Economics 2010b: Microeconomic Theory II

Michael Powell

The Moral Hazard Problem

# 1 Introduction to Contract Theory

One of the important issues that we touched on briefly in our discussion of general equilibrium theory is the idea of market incompleteness and its consequences. When markets are incomplete—either in the sense that we talked about last time or in the sense that consumption involves unpriceable externalities—equilibrium allocations may not be constrained-efficient, opening up scope for some sort of third-party intervention. It may be government intervention via a system of taxation or rules, or it may be private intervention by an entrepreneur who sets up governance institutions. We were able to make some high-level claims last time about what happens when there are these "market failures," but without imposing more structure on the problem, it is difficult to make specific claims about how they should be managed.

For the last three weeks of the class, we will zoom in and study micro situations in which it could be said that markets are incomplete. We will focus on what is referred to as the Principal–Agent problem in which there are two players, a Principal $P$ and an Agent $A$. The Principal needs the Agent to do something that she cannot do herself, so she hires the Agent and writes a contract that governs how the Agent will be paid. We can think of the Principal being an employer and the Agent an employee, where the Principal lacks the time or expertise to engage in production. We can think of the Principal being a patient and the Agent a doctor, where the doctor takes some actions that the patient does not know or understand. We can think of the Principal being a client and the Agent being a lawyer acting on the client's behalf. And so on.

When equilibrium outcomes arising from the Principal–Agent interaction are Pareto inefficient, we will say that there is a *moral hazard problem*, which is a term that originated in the insurance industry to describe situations in which someone increases their exposure to risk in response to buying insurance. Fundamentally, the moral hazard problem is a just an externality problem. Now, when we make a claim like "there are externalities, so outcomes will be inefficient" it is important to have in mind that whether or not externalities "matter" in the sense that they lead to Pareto inefficient equilibrium outcomes depends critically on the set of instruments parties have for managing those externalities: it depends on the contracting space. Over the next couple lectures, we will look at several different sources of *contractual frictions* that prevent the Principal and Agent from writing contracts with each other that result in Pareto optimal outcomes.

The first situation we will look at will occur when individual actions chosen by the Agent are not observed by the Principal but determine the distribution of a verifiable performance measure that can be written into a contract. The Agent may be more risk-averse than the Principal, so writing a high-powered contract on that noisy performance measure transfers risk onto the Agent and therefore leads to an inefficient allocation of risk between the two parties. As a result, there is a trade-off between incentive provision (and therefore what the Agent chooses to do) and inefficient risk allocation. This is the celebrated *risk–incentives trade-off*.

The second contracting friction that might arise is that an Agent is either liquidity-constrained or is subject to a limited-liability constraint. As a result, the Principal is unable to extract all the surplus the Agent generates and must therefore provide the Agent with *incentive rents* in order to motivate him. That is, offering the Agent a higher-powered contract induces him to work harder and therefore increases the total size of the pie, but it also leaves the Agent with a larger share of that pie. The Principal then, in choosing a contract, chooses one that trades off the creation of surplus with her ability to extract that surplus. This is the *motivation–rent extraction trade-off*.

A third contracting friction that might arise is that the Principal's objective simply cannot be written into a formal contract. Instead, the Principal has to rely on imperfectly aligned performance measures. Increasing the strength of a formal contract that is based on imperfectly aligned performance measures may motivate the Agent to work hard toward the Principal's objectives, but it may also motivate him to work hard toward objectives that either hurt the Principal or at least do not help her. This is known as the *multi-task problem* (Holmström and Milgrom, 1991), and failure to account for the effects of using distorted performance measures is sometimes referred to as *the folly of rewarding A while hoping for B* (Kerr, 1975).

Finally, there may be multiple Agents who work together to produce something for the Principal. Their individual contributions may not be observable, so contracts may only be able to be written on the final output. This inability to distinguish individual contributions is what is referred to as the *moral hazard in teams problem* (Holmström, 1982).

All of these sources of contractual frictions lead to similar results—under the optimal contract, the Agent (or Agents) chooses an action that is not jointly optimal from his and the Principal's perspective. But in different applied settings, different assumptions regarding what is contractible and what is not are more or less plausible. As a result, it is useful to master at least elementary versions of models capturing these four sources of frictions, so that you are well-equipped to use them as building blocks.

# 2 The Risk-Incentives Trade-off

I will begin with a pretty general description of the standard principal-agent model, but I will shortly afterwards specialize the model quite a bit in order to focus on a single point—the risk–incentives trade-off.

## 2.1 The Model

There is a risk-neutral Principal ($P$) and a risk-averse Agent ($A$). The Agent chooses an **effort level** $e \in \mathcal{E} \subset \mathbb{R}_+$ and incurs a cost of $c(e)$, where $c : \mathbb{R}_+ \to \mathbb{R}_+$ is strictly increasing and strictly convex. If $\mathcal{E}$ is an interval, we will say that **effort is continuous**, and if $\mathcal{E}$ consists of a finite number of points, we will say that **effort is discrete**. We will assume $0 \in \mathcal{E}$, and $c(0) = 0$. The effort level affects the distribution over **output** $y \in \mathcal{Y}$, with $y$ distributed according to cdf $F(\cdot|e)$. This output can be sold on the product market at price $p$, and the revenues $py$ accrue to the Principal.

The Principal does not have any direct control over the Agent, but what she can do is write a contract that influences what the Agent will do. In particular, she can write a contract $w \in \mathcal{W} \subset \{w : \mathcal{Y} \times \mathcal{E} \to \mathbb{R}\}$, where $\mathcal{W}$ is the **contracting space**. The contract determines a transfer $w(y, e)$ that she is compelled to pay the Agent if output $y$ is realized, and he chose effort $e$. If $\mathcal{W}$ does not allow for functions that depend directly on effort, we will say that **effort is noncontractible**, and abusing notation slightly, we will write the contractual payment the Principal is compelled to pay the Agent if output $y$ is realized as $w(y, e) = w(y)$ for all $e \in \mathcal{E}$. We will be assuming throughout that effort is noncontractible, but I wanted to highlight that it is a real restriction on the contracting space, and it is one that we will impose as a primitive of the model.

The Agent can decline to work for the Principal and reject her contract, pursuing his outside option instead. This outside option provides utility $\bar{u}$ to the Agent and $\bar{\pi}$ to the Principal. If the Agent accepts the contract, the Principal's and Agent's preferences are, respectively,

$$\begin{aligned} \Pi(w, e) &= \int_{y \in \mathcal{Y}} (py - w(y))\, dF(y|e) = E_y[py - w|e] \\ U(w, e) &= \int_{y \in \mathcal{Y}} u(w(y) - c(e))\, dF(y|e) = E_y[u(w - c(e))|e], \end{aligned}$$

where $u$ is increasing and weakly concave.

We have described the players, what they can do, and what their preferences are. We still need to describe the timing of the game that the players play, as well as the solution concept. Explicitly describing the timing of the model is essential to remove any ambiguity about what players know when they make their decisions. In this model, the timing of the game is:

1. $P$ offers $A$ a contract $w \in \mathcal{W}$. $w$ is commonly observed.

2. $A$ accepts the contract $(d = 1)$ or rejects it $(d = 0)$, in which case he receives $\bar{u}$, and the game ends. $d$ is commonly observed.

3. If $A$ accepts the contract, $A$ chooses effort level $e$ and incurs cost $c(e)$. $e$ is privately observed by $A$.

4. Output $y$ is drawn from distribution with cdf $F(\cdot | e)$. $y$ is commonly observed.

5. $P$ pays $A$ an amount $w(y)$. The payment is commonly observed.

A couple remarks are in order at this point. First, behind the scenes, there is an implicit assumption that there is a third-party contract enforcer (a judge or arbitrator) who can costlessly detect when agreements have been broken and costlessly exact harsh punishments on the offender.

Second, much of the literature assumes that the Agent's effort level is privately observed by the Agent and therefore refers to this model as the "hidden action" model. Ultimately, though, the underlying source of the moral-hazard problem is that contracts cannot be conditioned on relevant variables, not that the relevant variables are unobserved by the Principal. Many papers assume effort is unobservable to justify it being noncontractible. While this is a compelling justification, in our framework, the contracting space itself is a primitive of the model. Later in the course, we will talk a bit about the microfoundations for different assumptions on the contracting space.

Finally, let us describe the solution concept. A **pure-strategy subgame-perfect equilibrium** is a contract $w^* \in \mathcal{W}$, an **acceptance decision** $d^* : \mathcal{W} \to \{0, 1\}$, and an **effort choice** $e^* : \mathcal{W} \times \{0, 1\} \to \mathcal{E}$ such that, given the contract $w^*$, the Agent optimally chooses $d^*$ and $e^*$, and given $d^*$ and $e^*$, the Principal optimally offers contract $w^*$. We will say that the optimal contract **induces** effort $e^*$.

## 2.2 First-Best Benchmark

If we want to talk about the inefficiencies that arise in equilibrium in this model, it will be useful first to establish a benchmark against which to compare outcomes. In this model, a **feasible outcome** is a distribution over payments from the Principal to the Agent as well as an effort level $e \in \mathcal{E}$. We will say that an outcome is **Pareto optimal** if there is no other feasible outcome that both players weakly prefer and one player strictly prefers. If an effort level $e$ is part of a Pareto optimal outcome, we will say that it is a **first-best** effort level, and we will denote it by $e^{FB}$.

**Lemma 1.** The first-best effort level satisfies

$$e^{FB} \in \operatorname*{argmax}_{e \in \mathcal{E}} E\left[\left.py\right| e\right] - c\left(e\right).$$

**Proof of Lemma 1**. In any Pareto-optimal outcome, payments to the agent are deterministic. Since the Agent is risk averse, given an outcome involving stochastic payments to the Agent, there is another outcome in which the Agent chooses the same effort level and receives the certainty equivalent wage instead. This outcome yields the same utility for the Agent, and since the Agent is risk averse, the certainty equivalent payment is smaller in expectation, so the Principal is strictly better off. Next, given constant deterministic wages, any Pareto-optimal outcome must solve

$$\max_{w \in \mathbb{R}, e \in \mathcal{E}} E\left[\left.py\right| e\right] - w$$

subject to

$$u\left(w-c\left(e\right)\right)\geq\bar{u},$$

for some $\bar{u}$. In any solution to this problem, the constraint must bind, since $u$ is increasing. Moreover, since $u$ is increasing, it is invertible, so we can write

$$w=u^{-1}\left(\bar{u}\right)+c\left(e\right),$$

and therefore the first-best effort level must solve the problem specified in the Lemma.■

This Lemma shows that the first-best effort level maximizes expected revenues net of effort costs. If effort is fully contractible, so that the Principal could offer any contract $w$ that depended nontrivially on $e$, then the first-best effort would be implemented in equilibrium. In particular, the Principal could offer a contract that pays the Agent $u^{-1}\left(\bar{u}\right)+c\left(e^{FB}\right)$ if he choose $e^{FB}$, and pays him a large negative amount if he chooses any $e\neq e^{FB}$. That the first-best effort level can be implemented in equilibrium if effort is contractible is an illustration of a version of the *Coase Theorem*: if the contracting space is sufficiently rich, equilibrium outcomes will be Pareto optimal.

If effort is noncontractible, and $e^{FB}>0$, then equilibrium will not involve Pareto optimal outcomes. For an outcome to be Pareto optimal, it has to involve a deterministic wage payment to the Agent. But if the Agent's wage is independent of output, then it must also be independent of his effort level. He will therefore receive no benefit from choosing a costly effort level, and so he will choose $e=0<e^{FB}$. The question to which we will now turn is: what effort will be implemented in equilibrium when effort is noncontractible?

## 2.3 Equilibrium Effort

Since the Agent's effort choice affects the Principal's payoffs, the Principal would ideally like to directly choose the Agent's effort. But, she has only indirect control: she can offer different contracts, and different contracts may get the Agent to optimally choose different effort levels. We can think of the Principal's problem as choosing an effort level $e$ as well as a contract for which $e$ is *incentive compatible* for the Agent to choose and for which it is *individually rational* for the Agent to accept. As a loose analogy, we can connect the Principal's problem to the social planner's problem from general equilibrium theory. We can think of $e$ as analogous to an allocation the Principal would like to induce, and the choice of a contract as analogous to setting "prices" so as to decentralize $e$ as an equilibrium allocation.

Formally, the Principal offers a contract $w \in \mathcal{W}$ and "proposes" an effort level $e$ in order to solve

$$\max_{w \in \mathcal{W}, e \in \mathcal{E}} \int_{y \in \mathcal{Y}} (py - w(y)) \, dF(y|\, e)$$

subject to two constraints. The first constraint is that the agent actually prefers to choose effort level $e$ rather than any other effort level $\hat{e}$. This is the **incentive-compatibility constraint**:

$$e \in \underset{\hat{e} \in \mathcal{E}}{\text{argmax}} \int_{y \in Y} u(w(y) - c(\hat{e})) \, dF(y|\, \hat{e}).$$

The second constraint ensures that, given that the agent knows he will choose $e$ if he accepts the contract, he prefers to accept the contract rather than to reject it and receive his outside utility $\bar{u}$. This is the **individual-rationality constraint** or **participation constraint**:

$$\int_{y \in Y} u(w(y) - c(e)) \, dF(y|\, e) \geq \bar{u}.$$

At this level of generality, the model is not very tractable. We will need to impose more structure on it in order to highlight some its key trade-offs and properties.

**CARA-Normal Case with Affine Contracts** In order to highlight one of the key trade-offs that arise in this class of models, we will make a number of strong simplifying assumptions.

**Assumption A1 (CARA).** The Agent has CARA preferences over wealth and effort costs, which are quadratic:

$$u\left(w\left(y\right)-c\left(e\right)\right)=-\exp\left\{-r\left(w\left(y\right)-\frac{c}{2}e^{2}\right)\right\},$$

and his outside option yields utility $-\exp\left\{-r\bar{u}\right\}$.

**Assumption A2 (Normal Output)**. Effort shifts the mean of a normally distributed random variable. That is, $y\sim N\left(e,\sigma^{2}\right)$.

**Assumption A3 (Affine Contracts).** $\mathcal{W}=\left\{w:\mathcal{Y}\to\mathbb{R},w\left(y\right)=s+by\right\}$. That is, the contract space permits only affine contracts.

**Assumption A4 (Continuous Effort)**. Effort is continuous and satisfies $\mathcal{E}=\mathbb{R}_{+}$.

In principle, we should not impose exogenous restrictions on the *functional form* of $w\left(y\right)$. There is an important class of applications, however, that restrict attention to affine contracts, $w\left(y\right)=s+by$, and a lot of the basic intuition that people have for the comparative statics of optimal contracts come from imposing this restriction.

In many environments, an optimal contract does not exist if the contracting space is sufficiently rich, and situations in which the agent chooses the first-best level of effort, and the principal receives all the surplus can be arbitrarily approximated with a sequence of sufficiently perverse contracts (Mirrlees, 1974; Moroni and Swinkels, 2014). In contrast, the optimal affine contract often results in an effort choice that is lower than the first-best effort level, and the principal receives a lower payoff.

There are then at least three ways to view the exercise of solving for the optimal affine contract.

1. From an applied perspective, many pay-for-performance contracts in the world are affine in the relevant performance measure—franchisees pay a franchise fee and receive a constant fraction of the revenues their store generates, windshield installers receive a base wage and a constant piece rate, fruit pickers are paid per kilogram of fruit they pick. And so given that many practitioners seem to restrict attention to this class

of contracts, why not just make sure they are doing what they do optimally? Put differently, we can brush aside global optimality on purely pragmatic grounds.

2. Many pay-for-performance contracts in the world are affine in the relevant performance measure. Our models are either too rich or not rich enough in a certain sense and therefore generate optimal contracts that are inconsistent with those we see in the world. Maybe the aspects that, in the world, lead practitioners to use affine contracts are orthogonal to the considerations we are focusing on, so that by restricting attention to the optimal affine contract, we can still say something about how real-world contracts ought to vary with changes in the underlying environment. This view presumes a more positive (as opposed to normative) role for the modeler and hopes that the theoretical analogue of the omitted variables bias is not too severe.

3. Who cares about second-best when first-best can be attained? If our models are pushing us toward complicated, non-linear contracts, then maybe our models are wrong. Instead, we should focus on writing down models that generate affine contracts as the optimal contract, and therefore we should think harder about what gives rise to them. (And indeed, steps have been made in this direction—see Holmström and Milgrom (1987), Diamond (1998) and, more recently, Carroll (2013) and Barron, Georgiadis, and Swinkels (2017)) This perspective will come back later in the course when we discuss the Property Rights Theory of firm boundaries.

Given Assumptions $(A1)-(A3)$, for any contract $w(y)=s+by$, the income stream the agent receives is normally distributed with mean $s+be$ and variance $b^2\sigma^2$. His expected utility over monetary compensation is therefore a moment-generating function for a normally distributed random variable, (recall that if $X\sim N(\mu,\sigma^2)$, then $E[\exp\{tX\}]=\exp\left\{\mu t+\frac{1}{2}\sigma^2t^2\right\}$), so his preferences can be written as

$$E\left[-\exp\left\{-r\left(w(y)-c(e)\right)\right\}\right]=-\exp\left\{-r\left(s+be-\frac{r}{2}b^2\sigma^2-\frac{c}{2}e^2\right)\right\}.$$

We can take a monotonic transformation of his utility function ($f(x) = -\frac{1}{r}\log(-x)$) and represent his preferences as:

$$\begin{aligned} U(e,w) &= E[w(y)] - \frac{r}{2}Var(w(y)) - \frac{c}{2}e^2 \\ &= s + be - \frac{r}{2}b^2\sigma^2 - \frac{c}{2}e^2. \end{aligned}$$

The Principal's program is then

$$\max_{s,b,e} pe - (s + be)$$

subject to incentive-compatibility

$$e \in \operatorname*{argmax}_{\hat{e}} b\hat{e} - \frac{c}{2}\hat{e}^2$$

and individual-rationality

$$s + be - \frac{r}{2}b^2\sigma^2 - \frac{c}{2}e^2 \geq \bar{u}.$$

Solving this problem is then relatively straightforward. Given an affine contract $s + be$, the Agent will choose an effort level $e(b)$ that satisfies his first-order conditions

$$e(b) = \frac{b}{c},$$

and the Principal will choose the value $s$ to ensure that the Agent's individual-rationality constraint holds with equality. If it did not hold with equality, the Principal could reduce $s$, making herself better off without affecting the Agent's incentive-compatibility constraint, while still respecting the Agent's individual-rationality constraint. That is,

$$s + be(b) = \frac{c}{2}e(b)^2 + \frac{r}{2}b^2\sigma^2 + \bar{u}.$$

In other words, the Principal has to ensure that the Agent's total expected monetary compensation, $s + be(b)$, fully compensates him for his effort costs, the risk costs he has to bear if he accepts this contract, and his opportunity cost. Indirectly, then, the Principal bears these costs when designing an optimal contract.

The Principal's remaining problem is to choose the incentive slope $b$ to solve

$$\max_b pe(b) - \frac{c}{2}e(b)^2 - \frac{r}{2}b^2\sigma^2 - \bar{u}.$$

This is now an unconstrained problem with proper convexity assumptions, so the Principal's optimal choice of incentive slope solves her first-order condition

$$0 = p\underbrace{e'(b^*)}_{1/c} - c\underbrace{e^*(b^*)}_{b^*/c}\underbrace{e'(b^*)}_{1/c} - rb^*\sigma^2,$$

and therefore the optimal incentive slope satisfies

$$b^* = \frac{p}{1 + rc\sigma^2}.$$

Moreover, given $b^*$ and the individual-rationality constraint, we can back out $s^*$.

$$s^* = \bar{u} + \frac{1}{2}\left(rc\sigma^2 - 1\right)\frac{(b^*)^2}{c}.$$

Depending on the parameters, it may be the case that $s^* < 0$. That is, the Agent would have to pay the Principal if he accepts the job and does not produce anything.

Now, how does the effort that is induced in this optimal affine contract compare to the **first-best effort**? Using the result from Lemma 1, we know that first-best effort in this setting solves

$$\max_{e \in \mathbb{R}_+} pe - \frac{c}{2}e^2,$$

and therefore $e^{FB} = p/c$.

Even if effort is noncontractible, the Principal could in principle implement exactly this same level of effort by writing a contract only on output. To do so, she would choose $b = p$, since this would get the Agent to choose $e(p) = p/c$. Why, in this setting, does the Principal not choose such a contract? Let us go back to the Principal's problem of choosing the incentive slope $b$.

$$\max_b pe(b) - \frac{c}{2}e(b)^2 - \frac{r}{2}b^2\sigma^2 - \bar{u}$$

Often, when an economic model can be solved in closed form, we jump right to the solution. Only when a model cannot be solved in closed form do we typically stop to think carefully about what economic properties its solution must possess. I want to spend a couple minutes *partially* characterizing this model's solution, even though we already completely characterized it above, just to highlight how this kind of reasoning can be helpful in developing intuition that might generalize beyond the present setting. In particular, many fundamental features of models can be seen as a comparison of first-order losses or gains against second-order gains or losses, so it is worth going through this first-order–second-order logic. Suppose the Principal chooses $b = p$, and consider a marginal reduction in $b$ away from this value. The change in the Principal's profits would be

$$\begin{aligned} & \frac{d}{db}\left(pe(b) - \frac{c}{2}e(b)^2 - \frac{r}{2}b^2\sigma^2\right)\Bigg|_{b=p} \\ = & \underbrace{\frac{d}{db}\left(pe(b) - \frac{c}{2}e(b)^2\right)\Bigg|_{b=p}}_{=0} - rp\sigma^2 < 0. \end{aligned}$$

This first term is zero, because $b = p$ in fact maximizes $pe(b) - \frac{c}{2}e(b)^2$, since it induces the first-best level of effort. This is just an application of the envelope theorem you learned in Ec 2010a. The second term in this expression is strictly negative. This implies that, relative to the contract that induces first-best effort, a reduction in the slope of the incentive contract yields a first-order gain to the Principal resulting from a decrease in the risk costs the Agent bears, while it yields a second-order loss in terms of profits resulting from moving away from the effort level that maximizes revenues minus effort costs. The optimal contract balances the incentive benefits of higher-powered incentives with these risk costs, and these risk costs are higher if the Agent is more risk averse and if output is noisier.

This trade-off seems first-order in some settings (e.g., insurance contracts in health care markets, some types of sales contracts in industries in which individual sales are infrequent, large, and unpredictable) and for certain types of output. There are many other environments in which contracts provide less-than-first-best incentives, but the first-order reasons for these low-powered contracts seem completely different, and we will turn to these environments next week.

**Exercise 18 (Adapted from MWG 14.B.4)**. Suppose there are three possible effort levels, $\mathcal{E} = \{e_1, e_2, e_3\}$, and two possible output levels, $\mathcal{Y} = \{0, 10\}$, and the output price is $p = 1$. The probability that $y = 10$ conditional on each of the effort levels is given by the probability mass function $f(10|\, e_1) = 2/3$, $f(10|\, e_2) = 1/2$, and $f(10|\, e_3) = 1/3$. The Agent's effort cost function satisfies $c(e_1) = 5/3$, $c(e_2) = 8/5$, and $c(e_3) = 4/3$. Finally, the Agent's utility function is given by $u(w) = \sqrt{w}$, and his outside option yields utility $\bar{u} = 0$.

$(a)$ What is the optimal contract for the Principal when effort is contractible?

$(b)$ Show that if effort is noncontractible, and $\mathcal{W} = \{w : \mathcal{Y} \to \mathbb{R}\}$, then there is no contract $w$ for which the Agent will choose $e_2$. For what levels of $c(e_2)$ would there exist a contract $w$ under which the Agent would choose $e_2$?

$(c)$ What is the optimal contract when effort is noncontractible, and $\mathcal{W} = \{w : \mathcal{Y} \to \mathbb{R}\}$?

$(d)$ Suppose instead that $c(e_1) = \sqrt{8}$, and let $f(10|\, e_1) = x \in (0, 1)$. If effort is noncontractible, and $\mathcal{W} = \{w : \mathcal{Y} \to \mathbb{R}\}$, what is the optimal contract for the Principal as $x$ approaches 1? Is the level of effort implemented higher or lower than when effort is contractible?

**Exercise 19**. Suppose the Agent can allocate time to two different tasks. Let $e_i$ be the amount of time spent on task $i \in \{1, 2\}$. The Principal cares only about task 1 and obtains payoff $y = e_1 + \varepsilon$, where $\varepsilon \sim N(0, \sigma^2)$. The Agent, however, derives a benefit $v(e_2)$ from spending time on task 2. The Agent has CARA preferences with utility function

$$u(w, e_1, e_2) = -\exp\{-r[w - c(e_1 + e_2) + v(e_2)]\},$$

where $c(e_1 + e_2)$ is the cost of time, with $c'(\cdot) > 0$, $c''(\cdot) > 0$, and $c(0) = 0$. Assume also that $v'(\cdot) > 0$, $v''(\cdot) < 0$, and $v(0) = 0$, and that optimization with respect to $(e_1, e_2)$ results in an interior solution. Let $\bar{w}$ denote the wage the Agent receives from his outside option, so $\bar{u} = -\exp\{-r\bar{w}\}$.

$(a)$ What is the first-best outcome in this setting?

$(b)$ Suppose effort $e_1$ is noncontractible, and the Principal can write a contract that is an affine function of output and can also allow the Agent to engage in task 2 or not. Under these assumptions, what is the contracting space?

$(c)$ Suppose the Principal must pay the Agent $s = 1$ if $y = 0$. Will the Principal allow the Agent to engage in task 2? Compare this to your answer in part $(a)$. What if $s < 1$ is set exogenously? Find the difference in the Principal's utility under the two policies, as a function of $s$.