

A Note on Moral Hazard and Linear Compensation Schemes

Xiao Yu Wang*

July 18, 2013

Abstract

This note identifies a moral hazard environment in which a piecewise linear compensation scheme is optimal. Both the principal and the agent have CARA utility, mean output is increasing in the agent's noncontractible input, and output is distributed according to a Laplace distribution, which resembles a normal distribution (e.g. it is symmetric about the mean), but has fatter tails. The key property of the Laplace distribution is that the likelihood ratio is a piecewise constant, where the discontinuity occurs at the mean.

The value of this approach is twofold: first, a tractable, empirically-observed wage scheme emerges as the equilibrium in a simple static contracting model. Second, the optimal piecewise linear scheme cleanly separates insurance and incentive provision. The linearity at output levels away from the mean captures insurance, while the jump at the mean captures incentive provision. Hence, this model is well-suited for studying a wide variety of principal-agent problems in risky environments subject to moral hazard, such as the effect of risk and moral hazard considerations on employment relationships in developing economies.

JEL Classification Codes: C70, D86, O12

Keywords: Principal-Agent Problems; Moral Hazard; Linear Incentive Schemes; Insurance; Incentives

**I thank Gabriel Carroll, Sebastian Di Tella, Bengt Holmstrom, and Juan Pablo Xandri for helpful comments. Any errors are my own. Support from the National Science Foundation is gratefully acknowledged.*

Duke University. Email: xy.wang@duke.edu

1 Introduction

Applications of the simple static principal-agent problem under one-sided moral hazard date back to Zeckhauser's discussion of insurance policies in 1970 and Stiglitz's discussion of the institution of sharecropping in 1974, and this model has continued to be popular over the years. The key tradeoff of incentive provision and insurance provision by the equilibrium wage has inspired much research. For instance, Stiglitz (1974) suggested that tenant farmers are observed to pay their landlords rent as a share of realized harvest, rather than as a fixed amount, precisely because missing formal institutions imply that tenant farmers can be protected from risk only informally, through their rental contracts with their landlords. Because of the simultaneous need to provide incentives, however, landlords can only partially insure their tenants, and this results in the share contract.

While the standard principal-agent moral hazard model appears to have the advantage of being a simple and reasonable starting point for thinking about many interesting questions and settings, it can actually be quite challenging to solve with full rigor. In particular, the agent's global incentive compatibility constraint is often difficult to work with. The traditional approach is to replace this global IC constraint with its first-order condition, and solve the first-order problem using standard Lagrangean methods. This method is commonly called *the first-order approach*.

However, replacing the global IC condition with its first-order condition is a nontrivial substitution. A number of papers have sought to provide conditions for the validity of the first-order approach. Mirrlees (1979) and Rogerson (1983) presented the well-known pair of sufficient conditions for a model with a risk-averse principal and a risk-averse agent, the "monotone likelihood ratio" property (MLRP) and the "convexity of the distribution function" condition (CDFC). While the monotone likelihood ratio is a natural way of capturing the idea that higher effort should have a positive impact on output production, the convexity of the distribution function is prohibitively stringent. Virtually no commonly-used distributions in statistics or economics satisfy CDFC. Grossman and Hart (1983) proposed an alternative procedure by which the first-order approach could be avoided: they calculate the cost to the principal of the incentive scheme that would be used to implement each possible action of the agent, and then select the scheme that best balances costs with benefits. However, while elegant, this procedure was not easy to use in practice. Since the classic moral hazard model is largely a tool of applied theory, Jewitt (1988) returned to the first-order approach, but sought to relax the convexity of distribution function condition. Jewitt (1988) considers a setting with a risk-neutral principal and a risk-averse agent¹, where the agent's utility is additively separable in consumption and effort, and is able to somewhat relax the convexity of distribution function condition. Despite many efforts to identify a more convenient set of sufficient conditions for the validity of the first-order approach, a significant improvement over MLRP and CDFC for the contracting problem between a risk-averse principal and a risk-averse agent has not

¹Jewitt (1987) considers the case where both the principal and the agent are risk-averse, but a myriad of further assumptions is needed to gain traction on this problem, and the difficulty boils down to signing a Lagrange multiplier. So, no clean conditions emerge.

yet been found.

A related challenge of working with the moral hazard framework is that an overwhelming fraction of contracts observed in real-world settings are linear. However, linear contracts arise as the optimum in only a handful of specialized moral hazard models. Mirrlees famously showed that, when risk-averse principals and agents have CARA utility and output is normally-distributed, the linear contract is far from optimal, and in fact there exists a contract which arbitrarily approximates the first-best. This is a consequence of the infinite precision of a tail realization of output of a normal distribution of the effort exerted, and of the possibility of arbitrarily large punishments, but it is nevertheless bad news for theoretical predictions of linear contracting.

The environment of Holmstrom and Milgrom (1987) is perhaps the best-known environment in which linear contracts arise in equilibrium. They consider the problem of a risk-averse principal and agent with CARA utility, where the key to linearity in the discrete environment is a dynamic setting where the principal observes output at a coarser level than the agent. More specifically, the agent chooses an effort level in each period and observes the output that is realized, while the principal only observes the aggregate output produced at the end of the game, and pays based on that quantity.

However, there are important economic environments which are not well-described by the Holmstrom and Milgrom setup. For example, the intuition of the story is inherently dynamic, yet it is often used incongruously to justify the use of linear contracts in static models. Furthermore, in many settings, it is unrealistic to think that the agent observes the continuous output feedback of his private actions, and that moreover the principal cannot. More abstractly, it seems that contracts with linear components may arise for reasons unrelated to the principal wishing to apply the same incentive pressure regardless of the agent's past performance.

The upshot of all this is that we do not yet have a static, bare bones model of a moral hazard problem with a risk-averse principal and risk-averse agent, where output is distributed according to some natural but nonconvex distribution function (for example, the normal distribution is not convex), and where the equilibrium wage can be characterized rigorously without further assumptions on the environment. It would be particularly helpful if such a model yielded a simple contracting form in equilibrium, corresponding to empirical observations.

Why is it important to identify such a model? Many settings of economic interest can be captured by a framework with these elements. For example, principal-agent type relationships abound in developing countries, where individuals are risk-averse and heterogeneous in their risk-aversion, and formal insurance and credit markets are often missing. A classic example is sharecropping (Stiglitz (1974)). A risk-averse landlord hires a risk-averse farmer, and cannot perfectly monitor the farmer's use of inputs. In the absence of formal credit and insurance institutions, the landowner and farmer incorporate insurance provision as well as incentive provision into their rental contract. Thus, the institution of sharecropping serves an informal insurance purpose.

In this note, I provide a simple static model of a contracting relationship subject to one-sided moral hazard between a risk-averse principal and a risk-averse agent, both of whom have CARA

utility. Output is distributed according to a Laplace distribution, which resembles a normal distribution (e.g. it is symmetric about its mean), but, loosely, has fatter tails. The key characteristic of the Laplace distribution is that its likelihood ratio is a piecewise constant, with discontinuity at the mean.

I show that, although none of the known sets of sufficient conditions for the validity of the first-order approach can be applied to this setup, the first-order approach is nevertheless valid. I then show that the equilibrium can be fully characterized, and moreover, that the equilibrium wage is piecewise linear, a tractable and empirically-observed payment schedule.

The contribution of this note is therefore to move away from identifying sufficient conditions for the validity of the first-order approach in more general moral hazard frameworks, and to focus instead on characterizing the behavior of a more specialized but reasonable framework which can be used as a workhorse for basic moral hazard problems in risky environments when both the principal and the agent are risk-averse.

The rest of this note proceeds as follows. In the next two sections, I set up the model and work through the solution. I then discuss an application of this framework. Finally, I conclude.

2 The Model

The framework consists of the following elements:


The principal and the agent: both the principal r_1 and the agent r_2 are risk-averse with CARA utility $u(x) = -e^{-rx}$, $r > 0$. The principal owns one unit of physical capital but has infinite marginal cost of effort, while the agent owns no physical capital, but has a finite marginal cost of effort. The cost of effort for all agents is $c(a)$, $c(a) > 0$, $c'(a) > 0$, $c''(a) > 0$. The agent has an exogenously-given outside option denoted by $-e^{-v}$.

The risky environment: A principal-agent partnership can only produce positive output if one unit of physical capital is combined with human capital. Output is given by $R|a = \gamma a + \varepsilon$, where the riskiness of the environment is captured by $\varepsilon \sim \text{Laplace}(0, V)$, which has support $(-\infty, \infty)$. Hence, the effort of an agent, a , increases the mean but doesn't affect the variance of returns.

The Laplace distribution is a continuous probability distribution with location and scale parameters $\mu \in \mathbb{R}$ and $b \in \mathbb{R}^+$, respectively. The pdf of a random variable $X \sim \text{Laplace}(\mu, b)$ is:

$$f(x|\mu, b) = \begin{cases} \frac{1}{2b} e^{-\frac{(\mu-x)}{b}}, & x < \mu \\ \frac{1}{2b} e^{-\frac{(x-\mu)}{b}}, & x \geq \mu \end{cases}$$

The following graphic illustrates the pdf of the Laplace distribution for a variety of parameters (μ, b) :



Note that this resembles two back-to-back exponential distributions. In fact, if $Y \sim Laplace(0, V)$, then $|Y| \sim \exp(\frac{1}{V})$.

Information and commitment: The principal and the agent know each other's risk types. However, the agent's effort is not observable and not contractible.

A given principal and agent pair (r_1, r_2) observes the realized output of their partnership, and is able to commit to a return-contingent sharing rule $s(R)$, where R is realized output.

The equilibrium: An equilibrium consists of a wage $s(R)$ set by the principal satisfying the constraints of the problem, such that any other choice of feasible sharing rule by the principal would leave her worse off, as well as an effort level a chosen by the agent such that any other choice of effort would leave the agent worse off.

3 The Solution

The principal r_1 chooses wage $s(R)$ for the agent r_2 by solving the following problem:

$$\max_{s(R)} \int_{-\infty}^{\gamma a} -e^{-r_1(R-s(R))} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \int_{\gamma a}^{\infty} -e^{-r_1(R-s(R))} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR$$

such that:

$$\begin{aligned} (IR) : & \int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \\ & + \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR \geq -e^{-v} \\ (IC) : & a \in \arg \max_{a \in (0, \infty)} \int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \\ & + \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR \end{aligned}$$

It is clear that the IR constraint binds in equilibrium.

Since effort a is chosen from an open interval, all global maxima are stationary points, if they exist. Hence, the first-order problem (the problem wherein the global IC constraint is replaced by its first-order condition) is a relaxed problem: an optimum must be a stationary point, but a stationary point is not necessarily an optimum.

Differentiating the global IC constraint with respect to a yields the first-order condition of the global IC constraint:

$$\left[r_2 c'(a) - \frac{\gamma}{V} \right] \int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \left[r_2 c'(a) + \frac{\gamma}{V} \right] \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR = 0$$

Note that the compensation schedule $s(R)$ specifically *does not* depend on a , so differentiating the agent's expected utility from exerting effort a given $s(R)$ with respect to a does not require any assumptions about $s(R)$.

Solving the first-order problem yields the following wage schedule:

$$\begin{aligned} s(R < \gamma \hat{a}) &= \frac{r_1}{r_1 + r_2} R - \frac{r_1}{r_1 + r_2} \gamma \hat{a} + c(\hat{a}) + \frac{1}{r_2} v - \frac{1}{r_2} \log \left[\left(1 - \frac{r_1 r_2}{r_1 + r_2} V \right) \left(1 + \frac{r_2 c'(\hat{a}) V}{\gamma} \right) \right] \\ s(R > \gamma \hat{a}) &= \frac{r_1}{r_1 + r_2} R - \frac{r_1}{r_1 + r_2} \gamma \hat{a} + c(\hat{a}) + \frac{1}{r_2} v - \frac{1}{r_2} \log \left[\left(1 + \frac{r_1 r_2}{r_1 + r_2} V \right) \left(1 - \frac{r_2 c'(\hat{a}) V}{\gamma} \right) \right] \end{aligned}$$

where \hat{a} is the effort level "anticipated" by the principal.

Observe that if $\hat{a} = c'^{-1} \left(\frac{r_1}{r_1 + r_2} \gamma \right)$, then this compensation schedule is fully linear with slope $\frac{r_1}{r_1 + r_2}$; if $\hat{a} > c'^{-1} \left(\frac{r_1}{r_1 + r_2} \gamma \right)$, then this compensation schedule is piecewise linear with a discrete jump at $\gamma \hat{a}$, and slope $\frac{r_1}{r_1 + r_2}$ everywhere else, that is, at output levels $R \neq \gamma \hat{a}$.

What \hat{a} does the principal choose in equilibrium? Note that in equilibrium it must be that $a^* = \hat{a}$: that is, the principal would never pay for a level of effort higher than the one she anticipates, and an agent would never exert more effort than she is compensated for.

Now observe that for all $\hat{a} < \hat{a}_t$, for some threshold $\hat{a}_t > 0$, the unique stationary point of the agent r_2 's expected utility given effort exertion a and compensation scheme $s(R|\hat{a})$ is $a = \hat{a}$. Moreover, r_2 's expected utility given a and $s(R|\hat{a})$ is strictly concave in a , when $\hat{a} < \hat{a}_t$.

However, once $\hat{a} > \hat{a}_t$, where again, this \hat{a}_t is a threshold which will be rigorously characterized shortly, there are two stationary points of r_2 's expected utility given effort exertion a and compensation scheme $s(R|\hat{a})$: $a_1 = \hat{a}$ continues to be a stationary point, but $a_2 < \hat{a}$ is also a stationary point, and it can be seen that a_2 is the unique maximizer. The intuition is that, if the principal tries to induce "too much" effort, the agent will find it profitable to discretely deviate downwards.

Therefore, the equilibrium \hat{a} is $\hat{a} = \hat{a}_t$, where \hat{a}_t is the value of \hat{a} such that the second derivative of r_2 's expected utility given effort exertion a and compensation scheme $s(R|\hat{a})$ at $a = \hat{a}$ is precisely 0.

This yields the following expression characterizing $\hat{a}^* = \hat{a}_t$:

$$\left(c'(\hat{a}_t) - \frac{r_1}{r_1 + r_2}\gamma\right) \left(c'(\hat{a}_t) + \frac{\gamma}{r_2 V}\right) = \frac{1}{r_2}c''(\hat{a}_t) > 0$$

where it can clearly be seen that $\hat{a}_t > c'^{-1}\left(\frac{r_1}{r_1+r_2}\gamma\right)$, since $c(a)$ is increasing and convex, and the left-hand side is increasing in a .

But we know the agent's expected utility from exerting effort a given compensation schedule $s(R|\hat{a}^*)$ is strictly concave in a .

Therefore, the unique solution we found to the first-order problem is indeed the unique optimum.

This characterizes the equilibrium.

4 Incentive Provision and Insurance Provision

The piecewise linearity of the equilibrium wage schedule which emerges in this framework is a nice property for two key reasons. First, it's both tractable and realistic. Second, it neatly separates incentive provision from insurance provision: the linearity with slope $\frac{r_1}{r_1+r_2}$ at output levels above and below the mean captures efficient risk-sharing, while the discrete jump at the mean captures incentive provision. Because the likelihood ratio is a piecewise constant with discontinuity at the mean output level, the key information about effort contained in output realization is whether the realized output is above or below $R = \gamma\hat{a}$. Conditional on knowing that realized output is below (above) mean output, however, no output level below (above) the mean is more informative about effort than another output level below (above) the mean. Furthermore, because the likelihood ratio assumes two "symmetric" values, $-Q$ and Q , the slope is the same for output levels above and below the mean.

To see more clearly how the piecewise linear wage scheme of the second-best solution cleanly separates incentive and insurance provision, it will be helpful to solve two cases: (a) the contracting principal between a risk-neutral principal and a risk-averse agent, where effort is not contractible, and (b) the contracting problem between a risk-averse principal and a risk-averse agent, where effort is contractible. The solution to the first case will be a step function, with one fixed wage for low output and a discretely higher fixed wage for high output, and the solution to the second case will be a perfectly linear wage with slope $\frac{r_1}{r_1+r_2}$, to capture efficient risk-sharing.

4.1 Case (a): the contracting problem between a risk-neutral principal and a risk-averse agent; effort not contractible

The principal has utility $u(x) = x$, while the agent has utility $u(x) = -e^{-r_2x}$. Effort is not contractible. Hence, the principal chooses wage $s(R)$ for the agent r_2 by solving the following problem:

$$\max_{s(R)} \int_{-\infty}^{\gamma a} (R - s(R)) \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \int_{\gamma a}^{\infty} (R - s(R)) \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR$$

such that:

$$\begin{aligned}
(IR) : & \int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \\
& + \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR \geq -e^{-v} \\
(IC) : & a \in \arg \max_{a \in (0, \infty)} \int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \\
& + \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR
\end{aligned}$$

Replacing the global IC constraint with the first-order condition and writing the Lagrangean (letting λ be the multiplier on the binding IR constraint, and μ be the multiplier on the binding first-order condition):

$$\begin{aligned}
& \int_{-\infty}^{\gamma a} (R - s(R)) \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \int_{\gamma a}^{\infty} (R - s(R)) \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR + \\
& + \lambda \left[\int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR + e^{-v} \right] \\
& + \mu \left[\left[r_2 c'(a) - \frac{\gamma}{V} \right] \int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \left[r_2 c'(a) + \frac{\gamma}{V} \right] \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR \right]
\end{aligned}$$

Then differentiating pointwise with respect to $s(R)$ for $R < \gamma a$ yields:

$$\begin{aligned}
-1 + \lambda r_2 e^{-r_2[s(R)-c(a)]} + \mu \left[r_2 c'(a) - \frac{\gamma}{V} \right] r_2 e^{-r_2[s(R)-c(a)]} & = 0 \Rightarrow \\
r_2 \left(\lambda + \mu \left[r_2 c'(a) - \frac{\gamma}{V} \right] \right) e^{-r_2[s(R)-c(a)]} & = 1 \\
\frac{1}{r_2} \log \left(r_2 \left(\lambda + \mu \left[r_2 c'(a) - \frac{\gamma}{V} \right] \right) \right) + c(a) & = s(R | R < \gamma a)
\end{aligned}$$

Hence, $s(R)$ for $R < \gamma a$ is:

$$s(R | R < \gamma a) = \frac{1}{r_2} \log \left(r_2 \left(\lambda + \mu \left[r_2 c'(a) - \frac{\gamma}{V} \right] \right) \right) + c(a)$$

And, differentiating pointwise with respect to $s(R)$ for $R > \gamma a$ yields:

$$s(R | R > \gamma a) = \frac{1}{r_2} \log \left(r_2 \left(\lambda + \mu \left[r_2 c'(a) + \frac{\gamma}{V} \right] \right) \right) + c(a)$$

Hence, the equilibrium wage contract for the agent r_2 is a wage fixed at a certain level for $R < \gamma \hat{a}$, and a wage fixed at a discretely higher level for $R > \gamma \hat{a}$, where \hat{a} is the effort level anticipated by the principal. So, the wage looks like, "flat, jump at the anticipated mean level of output, and flat again".

4.2 Case (b): the contracting problem between a risk-averse principal and a risk-averse agent; effort contractible

Now, both the principal and the agent are risk-averse with CARA utility, but effort is contractible. The principal r_1 solves the following problem:

$$\max_{a, s(R)} \int_{-\infty}^{\gamma a} -e^{-r_1(R-s(R))} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \int_{\gamma a}^{\infty} -e^{-r_1(R-s(R))} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR$$

such that:

$$(IR) : \int_{-\infty}^{\gamma a} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{\frac{1}{V}[R-\gamma a]} dR + \int_{\gamma a}^{\infty} -e^{-r_2[s(R)-c(a)]} \frac{1}{2V} e^{-\frac{1}{V}[R-\gamma a]} dR \geq -e^{-v}$$

Letting λ denote the Lagrange multiplier on the IR constraint, differentiating pointwise with respect to $s(R)$ for $R < \gamma a$ yields:

$$-r_1 e^{-r_1(R-s(R))} + \lambda r_2 e^{-r_2[s(R)-c(a)]} = 0$$

And differentiating pointwise with respect to $s(R)$ for $R > \gamma a$ yields:

$$-r_1 e^{-r_1(R-s(R))} + \lambda r_2 e^{-r_2[s(R)-c(a)]} = 0$$

the same condition.

Hence:

$$s(R|\hat{a}) = \frac{r_1}{r_1 + r_2} R + \frac{r_2}{r_1 + r_2} c(\hat{a}) - \frac{1}{r_1 + r_2} \log \left(\lambda \frac{r_2}{r_1} \right)$$

a linear wage schedule with slope $\frac{r_1}{r_1 + r_2}$.

The key takeaway from both of these cases is that, in this model, it is possible to see the effect of both the principal and the agent being risk-averse on the equilibrium wage, *and* the effect of effort not being observable or contractible. The equilibrium wage schedule that emerges when both of these components are combined is the wage schedule which is linear at output levels away from the mean (as in Case (b)), and where there is a discontinuous jump at the mean (as in Case (a)).

5 An Example

What are the advantages to working with this framework as opposed to using a Holmstrom and Milgrom style justification for linear contracts in a CARA-Normal setting?

The Holmstrom and Milgrom story requires one to take quite a structured stand about the nature of the risky environment and differences in the way that the agent and the principal observe the output which stochastically depends on agent effort. Moreover, the Holmstrom and Milgrom justification is often applied to static models, despite the inherent dynamism of the story. By

contrast, this simple framework is quite unstructured—one needs only to believe that a better choice of inputs leads to a higher probability of good output realizations. The assumption of a Laplace distribution is about as arbitrary as the assumption of a normal distribution—the Laplace distribution resembles the normal distribution along key dimensions, namely, symmetry, and in any case, both distributional assumptions are often made for convenience.

For example, a large body of work in development economics has studied informal insurance institutions in developing economies. Sharecropping relationships are observed to play an important informal insurance role: a risk-averse landlord and a risk-averse tenant farmer agree to a rental contract dividing the realized harvest profits in a way that balances provision of incentives as well as protection from risk (Stiglitz (1974), Townsend and Mueller (), and many others). One approach to characterizing the equilibrium strength of informal insurance institutions, then, is to characterize the equilibrium network of sharecropping relationships and contracts (Wang (2012)).

Studying the equilibrium network of relationships naturally calls for methods from cooperative game theory, where the equilibrium matching is defined to be an assignment of landlords to farmers such that, loosely, there are no profitable deviations (for example, no landlord and farmer who are not assigned to each other can defy their assignment by pairing with each other and both becoming better off). In this agricultural setting, a landlord cannot perfectly monitor the inputs of her hired farmer, but as the farmer lives on the land of the landlord, both the landlord and the farmer are typically easily able to observe the growth of the crops. Moreover, the growth of crops does not respond immediately or in any obvious way to inputs—it is known that good input choice will increase the probability of a good harvest, but it is unrealistic to think of crop growth as providing continuous feedback about farmer input.

Hence, this framework is a more natural fit for the analysis of this question. Moreover, the clean separation of incentive provision and insurance provision is particularly suited to the study of endogenous relationship formation in this context, since partner choice depends exactly on the costs and benefits of productivity and insurance within each possible partnership.

For more details, please see Wang (2012).

6 Conclusion

In this note, I have identified and solved a simple, static model of moral hazard in which both the principal and the agent are risk-averse and have CARA utility, where the key innovation is the assumption of Laplace-distributed output. I show from first principles that the first-order approach is valid in this model, and that the equilibrium wage is piecewise linear with discontinuity at mean output. This compensation scheme is tractable and observed empirically. Moreover, the incentive provision component and the insurance provision component are clear: at output levels above and below the mean, the scheme is linear with slope $\frac{r_1}{r_1+r_2}$, capturing efficient risk-sharing, while at the mean output level, the wage jumps discretely, capturing incentive provision.

I discuss the advantages of working within this framework relative to working in a CARA-Normal framework and justifying the use of linear contracts with the story from Holmstrom and Milgrom (1987). This framework requires much less structured assumptions on the environment. Moreover, the dynamic intuition of the Holmstrom and Milgrom story is often incongruously applied to static models. Finally, the clean separation of incentives and insurance in the equilibrium wage arising in this framework is particularly useful for the study of a variety of economic settings, including the study of the endogenous formation of sharecropping relationships to characterize the equilibrium strength of informal insurance institutions in developing countries.

7 References

1. Grossman, S. and O. Hart. "An Analysis of the Principal-Agent Problem." *Econometrica* (51) 1983 7-45.
2. Holmstrom, B. and P. Milgrom. "Aggregation and Linearity." *Econometrica* (55) 1987 303-328.
3. Jewitt, I. "Justifying the First-Order Approach to Principal-Agent Problems." *Econometrica* (56) 1988 1177-1190.
4. Jewitt, I. "Risk Aversion and the Choice Between Risky Prospects: the Preservation of Comparative Statics Results." *Review of Economic Studies* (54) 1987 73-85.
5. Mirrlees, J. "The Theory of Moral Hazard and Unobservable Behavior." *Review of Economic Studies* (66) 1999 3-21.
6. Mirrlees, J. "The Implications of Moral Hazard for Optimal Insurance." Mimeo 1979.
7. Rogerson, W. "The First-Order Approach to Principal-Agent Problems." *Econometrica* (53) 1985 1357-1367.
8. Stiglitz 1974. "Incentives and Risk Sharing in Sharecropping." *Review of Economic Studies* (41) 1974 219-255.
9. Townsend, R. and R. Mueller. "Mechanism Design and Village Economies: From Credit to Tenancy to Cropping Groups." *Review of Economic Dynamics* (1) 1998 119-172.
10. Wang, X. "Risk, Incentives, and Contracting Relationships." Working Paper 2012.
11. Zeckhauser, R. "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory* (2) 1970 10-26.