CEO pay has a big impact on future stock returns

AUTHOR: Stephen F. O'Byrne, Shareholder Value Advisors

SHAREHOLDER VALUE ADVISORS

Executive pay has three basic objectives: provide strong incentives to increase shareholder value, retain key talent and limit shareholder cost. The conventional wisdom, widely embraced by boards, compensation consultants and proxy advisors, is that a company can achieve the three basic objectives if has a high percent of pay at risk and maintains target pay at the median of its peer companies. The high percent of pay at risk ensures a strong incentive, the target pay percentile retains key talent because the company doesn't allow target pay to fall below the median and the target pay percentile limits shareholder cost because the company doesn't allow target pay to rise above the median. In this paper, I will show that percent of pay at risk is a poor measure of incentive strength and that the target pay percentile is a poor measure of retention risk and shareholder cost. I will also show that there are better measures of incentive strength and of retention risk/compensation cost, and that these better measures have statistically and economically significant effects on future stock returns. Boards, compensation consultants, proxy advisors and institutional investors can use our future return models to assess the cost-efficiency of current and proposed executive pay plans. Active investors can use our pay dimensions to enhance their stock selection models. Hedge funds can use our models to create profitable long/short portfolios based on pay dimensions.

The conventional wisdom about executive pay design

ConocoPhillip's 2025 proxy statement illustrates the conventional wisdom. ConocoPhillips says that "a company must offer competitive compensation to attract and retain experienced, talented and motivated employees" and "employees in leadership roles are motivated to perform at their highest levels when performance-based pay is a significant portion of their compensation". ConocoPhillips notes that "target total compensation for each NEO is structured to target market competitive pay levels at approximately the 50th percentile" and highlights that "over 90 percent of the CEO's 2024 target pay and over 80 percent of the other NEOs' 2024 target pay was performance based." The proxy advisor Institutional Shareholder Services (ISS), like Conoco Phillips, prefers a high percent of pay at risk with target pay set at the peer company median regardless of past performance.

Public company pay practices over the past twenty years show an increasing embrace of the conventional wisdom. Percent of pay at risk for the median S&P 1500 CEO has increased from 74% in 2004 to 86% in 2024. (1) A 2024 paper found that the average deviation of CEO grant date pay from the median pay of similarly sized companies in the same industry has declined by 45 percent from 2007 to 2021. (2)

The negative consequences of target dollar pay

An important, but little recognized, weakness of the conventional wisdom is using target dollar pay without any adjustment for prior performance. This creates an inherent "performance penalty" in the determination of equity grant shares. An increase in the stock price is penalized with a reduction in grant shares, while a decline in the stock price is rewarded with an increase in grant shares. We can see this performance penalty in operation at ConocoPhillips. In 2021, when the stock price was $46.65, CEO Ryan Lance was granted 200,722 performance shares and 108,082 restricted shares. In 2024, when the stock price was $110.39, Lance was granted 90,781 performance shares and 48,883 restricted shares. In other words, while the stock price increased by 137%, Lance's grant shares declined by 55%. The value of his equity grant in 2024 was only 7% higher than it was in 2021, and his proxy statement total compensation was 3% lower in 2024 than it was in 2021 ($23.9 million vs $23.1 million).

The performance penalty also works to increase grant shares for poor performance. In 2015, when the stock price was $29.06, Marathon Oil CEO Lee Tillman received 256,591 stock options, 135,487 performance shares and 81,292 restricted shares. In 2018, when the stock price was $14.52, Tillman received 298,914 stock options, 284,091 performance shares and 170,455 restricted shares. When we use the Black-Scholes model to convert Tillman's stock options to an equivalent number of common shares, we find that Tillman's equivalent common shares increased from 277,174 in 2015 to 574,564 in 2018, an increase of 107%. With the increase in grant shares, Tillman's equity grant value in 2018 was 4% greater than it was in 2015 despite the 50% decline in the stock price. His proxy statement total compensation was 8% higher in 2018 than in 2015 ($12.2 million vs $11.3 million), despite a 50% decline in stock price.

Target dollar pay (unadjusted for prior performance) has three negative consequences. First, it weakens financial incentives for superior performance. Second, it undermines the alignment of cumulative pay and cumulative performance. Equity grant shares depend on the price path; cumulative performance with poor early performance pays much more than the same cumulative performance with good early performance. Third, it makes percent of pay at risk a poor proxy for incentive strength. With the same percent of pay at risk, an executive with bad early performance accumulates many more shares than an executive with good early performance. That makes the bad early performer's pay much more sensitive to company performance; in other words, it gives the bad early performer a much stronger incentive even though the bad early performer has the same percent of pay at risk as the good early performer. Percent of pay at risk is a poor measure of incentive strength because executives can have very different incentives with the same percent of pay at risk. (3)

A better set of pay dimension measures

The pay measures used by the conventional wisdom – percent of pay at risk as a proxy for incentive strength and target pay percent from market – don't provide good measures of a company's success in achieving the three basic objectives of executive pay. Percent of pay at risk doesn't provide a good measure of incentive strength because two executives with the same percent of pay at risk can have very different pay sensitivities to performance. Target pay percent from market doesn't provide a good measure of retention risk or shareholder cost because it doesn't adjust for performance. Target pay at market provides a very strong retention incentive for a 10th percentile performer but a very weak retention incentive for a 90th percentile performer. Target pay at market is high shareholder cost for a 10th percentile performer but low shareholder cost for a 90th percentile performer.

To improve on the conventional wisdom measures, we need to measure incentive strength, that is, pay sensitivity to relative performance, and we need to measure performance adjusted cost, that is, the pay premium at industry average performance. To do these two things in the most meaningful way we need to measure pay on a "mark to market" basis rather than rely on the grant date pay reported in the proxy's Summary Compensation Table. We need mark to market pay because changes in mark to market pay capture the incentive provided by unvested equity compensation. Happily, U.S. companies are now required to make a "Pay versus Performance" (PvP) disclosure that includes a five year history of mark to market pay, called "Compensation Actually Paid" or "CAP", for the CEO and for the average of the other Top 5 executives. (4)

The new CAP data shifts a huge data analysis burden from investors to companies but, unfortunately, it doesn't currently provide the long history needed to assess the impact of pay dimensions on future stock returns. We can't use the new PvP disclosures to estimate the impact of pay dimensions on future stock performance because we don't yet have any five year PvP disclosures followed by a multi-year return period. Fortunately, we have an alternative although it requires extensive programming and data analysis. We can use historical grant data to estimate mark to market pay for a long history period, then estimate historical pay dimensions from that estimated mark to market pay and finally, relate those historical pay dimensions to future period stock returns. With Standard & Poor's Execucomp database, we can estimate mark to market pay for 1992-2024, use that data to calculate pay dimensions for rolling five year periods and then assess the impact of those pay dimensions on future returns. Companies and investors should feel comfortable in applying our model of pay dimensions and future returns to the new PvP disclosures because we show, in the Appendix, that our mark to market pay estimates for 2020-2024 are similar to the mark to market pay values reported in the new PvP disclosures.

Estimated mark to market pay

We have estimated mark to market pay for executives reported in Execucomp for the years 1992-2024. Figure 1 shows 27 CEOs with five year mark to market pay exceeding $1 billion, inflation adjusted to 2024. This group includes many very noted executives, including Larry Bossidy of Honeywell, John Chambers of Cisco, Michael Eisner of Disney, Lou Gerstner of IBM, Steve Jobs of Apple, William McGuire of United Health Group, and Elon Musk of Tesla. 15 of these 27 executives earned their $1 billion in the five year periods ending in 1998-2002. Only 12 of them have earned $1 billion in a five year period ending after 2002.

Figure 1

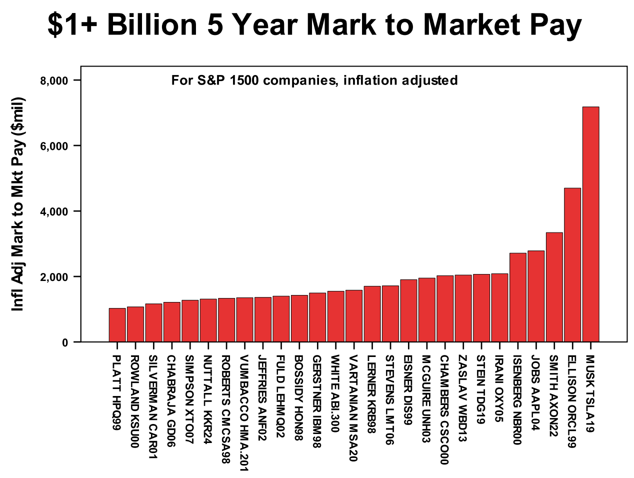

Figure 2

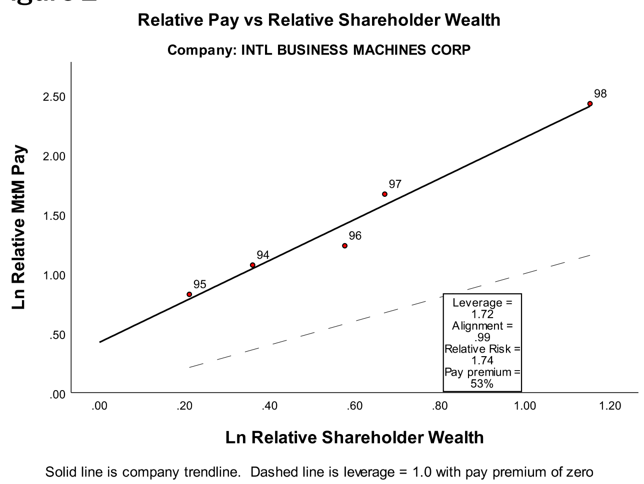

The vast majority of these $1+ billion gains from stock gains. For example, 91% of Lou Gerstner's $1.5 billion in inflation adjusted gains came from stock options.

How we measure pay dimensions

To measure the pay dimensions associated with these gains, we plot the natural log of relative mark to market pay against the natural log of (1 + relative TSR) and calculate the trendline. Figure 2 shows the scatterplot for 1994-1998 for IBM CEO Lou Gerstner. The vertical axis is the natural log of relative mark to market pay and the horizontal axis is the natural log of relative shareholder wealth. We use logarithms because they capture the assumption that a 1% change in relative shareholder wealth results in a constant percentage change in relative mark to market pay. Mark to market pay is calculated on a cumulative basis over the five year measurement period. Relative mark to market pay is mark to market pay divided by cumulative market pay. Market pay is average pay for the same position in companies of equal revenue in the same industry. Our methodology for estimating market pay is explained in the Appendix. We estimate market pay for the first year of the five year period and adjust it for subsequent executive pay inflation but not for subsequent revenue growth. We don't adjust for subsequent revenue growth because doing so obscures pay sensitivity to performance since shareholder return is often correlated with revenue growth. All figures are adjusted for inflation to 2024.

Inflation adjusted 1994 market pay for Gerstner is $20.7 million. It's based on pay for CEOs in Technology Hardware (GICS 452020) adjusted to IBM's inflation adjusted sales of $135 billion. Relative mark to market pay is intended to be a measure of pay relative to opportunity cost. To better match opportunity cost with utilization of resources (i.e., grant date pay), we match market pay to reflect the timing of Gerstner's grant date pay. Since Gerstner had a large equity grant in 1997, we allocate 56% of Gerstner's market pay to 1997 (without changing the five year present value of market pay).

How we measure expected shareholder wealth

Relative shareholder wealth is actual shareholder wealth divided by expected shareholder wealth. We calculate relative shareholder wealth for IBM for each of the five years 1994-1998 but we'll illustrate the calculation just using the full five year period. IBM had a huge five year return, increasing shareholder wealth, with dividends re-invested, by 665%. To compute IBM's relative shareholder wealth, we need to divide IBM's actual shareholder wealth by its expected shareholder wealth. Expected shareholder wealth is beginning shareholder wealth increased by IBM's expected return, taking account of risk and taking account of actual market and industry returns. IBM's cost of equity, 10.8% based on a four factor Fama-French model, provides the adjustment for risk, a five year expected return of +67% (= 1.108^5 – 1). This brings expected fifth year shareholder wealth up to 167% of beginning shareholder wealth but we still need to adjust for market and industry performance. We use market and industry betas for IBM, together with market and industry excess returns, to adjust for market and industry performance. The five year market excess return is 42.5% and IBM's beta to the market excess return is

Figure 3

Pay Leverage 1998-2024

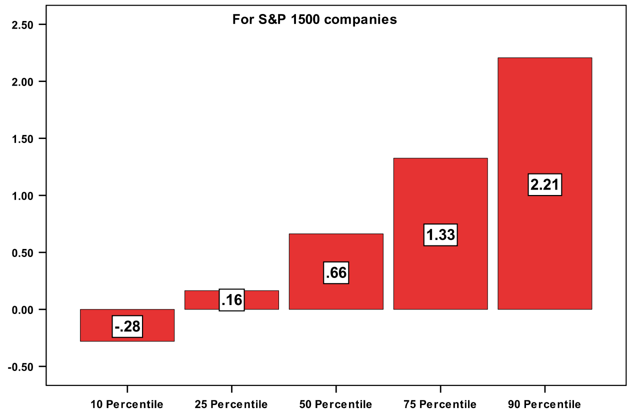

0.47, so the market excess return adds 20.0% (= 42.5% x 0.47) to IBM's expected shareholder wealth. The five year industry excess return is 47.3% and IBM's beta to the industry excess return is 0.55, so the industry excess return adds 26.0% to IBM's expected shareholder wealth. The total contribution of market and industry performance is 46.0% and when we add that (geometrically) to the 67% risk adjustment, we get a total expected return of 144% (= (1 + 67%) x (1 + 46%) – 1). This means that expected shareholder wealth is 244% of beginning shareholder wealth and IBM's five year excess return is (765%/244%) -1 = 217%.

Once we have calculated relative pay and relative performance for each year, we can calculate the trendline relating log relative pay to log relative TSR. As shown in Figure 2, the trendline slope, what we call pay leverage, is 1.72 for Gerstner. This means that a 1% increase in relative shareholder wealth is associated with a 1.72% increase in relative pay. This

Figure 5

CEO Percent Pay Premium 1998-2024

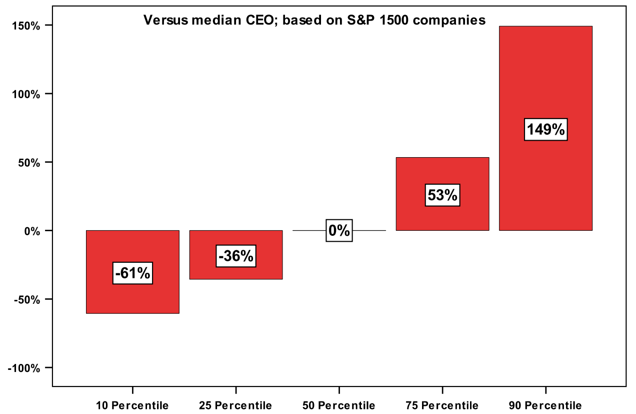

Figure 4

Pay Leverage 2020-2024

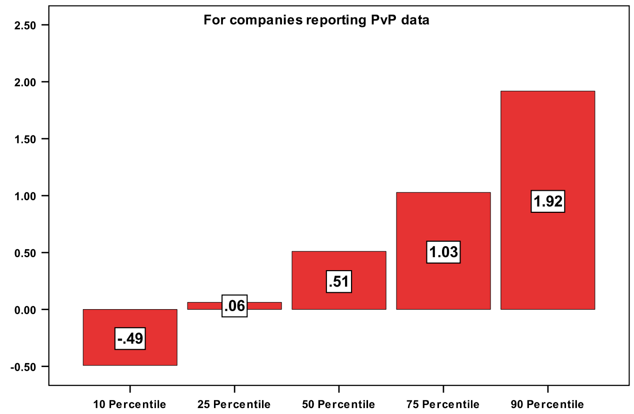

is a quite strong incentive. Figures 3 and 4 show that it's above the 75th percentile compared to S&P 1500 CEOs for the years 1998-2024 (1.33) and against the CEOs reported in the 2020-2024 PvP disclosures (1.03).

The pay premium at industry average performance

The intercept of the trendline is the ln pay premium at industry average performance. We can convert the ln pay premium to a percentage premium by taking the anti-log, i.e., percentage pay premium = exp(ln pay premium) – 1. Gerstner's percentage pay premium, shown in Figure 2, is 53%. The pay premium is measured relative to the sum of market pay. Market pay values stock and options based on their values at the time of grant. Those values represent present values, while what's incorporated in mark to market pay represents the future value of the stock or option. Since future values exceed

Figure 6

CEO Ln Pay Premium 1998-2024

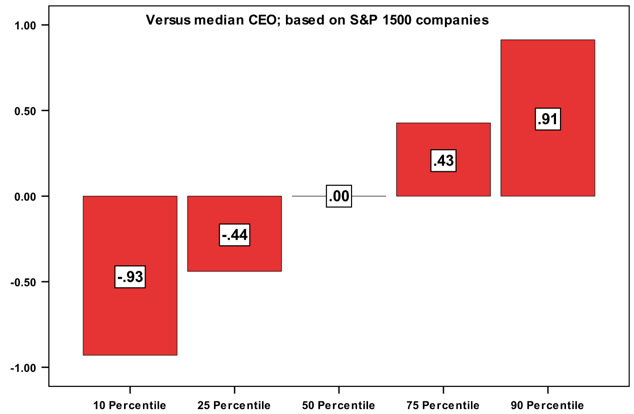

present values, on average, by the expected return on the security (stock or option), we would expect companies that set their target pay at market to show a mark to market pay premium at industry average performance. We can adjust for the expected accretion in equity compensation by subtracting the median company's ln pay premium from each company's ln pay premium. For Gerstner, that reduces his percent pay premium from 53% to 33%. Figure 5 shows that this puts him at about the 65th percentile. The 90th percentile percentage pay premium in Figure 5, +149%, is much higher, in absolute value, than the 10th percentile, -61%. This suggests that our models are biased because they show much higher positive pay premiums than negative pay premiums. It's important to remember that our models use logarithms because they fit the data much better than models using dollars. Figure 6 shows that the log pay premiums are almost perfectly symmetric with an equal balance of over-payment and under-payment.

Pay alignment and relative pay risk

The pay-performance trendline quantifies two additional pay dimensions. The correlation, what we call pay alignment, is the third dimension and it measures the degree to which ln relative pay and ln relative TSR move together. A correlation of 1.0 means that a positive one standard deviation change in ln relative TSR is always associated with a positive one standard deviation change in ln relative pay. A correlation of -1.0 means that a positive one standard deviation change in ln relative TSR is always associated with a negative one standard deviation change in ln relative pay. A correlation of zero means that changes in ln relative TSR and ln relative pay are unrelated to each other. Figure 2 shows that Gerstner has very high pay alignment, 0.99. The fourth dimension is relative pay risk which is the ratio of ln relative pay variability to ln relative TSR variability, or, more specifically, the ratio of the standard deviation of ln relative pay to the standard deviation of ln relative TSR. Gerstner's relative pay risk is 1.74. This means that his ln relative pay is 74% more variable than ln relative TSR. Pay leverage is the product of pay alignment and relative pay risk. [6]

Assessing the impact of pay on future stock returns

Our analysis of Gerstner's 1994-1998 pay shows that he has a strong incentive, with pay leverage of 1.72, but it also shows that he has a pay premium at industry average performance of +53%. We would expect that his strong incentive would have a positive effect on shareholder value, but his pay premium might have an offsetting negative effect on shareholder value. The dollar magnitude of his pay premium is tiny compared to the aggregate shareholder value of IBM, so his pay premium by itself would not cause any detectable decline in IBM shareholder value. But his pay premium might lead to higher pay for many other IBM employees – perhaps because Gerstner was trying to give them strong incentives too – and this indirect effect could have a detectable effect on IBM shareholder value.

Our models of future returns should be useful to three audiences

In this paper, we will present the results of models of future relative returns using pay factors that can be calculated from the new PvP disclosures. The initial dependent variable in the these models is the natural log of (1 + future 3 year relative return). These models should be useful to three audiences. The first audience is corporate directors, compensation consultants, proxy advisors and institutional investors. This audience can use the models to evaluate corporate pay programs by calculating the expected impact of the CEO's pay on future relative returns. A negative predicted return is a sign that top management pay should be re-evaluated. A positive predicted return is evidence that the program is aligned with shareholder value. The second audience is investors using factor models to estimate individual stock returns, and thereby, improve their stock selection. This audience can use the models by including the pay dimensions, or the predicted returns from the pay dimensions, as additional factors in their factor model. The third audience is hedge funds seeking trading gains from pay dimensions. This audience can use the model to achieve trading gains by going long companies with positive pay factors and short companies with negative pay factors.

Our models can be enhanced by taking account of pay dimensions calculated from grant date pay. They can also be made more useful for hedge funds by changing the dependent variable from the natural log of (1 + future 3 year relative return) to the percentage gross return to recognize that hedge funds can't invest in log returns and that it's difficult for investors to hold long and short positions in the relative return. These enhancements are beyond the scope of this paper.

We'll focus first on future three year returns because we believe that three years is a reasonable horizon for pay dimensions to affect stock returns. We expect pay dimensions to affect stock returns by changing management behavior and we anticipate that it takes time for pay dimensions to change management behavior and for changes in management behavior to affect the stock price. We'll see that our long-short portfolio simulations confirm our expectation that pay dimensions affect performance gradually. Our simulations show that companies with positive pay factors out-perform those with negative pay factors for one, two, three, four and five year periods but also show that the magnitude and statistical significance of the return differences increase with the duration of the future return period.

The variables in our models of future relative returns

Our models use four variables in addition to the pay dimensions discussed earlier. Two variables are company characteristics known to affect returns, i.e., company size and prior company relative return. The third variable is time, which captures differences in return due to overall market factors. The fourth variable is stock ownership. We include stock options, restricted stock and performance shares in stock ownership. For the median CEO, these incentive compensation holdings represent 55% of total stock ownership. We take account of stock ownership because it provides an additional source of financial incentive for the CEO. The pay dimension variables are pay leverage, pay alignment, the pay premium at industry average performance and relative pay risk.

The sample for our model of future 3 year returns is 26,199 five year periods ending in 1998-2021. We start with 1998 because it is the first year in which we have at least 500 cases. We end with 2021 because we need a subsequent three year return to assess the impact of the pay dimenions. To limit the impact of extreme outliers, we winsorize all our variables at the 1st and 99th percentiles. Winsorizing keeps all 26,199 cases in the sample but truncates variable values at 1st and 99th percentiles. For each of our variables, we need to decide what function of the variable has the best linear relationship with ln(1 + future three year relative return). For stock ownership, we first scale the variable by market pay. This gives us a better fit than stock ownership by itself. We then find that the natural log of (stock owned/market pay) fits the data better than the ratio itself. The log-log model implies that a given percentage change in the stock ownership multiple has a constant percentage effect on future relative shareholder wealth. The effect is small, a 1% increase in stock ownership multiple increases future relative shareholder wealth by only 7/1000 of 1%. We also test whether the incremental effect of stock ownership disappears at some high level of stock ownership multiple. We find that ln(stock owned/ market pay) has a higher correlation with ln(1 + future three year relative return) when we truncate the stock ownership multiple at 60x. This raises the coefficient of stock ownership from .007 to .009.

Our analysis of individual company pay (as shown in Figure 2) uses a log-log model where the pay premium at industry average performance is expressed in logs. We use the log pay premium as an explanatory variable in our model of ln(1 + future 3 year relative return) but test whether positive and negative pay premiums have similar effects on future returns and whether the incremental negative effect of a positive pay premium disappears at some high level of pay premium. We find that positive and negative pay premiums have different effects. Positive pay premiums have a negative (and statistically significant) effect on future returns but negative pay premiums do not have a negative or statistically significant coefficient for future returns. In other words, high pay is a bad thing, but low pay is not necessarily a good thing. Given this finding, we only use positive pay premiums as an explanatory variable. We find that ln positive pay premium has a higher correlation with ln(1 + future three year relative return) when the ln positive pay premium is truncated at 1. This makes the ln positive pay premium coefficient substantially more negative, -.054 vs -.020. A ln positive pay premium of 1 implies a percentage pay premium of 172%, so our truncation implies that CEO pay premiums in excess of 172% have no further negative effect on future returns.

Relative pay risk has a bigger impact than pay leverage

Our third variable is relative pay risk, which has a positive effect on ln(1 + future 3 year relative return). We test whether the incremental positive effect of relative pay risk disappears at some high level of relative pay risk. We find that relative pay risk has a higher correlation with ln(1 + future three year relative return) when relative pay risk is truncated at 4. This raises the coefficient of relative pay risk from .014 to .023. The truncation point (4) is more than double the relative pay risk (1.74) we found for Lou Gerstner in Figure 2. It is a surprise to me that relative pay risk – one of the two components of pay leverage (which equals pay alignment x relative pay risk) – has a much stronger impact on future stock returns than pay leverage. My tentative explanation for this is that relative pay risk has less "noise" than pay leverage. Relative pay risk is statistically significant for 80% of all cases while pay leverage is statistically significant for only 40% of all cases. Because it is less noisy, relative pay risk may have a bigger impact on executives' perceived incentive than pay leverage.

It's easier to explain differences in predicted returns if we express our variables in terms of differences from annual mean values. In this model, the dependent variable is ln(1 + future 3 year relative return) and the independent variables are the annual mean value of ln(1 + future 3 year relative return) and the differences from the annual variable mean for the other variables. This model explains 1.6% of the variation in ln(1 + future 3 year relative return) across 26,199 cases. The first variable, which captures the impact of differences in time, explains 0.8% of the variation. Company size and prior relative return explain an additional 0.5%. The incentive and cost variables explain an additional 0.3% of the variation. While the variance explained is small, the incentive and cost variables are statistically [7] and economically significant.

Calculation of the predicted future return for Lou Gerstner

Let's look at the calculation of the predicted ln 3 year relative return from pay factors, .0262, for Lou Gerstner at the end of 1998. The contribution from his stock ownership is .0238, the contribution from his relative pay risk is .0064 and the contribution from his positive pay premium is -.0040. Stock ownership includes incentive securities such as stock options, restricted stock and performance shares. For Gerstner at the end of 1998, his incentive securities were worth $539 million, or 84% of his total stock ownership of $638 million. His total stock ownership was 34.3 times his current sales market pay of $18.6 million and the log of his stock ownership multiple, 3.53, was 1.71 greater than the mean log stock ownership multiple of 1.82. When we multiply his log stock ownership difference, 1.71, by the log stock ownership multiple coefficient, .0142, we get the ln future return contribution from his stock ownership, .0238. The second incentive variable, relative pay risk, made a much smaller contribution to Gerstner's predicted ln 3 year relative return. His pay leverage of 1.74 was only 0.28 above the average relative pay risk of 1.47. When we multiply his relative pay risk difference, 0.28, by the relative risk coefficient, .0230, we get the ln future return contribution from his relative pay risk, 0.0064. Gerstner's ln positive pay premium of 0.422 was only 0.063 above the mean ln positive pay premium of 0.359. When we multiply his ln positive pay premium difference, 0.063, by the ln positive pay premium coefficient, -.0639, we get the ln future return contribution from his positive pay premium, -.004.

There are substantial differences in predicted future returns from stock & pay factors

A useful way to assess the economic significance of the incentive and cost variables is to look at the differences in predicted future 3 year relative return from pay factors between the 90th and 10th percentiles and between the 75th and 25th percentiles. Figure 7 shows that the predicted 3 year relative return at the 90th percentile is +4.6% vs -3.8% at the 10th percentile, a difference of 8.4 percentage points. The predicted relative return at the 75th percentile is +2.1% vs. -2.3% at the 25th percentile, a difference of 4.4 percentage points. No corporate director should approve a corporate pay plan without using a model like this to assess its prospective impact on future stock returns.

Relative pay risk has the biggest impact on future returns

For Lou Gerstner at the end of 1998, his accumulated incentive securities had a much bigger impact on his predicted future ln 3 year relative return than his relative pay risk. But that's unusual. We can assess the relative importance of stock ownership vs the other two components of the predicted future 3 year relative return – relative pay risk and the pay premium at industry average performance – by using the three components to explain the total predicted future 3 year relative return. When we do this, we find that relative pay risk accounts for 46% for the variation in the total predicted return, the pay premium accounts for an additional 33% and stock

Figure 7

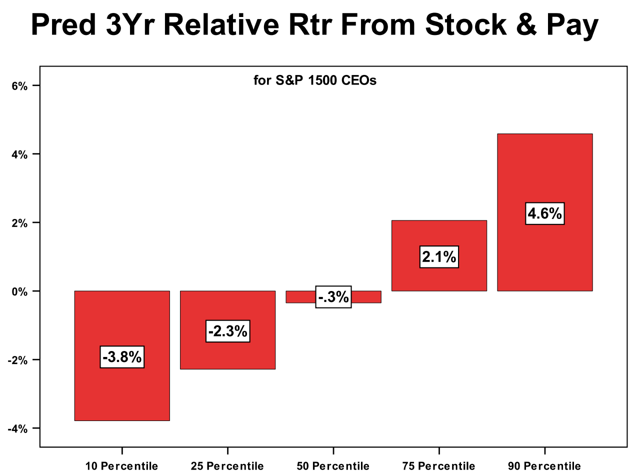

Figure 8

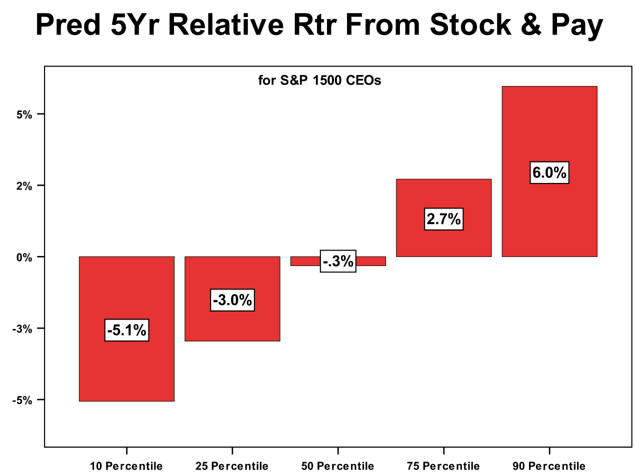

ownership accounts for the remaining 21%. This shows that the pay dimensions that emerge from our pay leverage analysis are extremely important in assessing the impact of incentive and cost variables on future stock returns.

Pay dimensions affect five year future returns

We gain additional perspective on the impact of incentive and cost variables if we analyze future 5 year relative returns. This model explains 1.3% of the variation in ln(1 + future 5 year relative return) across 21,696 cases. The incentive and cost variables explain 0.4% of the variation. Again, we find that the incentive and cost variables are statistically [8] and economically significant. Figure 8 shows that the predicted relative return at the 90th percentile is +6.0% vs -5.1% at the 10th percentile, a difference of 11.1 percentage points. The predicted relative return at the 75th percentile is +2.7% vs. -3.0% at the 25th percentile, a difference of 5.7 percentage points. As before, a useful way to assess the relative importance of stock ownership vs the other two components of the predicted 5 year relative return - relative pay risk and the pay premium at industry average performance – is to use the three components to explain the total predicted future 5 year relative return. When we do this, we find that relative pay risk accounts for 37% for the variation in the total predicted return, the pay premium accounts for additional 39% and stock ownership accounts for the remaining 24%. This shows again that the pay dimensions that emerge from our pay leverage analysis are extremely important in assessing the impact of incentive and cost variables on future stock returns.

The impact of the conventional pay dimensions: percent at risk & percent from market

Another model of future relative returns provides additional insight into the problems of the conventional wisdom. When we substitute percent of pay at risk for relative pay risk and the gross grant date pay premium for performance adjusted cost, we find two important results. First, percent of pay at risk, unlike relative pay risk, has a negative coefficient and is not statistically significant. This highlights a fundamental problem with the conventional wisdom used by companies and proxy advisors: they have no meaningful measure of incentive strength. Second, the gross grant date pay premium is statistically significant. When we split the gross grant date pay premium into two pieces, the component that is related to performance adjusted cost and a residual component, we find that the coefficient on the component related to performance adjusted cost is 60% larger (in absolute value) than the coefficient on the residual component. This confirms that performance adjusted cost is a better cost measure but also shows that investors provide some punishment for high pay that is due to incentives and superior performance, not the pay premium at industry average performance. Hopefully, this investor reaction will disappear as investors become more knowledgeable about better measures of incentive strength.

Returns from going long good pay factors and short bad pay factors

A demanding test of the economic significance of our incentive and cost variables is the payoff from going long companies with good pay factors and short companies with bad pay factors. Here, for brevity, we're treating stock ownership as a "pay factor". The basic concept for the long-short portfolios is to use the predicted return from our incentive and cost variables to select the long and short portfolios. To do a fair simulation, we need to make sure that our predicted returns don't make any use of data that's not known when the long and short portfolios are selected each year. To make sure that happens, we use predicted returns from a model of future 3 year returns ending before the portfolios are selected. For example, for the portfolios selected just after fiscal year 2006, we use data for the five year periods ending in 1999-2003 to predict subsequent three year returns. The last three year future return included in the model is the three year relative return for 2004-2006 and the long and short companies are selected three months after fiscal 2006 ends.

We construct our long and short portfolios each year in four steps. First, we calculate quintiles based on company size, measured by market equity value. Second, we calculate prior 3 year relative return quintiles within each size quintile. Third, within each of the 25 size/prior relative return cells, we calculate incentive predicted return terciles. The incentive predicted return is the predicted return from stock ownership and relative pay risk. Fourth, within each of the 75 size/prior relative return/incentive predicted return cells, we calculate cost predicted return terciles. The cost predicted return is the predicted return from performance adjusted cost. At the end, we have incentive predicted return terciles and cost predicted return terciles within each of the 25 size/prior relative return cells. Taking the sum of the two tercile ranks gives us 6 groups within each of the 25 size/prior relative return cells. Our short portfolio is companies in the bottom terciles for both incentive and cost. Our long portfolio is companies in the top terciles for predicted return from both incentive and cost.

When we look across all years, our long portfolio has a higher return than our short portfolio for 1 year, 2 years, 3 years, 4 years and 5 years. The mean differences for 1 and 2 years are 3.2 and 3.2 percentage points, but those differences are not statistically significant at the conventional 5% level. The mean differences for 3, 4 and 5 years are 7.7, 14.7 and 19.6 percentage points and those differences are all statistically significant. [9] We can make these returns comparable to other investment opportunities if we express them as annualized returns on the capital required to support the long-short hedge. FINRA requires a margin of 10% of the long position. [10] This means that an 7.7 percentage point 3 year return on the initial gross position implies a 77.0 percentage point 3 year return on the capital required to support the hedged portfolio. This represents an annualized return on capital of 21.0%. Similarly, a 19.6 percentage point 5 year return on the initial gross position implies a 196.0 percentage point 5 year return on the capital required to support the hedged portfolio. This represents an annualized return on capital of 24.2%.

A challenge for the construction of profitable long-short portfolios is that our models are based on log returns, not dollar returns. If investors were able to invest in log returns, the long-short portfolios would be much more profitable! The t-statistic for the differences in future 3 year ln returns is 4.6, while the t-statistic for the difference in future 3 year percentage returns is only 2.7.

The shortcomings of academic research on executive pay and future returns

A widely cited paper on CEO pay and future shareholder returns finds that "measures of Chief Executive Officer (CEO) excess compensation are negatively related to future firm returns and operating performance." [11] Our measure of excess compensation is the pay premium at industry average performance and we also find that excess compensation has a negative effect on future returns. But that's not the whole picture on executive pay. Incentives, measured by relative pay risk and stock ownership, have a positive effect on future stock returns. Cooper, Gulen and Rau believe they have taken incentives into account because their model of excess compensation uses stock return and return on assets as explanatory variables. But their model assumes that all companies have the same pay leverage. They make no effort to measure CEO pay leverage or wealth leverage at the individual company level. An additional flaw of their paper is that their pay model is mis-specified. Their model relates dollars of pay to dollars of sales, stock return, ROA and other variables and only explains 7% of the variation in CEO pay across 26,582 cases. Their model is mis-specified because log pay has a much stronger relationship to log sales than dollars of pay have to dollars of sales, and the log-log model has roughly equal errors as a function of size while the dollar-dollar model has massive heteroskedasticity. The standard error in a dollar-dollar model exceeds total pay for smaller companies. Cooper, Gulen and Rau divert the reader's attention from that problem by not reporting the standard error for their model.

Conclusion: directors, proxy advisors and investors need better measurement

Corporate directors, proxy advisors and institutional investors rely on poor measures of executive pay: percent of pay at risk as a proxy for incentive strength and target pay percentile as a proxy for retention risk and shareholder cost. We have shown that those are bad measures, that an individual company regression of relative pay on relative TSR provides much better measures: pay leverage, the pay premium at industry average performance, pay alignment and relative pay risk, and that relative pay risk and performance adjusted cost have statistically and economically significant effects on future stock returns. It's time for corporate directors, proxy advisors and institutional investors to acknowledge the need for better measurement of pay dimensions and the importance of making decisions about pay plan design, say on pay voting and stock selection on the basis of empirical evidence about pay dimensions and future firm performance.

APPENDIX

Our mark to market pay is a good proxy for Compensation Actually Paid

We estimate cumulative mark to market pay, using grant data, for each year of each five year period ending in 1998-2024. Our mark to market pay is similar to Compensation Actually Paid (CAP) in trying to measure pay based on the year end value of unvested or unexercised equity compensation. On the surface, CAP looks different because it doesn't provide cumulative mark to market pay. CAP provides an annual figure that includes the change in value during the year of unvested equity granted in prior years. Fortunately, this isn't a substantive difference because we can convert Compensation Actually Paid to cumulative pay simply by adding up the annual values.

Both methodologies are the same in their treatment of salary, annual bonus, restricted stock and performance shares. Salary and cash bonuses are included in mark to market pay based on their value when paid. The values of restricted and performance shares are finalized at vesting and no further adjustment is made. There are four differences in methodology between CAP and Shareholder Value Advisors (SVA). First, CAP includes changes in value of unvested equity granted prior to the five year measurement period, while SVA only includes the value of equity awards granted during the five year measurement period. This is a weakness of CAP because it mis-matches the pay and performance periods. This methodology difference can be largely eliminated because it is possible to use holdings data to estimate the value changes attributable to prior grants and then calculate an adjusted CAP that excludes the value changes attributable to grants made prior to the five year measurement period. The second methodology difference is that CAP includes actual multi-year cash incentive plan values based on the year of payment, while SVA includes estimated multi-year cash incentive plans values based on the year of grant. The SVA methodology should provide better pay dimension measures because it matches the pay and performance periods, only including in pay awards that are based on performance during the five year performance period. We don't try to adjust CAP for this difference, but, fortunately, multi-year cash plans are small enough so that cumulative adjusted CAP and SVA mark to market pay are highly correlated across the middle 90 percent of cases, as Figure 9 shows. The third major methodology difference is the treatment of stock options. CAP includes the expected value of the option in mark to market pay at the date of vesting and makes no adjustment for valuation changes

Figure 9

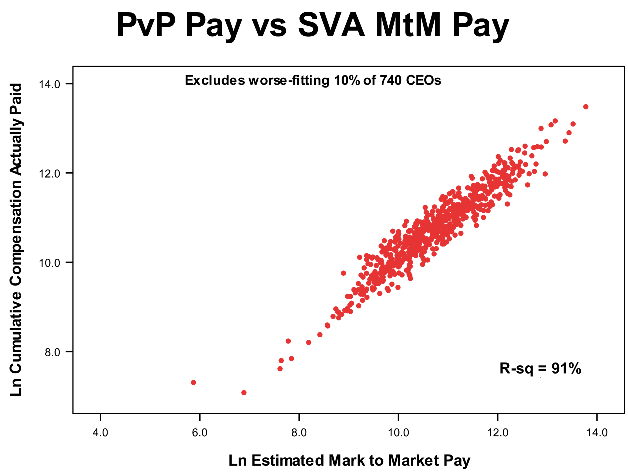

between vesting and exercise. SVA mark to market pay revalues the option after vesting up to the time of exercise. We assume that the option is exercised after vesting as soon as the stock price is double the exercise price. Our methodology should provide better pay dimension measures because vested options do provide a financial incentive until they are exercised.

Fourth, CAP reflects the company's actual vesting period and estimates of vesting multiples based on the company's complete performance data. SVA mark to market pay requires estimates of vesting period and vesting multiples.

How we estimate vesting periods

Our historical pay database, Standard & Poor's Execucomp database, does not include data on vesting periods, so we need to estimate vesting periods. We use data on grants and holdings to estimate vesting periods. We estimate the vesting period separately for stock options, restricted stock and performance shares. For stock options, we estimate average time to exercise, but, for brevity, we refer to the average time to exercise as the vesting period in this methodology summary. The estimated vesting period is the number of years that minimizes the difference between actual holdings and expected holdings, expressed as a percentage of actual holdings. Expected holdings are calculated from actual grants, the assumed vesting period and the assumption of pro-rata vesting. We evaluate the difference between actual and expected holdings for vesting periods of 2 to 8 years and pick the period with the smallest percent difference in absolute value. Across all top 5 executives, the average estimated vesting period is 3.2 years for stock options and performance shares and 3.9 years for restricted stock.

A major challenge in developing accurate vesting periods for each company/year is finding a way to adjust for the bias that comes from short service. For top 5 executives with 9+ years of service, the average estimated vesting periods are 3.6 years for stock options, 3.7 years for performance shares and 4.9 years for restricted stock. We limit the sample used to estimate the vesting period to the longest-serving top 5 executives in each company (using data for the current year and prior nine years), but this isn't perfect because the longest service period is less than 9 years for 60% of the companies granting performance shares. To improve our vesting period estimates for companies where the maximum service period is less than 9 years, we do a regression across companies where the dependent variable is the mean estimated vesting period and the independent variables are dummy variables for maximum service periods of 2 years to 9+ years. We use this regression to correct the bias created by a maximum service period less than 9 years. For example (using data from the 2015 performance share regression), the mean vesting period for companies with a maximum service period of only 4 years is 2.57 years, 1.14 years less than the mean vesting period for companies with a maximum service period of 9+ years. For the companies with a maximum service period of only 4 years, we increase the performance vesting period calculated from their maximum service executives by 1.14 years to offset the bias created by the company's short maximum service period.

When we adjust for short maximum service periods, the company average estimated vesting period is 3.5 years for stock options, 3.7 years for performance shares and 4.2 years for restricted stock. The modal vesting periods are 3 years for stock options and performance shares and 4 years for restricted stock but less than 37% of companies have the model vesting period.

How we estimate vesting multiples

Companies use a wide range of metrics to determine vesting multiples for performance share and performance cash plans. One study found that 72% of companies used stock price/TSR, 47% used earnings, 35% used accounting return measures, 30% used revenue, 20% used cash flow and lesser percentages used many other measures. [12] Execucomp does not report the metrics companies use for performance vesting. For simplicity and economy of effort, we decided, many years ago, to use relative TSR to estimate performance vesting multiples. We assume the performance share or performance cash award is made at the end of the fiscal year and we calculate the company's TSR percentile vs the company's GICS industry for each subsequent year until the end of the vesting period. We convert the company's TSR percentile into a performance vesting multiple using a common performance-payout schedule. The schedule we use provides a zero multiple below the 25th percentile, a 0.50 multiple at 25th percentile performance, a 1.0 multiple at 50th percentile performance and a 2.0 multiple at 75th percentile performance and above. We use linear interpolation between the 25th and 50th percentiles and between the 50th and 75th percentiles.

How we estimate market rates

Market rates for executives have been estimated using a regression relating log pay to log revenue for more than 70 years. Regular executive pay surveys were started after the Second World War. The best known, and most widely used, early survey was conducted for the American Management Association by Arch Patton, a leading compensation consultant and a partner of McKinsey & Company. Initially, Patton adjusted for differences in company size by reporting pay relative to total profits. By 1955, however, he had converted to showing pay relative to revenue using log scales. [13] A log pay – log revenue regression fits executive pay data much better than a dollar pay – dollar revenue regression. The log regression makes the assumption that a given percentage change in revenue is associated with a constant percentage change in pay. This assumption fits executive pay data far better than the assumption that a $1 billion change in revenue is associated with a constant dollar change in pay. Each additional $1 billion in revenue tends to be associated with a smaller dollar change in pay.

We use market pay to calculate relative pay and then use relative pay to estimate pay dimensions. The median coefficient of log revenue in our executive pay regressions by industry is 0.45 but 10% of the coefficients are below 0.32 and another 10% are above 0.58. We don't want noise in pay sensitivity to revenue to distort the relationship between pay and performance. We want pay sensitivity to revenue, within each industry, to be consistent across the top 5 executives and to be consistent across the five year time horizon we use for our pay leverage calculations. We have developed a methodology to estimate market rates that makes pay sensitivity to revenue constant across jobs and time (for five years) within each industry. We start by adjusting pay and revenue for inflation. We then express the log-log regression in terms of differences from job means using five years of historical data. The traditional regression is:

Ln inflation adjusted pay = $\alpha + \beta$ x ln inflation adjusted revenue

We re-express this equation using five year job means within each industry and differences from those job means:

Ln infl adj pay = job mean ln infl adj pay + β x (ln infl adj rev – job mean ln infl adj rev)

When we fit this regression to five years of inflation adjusted pay for top 5 executives within an industry, we get a constant log revenue slope for five years and for all top 5 executives.

1. Author's calculations from Standard & Poor's Execucomp database. Over the same period, median CEO percent of pay in equity increased from 46% to 67%.

2. Jochem, Torsten, Gaizka Ormazabal and Anjana Rajamani (2024), "Why Have CEO Pay Levels Become Less Diverse?", ECGI Finance Working Paper No 707/2020.

3. See Stephen F. O'Byrne, "Why and how U.S. executive pay should change", openaccessgovernment.org. It provides an example showing pay leverage of 0.61 for good early performance and 2.18 for bad early performance.

4. The new requirement was adopted in 2022 pursuant to the Dodd-Frank Act of 2010.

5. This assumption fits the data much better than the assumption that a $1 change in relative shareholder wealth results in a constant dollar change in relative mark to market pay.

6. Pay leverage is a single regression coefficient and any single regression coefficient is equal to correlation x (dependent variable standard deviation/independent variable standard deviation).

7. The t-statistics are 4.5 for stock ownership, -6.9 for the pay premium and 7.3 for relative risk. Using robust standard errors based on company, the coefficient/standard error ratios are 3.1 for stock ownership, -5.1 for the premium and 6.1 for relative risk.

8. The t-statistics are 4.4 for stock ownership, -6.8 for the pay premium and 6.4 for relative risk. Using robust standard errors based on company, the coefficient/standard error ratios are 2.6 for stock ownership, -4.4 for the pay premium and 4.7 for relative pay risk.

9. The t-statistics are 2.7 for 3 years, 3.7 for 4 years and 3.9 for 5 years.

10. See www.finra.org/rules-guidance/rulebooks/finra-rules/4210.

11. Cooper, Michael, Huseyin Gulen and P. Raghavendra Rau (2016), "Performance for pay? The relation between CEO incentive compensation and future stock price performance", available at ssrn.com/abstract=1572085. Quote is from the abstract. This paper has13 citations and 11,496 downloads from ssrn.com as of November 15, 2025.

12. Clearbridge Compensation Group, Long-Term Incentive Plan Report October 2023, available at clearbridgecomp.com.

13. Patton, Arch (1955), "Building On the Executive Compensation Survey", Harvard Business Review, Vol 33, No 3, May/June.

CONTACT
Shareholder Value Advisors, Inc.
21 Bonnie Way,
Larchmont, NY 10538

Main Phone Number: (914)-833-5891
General Inquiries: info@valueadvisors.com
Website: https://valueadvisors.com/