What Steve Ross Taught me about Contracting

May 8, 2018

Philip H. Dybvig

Olin School of Business, Washington University in Saint Louis, Campus Box 1133, One Brookings Drive, Saint Louis, MO 63130-4899, USA

Philip H. Dybvig is the Boatmen's Bancshares Professor of Banking and Finance at Washington University in Saint Louis, and the Dean of the Institute of Financial Studies at the Southwestern University of Finance and Economics in Chengdu.

What Steve Ross Taught me about Contracting
May 8, 2018

Steve Ross may be best known in finance for his work on asset pricing, but he is also the first person in economics or finance to study agency theory as a theory of incentive contracting, as for a portfolio manager or the CEO or other employee of a firm. This paper looks at some lessons we can learn from Steve's work in agency theory. Incentive contracting involves a trade-off between incentives and risk sharing. Traditional performance measures such as Jensen's alpha and the Sharpe measure can fail for market timers. And, convex compensation from call or put options, do not necessarily increase incentives for risk-taking. Under reasonable assumptions, adding a put option to equity compensation increases incentives for risk-taking, but adding a call option is ambiguous.

Steve Ross published over a hundred academic papers in which he made many important contributions to different fields of economics. One of the most important and influential contributions was in a short paper, Ross (1973), which was the first paper on agency in economics or finance. Steve pointed out that there are interesting economic implications for the legal principal of agency, concerning one person (the agent) doing work on behalf of another (the principal). This initial paper made several important contributions: (1) modelling the conflict of interest between the principal and agent as an optimal contracting problem of choosing the compensation contract for the agent, (2) labelling this conflict the "agency" problem, now standard terminology in economics and finance, (3) looking at the contracting problem from the perspective of the principal, and (4) the first-order approach (a technical approach still widely used today). In this paper, I will discuss some of the ideas about incentives we can learn from Steve's research and the literature it generated. I will also talk about Steve's related work on portfolio performance measurement and whether convex compensation gives managers an incentive to take on risk, as well as comment briefly on Steve's work on forensic finance. I should add that I do this from the perspective of a student of Steve's who owes him everything and also worked with Steve as a co-author.

AGENCY THEORY

The standard approach to agency started by Steve poses a variational problem. We will talk about the problem in words and write it down formally, but we will not solve it, which is harder than might be expected even in simple cases. Instead, we will talk about the economics and the lessons that can be learned.

First we describe the economic setting for the basic agency problem. The principal wants to hire the agent to help with a project. The output of the project depends both on the agent's effort, denoted by e, and random luck, denoted by ε. The output is generated according to a known function $Y(e, \varepsilon)$. The output is observable to everyone, and can be contracted upon, but the agent's effort e is not observed directly. Because effort is not directly observable, it is typically not clear to what extent the level of output is due to effort or luck. It is not possible to punish the agent for low effort directly (because that is not observable). Instead, we must be content to punish the agent for low output (which is observeable). This sort of punishment can generate incentives for efficient effort if the agent is risk neutral, since that would imply that the agent does not care about being punished unfairly some of the time so long as pay is fair on average. The interesting case is the case in which the agent is risk averse. If the agent is risk-averse, you can generate incentives to work hard and generate the socially efficient effort, but only by forcing the agent to take on a lot of risk, for which the agent must be compensated. This is the rub, and we can say quite generally that the traditional agency problem concerns the trade-off between incentives and risk-sharing. Good risk-sharing probably implies that the principal should take on a lot of risk, which unfortunately implies that the agent has little incentive to exert effort.

At this point, you may say that this is a cynical view of the world, and that professionals will often do the correct thing even if doing the wrong thing is more profitable in the short-run. After all, there is such a thing as professional pride, and their boss may notice they are not working hard or they may be concerned about their reputations if they are discovered slacking. These are all fair points, but this is the normal mode of analysis for economic modelling: we write down a very simple model in which people respond directly, almost robotically, to the incentives they face. Simplicity of the model makes it tractable so we can form definite conclusions that can be applied to practice. Steve used to say that finance is different than, say, mathematical economics: we make strong assumptions so we can derive testable conclusions, and that is important because we have work to do! One necessary part of this approach is that we approach practice with some humility, and we should think broadly about whether the assumptions make sense when we bring our model to practice. Steve argued in his work on forensic finance that rigid adherence to a single model or point-of-view is a common feature of many catastrophic financial failures.

Moving back to the economics of Ross's agency model, we can think of things proceeding in four stages. First, the principal offers the agent the engagement, with a fee schedule $\phi(\cdot)$ specifying how much the agent will be paid as a function of the success of the project measured by its realized output. Second, the agent chooses whether to take the engagement. Third, the agents chooses how much effort e to expend on the project. Finally, luck plays out and everyone observes the ouput $Y(e, \varepsilon)$ realized by the project. The proceeds are divided up according to the fee schedule: the agent receives the fee and the principal gets the output net of the fee. If we were modeling this as a game, we would have technical problems because there always some contracts the principal can offer for which there is no optimal response, but Steve's approach cleverly side-steps that issue. In particular, in his specification the principal chooses a plan for the agent and puts in as constraints the requirement that the agent does as planned.

We will not show how to solve Ross's agency problem, but instead we will discuss the economics. It is, however, useful to write down a formal problem of the sort Steve solved:

Problem 1 *Choose a fee schedule $\phi(Y)$ and effort e for the agent to maximize the principal's utility of consuming output net of the fee,*

$$(1) \quad E[U^P(Y(e^*, \varepsilon) - \phi(Y(e^*, \varepsilon))],$$

subject to the agent's participation constraint,

$$(2) \quad E[U^A(\phi(Y(e^*, \varepsilon)))] \geq U^A_R],$$

and the constraint that the agent is happy to do as the principal plans (incentive-compatibility):

> *The effort level e^* solves the agent's problem, which is to choose e to maximize*
>
> $$(3) \quad E[U^A(y(e, \varepsilon))].$$

This problem uses notation we have already discussed for effort e, fee schedule $\phi(Y)$, random luck ε, and output $Y = Y(e, \varepsilon)$. The utility functions ($U^P(\cdot)$

for the principal and $U^A(\cdot, \cdot)$ for the agent) represent preferences for risk-taking for both principal and agent, and preference for effort as well for the agent. The constant U_R^A is the reservation utility of the agent, which can be thought of how well off the agent would be taking another available offer, or it could also depend on the relative bargaining power of the principal and agent.

All of this is from the perspective of the principal, who is choosing the compensation contract and makes plans for what the agent will do, subject to the agent being happy to follow the plans. The agent's choices are represented in the constraints: the participation constraint says that the agent is made well enough off to to be able to follows the principal's plan and accept the engagement, and the incentive-compatibility constraint says the agent wants to choose the effort the principal intends.

The optimal fee schedule that solves Problem 1 can be viewed as trading off incentives and risk-sharing. Ideal incentives would cause the agent to expend the optimal level of output but would give the agent too much risk for which the agent would have to be compensated. Ideal risk-sharing would expose the agent to less risk so the agent doesn't have to be paid so much, but incentives would be weak so the effort would be suboptimal. Usually, we think the problem is that to give the agent enough of an incentive to work hard, but incentives to work too hard can also be inefficient. Also, I am talking about a trade-off between incentives and risk-sharing (which is usually the critical issue), but we also also need to avoid fee schedule that has poor incentives as well as poor risk-sharing, for example, if the fee depends negatively on the output.

The trade-off between incentives and risk-sharing can be thought of as being related to the two constraints of Problem 1. The incentives for effort are given by the incentive-compatibility constraint of maximization of (3) that says the agent wants to choose the effort level that the principal plans. Intuitively, maximizing (3) typically means that the more sensitive compensation is to output, the harder the agent will work. On the other hand, risk sharing shows up in the participation constraint (2). If the agent is risk-averse and does not like noise in compensation, the more random compensation is the more the principal has to pay. The principal cares about these things because they appear in the principal's expected utility (1). Increasing effort increases the output $Y(e, \varepsilon)$ which increases the principal's expected utility, while increasing the agent's compensation (to compensate the agent for taking on risk) reduces the residual $Y(e^*, \varepsilon) - \phi(Y(e^*, \varepsilon))$ the principal consumes, which reduces the principal's expected utility. The only time when optimal incentives corresponds with optimal risk-sharing is the case when the agent is risk-neutral. For this special case, the optimal incentive contract is similar to selling the project to the agent.

Hölmstrom (1979) has a nice paper extending Ross (1973). His paper reinforces the trade-off between incentives and risk-sharing. He shows that deviations from optimal risk-sharing will only depend on information that is affected at the margin by effort. This is because any dependence on uninformative information adds noise to compensation (and makes risk-sharing worse). Similarly, adding noise to information about performance makes the contracting less efficient if the agent is risk-averse and the information about performance is useful.

To make the discussion more concrete, consider the reasonable case of a risk-neutral principal and a risk-averse agent. This is probably a good approximation for a principal which is a big firm or a rich employer of a much smaller employee. Given that the employer is indifferent about taking on risk and the employee does not like risk, it is optimal risk-sharing for the employer to take on all the risk. However, this provides the employee no incentive at all to do any work. On the other hand, paying the employee one-for-one for the change in the value of the firm will give the employee an incentive to undertake a lot of effort. Unfortunately, this provides terrible risk sharing and the employee must be paid a lot to take on so much risk. Think about how much you would have to pay the CEO of a Fortune 500 company to take on all the risk of the change in value of the firm. The optimal fee structure is a compromise between these two, giving the employee more when the firm does well and less when the firm does badly, but with much less variation than the value of the entire firm.

PORTFOLIO PERFORMANCE EVALUATION

To the extent that portfolio performance measures are used for compensation or retention, choosing the optimal performance measure is also a form of agency problem. However, the academic literature on portfolio performance measurement is largely separate from the literature of agency in portfolio management.[1] My work on performance management with Steve, Dybvig and Ross (1985a,b), looks at properties of traditional performance measures, such as the Sharpe ratio and Jensen's alpha. The paper shows that (along the lines of Mayers and Rice (1979)), that optimal use of "security-specific" information will improve the Sharpe measure and Jensen's alpha, but optimal use of "market timing" information that is informative about the market's return may go the other direction.

The original motivation of these measures was based on intuition, not theory. For example, the mean-variance CAPM, like the one-factor APT, if we plot mean asset returns as a function of the sensitivity beta to the market (or other common factor in the APT), then all assets plot on a single line called the security market line (SML). This gives us an *intuition* that the SML gives the amount of expected return justified by the level of risk measured by beta, and that the deviation of excess returns from this line, known as Jensen's alpha because it was proposed by Mike Jensen, represents superior performance (if alpha is positive) or inferior performance (if alpha is negative). Note that this interpretation does not follow from the theory, since the theory implies all securities plot on the SML and that therefore all Jensen's alphas are zero, at least in expectation. We could have deviations from the security market line due to statistical measurement error in the realization, but of course that is random noise unrelated to superior performance. Indeed, the derivation of the CAPM and the SML assumes that all agents are optimizing and have the same beliefs about means and variances. Consequently, Jensen's alpha cannot measure superior performance because there is no possibility of superior performance in the model. A theoretical justification or refutation of the traditional measures requires a model, like the one in Mayers and Rice (1979) or Dybvig and Ross (1985a), that admits the possibility of superior

[1] A notable exception is Dybvig, Farnsworth, and Carpenter (2010), which is technically a little complicated. That paper shows that a linear contract gives good incentives for information-gathering, but does not give incentives not to be a closet indexer.

performance.

The most striking result in Dybvig and Ross (1985a) is a sensible example in which the traditional measures, Jensen's alpha and the Sharpe ratio, can misidentify a successful market timer as an inferior performer. The assumptions of the example are

1. There is a riskless asset and a single risky asset (which we can think of as the market).
2. The payoff of the risky asset is normally distributed.
3. The informed agent whose performance we are considering maximizes expected utility given constant absolute risk aversion (exponential utility).
4. The informed agent's information has the form of a private signal that is joint normally distributed with the payoff on the risky asset.
5. The informed observer bases estimates on the unconditional expectation of the statistics, abstracting from statistical measurement error.

Having a single risky asset gives this analysis the flavor of market timing rather than security selection. The question in the end is whether the traditional measures correctly identify superior performance by this agent who definitely has superior information and is using it optimally. Assuming constant absolute risk aversion (CARA utility) and normally distributed returns is a common tractible approach to mean-variance analysis, and was assumed to give the measures as good a fighting chance to do well. Joint normality of the signal implies that the risky asset return is normally distributed both unconditionally and conditional on the the agent's information. Similarly, abstracting from measurement error is biased in favor of the measures, since measurement error can only cause problems.

In notation,[2] we have a riskless asset with constant rate of return r and a risky asset with random excess return x with mean π (the risk premium) and

[2]This paper uses different notation than the Dybvig and Ross (1985a) Section I, but the example is equivalent, except that this paper assumes w.l.o.g. that the risky asset is the benchmark ($\alpha = 1$).

standard deviation σ_x. The informed agent receives a signal s with mean 0 and standard deviation σ_s. The signal s and excess return x are joint normally distributed with covariance σ_{sx}. The portfolio choice $\theta(s)$ maximizes expected utility $E[-exp(-A\theta(s)x)]$, where the constant $A > 0$ is the coefficient of absolute risk aversion. Given joint normality, the expectation of x given s is given by the linear regression $\mu_{x|s} = \pi + s\sigma_{sx}/\sigma_s^2$, and the variance if x given s is $\sigma_{x|s}^2 = \sigma_x^2 - \sigma_{xy}^2/\sigma_s^2$. The optimal portfolio choice is $\theta(s) = \mu_{x|s}/(\sigma_{x|s}^2 A) = (\pi + s\sigma_{sx}/\sigma_s^2)/((\sigma_x^2 - \sigma_{xy}^2/\sigma_s^2)A)$. Note that in this expression everything is constant except for s, so that the portfolio choice depends linearly on the signal s. We set initial wealth to 1 to save on notation, since this choice will not affect our conclusion.

Jensen's alpha, the excess return above the security market line, is measured by the observer as

$$(4) \qquad \alpha = (E[r + \theta(s)x] - r - \frac{cov(\theta(s)x, x)}{var(x)} E[x],$$

the expected return on the portfolio less the riskfree rate less a risk premium computed as the exposure to market risk (beta) times the expected return on the market. The observer will normally use sample values to calculate the expectations, variances, and covariances, but we will abstract from the problems with sampling and use the population expectations. After some algebra (more details in Dybvig and Ross (1985a)), we find that

$$\begin{aligned}
(5) \qquad \alpha &= \sigma_s^2 \frac{sigma_x^2 - \pi^2}{\sigma_{x|s}^2 A \sigma_x^2} \\
(6) \qquad &< 0,
\end{aligned}$$

provided $\pi^2 > \sigma_x^2$. This shows that in general the informed manager who is behaving optimally (and faces a mean-variance problem conditional on s) has a negative α, i.e., plots below the SML.

The reason for the failure of the performance measures is related to measurement issues and the statistical properties of the estimates from the perspective of an uninformed observer. Although we are abstracting from measurement errors, we are assuming that the observer uses sample moments which look different from the conditional moments the informed agent sees.[3] We can illustrate this point using a simpler example. If we are looking at nominal returns, the return of a 3-month Treasury Bill over its three months to maturity is riskless. However, if we compute the variance of its returns, we will get a positive variance because the return is moving around. However, this risk does not represent variance from the perspective of the investor, because conditional on your information at the start of a period you know what will be the yield of the Treasury Bill over the three months to maturity. This is a problem with the variances we compute using the natural measures (which includes both the conditional variance the investor sees and the variance in conditional means). There is also a problem with the beta coefficient: because the amount of stock holding is higher when the return is expected to be good, the estimated beta is higher than the average beta of the portfolio.

One particularly vexing problem with this example is that for the same parameter values, the agent's portfolio plots inside the mean-variance frontier for constant (no-information) portfolios (see Dybvig and Ross (1985a) for the algebra). The reason this can happen is interesting. Because returns of constant portfolios are normally distributed, the CARA agent has mean-variance preferences for these portfolios. Also, given the signal s, all portfolios are normally distributed. However, if we look at the manager's portfolio unconditionally, it is not normally distributed. In particular, it is $\theta(s)x$, where θ is linear in s and s and x are joint normal. The correlated part of s and x means there is a normal-squared (χ^2) term that the informed agent likes more than a normal random variable with the same mean and variance. Although the informed agent's portfolio is dominated by some constant portfolios in mean and variance, it is still preferred.

In my view, our understanding of the implications of this example is still incomplete. Although it seems to refute the usefulness of Jensen's alpha, we are not sure whether constructing such an example is dependent on unreasonable parameter values we are unlikely to find in practice.

---

[3] In this simple example, a more sophisticated observer can back out $\theta(s) = (\theta(s)x)/x$ and then can back out $s = (\theta(s)((\sigma_x^2 - \sigma_{xy}^2/\sigma_s^2)A) - \pi)\sigma_s^2/\sigma_{sx}$, but of course that is only because of the very simple structure of the example.

INCENTIVES FOR RISK-TAKING

It's very common for people talking about executive compensation or compensation of portfolio managers to be concerned that convex compensation as from options may induce excessive risk-taking. Convex compensation means that the shape is curved upwards so that the pay-performance sensitivity is larger at larger values of underlying performance; a differentiable function $g(z)$ is convex if $g''(z) > 0$. In practice, convex compensation may come from bonuses for superior performance without penalties for inferior performance, granting of call options, or granting of puts on top of option grants. Exhibit 1 gives several examples. One reason people suspect that convex compensation encourages risk-taking is that the value of an option with a convex payoff usually increases in the volatility of the underlying. Another reason is that the classic results of Arrow (1965) and Pratt (1964) show that an increasing convex transform of a utility function makes the agent less adverse to taking on risk. However, Ross (2004) showed us that all of this intuition is mostly off the mark.

If an agent maximizes expected utility $E[u(w)]$, it is a standard result that an agent who maximizes instead $E[g(u(w))]$, where g is an increasing linear function of the form $g(z) = a+bz$ where $b > 0$, does not change preferences at all, since $E[g(u(w))] = g(E[u(w)])$, and maximizing an increasing function of the objective does not change the choice problem. Furthermore, Arrow and Pratt showed, using any one of several equivalent reasonable definitions of more risk averse, maximizing a convex increasing function g of utility makes the agent less risk averse, while maximizing a concave increasing function (curved downward instead of upward), makes the agent more risk averse. If the functions involved are smooth enough, the local absolute risk aversion of u is $-u''(w)/u'(w)$, while the local absolute risk aversion of $v(w) \equiv g(u(w))$ is $-u''(w)/u'(w) - u'(w)g''(u(w))/g'(u(w))$. Maintaining the assumption that $v'(w) > 0$ and $u'(w) > 0$ (both utility functions prefer more to less), we have that $g(u(w))$ is less (respectively more) risk averse than $u(w)$ if and only if $g'' > 0$ (respectively $g'' < 0$).

The problem we are considering seems very similar: we want to know which is more risk averse, $u(w)$ or $V(w) \equiv u(\phi(w))$ where $\phi(w)$ is the fee schedule for compensation. However, the analytics are completely different. Think first about the linear case. While we know that the risk aversion of $v(w) \equiv g(u(w))$ is exactly the same as the risk aversion of $u(w)$ when g is linear, the risk aversion of $V(w) \equiv u(\phi(w))$ is different from the risk aversion of $u(w)$ except in degenerate cases when $\phi(w)$ is linear. There are two reasons for this: one is that the argument to the utility function $u()$ takes on a different value in $V(w) = u(\phi(w))$. Another reason, is that risk averison depends on the sensitivity of wealth to w. Changing the mean payoff affects expected utility to first order, but changing the standard deviation affects expected utility to the second order. This is the mathematics implicit in the theory of diversification, which says to first order we do not care about small risks. Reducing the pay-performance sensitivity $\phi'(w)$ reduces the benefit proportionally, but it reduces the damage from risk more than proportionately (as a quadratic).

Overall, Steve tells us that there are three things that affect the desire of the manager to take risk. First, risk-taking declines when the compensation is more sensitive to the underlying. Second, risk-taking may change (in either direction) because the convex compensation moves you to a different wealth level where your risk-taking preferences are different. Third, and this is the main intuition people have had in the past, more convex compensation (pay-performance sensitivity increasing with performance) tends to increase risk-taking. Since there are three different effects, it is only in special cases that we know the direction of the effect.

For example, if we add call options to equity compensation, the first effect may dominate and the increased pay-performance sensitivity may make the agent avoid risk. If we are willing to assume (reasonably[4]) that absolute risk aversion is decreasing, then we do have one positive result that adding a put option to equity compensation increases incentives for risk-taking. This is because all three effects go in the same direction: (1) pay-performance sensitivity is less (when the option is in the money) or the same (when the option is out of the money), (2) the additional compensation moves us to higher wealth levels where the risk aversion is less, and (3) compensation is more convex.

[4]Decreasing absolute risk aversion means that you have to compensate a wealthy person less than a poor person to take on the same dollar amount of risk. Algebraically, it says that $d(-u''(w)/u'(w))/dw < 0$.

CONCLUSION

Steve Ross may be best known in finance for his work on asset pricing, but he is also the first person in economics or finance to study agency theory as a theory of incentive contracting, as for a portfolio manager or the CEO or other employee of a firm. This paper looks at some lessons we can learn from Steve's work in agency theory. Incentive contracting involves a trade-off between incentives and risk sharing. Traditional performance measures such as Jensen's alpha and the Sharpe measure can fail for market timers. And, convex compensation, as from call or put options, does not necessarily increase incentives for risk-taking.

REFERENCES

Admati, Anat, and Stephen A. Ross, 1985, "Measuring Investment Performance in a Rational Expectations Equilibrium Model," *Journal of Business* **58**, 1–26.

Arrow, Kenneth J., 1965, *Aspects of the theory of risk-bearing*, Academic Bookstore, Helsinki.

Dybvig, Philip H. Farnsworth, Heber K., and Carpenter, Jennifer N., 2010, "Portfolio Performance and Agency," *Review of Financial Studies* **23**, 1–23.

Dybvig, Philip H. and Stephen A. Ross, 1985a, "The Analytics of Performance Measurement Using a Security Market Line," *Journal of Finance* **40**, 401–416.

Dybvig, Philip H. and Stephen A. Ross, 1985b, "Differential Information and Performance Measurement Using a Security Market Line," *Journal of Finance* **40**, 383–399.

Hölmstrom, Bengt, 1979, Moral Hazard and Observability, 1979, *Bell Journal of Economics* **10**, 74–91.

Mayers, David, and Edward M. Rice, 1979, “Measuring Portfolio Performance and the Empirical Conteent of Asset Pricing Models,” *Journal of Financial Economics* **7**, 3–28.

Pratt, John W., 1964, “Risk Aversion in the Small and in the Large,” *Econometrica* **32**, 122–136.

Ross, Stephen A., 1973, “The Economic Theory of Agency: The Principal’s Problem,” *American Economic Review* **63**, 134–139.

Ross, S.A., 2004, “Compensation, Incentives, and the Duality of Risk Aversion and Riskiness,” *Journal of Finance* **59**.

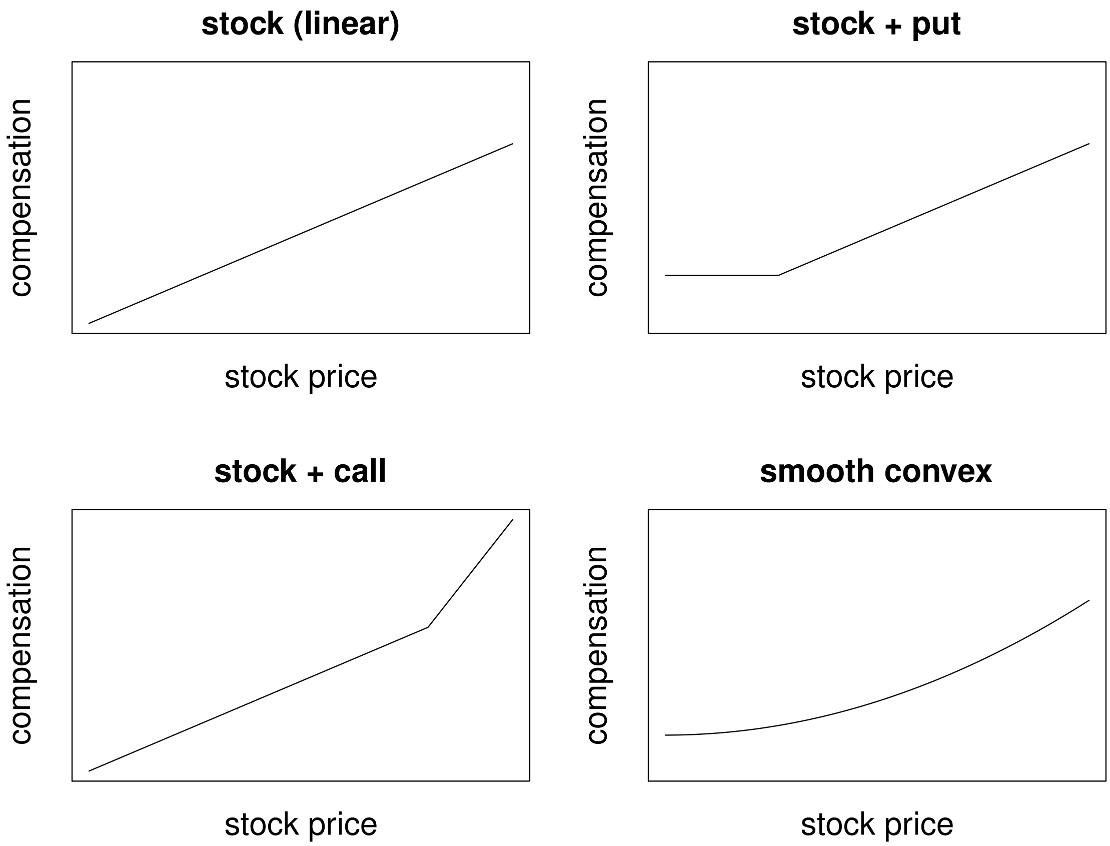

Exhibit 1: This graph illustrates linear compensation and several examples of convex compensation. While the underlying is labeled as the stock price, in another compensation it could be the portfolio payoff or earnings. While the expected value of convex compensation increases when volatility increases, Ross shows that it does not necessarily increase incentives for risk-sharing over linear compensation.