

ORIGINAL ARTICLE

Why has academic research had so little impact on US executive pay practice?

Stephen F. O'Byrne

Shareholder Value Advisors

Correspondence

Stephen F. O'Byrne, Shareholder Value Advisors.

Email: sobyrne@valueadvisors.com

INTRODUCTION

The now voluminous academic literature on executive pay shows surprisingly little awareness of the history of US public company pay practices. There is today a nearly universal misconception, even among finance and accounting scholars, that there was little incentive pay at US public companies until stock option use exploded in the 1980s. What's more, there has been virtually no recognition of the widespread use of incentive plan fixed-sharing formulas by US public companies in the first half of the 20th century. General Motors, as the most famous example, had a pay plan that made the bonus shared by the entire senior management team (nearly 3000 managers in 1929) equal to 10% of GM's total after-tax profit in excess of a 7% return on capital. GM continued to use this formula, with only slight variations, from its start in 1918 until 1982, when the company's drop in profitability reduced the bonus pool to zero.

Also conspicuously absent in the academic literature on executive pay is recognition of the reduction in effective CEO pay-for-performance that has been accomplished (presumably unintentionally and unknowingly) by the adoption of so-called "competitive pay" practices starting in the late 1960s. The post-World War II development of *modern* human resource management, with its reliance on private compensation surveys designed to establish the "going rate" for each "benchmark" job, has led to the replacement of the fixed-sharing practices, often based on EVA-like formulas (like GM's above), that were in fact the rule among large public companies before World War II. The competitive pay policies that began to take hold in the 60s and 70s aimed to establish targeted levels of pay (measured in dollars) as opposed to fixed sharing percentages, which had the effect of making (cumulative) pay over longer periods of time largely independent of company performance. What corporate boards and HR departments came to care about most was that their top execs had pay plans that promised to provide, *at the start of each new year*, total pay equal to at least 50th percentile pay for the job.

Both the near-complete replacement of GM-type sharing formulas in public companies by such targeted dollar pay and the

resulting reduction in pay for performance (which I have long tracked using a measure called "managerial wealth leverage") have gone largely unnoticed by both academics and practitioners. Stated as briefly as possible, "wealth leverage" is the ratio of the percentage change in an executive's total wealth—stock and options holdings *plus* the present value of expected future pay—to the percentage change in shareholder wealth.¹

To see the difference between old-fashioned (pre-WW II) fixed sharing and today's competitive pay practices, let's consider the effect of a 5% increase in a company's profitability and value. For a CEO paid with a fixed percentage (or "share") of current and future EVA, such an improvement would increase the CEO's expected future pay by the same 5%. But for the CEO of a company run according to competitive pay practices, that same 5% improvement would be associated with a significantly smaller change in expected future pay. Why? Because of the annual recalibration built into competitive pay policy that effectively *penalizes superior performance* (by raising performance targets and reducing equity grant shares) and *rewards substandard performance* (by reducing targets and *increasing* equity grant shares).

So, for example, if a company's return on capital increases by 200 basis points (and is maintained at the higher level), an executive's bonus under the old GM bonus formula would have increased not only in the current year, but in all future years. But in the case of competitive pay policy, the executive's bonus increases only in the current year. In the years that follow, the company raises its EVA target, reducing the management team's share of EVA, to bring the expected bonus back down to the level of the target bonus. And so, whereas under the GM sharing formula a large component of an executive's lifetime expected wealth—in fact, the entire present value of his or her expected future pay—was highly sensitive to *current* performance, under the competitive pay practices that prevail in most large US companies today that large future expected wealth is largely unaffected by current performance.

¹ See O'Byrne, Stephen F., and S. David Young, "Top Management Incentives and Corporate Performance." *Journal of Applied Corporate Finance* 17(4), Fall 2005.

The most notable effect of this reduction in wealth leverage is that even the most successful US public company CEOs end up making far less than, say, their counterparts in private equity, where payment takes the form mainly of large, upfront fixed grants of stock—while the CEOs of the worst-performing public companies end up making far more.

But if these incentive-reducing effects have been largely ignored by even academic studies of US public company CEO pay, there has been at least one exception I'm aware of. In 2012, four academics—Alex Edmans, Xavier Gabaix, Tomasz Sadzik and Yuliy Sannikov (henceforth “EGSS”)—published in the *Journal of Finance* an article called “Dynamic CEO Compensation” that, in presenting a model of executive wealth they called the “Dynamic Incentive Account,” or DIA for short, exposed these unintended and (again, presumably) unwanted effects of current competitive pay policy practices.² As discussed in more detail below, the DIA, or what might be thought of as a *career pay account*, is an ingenious heuristic designed to help practitioners (and academics) think about the strength of management incentives.

As discussed below, this pioneering paper by EGSS could (and in my view should) have initiated a dramatic change in academic and practitioner thinking about optimal executive pay. But it has had no discernable impact on either more recent academic research or current executive pay practice. Companies continue to embrace competitive pay policy, with its core premise that ensuring a high percentage of “at risk” pay is tantamount to providing a strong incentive to create shareholder value. What's more, by 2020, Edmans himself had abandoned his advocacy of the DIA and taken the position that the academic theory of executive compensation needed to take account of other factors, notably considerations of pay equity and “fairness.” In the pages that follow, I'll try to explain why academic theory and research on CEO pay have had so little effect on executive pay practice, while offering a number of suggestions for compensation theorists intent on making their work more relevant to practitioners.

EGSS'S DYNAMIC INCENTIVE MODEL CHALLENGES THE CONVENTIONAL WISDOM SUPPORTING COMPETITIVE PAY POLICY

EGSS derive the optimal compensation plan using a “principal-agent” model that assumes that corporate directors are aiming to maximize shareholder value net of the cost of CEO compensation and subject to two constraints: a participation or “individual rationality” constraint that says that the contract is attractive enough to retain the CEO, and an “incentive compatibility” constraint that says that any deviation from target effort reduces the CEO's utility. They arrive at a number of conclusions—that the proper measure of incentive strength is wealth leverage (pretty much as I defined it earlier), that “the optimal contract exhibits memory (that is, current pay depends on past output)”³ and that the optimal contract can be achieved through a DIA.

The DIA, again, should be thought of as a kind of career pay account whose value is meant to reflect (at all times) the present value of the CEO's expected future market pay as *adjusted* to reflect the CEO's cumulative performance to date. What's more, and critically important, the value of the account is rebalanced *continuously* to reflect changes in company performance and value with the aim of maintaining a fixed, targeted percentage of equity in the company. In each year of the plan, a percentage of the DIA—as given by a formula derived in the paper—is paid to the CEO in cash. But the calculation of the DIA account effectively makes the annual cash payment and the change in the DIA (prior to the cash payment) a direct function of cumulative company performance since the start of the DIA. And the DIA plan also makes all expected *future* pay—the annualized value of the DIA balance—a direct function of cumulative company performance since the outset of the plan.

To be sure, the DIA concept was not entirely ignored when it was published in 2012. But despite its recognition as a runner-up for that year's *Harvard Business Review*/McKinsey Management Innovation of the Year Award, the concept has had almost no academic follow-up. Nor has it had any detectable influence on executive pay practices in public companies, where total compensation typically depends on performance in only the current year but, at most, the past two to three years. What's more, as I will show later, the setting by boards of the typical CEO's targeted pay—which is the main driver of his or her expected future pay—also turns out to be largely unaffected by past performance.

EXPLAINING THE RISE OF COMPETITIVE PAY PRACTICES

How and why did this happen? For the past 50 years or so, US public company directors, compensation consultants, and proxy advisors have been nearly unanimous in embracing the conventional wisdom that competitive pay policy, with a high percentage of “variable,” or “at risk,” pay, does a reasonably good job of achieving three basic objectives of executive pay: (1) providing strong incentives to increase shareholder value while (2) limiting retention risk and (3) managing shareholder cost. This conventional wisdom holds that a high percentage of pay at risk in and of itself ensures a strong pay-for-performance incentive—and that setting target pay no lower than the 50th percentile limits retention risk and that setting target pay no higher than the 50th percentile and providing performance conditions for equity compensation can at least be defended as managing company cost.

The main flaw in this conventional wisdom, as our wealth leverage example above was meant to show, is the premise that a high percentage of pay at risk necessarily creates a strong incentive for superior performance over longer periods of time. The beauty of the EGSS analysis was its demonstration of how the DIA plan effectively makes an executive's expected cumulative future pay a direct function of the entire series of interim changes in periodic performance. As EGSS puts it, their system actually remembers, and rewards, superior past performance, and penalizes substandard performance.

But since publication of the article, and much to my disappointment, Edmans has retreated from the innovative DIA

² Edmans, Alex, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov. 2012. “Dynamic CEO Compensation.” *Journal of Finance* LXVII(5).

³ Edmans, op. cit., p. 1608.

concept. In 2020, he wrote a book for business and public policy leaders called *Grow The Pie: How Great Companies Deliver Both Purpose and Profit* that, although including a chapter on incentives, fails to even mention the DIA. It instead simply falls back on broad generalities about the desirability of restricted stock with long vesting periods.⁴

What's more, during the period of 2020 and 2021, Edmans collaborated with Tom Gosling and Dirk Jenter in a survey of U.K. directors and investors designed to “understand the objectives and constraints of contract design” and “to guide future compensation theories.” As reported in their recent paper “CEO Compensation: Evidence From the Field,”⁵ one of the key conclusions of the “EGJ” survey is that academic research needs to take account of the role of “fairness considerations” in CEO pay setting.

My own explanation of public company CEO pay has little to do with questions of fairness, but a lot with what I like to call board amnesia and ignorance about once highly effective US corporate pay practices. In my experience, most public company directors—generally with the best of intentions—have failed to recognize that the percentage of pay at risk is not a meaningful measure of incentive strength and, in many cases, fails to provide CEOs with a strong financial incentive for long-run shareholder success. To have an impact on public company practice, academic economists should consider steps like the following:

- (1) getting directors and less sophisticated institutional investors to see that wealth leverage, not percentage of pay at risk, is the right measure of incentive strength (hedge funds and PE investors appear to have recognized this);
- (2) developing a workable method for estimating wealth leverage;
- (3) demonstrating the incompatibility of competitive pay practices with longer-run, cumulative pay-for-performance; and
- (4) testing the extent to which competitive pay policy is really necessary to retain talented executives, as so many directors believe.

BAD MEASURES OF WEALTH AND INCENTIVE STRENGTH HAVE DIVERTED ATTENTION FROM THE DYNAMICS OF TARGET PAY

Michael Jensen and Kevin Murphy's famous 1990 paper on executive pay, “CEO Incentives: It's Not How Much You Pay, But How,” probably the most cited in the academic literature,⁶ reported finding that, during the years 1974–1988, a \$1000 increase in shareholder wealth increased executive wealth by only

two or three dollars. Their measure of incentive strength—the dollar change in CEO wealth per \$1000 change in shareholder wealth—effectively assumed that each additional dollar of wealth makes the same contribution to the CEO's well-being (or “utility”).

EGSS, by contrast, made a compelling case that “percentage-to-percentage” measures give a better picture of CEO incentives because CEO effort tends to have multiplicative (i.e., percentage) “effects on both firm value and CEO utility.”⁷ Some later academic studies use the percentage (or log) change in CEO wealth as a measure of incentive strength, but then fail to estimate expected future pay in ways that reflect the widespread use of competitive pay policy. Perhaps the best example of this approach is Glenn Hubbard and Charles Himmelberg's widely circulated paper that simply assumes that the present value of expected future pay is three times current pay.⁸

The use of such an assumption overstates wealth leverage in two ways. It assumes, contrary to competitive pay policy, that higher current pay automatically means higher expected future pay—and that expected future pay is a very low multiple of current pay. In practice, exceptional current pay tends to be followed by *lower future pay*, thanks to the competitive pay dynamics cited above. The failure of academic studies to understand and reflect the dynamics of expected future pay have contributed to the widely accepted belief that the dramatic growth in equity compensation and executive pay levels after 1980 had greatly strengthened management incentives. And although there was some truth to this perception, there was an important qualification that went largely unnoticed. Yes, the growth in equity compensation starting in the '80s worked to increase both the value of their current equity holdings and the present value of expected future pay. But for many public company CEOs, because the second component was so much larger than the first, and because it had what amounted to *zero wealth leverage*, such larger equity awards failed to increase wealth leverage (or the sensitivity of cumulative pay to cumulative performance) and hence incentives to create shareholder wealth (as compared to those provided by the pre-WWII fixed-sharing bonus schemes).

FAIRNESS DOESN'T HELP US UNDERSTAND EXPLAIN WHY COMPANIES EMBRACE COMPETITIVE PAY POLICY

The DIA, unlike conventional practice, is designed to ensure that poor performance leads to lower expected future pay, and superior performance to higher targeted cumulative pay. To see how this works, let's consider the case of a CEO with annual market pay of \$1 million, an expected tenure of 10 years, and a targeted wealth leverage of 0.6—that is, for every 1% increase (or drop) in his or her shareholders' wealth, the CEO experiences (immediately) a 0.6% increase (or drop) in expected lifetime pay.

If we use an 8% discount rate, the CEO in such a case can be shown (using EGSS's methods) to have a beginning DIA balance

⁴ An example of Edmans' general counsel in the book is “The best way to make a leader accountable to the long-term stock return is to cut her salary, which she receives irrespective of performance, and pay her more in shares.” Edmans, Alex. 2020. *Grow The Pie: How Great Companies Deliver Both Purpose and Profit*. Cambridge, UK: Cambridge University Press, p. 106.

⁵ Edmans, Alex, Tom Gosling, and Dirk Jenter. “CEO Compensation: Evidence From the Field.” <https://ssrn.com/abstract=3877391>, pp. 2,8.

⁶ Jensen, Michael C., and Kevin J. Murphy. 1990. “Performance Pay and Top-Management Incentives.” *Journal of Political Economy* 98(2): 225–84, and “CEO Incentives: It's Not How Much You Pay, But How” *Harvard Business Review*, May/June 1990. JSTOR lists 10,458 citations for the JPE paper while the *Journal of Finance* lists only 170 citations for Dynamic CEO Compensation.

⁷ EGSS, p. 1619.

⁸ Himmelberg, Charles D., and R. Glenn Hubbard, “Incentive Pay and the Market for CEOs – An Analysis of Pay for Performance Sensitivity.” ssrn.com/abstract=236089, March 2000.

of \$6.7 million, consisting of \$4.0 million in company stock and the rest in cash. In the event the stock declines by 25%, the DIA balance declines—thanks to continuous rebalancing designed to maintain the stock-to-cash ratio at 1.5:1—to \$6.0 million. And this change in turn has the effect of reducing the CEO's future expected *annual* pay—that is, target pay—from \$1.0 million to \$0.84 million.⁹

As this example is meant to show, the DIA reduces target (or expected future) pay whenever the stock price declines. But when EGJ asked the UK directors in their survey if they had ever significantly reduced target pay for an incumbent CEO, only 23% said yes,¹⁰ leading EGJ to conclude that "...investors and especially directors believe that CEOs need to be paid at competitive levels, even absent any recruitment or retention concerns, because the failure to do so would be viewed as unfair and undermine the CEO's intrinsic motivation."¹¹ EGJ appear to be suggesting that, although cutting pay is the right thing to do in this case, few directors actually do so out of concern it would be "unfair." Another possible interpretation, however, is that the UK directors believe that cutting pay is the wrong thing to do under such circumstances, and that deviating from the company's competitive pay policy would be not only unfair, but in the long run counterproductive.¹²

And as we noted earlier, this finding of the EGJ survey reflects the conventional wisdom of directors and investors that competitive pay policy with a high percentage of pay at risk achieves the three main objectives of executive pay. The gap between theory and practice thus should not be seen as a failure of the theory to recognize the role of considerations of fairness in compensation decisions, but rather of academics (and compensation consultants) to convince directors and investors that percentage of pay at risk is neither a reliable indicator nor a guarantee of incentive strength.¹³

But to implement EGSS's concept of executive pay as a career pay account, one would require estimates of the following: (1) the values of current stock and option holdings; and (2) the present value of expected future pay—and also estimates of the sensitivities of both to changes in relative shareholder wealth. Estimating the value and sensitivity of stock and option holdings should be straightforward, if sometimes complicated in the actual calculation. The value and sensitivity of expected future pay is more challenging. In a 2010 study, David Young and I used the sensitivity of current year [t] pay to earlier years' [t-x] performance to estimate the sensitivity of a future year's [t+x] pay to current year performance.¹⁴ To cite just one example from study, our analysis of Exxon CEO Rex Tillerson's wealth at the end of 2006 showed that although Tillerson had fully 74% of his current pay at risk,

his effective wealth leverage was only 0.35—a sign that at-risk pay is at best an unreliable proxy for wealth leverage.¹⁵

The explanation of this difference between at-risk pay and wealth leverage turns on the recognition that Tillerson's 2006 stock and option holdings in Exxon of \$42 million then represented only 20% of his expected total wealth, which was a relatively small fraction of the present value of his then expected future pay (which we estimated at \$166 million). And although the wealth leverage of his stock and options was 1.32, our estimate of the leverage of his expected future pay was a mere 0.11, bringing his total wealth leverage down to 0.35. Why was it just 0.11? Because of the loose connection between Exxon's stock awards to its CEO's and the company's past performance during Tillerson's tenure to date (as reflected in our regressions of Tillerson's pay on lagged company performance).

In other words, our analysis showed Tillerson as having a weak wealth incentive even though his equity holdings would qualify as "substantial" in EGJ's view. It also shows that strong incentives from continuous adjustments in expected future pay are required for strong wealth leverage.

The unfortunate reality, however, is that David's and my wealth leverage analysis—our version of EGSS's career pay account—has never gained a wide audience. And both David and I have been disappointed by EGSS's inability to attract followers of their own work in presenting corporate pay practices as a career pay account.

FAIRNESS IS NOT NEEDED TO EXPLAIN WHY COMPANIES PROVIDE INCENTIVE PAY

EGJ offer two main reasons for their argument that incentive pay is—and presumably should continue to be—motivated by fairness considerations. First, the majority of directors and investors "view financial incentives as less important than intrinsic motivation and the CEO's personal reputation."¹⁶ EGJ's argument is that incentive pay is not necessary to motivate the CEO's effort, as compensation theory assumes, so its justification must lie elsewhere. But the complexity of CEOs' motives and concerns should not be allowed to obscure the role of financial incentives. Even the most talented and reputable CEOs are liable to become excessively focused on goals that conflict with the financial well-being of their shareholders, such as expanding the size of the company or embracing technical innovation for innovation's sake. The blind spot in EGJ's argument may well be that directors' view of financial incentives as relatively unimportant could simply reflect the directors' inability (or unwillingness) to recognize that the company's performance incentives, as measured by percentage of pay at risk, have been ineffective—and that competitive pay practices are at the heart of the problem they are supposed to be addressing.

The second main reason, in EGJ's view, that incentive pay practices tend to be motivated by fairness consideration is that such pay constitutes a trivial fraction of changes in equity value. EGJ's argument effectively boils down to the proposition that since

⁹ A 10-year annuity of \$0.84 million has a present value of \$6.0 million.

¹⁰ EGJ, p. 26.

¹¹ EGJ, pp. 41–42.

¹² This conception of "fairness" simply interprets it as intellectual consistency with no substantive policy content.

¹³ Competitive pay policy, that is, target dollar pay, creates a systematic performance penalty in that superior performance is penalized with fewer equity grant shares, while poor performance is rewarded with more equity grant shares. A high percentage of pay at risk does nothing to mitigate the performance penalty. See O'Byrne, Stephen F., and E. Mark Gressle. 2013. "How 'Competitive Pay' Undermines Pay for Performance (and What Companies Can Do to Avoid That)." *Journal of Applied Corporate Finance* 25(2), Spring 2013.

¹⁴ O'Byrne, Stephen F., and S. David Young. 2010. "What Investors Need to Know About Executive Pay." *Journal of Investing* Spring 2010.

¹⁵ O'Byrne and Gressle provide a more systematic demonstration, using 1000 five year Monte Carlo simulations, that percent of pay at risk is a poor proxy for wealth leverage.

¹⁶ EGJ, p. 42. "Intrinsic motivation" and "personal reputation" are responses to the survey question, "What motivates your CEO to perform strongly?"

incentive pay is not really necessary to provide financial incentives for the CEO, its justification must lie elsewhere. This argument is based on EGJ's misguided embrace of the old Jensen & Murphy view that it is the dollar change in CEOs' wealth, and not wealth leverage itself, that is the proper measure of incentive strength. This assumption seems clear in EGJ's statements that "almost all CEOs have substantial equity holdings" and that changes in pay are "so small" compared to those from equity holdings that incentives from pay can be "ignored."¹⁷

But contrary to EGJ's suggestions, our concept of wealth leverage, as illustrated in the analysis of Tillerson's holdings above, shows that the incentive effects of equity holdings can depend heavily on the present value of *expected future* pay, which may be large relative to holdings. And it is likely to be the CEOs' cumulative expected future pay, and how it relates to shareholder returns, that determines the real incentive for performance.

THE SEC'S NEW PAY DISCLOSURE RULE (RAISING THE POSSIBILITY OF A MORE PRACTICAL WAY TO ESTIMATE WEALTH LEVERAGE)

But if neither DIA nor wealth leverage analysis has had a discernible effect on public company thinking about incentive strength, general interest in incentive analysis based on historical data should increase markedly with the SEC's new executive pay disclosure requirement. Set to take effect this year (2023), public companies will be required to disclose their "Compensation Actually Paid"—which includes the year end value of current equity compensation as well as changes in the value of prior unvested equity grants for the past three years. And starting in 2025, companies will be required to provide the same "mark to market" estimates for the past 5 years.

For analysts of CEO pay like me, the new requirements may prove especially useful since, by eliminating the need to come up with such estimates on our own,¹⁸ the disclosures could make it fairly straightforward to take a somewhat simpler approach to estimating wealth leverage that I call *pay leverage*. The main difference between the two approaches to quantifying incentive strength is that whereas estimates of wealth leverage rely heavily on *projections* of future pay and shareholder returns (based on past data), calculations of pay leverage can be done entirely using just the actual historical pay data.

More specifically, in my own pay leverage analysis, I have used a CEO's historical pay data to estimate his or her cumulative "mark to market" pay for the most recent ten years. As a measure of incentive strength, I plot the CEO's *relative* pay (on a log scale) on the vertical axis against the company's *relative* performance—that is, how the company's shareholder returns compared to its peers—on the horizontal axis.¹⁹ An executive's relative pay for

a given year is his or her cumulative "mark to market" pay divided by the cumulative future value of his or her market pay. A company's relative performance is cumulative shareholder wealth divided by cumulative peer company shareholder wealth.²⁰ Mark-to-market pay values all equity compensation based on the stock price at the end of the cumulative measurement period, and estimates of percentage vesting.²¹ The future value of market pay adjusts market pay, which is a present value number, for the expected accretion of equity compensation over time.

What emerges from this analysis, then, is an estimate of my measure called *pay leverage*. Calculated as the ratio of the (log) change in cumulative relative pay to the (log) change in relative shareholder wealth, pay leverage is also the slope of the regression trendline. For small changes, this is equal to the ratio of the percentage change in cumulative relative pay to the percentage change in relative shareholder wealth. Other useful outputs from this regression analysis are the correlation, which we interpret as "pay alignment," and the intercept, which represents the "pay premium" for average performance—that is, the amount by which an executive's total pay exceeds market average pay when the company has average performance for its industry.²²

This kind of pay leverage analysis can be used to answer the question, "is there a simple pay plan that provides a perfect (1.0) correlation of an executive's *relative* pay and the company's relative shareholder performance with a slope of 1 and an intercept of zero?" The short answer is that there is a simple "perfect" pay plan that uses only annual grants of performance shares,²³ but it differs from conventional practice in three important ways: (1) target pay is not market pay, but market pay adjusted for the company's relative performance from the start of the plan; (2) the vesting multiple for each grant identifies and removes the impact of industry performance (by using $1/(1 + \text{the industry return from the date of grant})$ as the multiple); and (3) all cash paid out prior to retirement is treated as a draw against the value of the performance shares. The vesting multiple widely used in practice is $[1 + \text{relative TSR}]$, but this leverages the industry component of the stock return instead of taking it out of pay.²⁴

To sum up their differences, then, whereas wealth leverage looks forward in time and tries to capture the relationship between future changes in pay and performance, pay leverage looks backward in time to capture the relationship between past changes in pay and performance. Because their calculations are so different,

frequently reward their CEOs for good market and industry performance. Market pay can be thought of the normalized pay of peer CEOs. If actual peer pay were used to compute relative pay, matching relative pay to relative performance would require paying for market and industry performance just because peer companies do so.

²⁰ Cumulative peer company shareholder wealth is peer company shareholder wealth weighted by beginning shareholder wealth. Peer company shareholder wealth adjusted for the subject company's "industry beta" provides a more accurate estimate of management's contribution to shareholder wealth.

²¹ I have estimated vesting percentages using relative shareholder return as a proxy for the company's actual vesting measures. The pay reported for the new SEC disclosure requires companies to estimate vesting percentages.

²² Ten-year pay leverage regressions are shown in O'Byrne, Stephen F. 2018. "Say on Pay – Is it Needed? Does It Work?" *Journal of Applied Corporate Finance* 30(1), Winter 2018.

²³ See O'Byrne and Gressle.

²⁴ $\text{Stock value}/\text{initial stock price} = 1 + \text{TSR} = (1 + \text{relative TSR}) \times (1 + \text{industry TSR}) = (1 + r\text{TSR})(1 + i\text{TSR})$, so $\text{stock value} \times [1/(1 + i\text{TSR})] = \text{initial stock price} \times (1 + r\text{TSR})$. This is the "perfect" pay plan. Conventional practice is $\text{stock value} \times (1 + r\text{TSR}) = \text{initial stock price} \times (1 + i\text{TSR}) \times (1 + r\text{TSR}) \times (1 + r\text{TSR})$, so conventional practice leverages the industry return, $1 + i\text{TSR}$, by a factor of $(1 + r\text{TSR})^2$.

¹⁷ EGJ, pp. 23–24.

¹⁸ Although the SEC uses the term "Compensation Actually Paid," the pay to be disclosed is more accurately characterized as "realizable" or "mark-to-market" pay. And though the SEC requires disclosure of annual, not cumulative, mark to market pay, the annual figures can be easily summed to get cumulative pay. The SEC does not require companies to estimate market pay for the CEO but market pay can be estimated from peer company pay and size data.

¹⁹ Our normative concept is that relative pay should match relative performance. We use the market pay, not the actual pay, of peer CEOs, to compute relative pay because peer companies

we would not expect pay leverage to be a perfect proxy for wealth leverage. They are the same in the special case of the “perfect” pay plan,²⁵ but when applied to the typical US public company that embraces competitive pay policy, estimates of pay leverage will overstate wealth leverage. This is almost certain to be true because, as we saw earlier in the case of Exxon’s Rex Tillerson, estimates of pay leverage ignore the present value of expected future pay that, in our public companies, tends to have wealth leverage of zero.²⁶

The great advantage of pay leverage is that it is a replicable calculation from historical pay and performance data, not an easily challenged projection of future pay and performance data.²⁷ And the upward bias in pay leverage can be reduced by using long periods of historical data. A ten-year analysis of S&P 1500 companies shows, for example, that nearly two thirds (63%) of CEOs have low alignment of relative pay and performance, and that some 87% of such companies have pay-for-performance deficiencies that include outsized retention risk and high cost to shareholders.²⁸

THE CRUX OF THE CEO PAY PROBLEM: DIRECTORS ARE NO LONGER OWNERS

Today’s directors of public companies are largely paid labor providers, not as in the past the stewards of substantial capital investments in the companies they are supposed oversee. In a paper on “The Evolution of Executive Pay Policy at General Motors 1918–2008,” David Young and I traced the evolution of the modern director from capital steward to hired hand.²⁹ Whereas the median GM director in 1947 held \$1.65 million in stock and received an annual director’s fee of \$900, the median GM director in 1977 held \$34,000 in stock and received an annual director’s fee of \$47,000. For the median GM director in 1947, the annual expected return on the director’s stock, assuming a 10% expected return, was 183 times the director’s annual fee for service, so the director’s fee would cover the loss of the expected return on the director’s stock for only 2 days.³⁰ For the median

²⁵ Executive wealth is the sum of cumulative pay and the present value of expected future pay. The “perfect” pay plan makes cumulative pay equal to $(1 + rTSR) \times$ cumulative market pay and the present value of expected future pay equal to $(1 + rTSR) \times$ the present value of future market pay, so executive wealth is equal to $(1 + rTSR) \times$ [cumulative market pay + the present value of future market pay], or $(1 + rTSR) \times$ expected executive wealth. This means that the change in executive wealth is equal to the change in relative shareholder wealth, so wealth leverage, like pay leverage, is equal to one. Slightly more complicated versions of the “perfect” pay plan allow for pay leverage greater than or less than 1.0. See O’Byrne, Stephen F. 2012. “Achieving Pay for Performance.” *Conference Board Director Notes* 4(24), December 2012.

²⁶ Wealth leverage is a weighted average of pay leverage and the wealth leverage of expected future pay. Since the wealth leverage of expected future pay is zero when a company pursues competitive pay policy, wealth leverage < pay leverage for positive pay leverage and the absolute value of wealth leverage is less than the absolute value of pay leverage for negative pay leverage.

²⁷ This difference between pay leverage and wealth leverage as measures of incentive strength is similar to the difference between EVA and discounted cash flow projections as measures of shareholder value. EVA is an incomplete measure but much more widely used in cross company analyses because it’s a replicable calculation from historical data.

²⁸ See O’Byrne, Stephen F. 2018. “Say on Pay: Is It Needed? Does It Work?” *Journal of Applied Corporate Finance* 3(1), Winter 2018. Using 2007–2016 data for S&P 1500 CEOs, 63% have low alignment ($r^2 < 50\%$), 28% have high retention risk (i.e., a -33% pay “premium” at average performance), 27% have high pay risk (i.e., pay volatility $> 1.5 \times$ stock volatility) and 24% have high cost (a 50%+ pay premium at average performance).

²⁹ O’Byrne, Stephen F., and S. David Young. 2017. “The Evolution of Executive Pay Policy at General Motors, 1918–2008.” *Journal of Applied Corporate Finance* 29(1), Winter 2017 and O’Byrne, Stephen F. “Measuring and Improving Pay for Performance.” Chapter 39, *Handbook of Board Governance 2nd Edition*, Richard W. LeBlanc, editor, John Wiley & Sons, 2020.

³⁰ $10\% \times \$1.65 \text{ million} / \$900 = 183$; $10\% \times \$34,000 / \$47,000 = .07$.

GM director in 1977, the annual expected return on the director’s stock, again assuming a 10% expected return, was 0.07 times the director’s annual fee for service, implying that the director’s fee would cover the loss of the expected return on the director’s stock for almost 14 years.

Like today’s directors of US public companies, the 159 “investors” who responded to EGJ’s survey of UK companies are also not close matches for the “principal” of standard economic models. Over half (52%) of them are governance specialists who are rarely compensated on investment performance, and, as EGJ acknowledge, such respondents “may be more sensitive to the societal consequences of pay” than the consequences for the company.³¹ Only 26% of the “investors” responding to EGJ’s survey are fund managers, and even these respondents would be holding substantially diversified portfolios that would make them a weak match for the “principals” of standard economic models.

PAYING FOR INDUSTRY PERFORMANCE MAY REFLECT BARGAINING POWER, NOT FAIRNESS

EGJ see fairness considerations in directors’ endorsement of the statement that “the CEO should benefit from an industry upswing since investors and stakeholders do.” Investment management pay can be used to provide some insight on the dynamics of payment for industry performance in a somewhat different kind of principal-agent relationships. Investment managers are typically hired based on an evaluation of their relative performance, but at the same time paid with formulas like “2 and 20” that give managers a share of the total investment gain, including the portion due to industry performance. It’s possible that investors want hedge fund managers to share in gains due to industry performance because it’s fair—but it’s also possible that investment managers have the power to resist a carve-out of the gains attributable to industry performance alone.

There is some evidence to support the latter argument. When Don Raymond was the chief investment strategist of the Canada Pension Plan Investment Board (CPPIB), he developed an improved fee structure for external managers that paid managers a performance fee based on their value added relative to a benchmark and an asset fee that was a draw against the performance fee.³² “The fee structure...was not an attempt to reduce the amount of fees paid but to design a fee structure that would better align interests.” Raymond himself told me in 2013 that CPPIB had succeeded in getting about half of its external managers to accept the new fee structure, but the other half had refused to accept it. CPPIB retained many of the managers who refused to accept the new fee structure because their superior performance outweighed the expected benefits of the better fee structure. Raymond’s experience suggests that payment for industry performance is attributable more to power dynamics than considerations of fairness.

³¹ EGJ, p. 14.

³² Raymond, Donald M. 2008. “Paying (Only) for Skill (Alpha) – A Practical Approach”, *CFA Conference Proceedings*, June 2008.

A CHALLENGE FOR COMPENSATION THEORY: EXPLAINING CORPORATE AND HEDGE FUND PAY

One major shortcoming of principal-agent models applied to executive pay has been the failure to create a model of the investor principal that applies to both the corporate executive agent and the investment manager agent while nevertheless explaining the big differences in structure and pay levels for seemingly similar principal-agent relationships. Hedge funds have simple, formula-based compensation that leads to very high pay levels for successful managers. Institutional Investor's "Rich List" of the highest earning hedge fund managers includes seven people making over \$1 billion in 2022 and two making more than \$3 billion.³³ By contrast, the new CEO of Amazon, Andrew Jassy, received a stock grant worth \$212 million on his promotion to CEO in 2021.³⁴ And since this grant vests over a ten-year period and is intended to be his equity compensation for the next 10 years, his annualized compensation is around \$21 million when including his salary of \$317,500, though it jumps to over \$33 million when we include the expected annual return on his retained shares from prior stock grants. But even the larger of these two numbers is a mere 3% of the \$1 billion received by the top seven hedge fund managers.³⁵ Explaining the difference between corporate and hedge fund pay might lead theorists to moderate the standard model's assumption that the CEO's pay has no effect on the pay of other members of the organization.

My experience suggests that it's rare for large organizations to adopt an incentive plan for a CEO without including other employees in the plan. And this in turn suggests that public company CEO pay may effectively be constrained by the costs of extending supporting incentives to a much broader employee population. Since hedge funds are much smaller organizations than operating companies, the cost of extending the leader's incentive to a broader employee population is much less for hedge funds and would operate as a much weaker constraint on leader pay. This possibility alone, quite apart from the difference in the profitability per individual executive, could account for the extraordinary premium of hedge leader pay over CEO pay.

ANOTHER CHALLENGE FOR COMPENSATION THEORY: CONTROLLABLE VS UNCONTROLLABLE RISK

The simplifying assumption that the CEO's pay has no impact on any other employee's pay may have contributed to the perception that the central problem in CEO pay is the trade-off between effort incentives and risk. Since the CEO's pay, as EGJ recognize,

"is a small percentage of firm value for most firms," there is little likelihood that cost will constrain CEO incentive plan design provided all other employees are unaffected. If cost is irrelevant, risk appears to be the only constraint on incentive plan design. The standard model implies that incentives will be weaker when risk is high because the agent's risk aversion increases the compensation cost borne by the principal.

But as Canice Prendergast has shown using data for executives, sharecroppers and franchisees, strong incentives tend to be more common in high volatility environments.³⁶ Prendergast argues that reducing pay risk for the agent is practical only when the principal can monitor the agent's efforts at reasonable cost. In cases where the cost of input monitoring is high and managerial discretion is great (and effective use of it is critical to the success of the business)—and where there is a lot of uncertainty about the business itself (and how the agent should respond to it), principals will prefer to give their agents incentive pay based on output. Prendergast stresses the importance of distinguishing between the variability that is attributable to the executive (or agent's) responsibility for decision making and the variability in the uncertainty of the environment.

If we think of the CEO's performance as the sum of industry and personal performance, we can see that the standard model is assuming that the CEO is equally risk averse to all sources of variability, both the variability of industry performance that is clearly beyond the agent's control and the variability of the CEO's personal performance.³⁷ But the CEO is much more likely to demand a pay premium for exposure to industry risk than to the variability of his own performance. Compensation theory needs to distinguish, as Prendergast does, between controllable and uncontrollable risk.

CONCLUSION: A NEW AGENDA FOR COMPENSATION THEORISTS

First, theorists need to recognize that CEO pay is typically part of a management team—and sometimes a company-wide employee—pay plan, and so the shareholder benefit of CEO "effort" needs to be offset by the cost to shareholders of the entire management team pay package. The cost of management team pay, together with possible differences in the risk preferences of the management team members and the need to distinguish between controllable and uncontrollable risk, may lead theorists to the conclusion that cost is a far more significant constraint in pay plan design than risk.

Considerations of fairness, in my experience, are a key reason why CEO and management team incentive plans are similar, but that doesn't mean that fairness needs to be a parameter of compensation theory. It may be enough to make an assumption about the sensitivity of management team pay to CEO pay.

Second, theorists should try to distinguish between controllable and uncontrollable risks. If the CEO's exposure to uncontrollable

³³ Ken Griffin of Citadel, \$4.1 billion, and Israel Englander of Millenium Management, \$3.2 billion. *Institutional Investor*, "The Rich List: The 22nd Annual Ranking of the Highest-Earning Hedge Fund Managers", March 7, 2023.

³⁴ Amazon proxy statement April 14, 2022. Jassy's stock grant was 61,000 shares valued at the grant date stock price of \$3474 per share.

³⁵ Jassy's salary is from Amazon's 2023 proxy statement. Hedge fund pay reported in the press typically includes income from prior pay re-invested in the fund. At the end of 2021, Jassy held 33,729 shares in addition to that year's 61,000 share grant. The expected annual expected return on those shares is \$11.7 million, based on the 2022 grant date stock price and a 10% expected return assumption.

³⁶ Prendergast, Canice. 2002. "The Tenuous Trade-off between Risk and Incentives." *Journal of Political Economy* 110(5).

³⁷ See Gibbs, Michael. "Designing Incentive Plans: New Insights from Academic Research." *WorldatWork Journal*, December 2012.

risk puts a much bigger cost burden on the principal, companies may find it cost-effective to introduce relative performance measures designed to eliminate the effects of uncontrollable risk.

Third, theorists should try to develop an agency model that is sufficiently general to explain the differences between operating company pay and hedge fund pay. The model needs to explain the difference in structure and the difference in pay level. It may be that the difference in pay level is due to the size and cost of the management team. Higher operating company CEO pay has the potential to raise the pay of thousands of employees while higher hedge fund CEO pay potentially raises the pay of just a few dozen employees.

Fourth, theorists need to develop models that enforce a “participation constraint” every year—for example, recalibrating performance targets and grant shares to provide competitive pay every year, as directors appear to believe they must. We could compare these models with models like EGSS’s that enforce the

participation constraint only at the start, and then let empiricists test the validity of the major premise of the conventional wisdom.

Fifth, theorists need to help empiricists (and practitioners!) by making a serious effort to express conventional pay practices in terms of their model. Specifically, and most urgently, they need to develop a methodology for estimating the sensitivity of expected future pay to current performance. My hope is that they begin to recognize that long horizon pay leverage is a much needed addition to academic theory on executive pay, one that is much more pressing than “fairness considerations.”

How to cite this article: O’Byrne, Stephen F. 2023. “Why has academic research had so little impact on US executive pay practice?” *Journal of Applied Corporate Finance* 1–8. <https://doi.org/10.1111/jacf.12564>