Are CEOs paid extra for riskier pay packages?*

Ana Albuquerque Rui Albuquerque Mary Ellen Carter
Qi (Flora) Dong

November 2019

Abstract

A fundamental hypothesis in optimal contracting models is that incentive pay is expensive to the principal because risk averse agents require extra pay to accept riskier pay packages. We test this hypothesis on U.S. CEO compensation data using a variety of datasets and empirical approaches. We find that CEOs with riskier pay packages are paid more on average, and that proxies for risk aversion affect the sensitivity of pay to variance in pay. The elasticity of pay to variance in pay appears to be economically small, suggesting, among other things, that CEO pay packages are saturated with incentives.

Keywords: CEO pay, contract theory, moral hazard, participation constraint, risk aversion, Incentive Lab, realized variance, ARCH.
JEL Classification: D81, G30, J33, M52.

*We thank Joshua D. Anderson, Kee-Hong Bae, Philip Bond, Pierre Chaigneau, Eitan Goldman, Luis Marques, Yihui Pan, Jeff Pontiff, Jean-Charles Rochet, Jun Yang, and Wenyu Wang for their comments as well as seminar participants at Boston College, Boston University, Indiana University, and at the conferences, 2019 FIRS, joint McGill-HEC Montreal, Telfer Annual Conference on Accounting and Finance, and 2019 Western Finance Association. We thank Michael Guo for excellent research assistance in getting detailed information about compensation contracts from Proxy Statements. Ana Albuquerque: Boston University, albuquea@bu.edu. Rui Albuquerque: Boston College, ECGI, CEPR, rui.albuquerque@bc.edu. Mary Ellen Carter: Boston College, maryellen.carter@bc.edu. Flora Dong: Kennesaw State University, qdong2@kennesaw.edu. Rui Albuquerque gratefully acknowledges financial support from the European Union Seventh Framework Programme (FP7/2007- 2013) under grant agreement PCOFUND-GA-2009-246542 and from the Foundation for Science and Technology-FCT under grant PTDC/IIM-FIN/2977/2014. The usual disclaimer applies.

1 Introduction

A fundamental hypothesis in moral hazard models is that risk averse CEOs require extra pay for riskier pay packages. Firms providing incentive pay as a way to reduce principal-agent conflicts incur costs as incentive pay imposes additional risk on the CEO. This is a fundamental hypothesis in the sense that it is born out of the participation constraint and requires only information about the agent's expected utility. Other model predictions, including the much studied sensitivity of incentive pay to stock return volatility that we discuss below, not only rely on the participation constraint but also on other assumptions including the production function, the number of performance metrics used to incentivize the agent, and the principal's objective function. The hypothesis we study is fundamental also because agency theory continues to be the work-horse model of much of the empirical studies in executive compensation.[1]

In static principal-agent moral hazard models, a feasible contract between a principal and an agent prescribes an action by the agent and a payment from the principal which, among other things, must respect the agent's participation constraint. According to this constraint, the agent's expected utility under the contract equals the agent's utility under an alternative opportunity (Grossman and Hart, 1983, prove that the participation constraint binds under general conditions). A general implication of the participation constraint is that conditional mean pay must equal the conditional variance of pay times risk aversion.[2] A parallel to this relation exists in some asset pricing models where expected returns are related to the variance of returns times risk aversion. Just as in the asset pricing literature, the trade off between mean pay and variance of pay is not a causal relation because both variables are determined in

[1]Conyon, Core and Guay (2011) call this hypothesis the "central tenet of agency theory and contracting" (p. 404). We discuss below how our paper relates to Conyon, Core and Guay (2011) and Fernandes, Ferreira, Matos, and Murphy (2013). We also discuss below why the trade off between CEO pay and risk in CEO pay is more fundamental than the question about how *firm* risk is linked to *incentive* pay.

[2]This statement is exact up to a constant of proportionality with mean-variance utility. It is common to assume exponential utility and normally distributed shocks, which results in a preference specification equivalent to mean-variance utility (see Holmstrom and Milgrom, 1987). We discuss more general preference specifications below.

equilibrium. Rather, it defines a structural relation between the two variables: any shock to one of the variables must be met with a shock to the other that preserves their structural relation. It is this structural relation that we test.

We use three different empirical approaches to test the hypothesis that CEOs with more volatile pay receive higher pay on average. The first empirical approach uses Incentive Lab's detailed contract information on the relation between performance metrics and performance-based compensation (i.e., cash bonus, equity, and options grants) in actual CEO contracts. The main advantage of these rich data is that we can evaluate simultaneously, through a simulation exercise, the beginning-of-the-year expected value and variance of the CEO's end-of-year pay. We are therefore able to explicitly consider the underlying sources of the riskiness of CEO pay packages. Consider, for example, the simplest contract that includes salary plus a cash bonus grant. The bonus grant may have a *threshold* payout of 100% of base salary, a *target* payout of 200% of base salary, and a *maximum* payout of 400% of base salary. The contract defines a metric, say net sales, and performance levels that determine threshold, target, and maximum payouts. By simulating the performance metric,we are able to construct mean and variance of total pay.[3] To peek into the full complexity of this simulation exercise note that CEO contracts may have grants of bonus, stock, and options in any given year, and in fact, multiple grants of each kind, possibly with multiple performance metrics across grants and even also in the same grant, where performance metrics may or may not all have to be met to yield a payout, and finally, that these characteristics change over time even for the same firm-CEO pair. In our regressions, we find a positive and statistically significant elasticity of conditional mean of pay to conditional variance of pay.

The second empirical approach is to model the conditional variance of pay with realized variance of past CEO pay in the spirit of Schwert (1989) and Andersen and Bollerslev (1998). Using realized variance has the benefit that it allows for the uni-

[3]Bizjak, Kalpathy, Li, and Young (2018) state that compensation consultants often use simulations when presenting the valuation of the awards to the board of directors. So, it is conceivable that compensation committees at the beginning of the year use a similar approach to evaluate whether enough pay is being offered on average to the CEO to compensate for the risk in her compensation package.

verse of ExecuComp firms to be used, increasing the sample size considerably in the cross section and time series. However, it comes at the cost of having to assume that contract parameters are time invariant for each CEO, an assumption that is nonetheless standard in panel data studies estimating pay-for-performance parameters. The results using realized variance of total CEO pay are broadly consistent with the results using Incentive Lab data that explicitly deal with time variation in contract parameters, though the point elasticities are somewhat lower.

The third empirical approach to test the hypothesis that CEOs require higher pay for higher risk in their pay uses a variant of Engle's (1982) autoregressive conditional heteroskedasticity (ARCH) model to jointly estimate the mean of pay and the volatility of pay.[4] The ARCH model estimates an equation for the level of pay and another equation for the volatility of pay. This simultaneous determination of mean and variance of pay is consistent with the theory and constitutes an advantage over the realized volatility approach. We minimize the limitation of this approach of requiring a large time series of data by estimating the model in the full panel of ExecuComp firms. We use an ARCH-in-mean model where the conditional volatility of pay enters the mean equation of pay. The point estimates of the elasticity of total pay to variance of pay are positive and statistically significant.

Across the three approaches, we find estimates of elasticity of average pay to variance of pay ranging from 0.10 to 0.20. These estimates imply coefficients of risk aversion between 1.2 and 1.4, arguably low values. To understand the economic significance of this elasticity, we conduct a back-of-the envelope calculation where we ask what is the share of pay at risk in total pay implied by the estimated elasticity if all the variation in pay is driven by variation in incentives. This calculation reveals an implied share of pay at risk between 20% and 40%. This range contrasts significantly with the 75% average pay at risk in total pay in our Incentive Lab data, suggesting

[4]To the best of our knowledge this is the first paper that uses ARCH modeling to study CEO compensation data. We follow a long tradition in economics of using ARCH models to explain the time series behavior of economic variables, from inflation, in the path breaking study of Engle (1982), to GDP growth in Ramey and Ramey (1995), and to stock returns in Bollerslev, Engle, and Wooldridge (1988). The last two papers, like ours, model the conditional mean of the dependent variable as a function of its conditional variance.

that CEOs are saturated with incentives in their pay packages.[5]

We study several reasons for this discrepancy. First, we consider the possibility that risk aversion varies in the cross section. Haubrich (1994) calibrates the agency model and shows that equity incentives increase sharply as CEO risk aversion decreases. This suggests that a low estimated elasticity could be due to the presence of a few CEOs with really low risk aversion. We use several proxies for risk aversion used in the literature and find evidence consistent with higher risk aversion being associated with a higher elasticity as predicted. However, the effect is not economically significant to cover the gap identified in the back-of-the-envelope calculation.

Second, we consider a wide array of factors that may affect the relation between mean pay and variance of pay: CEO's outside market opportunities Oyer (2004); CEO preference for positively skewed payouts (Hemmer, Kim and Verrecchia, 2000, Ross, 2004, and Chaigneau, 2015); CEO overconfidence (Malmendier and Tate, 2005 and 2008); CEO power (Bebchuk and Fried, 2003); and, shocks to the CEO's marginal disutility of effort (Laffont and Martimort, 2002). We find that these characteristics do not significantly affect the previous finding of a low elasticity of mean pay to volatility in pay.

Third, an interpretation for our finding is that CEOs appear to be saturated with incentives. While Conyon (2006) suggests that in the 1990s growth in pay accompanies growth in incentive pay as predicted by the agency model, Murphy and Jensen (2018) argue that the growth in incentive pay in the last two decades seems not to be driven by the provision of economic incentives. Instead, Murphy and Jensen (2018) point to the deductibility rules of Section 162(m) of the IRS code (see also Rose and Wolfram, 2002), which have kept salaries capped at around $1 million for CEOs and restricted growth in pay to occur solely through incentive pay.[6] Our evidence

[5]Dittmann and Maug (2007) assume that risk in pay comes only from stock and options. In calibrated versions of their model, they find that many firms would benefit from increasing incentive pay to their CEOs. We examine the full spectrum of compensation risk, which not only arises from equity incentives (i.e., stocks and options), but also from cash incentives (i.e., bonuses).

[6]According to Rose and Wolfram (2002), Section 162(m) of the IRS Code appears to have kept CEO salaries capped at around $1 million over the last two decades causing the growth in pay to come from incentive pay, which was the component of pay not affected by the Section. As a result, the tax code may have created an inefficiency in pay by overexposing CEOs to risk for which they were not compensated.

suggests that while there is a link between the level of incentive pay (and volatility in pay) and mean pay, the economic magnitude of this link appears too low to be explained by agency theory considerations alone.

In their seminal work, Conyon et al. (2011) and Fernandes et al. (2013) are the first to test the prediction that the provision of incentives is expensive. Our test differs from theirs in several respects. First, we simulate bonus, equity and options grants, and from the simulation calculate expected pay and variance of pay. Their estimate of a risk premium from incentives excludes volatility in pay from bonus grants, which is important in light of changes that have lead firms to move away from option grants and to bonus grants. Second, they do not have available to them the wealth of contractual data that we have, which leads them to make assumptions that we do not require. To estimate the risk premium they require additional assumptions about the risk in the CEOs' outside opportunities. They assume that the CEOs' outside opportunity is fully diversified. This assumption may lead to an overestimation of the risk premium. In particular, if a CEO's outside opportunity is to join another firm and be equally under-diversified, then the risk premium should be small. Our simulation of variance does not require this additional assumption. Third, because our sample focuses on one country, all CEOs are exposed to the same legal, taxation, and economic environment, characteristics that may be hard to control in cross-country studies such as theirs.

There is a large literature linking *firm volatility* to equity incentives in CEO pay, referred to as a risk-return trade off. This literature has produced somewhat inconclusive results (see for example Aggarwal and Samwick, 1999, Core and Guay, 2002, and Prendergast, 2002). The risk and reward trade-off studied in this paper is *not equivalent* to the hypothesis that equity incentives decrease with firm stock return volatility, partly because the trade-off we address is about mean and volatility of CEO pay, not the volatility of firm stock returns, nor the equity incentive component of pay alone. The distinction is important. Cheng, Hong, and Scheinkman (2016) argue that there may be benefits from higher firm volatility, specifically that firms with higher volatility may also be more productive. In their model, there is a trade off between mean pay and volatility of pay implied by the agent's participation constraint as in our paper. However, whether equity incentives increase or decrease in firm return volatility in their model depends on how firm return volatility affects firm productivity (see Raith, 2003, for another model where incentives vary positively with firm risk and yet total pay co-moves positively with variance in pay). In conclusion, a positive association between mean pay and variance of pay in agency models does not necessitate or is implied by a positive association between firm volatility and equity incentives.

The paper proceeds with a formal statement of the main hypothesis in the next section and the design of the empirical tests in Section 3. Section 4 presents the data and section 5 presents the results on the main hypothesis, and offers an array of alternative specifications and robustness tests. Section 6 concludes with directions for future work. The Appendix contains details associated with the simulation exercise using Incentive Lab data, and the definition of variables used in the empirical tests.

2 Hypothesis Development

In the benchmark, static, moral hazard model the principal (i.e., shareholder) maximizes expected operating profits net of the compensation paid to the agent (i.e., CEO), w, subject to an incentive compatibility constraint and a participation constraint. Our paper focuses on the participation constraint

$$E\left(U\left(w,e\right)\right) \geq \bar{U}. \tag{1}$$

In this expression, $U\left(.\right)$ is the agent's utility, e is the agent's effort, $\bar{U}$ is the agent's utility under her best outside employment opportunity, and E is the expectations operator. In the static version of Holmstrom and Milgrom (1987) with normal shocks, exponential utility and separability between consumption and effort, $E\left(U\left(w,e\right)\right)$ can be replaced by the certainty equivalent so (1) becomes

$$E\left(w\right) - \frac{\gamma}{2}V\left(w\right) - \text{ cost of effort } = \bar{u}, \tag{2}$$

where γ is the constant absolute risk aversion coefficient, $V\left(w\right)$ is the variance of pay, and $\bar{u} = \log(\bar{U})$. A similar relation holds in the dynamic model of Holmstrom and Milgrom (1987) as well. Inserting time subscripts and moving all terms to the right-hand side except for $E(w)$,we get

$$E_t(w_t) = \frac{\gamma}{2} V_t(w_t) + \text{ cost of effort}_t + \bar{u}_t. \tag{3}$$

The information set used to compute the conditional moments refers to information available at the time when contracts are written, i.e. the beginning of period t, and w_t is the total pay realization through period t.

To test our main hypothesis, we develop below three approaches to estimating $V_t(w_t)$ and to deal with $E_t(w_t)$. In the first approach, we are able to estimate $E_t(w_t)$ and so our regression specification is

$$E_t(w_t) = \lambda\sigma_t^2 + X_t'\beta + \varepsilon_t. \tag{4}$$

In this regression model, ε_t is a measurement error, $V_t(\varepsilon_t) = V_t(w_t) \equiv \sigma_t^2$, and drivers of the cost of effort and of outside opportunities are captured by the vector X_t and the slopes β.

For our two other approaches, where we cannot rely on an estimate of $E_t(w_t)$ for our regression, define the error term, ε_t, as the unpredictable residual in pay given information available at the beginning of period t

$$\varepsilon_t \equiv w_t - E_t(w_t). \tag{5}$$

By construction, $E_t(\varepsilon_t) = 0$ and the variance of ε_t conditional on beginning of period t information is $V_t(\varepsilon_t) = V_t(w_t) \equiv \sigma_t^2$ (see Taylor, 2013, for a similar specification of the residual). Using the information about $E_t(w_t)$ contained in the participation constraint (3), we obtain our regression specification

$$w_t = \lambda\sigma_t^2 + X_t'\beta + \varepsilon_t. \tag{6}$$

Formally, the main hypothesis is

$$H_0 : \lambda > 0. \tag{7}$$

The null hypothesis is a joint test of Grossman and Hart's participation constraint and of risk aversion.[7] As a caveat, it is concievable that other models also deliver the same null hypothesis. Thus, rejecting the null hypothesis is more informative about the agency model than not rejecting it.

The main hypothesis we test relies on the fact that if the participation constraint is to hold at all times, then an increase in the conditional volatility of pay must be met with an increase in the conditional mean of pay, all else equal. That is, higher volatility in pay that comes with more incentive pay is expensive and requires higher average total pay. This is not a test of a causal relation, but rather of a structural relation given by the model, because when pay parameters are decided they impact both mean pay and variance of pay in a way that is constrained by this relation.

Our exercise is related to the literature that tries to identify the sensitivity of pay to various firm performance metrics. That literature assumes a particular specification for how pay depends on firm performance based on the solution to an optimal contracting model like the Grossman and Hart (1983) principal-agent model. Instead, hypothesis (7) does not rely on fully solving the model because it is obtained solely from the participation constraint. We see this as a benefit to our test that we can be agnostic with respect to many aspects of the model in order to formulate our main hypothesis. Our hypothesis is, in this sense, a more fundamental result of the agency model.

3 Empirical Strategy

We use three approaches to test (7), each with a different way to estimate σ_t^2.

[7]The null hypothesis was derived also with the assumptions of exponential utility and of normality of shocks. Absent these assumptions, a Taylor series expansion of utility as a function of pay shows that a trade off is present provided the utility function displays concavity–a standard assumption for utility maximization–though the utility function may also put weight on higher moments of pay as we discuss in Section 5.3 below.

3.1 Conditional variance using simulated variance from contract data

We use detailed contract information available at the beginning of each year to simulate expected pay and variance of pay for the year for each CEO. There are two main inputs for this simulation exercise: (i) compensation contract information from Incentive Lab describing the relation between contracted performance metrics and the corresponding performance-based compensation (i.e., cash bonus, equity and options grants) and (ii) Compustat data on lagged realizations of the performance metrics. We describe the procedure next leaving the details to the Appendix.[8]

The Incentive Lab contract information is presented at the firm-year-grant-metric level. For each performance metric used, Incentive Lab gives the threshold, target, and maximum level of the performance metric, and the threshold, target, and maximum level of the corresponding performance-based compensation. The CEO earns no performance-based compensation when actual firm performance is below the threshold and earns the maximum amount of performance-based compensation when actual firm performance is above its maximum. When the performance metric falls between its threshold and the maximum, the CEO earns performance-based compensation in an amount between its threshold and the maximum: following firm policies disclosed in the proxy statements (DEF 14A), we fit a piece-wise linear function between the threshold, the target and the maximum to determine the award amount.

To simulate pay for a CEO in a given year, we first simulate the performance metrics used by the CEO's firm in all the grants awarded in that year. It is possible for firms to use more than one performance metric for a given grant and to award several grants to the same CEO in a given year. We consider all metrics used for a given firm-year-CEO and simultaneously simulate all metrics for that year accounting for the joint distributional properties of the metrics. In particular, we assume a joint normal distribution for all performance metrics used. For our main results, we set the mean of the joint normal distribution equal to the last year's value of the respective

[8]Holden and Kim (2017) offer valuation formulas for performance equity grants. Because we consider also bonus and options plans and need to obtain measures of conditional volatility of pay, we have to use simulation methods.

performance metrics. We set the covariance matrix of the joint normal distribution equal to the sample covariance matrix of the performance metrics using four years of data prior to the grant year.[9] We then simulate 10,000 firm-year-grant-metric observations.

We calculate simulated compensation by fitting the simulated performance metrics to the compensation contracts. Since performance is simulated at the firm-year-CEO-grant-metric level, we calculate the simulated compensation at the firm-year-CEO-grant-metric level. We then aggregate the metric-level compensation into the grant level based on information in Incentive Lab about the relation between the various performance metrics. Compensation contracts are either separable contracts or non-separable contracts. Separable contracts allow CEOs to earn part of the bonus, equity or options grant even though some of the performance metrics do not meet their goal threshold, while non-separable contracts result in zero payout if any of the performance metric threshold is not met. Further, following Incentive Lab, we add the equally weighted pay from all metrics in separable contracts to get total simulated pay at the grant level. For a CEO with more than one grant in a given year, we add simulated pay from all her grants. We add the salary to the simulated pay values at the firm-year-CEO level and calculate the mean, variance, and skewness across the simulated values. These are respectively, the conditional mean pay, the conditional variance of pay and the conditional skewness in pay.

3.2 Conditional variance using realized variance of pay

In the second empirical approach, we estimate σ_t^2 using past CEO-firm pay data. Specifically, we use the last 5 years of w_t to compute realized variance of pay as

$$RealizedVariance_t = \frac{1}{4}\sum_{s=0}^{4}\left(w_{t-s} - \bar{w}_t\right)^2, \tag{8}$$

[9]To simulate the value of option grants we estimate the volatility of stock returns using the last five years of monthly data and cap volatility by the average volatility across all simulation years. When stock price is used as a performance metric, we take the residual variance of log price to lagged sales regressed on a time trend. The ratio of price to sales is adjusted for stock splits using lagged *ajex*.

where $\bar{w}_t$ is the 5-year sample mean. This estimator of the 5-year conditional volatility of pay is similar to Schwert's (1989) estimate of conditional monthly return volatility that uses daily data and to Andersen and Bollerslev's (1998) estimate of conditional daily return volatility that uses intraday data. The 5-year conditional volatility is a smooth function of past 1-year conditional volatilities, which may introduce an upward bias on the estimated slope coefficient in equation (6). Preempting our results, the fact that we estimate a small slope coefficient, suggests that this bias is not severe.

If pay is a function of stock returns alone, then this estimator is a consistent estimator of conditional volatility of pay. To see this point, evaluate the estimator in equation (8) applied to data generated by a model similar to Holmstrom and Milgrom (1987) where pay is a linear function of the firm's stock return, r_t, that is $w_t = m_o + m_1 r_t$, and m_0 and m_1 are optimal contract parameters. Then

$$\mathit{RealizedVariance}_t = m_1^2 \frac{1}{4} \sum_{s=0}^{4} \left(r_{t-s} - \bar{r}_t\right)^2 .$$

Andersen and Bollerslev (1998) show that under general properties for stock returns, the estimator above converges to the conditional variance of pay in the model, i.e., $m_1^2 V_t\left(r_t\right)$, where $V_t\left(r_t\right)$ is the conditional volatility of stock returns, if we are allowed to sample returns at increasingly higher frequencies.

We expect realized variance to work well as an estimator under the null that pay evolves linearly with stock returns. Intuitively, if contract parameters are time invariant, any variation in pay is due to variation in the level of performance metrics, which can be captured with past data. However, note that even in this simple example, there is more information in the realized volatility of pay than there is in the realized volatility of stock returns, because of the presence of m_1^2 that is firm specific (see, for example, Aggarwal and Samwick, 1999). More generally, the two may not even be proportional to each other as realized variance of pay entails the variance of other performance metrics and their covariances.

The main advantage of using realized variance over the Incentive Lab simulated variance is the significanty larger number of observations that come with ExecuComp. The main disadvantage over simulated variance is that realized variance is potentially a less efficient way to estimate ex-ante volatility if parameters are time variant. The

Incentive Lab simulation approach is particularly versatile along this dimension as we obtain estimates of the variance of pay that condition on firm-year-CEO actual contract parameters. While it is unclear how much of a constraint this is for the realized variance approach, we note that assuming time invariant contract parameters is the standard assumption in empirical models of CEO pay that use panel data regressions.[10]

3.3 Conditional variance using ARCH model of variance of pay

The last empirical approach to estimating σ_t^2 uses the autoregressive conditional heteroskedasticity (ARCH) model. To the best of our knowledge, this is the first paper that estimates an ARCH model for CEO pay. Empirically, we assume that variance of pay can be modeled using

$$\sigma_t^2 = \alpha + \sum_{j=1,\dots,p} \delta_j \varepsilon_{t-j}^2. \tag{9}$$

with the parameters $\alpha, \delta_j \geq 0$ and $j = 1, ..., p$ indexes the number of ARCH terms. We estimate equations (6) and (9) as an ARCH-in-mean model. The estimation uses pooled data and so the parameters α and δ_j in the volatility equation and the parameters in the mean equation (6) are assumed identical across firms. The estimation of these models is done in an unrestricted fashion and we check ex-post the non-negativity constraints on the variance-equation parameters, $\alpha, \delta_j \geq 0$.

The empirical approaches using Incentive Lab data and the ARCH model have the advantage over the realized variance approach of not requiring the assumption of time invariant contract parameters. One advantage of realized variance and ARCH over simulated variance is that by using TDC1 they accounts for the possibility that firms change compensation parameters in a discretionary fashion in the middle of the year (Kim and Yang, 2014).

[10]Some evidence in rigidity in contract parameters can be found in Shue and Townsend (2015).

4 Data

We use two main datasets, Incentive Lab by Institutional Shareholder Services and ExecuComp. Incentive Lab contains detailed compensation contract information for the 750 largest U.S. firms collected from proxy statements (DEF 14A) for CEOs and other executives over the period 1998-2016. ExecuComp contains the largest 1,500 firms in the U.S. in any given year, from 1992 to 2016. For both datasets, we limit our sample to CEOs and drop part-year CEOs. In addition, the analysis uses financial data from Compustat, stock return data from CRSP, data on board of directors from Institutional Shareholder Services, and institutional holdings data from Thomson Reuters Institutional (13F) Holdings. The variables used are described in a table in the Appendix.

In using the Incentive Lab data, we restrict the sample period to 2006-2016 because data collection is very sparse and incomplete before 2006, the year when mandatory disclosure of compensation contracts started. We restrict attention to contracts that use absolute performance metrics only (contract details for relative performance goals are generally insufficient in Incentive Lab) and to contracts with quantitative performance metrics (there is less data on qualitative performance metrics such as customer satisfaction for us to conduct a simulation exercise). We include bonus, equity compensation and options grants (we include both performance-vesting and time-vesting options grants).

We identify the specific performance metrics used in each contract. In addition to providing the name of the performance metric, Incentive Lab contains information on whether it is scaled (either by shares outstanding or by sales) or is expressed as a growth rate. It also provides additional textual information to more precisely describe the metric (e.g., when Incentive Lab field "metric" has the value of "Cashflow", Incentive Lab clarifies whether it is operating cash flow, free cash flow or net cash flow). We use this detailed information for each compensation contract. Despite the large volume of data, these data are incomplete. To define our *clean* sample, we drop firms that have actual restricted stock or bonus payments but for which there is no information in Incentive Lab. We construct an *alternative* sample that also includes observations for which we do not have complete contract information for the CEO for a given year (e.g., we may have bonus contract details but not restricted stock details despite observing that the CEO was paid some restricted stock, as well as bonus, in that year).[11]

[Sample selection numbers in the next two paragraphs need updating]

From Incentive Lab, we obtain 46,292 compensation contracts at the firm-year-grant-metric level for bonus and restricted stock grants. We are left with 4,454 compensation contracts at the metric level after excluding contracts with some missing values for the performance metrics or payouts, contracts with option grant information for the same year, and contracts where information on actual compensation is not available (i.e. salary data). The 4,454 compensation contracts at the *firm-year-grant-metric* level consist of 2,750 bonus-metric contracts and 1,704 restricted stock-metric contracts. These contracts aggregate at the *firm-year-grant* level to yield 2,954 total contracts, consisting of 1,714 bonus contracts and 1,240 restricted stock contracts. We refer to this sample as the "alternative" sample. The alternative sample has 2,127 *firm-year* observations with data available in Compustat on past performance data required for the simulation.

Our goal is to simulate total compensation for a given year, which requires all grants to be available for each year. In many instances, Incentive Lab will reports complete data on one form of compensation while reporting incomplete data or entirely omitting data on another. We are able to identify the existence of incomplete data because Incentive Lab also gives information in a separate file on all the grants paid out. We exclude firm-years with incomplete compensation contract information from the alternative sample to construct the "clean" sample. After this exclusion, we are left with 692 compensation contracts at the *firm-year-grant-metric* level, which consists of 344 bonus-metric contracts and 348 restricted stock-metric contracts. These contracts aggregate to yield 507 total contracts at the *firm-year-grant* level, consisting of 233 bonus contracts and 274 restricted stock contracts. The

[11] Bettis et al. (2018) also simulate pay from contract data in Incentive Lab. Their focus is on the valuation of equity grants only, while ours is on estimating annual outcomes of all performance grants. Therefore, our data requirements are more significant, resulting in a smaller sample size, and our simulation procedure is more complex.

clean sample has 285 *firm-year* observations with data available in Compustat on past performance required for the simulation. The online appendix presents detailed results of the sample selection procedure.[12]

Table 1 presents the distribution of performance metrics used in the compensation contracts at the metric level (Panel A) and descriptive statistics for the number of performance metrics per grant/year and the number of grants per CEO/year (Panel B). For both Panels A and B, Columns 1 to 4 present results for the clean sample and Columns 5 to 8 present results for the alternative sample.

[Table 1 about here.]

Across all awarded contracts, the use of accounting-based performance metrics dominates that of stock price-based performance metrics. For the clean (alternative) sample, 98% (98%) of the performance metrics are accounting-based. Among accounting-based performance metrics, EPS, sales, and operating income are the three most commonly used performance metrics in both the clean and alternative samples. Based on the clean sample, on average, each bonus (restricted stock and options) contract uses 1.5 (1.3 and 1) performance metrics, and each CEO receives 1.10 (1.12 and 2) grants per year. The maximum number of performance metrics used in bonus (restricted stock and options) contracts is 3 (3 and 1). These numbers are similar to those in the alternative sample.

Turning now to the ExecuComp sample, our main compensation variable is CEO total annual compensation flow (TDC1). In a small number of instances, TDC1 is less than 30% of Total_SEC (firm's total compensation disclosed in SEC filings that uses subjective assessment of the value of restricted stock and options), and we replace TDC1 by Total_SEC. This affected 31 observations in the sample.

[12] Omitting grants with relative performance metrics affects 45 observations in the clean sample (35 of which Incentive Lab has complete data on the grants), and 541 observations in the alternative sample (413 of which Incentive Lab has complete data on the grants). Omitting data on performance-vested options affects zero observations in the clean sample and only 2 observations in the alternative sample of which we do not have complete data on. Omitting data on time-vested options affects 28 observations in the clean sample and 975 observations in the alternative sample.

Table 2 provides descriptive statistics for our sample firms. The average (median) of the logarithm of CEO total annual compensation flow (TDC1) is 7.91 (7.92) close to the average (median) logarithm of simulated CEO pay of 8.15 (8.16) in the clean sample.[13] We use the logarithm of one plus total annual compensation in the empirical analysis to mitigate the effect of skewness in compensation and not be affected by units of measurement.

In robustness tests, we use the inside wealth variable in Coles et al. (2006) as an alternative compensation variable. In dynamic models (e.g., Albuquerque and Hopenhayn, 2004, and DeMarzo and Fishman, 2007) lifetime utility under the contract depends on wealth, the state variable of the system. Intuitively, future utility promises in the form of restricted stock or non-vested option that were granted in the past all contribute to incentivizing management currently, but are costly as they create volatility in the consumption stream. For example, we drop part-year executives to avoid, among other things, a confounding effect of large initial signing grants, but capture some of the incentives provided by these initial grants when we run our tests using CEO inside wealth.

[Table 2 about here.]

The mean three-year lagged stock return (assuming dividends reinvested) is 11 percent. The average firm has a market value in logarithms of 7.45, slightly higher than the median value, consistent with our sample being skewed towards larger firms. Sample firms have 56 percent of board members hired by the CEO (coopt) on average, and have 68 percent of average institutional ownership. The CEOs in our ExecuComp sample are on average of 56 years old, and stay in that role for about 8.2 years. About 10 percent of our sample CEOs are founders of their firms. The mean (median) firm stock return volatility (i.e., variance of stock returns over the last 36 months) is

[13]In the clean sample, 21.75% of firm-years have actual compensation for bonus and restricted stock equal to 0, while in our simulation only 5.61% of firm-years have simulated compensation for bonus and restricted stock equal to 0 (untabulated). The small sample may explain our inability to generate a large number of zeros. We do better in matching these zeros in the alternative sample. For the alternative sample, 3.57% of firm-years have actual compensation for bonus and restricted stock equal to 0, while 6.49% of firm-years have simulated compensation for bonus and restricted stock equal to 0 (untabulated). [These numbers need to be edited.]

0.11 (0.10). The mean (median) log CEO pay volatility (i.e., measured with realized volatility) is 13.80 (13.89) in the ExecuComp sample, whereas the simulated log CEO pay volatility (i.e., measured by the log of the variance of simulated compensation) is 14.44 (14.83) in the clean sample of the Incentive Lab data.

[Figure 1 about here.]

Figure 1 plots the cross sectional means of TDC1 and simulated pay from Incentive Lab over time for the firms in the clean sample (top left panel), and the means of bonus from ExecuComp and simulated (top right panel), the means of restricted stock grants from ExecuComp and simulated (bottom left panel), and the means of option grants from ExecuComp and simulated (bottom right corner). Overall, the simulation procedure does well in capturing the value of bonus, options, and total pay, but has difficulty in capturing the value of restricted stock as disclosed in ExecuComp. Because ExecuComp bonus is the realized value of compensation and simulated bonus is its expected value, it is natural to expect yearly deviations that wash away with a large enough sample. As with ExecuComp, we use Black-Scholes to assess the value of option grants. We differ from ExecuComp because we simulate the stock price at the end of the year, whereas ExecuComp takes the realized end-of-year price. As with bonus, therefore, we expect yearly deviations to wash away in a large sample.

Bettis et al. (2018) also document that their simulated fair values of restricted stock grants using Incentive Lab data are lower than the values reported in financial statements. To better understand the nature of this gap, we inspected several firms where the discrepancy between simulated and reported values were largest. We found several reasons for this discrepancy. First, ExecuComp sometimes reports the value of restricted stock and options all as restricted stock. This does not affect TDC1. It could affect the value of options, but the fact that we have many more observations of firms paying options than paying equity grants dilutes this effect when we compare average simulated value of option grants to average reported values. Second, we noted that sometimes companies acknowledge forfeiting the value of equity grants in DEF14A, but these end up being reported anyway in ExecuComp. Third, as Kim and Yang (2014) have noted, firms appear to adjust performance metrics to boost executive bonus compensation. While this issue does not appear to be a major concern in our simulation for bonus, it could still play a role in restricted stock. In one of our robustness tests, we use the realized value of the performance metrics as their mean in the simulation.

Panel A in Figure 2 plots the percentage of firms in the clean sample that offer any of the compensation grants. Options were more popular earlier in the sample, whereas bonus and equity grants became more popular later in the sample. These patterns in the clean sample are very similar to those displayed in the ExecuComp sample (panel B). The Execucomp sample shows a somewhat faster increase in the percentage of firms that award cash bonuses in the earl part of the sample.

[Figure 2 about here.]

Figure 3 reproduces Figure 1, but now using the alternative sample. In the alternative sample, we have firms with incomplete data in Incentive Lab on one or more compensation grants and this mostly explains the gap between simulated values and reported values. The figure indicates that this gap appears constant through the sample period. To further document the relevance of incomplete data, we plot in Figure 4 the cross sectional mean of total simulated pay when we drop, in turn, firms with incomplete data in one or more of the compensation grants. After we exclude firms for which we have incomplete data in Incentive Lab on bonus, equity and option grants, the simulated value of total pay in the alternative sample approaches the reported value, though there remains a downward bias. We believe this remaining bias could be due to a selection bias, that the firms we drop when evaluating simulated pay because of incomplete information are firms that pay more to their CEOs. It is also possible that Incentive Lab reports only data on one grant in bonus, or equity, or options, when in fact the firm awarded more than that one grant.

[Figures 3 and 4 about here.]

Figure 5 plots simulated and reported values of total compensation across industries for the clean sample and the alternative sample. As in Figures 1 and 3, there is a gap between total reported pay and total simulated pay, but there is not apparent difference in this gap across industries. Overall, our simulated data appears to capture on average some relevant patterns present in the reported data.

[Figure 5 about here.]

5 Results

5.1 Main results on risk and reward in pay

Table 3, panel A contains the results of panel regressions of the logarithm of simulated mean pay on the logarithm of simulated variance of pay using ordinary least squares. Columns 1 through 2 report the results using the clean sample and columns 5 through 8 report the results using the alternative sample. In odd numbered columns the regressions have no fixed effects and in even numbered columns the regressions have firm and year fixed effects. Standard errors are clustered by firm when no fixed effects are present and by firm and year when fixed effects are present. We also present results for the subsample of firms that do not award option grants as part of compensation.

The coefficient on the logarithm of simulated variance of pay describes the average risk and reward trade off. This coefficient is positive and statistically significant at 1% level across all specifications consistent with hypothesis (7). When fixed effects are included, the coefficient estimates are 0.091 and 0.151 in the clean sample and 0.108 and 0.137 in the alternative sample. This estimate is somewhat lower than when no fixed effects are present, which partly explains the drop in the t-statistics. Overall, the alternative sample produces slightly smaller estimates of the risk and reward trade off in pay, but the statistical significance increases dramatically in the alternative sample with the larger number of observations. The elasticity of pay to variance of pay is slightly higher when we include firms that also pay with options. We discuss this finding below.[14]

[Table 3 about here.]

[14]This finding contrasts with Hayes et al. (2012) who argue that the convexity imbedded in option pay is not used to reduce risk-related agency problems between CEOs and shareholders.

The higher estimated coefficient in the regressions without firm fixed effects in panel A of Table 3 reveals that the trade-off between variance of pay and mean pay comes both from *across firms* and *within firms.* The cross sectional finding of a large effect across firms can also be seen in panel B of Table 3, where we report cross-sectional regressions and a Fama-MacBeth estimate of the average effect.[15] This panel shows overwhelming evidence of a positive relation between mean pay and variance of pay across firms for all cross sections for both the clean and alternative samples. The magnitude of the trade off has some temporal dispersion between 0.11 and 0.43, with a mean estimated coefficient close in magnitude to the pooled regression estimates without firm fixed effects of about 0.20. One concern in studies of CEO pay over time is that significant regulatory changes may lead to structural breaks in the model, invalidating the analysis. The evidence in panel B of Table 3 significantly dispels this concern by showing a consistent positive estimate over time of the elasticity when using only cross sectional variation in the data.

Table 4 repeats the same exercise but using realized variance as a measure of the conditional volatility of pay from ExecuComp data. The dependent variable in columns 1 and 2 is the logarithm of TDC1 and in columns 3 and 4 is the logarithm of CEO wealth. Columns 1 and 3 have no fixed effects and columns 2 and 4 have firm and year fixed effects. As in Table 3, standard errors are clustered by firm when no fixed effects are present and by firm and year when fixed effects are present. In the specifications without fixed effects, the estimated elasticity describing the risk and reward trade off in pay is positive and only slightly higher than in the alternative sample using Incentive Lab data. Introducing fixed effects as shown in columns 2 and 4, however, lowers the magnitude of the risk and reward trade off to levels that are below those in Table 3. Across all specifications, the estimated coefficients using realized variance of pay are significant at 1%. In the appendix, we report on cross sectional regressions using realized variance. The results are qualitatively the same as in Table 3, panel B.

[Table 4 about here.]

[15] The standard error on the estimate is computed with the bias correction proposed in Pontiff (1996) that accounts for time series correlation of the residuals.

Table 5 presents the results with the ARCH-in-mean model. Contrary to the specifications in Tables 3 and 4, the dependent variable in Table 5 is not the logarithm of pay but the dollar value of pay. This is done so that the residual variance is the conditional variance of pay, to be consistent with the measures of variance used above, otherwise the residual variance would be the conditional variance of the logarithm of pay. The main independent variable remains the logarithm of the variance of pay. To interpret the coefficient on variance of pay as an elasticity, we divide the estimated coefficient by the mean of pay. Columns 1 and 2 report the results using TDC1 and columns 3 and 4 report the results using CEO wealth. Columns 1 and 3 have no fixed effects and columns 2 and 4 have industry and year fixed effects. We use industry fixed effects as opposed to firm fixed effects, because the ARCH estimation cannot handle the many firm-specific indicator variables.

Using TDC1, the elasticity of mean pay to variance of pay is positive and statistically significant at 1%. In column 2, that coefficient is 150, and dividing by the mean of TDC1 of $\$4,690$ (the unit of TDC1 is thousands) gives an elasticity of 3.2%. This estimate is remarkably similar to the effect uncovered using realized variance (in Table 4, with firm and year fixed effects). Using CEO wealth, the estimated elasticity is negative when no fixed effects are used and insignificant with fixed effects. The ARCH coefficients across the four specifications in Table 5 are all positive as required so that variance is positive throughout.

[Table 5 about here.]

The evidence of a low elasticity presented above is especially surprising given potential contractual features linked to option grants that push the estimates upward. First, increases in firm return volatility increase average pay through the higher value of options, and increase the variance of pay through the higher variance in the value of options. Thus, options are a natural mechanism for firms to use to induce a positive relation between conditional volatility of pay and conditional mean of pay in response to changing volatility of firm stock returns. Second, Black-Scholes option values (used in TDC1 and in our simulation of option grants) overstate the value of options for underdiversifed CEOs and understate their risk.

5.2 Economic significance of the estimated elasticity

The estimated elasticity of pay to variance of pay can be transformed into an estimate of the risk aversion coefficient. Estimating the model using the logarithms of mean pay and variance of pay is useful to reduce the skewness in these variables and to facilitate interpretation of the estimates as elasticities. However, to our knowledge, this empirical specification is not a direct representation of any preference specification, which constrains our ability to report on an implied estimate of risk aversion. The closest we can get to an estimate of risk aversion is to assume that CEOs have constant relative risk aversion preferences and that pay is log normally distributed. In that case, we can show that the expected log pay (which is what we model) equals the volatility of log pay (we instead model the log of the volatility of pay) times $(\gamma - 1)/2$, where $\gamma > 0$ is risk aversion. An elasticity of 0.2 translates to an estimate of risk aversion of 1.4. Alternatively, the implied elasticity under the these assumptions is 1 if risk aversion equals 3. To put things into perspective, asset pricing studies usually assume the coefficient of relative risk aversion to be 5 or higher, but it is possible that CEOs have lower risk aversion than the marginal investor. Indeed, Taylor (2013) estimates risk aversion to be 2.8 for CEOs, still substantially higher than our estimate.

Another way to assess the economic significance of the estimated elasticity value, is through a back-of-the-envelope calculation that uses the fact documented by Murphy and Jensen (2018) that growth in incentive pay dominated growth in pay in the last two decades. We ask what is the share of pay at risk in total pay implied by the estimated elasticity if all the variation in pay is driven by variation in incentives. Consider a contract where pay depends linearly on incentives,

$$w_t = c_0 + c_1 x_t,$$

where pay, w_t, is base salary, c_0, plus a cash bonus, $c_1 x_t$, c_1 is a parameter that describes the level of incentives and x_t is some payout based on a performance metric. The variance of pay is $V(w_t) = c_1^2 V(x_t)$ and the expected value of pay is $E(w_t) = c_0 + c_1 E(x_t)$. Consider now a change in incentives, c_1. Assuming for simplicity that the mean and variance of the performance metric do not change with the change in incentives, then[16]

$$\begin{aligned}\frac{\%\Delta E(w_t)}{\%\Delta c_1} &= \frac{c_1 E(x_t)}{E(w_t)},\\ \frac{\%\Delta V(w_t)}{\%\Delta c_1} &= 2,\end{aligned}$$

where the notation $\%\Delta y$ refers to a percentage change in the variable y. Increasing incentives by 1%, leads to an increase of 2% in the volatility of pay, but only increases mean pay by $\frac{c_1 E(x_t)}{E(w_t)}\%$. This can be summarized by taking the ratio of the two expressions above,

$$\frac{\%\Delta E(w_t)}{\%\Delta V(w_t)} = \frac{1}{2}\frac{c_1 E(x_t)}{E(w_t)}. \tag{10}$$

Since we estimate an elasticity of mean pay to variance of pay between 0.1 and 0.2 in panel A of Table 3, the implied share of pay at risk in total pay should be between 20% and 40%. Instead, the average percentage of incentive pay in total pay (i.e., $c_1 E(x_t)/E(w_t)$) in Incentive Lab contracts is 75% (untabulated). This suggests that there is too much incentive pay: CEOs appear saturated with incentives without being fully compensated for the risk that they are taking. This evidence offers a quantification to the discussion in Murphy and Jensen (2018) regarding how the growth in pay at risk in the last two decades is unrelated to provision of economic incentives and risk and reward arguments.

Continuing with the exercise, if CEO incentives were to double (i.e., c_1 doubles), then mean pay should increase by 75% (equal to $\frac{1}{2} \times 0.75 \times 200\%$), but instead and according to our estimated elasticity mean pay increases only by at most 40% (equal to $0.2 \times 200\%$). In summary, while we find a statistically strong qualitative result, the quantitative magnitude of the effect appears small to explain large differences in pay in the cross-section of U.S. CEOs.

[16] In this calculation we assume that $\%\Delta V(x_t)/\%\Delta c_1 = \%\Delta E(x_t)/\%\Delta c_1 = 0$, which have opposing, potentially neutralizing effects. Most contracting models assume $\%\Delta V(x_t)/\%\Delta c_1 = 0$. If we allow $\%\Delta E(x_t)/\%\Delta c_1 > 0$, as greater incentives induce greater effort, the puzzle to be presented deepens.

5.3 Cross-sectional variation in risk aversion

Next we consider several proxies of risk aversion and test a modification of hypothesis (7) by allowing the sensitivity of pay to variance of pay to depend on these proxies. The work of Haubrich (1994) showing that in the agency models of Grossman and Hart (1983) and Holmstrom and Milgrom (1987) equity incentives increase sharply as CEO risk aversion decreases suggests that the low elasticity may be driven by the presence of a few low risk averse CEOs.

The proxies we consider are CEOs' early-life exposure to fatal disasters (Bernile et al., 2017), whether the CEO possesses a private pilot license (Cain and McKeon, 2016), whether the CEO is a depression baby (Malmendier and Nagel, 2011), female CEO (Borghans et al., 2009), and CEO's marital status (Roussanov and Savor, 2014).

We also include in the regressions simulated skewness in pay and skewness in pay interacted with risk aversion. Hemmer, Kim and Verrecchia (2000), Ross (2004), and Chaigneau (2015) using a more general utility specification than the mean-variance utility that we used to motivate the preceding tests predict a preference for positive skewness, besides a disutility to volatility associated with risk aversion. A *prudent CEO* has a preference for positive skewness in pay if the CEO dislikes downside risk, i.e., the third partial derivative of utility with respect to pay is positive (Chaigneau, 2015). A prudent CEO requires less mean pay if awarded a contract with positively skewed payouts, say through option grants. In addition, if skewness in pay is positively related to the variance of pay, then a strong enough effect of volatility of pay on skewness can introduce a downward bias in the relation between mean and volatility of pay due to an omitted variable. The effect of skewness of pay on mean pay is not unambiguous. Agren (2006) shows that loss averse investors have a preference for negative skewness, which predicts that skewness should instead have a positive association with mean pay.[17]

The results are presented in Table 6 using the alternative sample from Incentive Lab with firm and year fixed effects. The table shows that all the estimated coefficients associated with the interaction of the risk aversion proxy and the variance of pay

[17]Interestingly, Dittmann et al (2010) show that options are still optimal with loss averse agents.

display the predicted sign. With one of the proxies for risk aversion, we do not find significance. We also perform a one-sided t-statistic and find significant estimates in all proxies for risk aversion. However, accounting for the cross sectional variation in risk aversion does not significantly increase the elasticity to those CEOs with higher risk aversion.

[Table 6 about here.]

The coefficient on skewness is estimated to be positive and statistically significant in all but one specification. A positive coefficient in skewness is consistent with a preference for negative skewness, i.e., CEOs demand higher pay if pay is positively skewed. The coefficient with the interaction of skewness and risk aversion displays the same sign as that of the interaction of variance and risk aversion, suggesting that higher risk aversion is associated with greater compensation for positive skewness in pay.

In the online appendix, we repeat this exercise using realized variance and ExecuComp data. We find that CEOs' early-life exposure to disasters with medium fatalities is significant and with a point estimate about half the size of that in Table 6. The point estimates on the coefficients on the other measures of risk aversion interacted with realized volatility carry the predicted sign (except for when risk aversion is proxied by the gender of the CEO) but are lower, which partly explains the statistically insignificant results.

5.4 Alternative hypotheses

In this section, we consider several extensions of the basic Grossman and Hart model that could explain the finding of a low elasticity of mean pay to variance of pay. Following Oyer (2004), Laffont and Martimort (2002) and Nickerson (2017), CEO's outside opportunities may be type dependent, correlating positively with CEO pay. If the CEO's outside opportunities happen to covary negatively with the variance in pay, then mean and variance of pay may be negatively related causing a downward bias in the elasticity due to an omitted variable problem. Following Himmelberg and

Hubbard (2000), and Oyer (2004), we use lagged values of stock performance of the firm's industry, and lagged own stock performance, respectively, to describe outside opportunities.[18]

An entrenched CEO may be able to guarantee an expected utility under the optimal contract that is above her reservation utility, generating slack in the participation constraint that depends on the level of CEO entrenchment. If greater slack comes with high volatility of pay, and this effect is strong enough, then mean and variance of pay would be positively related in ways that do not reflect the risk and reward trade off that we investigate. For entrenchment proxies, we use Coles et al. (2014) co-opted board measure, as well as the percentage ownership by institutional investors.

If the utility function is nonseparable in consumption and effort, then incentives may be provided by reducing the marginal cost of effort. For example, higher CEO pay creates status enjoyed by the CEO that reduces the cost of effort for the CEO. In such cases, the participation constraint is not binding and the risk and reward trade-off in pay may not be directly implied by the participation constraint (Laffont and Martimort, 2002). For cost of effort proxies, we use CEO age and log of CEO tenure, the volatility of stock returns, a dummy for when the CEO is the founder, and the lagged value of the firm's market capitalization.

We control for CEO overconfidence. There is evidence that CEOs overestimate the performance of their investments while underestimating the risks (e.g., Dittrich et al., 2005, Huang and Kisgen, 2013, Kolasinski and Li, 2013, and Malmendier and Tate, 2005 and 2008). This overconfidence can be used by the shareholders to save on the costs of incentive provision by offering contracts that are incentive-intensive (Gervais, Heaton and Odean, 2011). We use the Humphery et al. (2014) overconfidence indicator that is based on whether the CEO keeps deep-in-the money options that have vested.

[Table 7 about here.]

Laffont and Martimort (2002) show that in moral hazard models if the agent is

[18] In untabulated results, we control outside opportunities using industry times year fixed effects and find similar results to those shown above.

risk neutral and there is a limited liability constraint that sets a lower bound to pay, then this new constraint together with the incentive compatibility constraint imply the participation constraint. That is, if the limited liability constraint is binding, then the principal is constrained in her ability to induce effort and the participation constraint may be slack. In this case the risk-return trade-off breaks down. This explanation is less realistic unless one wants to dismiss the long standing assumption of risk averse CEOs. Accordingly, we do not explore it further.

Table 7 presents the results of testing the null hypothesis in (7), controlling for these various alternative hypotheses, using OLS on the Incentive Lab alternative sample.[19] The main result to note from the table is that the inclusion of the various controls does not significantly change the small economic magnitude of the sensitivity of mean pay to variance of pay. Focusing on column 7, high skewness in pay is associated with higher pay consistent with loss aversion utility (Agren, 2006). A positive coefficient associated with skewness is also consistent with the view that firms use stock options to provide excess pay to CEOs (Bebchuk and Fried, 2004, made this argument for when options were not expensed). CEO outside opportunities do not appear to affect mean pay nor do the entrenchment proxies. CEO age, and firm size both lead to higher pay, whereas overconfidence is associated with lower pay.

5.5 Robustness analysis

Table 8 repeats the analysis with controls when using realized variance and ARCH volatility. In column 1, the elasticity of pay to variance of pay is halved relative to the model without controls but remains significant at the 1% level. Consistent with the results in Table 7, firm size enters positively, but differently from Table 7 overconfidence is positively related to pay and CEO age is negatively related to pay. In column 2, we replace realized variance by simulated variance and find that the coefficient on variance of pay is remarkably stable. Finally, in column 3, we again

[19] In addition to the control variables shown in the paper, we have also used IPO activity, average industry ROA, and median of peer pay as proxies for outside opportunities, a concentration index of institutional holdings for governance, and a concentration index for business segments and book-to-market value for job complexity. These variables are not significant and were dropped from the estimation. The online appendix reports results with CEO wealth.

use realized variance, but match the sample to firms for which we have Incentive Lab data. The coefficient on the elasticity decreases in magnitude and is not statistically significant, but because the number of observations also decreases by about one third, it is difficult to draw conclusions regarding any possible selection bias.

The last two columns repeat the exercise in column 7 of Table 7 with the ARCH-in-mean model for the whole sample of ExecuComp firms (column 4) and for a matched sample of Incentive Lab firms (column 5). The estimated elasticity is either unchanged (column 4) or slightly higher (column 5). There is no evidence of selection bias of the Incentive Lab sample when using the ARCH-in-mean model.

For columns 1, 3, 4, and 5, that use realized volatility or ARCH volatility and ExecuComp data, realized skewness enters positively, consistent with Taylor (2013) who finds that firms pay less on average to CEOs with contracts that offer downward rigidity in pay.

The online appendix adds a proxy for risk of turnover to these regressions. We find that the prospect of forced turnover is associated with less pay to the CEO, contrasting with evidence in Peters and Wagner (2014). We also find that the estimated elasticity of pay to variance of pay is virtually unchanged. The reason for the similarity in results with regards to the elasticity is the small unconditional probability of CEO forced turnover in the US (Peters and Wagner, 2014, show that it is just under 3%).

We also report in the online appendix the results from simulating mean pay and the variance of pay using time-t variables as the mean in the conditional distribution of the performance metrics. Using time-t variables presumes that the board had perfect foresight when designing the contracts, which is unrealistic but constitutes an upper bound on the boards information. The results usng the simulation in the paper constitute a lower bound on the boards information as they rel solely on past information. Overall, the results on the elasticty of pay to variance of pay are largely unchanged.

In summary, although we cannot fully rule out the possibility that the low compensation for risk in the cross section of U.S. CEOs is not subject to an omitted variable bias, we are comforted by the result that our conclusion does not change after inclusion of a long list of controls in the regressions and of several robustness tests.

6 Conclusion

This paper examines a robust implication of a model widely used to describe CEO pay setting. The mainstream agency model of Grossman and Hart (1983) has a fundamental prediction imbedded in the model's participation constraint that says that mean pay and the volatility of pay are positively related. Using a variety of methods to test this prediction and a variety of datasets for U.S. CEOs, including painstaking simulations using Incentive Lab CEO-contract data, we find a positive association between mean pay and variance of pay. However, the compensation for risk in pay appears small to be fully consistent with model. Using a back-of-the-envelope calculation, we find that CEOs' pay packages appear saturated with incentive pay. Allowing for the elasticity of pay to variance of pay to vary with proxies for risk aversion, increases the elasticity for more risk averse CEOs but not in an economically significantly manner. We discuss possible reasons for this finding including the substantial growth in pay at risk in the last two decades in the US.

There are other model departures that are worth studying in future research within the context of optimal contracting models. One is to introduce dynamic considerations. For example, models of career concerns (Axelson and Baliga, 2009, Gibbons and Murphy, 1992, or Noe and Rebello, 2012) will have an implicit trade off between mean and volatility of consumption streams as opposed to current pay. Another is agent heterogeneity. Models with adverse selection in agent type or models of assortative matching (e.g., Tervio, 2008, and Edmans, Gabaix and Landier, 2009) generally predict that the participation constraint holds only for one agent. If CEO compensation data come from a cross section where the participation constraint holds only for a few of the CEOs, then there is a downward bias in the risk and reward trade off.

7 Appendix

This appendix further details on the simulation exercise using Incentive Lab data. It also contains a table with the definitions of variables used in the paper.

7.1 Simulation using the Incentive Lab data

We use information available at the beginning of the year to simulate expected pay and variance of pay for the current year. We take two inputs for the simulation: (i) compensation contract information from Incentive Lab, which describes the relationship between the chosen performance metric (metrics) and the corresponding performance-based compensation (i.e., cash bonus or equity grants), and (ii) actual performance in the past five years from Compustat, the mean and variance (covariance) of which are used to estimate simulated performance for the current year. We then fit the simulated performance from (ii) to the compensation contracts estimated in (i) to generate simulated pay. Since we simulate 10,000 times for each firm-year-CEO-grant, we can calculate the expected pay and variance of pay from the 10,000 simulations. We describe details of the procedures below.

Compensation contract fitting. We estimate the compensation contracts using the Incentive Lab data. Incentive Lab data provides information on: (i) the performance metrics used in the compensation contracts (variable name: "metric"), (ii) the threshold, target, and maximum level of each performance metric (variable names: "goalThreshold", "goalTarget", and "goalMax"), and (iii) the threshold, target, and maximum level of the compensation (variable names: "nonEquityThreshold", "nonEquityTarget", and "nonEquityMax" for bonus, and "equityThreshold", "equityTarget", and "equityMax" for equity grants).

When actual firm performance is below the threshold of the performance metric indicated in the contract, the CEO does not earn any performance-based compensation; when actual firm performance equals the target performance metric indicated in the contract, the CEO earns the target amount of performance-based compensation; when actual firm performance is above the maximum of the performance metric indicated in the contract, the CEO earns the maximum amount of performance-based compensation; when actual performance falls between the threshold and the maximum of the performance metric indicated in the contract, the CEO earns performance-based compensation in the amount between the target and the maximum of the performance-based compensation indicated in the contract.

For firms with no missing values of the contracts details (i.e., threshold, target, and maximum for the performance metric and the performance-based compensation), we can fit the compensation contracts using either (i) piece-wise linear estimation (i.e., two linear slopes: one between the threshold and the target, the other between the target and the maximum), or (ii) quadratic estimation. For firms with missing values of the target performance metric or target compensation, we have to fit the compensation contracts using the linear estimation (i.e., one linear slope between the threshold and the maximum). We drop firms with missing values of the threshold or the maximum, because the missing values make it impossible to estimate the contracts.

We implement the contract estimations in four steps. In the first step, we construct a sample of compensation contracts that meets the following three criteria: (i) using absolute performance metrics only, (ii) including cash and equity compensation contracts only, and (iii) including contracts for CEOs only. In particular, we start with the Absolute Performance Goals Data (referred to as "GpbaAbs" by Incentive Lab) to get all compensation contracts using absolute performance metrics. We drop firms that also use relative performance metrics (in addition to absolute performance metrics, i.e., the variable "numRelative" has a positive value). We then limit the sample to cash and equity compensation contracts by merging the Grants of Plan-Based Awards Table (referred to as "GpbaGrant"): we keep contracts where the "AwardType" variable in GpbaGrant has the value of "cashShort", "cashLong", "Option", or "rsu". Next, we limit the sample to include contracts for CEOs only by merging the Participant Data by Fiscal Year (referred to as "ParticipantFY"): we keep contracts where the "currentCEO" variable in ParticipantFY has the value of one.

In the second step, we classify the compensation contracts into three groups: (i) firm-years with bonus contracts only, (ii) firm-years with restricted stock contracts only, and (iii) firm-years with both bonus and restricted stock contracts. We separately examine the three groups because we need to ensure the contract details are available for simultaneously simulating all performance metrics for a given firm-year-CEO. Some firms have actual restricted stock grants without disclosing the restricted stock contract details (either not listed in GpbaAbs or showing missing values of the contract details in GpbaAbs); these firms are dropped in our main sample (i.e., the clean sample) for the "bonus only" firms; we keep these firms in an alternative sample (i.e., the alternative sample) in robustness check. Similarly, when constructing the clean sample for the "restricted stock only" firms, we drop firms that have actual bonus compensation without disclosing the bonus contract details. When constructing the clean sample for the "bonus and restricted stock" firms, (i) we first combine all bonus contracts with all restricted stock contracts, (ii) we then drop firms in the "bonus only" sample and the "restricted stock only" sample, and (iii) next we keep firms with both bonus contracts and restricted stock contracts available in the same year.

In the third step, we pinpoint the specific performance metrics used in each contract. In particular, "GpbaAbs" has five relevant variables for this task: (i) the variable "metric" lists the name of the performance metric, (ii) the indicator variable "metricIsPerShare" describes whether the performance metric is scaled by the number of common stocks; (iii) the indicator variable "metricIsMargin" describes whether the performance metric is scaled by sales; (iv) the indicator variable "metricIsGrowth" describes whether the performance metric is measured as the growth rate; and (v) the variable "metricOther" provides additional textual information about the performance metric. For example, when "metric" has the value of "Cashflow", several possibilities exist: if all three indicator variables equal zero, it means the performance metric used in the contract is the dollar amount of cash flow; if "metricIsPerShare" equals one, "metricIsMargin" equals zero, and "metricIsGrowth" equals one, it means the performance metric used in the contract is the growth rate of cash flow per share. In addition, the textual description in "metricOther" may indicate it is operating cash flow or free cash flow rather then net cash flow. We consider all possible combinations of the indicator variables as well as the additional information in the textual description from "metricOther" to pinpoint the performance metric used in each compensation contract.

In the fourth step, we fit the contract using linear or quadratic estimation. Specifically, for firms with no missing values for the contract details, i.e., firms with all three pairs of data points available: the threshold x1 and y1, the target x2 and y2, and the maximum x3 and y3 (x refers to performance and y refers to compensation), we use both methods to fit the same contract: piece-wise linear and quadratic. For firms with missing values for the contract target, i.e., firms with only two pairs of data points available: the threshold x1 and y1, and the maximum x2 and y2, we use the linear method to fit the contract. Once the contracts are estimated, we can then apply the simulated performance to get simulated compensation. We present results from the linear estimation in the paper, and results from the quadratic estimation in the Online Appendix.

Performance Simulation. We simulate current year performance using actual performance in the past five years from Compustat. The Incentive Lab contract information is presented at the firm-year-grant-metric level. It is possible for firms to use more than one performance metric for a given grant (contract). It is also possible for firms to set up several grants (contracts) for the same CEO in a given year. We consider all metrics used for a given firm-year-CEO and simultaneously simulate all metrics for that year. In particular, for each CEO and year, we assume a multivariate normal distribution for all performance metrics used for a given CEO across all contracts; we set the mean of the joint normal distribution equal to the actual mean in the past five years (i.e., year t-5 to t-1), and set the covariance matrix for the joint normal distribution equal to the covariance matrix calculated from the actual values of the performance metrics in the past five years. Using these assumptions, we run 10,000 simulations for each firm-year-grant-metric, which provides simulated performance for estimating simulated compensation.

In our main test, we convert the performance metrics stated in dollar amount into scaled variables to make the covariance matrix comparable with other scaled metrics (i.e., metrics expressed as a rate or ratio such as growth rate, margin, per share value, ROA, etc.). In particular, when the performance metric is the dollar amount of sales, we simulate the firm's sales growth rate, and get the dollar amount of simulated sales as simulated sales_t $=\text{sales}_{t-1}\times$(1+simulated sales growth rate_t); when the performance metric is operating income, profits before tax, net income, cash flow, etc., which can have negative values in the past five years, we simulate the corresponding performance scaled by lagged total assets, and get the dollar amount of the simulated performance as: simulated performance_t =total $\text{assets}_{t-1}\times$simulated scaled performance_t.

In a robustness check, we do not use the conversion, and simulate the dollar amount of performance metrics directly. We obtain similar results whether the simulated performance is scaled or not.

While bonus contracts are written on the dollar value of cash payment, equity grants are written on the number of shares granted. Thus we need a price estimate to convert the simulated number of shares granted to the dollar value. Because price is related to the accounting performance, we avoid simulating price directly; instead, we simulate the price to lagged sales ratio to get simulated price. In particular, simulated price_t $=\text{sales}_{t-1}\times$simulated price to lagged sales ratio_t. Because of the price estimates, for all CEOs with restricted stock grants (restricted stock contracts only or both bonus and restricted stock contracts), the covariance matrix for the joint normal distribution includes the price to lagged sales ratio as an additional input variable in addition to the actual performance metrics used in the compensation contracts.

Compensation Simulation. We calculate simulated compensation by fitting the simulated performance to the estimated compensation contracts. Since the simulated performance is conducted at the firm-year-CEO-grant-metric level, we first calculate the simulated compensation at the firm-year-CEO-grant-metric level. We then collapse the metric level compensation into the grant level compensation based on information in the variable "performanceGrouping", which describes the relationship between the various performance metrics.

The compensation contracts can be described in two overall patterns: (i) separable contracts, and (ii) non-separable contracts. While separable contracts allow CEOs to earn part of the bonus (or equity grant) when some of the performance metrics are not met, non-separable contracts result in zero bonus (or equity grant) if any of the performance metric is not met.

Incentive Lab indicates that the performance metrics in the separable contracts are equal weighted. So take the example of a separable contract with three performance metrics, each metric is worth one third of the total compensation indicated in that contract. As a result, we assign the weight of one third to each simulated pay at the metric level, and add the weighted pay from all three metrics to get total simulated pay at the grant level. For CEOs with more than one grant in a given year, we add simulated pay from all grants for a given CEO. As explained before, if a contract is separable, it is possible for a CEO to miss some performance metrics and still earn some performance-based compensation.

For non-separable contracts, we impose an additional requirement for consolidating the metric level simulated pay to the grant level simulated pay: if any of the simulated performance metric does not meet the goal threshold set in the contract, then the total grant level simulated pay is zero.

Once we have 10,000 simulated pay at the firm-year-CEO level, we can calculate the mean, variance, and skewness of the simulated pay from the 10,000 simulated results for each firm-year-CEO. Since we have excluded firms with other performance-based compensation from our sample, expected total pay for the current year using information available at the beginning of the year equals salary plus mean simulated pay from the procedures described above. Since salary is constant for a given year, expected variance of total pay equals variance of simulated pay, and skewness of total pay equals skewness of simulated pay.

7.2 Variable definitions

[Table with variable definitions here.]

References

[1] Adams, R., Almeida, H., and Ferreira, D., 2005, Powerful CEOs and Their Impact on Corporate Performance, *The Review of Financial Studies* 18, 1403-1432.

[2] Aggarwal, R. K., and Samwick, A. A., 1999, The Other Side of the Trade-Off: The Impact of Risk on Executive Compensation, *Journal of Political Economy* 107, 65-105.

[3] Agren, M., 2006, Prospect Theory and Higher Moments, working paper Uppsala University.

[4] Albuquerque, R., and Hopenhayn, H., 2004, Optimal Lending Contracts and Firm Dynamics, *The Review of Economic Studies* 71, 285-315.

[5] Albuquerque, R., 2012, Skewness in Stock Returns: Reconciling the Evidence on Firm Versus Aggregate Returns, *The Review of Financial Studies* 25, 1630-1673.

[6] Allgood, S., and Farrell, K. A., 2003, The Match between CEO and Firm, *Journal of Business* 76, 317-341.

[7] Axelson, U., and Baliga, S., 2009, Liquidity and Manipulation of Executive Compensation Schemes, *Review of Financial Studies* 22, 3907–3939.

[8] Barndorff-Nielsen, O. E., and Shephard, N., 2002a, Estimating Quadratic Variation Using Realized Variance, *Journal of Applied Econometrics* 17 (special Issue), 457-477.

[9] Barndorff-Nielsen, O. E., and Shephard, N., 2002b, Econometric analysis of realized volatility and its use in estimated stochastic volatility models, *Journal of the Royal Statistical Society*, Series B 64, 253-280.

[10] Bebchuk, L. A., and Fried, J. M., 2003, Executive Compensation as an Agency Problem, *Journal of Economic Perspectives* 17, 71–92.

[11] Bebchuk, L. A., and Fried, J. M., 2004, Pay Without Performance: The Unfulfilled Promise of Executive Compensation. Harvard University Press, Cambridge.

[12] Bebchuk, L. A., and Grinstein, Y., 2005, The Growth of Executive Pay, *Oxford Review of Economic Policy* 21, 283-303.

[13] Bertrand, M., and Mullainathan, S., 2001, Are CEOs Rewarded for Luck? The Ones Without Principals Are, *The Quarterly Journal of Economics* 116, 901-932.

[14] Bettis, J. C., Bizjak, J., Coles, J., and Kalpathy, S., 2018, Performance-Vesting Provisions in Executive Compensation, Working paper Arizona State University.

[15] Bizjak, J., Kalpathy, S., Li, Z., and Young, B., 2018, The Role of Peer Firm Selection in Explicit Relative Performance Awards, available at SSRN: https://ssrn.com/abstract=2833309.

[16] Bollerslev, T., Engle, R. F., and Wooldridge, J. M., 1988, A Capital Asset Pricing Model with Time-Varying covariances, *Journal of Political Economy* 96, 116-131.

[17] Cai, J., and Vijh, A. M., 2005, Executive Stock and Option Valuation in a Two-state Variable Framework, *Journal of Derivatives* 12, 9–27.

[18] Chaigneau, P., 2015, Risk Aversion, Prudence, and Compensation, *European Journal of Finance* 21, 1357-1373.

[19] Cheng, Ing-Haw, Hong, H., and Scheinkman, J. A., 2015, Yesterday's Heroes: Compensation and Risk at Financial Firms, *Journal of Finance* 70, 839–879.

[20] Choi, I., 2001, Unit root tests for panel data, *Journal of International Money and Finance* 20, 249–272.

[21] Coles, J. L., Daniel, N. D., and Naveen, L., 2006, Managerial Incentives and Risk-Taking, *Journal of Financial Economics* 79, 431-468.

[22] Coles, J. L., Daniel, N. D., and Naveen, L., 2014, Co-Opted Boards, *The Review of Financial Studies* 27, 1751–1796.

[23] Conyon, M., Core, J., and Guay, W., 2011, Are U.S. CEOs Paid More than U.K. CEOs? Inferences from Risk-Adjusted Pay, *Review of Financial Studies* 24, 402–38.

[24] Cuñat, V., and Guadalupe, M., 2009, Executive Compensation and Competition in the Banking and Financial Sectors, *Journal of Banking and Finance* 33, 439-474.

[25] DeMarzo, P. M., and Fishman, M. J., 2007, Agency and Optimal Investment Dynamics, *Review of Financial Studies* 20, 151–188.

[26] Dittmann, I., and Maug, E., 2007, Lower Salaries and no Options? On the Optimal Structure of Executive Pay, *Journal of Finance* 62, 303–343.

[27] Dittmann, I. , Maug, E., and Spalt, O., 2010, Sticks or Carrots? Optimal CEO Compensation when Managers Are Loss Averse, *The Journal of Finance* 65: 2015-2050.

[28] Dittrich, D.A.V., Guth, W., and Maciejovsky, B., 2005, Overconfidence in Investment Decisions: An Experimental Approach, *European Journal of Finance* 11, 471–491.

[29] Edmans, A., Gabaix, X., and Landier, A., 2009, A Multiplicative Model of Optimal CEO Incentives in Market Equilibrium, Review of Financial Studies 22, 4881-4917.

[30] Engle, R. F., 1982, Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica* 50, 987-1007.

[31] Fernandes, N., Ferreira, M. A., Matos, P., Murphy, K. J., 2013, Are U.S. CEOs Paid More? New International Evidence, *Review of Financial Studies* 26, 323–367.

[32] Gervais, S., Heaton, J.B., Odean, T., 2011, Overconfidence, Compensation Contracts, and Capital Budgeting, *Journal of Finance* 66, 1735–1777.

[33] Gibbons, R., and Murphy, K. J., 1992, Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence, *Journal of Political Economy* 100, 468-505.

[34] Grossman, S. J., and Hart, O. D., 1983, An Analysis of the Principal- Agent Problem, *Econometrica* 51, 7-45.

[35] Hall, B., and Murphy, K. J., 2002, Stock Options for Undiversified Executives, *Journal of Accounting and Economics* 33, 3–42.

[36] Haubrich, J. G., 1994, Risk Aversion, Performance Pay, and the Principal-Agent Problem, *Journal of Political Economy* 102, 258–276.

[37] Hayes, R. M., Lemmon, M., and Qiu, M., 2012, Stock Options and Managerial Incentives for Risk Taking: Evidence from FAS 123R, *Journal of Financial Economics* 105, 174-190.

[38] Hemmer, T., Kim, O., and Verrecchia, R. E., 2000, Introducing Convexity into Optimal Compensation Contracts, *Journal of Accounting and Economics* 28, 307–327.

[39] Himmelberg, C. P., and Hubbard, R. G., 2000, Incentive Pay and the Market for CEOs: An Analysis of Pay-For-Performance Sensitivity, Working paper Columbia Business School.

[40] Holden, C. W., and Kim, D. S., 2017, Performance Share Plans: Valuation and Empirical Tests, *Journal of Corporate Finance* 44, 99-125.

[41] Holmstrom, B., and Milgrom, P., 1987, Aggregation and Linearity in the Provision of Intertemporal Incentives, *Econometrica* 55, 303-28.

[42] Huang, J., and Kisgen, D., 2013, Gender and Corporate Finance: Are Male Executives Overconfident Relative to Female Executives?, *Journal of Financial Economics* 108, 822-839.

[43] Humphery-Jenner, M., Lisic, L. L., Nanda, V., and Silveri, S., 2016, Executive Overconfidence and Compensation Structure, *Journal of Financial Economics* 119, 533-558.

[44] Jensen, M. C., and Murphy, K. J., 1990, Performance Pay and Top-Management Incentives, *Journal of Political Economy* 98, 225-264.

[45] Kim, D. S., and Yang,J., 2014, Beating the Target: Performance Management around The Annual Incentive Target, working paper Indiana University.

[46] Kolasinski, A.C., Li, X., 2013, Can Strong Boards and Trading Their Own Firm's Stock Help CEOs Make Better Decisions? Evidence from Acquisitions by Overconfident CEOs, *Journal of Financial and Quantitative Analysis* 48, 1173–1206.

[47] Laffont, J.-J., and Martimort, D., 2002, The Theory of Incentives. Princeton University Press (Princeton, NJ).

[48] Lambert, R. A., Larcker, D. F. , and Verrecchia, R. E. , 1991, Portfolio Considerations in Valuing Executive Compensation, *Journal of Accounting Research* 29, 129-149.

[49] Malmendier, U., Tate, G., 2005, CEO Overconfidence and Corporate Investment, *Journal of Financial Economics* 60, 2661–2700.

[50] Malmendier, U., Tate, G., 2008, Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction, *Journal of Financial Economics* 89, 20–43.

[51] Murphy, K. J., and Jensen, M. C., 2018, The Politics of Pay: The Unintended Consequences of Regulating Executive Compensation, Working paper Harvard Business School.

[52] Nickerson, J., 2017, Market Forces and CEO Pay: Shocks to CEO Demand Induced by IPO Waves, Working paper Boston College.

[53] Noe, T. H., and Rebello, M. J., 2012, Optimal Corporate Governance and Compensation in a Dynamic World, *Review of Financial Studies* 25, 480–521.

[54] Oyer, P., 2004, Why do Firms Use Incentives that Have No Incentive Effects?, *Journal of Finance* 59, 1619–1650.

[55] Peters, F. S., and Wagner, A. F., 2014, The Executive Turnover Risk Premium, *Journal of Finance* 69, 1529-1563.

[56] Pontiff, J., 1996, Costly Arbitrage: Evidence from Closed-End Funds, *Quarterly Journal of Economics* 111, 1135-1151.

[57] Prendergast, C., 2002, The Tenuous Trade-Off Between Risk and Incentives, *Journal of Political Economy* 110, 1071–1102.

[58] Raith, M., 2003, Competition, Risk, and Managerial Incentives, *The American Economic Review* 93, 1425-1436.

[59] Ramey, G., and Ramey, V. A., 1995, Cross-Country Evidence on the Link Between Volatility and Growth, *The American Economic Review* 85, 1138-1151.

[60] Rose, N. L., and Wolfram, C., 2002, Regulating Executive Pay: Using the Tax Code to Influence Chief Executive Officer Compensation, *Journal of Labor and Economics* 20, S138-S174.

[61] Ross, S., 2004, Compensation, Incentives, and the Duality of Risk Aversion and Riskiness, *Journal of Finance* 59, 207-225.

[62] Shue, K., and Townsend, R., 2015, Growth through Rigidity: An Explanation for the Rise in CEO Pay, Working paper University of Chicago.

[63] Taylor, L. A., 2013, CEO Wage Dynamics: Estimates from a Learning Model, *Journal of Financial Economics* 108, 79-98.

[64] Tervio, M., 2008, The Difference That CEOs Make: An Assignment Model Approach, *The American Economic Review* 98, 642-668.

[65] Whitehead, A., 2002, Meta-Analysis of Controlled Clinical Trials. Wiley (Chichester, U.K.).

APPENDIX
Variable Definitions

Total compensation (TDC1)	=	Total annual compensation flow is calculated as the sum of salary, bonus, other annual compensation (e.g., gross-ups for tax liabilities, perquisites, preferential discounts on stock purchases), long-term incentive payouts, restricted stocks granted during the year (determined as market value of the date of the grant), the value of stock options granted (estimated using the Black-Scholes formula or total grant-date present value of options awarded when Black-Scholes is not available), and all other compensation (e.g., payouts for cancellation of stock options, 401K contributions, signing bonuses, tax reimbursements) before 2006. After 2006, annual compensation flow is calculated as the sum of salary, bonus, non-equity incentive plan compensation, the grant-date fair value of option awards, the grant-date fair value of stock awards, and other compensation.
Log of TDC1	=	The natural logarithm of total compensation (TDC1).
CEO Inside Wealth	=	Value of the CEO's stock and option portfolio (in $000s) from Coles, et al. (2006) plus salary, bonus, and other annual compensation (othcomp) before 2006; or value of the CEO's stock and option portfolio plus salary, bonus, non-equity incentive plan compensation, and other compensation after 2006.
Simulated mean pay	=	Salary plus the mean value of the total of simulated bonus and simulated restricted stock from 10,000 simulations for each firm-year-CEO.
3-Year Stock Return	=	The 3-year total return to shareholders, including the monthly reinvestment of dividends.
Firm Volatility	=	The standard deviation of monthly stock returns calculated over months $t-37$ to $t-1$.
Average Industry Return	=	Average of all the firms' annual stock return that are in the same industry, defined as firms in the same 4-digit Global Industry Classification System (GICS).
Log of Market Capitalization	=	Logarithm of the market capitalization, calculated as number of shares outstanding multiplied by the firm's stock price at the end of fiscal year.
Overconfidence	=	Indicator variable equal to one if the CEO has held options for at least two years in a row that are deep in the money, where deep in the money is defined as when the average value per option is at least 67% of the option strike price, zero otherwise.
Co-Opted Board (Coopt)	=	Indicator variable equal to one if the number of directors hired after the CEO took office is above the sample mean, zero otherwise.
CEO Age (Age)	=	The age of the CEO while in office.
Log of CEO Tenure	=	Natural logarithm of the number of years the CEO has been in office at the firm.
CEO is Founder (Founder)	=	Indicator variable equal to one if the title CEO indicates that the CEO is also the founder of the firm, zero otherwise.

Institutional Holdings Percent	=	Percentage of the firm's shares outstanding that are owned by all institutional investors. This is obtained by Thomson Reuters Institutional (13f) Holdings – Stock Ownership (variable "instown_perc").
Realized Variance of CEO Pay	=	The natural logarithm of the variance of CEO total pay flow (TDC1) calculated over the last 5 years ($t-5$ to t).
Realized Variance of CEO Wealth		The natural logarithm of the variance of CEO inside wealth calculated over the last 5 years ($t-5$ to t).
Simulated Variance of CEO Pay	=	The natural logarithm of the variance of simulated bonus plus simulated restricted stock from the 10,000 simulations for each firm-year-CEO.
Realized Skewness of CEO Pay	=	Skewness of CEO total pay flow (TDC1) calculated over the last 5 years ($t-5$ to t).
Realized Skewness of CEO Wealth	=	Skewness of CEO inside wealth calculated over the last 5 years ($t-5$ to t).

Figure 1. Clean Sample Simulated Pay Versus Actual Pay

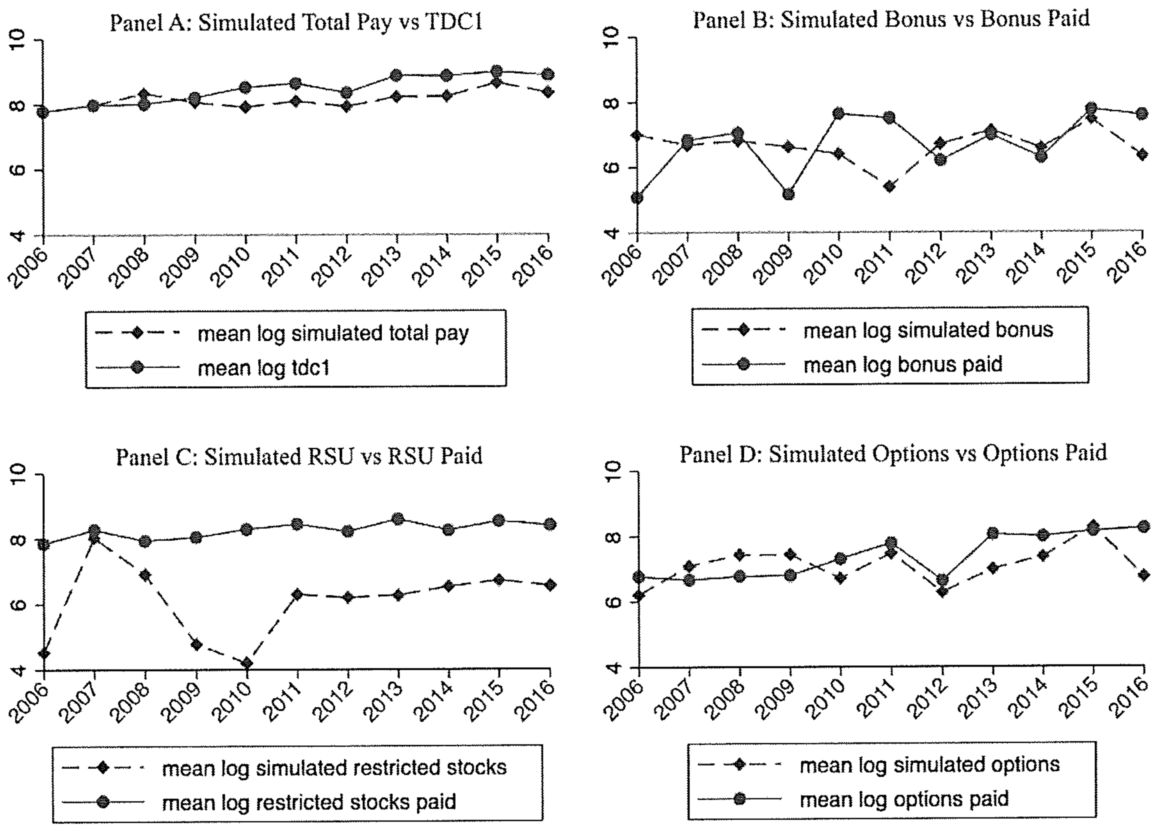

Figure 2. Clean Sample Grant Frequency Incentive Lab Versus ExecuComp

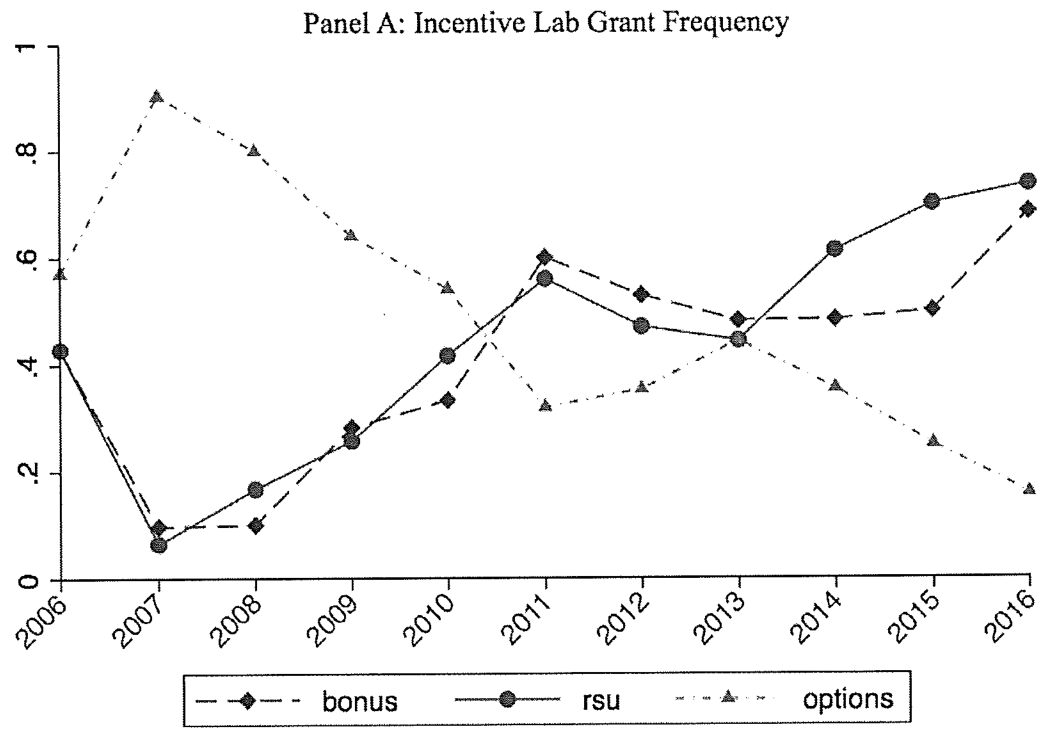

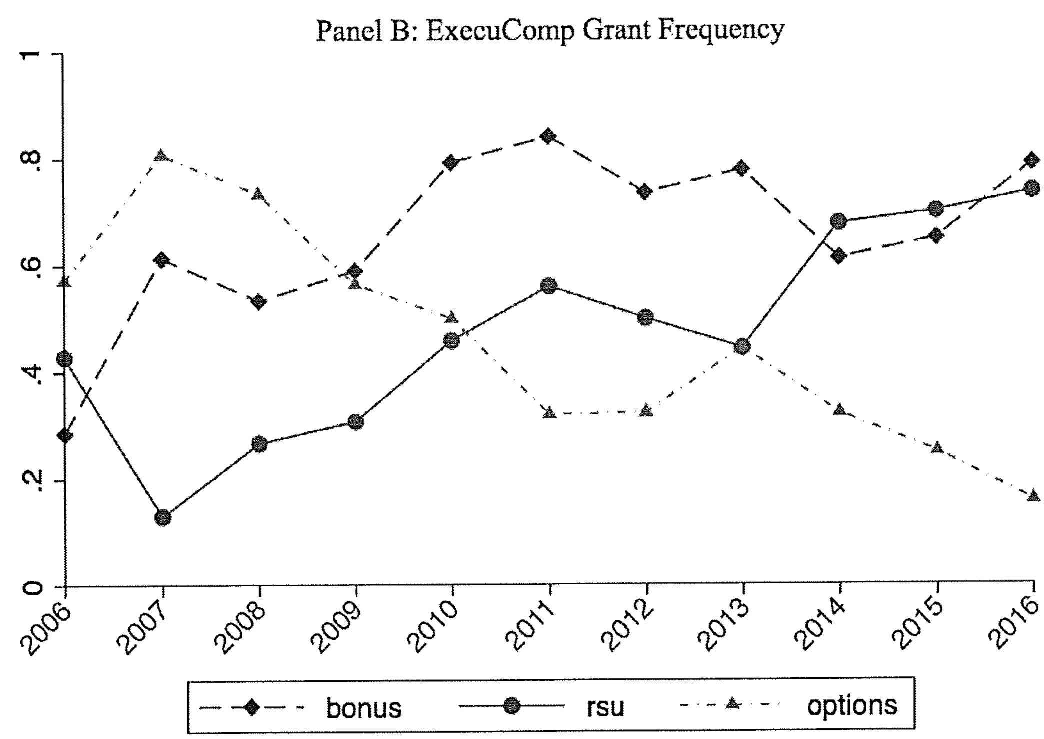

Figure 3. Alternative Sample Simulated Pay Versus Actual Pay

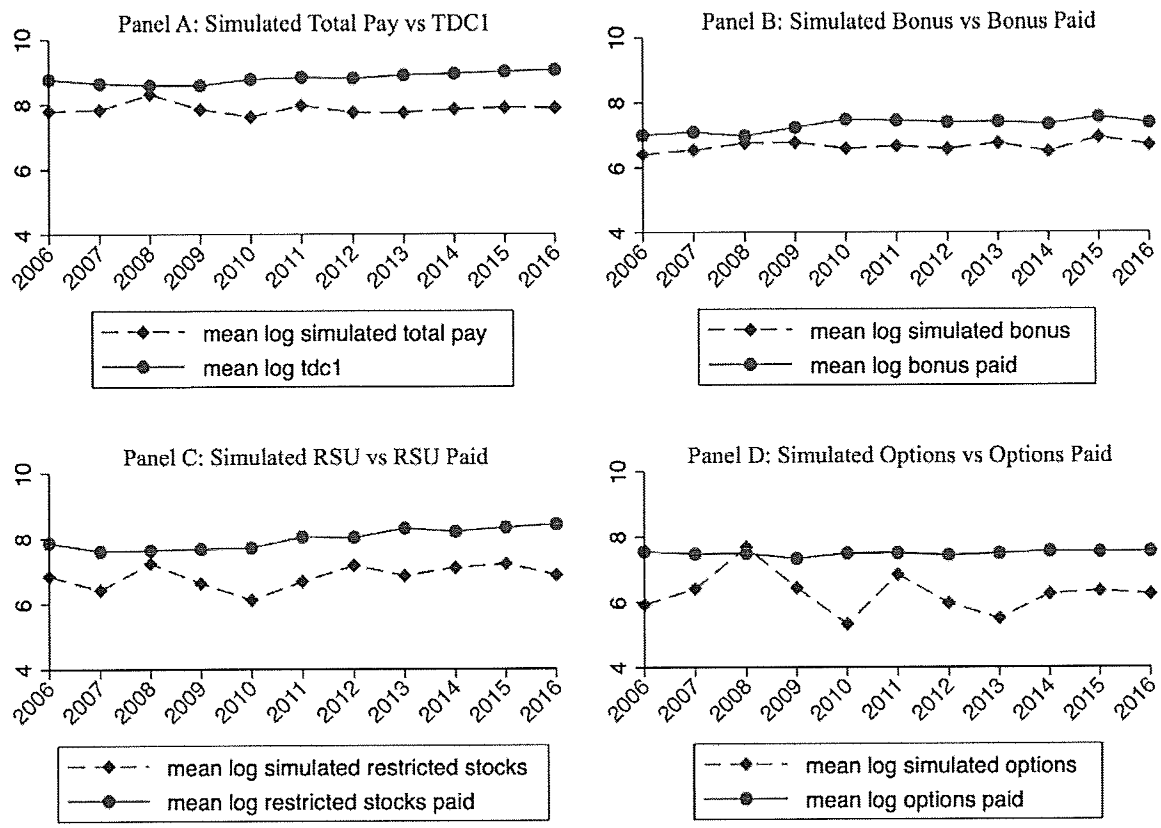

Figure 4. Alternative Sample Simulated Pay Versus Actual Pay with Exclusions

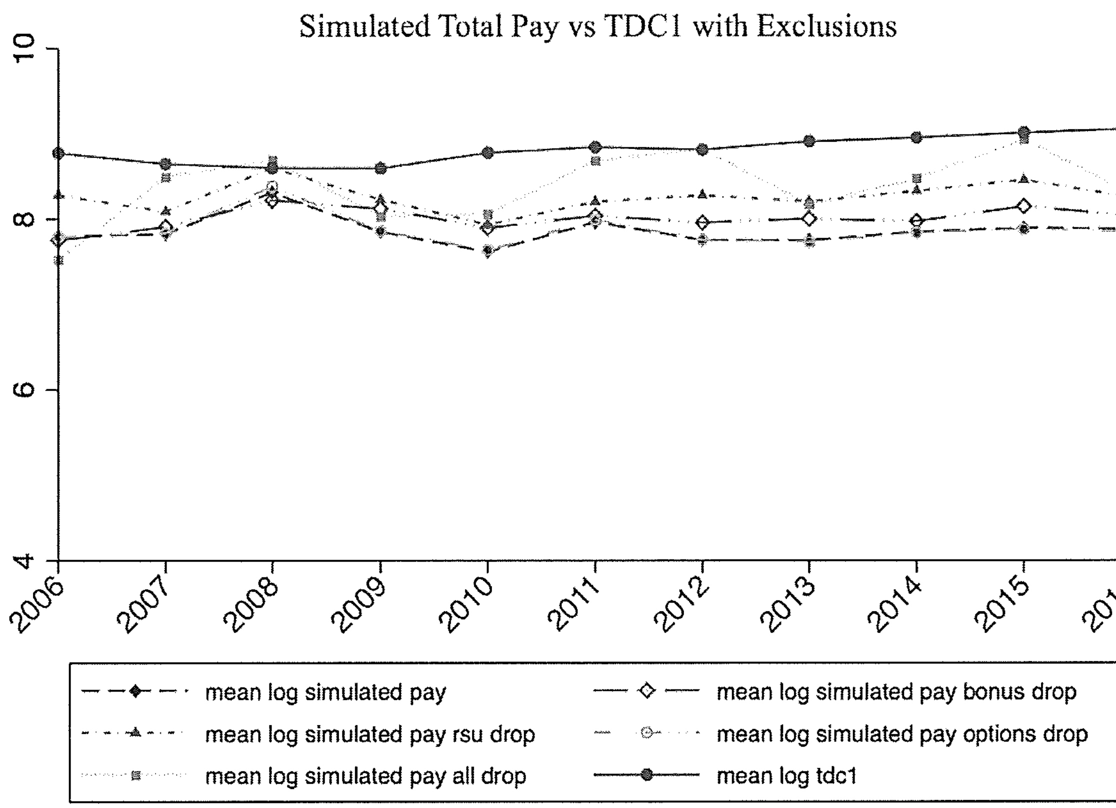

Figure 5. Simulated Pay Versus Actual Pay by GICS 4-digit Industry

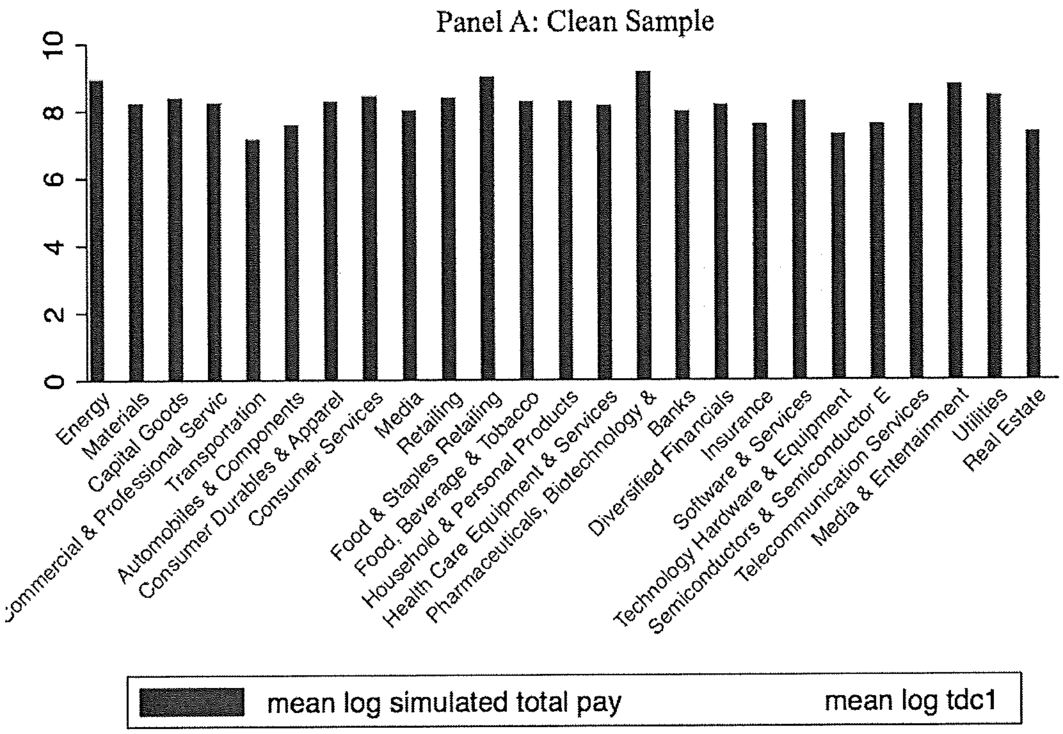

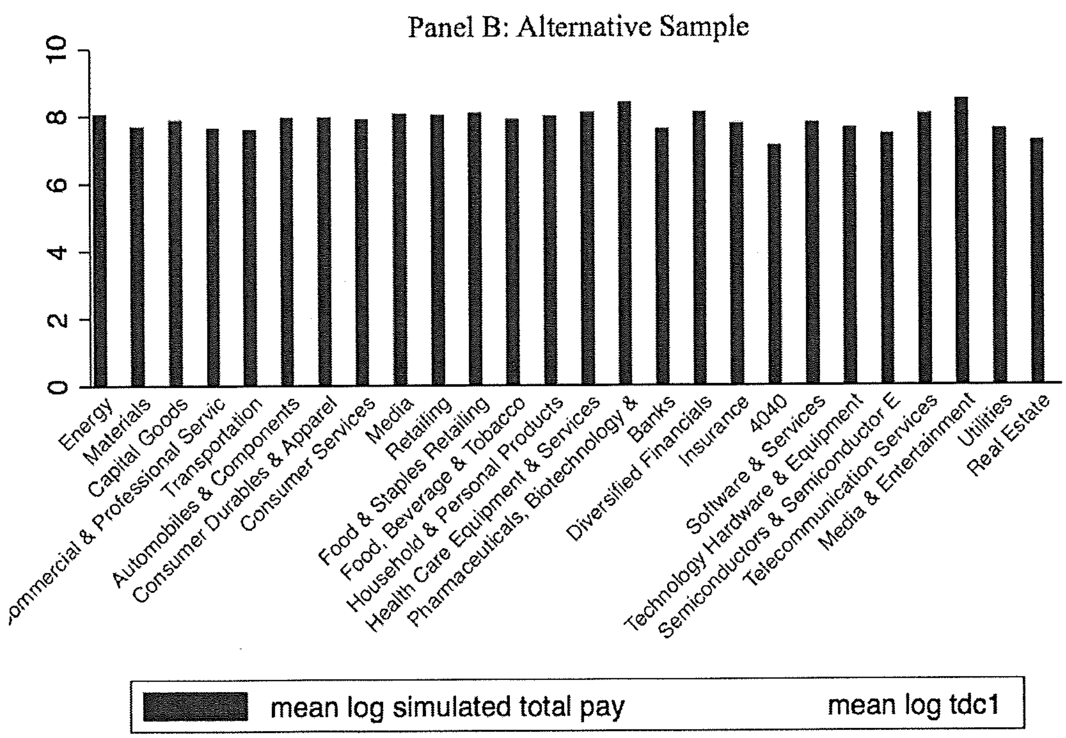

Table 1. Frequency Distribution of Performance Metrics

Panel A reports the frequency distribution of the performance metrics used in compensation contracts at the metric level. Panel B reports descriptive statistics for the number of performance metrics per grant/year, and the number of grants per CEO/year. For both Panel A and Panel B, Columns 1 to 4 (5 to 8) present results for the clean (alternative) sample.

Panel A. Metric Level Information

	Clean Sample				Alternative Sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Metric	Bonus	Restricted Stock	Options	Combined	Bonus	Restricted Stock	Options	Combined
Book Value	0	1	0	1	0	2	0	2
Cashflow	13	15	0	28	223	81	0	304
EBIT	7	5	0	12	60	12	0	72
EBITDA	22	13	0	35	160	78	2	240
EBT	16	16	0	32	87	30	0	117
EPS	37	33	0	70	414	287	5	706
Earnings	6	11	0	17	146	47	0	193
FFO	2	1	0	3	22	5	0	27
Operating Income	22	15	0	37	252	94	0	346
Profit Margin	3	3	0	6	53	26	0	79
ROA	1	10	0	11	32	29	0	61
ROE	9	11	0	20	84	83	1	168
ROI	0	1	0	1	7	5	0	12
ROIC	6	13	0	19	88	140	0	228
Sales	32	11	0	43	371	184	1	556
Stock Price	0	7	0	7	10	72	5	87
Time	0	0	290	290	0	0	5,031	5,031
Total (metric level)	**176**	**166**	**290**	**632**	**2,009**	**1,175**	**5,045**	**8,229**

Table 1. Panel B. Grant Level and CEO Level Information

	Clean Sample				Alternative Sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Bonus	Restricted Stock	Options	Combined	Bonus	Restricted Stock	Options	Combined
Number of performance metrics per grant/year								
Min	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Mean	1.43	1.25	1.00	1.16	1.61	1.37	1.00	1.15
Std. Dev.	0.62	0.51	0.00	0.43	0.74	0.61	0.01	0.45
Skewness	1.11	1.97	.	2.76	1.19	1.50	71.00	3.40
Max	3.00	3.00	1.00	3.00	5.00	4.00	2.00	5.00
Number of grants per CEO/year								
Min	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Mean	1.10	1.12	2.00	1.90	1.08	1.10	1.14	1.37
Std. Dev.	0.38	0.37	2.92	2.15	0.32	0.34	0.72	0.86
Skewness	5.03	3.30	3.05	4.03	5.07	3.81	10.29	5.75
Max	4.00	3.00	12.00	12.00	4.00	3.00	12.00	12.00

Table 2. Descriptive Statistics

N is the number of observations, and *Px* is the percentile *x* value of the sample distribution, with *x*=1, 25, 50 (median), 75, and 99. The appendix gives detailed definitions of each variable, data source and time availability.

VARIABLES	N	Mean	Std. Dev.	P1	P25	P50	P75	P99
ln of simulated CEO total pay (clean sample with options)	287	8.15	1.04	6.17	7.39	8.16	8.77	11.23
ln of total compensation (TDC1)	37,322	7.91	1.07	5.30	7.16	7.92	8.68	10.34
ln of CEO wealth	29,567	9.90	1.44	6.73	8.93	9.83	10.79	13.97
3-year shareholder return	34,963	0.11	0.25	-0.53	-0.03	0.10	0.23	1.02
Firm return volatility	35,936	0.11	0.06	0.04	0.07	0.10	0.14	0.34
Average industry return	37,319	0.16	0.24	-0.41	0.01	0.15	0.30	0.85
ln of market value	32,836	7.45	1.64	3.63	6.35	7.35	8.47	11.67
Overconfidence indicator	37,320	0.34	0.47	0.00	0.00	0.00	1.00	1.00
Co-opted board (coopt)	20,905	0.56	0.32	0.00	0.29	0.55	0.88	1.00
CEO age (age)	36,495	55.82	7.38	39.00	51.00	56.00	60.00	76.00
CEO tenure	34,731	8.23	7.39	0.92	2.92	5.92	10.92	35.92
Founder indicator (founder)	37,322	0.11	0.32	0.00	0.00	0.00	0.00	1.00
Percent of institutional ownership	29,852	0.68	0.22	0.07	0.54	0.70	0.83	1.13
Simulated variance of CEO pay (ln of variance(simulated pay)) (clean sample with options)	287	14.44	3.09	4.88	12.95	14.83	16.25	21.88
Realized variance of CEO pay (ln of variance(TDC1))	24,813	13.80	2.53	6.75	12.34	13.89	15.45	18.66
Realized variance of CEO wealth (ln of variance(wealth))	17,147	18.04	3.20	10.69	16.15	18.02	19.90	25.52
Simulated skewness of CEO pay (skewness(simulated pay)) (clean sample with options)	287	1.45	3.41	-2.13	0.15	0.82	1.66	28.30
Realized skewness of CEO pay (skewness(TDC1))	24,794	0.29	0.70	-1.33	-0.21	0.31	0.84	1.49

Table 3. Panel A. The Risk and Reward trade off in Pay Using Simulated Conditional Volatility

This table present results from regressions of the natural log of simulated compensation using Incentive Lab data and compensation contract information available at the beginning of each year. Columns 1 and 2 (3 and 4) report results using the clean sample for firms that do not (do) use stock options for CEO compensation. Columns 5 and 6 (7 and 8) report results using the alternative sample for firms that do not (do) use stock options for CEO compensation. Columns 1, 3, 5, and 7 (2, 4, 6, and 8) report results without (with) firm and year fixed effects included. Robust t-statistics are reported in parentheses, clustered by firm (firm and year) in Columns 1, 3, 5, and 7 (2, 4, 6, and 8). Significance at 1%, 5%, and 10% is indicated by ***, **, and *, respectively.

	Clean Sample (No Options)		Clean Sample (With Options)		Alternative Sample (No Options)		Alternative Sample (With Options)	
VARIABLES	(1) log Simulated Mean Pay	(2) log Simulated Mean Pay	(3) log Simulated Mean Pay	(4) log Simulated Mean Pay	(5) log Simulated Mean Pay	(6) log Simulated Mean Pay	(7) log Simulated Mean Pay	(8) log Simulated Mean Pay
Log Simulated Variance of Pay	0.148***	0.091***	0.212***	0.151***	0.140***	0.108***	0.171***	0.137***
	(5.40)	(4.24)	(7.52)	(3.89)	(11.19)	(11.04)	(23.25)	(14.62)
Constant	6.080***		5.165***		6.166***		5.641***	
	(15.01)		(12.32)		(35.51)		(56.79)	
Observations	143	99	287	204	1,671	1,508	5,202	5,072
Adj. R-squared	0.505	0.668	0.592	0.694	0.390	0.689	0.481	0.694
Firm and Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Cluster s.e.	Firm	Firm/Year	Firm	Firm/Year	Firm	Firm/Year	Firm	Firm/Year

Table 3. Panel B. Fama-MacBeth Regressions

VARIABLES (RHS variance)	(1) log of Simulated Pay (Clean Sample With Options)	(2) log of Simulated Pay (Alternative Sample With Options)
CROSS SECTION		
2006	0.431	0.171
2007	0.356	0.179
2008	0.247	0.242
2009	0.276	0.200
2010	0.209	0.133
2011	0.112	0.198
2012	0.120	0.145
2013	0.154	0.124
2014	0.234	0.155
2015	0.189	0.181
2016	0.272	0.151
Average slope	0.24***	0.17***
T-stat (corrected)	3.29	6.13
Observations	287	5,202
Number of groups	11	11

Table 4. The Risk and Reward trade off in Pay Using Realized Conditional Volatility

This table presents results from regressions of the natural log of TDC1 (Columns 1 and 2) and the natural log of CEO wealth (Columns 3 and 4). The measures of conditional variance of pay are lagged to reflect the information known at the beginning of the period and are based on TDC1 (Columns 1 and 2) or CEO wealth (Columns 3 and 4). Panel A reports robust *t*-statistics in parentheses. Panel B reports Fama-Macbeth regressions with standard errors on the average slope corrected according to Pontiff (1996). Significance at 1%, 5%, and 10% is indicated by ***, **, and *, respectively.

VARIABLES	(1) log of TDC1	(2) log of TDC1	(3) log of CEO wealth	(4) log of CEO wealth
Lag log realized var(TDC1)	0.246***	0.046***		
	(38.92)	(6.29)		
Lag log realized var(CEO Wealth)			0.260***	0.099***
			(9.20)	(5.32)
Constant	4.789***		5.526***	
	(54.75)		(10.43)	
Observations	16,769	16,522	11,971	11,744
Adjusted R-squared	0.405	0.760	0.366	0.811
Firm and Year FE	NO	YES	NO	YES
Cluster s.e.	Firm	Firm/Year	Firm	Firm/Year

Table 5. The Risk and Reward trade off in Pay Using ARCH Conditional Volatility

The table presents estimates of ARCH-in-mean models on TDC1 and CEO wealth (TDC1 in columns 1 and 2, and CEO wealth in columns 3 and 4). The estimations assume an ARCH(p) model for the conditional heteroskedasticity; industry fixed effects are from one-digit SIC; the ARCH-in-mean term is the natural logarithm of the estimated variance of the left-hand side variable; t-statistics are computed using White robust standard errors; the residuals follow a student-t distribution and the priming values are obtained from the estimated variance of the residuals from OLS. Significance at 1%, 5%, and 10% is indicated by ***, **, and *, respectively.

VARIABLES	(1) TDC1	(2) TDC1	(3) CEO wealth	(4) CEO wealth
Lag log var(TDC1)	128.0***	150.0***		
	(12.02)	(12.41)		
Lag log var(CEO Wealth)			-96.9***	-27.6
			(-5.48)	(-1.46)
Constant	-87.0	-2242.4***	10843***	44033***
	(-0.52)	(-3.39)	(26.49)	(18.42)
Industry and Year FE	NO	YES	NO	YES
ARCH(1) coefficient	2.16***	2.42***	1.68***	1.72***
	(22.33)	(20.10)	(34.55)	(34.20)
ARCH(2) coefficient			0.00005	0.00001
			(0.62)	(0.36)
ARCH(3) coefficient			0.0002**	0.00006
			(2.40)	(1.55)
ARCH constant	2.57***	2.37***	36.5***	28.5***
(in millions)	(18.67)	(16.26)	(15.41)	(13.76)
Observations	37,322	37,322	29,567	29,567

Table 6. Cross-Sectional Analysis Using Simulated Conditional Volatility for Risk Aversion

The table presents results examining the cross-sectional variation in the risk and reward trade off in pay based on five risk aversion proxies using all the firms in the alternative sample: (i) medium fatality relative to low and high fatality CEOs (Column 1), (ii) pilot CEOs relative non-pilot CEOs (Column 2), (iii) CEOs born in the great depression relative to other CEOs (Column 3), (iv) female CEOs relative to male CEOs, and (v) married CEOs relative to unmarried CEOs. The dependent variable is the natural log of simulated compensation (i.e., expected annual compensation at the end of the year) based on the compensation contract information available at the beginning of the year from Incentive Lab. The measures of conditional volatility of pay are obtained from the simulation. All regressions include firm and year fixed effects. Robust t-statistics are reported in parentheses, clustered by firm and year. Significance at 1%, 5%, and 10% (2-sided) is indicated by ***, **, and *, respectively.

	Low Risk Aversion		High Risk Aversion		
	(1)	(2)	(3)	(4)	(5)
VARIABLES	MediumFatality	Pilot	Depression	Female	Married
Log simulated variance	0.226***	0.163***	0.158***	0.161***	0.129***
	(10.82)	(15.37)	(14.48)	(14.60)	(7.60)
Simulated skewness	0.030***	0.015***	0.016***	0.016***	0.008
	(4.87)	(6.06)	(4.68)	(6.35)	(1.76)
Risk aversion	1.598***	0.588*	-0.358	-0.636	-0.628*
	(4.08)	(2.15)	(-1.41)	(-1.67)	(-1.97)
Risk aversion*variance	**-0.099*****	**-0.050****	**0.038****	**0.040**	**0.047***
	(-3.93)	(-2.39)	(2.33)	(1.43)	(2.03)
Risk aversion *skewness	-0.020**	-0.007	0.056	0.010	0.012**
	(-2.94)	(-1.36)	(1.37)	(1.36)	(2.26)
Expected sign on Risk aversion*variance coefficient	-	-	+	+	+
1-sided p-value statistic	0.00	0.02	0.02	0.09	0.04
Observations	1,211	5,054	3,962	4,570	3,267
Adjusted R-squared	0.732	0.709	0.705	0.704	0.737

Table 7. Alternative Hypotheses Using Simulated Conditional Volatility

The table evaluates several alternative hypotheses using panel regressions and simulated conditional variance of pay. The dependent variable is the natural log of simulated compensation based on the compensation contract information available at the beginning of the year from Incentive Lab. The measures of conditional volatility of pay and skewness of pay are obtained from the simulation. All regressions include firm and year fixed effects. Robust *t*-statistics are reported in parentheses, clustered by firm and year. Significance at 1%, 5%, and 10% is indicated by ***, **, and *, respectively.

	Dependent variable is log simulated mean pay						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log simulated variance	0.137***	0.158***	0.137***	0.133***	0.133***	0.136***	0.165***
	(14.56)	(15.80)	(13.69)	(12.46)	(13.60)	(14.19)	(12.83)
Simulated skewness		0.015***					0.018***
		(6.92)					(6.80)
Lag 3-year stock return			0.100				0.123
			(1.25)				(1.39)
Lag avg industry return			0.101				0.086
			(1.16)				(0.90)
Coopt				0.188**			0.054
				(3.01)			(0.77)
Institutional holding %				0.086			0.089
				(0.80)			(0.97)
Firm return volatility					0.612		0.593
					(1.13)		(1.25)
Founder					0.045		-0.010
					(0.87)		(-0.16)
Age					0.007*		0.006*
					(1.90)		(1.89)
Log CEO tenure					0.051**		0.035
					(2.33)		(0.98)
Log lag market value		0.168***	0.143***	0.216***	0.172***	0.165***	0.181***
		(5.86)	(4.97)	(6.10)	(6.16)	(5.35)	(6.80)
Overconfidence						0.023	-0.062*
						(0.86)	(-2.15)
Observations	5,054	4,964	4,857	3,961	4,830	4,964	3,804
Adjusted R-squared	0.694	0.720	0.706	0.708	0.705	0.706	0.723

Table 8. Alternative Hypotheses Using ExecuComp Sample

This table evaluates several alternative hypotheses. Columns (1)—(3) report results from panel regressions for log TDC1. Columns (4)—(5) report results from ARCH estimation for TDC1. Realized volatility of pay is used in Columns (1) and (3); simulated volatility of pay is used in Column (2); and ARCH volatility of pay is used in Columns (4) and (5). The full ExecuComp sample is used in Columns (1), (2), and (4); ExecuComp / Incentive Lab matched sample is used in Columns (3) and (5). The measure of realized volatility of pay is computed as the lagged sample variance of last five year pay and is based on TDC1. The measure of simulated volatility of pay is obtained from the simulation exercise, using compensation contract information available at the beginning of the year from Incentive Lab. All regressions include firm and year fixed effects. Robust *t*-statistics are reported in parentheses, clustered by firm and year. Significance at 1%, 5%, and 10% is indicated by ***, **, and *, respectively.

	Panel Regression			ARCH Estimation	
	(1)	(2)	(3)	(4)	(5)
VARIABLES	Log of TDC1	Log of TDC1	Log of TDC1	TDC1	TDC1
Log variance of pay	0.024***	0.018***	0.009	143.4***	253.2***
	(3.24)	(4.15)	(1.26)	7.33	4.01
Skewness of pay	-0.042***	0.003**	-0.002	-73.43*	-61.77
	(-3.15)	(2.80)	(-0.14)	(-1.76)	(-0.59)
Lag 3-year stock return	0.213***	0.148**	0.162	-1214***	-878.9**
	(3.39)	(2.32)	(1.64)	(-9.15)	(-2.09)
Lag avg industry return	0.123***	0.094*	0.103**	92.77	361.4
	(3.79)	(1.94)	(2.51)	0.79	1.00
Coopt	0.153**	-0.006	0.015	727.1***	424.7
	(2.54)	(-0.11)	(0.13)	6.02	1.05
Institutional holding %	0.147	-0.026	-0.032	649.5***	600.7
	(1.39)	(-1.03)	(-0.28)	4.11	1.41
Firm return volatility	-0.260	-0.846**	-1.054*	5397***	9217***
	(-0.93)	(-2.33)	(-2.12)	9.05	4.14
Founder	-0.092**	-0.092*	-0.056	-458.5***	-831.0***
	(-2.43)	(-2.21)	(-1.43)	(-5.87)	(-3.03)
Age	-0.007**	0.003	0.005	7.94*	36.88**
	(-2.17)	(0.99)	(0.96)	1.74	2.36
Log CEO tenure	0.004	0.071*	0.024	-240.3***	453.0**
	(0.13)	(2.23)	(0.46)	(-3.51)	2.00
Log lag market value	0.166***	0.171***	0.177***	1785***	2492***
	(6.38)	(5.91)	(4.23)	35.87	31.68

Overconfidence	0.072**	-0.017	0.000	20.93	-253.8
	(2.78)	(-0.60)	(0.01)	0.34	(-1.63)
ARCH(1) coefficient				1.58***	1.26***
				13.79	8.25
ARCH constant				2.62***	3.56***
(in millions)				13.07	6.75
Observations	10,346	3,804	2,467	10,530	2,591
Adjusted R-squared	0.783	0.800	0.819		