996). "Political Constraints on Executive Compensation: Evidence from the Electric Utility Industry." *RAND Journal of Economics* 165–182. https://doi.org/10.2307/2555797.

Kaplan, Steven N. and Bernadette A. Minton (2012). "How Has CEO Turnover Changed?" *International Review of Finance* 12(1): 57–87. https://doi.org/10.1111/j.1468-2443.2011.01135.x.

Kleibergen, Frank and Richard Paap (2006). "Generalized Reduced Rank Tests Using the Singular Value Decomposition." *Journal of Econometrics* 97–126. https://doi.org/10.1016/j.jeconom.2005.02.011.

Kline, Patrick (2008). "Understanding Sectoral Labor Market Dynamics: An Equilibrium Analysis of the Oil and Gas Field Services Industry." *Mimeo*.

Lewellen, Stefan (2015). *Executive Compensation and Industry Peer Groups*. Working Paper, London Business School, https://sites.google.com/site/stefanlewellen/.

Lustig, Hanno, Chad Syverson, and Stijn Van Nieuwerburgh (2011). "Technological Change and the Growing Inequality in Managerial Compensation." *Journal of Financial Economics* 99(3): 601–627. https://doi.org/10.1016/j.jfineco.2010.09.007.

Morse, Adair, Vikram Nanda, and Amit Seru (2011). "Are Incentive Contracts Rigged by Powerful CEOs?" *Journal of Finance* 66(5): 1779–1821. https://doi.org/10.1111/j.1540-6261.2011.01687.x.

Muehlegger, Erich and Richard L. Sweeney (2019). "Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom." *NBER Working Paper 24025*.

Murphy, Kevin J. (1999). "Executive Compensation." *Handbook of Labor Economics* 3: 2485–2563. https://doi.org/10.1016/S1573-4463(99)30024-9.

Murphy, Kevin J. (2013). "Executive Compensation: Where We Are, and How We Got There." *Handbook of the Economics of Finance* 211–356. https://doi.org/10.1016/B978-0-44-453594-8.00004-5.

Noe, Thomas H. and Michael J. Rebello (2012). "Optimal Corporate Governance and Compensation in a Dynamic World." *Review of Financial Studies* 25(2): 480–521. https://doi.org/10.1093/rfs/hhr131.

Oyer, Paul (2004). "Why Do Firms Use Incentives That Have No Incentive Effects?" *Journal of Finance* 59(4): 1619–1649. https://doi.org/10.1111/j.1540-6261.2004.00674.x.

Rose, Nancy L. and Catherine Wolfram (2002). "Regulating Executive Pay: Using the Tax Code to Influence Chief Executive Officer Compensation." *Journal of Labor Economics* 20(2): S138–S175. https://doi.org/10.1086/338677.

Tervio, Marko (2008). "The Difference that CEOs Make: An Assignment Model Approach." *American Economic Review* 98(3): 642–68. https://doi.org/10.1257/aer.98.3.642.

Wade, James B., Joseph F. Porac, and Timothy G. Pollock (1997). "Worth, Words, and the Justification of Executive Pay." *Journal of Organizational Behavior* 18(SI): 641–664. https://onlinelibrary.wiley.com/doi/10.1002/%28SICI%291099-1379%28199711%2918%3A1%2B%3C641%3A%3AAID-JOB910%3E3.0.CO%3B2-M.

Yergin, Daniel (2011). *The Quest: Energy, Security, and the Remaking of the Modern World*. Penguin.