Aggregation and Convexity in the Provision of Dynamic Incentives

Thomas Hemmer*

May, 2017

Abstract

In this paper I identify an alternative preference structure that preserves most of the cherished simplicity of the formulation of the Principal-Agent problem pioneered by Holmström and Milgrom (1987). The main advantage of my approach is in the structure of the optimal contract which adds a convex component to their optimal linear contract. This provides new opportunities to revisit empirical predictions and studies based off of their linear formulation and to demonstrate how the empirical irregularities may be at least partially explained by this one additional component identified here.

*This 0.5^{th} preliminary, incomplete, rough and flawed draft has been prepared exclusively for the purpose of discussion July at the Purdue University Theory Conference, May 2017. Please do not circulate further or quote.

1 Introduction

The seminal paper by Holmström and Milgrom (1987) established elegantly that more complex contracting problems may well have the simpler solutions. First, by allowing the agent full control over the output distribution as opposed to limit his control to one or a few central parameters, they showed that when it comes to implementing a particular distribution, the principal's hands are tied: the contract that implements a given distribution is unique and thus not subject to the complications of introduced by the standard minimization needed to find the optimal contract that implements an action in more limited choice situations, such as Grossman and Hart (1983). While the contract-action combination that solves the principal's over-arching problem of expected utility maximization of course may still not have a unique solution, finding a contract-action pair that does solve the principal's problem does arguably get significantly simpler when all contract-action pairs are unique.

Second, Holmström and Milgrom (1987) proceeded to demonstrate that when the (principal and) agent has CARA preferences over aggregate consumption net of the aggregate cost of effort here denominated in the units of compensation, a solution to a dynamic multi-period extension of this problem to one where the agent observes past performance prior to his next action is for the principal to implement the same action using the same unique contract each and every period. Moreover, implementing this solution becomes particularly simple: the contract can be written as a linear combination of the aggregate balances of a set of enumeration accounts that track the (relative) frequency of a particular outcome being realized. In the limiting case of continuos time these balances become normal distributed yielding the added bonus of easy to obtain closed form approximate solution to the optimal dynamic contract written on such aggregate performance measures.

The key to the latter part is of course stationarity and this is where the restrictions on the parties preferences are crucial. As Holmström and Milgrom (1987) show, their particular format lends itself to decomposing the principal's multi-period problem into as many one-period problems as there are sub-periods. Moreover, the solution to each of them are independent of the solution to either of the remaining problems and thus is the same as the solution to the single period problem. Without this feature the solution become path dependent and the aggregation (and linearity) result then no longer holds. Those that have exploited the tractability of this framework, weather in its fundamental form or relying on the so-labeled "ad-hoc principle," have thus also been bound by the particular multiplicatively-separable, negative-exponential preference set-up.

While the linearity in normal distributed performance measures is a boon to tractability that has been instrumental in extending P-A models in ways that otherwise are off limits, as always there are no entirely free lunches. Some argue that while, when it comes to the nature of empirically observed contracts, linearity may be a good first (local) approximation, it is not the whole story. Many contracts have significant convex features, such as options for example. Also, the unlimited downside not just in utility space but in terms of the actual cash transfers that the linear contracts entertain in a normally distributed world does not seem to be entirely descriptive either. Wealth concerns and legal limitations may be somewhat incompatible with this from a more descriptive vantage point.

Perhaps unsurprisingly, then, there have been ongoing efforts to develop an alternative framework that, while still offering the coveted tractability, also could provide for the analysis of settings where linearity itself limits the scope of the issues that can be satisfactorily addressed. Edmans and Gabaix (2011), for example, modified the assumptions of Holmström and Milgrom (1987) in several significant ways to this end. First, they allow for any preference representation consistent with A1 in Grossman and Hart (1983).[1] Second, they restrict the underlying uncertainty to a particular class of probability distributions over which the agent has only limited control. Third, the order of play is partially reversed - in their setting the agent observes the state of nature *before* choosing his action. Fourth, the consequences of the agent's action(s) is assumed to be deterministic such that period output is simply the sum of resolved uncertainty and the agents (informed) choice.

[1] The preference representation of Holmström and Milgrom (1987) is clearly a special case hereof.

The first deviation clearly improves generality. The other three are designed to ensure that the control problem remains rich enough that the uniqueness feature of Holmström and Milgrom (1987) is not lost while gaining some control over the nature of equilibrium output distributions. What *is* lost by Edmans and Gabaix (2011), however, is the all important stationarity result. Because of this and because little can be said in their setting without this critical property, they eventually resort to imposing a set of assumptions that guarantees that the most productive action as *always* strictly preferred and thus will always be implemented by the principal.

The proposed benefits of all of this is about the nature of the tractable contract that implements this “most productive action." Such contracts tend to be convex and thus lend themselves to types of inquiry contract linearity may preclude. The drawbacks are not insignificant, however. The reverse order of play is not merely a technicality but a very real and potentially serious limitation of this analysis. If it is to be thought of as economically relevant, it must be that it is a fair characterization of a meaningful sub-set of firms in the economy. However, if firms generally can hire a manager that will be able to observe their resolution of uncertainty perfectly at any given point in time, there must be an outsized role and market for monitors not present in standard P-A set-ups and not considered by Edmans and Gabaix (2011).

In addition, with the strong assumptions needed to guarantee the stationarity in their setting, deviations two to four above, while sufficient, are really not necessary. Holmström and Milgrom (1987) extends easily to alternative preference representations and, in particular, to convex tractable contracts when it can be assumed that the same action is dominant in each period regardless of history. This is of course true whether a particular action is assumed dominant directly or indirectly such as via, for example, the “High Effort Principle,” Edmans and Gabaix (2011). As of now, however, no framework exists that avoids either the multiplicatively separable negative exponential (CARA) specification and the resulting linearity of Holmström and Milgrom (1987) or, in its place, the drawbacks of additional assumptions and restrictions then needed to make the contract tractable. This paper establishes one such arguably straight-forward extension of Holmström and Milgrom (1987) that achieves both.

The observation behind the extension developed here is a relative simple one but one that appears to have been entirely ignored in the literature: CARA somewhat loosely implies that the pecuniary loss (the "risk-premium" associated with a local "fair" permutation in *cash*-space is insensitive to the level of wealth. But, as made precise later, there similarly exists a particular CRRA utility function for which the pecuniary loss of a particular fair permutation in *utility*-space is also independent of wealth. Accordingly, as I will show generally, just as the CARA set-up of Holmström and Milgrom (1987) where cost of effort is denominated in cash leads to optimal effort-stationarity over time, so does this particular CRRA set-up when effort costs are denominated in utiles. The latter is, of course, consistent with the standard additively separable utility function underlying classical papers such as Holmström (1979) a.m.o.

The benefits here are that optimal contracts, while remaining highly "tractable" are strictly convex. Indeed, the model set-up presented and analyzed here extends the original Holmström and Milgrom (1987) framework to such cases without giving purchase on the simplicity of their solution. In particular, convexity and simplicity is achieved without any of the baggage of Edmans and Gabaix (2011). Thus, because it is possible to obtain simple closed form solutions to the contracting problem that invites convexity of optimal aggregate performance-based contracts, the relevant comparative statics that can be obtained here can also be compared and contrasted directly with those obtained from Holmström and Milgrom (1987).

I proceed as follows. First I introduce the model modifications that will preserve all features of Holmström and Milgrom (1987) except the linearity. Second I will introduce the specific structure and establish formally the resulting time- and outcome-independence of the optimal sub-period effort-contract pairing when the optimal contract is represented in utility space. I will then elaborate on the properties of the optimal contract when evaluated in cash compensation space. Finally I conclude by offering some suggestions about potentially interesting applications of this framework.

2 A Single-Period Model

Key to the aggregation result of Holmström and Milgrom (1987) was to first establish (Theorem 2) for their multiplicatively separable CARA preference representation, the solution to the principal's problem is independent of the agent's personal wealth or, equivalently, of the RHS of the standard IR-constraint. This is, of course, critical to the simplicity of the multi-period extension because net wealth accumulated during past periods (or to be accumulated in the future) thus has no impact on the present. The principal's problem for each sub-period therefore can be solved individually as one-period problems and all have the exact same solution.

The approach I take here is similar in terms of sequencing, but the nature of the additively separable preference representation used here leads to a fundamental difference between the single and multi-period models that is instructive. To that end, assume that the principal is risk-neutral and care only about terminal output, Π net of the agent's compensation. The final output depends on the agent's actions during the contracting horizon but is not directly observable/contractible during the relevant time-frame. Instead, the principal observes a set of $N+1$ informative signals, $x_i \in x$, $i = 0, ..., N$. share s_i, The signals here are ordered from lowest to highest and wlog I normalize $x_0 = 0$. The agent's action choice is of full dimension in the sense that he chooses directly the probability of x_i, $p_i \in p$ with the constraint that $p_0 = 1 - \sum_{i=1}^{N} p_i$.

As mentioned previously, the agent here is risk and effort averse as represented by the following additively separable utility function defined over end-of-period consumption, y,

$$H(y, p) = u(y) - v(p),$$

where $u' > 0$ and risk aversion implies $u" < 0$.Using $v_i(p)$ and $v_{ik}(p)$ to denote the partial derivative of $v(p)$ w.r.t. p_i and the cross-partial of $v_i(p)$ w.r.t. p_k respectively, I assume that $v_i(p) > 0$ and $v_{ii}(p) > 0$, $i = 1, ..., N$ while for simplicity I assume that $v_0(p) = v_{ik}(p) = 0$ $\forall p$, $i, k = 1, ..., N$.

With this, it is straight forward to confirm that the uniqueness of the contract that implements a particular p extends readily to the additively separable case considered here. The agent's expected utility under a given contract is given as

$$\sum_{i=0}^{N} u(s(x_i)) p_i - v(p) = R \tag{1}$$

and assuming p is the interior,[2] we have the agent's first-order condition as

$$u(s(x_i)) - u(s(x_0)) = v_i(p) - v_0(p) \tag{2}$$

where uniqueness then follows from strict convexity of $v(p)$. Then using (1) and (2) we have

$$u(s(x_i)) = R + v(p) + (1 - p_i) v_i(p) - \sum_{k \neq i} p_k v_k(p).$$

The principal is kept risk-neutral throughout. Letting $z(.)$ be the inverse of $u(y)$ such that $z(u(y)) \equiv y$, the standard and familiar formulation of the principal's (first-order) one-period problem thus becomes

$$\max_{p} \quad p'x - \sum_{i=0}^{N} p_i z \left(R + v(p) + (1 - p_i) v_i(p) - \sum_{k \neq i} p_k v_k(p) \right)$$

with first-order conditions

$$x_i = (s_i - s_0) + \sum_{i=0}^{N} p_i G' \left(R + v(p) + (1 - p_i) v_i(p) - \sum_{k \neq i} p_k v_k(p) \right) \times \left[(1 - p_i) v_{ii} - \sum_{k \neq i} p_k v_{ki}(p) \right].$$

[2] I will introduce simple assumptions below that guarantees this.

Obviously, then, here the role of R depends on the cross-partials of the cost function. However, for the arguably neutral case where $v_{ki}(p) = 0$, $\forall k \neq i$, all terms on the RHS are increasing in R due to the convexity of $G(\cdot)$. Accordingly, the marginal cost of p_i is then increasing in R for all p_i, $i = 1, ..., N$, and the optimal p is therefore is decreasing in R for *any* concave $u(\cdot)$.

The broader message is, however, that there is no (obvious) counterpart to Theorem 4 in Holmström and Milgrom (1987) in the case of additively separable preferences such as those introduced above: the optimal action generally depends directly on the agent's wealth and his expected utility requirement. While that may at first seem to rule out getting tractability in the multi-period extension outside of out-right exogenously imposing one action always to be the dominant one regardless, it is worth noting that the assumptions in Holmström and Milgrom (1987) is sufficient to guarantee that the optimal contract is a simple solvable function of aggregate performance. They are not necessary, however, as I will proceed to establish.

Before doing so, the following insight may prove instructive. Let w will be some (constant) level of wealth of the agent and consider now offering the agent a (cash) lottery, $\widetilde{\epsilon}$, that is actuarially fair in *utility space*. That is

$$E\left[u\left(w+\widetilde{\epsilon}\right)\right] = u(w).$$

Let then $\widetilde{\xi}$ represent the induced variation in the agent's utility by this lottery so that for any given realization of $\widetilde{\epsilon}$, the corresponding realization of ξ is determined as

$$\xi \equiv u(w+\epsilon) - u(w).$$

Now, since the agent is risk-averse, the exist a strictly positive number, δ, such that

$$E\left[z\left(u(w)+\widetilde{\xi}\right)\right] = w + \delta,$$

where δ then is the actuarial cash value of the expected-utility-neutral lottery $\widetilde{\epsilon}$. Using a Taylor-expansion of the LHS, we have

$$E\left[z\left(u\left(w\right)\right)+\widetilde{\xi}z'\left(u\left(w\right)\right)+\frac{1}{2}\widetilde{\xi}^{2}z"\left(u\left(w\right)\right)\right]=w+\delta$$

or

$$\frac{1}{2}var\left(\widetilde{\xi}\right)z"\left(u\left(w\right)\right)=\delta.$$

Notice that the actuarial value depends only on $z"\left(\cdot\right).$ Accordingly, iff $z"=k>0$, where k is some constant and the last inequality follows from the agent being strictly risk-averse, δ is independent of the individual's level of (expected) utility. Then, integrating twice over k recovers

$$z\left(u\left(w\right)\right)=a+b\left(u\left(w\right)\right)^{2}, \quad b>0,$$

so that

$$u\left(w\right)=\left(\frac{w-a}{b}\right)^{1/2}, \qquad a<w.$$

In other words, the expected cash value of a lottery acceptable to a risk-averse agent is independent of the agent's current wealth iff that agent has a power utility function where the power is one-half. As I will proceed to show, this feature provides the a key component for the optimal action to be time- and outcome invarriant in the case of additively separable preferences. For the remainder of this paper I therefore rely on the following specific preference representation for the agent:

$$H\left(y,p\right)=\left(\textstyle\sum_{t=1}^{T}y_{t}\right)^{1/2}-\textstyle\sum_{t=1}^{T}v\left(p_{t}\right), \quad T>1. \tag{3}$$

3 Multi-Period Extension

To extend the single-period model to a continuos time/outcome version, I proceed in two steps: I first identify the fundamental problem involved in multi-period extensions of the basic model when deviating from the multiplicatively separable CARA preference structure of Holmstrom and Milgrom (1987). I then identify a production and measurement structure set-up that will suffice in terms of eliminating the problem at hand while at the same time not changing the nature of the principal's dynamic problem while at the same time enriching the nature of the contract that solves the principal's problem.

The first point to be made here is that unlike in the multiplicatively-separable case, the fact that cost of implementing a particular action is also wealth independent in the additively separable "square root" case pursued here does *not* logically imply that effort is also necessarily time and outcome independent here.

Lemma 1 *For the multi-period version of the model introduced here, it is not the optimal strategy to make the agent's future actions independent of past outcome realizations.*

Proof. Consider first the solution to the principal's problem here with the additional restriction that for some $t = 1, ...m - 1$, p_τ, $\tau = t + 1, ..., m$ be independent of X_t. Let $\overline{p}_{t+1}$ denote the the action that solves the principal's thus restricted problem for period $t + 1$. Based on Lemma 1 it is straight forward to verify that $u\left(\overline{s}_t\left(X_t,\right) + \widehat{s}_2\left(x_{2i}\right)\right) - u\left(\widehat{s}_1\left(x_{10}\right) + \widehat{s}_2\left(x_{2i}\right)\right) = u\left(\widehat{s}_1\left(x_{1i}\right) + \widehat{s}_2\left(x_{21}\right)\right) - u\left(\widehat{s}_1\left(x_{1i}\right) + \widehat{s}_2\left(x_{20}\right)\right), i = 0, 1$. Now consider modifying this contract so that $\widetilde{s}_2\left(x_{10}, x_{20}\right) \equiv \widehat{s}_2\left(x_{20}\right) - \delta s_1$, $\widetilde{s}_2\left(x_{10}, x_{21}\right) \equiv \widehat{s}_2\left(x_{21}\right) + \delta s_1$, $\widetilde{s}_2\left(x_{11}, x_{20}\right) \equiv \widehat{s}_2\left(x_{20}\right) + \delta s_1$ and $\widetilde{s}_2\left(x_{11}, x_{21}\right) \equiv \widehat{s}_2\left(x_{21}\right) + \delta s_1$, where s_1 is a positive constant and $\delta \in R$. First verify that the derivative of the expected compensation w.r.t. δ evaluated at $\delta = 0$ is also zero. Second, verify that the derivative of expected second period effort w.r.t. δ is also zero at $\delta = 0$. Finally verify that the derivative of first period effort w.r.t. δ evaluated at $\delta = 0$ is strictly positive. By adding some outcome based variation in the restricted contract $\widehat{s}_t\left(x_t\right)$, the principal can achieve the same first period incentives at a lower cost as the added incentives provided by second period work-load variations allow the principal to reduce the compensation risk associated with first period output. ■

While this difference between the multiplicative and additive separable representations is perhaps interesting in its own right, the desire to make future actions conditioned on past outcomes is of course detrimental to the ability to write the optimal contract as a simple function of aggregate performance. At first glance it would seem that adding structure to prevent this natural economic demand for path dependent pay and economic activity from playing out would be somewhat unappealing and not that much different from assuming that there is a dominant action to be implemented each and every sub-period. A more reasonable way of looking at the difference between this and the result of Holmstrom and Milgrom (1987) is to note that production volatility generally is considered as a negative for a variety of practical reasons.

One clear such reason is because it requires costly slack capacity being kept in reserve. While the cost of capacity is not relevant in their formulation and thus can be safely ignored, it is the implicit assumption that the (marginal) cost of capacity is zero (or sufficiently low) that is one reason the optimal solution is path-dependent here. While adding "sufficiently large" capacity costs here could recover "effort stationarity," I'll pursue a much simple approach that achieves the same objective here. Rather than acquiring capacity, I'll simply assume that the principal at the start of the contracting horizon must acquire a (for the horizon) fixed technology that must be compatible, as to be defined in the following, with the action(s) chosen by the agent.

As in Holmstrom and Milgrom (1987), first I make a distinction between value creation and measures thereof. The publicly observable construct x introduced above is thus, hereafter, simply a measure that reflect value creating activities like an accounting measure. The actual value create in each sub-period, π_t hereafter, is not assumed to be observable during the relevant horizon. It's expected value, Π_t, however, depends on both the firms chosen technology. Specifically, letting $\phi_t \subseteq [0, 1]$ be the pre-determined technology for period t. I then assume that $d\Pi_t/dp_t = d\Pi_t/d\phi_t > 0$ for $p_t = \phi_t$ and zero otherwise. In other words, if the principal wishes to increase p_t he must also install a higher ϕ_t at the start of the horizon. With this added structure I can state the following Lemma:

Lemma 2 *Let the agent's utility function be given by (3). Then the optimal solution is for the principal to the same $\{p_t, \phi_t\}$ pair $\forall t = 1, ..., T$.*

Proof. *When the action cannot be contingent of realized performance, absence of wealth-effects is sufficient for the optimal action to be time independent as well.* ■

At first glance the assumption on the link of the technology, productive action and output may seems somewhat heavy-handed. There are several counter-points to that, however. First note that it is not sufficient to ensure time-independence of the optimal effort alone - this also takes the properties of the utility function 3. In Holmstrom and Milgrom (1987), their assumption that the agent's (convex) cost of effort is denominated in cash rules out any potential benefits of using future actions to incentivize current actions. Adding then CARA ensures that the principal implements the same particular action in every sub-period. In the additively separable setting I analyze here, as long as *some* feature of the model, such as a "sufficiently" convex capacity cost or, much simpler, just having to pre-comitt to the type of technology I rely on, renders outcome dependent actions unattractive, time independence is assured by 3.

Second, from an applied perspective, and the exercise here is aimed at producing applicability, nothing is lost relative to Holmstrom and Milgrom (1987). Both the optimal action and, thus, the properties of the optimal contract, of course, still depends on the particular exogenous properties of the problem and as such is amenable to the full set of comparative statics. That contrasts sharply with the approach of Edmans and Gabaix (2011) that requires a singular preferred action be present regardless of the preferences of the agent(!) and the parameters of the environment in which the agent is operating. From these vantage points, the assumption that the actions match the technology/strategy chosen by the principal, seems quite benign.[3]

4 Brownian Approximation and Applications

The contract derived in the previous section has the same qualitative property as that obtained by Holmström and Milgrom (1987): it can be written as a simple function of aggregate account balances at the end of the contracting horizon. The only difference is that in the case analyzed here the contract is linear in the account balances in *utility space* rather than in terms of actual cash disbursements. As a result, here the contract is actually convex in cash space due to the convexity of the inverse utility function. Importantly, the particular utility specification behind the results obtained here lend itself well to obtaining simple-to-calculate closed form convex contracts using the same "Brownian Approximation" pursued in the latter part of Holmström and Milgrom (1987).

For simplicity I here I continue to focus on the approximation of the simple binomial case, as extensions are straightforward and left to the reader. To extend the single-period model to a continuos time/outcome version, I also rely heavily on the methodology provided by Hellwig and Schmidt (2002) with only minor departures to ensure consistency with the alternative preference representation and added structure I have introduced above. I proceed in two steps: I first subdivide the one period model into a discrete time version with m sub-periods of length $\Delta = 1/m$ to identify specific conditions under which the optimal action for each sub-period is the same. Subsequently I proceed to show that, in the limit as $\Delta \to 0$, becomes indistinguishable from a one-shot model where the principal incentivizes the agent to take a costly action once that determines the mean and variance of an (approximately) normally distributed performance measure via a contract that is a linear function of this performance measure in utility space and thus convex in cash-space.

Since I here I restrict attention to the case of $N = 1$ it allows me to let $p_1 \equiv p$ and $p_0 \equiv$

[3] Note again that here the assumptions are sufficient to deliver the desired tractability In both Holmstrom and Milgrom (1987) and Edmans and Gabaix (2011), the assumptions are necessary.

$1 - p$. The fundamental uncertainty embedded in output is characterized by the exogenous probability $\widehat{p}$ (> 0) which is the probability of a positive outcome when no (unobservable) effort is provided. For the base-case where $m = 1$ let $\kappa \equiv x_1 - x_0$, where $x_1 > \widehat{\mu} > x_0$, where $\widehat{\mu}$ is the the expected output associated with $\widehat{p}$. Denote by $\widehat{\mu}^{\Delta} \equiv \Delta\widehat{\mu}$, $x_1^{\Delta} \equiv \Delta^{1/2}x_1$ and $x_0^{\Delta} \equiv \Delta^{1/2}x_0$, so that $\kappa^{\Delta} = \Delta^{1/2}\kappa$, the values of the corresponding variables in a given sub-period of length Δ.Let the expected output associated with $\widehat{p}$ That is, $\widehat{p} \equiv (\widehat{\mu} - x_0)/\kappa$. From the prior sections it already follows that as in Holmström and Milgrom (1987) an optimal solution to the principal's problem here is to induce the agent to select the same action, denoted p^{Δ} $(\geq \widehat{p})$, in each of the m sub-periods.

With only a slight departure from Hellwig and Schmidt (2002), the personal cost of p^{Δ} to the agent is here assumed to take the form

$$c^{\Delta}\left(p^{\Delta}\right) = \Delta c\left(\frac{p^{\Delta} - \widehat{p}}{\Delta^{1/2}}\right).$$ [4]

This ensures that the cost of enhancing expected performance *beyond* $\widehat{\mu}$ is independent of the number of sub-periods, m, which of course also implies that the cost of implementing a particular μ is independent of m here. Moreover, suppose the principal wants to implement a particular $\mu \geq \widehat{\mu}$ for the entire horizon by implementing $\mu^{\Delta} \equiv \Delta\mu$ in each of the m sub-periods, he must implement in each sub-period $p^{\Delta} = \widehat{p} + \Delta^{1/2}(p - \widehat{p})$. The cost of that per sub-period is simply, then, $\Delta c(p - \widehat{p})$ and the total cost for the entire horizon thus is simply equal to $c(p - \widehat{p})$ independent of m.[5] Further, for the principal to implement a particular p^{Δ} or, equivalently, a particular μ^{Δ} in any given sub-period, the contract must be such that

[4] Hellwig and Schmidt (2002) rely on the cost-function

$$c^{\Delta}\left(p^{\Delta}\right) = \Delta c\left(\widehat{p} + \frac{p^{\Delta} - \widehat{p}}{\Delta^{1/2}}\right).$$

[5] Please note that the notation $\mu^{\Delta} \equiv \Delta\mu$ is a deviation from the similar notation in Hellwig and Schmidt (2002). In their notation, μ^{Δ} represents the same construct as μ does in this paper.

the desired p^{Δ} solves

$$u\left(x_1^{\Delta}\right)-\left(x_0^{\Delta}\right)=\Delta^{1/2}c'\left(\frac{p^{\Delta}-\widehat{p}}{\Delta^{1/2}}\right).$$

Since the marginal benefit in each sub-period is easily verified to be simply $\Delta^{1/2}\kappa$ here, the (first-best) marginal cost/benefit trade-off is also independent of the number of sub-periods with this particular structure.

Consider now the utility-performance-sensitivity in any given sub period,

$$\beta^{\Delta}\equiv\left(u\left(x_1^{\Delta}\right)-\left(x_0^{\Delta}\right)\right)/\Delta^{1/2}\kappa=c'/\kappa.$$

Since β^{Δ} thus is also independent of m I'll drop the superscript and simply refer to this construct as β. Now, let ω represent the agent's reservation utility and let, with a bit abuse of notation convention, $\mu^{\nabla}\equiv\mu-\widehat{\mu}$ The second-best cost of implementing a particular μ over and above the first-best cost of procuring μ directly. then can be found as

$$\begin{aligned}\Omega_{\mu} &\equiv E\left[\left(w+c\left(\mu^{\nabla}\right)+\widetilde{\xi}_{\mu}\right)^{2}\right]-\left(w+c\left(\mu^{\nabla}\right)\right)^{2}\\ &= E\left[\widetilde{\xi}_{\mu}^{2}\right]=\sigma_{\mu}^{2},\end{aligned}$$

where $\widetilde{\epsilon}_{\mu}$ is the mean-zero variation in utility-space that (uniquely) implements μ and σ_{μ}^{2} is the variance of $\widetilde{\xi}_{\mu}$. With $\sigma_{\mu}^{2}=\beta^{2}\sigma_{x}^{2}$, we thus have

$$\begin{aligned}\Omega_{\mu} &= \left[c'\left(\mu^{\nabla}\right)\right]^{2}\widehat{p}\left(1-\widehat{p}\right)\\ &\equiv \left[c'\left(\mu^{\nabla}\right)\right]^{2}\widehat{\sigma}_{x}^{2}.\end{aligned}$$

independent of m as well. The (again, risk neutral) principal's problem then simply reduces to solving

$$\max_{\mu}\quad \Pi\left(\mu^{\nabla}\right)-\left(w+c\left(\mu^{\nabla}\right)\right)^{2}-\left[c'\left(\mu^{\nabla}\right)\right]^{2}\widehat{\sigma}_{x}^{2}.$$

This implies that μ^{∇} solves

$$\frac{\Pi'\left(\mu^{\nabla}\right)}{2\left(w+c\left(\mu^{\nabla}\right)+c''\left(\mu^{\nabla}\right)\widehat{\sigma}_x^2\right)}=c'\left(\mu^{\nabla}\right), \tag{4}$$

and that, from (Z),

$$\beta^{\nabla}=\frac{\Pi'\left(\mu^{\nabla}\right)}{2\kappa\left(w+c\left(\mu^{\nabla}\right)+c''\left(\mu^{\nabla}\right)\widehat{\sigma}_x^2\right)}$$

so that the fixed part of the agent's utility, $\alpha^{\nabla}\equiv\alpha\left(\mu^{\nabla},w\right)$ hereafter, is determined as

$$\alpha^{\nabla}=w+c\left(\mu^{\nabla}\right)-\frac{\mu^{\nabla}\Pi'\left(\mu^{\nabla}\right)}{2\kappa\left(w+c\left(\mu^{\nabla}\right)+c''\left(\mu^{\nabla}\right)\widehat{\sigma}_x^2\right)}.$$

5 Applications

The point of developing this companion framework to Holmstrom and Milgrom (1987) is to, without loosing the indispensable tractability, provide slightly richer contracts that may be helpful in informing some of the discrepancies between the normative implications of the linear contracts and the empirical evidence accumulated. Furthermore, utilizing the approach of Hellwig and Schmidt (2002) to achieve the effort-independent variance result directly as the limiting case of the multinomial model adds additional structure that also have specific implications for the kind of predictions one can obtain here.

To facilitate specificity and tractability, as the base-line let $\Pi\left(\mu^{\nabla}\right)=\left(\mu^{\nabla}\right)^2/2=\left(\kappa p^{\nabla}\right)^2/2$ and $c\left(\mu^{\nabla}\right)=\left(p^{\nabla}\right)^2/2$. Then, using (FOC) we have

$$\frac{\kappa^2 p^{\nabla}}{\left(w+\left(p^{\nabla}\right)^2/2+\widehat{\sigma}_x^2\right)}=p^{\nabla}$$

or

$$p^{\nabla}=\sqrt{2\left(\kappa^2-w-\widehat{\sigma}_x^2\right)} \tag{5}$$

so that

$$\mu^{\nabla} = \kappa\sqrt{2\left(\kappa^2 - w - \widehat{\sigma}_x^2\right)} \tag{6}$$

while

$$\beta^{\nabla} = \frac{\sqrt{2\left(\kappa^2 - w - \widehat{\sigma}_x^2\right)}}{\kappa} \tag{7}$$

and

$$\begin{aligned}\alpha^{\nabla} &= w + p^2/2 - \beta\mu \\ &= w + \left(\kappa^2 - w - \widehat{\sigma}_x^2\right) - 2\left(\kappa^2 - w - \widehat{\sigma}_x^2\right) \\ &= 2w + \widehat{\sigma}_x^2 - \kappa^2. \end{aligned} \tag{8}$$

This slightly richer yet still relatively simple structure leads to a number of straight forward but arguably more subtle comparative statics than the standard ones based off of Holmström and Milgrom (1987).

5.1 Determinants of Convexity

An obvious issue to be investigated is the driver(s) of the optimal contract's "convexity." The way I'll proceed here is to look specifically at the role of productivity, κ, in determining the weights on the linear and on the convex pieces, $2\alpha\beta$ and β^2 respectively. To streamline this further, define $\overline{p} \equiv 1 - \widehat{p}$, let $w = \overline{p}^2/2$ and $\kappa \in \left(\overline{p}^2/2 + \widehat{\sigma}_x^2; \overline{p}^2 + \widehat{\sigma}_x^2\right)$. The lower bound on κ result from $p \geq 0$ with the upper bound imposed by $p \leq 1$. While this obviously is an extreme set-up with all the down sides that comprises, it has the advantage of delivering some crispness that could be considered helpful.

To see this note that at κ close to its lower bound here, β is (close to zero) while α is (close to) w. Moreover, we have

$$\frac{d\beta}{d\kappa} \to \infty$$

as κ approaches its lower bound. Accordingly, at the low end of the feasible productivity range, increased incentives in response to increased productivity is provided in the form of the linear component. The idea here would be that in (very) low-skill/productivity situations, incentives are provided in the from of piece rates: compensating based on the numbers produced, which does not seem ad odds with casual empiricism. Sharecroppers, newsboys, table servers all seem to fit with this.

On the other end of the productivity spectrum it is an entirely different story here, however. With the parametrization here, $\alpha \to 0$ and since

$$\frac{d\alpha}{d\kappa} < 0,$$

the weight on the linear component, $2\alpha\beta$, is strictly decreasing as κ increases towards its upper bound. Both β and $d\beta/d\kappa$ are, in contrast both strictly positive, so that at high levels of productivity, higher productivity implies substituting out of linear incentives and into convex ones. This seems, at least on the surface, consistent with options being more popular for top executives and more prevalent in strong growth environments such as the tech sector of the economy.

5.2 The Role of "Risk."

This slightly richer yet still relatively simple structure leads to a number of straight forward but arguably more subtle comparative statics than the standard ones based off of Holmström and Milgrom (1987). For example, it is immediately obvious from the above first-order condition that this set-up yields the same predicted inverse relation between "risk" and "effort" as does that of Holmström and Milgrom (1987). Consider thus the properties of the optimal cash compensation contract

$$s(X) = \alpha^2 + 2\alpha\beta X + \beta^2 X^2$$

Note first that pay-for-performance sensitivity

$$\frac{ds(X)}{dX} = 2\alpha\beta + 2\beta^2 X$$

and the expected (or average) PPS is thus

$$E\left[\frac{ds(X)}{dX}\right] = 2\alpha\beta + 2(\beta\mu)\beta$$

which obviously here depends on σ^2 only through its effect on μ since here α, β and μ are mechanically linked in equilibrium. Specifically, from the Principal's problem, we have

$$\frac{d\beta}{d\widehat{\sigma}_x^2} < 0$$

and

$$\frac{d(\beta\mu)}{d\widehat{\sigma}_x^2} = -2\frac{d\alpha}{d\widehat{\sigma}_x^2} = 1.$$

Accordingly,

$$\frac{dE\left[\frac{ds(X)}{dX}\right]}{d\widehat{\sigma}_x^2} < 0$$

which is, of course, consistent with the prediction provided by Holmstrom and Milgrom (1987). Estimating the PPS by regressing the change in compensation on changes in compensation, and then regressing the estimated PPS on "risk" should thus lead to the predicted inverse relation as long as there are no controls for average performance.

This is, of course, the equivalent of regressing compensation, $s(X)$, on performance X. Let $A \equiv \alpha^2$, $B \equiv 2\alpha\beta$ and $C \equiv \beta^2$. The regression coefficients, φ_0 and φ_X from a standard OLS-regression of compensation on performance are then the solution to the following problem:

$$\min_{\delta_0,\delta_X} \quad E\left[\left(A + BX + CX^2 - \varphi_0 - \varphi_X X\right)^2\right]$$

which has first-order conditions

$$A - \varphi_0 + (B - \varphi_X)\,\mu + C\left(\sigma^2 + \mu^2\right) = 0,$$
$$(A - \varphi_0)\,\mu + (B - \varphi_X)\left(\sigma^2 + \mu^2\right) + C\left(3\mu\sigma^2 + \mu^3\right) = 0,$$

so that

$$\varphi_X = 2C\mu + B.$$

Note, however, that κ^2 enters in the expressions for α, β and μ exactly the same way as $\widehat{\sigma}_x^2$. Without any further analysis it immediately clear from the above that$d\varphi_X/d\kappa^2 = -d\varphi_X/d\widehat{\sigma}_x^2 > 0$. Accordingly, since $\sigma_x^2 = \kappa^2\widehat{\sigma}_x^2$, if the PPS is estimated as the coefficient on performance from an OLS-regression of compensation on performance, it may just as well increase as decrease in σ_x^2. It should thus come at no surprise either, that the substantial empirical literature that focus on this relation has been inconclusive.

Suppose, however, it is possible to empirically separate the portion of risk that is fundamental from the part that is proportional to the productivity (squared), the model does predict a negative relation between fundamental risk and PPS. Moreover, there is an interaction-effect between the two pieces of overall risk in terms of the convexity of the optimal contract. Using the parameter values of the prior application, it is easily verified that for very high values of κ, $d\alpha\beta/d\widehat{\sigma}_x^2 > 0$ while $d\beta^2/d\widehat{\sigma}_x^2 < 0$. In contrast, $d\alpha\beta/d\widehat{\sigma}_x^2$ and $d\beta^2/d\widehat{\sigma}_x^2$ are both strictly negative for low values of κ. In other words, while increasing fundamental risk decreases (average) PPS regardless, in high productivity environments the linear component actually increases while at the low productivity side both risky components are scaled back.

5.3 Performance Measure Design

The final implication pursued here is that of the of an optimal performance measure. What is of particular interest here is the dynamic properties of the aggregate performance measure X. The main insight here is that while the contract can be written on aggregate performance, the *efficiency* of the contract depends on the time-series properties of performance holding all else constant. That is, performance measures that are indistinguishable on all dimensions in the aggregate are not the same in their dynamic properties and therefore not from the perspective of the second-best. Aggregates are sufficient, but not all aggregates are "sufficiently" good.

For expositional ease, consider the semi-parametric mode introduced above as a benchmark, with a bit of additional structure to the specifics of the performance measure, x. Specifically, introduce a separate productivity parameter, say γ, so that the probability of x_1^Δ, $Pr\left(x_1^\Delta|\widehat{p}, p, \gamma\right) \equiv p_\gamma^\Delta = \widehat{p} + \Delta^{1/2}\gamma\left(p - \widehat{p}\right)$. Then consider any class of performance measures, $\overline{X}$, for which $i)$ the variance $\kappa^2\widehat{p}\left(1-\widehat{p}\right) = \underline{\sigma}^2$, $ii)$ the sensitivity to the agent's action $d\mu^\Delta/dp^\Delta = \underline{\delta}$, and the mean is the same given the agent's action which, given $ii)$, is implied by $iii)$ $\widehat{\mu} = \underline{\widehat{\mu}}$. Notice that $i)$ implies

$$\kappa = \sqrt{\frac{\underline{\sigma}_x^2}{\widehat{p}\left(1-\widehat{p}\right)}},$$

$ii)$ implies

$$\gamma = \underline{\delta}/\kappa,$$

while $iii)$ implies

$$x_0 = \underline{\widehat{\mu}} - \widehat{p}\kappa.$$

In other words while $i)$ imposes an implicit relation between κ and $\widehat{p}$, the free parameters γ and x_0 can always be adjusted to satisfy $ii)$ and $iii)$. Importantly, it is easily verified that as a result of this, changing $\widehat{p}$ while satisfying $i)$, $ii)$ and $iii)$ does not affect the structure of the solution to the principal's problem as represented by (5)-(8) here. What does change, however, is $\widehat{\sigma}_x^2$ which is maximized for $\widehat{p} = .5$ and is approaching zero as $|\widehat{p} - .5| \rightarrow .5$. This is important for two reasons: first while $\widehat{\sigma}_x^2$ is a fundamental driver of contracting efficiency variations in $\widehat{\sigma}_x^2$ is not easily identified identified by the properties aggregate performance.

Where $\widehat{\sigma}_x^2$ is more easily detected is in the time-series properties of performance: if performance evolves symmetrically around its mean $\widehat{\sigma}_x^2$ is (close to being) at it's max. On the other hand a measure with more frequent smaller gains and fewer but bigger losses or vice versa, fewer but larger gains but more frequent smaller losses indicate lower values of $\widehat{\sigma}_x^2$.

This seems particularly interesting from an accounting perspective. Under something akin to clean surplus accounting the properties of *aggregate* earnings over a reasonable horizon is fundamentally unaffected by the particular approach to accounting measurement chosen by the firm while the short run properties by definition are directly affected. In terms of this model, this is exactly the consequence of being able to choose a $\widehat{p}$ without being able to alter the *aggregate* mean, variance or sensitivity to effort over a reasonable horizon. Moreover, choices of $\widehat{p} \neq .5$ have reasonably straight forward interpretation in terms of measurement bias: few "big" write-offs in conjunction with being able to produce smaller positive surprises over longer horizons (think "big bath" behavior in the extreme) is generally considered a feature of aggressive reporting while making overly cautious lone loss provisions for many periods that on average then all reverse in the period of repayment is generally considered conservative accounting. In the model, either is strictly preferred to the neutral case.

This, of course also brings about some suggestions for empirical inquiry. While γ generally is unobservable, holding aggregate (average) performance and performance variance constant, the degree of asymmetries in periodic performance (gains relative to losses and positive relative to negative short-run returns, for example) provides a measure of $\widehat{\sigma}_x^2$. This, in turn, provide a link between the time-series properties of measured performance, the convexity of contracts and their average PPS. While obviously only very simple and suggestive, these examples at least suggest that the framework here proposed, may have some promise in generating specific predictions along a number of such dimensions.

6 Conclusion

Quite a bit too early for that. The paper is still very much work-in-process.