1 Appendix: For Online Publication

1.1 Optimal contract with transfers

The set-up. In this section, I discuss the optimal contract when the principal can make transfers to the agent. I assume that the principal has full commitment power, that is, she can write a contract specifying both an experimentation policy π and a transfer scheme c at the outset of the game. I also assume that the agent is protected by limited liability so only non-negative transfers from the principal to the agent are allowed. An experimentation policy π is defined in the same way as before. A transfer scheme c offered by the principal is a non-negative, non-decreasing process $\{c_t\}_{t\geq 0}$, which may depend only on the history of events up to t, where c_t denotes the cumulative transfers the principal has made to the agent up to, and including, time t.[1] Let Π^* denote the set of all possible policy and transfer scheme pairs.

For any policy and transfer scheme pair (π, c) and any prior p, the principal's and the agent's payoffs are respectively

$$U_\alpha(\pi,c,p) = \mathbf{E}\left[\int_0^\infty re^{-rt}\left[(1-\pi_t)s_\alpha + \pi_t\lambda^\omega h_\alpha\right]\mathrm{d}t + \int_0^\infty re^{-rt}\mathrm{d}c_t \,\middle|\, \pi, c, p\right]$$
$$U_\rho(\pi,c,p) = \mathbf{E}\left[\int_0^\infty re^{-rt}\left[(1-\pi_t)s_\rho + \pi_t\lambda^\omega h_\rho\right]\mathrm{d}t - \int_0^\infty re^{-rt}\mathrm{d}c_t \,\middle|\, \pi, c, p\right].$$

For a given policy and transfer scheme pair (π, c), I define $\mathbf{t}^1(\pi,c)$ and $\mathbf{t}^0(\pi,c)$ as follows:

$$\mathbf{t}^1(\pi,c) \equiv \mathbf{E}\left[\int_0^\infty re^{-rt}\mathrm{d}c_t \,\middle|\, \pi,c,1\right] \text{ and } \mathbf{t}^0(\pi,c) \equiv \mathbf{E}\left[\int_0^\infty re^{-rt}\mathrm{d}c_t \,\middle|\, \pi,c,0\right]. \tag{1}$$

I refer to $\mathbf{t}^1(\pi,c)$ (resp. $\mathbf{t}^0(\pi,c)$) as the expected transfer in state 1 (resp. state 0).

Lemma 1 (A policy and transfer scheme pair as four numbers)**.**
For a given policy and transfer scheme pair $(\pi,c)\in\Pi^*$ and a given prior $p\in[0,1]$, the principal's and the agent's payoffs can be written as

$$\begin{aligned}
U_\alpha(\pi,c,p) =& p\left[\left(\lambda^1 h_\alpha - s_\alpha\right)\mathbf{w}^1(\pi) + \mathbf{t}^1(\pi,c)\right] \\
&+ (1-p)\left[\left(\lambda^0 h_\alpha - s_\alpha\right)\mathbf{w}^0(\pi) + \mathbf{t}^0(\pi,c)\right] + s_\alpha \\
U_\rho(\pi,c,p) =& p\left[\left(\lambda^1 h_\rho - s_\rho\right)\mathbf{w}^1(\pi) - \mathbf{t}^1(\pi,c)\right] \\
&+ (1-p)\left[\left(\lambda^0 h_\rho - s_\rho\right)\mathbf{w}^0(\pi) - \mathbf{t}^0(\pi,c)\right] + s_\rho.
\end{aligned}$$

[1] Formally, the number of successes achieved up to, and including, time t defines the point process $\{N_t\}_{t\geq 0}$. Let $\mathcal{F} := \{\mathcal{F}_t\}_{t\geq 0}$ denote the filtration generated by the process π and N_t. The process $\{c_t\}_{t\geq 0}$ is $\mathcal{F}$-adapted.

Proof. The agent's payoff given $(\pi, c) \in \Pi^*$ and prior $p \in [0,1]$ is

$$\begin{aligned} U_\alpha(\pi, c, p) =& p\mathbf{E}\left[\int_0^\infty re^{-rt}\pi_t\left(\lambda^1 h_\alpha - s_\alpha\right)\mathrm{d}t + \int_0^\infty re^{-rt}\mathrm{d}c_t \,\middle|\, \pi, c, 1\right] \\ &+ (1-p)\mathbf{E}\left[\int_0^\infty re^{-rt}\pi_t\left(\lambda^0 h_\alpha - s_\alpha\right)\mathrm{d}t + \int_0^\infty re^{-rt}\mathrm{d}c_t \,\middle|\, \pi, c, 0\right] + s_\alpha \\ =& p\left[\left(\lambda^1 h_\alpha - s_\alpha\right)\mathbf{w}^1(\pi) + \mathbf{t}^1(\pi, c)\right] \\ &+ (1-p)\left[\left(\lambda^0 h_\alpha - s_\alpha\right)\mathbf{w}^0(\pi) + \mathbf{t}^0(\pi, c)\right] + s_\alpha. \end{aligned}$$

The principal's payoff can be rewritten similarly. ■

Lemma 1 shows that all payoffs from implementing (π, c) can be written in terms of its expected resource and expected transfer pairs. Instead of working with a generic policy/transfer scheme pair, it is without loss of generality to focus on its expected resource and expected transfer pairs. The image of the mapping $(\mathbf{w}^1, \mathbf{w}^0, \mathbf{t}^1, \mathbf{t}^0) : \Pi^* \to [0,1]^2 \times [0,\infty) \times [0,\infty)$ can be interpreted as the new contract space when transfers are allowed. The following lemma characterizes this contract space.

Lemma 2. *The image of the mapping* $(\mathbf{w}^1, \mathbf{w}^0, \mathbf{t}^1, \mathbf{t}^0) : \Pi^* \to [0,1]^2 \times [0,\infty) \times [0,\infty)$, *denoted* Γ^*, *satisfies the following condition*

$$\mathrm{int}(\Gamma \times [0,\infty)^2) \subset \Gamma^* \subset \Gamma \times [0,\infty)^2.$$

Proof. The relation $\Gamma^* \subset \Gamma \times [0,\infty)^2$ is obviously true. Hence, I only need to show that $\mathrm{int}(\Gamma \times [0,\infty)^2) \subset \Gamma^*$. Given that Γ^* is a convex set, I only need to show that Γ^* is a dense set of $\mathrm{int}(\Gamma \times [0,\infty)^2)$: For any $(w^1, w^0, t^1, t^0) \in \mathrm{int}(\Gamma \times [0,\infty)^2)$ and any $\epsilon > 0$, there exists (π, c) such that the Euclidean distance $\|(w^1, w^0, t^1, t^0) - (\mathbf{w}^1, \mathbf{w}^0, \mathbf{t}^1, \mathbf{t}^0)(\pi, c)\|$ is below ϵ. Pick any point $(w^1, w^0) \in \mathrm{int}(\Gamma)$, there exists a bundle $(\tilde{w}^1, \tilde{w}^0) \in \mathrm{int}(\Gamma)$ and a small number Δ such that $(w^1, w^0) = (1-\Delta)(\tilde{w}^1, \tilde{w}^0) + \Delta(1,1)$. The policy is as follows. With probability $1-\Delta$, the agent implements a policy that is mapped to $(\tilde{w}^1, \tilde{w}^0)$. With probability Δ, the agent implements the policy that is mapped to $(1,1)$. In the latter case, the agent allocates the unit resource to R all the time. Transfers only occur in the latter case. Here, I construct a transfer scheme such that the expected transfer is arbitrarily close to (t^1, t^0). Let p_t denote the posterior belief that $\omega = 1$. Given that all the resource is directed to R, p_t converges in probability to 1 conditional on state 1 and p_t converges in probability to 0 conditional on state 0. This implies that $\forall \tilde{\epsilon} > 0$, $\exists \tilde{t}$ such that for all $t \geq \tilde{t}$, I have $\Pr(|p_t - 1| > \tilde{\epsilon} \mid \omega = 1) < \tilde{\epsilon}$ and $\Pr(|p_t - 0| > \tilde{\epsilon} \mid \omega = 0) < \tilde{\epsilon}$. The transfer scheme is to make a transfer of size $t^1/(\Delta r e^{-r\tilde{t}})$ at time $\tilde{t}$ if $p_{\tilde{t}} > 1 - \tilde{\epsilon}$ and make a transfer of size

$t^0/(\Delta re^{-r\tilde{t}})$ if $p_{\tilde{t}} < \tilde{\epsilon}$. The expected transfer conditional on state 1 is

$$\begin{aligned} &\Delta re^{-r\tilde{t}}\left[\Pr(p_{\tilde{t}} > 1-\tilde{\epsilon} \mid \omega = 1)\frac{t^1}{\Delta re^{-r\tilde{t}}} + \Pr(p_{\tilde{t}} < \tilde{\epsilon} \mid \omega = 1)\frac{t^0}{\Delta re^{-r\tilde{t}}}\right] \\ =&\Pr(p_{\tilde{t}} > 1-\tilde{\epsilon} \mid \omega = 1)t^1 + \Pr(p_{\tilde{t}} < \tilde{\epsilon} \mid \omega = 1)t^0. \end{aligned}$$

Given that $1-\tilde{\epsilon} < \Pr(p_{\tilde{t}} > 1-\tilde{\epsilon} \mid \omega = 1) \leq 1$ and $0 \leq \Pr(p_{\tilde{t}} < \tilde{\epsilon} \mid \omega = 1) < \tilde{\epsilon}$, the expected transfer conditional on state 1 is in the interval $(t^1 - \tilde{\epsilon}t^1, t^1 + \tilde{\epsilon}t^0)$. Similarly, the expected transfer conditional on state 0 is in the interval $(t^0 - \tilde{\epsilon}t^0, t^0 + \tilde{\epsilon}t^1)$. As $\tilde{\epsilon}$ approaches zero, the constructed transfer scheme is arbitrarily close to (t^1, t^0). ■

Lemma 2 says that any $(w^1, w^0, t^1, t^0) \in \Gamma \times [0,\infty) \times [0,\infty)$ is virtually implementable: for all $\epsilon > 0$, there exist a (π, c) such that $(\mathbf{w}^1, \mathbf{w}^0, \mathbf{t}^1, \mathbf{t}^0)(\pi, c)$ is ϵ-close to (w^1, w^0, t^1, t^0). To proceed, I treat the set $\Gamma \times [0,\infty) \times [0,\infty)$, the closure of Γ^*, as the contract space.

Based on Lemma 1, I can write players' payoffs as functions of (w^1, w^0, t^1, t^0). To simplify exposition, I assume that $s_\alpha - \lambda^0 h_\alpha = s_\rho - \lambda^0 h_\rho$. The method illustrated below can be easily adjusted to solve for the optimal contract when $s_\alpha - \lambda^0 h_\alpha \neq s_\rho - \lambda^0 h_\rho$. Without loss of generality, I further assume that $s_\alpha - \lambda^0 h_\alpha = 1$. The principal's and the agent's payoffs given (w^1, w^0, t^1, t^0) and type θ are then respectively

$$\frac{\theta}{1+\theta}\left(\eta_\rho w^1 - t^1\right) - \frac{1}{1+\theta}\left(w^0 + t^0\right) \text{ and } \frac{\theta}{1+\theta}\left(\eta_\alpha w^1 + t^1\right) - \frac{1}{1+\theta}\left(w^0 - t^0\right).$$

Based on Lemma 1 and 2, I reformulate the contract problem. The principal simply offers a direct mechanism $(w^1, w^0, t^1, t^0) : \Theta \to \Gamma \times [0,\infty) \times [0,\infty)$, called a contract, such that

$$\begin{aligned} \max_{w^1,w^0,t^1,t^0} \quad & \int_\Theta \left(\frac{\theta}{1+\theta}\left(\eta_\rho w^1(\theta) - t^1(\theta)\right) - \frac{1}{1+\theta}\left(w^0(\theta) + t^0(\theta)\right)\right) \mathrm{d}F(\theta), \\ \text{s.t.} \quad & \theta\left(\eta_\alpha w^1(\theta) + t^1(\theta)\right) - w^0(\theta) + t^0(\theta) \geq \\ & \theta\left(\eta_\alpha w^1(\theta') + t^1(\theta')\right) - w^0(\theta') + t^0(\theta'), \quad \forall \theta, \theta' \in \Theta. \end{aligned}$$

Given a direct mechanism $(w^1(\theta), w^0(\theta), t^1(\theta), t^0(\theta))$, let $U_\alpha(\theta)$ denote the payoff that the agent of type θ gets by maximizing over his report, *i.e.*, $U_\alpha(\theta) = \max_{\theta' \in \Theta}\{\theta\left(\eta_\alpha w^1(\theta') + t^1(\theta')\right) - w^0(\theta') + t^0(\theta')\}$. As the optimal mechanism is truthful, $U_\alpha(\theta)$ equals $\theta(\eta_\alpha w^1(\theta) + t^1(\theta)) - w^0(\theta) + t^0(\theta)$ and the envelope condition implies that $U'_\alpha(\theta) = \eta_\alpha w^1(\theta) + t^1(\theta)$. Incentive compatibility of (w^1, w^0, t^1, t^0) requires that, for all θ

$$\left(t^0\right)'(\theta) = -\theta\left(\eta_\alpha\left(w^1\right)'(\theta) + \left(t^1\right)'(\theta)\right) + \left(w^0\right)'(\theta), \tag{2}$$

whenever differentiable, or in integral form,

$$t^0(\theta) = U_\alpha(\overline{\theta}) - \int_\theta^{\overline{\theta}} \left(\eta_\alpha w^1(s) + t^1(s)\right) \mathrm{d}s - \theta \left(\eta_\alpha w^1(\theta) + t^1(\theta)\right) + w^0(\theta).$$

Incentive compatibility also requires $\eta_\alpha w^1 + t^1$ to be a nondecreasing function of θ. Thus, (2) and the monotonicity of $\eta_\alpha w^1 + t^1$ are necessary for incentive compatibility. As is standard, these two conditions are also sufficient.

The principal's payoff for a fixed θ is denoted $U_\rho(\theta)$

$$U_\rho(\theta) = \frac{\theta}{1+\theta} \left(\eta_\rho w^1(\theta) - t^1(\theta)\right) - \frac{1}{1+\theta} \left(w^0(\theta) + t^0(\theta)\right).$$

The principal's problem is to maximize $\int_\Theta U_\rho(\theta)\mathrm{d}F$ subject to (i) (2) and the monotonicity of $\eta_\alpha w^1 + t^1$; (ii) the feasibility constraint $(w^1(\theta), w^0(\theta)) \in \Gamma, \forall\theta \in \Theta$; and (iii) the limited liability constraint (LL constraint, hereafter) $t^1(\theta), t^0(\theta) \geq 0, \forall\theta \in \Theta$. I denote this problem by $\mathcal{P}$. Substituting $t^0(\theta)$ into the objective, I rewrite $\int_\Theta U_\rho(\theta)\mathrm{d}F$ as

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[\frac{(\eta_\alpha + \eta_\rho)\theta w^1(\theta)}{1+\theta} - \frac{2w^0(\theta)}{1+\theta} + \frac{\left(\eta_\alpha w^1(\theta) + t^1(\theta)\right) H(\theta)}{f(\theta)}\right] f(\theta)\mathrm{d}\theta - U_\alpha(\overline{\theta})H(\overline{\theta}),$$

where $h(\theta) = \dfrac{f(\theta)}{1+\theta}$ and $H(\theta) = \displaystyle\int_{\underline{\theta}}^{\theta} h(s)\mathrm{d}s$.

I then define a relaxed problem $\mathcal{P}'$ which differs from $\mathcal{P}$ in two aspects: (i) the monotonicity of $\eta_\alpha w^1 + t^1$ is dropped; and (ii) the feasibility constraint is replaced with $w^0(\theta) \geq \beta^{\text{se}}(w^1(\theta)), \forall\theta \in \Theta$, where $\beta^{\text{se}}(\cdot)$ characterizes the southeast boundary of Γ. If the solution to $\mathcal{P}'$ is monotone and satisfies the feasibility constraint, it is also the solution to $\mathcal{P}$. The problem $\mathcal{P}'$ can be transformed into a control problem with the state $\boldsymbol{s} = (w^1, w^0, t^1, t^0)$ and the control $\boldsymbol{y} = (y^1, y^0, y_t^1)$. The associated costate is $\boldsymbol{\gamma} = (\gamma^1, \gamma^0, \gamma_t^1, \gamma_t^0)$. The law of motion is

$$\left(w^1\right)' = y^1, \left(w^0\right)' = y^0, \left(t^1\right)' = y_t^1, \left(t^0\right)' = y^0 - \theta\left(\eta_\alpha y^1 + y_t^1\right). \tag{3}$$

For problem $\mathcal{P}'$, I define a Lagrangian

$$\begin{aligned}\mathcal{L}(\boldsymbol{s}, \boldsymbol{y}, \boldsymbol{\gamma}, \boldsymbol{\mu}, \theta) = &\left[\frac{(\eta_\alpha + \eta_\rho)\theta w^1(\theta)}{1+\theta} - \frac{2w^0(\theta)}{1+\theta} + \frac{\left(\eta_\alpha w^1(\theta) + t^1(\theta)\right) H(\theta)}{f(\theta)}\right] f(\theta) \\ &+ \gamma^1(\theta)y^1(\theta) + \gamma^0(\theta)y^0(\theta) + \gamma_t^0(\theta)\left[y^0(\theta) - \theta\left(\eta_\alpha y^1(\theta) + y_t^1(\theta)\right)\right] \\ &+ \gamma_t^1(\theta)y_t^1(\theta) + \mu_t^1(\theta)t^1(\theta) + \mu_t^0 t^0 + \mu(\theta)\left[w^0(\theta) - \beta^{\text{se}}(w^1(\theta))\right],\end{aligned}$$

where $\boldsymbol{\mu} = (\mu_t^1, \mu_t^0, \mu)$ are multipliers associated with the LL and feasibility constraints. From now on, the dependence of $(\boldsymbol{s}, \boldsymbol{y}, \boldsymbol{\gamma}, \boldsymbol{\mu}, f, h, H)$ on θ is omitted when no confusion arises. Given any θ, the control maximizes the Lagrangian. The first-order conditions are,

$$\frac{\partial \mathcal{L}}{\partial y^1} = \gamma^1 - \eta_\alpha \theta \gamma_t^0 = 0, \quad \frac{\partial \mathcal{L}}{\partial y^0} = \gamma^0 + \gamma_t^0 = 0, \quad \frac{\partial \mathcal{L}}{\partial y_t^1} = \gamma_t^1 - \theta \gamma_t^0 = 0. \tag{4}$$

The costate variables are continuous and have piecewise-continuous derivatives,

$$\begin{aligned} \dot{\gamma}^1 &= -\frac{\partial \mathcal{L}}{\partial w^1} = -\left[\frac{\theta}{1+\theta}(\eta_\alpha + \eta_\rho) + \frac{\eta_\alpha H}{f}\right] f + \mu \left(\beta^{\text{se}}\right)' (w^1), \\ \dot{\gamma}^0 &= -\frac{\partial \mathcal{L}}{\partial w^0} = \frac{2}{1+\theta} f - \mu, \quad \dot{\gamma}_t^1 = -\frac{\partial \mathcal{L}}{\partial t^1} = -H - \mu_t^1, \quad \dot{\gamma}_t^0 = -\frac{\partial \mathcal{L}}{\partial t^0} = -\mu_t^0. \end{aligned} \tag{5}$$

This is a problem with free endpoint and a scrap value function $\Phi(\bar{\theta}) = -U_\alpha(\bar{\theta}) H(\bar{\theta})$. Therefore, the costate variables must satisfy the following boundary conditions,

$$\begin{aligned} (\gamma^1(\underline{\theta}), \gamma^0(\underline{\theta}), \gamma_t^1(\underline{\theta}), \gamma_t^0(\underline{\theta})) &= (0, 0, 0, 0), \\ (\gamma^1(\bar{\theta}), \gamma^0(\bar{\theta}), \gamma_t^1(\bar{\theta}), \gamma_t^0(\bar{\theta})) &= \left(\frac{\partial \Phi}{\partial w^1}, \frac{\partial \Phi}{\partial w^0}, \frac{\partial \Phi}{\partial t^1}, \frac{\partial \Phi}{\partial t^0}\right)\Bigg|_{\theta = \bar{\theta}} \\ &= \left(-\eta_\alpha \bar{\theta} H(\bar{\theta}), H(\bar{\theta}), -\bar{\theta} H(\bar{\theta}), -H(\bar{\theta})\right). \end{aligned} \tag{6}$$

Also, the following slackness conditions must be satisfied,

$$\begin{aligned} &\mu_t^1 \geq 0, t^1 \geq 0, \mu_t^1 t^1 = 0; \quad \mu_t^0 \geq 0, t^0 \geq 0, \mu_t^0 t^0 = 0; \\ &\mu \geq 0, w^0 - \beta^{\text{se}}(w^1) \geq 0, \mu\left(w^0 - \beta^{\text{se}}(w^1)\right) = 0. \end{aligned} \tag{7}$$

Lemma 3 (Necessity and sufficiency)**.**
Let $\boldsymbol{y}^*$ be the optimal control and $\boldsymbol{s}^*$ the corresponding trajectory. Then there exist costate variables $\boldsymbol{\gamma}^*$ and multipliers $\boldsymbol{\mu}^*$ such that (3)–(7) *are satisfied. Conversely,* (3)–(7) *are also sufficient since the Lagrangian is concave in $(\boldsymbol{s}, \boldsymbol{y})$.*

Based on the sufficiency part of Lemma 3, I only need to construct $(\boldsymbol{s}, \boldsymbol{y}, \boldsymbol{\gamma}, \boldsymbol{\mu})$ such that the conditions (3)–(7) are satisfied. In what follows, I first describe the optimal contract and then prove its optimality.

Optimal contract: description. I identify a bundle (w^1, w^0) on the southeast boundary of Γ with the derivative $\left(\beta^{\text{se}}\right)'(w^1)$ at that point. If the optimal contract only involves bundles on the southeast boundary, the trajectory $\left(\beta^{\text{se}}\right)'(w^1(\theta))$ uniquely determines the trajectory $(w^1(\theta), w^0(\theta))$ and vice versa. Figure 1 illustrates three important trajectories of $\left(\beta^{\text{se}}\right)'(w^1(\theta))$ which determine the optimal contract under certain regularity conditions. The x axis is the agent's type. The y axis indicates the slope of the tangent line at a certain bundle on the southeast boundary. The thick-dashed line (labeled T_2) corresponds to the slope (or the bundle) preferred by the agent for any given θ. The thin-dashed line (labeled T_3) shows the bundle preferred by the principal if she believes that the agent's type is above θ. The thin-solid line (labeled T_1) is the bundle given by the following equation

$$(\beta^{\text{se}})'\left(w^1(\theta)\right) = \frac{(\eta_\alpha + \eta_\rho)\theta}{2 - \frac{(1+\theta)H}{\theta f} + \frac{1+\theta}{\theta^2 f}\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s}. \tag{8}$$

Loosely speaking, it is the bundle that the principal would offer if the LL constraint were not bound. Besides these three trajectories, the dotted line shows the bundle preferred by the principal for any given θ. Let θ^* denote the type at which T_1 and T_2 intersects and θ_p the type at which T_2 and T_3 intersects. Equations (9) and (10) gives the formal definition of θ^* and θ_p. It is easy to verify that $\theta^* > \underline{\theta}$ and $\theta_p < \overline{\theta}$. Moreover, θ^* increases and θ_p decreases in η_α/η_ρ.

$$\theta^* := \sup\left\{\hat{\theta} \in \Theta : \frac{(\eta_\alpha + \eta_\rho)\theta}{2 - \frac{(1+\theta)H}{\theta f} + \frac{1+\theta}{\theta^2 f}\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s} < \eta_\alpha\theta, \ \ \forall\theta \le \hat{\theta}\right\}, \tag{9}$$

$$\theta_p := \inf\left\{\hat{\theta} \in \Theta : \frac{\eta_\rho \int_{\theta}^{\overline{\theta}} sh(s)\mathrm{d}s}{H(\overline{\theta}) - H(\theta)} \le \eta_\alpha\theta, \ \ \forall\theta \ge \hat{\theta}\right\}. \tag{10}$$

When the bias η_α/η_ρ is small, $\theta^* < \theta_p$. The optimal contract (the thick-solid line) consists of three separate segments, *i.e.*, $[\underline{\theta}, \theta^*]$, $[\theta^*, \theta_p]$, and $[\theta_p, \overline{\theta}]$. (See Figure 1.) When the agent's type is below θ^*, the equilibrium allocation is given by (8), which lies between that optimal for the principal $((\beta^{\text{se}})'(w^1(\theta)) = \eta_\rho\theta)$ and that optimal for the agent $((\beta^{\text{se}})'(w^1(\theta)) = \eta_\alpha\theta)$. As θ increase, the contract bundle shifts toward the agent's preferred bundle, with a corresponding decrease in the transfer payments. When $\theta \in [\theta^*, \theta_p]$, the bundle that is preferred by the agent is offered and no transfers are made. It is as if the agent is free to choose any experimentation policy. For types above θ_p, the agent always chooses the bundle preferred by type θ_p. There is, effectively, pooling over $[\theta_p, \overline{\theta}]$.

When the bias η_α/η_ρ is large, $\theta^* > \theta_p$. The optimal contract (the thick-solid line) consists of only two segments which are denoted $[\underline{\theta}, \tilde{\theta}_p]$ and $[\tilde{\theta}_p, \overline{\theta}]$. (See Figure 2.) When $\theta \in [\underline{\theta}, \tilde{\theta}_p]$, the equilibrium allocation is between the principal's preferred bundle and the agent's preferred one. The contract bundle shifts toward the agent's preferred one as the type increases with a corresponding decrease in the transfers. When $\theta \in [\tilde{\theta}_p, \overline{\theta}]$, all types are pooled. The pooling bundle specifies a lower level of experimentation than what the principal prefers given the pooling segment $[\tilde{\theta}_p, \overline{\theta}]$. There is no segment in which the agent implements his preferred bundle. When the bias η_α/η_ρ is sufficiently large, the principal prefers to pool all types.

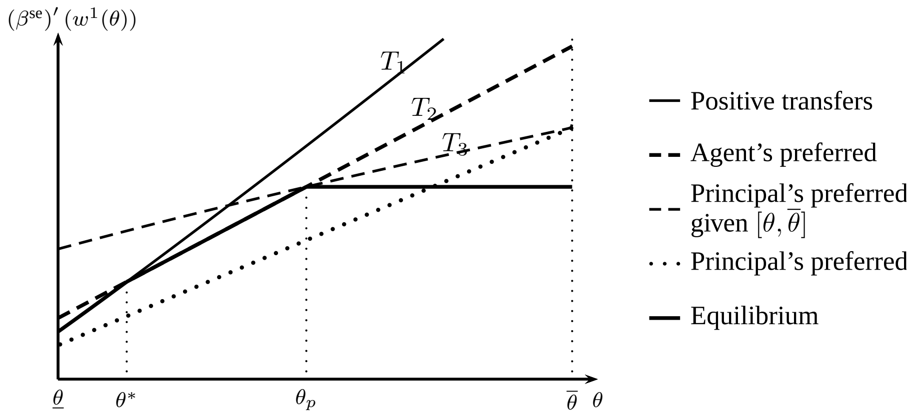

Figure 1: Equilibrium allocation: three segments

Parameters are $\eta_\alpha = 1, \eta_\rho = 4/5, \underline{\theta} = 1, \overline{\theta} = 3, f(\theta) = 1/2$.

One immediate observation is that the most pessimistic type's policy is socially optimal. The contract for type $\underline{\theta}$ is chosen such that $(\beta^{se})'(w^1(\underline{\theta})) = (\eta_\alpha + \eta_\rho)\underline{\theta}/2$. A key feature is that transfers only occur in state 1. The intuition can be seen by examining a simple example with three types $\theta_l < \theta_m < \theta_h$. To prevent the medium type from mimicking the high type, the principal compensates the medium type by promising him a positive transfer. This transfer promise makes the medium type's contract more attractive to the low type. To make the medium type's transfer less attractive to the low type, the principal concentrates all the payments in state 1 as the low type is less confident that the state is 1. Whenever the principal promises type θ a positive transfer, she makes type θ's contract more attractive to a lower type, say $\theta' < \theta$. As type θ' is less confident that the state is 1 than type θ, type θ' does not value transfers in state 1 as much as type θ is. Therefore, the most efficient way to make transfers is to condition on state being 1.

Optimal contract: proof. I start with the case when the bias is small and the contract has three segments. The proof is constructive. I first determine the trajectory of the costate γ^0, which pins down $\gamma^1, \gamma^1_t, \gamma^0_t$ according to (4). Then I determine the trajectories of $\mu, \mu^1_t, \mu^0_t, w^1$ based on (5). The trajectories of w^0, t^1, t^0 then follow.

On the interval $[\underline{\theta}, \theta^*]$, $t^1 > 0$ and $t^0 = 0$, so the LL constraint $t^1 \geq 0$ does not bind. Therefore, I have $\mu^1_t = 0$ and $\dot{\gamma}^1_t = -H$. Combined with the boundary condition, this implies that $\gamma^1_t = -\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s$. From (4), we know that $\gamma^0 = \int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s/\theta$ and $\gamma^1 = -\eta_\alpha \int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s$. Substituting $\dot{\gamma}^1$ and $\dot{\gamma}^0$ into (5), I have

$$\mu = \frac{2f}{1+\theta} - \frac{H(\theta)}{\theta} + \frac{\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s}{\theta^2} \text{ and } (\beta^{se})'(w^1(\theta)) = \frac{(\eta_\alpha + \eta_\rho)\theta}{2 - \frac{(1+\theta)H}{\theta f} + \frac{1+\theta}{\theta^2 f}\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s}.$$

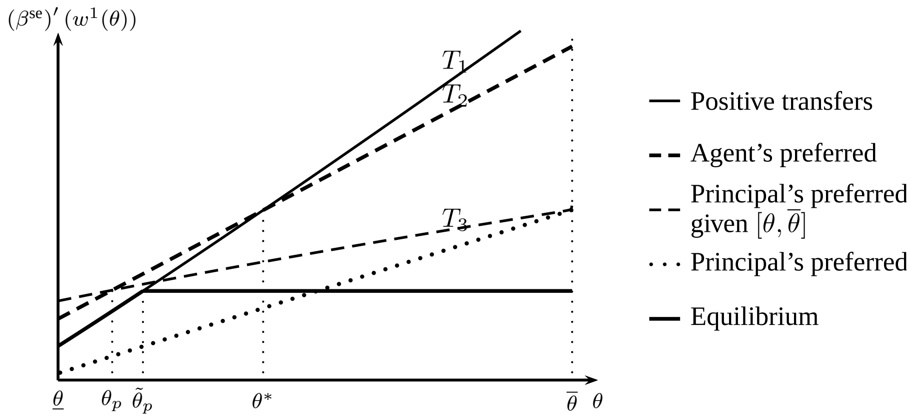

Figure 2: Equilibrium allocation: two segments

Parameters are $\eta_\alpha = 1, \eta_\rho = 3/5, \underline{\theta} = 1, \overline{\theta} = 3, f(\theta) = 1/2$.

Since $\mu_t^0 = -\dot{\gamma}_t^0 = \dot{\gamma}^0$, I have $\mu_t^0 = -\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s/(\theta^2) + H(\theta)/\theta$, which is always positive.

On the interval $[\theta^*, \theta_p]$, the type θ is assigned his most preferred bundle. Transfers t^1 and t^0 both equal zero. Therefore, I have $(\beta^{se})'(w^1(\theta)) = \eta_\alpha\theta$. From (4), we know that $\gamma^0 = -\dot{\gamma}^1/\eta_\alpha - \theta\dot{\gamma}^0$. Substituting $\dot{\gamma}^1$, $\dot{\gamma}^0$ and $(\beta^{se})'(w^1(\theta)) = \eta_\alpha\theta$, I have

$$\gamma^0 = H - \frac{\theta}{1+\theta}\frac{\eta_\alpha - \eta_\rho}{\eta_\alpha} f. \tag{11}$$

Combining (5) and (11), I have

$$\begin{aligned}
\mu &= \frac{f}{1+\theta}\left(1 + \frac{\eta_\alpha - \eta_\rho}{\eta_\alpha}\frac{1}{1+\theta} + \frac{\eta_\alpha - \eta_\rho}{\eta_\alpha}\frac{f'}{f}\theta\right) \\
\mu_t^0 &= \frac{f}{1+\theta}\left(1 - \frac{\eta_\alpha - \eta_\rho}{\eta_\alpha}\frac{1}{1+\theta} - \frac{\eta_\alpha - \eta_\rho}{\eta_\alpha}\frac{f'}{f}\theta\right) \\
\mu_t^1 &= \frac{\theta}{1+\theta} f\left(1 - \frac{\eta_\alpha - \eta_\rho}{\eta_\alpha}\frac{2+\theta}{1+\theta} - \frac{\eta_\alpha - \eta_\rho}{\eta_\alpha}\frac{f'}{f}\theta\right).
\end{aligned}$$

The multipliers μ, μ_t^0 and μ_t^1 have to be weakly positive, which requires that

$$\frac{\eta_\alpha}{\eta_\alpha - \eta_\rho} \geq -\frac{1}{1+\theta} - \frac{f'}{f}\theta, \quad \forall\theta \in [\theta^*, \theta_p] \tag{A1}$$

$$\frac{\eta_\alpha}{\eta_\alpha - \eta_\rho} \geq \frac{2+\theta}{1+\theta} + \frac{f'}{f}\theta, \quad \forall\theta \in [\theta^*, \theta_p]. \tag{A2}$$

Note that Assumption A1 is the same as the main assumption in the delegation case.

On the interval $[\theta_p, \overline{\theta}]$, all types choose the same bundle, the one preferred by type θ_p. Transfers

t^1 and t^0 both equal zero. The threshold of the pooling segment θ_p satisfies the following condition,

$$(\beta^{\text{se}})'(w^1(\theta_p)) = \theta_p \eta_\alpha = \frac{\eta_\rho \int_{\theta_p}^{\overline{\theta}} \theta h(\theta) \mathrm{d}\theta}{H(\overline{\theta}) - H(\theta_p)}.$$

I first check that the boundary condition $\gamma^1(\overline{\theta}) = -\eta_\alpha H(\overline{\theta})\overline{\theta}$ is satisfied. Over the interval $[\theta_p, \overline{\theta}]$, I have

$$\dot{\gamma}^1 = -\left[\frac{\theta}{1+\theta}(\eta_\alpha + \eta_\rho) f + \eta_\alpha H\right] + \mu\, (\beta^{\text{se}})'(w^1(\theta_p)).$$

Given the definition of θ_p, it is easy to verify that

$$\begin{aligned}
(\beta^{\text{se}})'(w^1(\theta_p)) \int_{\theta_p}^{\overline{\theta}} \mu \mathrm{d}\theta &= (\beta^{\text{se}})'(w^1(\theta_p)) \int_{\theta_p}^{\overline{\theta}} \left(\frac{2f}{1+\theta} - \dot{\gamma}^0\right) \mathrm{d}\theta \\
&= \eta_\rho \int_{\theta_p}^{\overline{\theta}} \theta h(\theta) \mathrm{d}\theta - (\eta_\alpha - \eta_\rho) \frac{\theta_p^2}{1+\theta_p} f(\theta_p).
\end{aligned}$$

Therefore, I have

$$\begin{aligned}
\gamma^1(\overline{\theta}) - \gamma^1(\theta_p) &= -\eta_\alpha \int_{\theta_p}^{\overline{\theta}} \theta h(\theta) d\theta - \eta_\alpha \int_{\theta_p}^{\overline{\theta}} H(\theta) \mathrm{d}\theta - (\eta_\alpha - \eta_\rho) \frac{\theta_p^2}{1+\theta_p} f(\theta_p) \\
&= -H(\overline{\theta})\overline{\theta}\eta_\alpha - \gamma^1(\theta_p).
\end{aligned}$$

Therefore, the boundary condition $\gamma^1(\overline{\theta}) = -\eta_\alpha H(\overline{\theta})\overline{\theta}$ is satisfied. The slackness condition $\mu \geq 0$ and $\mu_t^0 \geq 0$ requires that $0 \leq \dot{\gamma}^0 \leq 2f/(1+\theta)$. This is equivalent to the condition that $0 \leq \gamma^0(\overline{\theta}) - \gamma^0(\theta_p) \leq 2\left(H(\overline{\theta}) - H(\theta_p)\right)$, which is satisfied iff

$$\frac{\eta_\alpha}{\eta_\alpha - \eta_\rho} \geq \frac{\frac{\theta_p}{1+\theta_p} f(\theta_p)}{H(\overline{\theta}) - H(\theta_p)}. \tag{A3}$$

The slackness condition $\mu_t^1 \geq 0$ requires that $\dot{\gamma}_t^1 \leq -H$. This is equivalent to the condition that

$$\gamma_t^1(\overline{\theta}) - \gamma_t^1(\theta_p) = -\overline{\theta} H(\overline{\theta}) + \theta_p H(\theta_p) - \frac{\theta_p^2}{1+\theta_p} \frac{\eta_\alpha - \eta_\rho}{\eta_\alpha} f(\theta_p) \leq -\int_{\theta_p}^{\overline{\theta}} H(s) \mathrm{d}s,$$

which is always satisfied.

To sum, if Assumptions A1, A2 and A3 hold, the constructed trajectory solves $\mathcal{P}'$. If the trajectory w^1 defined by (8) is weakly increasing over $[\underline{\theta}, \theta^*]$, the monotonicity of $\eta_\alpha w^1 + t^1$ is satisfied.[2] Therefore, the constructed trajectory solves $\mathcal{P}$ as well.

When the bias is large and the contract has two segments, the proof is similar to the previous case. So, I mainly explain how to pin down the threshold $\tilde{\theta}_p$. When $\theta \in [\underline{\theta}, \tilde{\theta}_p]$, the LL constraint $t^1 \geq 0$ does not bind. The costate is derived in the same way as in the previous case when $\theta \in [\underline{\theta}, \theta^*]$. This implies that $\gamma^1(\tilde{\theta}_p) = -\eta_\alpha \int_{\underline{\theta}}^{\tilde{\theta}_p} H(s)\mathrm{d}s$. On the other hand, the threshold $\tilde{\theta}_p$ is chosen so that the boundary condition $\gamma^1(\overline{\theta}) = -\eta_\alpha H(\overline{\theta})\overline{\theta}$ is satisfied. This means that $\int_{\tilde{\theta}_p}^{\overline{\theta}} \dot{\gamma}^1 \mathrm{d}\theta = -\eta_\alpha H(\overline{\theta})\overline{\theta} + \eta_\alpha \int_{\underline{\theta}}^{\tilde{\theta}_p} H(s)\mathrm{d}s$. Substituting $\dot{\gamma}^1$ and simplifying, I obtain that $(\beta^{\text{se}})'(w^1(\tilde{\theta}_p))$ must satisfy the following condition

$$(\beta^{\text{se}})'(w^1(\tilde{\theta}_p)) = \frac{\eta_\rho \int_{\tilde{\theta}_p}^{\overline{\theta}} \theta h(\theta)\mathrm{d}\theta - \eta_\alpha \tilde{\theta}_p H(\tilde{\theta}_p) + \eta_\alpha \int_{\underline{\theta}}^{\tilde{\theta}_p} H(s)\mathrm{d}s}{H(\overline{\theta}) - 2H(\tilde{\theta}_p) + \frac{\int_{\underline{\theta}}^{\tilde{\theta}_p} H(s)\mathrm{d}s}{\tilde{\theta}_p}} \tag{12}$$

At the same time, $\tilde{\theta}_p$ must also satisfy (8). Equation (8) and (12) determines the threshold $\tilde{\theta}_p$. Since $(\beta^{\text{se}})'(w^1(\tilde{\theta}_p)) < \eta_\alpha \tilde{\theta}_p$, (12) implies that

$$(\beta^{\text{se}})'(w^1(\tilde{\theta}_p)) < \frac{\eta_\rho \int_{\tilde{\theta}_p}^{\overline{\theta}} \theta h(\theta)\mathrm{d}\theta}{H(\overline{\theta}) - H(\tilde{\theta}_p)}.$$

This shows that over the pooling region $[\tilde{\theta}_p, \overline{\theta}]$ the agent is asked to implement a bundle with less experimentation than what the principal prefers given that $\theta \in [\tilde{\theta}_p, \overline{\theta}]$.

[2] Given that t^0 is constantly zero, I have $\eta_\alpha \left(w^1\right)'(\theta) + \left(t^1\right)'(\theta) = \left(w^0\right)'(\theta)/\theta$. Therefore, the monotonicity of w^0 implies the monotonicity of $\eta_\alpha w^1 + t^1$. Given that only boundary bundles (w^1, w^0) are assigned, the monotonicity of w^1 suffices.