DYNAMIC DELEGATION OF EXPERIMENTATION

YINGNI GUO

ONLINE APPENDIX

B1. A special case: two types

In this subsection, I study the delegation problem with binary types, high type θ_h and low type θ_l. Let $q(\theta)$ denote the probability that the agent's type is θ. Formally, I solve for $(w^1, w^0) : \{\theta_l, \theta_h\} \to \Gamma$ such that

$$\max_{w^1,w^0} \sum_{\theta\in\{\theta_l,\theta_h\}} q(\theta)\left(\frac{\theta}{1+\theta}\eta_\rho w^1(\theta) - \frac{1}{1+\theta}w^0(\theta)\right),$$
$$\text{subject to} \quad \theta_l\eta_\alpha w^1(\theta_l) - w^0(\theta_l) \geq \theta_l\eta_\alpha w^1(\theta_h) - w^0(\theta_h),$$
$$\theta_h\eta_\alpha w^1(\theta_h) - w^0(\theta_h) \geq \theta_h\eta_\alpha w^1(\theta_l) - w^0(\theta_l).$$

For ease of exposition, I refer to the contract for the low (high) type agent as the low (high) type contract and the principal who believes to face the low (high) type agent as the low (high) type principal. Let $(w^{1*}(\theta_l), w^{0*}(\theta_l))$ and $(w^{1*}(\theta_h), w^{0*}(\theta_h))$ denote the equilibrium bundles. The optimum is characterized as follows.

PROPOSITION 8 (Two types):
Suppose that $(r+\lambda^1)\theta_l\eta_\rho/r > 1$. There exists a $b' \in (1, \theta_h/\theta_l)$ such that

1.1 If $\eta_\alpha/\eta_\rho \in [1, b']$, the principal's preferred bundles are implementable.

1.2 If $\eta_\alpha/\eta_\rho \in (b', \theta_h/\theta_l)$, separating is optimal. The low type contract is a stopping-time policy, with the stopping time between $\tau_\rho(\theta_l)$ and $\tau_\alpha(\theta_l)$. The low type's IC constraint binds, and the high type's does not.

1.3 If $\eta_\alpha/\eta_\rho \geq \theta_h/\theta_l$, pooling is optimal.

In all cases, the optimum can be attained using bundles on the boundary of Γ.

The presumption $(r+\lambda^1)\theta_l\eta_\rho/r > 1$ ensures that both the low type principal's preferred stopping time $\tau_\rho(\theta_l)$ is strictly positive. The degenerate cases in which $\tau_\rho(\theta_h) > \tau_\rho(\theta_l) = 0$ or $\tau_\rho(\theta_h) = \tau_\rho(\theta_l) = 0$ yield similar results to proposition 8, and thus are omitted.

Proposition 8 describes the optimal contract as the bias level varies. According to result (*1.1*), if the bias is low enough, the principal simply offers her preferred policies given θ_l and θ_h. This is incentive compatible because, at a low bias level, the low type agent prefers the low type principal's preferred bundle instead of the high type principal's. Consequently the principal pays no information rents. This result does not hold with a continuum of types. The principal's preferred bundles are two points on the southeast boundary of Γ with binary types, but they become an interval on the southeast boundary with a continuum of types in which case lower types are strictly better off mimicking higher types.

The result (*1.2*) corresponds to medium bias level. As the bias has increased, offering the principal's preferred policies is no longer incentive compatible. Instead, both the low type contract and the high type one deviate from the principal's preferred policies. The low type contract is always a stopping-time policy while the high type contract takes one of three possible forms: stopping-time, slack-after-success or delay policies.[26] One of the latter two forms is assigned as the high type contract if the agent's type is likely to be low, and his bias is relatively large. All three forms are meant to impose a significant cost—excessive experimentation, constrained exploitation of success, or delay in experimentation—on the high type contract so as to deter the low type agent from misreporting. However, the principal can more than offset the cost by effectively shortening the low type agent's experimentation. In the end, the low type agent over-experiments slightly and the high type contract deviates from the principal's preferred policy $(w^1_\rho(\theta_h), w^0_\rho(\theta_h))$ as well. One interesting observation is that the optimal contract can take a form other than a stopping-time policy.

If the bias is even higher, as shown by result (*1.3*), pooling is preferable. The condition $\eta_\alpha/\eta_\rho \geq \theta_h/\theta_l$ has an intuitive interpretation that the low type agent prefers to experiment longer than even the high type principal. The screening instruments utilized in result (*1.2*) impair the high type principal's payoff more than the low type agent's. As a result, the principal is better off offering her uninformed preferred bundle. For fixed types, the prior probabilities of the types do not affect whether it is better to pool or separate. Only the bias level does.

I make two observations. First, the principal chooses to take advantage of the agent's private information unless the agent's bias is too large. This result also applies to the continuous type case. Second, the optimal contract can be tailored to the likelihood of the two types. For example, if the type is likely to be low, the principal designs the low type contract close to her low type bundle and purposefully makes the high type contract less attractive to the low type agent. Similarly, if the type is likely to be high, the principal starts with a high type contract close to her high type bundle without concerning about the low type's over-experimentation. This "type targeting", however, becomes irrelevant when the principal faces a continuum of types and has no incentives to target certain types.

PROOF OF PROPOSITION 8:

Let α_l (or α_h) denote the low (or high) type agent and ρ_l (or ρ_h) the low (or high) type principal. Given that $\theta_h > \theta_l$ and $\eta_\alpha > \eta_\rho$, the slopes of players' indifference curves are ranked as follows

$$\theta_h\eta_\alpha > \max\{\theta_h\eta_\rho, \theta_l\eta_\alpha\} \geq \min\{\theta_h\eta_\rho, \theta_l\eta_\alpha\} > \theta_l\eta_\rho.$$

[26] Here, I give an example in which the high type contract is a slack-after-success policy. Parameters are $\eta_\alpha = 6, \eta_\rho = 1, \theta_l = 3/2, \theta_h = 19, r = \lambda^1 = 1$. The agent's type is low with probability 2/3. The optimum is $(w^{1*}(\theta_h), w^{0*}(\theta_h)) \approx (0.98, 1)$ and $(w^{1*}(\theta_l), w^{0*}(\theta_l)) \approx (0.96, 0.79)$.

Let ICL and ICH denote α_l's and α_h's IC constraints. Let I_{α_l} denote α_l's indifference curves. If α_l prefers $(w_\rho^1(\theta_l), w_\rho^0(\theta_l))$ to $(w_\rho^1(\theta_h), w_\rho^0(\theta_h))$, the optimum is

$$\{(w_\rho^1(\theta_l), w_\rho^0(\theta_l)), (w_\rho^1(\theta_h), w_\rho^0(\theta_h))\}.$$

This is true when the slope of the line connecting $(w_\rho^1(\theta_l), w_\rho^0(\theta_l))$ and $(w_\rho^1(\theta_h), w_\rho^0(\theta_h))$ is greater than $\theta_l\eta_\alpha$. This condition is satisfied when η_α/η_ρ is bounded from above by

$$b' \equiv \frac{\theta_h(\lambda^1 + r)\left({\theta_h}^{\frac{r}{\lambda^1}} - {\theta_l}^{\frac{r}{\lambda^1}}\right)}{r\left({\theta_h}^{\frac{r+\lambda^1}{\lambda^1}} - {\theta_l}^{\frac{r+\lambda^1}{\lambda^1}}\right)}.$$

If this condition does not hold, at least one IC constraint binds. I explain how to find the optimal bundles as follows.

1) ICL binds. Suppose not. It must be the case that ICH binds and that the principal offers two distinct bundles $(w^1(\theta_l), w^0(\theta_l)) < (w^1(\theta_h), w^0(\theta_h))$ which lie on the same indifference curve of α_h. Given that $\theta_h\eta_\alpha > \max\{\theta_h\eta_\rho, \theta_l\eta_\rho\}$, both ρ_h and ρ_l strictly prefer $(w^1(\theta_l), w^0(\theta_l))$ to $(w^1(\theta_h), w^0(\theta_h))$. The principal is strictly better off by offering a pooling bundle $(w^1(\theta_l), w^0(\theta_l))$. A contradiction. Hence, ICL binds.

2) If $\theta_h\eta_\rho \leq \theta_l\eta_\alpha$, the optimum is pooling. Suppose not. Suppose that the principal offers two distinct bundles $(w^1(\theta_l), w^0(\theta_l)) < (w^1(\theta_h), w^0(\theta_h))$ which are on the same indifference curve of α_l. Given that $\theta_l\eta_\rho < \theta_h\eta_\rho < \theta_l\eta_\alpha$, α_l's indifference curves are steeper than ρ_h's and ρ_l's. Both ρ_h and ρ_l strictly prefer $(w^1(\theta_l), w^0(\theta_l))$ to $(w^1(\theta_h), w^0(\theta_h))$. The principal is strictly better off by offering a pooling bundle $(w^1(\theta_l), w^0(\theta_l))$. A contradiction. If $\theta_h\eta_\rho = \theta_l\eta_\alpha$, ρ_h has the same indifference curves as α_l. If $\{(w^1(\theta_l), w^0(\theta_l)), (w^1(\theta_h), w^0(\theta_h))\}$ is optimal, it is optimal for the principal to offer a pooling contract $(w^1(\theta_l), w^0(\theta_l))$.

3) If $\theta_h\eta_\rho > \theta_l\eta_\alpha$, the optimum are on the boundary of Γ. Suppose not. Suppose that $(w^1(\theta_l), w^0(\theta_l))$ or $(w^1(\theta_h), w^0(\theta_h))$ is in the interior. The indifference curve of α_l going through $(w^1(\theta_l), w^0(\theta_l))$ intersects the boundary at $(\tilde{w}^1(\theta_l), \tilde{w}^0(\theta_l))$ and $(\tilde{w}^1(\theta_h), \tilde{w}^0(\theta_h))$ such that $\tilde{w}^1(\theta_l) < \tilde{w}^1(\theta_h)$. Given that $\theta_h\eta_\rho > \theta_l\eta_\alpha > \theta_l\eta_\rho$, ρ_h prefers $(\tilde{w}^1(\theta_h), \tilde{w}^0(\theta_h))$ to $(w^1(\theta_h), w^0(\theta_h))$ and ρ_l prefers $(\tilde{w}^1(\theta_l), \tilde{w}^0(\theta_l))$ to $(w^1(\theta_l), w^0(\theta_l))$. The principal is strictly better off by offering $(\tilde{w}^1(\theta_l), \tilde{w}^0(\theta_l))$ and $(\tilde{w}^1(\theta_h), \tilde{w}^0(\theta_h))$. Therefore, the optimal bundles are on the boundary. The problem is reduced to locate the low type agent's indifference curve on which $(w^{1*}(\theta_l), w^{0*}(\theta_l))$ and $(w^{1*}(\theta_h), w^{0*}(\theta_h))$ lie. This indifference curve must be between the indifference curves of α_l which go through $(w_\rho^1(\theta_l), w_\rho^0(\theta_l))$ and $(w_\rho^1(\theta_h), w_\rho^0(\theta_h))$.

■

B2. Discussion and necessity of assumption 1

In this subsection, I show that no x_p-cutoff contract is optimal for any $x_p \in \Theta$ if assumption 1 does not hold. The x_p-cutoff contract is defined as $(w^1(\theta), w^0(\theta)) = (w^1_\alpha(\theta), w^0_\alpha(\theta))$ for $\theta < x_p$ and $(w^1(\theta), w^0(\theta)) = (w^1_\alpha(x_p), w^0_\alpha(x_p))$ for $\theta \geq x_p$. The x_p-cutoff contract is denoted $(w^1_{x_p}, \underline{w}^0_{x_p})$.

Define the Lagrangian functional associated with $\mathcal{P}$ as

$$\text{(B1)}\quad L(w^1, \underline{w}^0 \mid \Lambda^{\text{se}}, \Lambda^{\text{nw}}) = \underline{\theta}\eta_\alpha w^1(\underline{\theta}) - w^0(\underline{\theta}) + \eta_\alpha \int_{\underline{\theta}}^{\overline{\theta}} w^1(\theta)G(\theta)\mathrm{d}\theta$$
$$+ \int_{\underline{\theta}}^{\overline{\theta}} \left(\theta\eta_\alpha w^1(\theta) - \eta_\alpha \int_{\underline{\theta}}^{\theta} w^1(\tilde{\theta})\mathrm{d}\tilde{\theta} - \underline{\theta}\eta_\alpha w^1(\underline{\theta}) + \underline{w}^0 - \beta^{\text{se}}(w^1(\theta))\right) \mathrm{d}\Lambda^{\text{se}}$$
$$+ \int_{\underline{\theta}}^{\overline{\theta}} \left[\beta^{\text{nw}}(w^1(\theta)) - \left(\theta\eta_\alpha w^1(\theta) - \eta_\alpha \int_{\underline{\theta}}^{\theta} w^1(\tilde{\theta})\mathrm{d}\tilde{\theta} - \underline{\theta}\eta_\alpha w^1(\underline{\theta}) + \underline{w}^0\right)\right] d\Lambda^{\text{nw}},$$

where the function $\Lambda^{\text{se}}, \Lambda^{\text{nw}}$ are the Lagrange multiplier associated with constraints (9) and (10). I first show that if $(w^1_{x_p}, \underline{w}^0_{x_p})$ is optimal for some x_p, there must exist some Lagrange multipliers $\tilde{\Lambda}^{\text{se}}, \tilde{\Lambda}^{\text{nw}}$ such that $L(w^1, \underline{w}^0 \mid \tilde{\Lambda}^{\text{se}}, \tilde{\Lambda}^{\text{nw}})$ is maximized at $(w^1_{x_p}, \underline{w}^0_{x_p})$. Since any x_p-cutoff contract is continuous, I can restrict attention to the set of continuous contracts

$$\hat{\Phi} \equiv \left\{w^1, \underline{w}^0 \mid w^1 : \Theta \to [0,1],\ w^1 \text{ is nondecreasing and continuous},\ \underline{w}^0 \in [0,1]\right\}.$$

LEMMA 7 (Lagrangian—necessity):
If $(w^1_{x_p}, \underline{w}^0_{x_p})$ solves $\mathcal{P}$, there exist nondecreasing functions $\tilde{\Lambda}^{se}, \tilde{\Lambda}^{nw} : \Theta \to \mathbf{R}$ such that

(B2)
$$0 = \int_{\underline{\theta}}^{\overline{\theta}} \left(\theta\eta_\alpha w^1_{x_p}(\theta) - \eta_\alpha \int_{\underline{\theta}}^{\theta} w^1_{x_p}(\tilde{\theta})\, d\tilde{\theta} - \underline{\theta}\eta_\alpha w^1_{x_p}(\underline{\theta}) + \underline{w}^0_{x_p} - \beta^{se}(w^1_{x_p}(\theta))\right) d\tilde{\Lambda}^{se}$$
$$+ \int_{\underline{\theta}}^{\overline{\theta}} \left[\beta^{nw}(w^1_{x_p}(\theta)) - \left(\theta\eta_\alpha w^1_{x_p}(\theta) - \eta_\alpha \int_{\underline{\theta}}^{\theta} w^1_{x_p}(\tilde{\theta})\, d\tilde{\theta} - \underline{\theta}\eta_\alpha w^1_{x_p}(\underline{\theta}) + \underline{w}^0_{x_p}\right)\right] d\tilde{\Lambda}^{nw}.$$

Furthermore, it is the case that

$$L(w^1_{x_p}, \underline{w}^0_{x_p} \mid \tilde{\Lambda}^{se}, \tilde{\Lambda}^{nw}) \geq L(w^1, \underline{w}^0 \mid \tilde{\Lambda}^{se}, \tilde{\Lambda}^{nw}),\ \ \forall (w^1, \underline{w}^0) \in \hat{\Phi}.$$

PROOF:

I first introduce the problem studied in section 8.4 of Luenberger, 1996, p. 217: $\max_{x \in X} Q(x)$ subject to $x \in \Omega$ and $J(x) \in P$, where Ω is a convex subset of the vector space X, $Q : \Omega \to \mathbf{R}$ and $J : \Omega \to Z$ are both concave; where Z is a normed vector space, and P is a nonempty positive cone in Z. To apply Theorem 1 in

Luenberger, 1996, p. 217, set

$$\begin{aligned}
&X = \{w^1, \underline{w}^0 \mid \underline{w} \in \mathbf{R} \text{ and } w^1 : \Theta \to \mathbf{R}\},\\
&\Omega = \hat{\Phi},\\
&Z = \{z \mid z : \Theta \to \mathbf{R}^2 \text{ with } \sup_{\theta\in\Theta} \|z(\theta)\| < \infty\},\\
&\quad \text{with the norm} \|z\| = \sup_{\theta\in\Theta} \|z(\theta)\|,\\
&P = \{z \mid z \in Z \text{ and } z(\theta) \geq (0,0),\ \forall\theta \in \Theta\}.
\end{aligned}$$

I let the objective function in (OBJ) be Q and the left-hand side of (9) and (10) be defined as J. It is easy to verify that both Q and J are concave. This result holds because the hypotheses of Theorem 1 in Luenberger, 1996, p. 217 are met. ∎

My next result shows that no x_p-cutoff contract is optimal if assumption 1 fails.

PROPOSITION 9:
If assumption 1 does not hold, then no x_p-cutoff contract is optimal for any $x_p \in \Theta$.

PROOF:
The proof proceeds by contradiction. Suppose that $(w^1_{x_p}, \underline{w}^0_{x_p})$ is optimal for some $x_p \in \Theta$. According to lemma 7, there exist nondecreasing $\tilde{\Lambda}^{\mathrm{se}}, \tilde{\Lambda}^{\mathrm{nw}}$ such that the Lagrangian (B1) is maximized at $(w^1_{x_p}, \underline{w}^0_{x_p})$ and (B2) holds. This implies that $\tilde{\Lambda}^{\mathrm{nw}}$ is constant so the integral related to $\tilde{\Lambda}^{\mathrm{nw}}$ can be dropped. Without loss of generality I set $\tilde{\Lambda}^{\mathrm{se}}(\overline{\theta}) = 1$. Integrating the Lagrangian by parts yields

$$L(w^1, \underline{w}^0 \mid \tilde{\Lambda}^{\mathrm{se}}) = \left(\underline{\theta}\eta_\alpha w^1(\underline{\theta}) - \underline{w}^0\right) \tilde{\Lambda}^{\mathrm{se}}(\underline{\theta}) + \int_{\underline{\theta}}^{\overline{\theta}} \left(\theta\eta_\alpha w^1(\theta) - \beta^{\mathrm{se}}(w^1(\theta))\right) \mathrm{d}\tilde{\Lambda}^{\mathrm{se}}$$
$$+ \eta_\alpha \int_{\underline{\theta}}^{\overline{\theta}} w^1(\theta) \left[\tilde{\Lambda}^{\mathrm{se}}(\theta) - (1 - G(\theta))\right] \mathrm{d}\theta.$$

Then, I establish the necessary first-order conditions for $L(w^1, \underline{w}^0 \mid \tilde{\Lambda}^{\mathrm{se}})$ to be maximized at x_p-cutoff rule and show that they cannot be satisfied if assumption 1 fails.

Let $a, b \in \Theta$ be such that $a < b < \theta_p$ and $1-G(a) > 1-G(b)$ (so assumption 1 does not hold). It is easy to verify that the Gâteaux differential $\partial L(w^1_{x_p}, \underline{w}^0_{x_p}; w^1, \underline{w}^0 \mid \tilde{\Lambda}^{\mathrm{se}})$ exists for any $(w^1, \underline{w}^0) \in \hat{\Phi}$. I want to show that a necessary condition that $(w^1_{x_p}, \underline{w}^0_{x_p})$ maximizes $L(w^1, \underline{w}^0 \mid \tilde{\Lambda}^{\mathrm{se}})$ over $\hat{\Phi}$ is that

(B3) $$\partial L(w^1_{x_p}, \underline{w}^0_{x_p}; w^1, \underline{w}^0 \mid \tilde{\Lambda}^{\mathrm{se}}) \leq 0,\ \forall (w^1, \underline{w}^0) \in \hat{\Phi},$$

(B4) $$\partial L(w^1_{x_p}, \underline{w}^0_{x_p}; w^1_{x_p}, \underline{w}^0_{x_p} \mid \tilde{\Lambda}^{\mathrm{se}}) = 0.$$

If $(w^1_{x_p}, \underline{w}^0_{x_p})$ maximizes $L(w^1, \underline{w}^0 \mid \tilde{\Lambda}^{\mathrm{se}})$, then for any $(w^1, \underline{w}^0) \in \hat{\Phi}$, it must be true that

$$\frac{d}{d\epsilon}L((w^1_{x_p},\underline{w}^0_{x_p})+\epsilon((w^1,\underline{w}^0)-(w^1_{x_p},\underline{w}^0_{x_p}))\mid\tilde{\Lambda}^{\text{se}})\Big|_{\epsilon=0}\leq 0.$$

Hence, $\partial L(w^1_{x_p},\underline{w}^0_{x_p};(w^1,\underline{w}^0)-(w^1_{x_p},\underline{w}^0_{x_p})\mid\tilde{\Lambda}^{\text{se}})\leq 0$. Setting $(w^1,\underline{w}^0)=(w^1_{x_p},\underline{w}^0_{x_p})/2\in\hat{\Phi}$ yields $\partial L(w^1_{x_p},\underline{w}^0_{x_p};w^1_{x_p},\underline{w}^0_{x_p}\mid\tilde{\Lambda}^{\text{se}})\geq 0$. By the definition of $(w^1_{x_p},\underline{w}^0_{x_p})$, there exists $\epsilon>0$ sufficiently small such that $(1+\epsilon)(w^1_{x_p},\underline{w}^0_{x_p})\in\hat{\Phi}$. Setting $(w^1,\underline{w}^0)$ to be $(1+\epsilon)(w^1_{x_p},\underline{w}^0_{x_p})$ yields $\partial L(w^1_{x_p},\underline{w}^0_{x_p};w^1_{x_p},\underline{w}^0_{x_p}\mid\tilde{\Lambda}^{\text{se}})\leq 0$. Together, (B3) and (B4) obtain.

The last step is to show that there exists no $\tilde{\Lambda}^{\text{se}}$ that satisfies the first-order conditions (B3) and (B4). Here, I use the same approach as in the proof of proposition 4 in Amador, Werning and Angeletos (2006). The Gâteaux differential $\partial L(w^1_{x_p},\underline{w}^0_{x_p};w^1,\underline{w}^0\mid\tilde{\Lambda}^{\text{se}})$ is similar to (A5) with θ_p replaced by x_p. Conditions (B3) and (B4) imply that $\tilde{\Lambda}^{\text{se}}(\underline{\theta})=0$. Integrating the Gâteaux differential by parts yields

$$\partial L(w^1_{x_p},\underline{w}^0_{x_p};w^1,\underline{w}^0\mid\tilde{\Lambda}^{\text{se}})=\chi(\underline{\theta})w^1(\underline{\theta})+\int_{\underline{\theta}}^{\overline{\theta}}\chi(\theta)\mathrm{d}w^1(\theta),\tag{B5}$$

with

$$\chi(\theta)\equiv\eta_\alpha\int_\theta^{\overline{\theta}}\left[\tilde{\Lambda}^{\text{se}}(\tilde{\theta})-(1-G(\tilde{\theta}))\right]\mathrm{d}\tilde{\theta}+\eta_\alpha\int_{\max\{x_p,\theta\}}^{\overline{\theta}}(\tilde{\theta}-x_p)\mathrm{d}\tilde{\Lambda}^{\text{se}}(\tilde{\theta}).$$

By condition (B3), it follows that $\chi(\theta)\leq 0$ for all θ. Condition (B4) implies that $\chi(\theta)=0$ for $\theta\in[\underline{\theta},x_p]$. It follows that $\tilde{\Lambda}^{\text{se}}(\theta)=1-G(\theta)$ for all $\theta\in(\underline{\theta},x_p]$. This implies that $x_p\leq b$ otherwise the associated multiplier $\tilde{\Lambda}^{\text{se}}$ would be decreasing. Integrating by parts the second term of $\chi(\theta)$, I obtain

$$\chi(\theta)=\int_\theta^{\overline{\theta}}G(\tilde{\theta})\mathrm{d}\tilde{\theta}+(\theta-x_p)(1-\tilde{\Lambda}^{\text{se}}(\theta)),\ \forall\theta\geq x_p.$$

By definition of θ_p, there must exist a $\theta\in[x_p,\theta_p)$ such that the first term is strictly positive; since $\tilde{\Lambda}^{\text{se}}(\theta)\leq 1$, the second term is nonnegative. Hence $\chi(\theta)>0$, contradicting the necessary conditions. ■

B3. Lévy processes and Lévy bandits

Here, I extend the analysis to the more general Lévy bandits (Cohen and Solan (2013)). The risky task's payoff is driven by a Lévy process whose Lévy triplet depends on an unknown binary state. In what follows, I start with a reminder about Lévy processes and Lévy bandits. Then, I show that the optimality of the cutoff rule and its time consistency property generalize to Lévy bandits.

A Lévy process $L=(L(t))_{t\geq 0}$ is a continuous-time stochastic process that (i) starts at the origin: $L(0) = 0$; (ii) admits càdlàg modification;[27] (iii) has stationary independent increments. Examples of Lévy processes include a Brownian motion, a Poisson process, and a compound Poisson process.

Let (Ω, P) be the underlying probability space. For every Borel measurable set $A \in \mathcal{B}(\mathbf{R} \setminus \{0\})$, and every $t \geq 0$, let the Poisson random measure $N(t, A)$ be the number of jumps of L in the time interval $[0, t]$ with jump size in A: $N(t, A) = \#\{0 \leq s \leq t \mid \Delta L(s) \equiv L(s) - L(s-) \in A\}$. The measure ν defined by

$$\nu(A) \equiv \mathbf{E}[N(1, A)] = \int N(1, A)(\omega) \mathrm{d}P(\omega).$$

is called the Lévy measure of the process L.

I focus on Lévy processes that have finite expectation for each t. For a fixed Lévy process L, there exists a constant $\mu \in \mathbf{R}$, a Brownian motion $\sigma Z(t)$ with standard deviation $\sigma \geq 0$, and an independent Poisson random measure $N_\nu(t, dh)$ with the associated Lévy measure ν such that, for each $t \geq 0$, the Lévy-Itō decomposition of $L(t)$ is

$$L(t) = \mu t + \sigma Z(t) + \int_{\mathbf{R} \setminus \{0\}} h \tilde{N}_\nu(t, \mathrm{d}h),$$

where $\tilde{N}_\nu(t, A) \equiv N_\nu(t, A) - t\nu(A)$ is the compensated Poisson random measure.[28] Hence, a Lévy process L is characterized by a triplet $\langle \mu, \sigma, \nu \rangle$.

The agent operates a two-armed bandit in continuous time, with a safe task S that yields a known flow payoff s_i to player i, and a risky task R whose payoff, depending on an unknown state $x \in \{0, 1\}$, is given by the process L^x. For ease of exposition, I assume that both players derive the same payoff from R but different payoffs from S. For a fixed state x, L^x is a Lévy process characterized by the triplet $\langle \mu^x, \sigma^x, \nu^x \rangle$. For an arbitrary prior p that the state is 1, I denote by P_p the probability measure over the space of the realized paths.

I keep the same assumptions (A1–A6) on the Lévy processes L^x as in Cohen and Solan (2013) and modify A5 to ensure that both players prefer to use R in state 1 and S in state 0. That is, $\mu^1 > s_i > \mu^0$, for $i \in \{\alpha, \rho\}$.[29] Let $\eta_i = (\mu^1 - s_i)/(s_i - \mu^0)$ denote player i's net gain from the experimentation. I assume that the agent gains more from the experimentation, i.e., $\eta_\alpha > \eta_\rho$.[30]

A (pure) allocation policy is a nonanticipative stochastic process $\pi = \{\pi_t\}_{t \geq 0}$.

[27] It is continuous from the right and has limits from the left.

[28] Consider a set $A \in \mathcal{B}(\mathbf{R} \setminus \{0\})$ and a function $f : \mathbf{R} \to \mathbf{R}$. The integral with respect to a Poisson random measure $N(t, A)$ is defined as $\int_A f(h) N(t, \mathrm{d}h) = \sum_{s \leq t} f(\Delta L(s)) \mathbb{1}_A(\Delta(L(s)))$.

[29] The assumptions are (A1) $\mathbf{E}[(L^x)^2(1)] = (\mu^x)^2 + (\sigma^x)^2 + \int h^2 \nu^x(\mathrm{d}h) < \infty$; (A2) $\sigma^1 = \sigma^0$; (A3) $|\nu^1(\mathbf{R} \setminus \{0\}) - \nu^0(\mathbf{R} \setminus \{0\})| < \infty$; (A4) $|\int h(\nu^1(\mathrm{d}h) - \nu^0(\mathrm{d}h))| < \infty$; (A5) $\mu^0 < s_\alpha < s_\rho < \mu^1$; (A6) For every $A \in \mathcal{B}(\mathbf{R} \setminus \{0\})$, $\nu^0(A) < \nu^1(A)$. Assumption (A1) states that both L^1 and L^0 have finite quadratic variation. It follows that both have finite expectation. Assumptions (A2) to (A4) ensure that players cannot distinguish between the two states in any infinitesimal time. Assumption (A5) states that the expected payoff rate of R is higher than that of S in state 1 and lower in state 0. The last assumption (A6) requires that jumps of any size h, both positive or negative, occur more often in state 1 than in state 0. Consequently, jumps always provide good news, and increase the posterior belief of state 1.

[30] The results generalize to the case in which, for a fixed state x, the drift term of the Lévy process L^x differs for the principal and the agent, as long as the relation $\eta_\alpha > \eta_\rho$ holds.

Here, $\pi_t \in [0,1]$ (resp. $1-\pi_t$) may be interpreted as the fraction of the resource in the interval $[t, t+dt)$ that is devoted to R (resp. S), which may depend only on the history of events up to t.[31] The space of all policies, including randomized ones, is denoted Π. (See footnote 10.)

Player i's payoff given a policy $\pi \in \Pi$ and a prior belief $p \in [0,1]$ that the state is 1 is

$$U_i(\pi, p) \equiv \mathbf{E}\left[\int_0^\infty re^{-rt}\left[\mathrm{d}L^x\left(\int_0^t \pi_s ds\right) + (1-\pi_t)\, s_i \mathrm{d}t\right] \,\Big|\, \pi, p\right].$$

Over an interval $[t, t+dt)$, if the fraction π_t of the resource is allocated to R, the expected payoff increment to player i conditional on x is $[(1-\pi_t)s_i + \pi_t \mu^x]dt$. By the Law of Iterated Expectations, player i's payoff can be written as the discounted sum of the expected payoff increments

$$U_i(\pi, p) = \mathbf{E}\left[\int_0^\infty re^{-rt}\left[\pi_t \mu^x + (1-\pi_t)s_i\right]\mathrm{d}t \,\Big|\, \pi, p\right].$$

For a fixed policy π, I define $W^1(\pi)$ and $W^0(\pi)$ as follows:

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

Then, player i's payoff can be written as

$$U_i(\pi, p) = p\left(\mu^1 - s_i\right)W^1(\pi) + (1-p)\left(\mu^0 - s_i\right)W^0(\pi) + s_i.$$

Let Γ denote the image of the mapping $(W^1, W^0) : \Pi \to [0,1]^2$, referred to as the feasible set. The following lemma characterizes the southeast boundary of Γ.

LEMMA 8: *There exists $a^* > 0$ such that the southeast boundary of Γ is given by*

$$\{(w^1, w^0) \mid w^0 = 1 - (1-w^1)^{a^*/(1+a^*)}, w^1 \in [0,1]\}$$

PROOF:

I want to show that the maximum in (5) is achieved by a lower-cutoff policy when $p_1 \geq 0, p_2 \leq 0$. If $p_1 \geq 0, p_2 \geq 0$ ($p_1 \leq 0, p_2 \leq 0$), the maximum is achieved by the policy which directs all resources to R (S). If $p_1 > 0, p_2 < 0$, according to Cohen and Solan (2013), the maximum is achieved by a lower-cutoff Markov policy which directs all resource to R if the posterior belief is above the cutoff and to S if below. The cutoff belief, denoted p^*, satisfies the equation $p^*/(1-p^*) = a^*/(1+a^*)$, where a^* is the positive root of Equation 6.1 in Cohen and Solan (2013). Let $K(p) \equiv \max_{w \in \Gamma}(p_1, p_2) \cdot w$. If $|p_1|/(|p_1|+|p_2|) \leq p^*$, $K(p)$ equals zero. If $|p_1|/(|p_1|+|p_2|) > p^*$, I obtain $K(p) = -p_2\left(\frac{-p_2 a^*}{p_1(1+a^*)}\right)^{a^*}/(a^*+1) + p_1 + p_2$. It is easy to verify that the

[31] Suppose the process L is a Lévy process L^1 with probability $p \in (0,1)$ and L^0 with probability $1-p$. Let $\mathcal{F}_s^L$ be the sigma-algebra generated by the process $(L(t))_{t \leq s}$. Then it is required that the process π satisfies that $\{\int_0^t \pi_s \mathrm{d}s \leq t'\} \in \mathcal{F}_{t'}^L$, for any $t, t' \in [0, \infty)$.

functional form of the southeast boundary is

$$\beta^{\mathrm{se}}(w^1) = 1 - (1 - w^1)^{\frac{a^*}{a^*+1}}, w^1 \in [0, 1].$$

$\blacksquare$

Given lemma 8, the proof of proposition 2, which only relies on the properties of the southeast boundary of the feasible set, applies directly to the current setting. Therefore, the cutoff rule is optimal.

PROPOSITION 10 (Lévy bandits—sufficiency):
The cutoff rule is optimal if assumption 1 holds.

For every prior $p \in [0, 1]$ that the state is 1, the probability measure P_p satisfies $P_p = pP_1 + (1-p)P_0$. An important auxiliary process is the Radon-Nikodym density, given by

$$\psi_t \equiv \frac{d(P_0 \mid \mathcal{F}_{K(t)})}{d(P_1 \mid \mathcal{F}_{K(t)})}, \quad \text{where } K(t) = \int_0^t \pi_s \mathrm{d}s \text{ and } t \in [0, \infty).$$

According to lemma 1 in Cohen and Solan (2013), if the prior belief is p, the posterior belief at time t is given by

$$p_t = \frac{p}{p + (1-p)\psi_t}.$$

The agent of type θ updates his belief about the state. He assigns odds ratio θ/ψ_t to the state being 1, referred to as his type at time t. Let $\underline{\theta}_t = \max\{\underline{\theta}, \theta^*_\alpha \psi_t\}$. Recall that θ^*_α denotes the odds ratio at which the agent is indifferent between continuing and stopping. At time t, only those types above $\underline{\theta}_t$ remain. The principal's updated belief about the agent's type distribution, in terms of his type at time 0, is given by the density function

$$f_t^0(\theta) = \begin{cases} \frac{[p(\theta)+(1-p(\theta))\psi_t]f(\theta)}{\int_{\underline{\theta}_t}^{\overline{\theta}}[p(\theta)+(1-p(\theta))\psi_t]f(\theta)\mathrm{d}\theta}, & \text{if } \theta \in [\underline{\theta}_t, \overline{\theta}], \\ 0, & \text{otherwise.} \end{cases}$$

The principal's belief about the agent's type distribution, in terms of his type at time t, is given by the density function

$$f_t(\theta) = \begin{cases} f_t^0(\theta\psi_t)\psi_t, & \text{if } \theta \in [\underline{\theta}_t/\psi_t, \overline{\theta}/\psi_t], \\ 0, & \text{otherwise.} \end{cases}$$

I prove that continuing the cutoff rule is optimal by showing two things. First, given the distribution f_t at time t, the threshold of the top pooling segment is θ_p/ψ_t. Second, if assumption 1 holds for $\theta \leq \theta_p$ under distribution f, then it holds for $\theta \leq \theta_p/\psi_t$ under f_t. The detailed proof is similar to that of proposition 5 and hence omitted.

PROPOSITION 11 (Lévy bandits—time consistency):
If assumption 1 holds, the cutoff rule is time-consistent.

B4. Optimal contract with transfers

The set-up

In this subsection, 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 nonnegative 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 nonnegative, nondecreasing 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.[32] 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:

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

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 9 (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

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

[32] Formally, the number of successes achieved up to, and including, time t defines the point process $\{N_t\}_{t\geq 0}$. Let $\mathcal{F} \equiv \{\mathcal{F}_t\}_{t\geq 0}$ 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

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

The principal's payoff can be rewritten similarly. ■

Lemma 9 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 $(W^1, W^0, \mathbf{t}^1, \mathbf{t}^0) : \Pi^* \to [0,1]^2 \times [0,\infty) \times [0,\infty)$ can be interpreted as the new contract space when transfers are allowed. The following lemma characterizes this contract space.

LEMMA 10: *The image of the mapping* $(W^1, W^0, \mathbf{t}^1, \mathbf{t}^0) : \Pi^* \to [0,1]^2 \times [0,\infty) \times [0,\infty)$*, 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 $\|(w^1, w^0, t^1, t^0) - (W^1, W^0, \mathbf{t}^1, \mathbf{t}^0)(\pi, c)\|$ 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 $t^1/(\Delta re^{-r\tilde{t}})$ 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

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

Given that $1-\tilde{\epsilon} < \Pr(p_{\tilde{t}} > 1-\tilde{\epsilon} \mid \omega=1) \le 1$ and $0 \le \Pr(p_{\tilde{t}} < \tilde{\epsilon} \mid \omega=1) < \tilde{\epsilon}$, 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 10 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 $(W^1, W^0, \mathbf{t}^1, \mathbf{t}^0)(\pi, c)$ 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 9, 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 9 and 10, 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

$$\max_{w^1, w^0, t^1, t^0} \int_\Theta \left(\frac{\theta}{1+\theta}\left(\eta_\rho w^1(\theta) - t^1(\theta)\right) - \frac{1}{1+\theta}\left(w^0(\theta) + t^0(\theta)\right)\right) \mathrm{d}F(\theta),$$

subject to the IC constraints.

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

$$\text{(B6)} \qquad \left(t^0\right)'(\theta) = -\theta\left(\eta_\alpha\left(w^1\right)'(\theta) + \left(t^1\right)'(\theta)\right) + \left(w^0\right)'(\theta),$$

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, (B6) 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) (B6) 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 (hereafter, LL constraint) $t^1(\theta), t^0(\theta) \geq 0, \forall\theta \in \Theta$. 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 $w^0(\theta) \geq \beta^{\mathrm{se}}(w^1(\theta)), \forall\theta \in \Theta$, where $\beta^{\mathrm{se}}(\cdot)$ 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 $\boldsymbol{y} = (y^1, y^0, y_t^1)$. For problem $\mathcal{P}'$, I define a Lagrangian

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

where $\boldsymbol{\gamma} = (\gamma^1, \gamma^0, \gamma_t^1, \gamma_t^0)$ are the associated costate variables and $\boldsymbol{\mu} = (\mu_t^1, \mu_t^0, \mu)$ are multipliers associated with the LL and feasibility constraints. The law of motion is

$$\text{(B7)} \qquad \left(w^1\right)' = y^1, \left(w^0\right)' = y^0, \left(t^1\right)' = y_t^1, \left(t^0\right)' = y^0 - \theta\left(\eta_\alpha y^1 + y_t^1\right).$$

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,

$$\text{(B8)} \qquad \frac{\partial \mathcal{L}}{\partial y^1} = \gamma^1 - \eta_\alpha \theta \gamma_t^0 = 0, \quad \frac{\partial \mathcal{L}}{\partial y^0} = \gamma^0 + \gamma_t^0 = 0, \quad \frac{\partial \mathcal{L}}{\partial y_t^1} = \gamma_t^1 - \theta \gamma_t^0 = 0.$$

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

$$\text{(B9)} \qquad \begin{aligned} \dot{\gamma}^1 &= -\frac{\partial \mathcal{L}}{\partial w^1} = -\left[\frac{\theta}{1+\theta}(\eta_\alpha + \eta_\rho) + \frac{\eta_\alpha H}{f}\right] f + \mu \left(\beta^{\text{se}}\right)' (w^1), \\ \dot{\gamma}^0 &= -\frac{\partial \mathcal{L}}{\partial w^0} = \frac{2}{1+\theta} f - \mu, \quad \dot{\gamma}_t^1 = -\frac{\partial \mathcal{L}}{\partial t^1} = -H - \mu_t^1, \quad \dot{\gamma}_t^0 = -\frac{\partial \mathcal{L}}{\partial t^0} = -\mu_t^0. \end{aligned}$$

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

$$\text{(B10)} \qquad \begin{aligned} (\gamma^1(\underline{\theta}), \gamma^0(\underline{\theta}), \gamma_t^1(\underline{\theta}), \gamma_t^0(\underline{\theta})) &= (0,0,0,0), \\ (\gamma^1(\overline{\theta}), \gamma^0(\overline{\theta}), \gamma_t^1(\overline{\theta}), \gamma_t^0(\overline{\theta})) &= \left.\left(\frac{\partial \Phi}{\partial w^1}, \frac{\partial \Phi}{\partial w^0}, \frac{\partial \Phi}{\partial t^1}, \frac{\partial \Phi}{\partial t^0}\right)\right|_{\theta = \overline{\theta}} \\ &= \left(-\eta_\alpha \overline{\theta} H(\overline{\theta}), H(\overline{\theta}), -\overline{\theta} H(\overline{\theta}), -H(\overline{\theta})\right). \end{aligned}$$

Also, the following slackness conditions must be satisfied,

$$\text{(B11)} \qquad \begin{aligned} &\mu_t^1 \geq 0, t^1 \geq 0, \mu_t^1 t^1 = 0; \quad \mu_t^0 \geq 0, t^0 \geq 0, \mu_t^0 t^0 = 0; \\ &\mu \geq 0, w^0 - \beta^{\text{se}}(w^1) \geq 0, \mu \left(w^0 - \beta^{\text{se}}(w^1)\right) = 0. \end{aligned}$$

LEMMA 11 (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* (B7)–(B11) *are satisfied. Conversely,* (B7)–(B11) *are also sufficient since the Lagrangian is concave in* $(\boldsymbol{s}, \boldsymbol{y})$.

Based on the sufficiency part of lemma 11, I only need to construct $(\boldsymbol{s}, \boldsymbol{y}, \boldsymbol{\gamma}, \boldsymbol{\mu})$ such that the conditions (B7)–(B11) 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 B1 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

$$(\text{B12}) \qquad (\beta^{\text{se}})'(w^1(\theta)) = \frac{(\eta_\alpha + \eta_\rho)\theta}{2 - \frac{(1+\theta)H}{\theta f} + \frac{1+\theta}{\theta^2 f}\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s}.$$

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 (B13) and (B14) 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.

$$(\text{B13}) \qquad \theta^* \equiv \sup\left\{\hat{\theta} \in \Theta : \frac{(\eta_\alpha + \eta_\rho)\theta}{2 - \frac{(1+\theta)H}{\theta f} + \frac{1+\theta}{\theta^2 f}\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s} < \eta_\alpha\theta, \ \ \forall\theta \le \hat{\theta}\right\},$$

$$(\text{B14}) \qquad \theta_p \equiv \inf\left\{\hat{\theta} \in \Theta : \frac{\eta_\rho \int_{\theta}^{\overline{\theta}} sh(s)\mathrm{d}s}{H(\overline{\theta}) - H(\theta)} \le \eta_\alpha\theta, \ \ \forall\theta \ge \hat{\theta}\right\}.$$

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 B1.) When the agent's type is below θ^*, the equilibrium allocation is given by (B12), 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 B2.) 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.

One immediate observation is that the most pessimistic type's policy is socially optimal. The contract for type $\underline{\theta}$ is chosen such that $(\beta^{\text{se}})'(w^1(\underline{\theta})) = (\eta_\alpha + \eta_\rho)\underline{\theta}/2$. 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

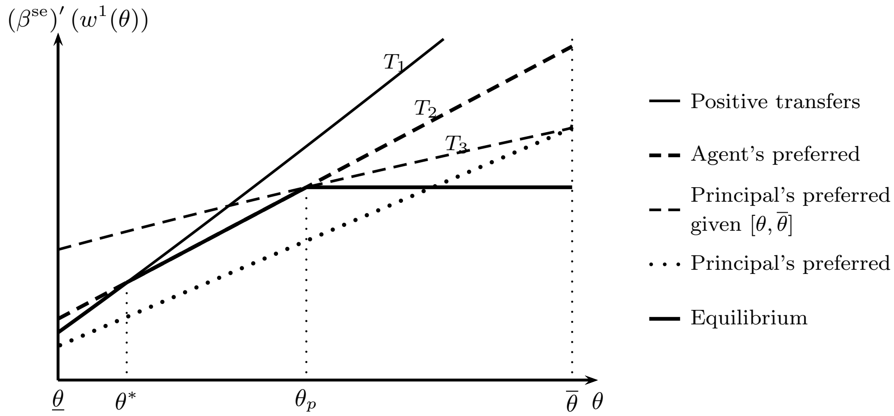

Figure B1. Equilibrium allocation: three segments

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

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.

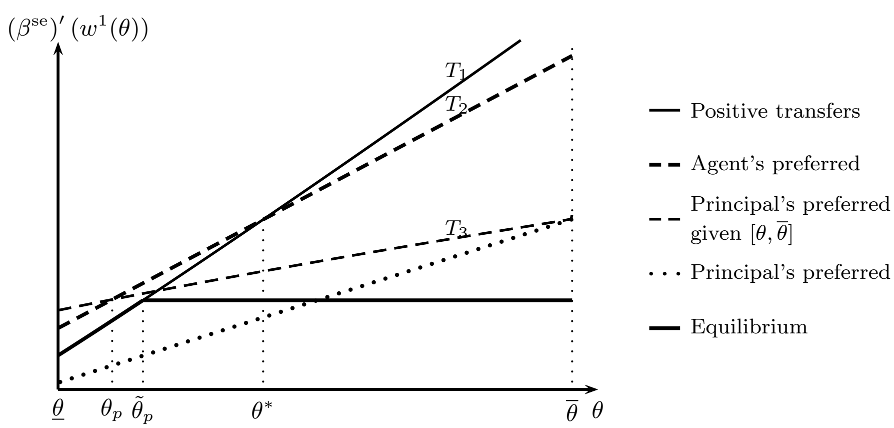

Figure B2. Equilibrium allocation: two segments

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

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 $\gamma^1, \gamma_t^1, \gamma_t^0$ according to (B8). Then I determine the trajectories of $\mu, \mu_t^1, \mu_t^0, w^1$ based on (B9). 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 $\mu_t^1 = 0$ and $\dot{\gamma}_t^1 = -H$. Combined with the boundary condition, this implies that $\gamma_t^1 = -\int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s$. From (B8), we know that

$$\gamma^0 = \int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s/\theta \text{ and } \gamma^1 = -\eta_\alpha \int_{\underline{\theta}}^{\theta} H(s)\mathrm{d}s.$$

Substituting $\dot{\gamma}^1$ and $\dot{\gamma}^0$ into (B9), I have

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

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 $(\beta^{\mathrm{se}})'(w^1(\theta)) = \eta_\alpha\theta$. From (B8), we know that $\gamma^0 = -\dot{\gamma}^1/\eta_\alpha - \theta\dot{\gamma}^0$. Substituting $\dot{\gamma}^1$, $\dot{\gamma}^0$ and $(\beta^{\mathrm{se}})'(w^1(\theta)) = \eta_\alpha\theta$, I have

$$\gamma^0 = H - \frac{\theta}{1+\theta}\frac{\eta_\alpha - \eta_\rho}{\eta_\alpha} f. \tag{B15}$$

Combining (B9) and (B15), 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

$$\frac{\eta_\alpha}{\eta_\alpha-\eta_\rho} \geq -\frac{1}{1+\theta}-\frac{f'}{f}\theta, \quad \forall \theta \in [\theta^*, \theta_p] \tag{B16}$$

$$\frac{\eta_\alpha}{\eta_\alpha-\eta_\rho} \geq \frac{2+\theta}{1+\theta}+\frac{f'}{f}\theta, \quad \forall \theta \in [\theta^*, \theta_p]. \tag{B17}$$

Note that assumption (B16) 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,

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

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

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

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

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

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

$$\frac{\eta_\alpha}{\eta_\alpha-\eta_\rho} \geq \frac{\frac{\theta_p}{1+\theta_p}f(\theta_p)}{H(\overline{\theta})-H(\theta_p)}. \tag{B18}$$

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 (B16), (B17) and (B18) hold, the constructed trajectory solves $\mathcal{P}'$. If the trajectory w^1 defined by (B12) is weakly increasing over $[\underline{\theta}, \theta^*]$, the monotonicity of $\eta_\alpha w^1 + t^1$ is satisfied.[33] 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 $(\beta^{\mathrm{se}})'(w^1(\tilde{\theta}_p))$ must satisfy the following condition

$$(\beta^{\mathrm{se}})'(w^1(\tilde{\theta}_p)) = \frac{\eta_\rho \int_{\tilde{\theta}_p}^{\overline{\theta}} \theta h(\theta)\mathrm{d}\theta - \eta_\alpha \tilde{\theta}_p H(\tilde{\theta}_p) + \eta_\alpha \int_{\underline{\theta}}^{\tilde{\theta}_p} H(s)\mathrm{d}s}{H(\overline{\theta}) - 2H(\tilde{\theta}_p) + \frac{\int_{\underline{\theta}}^{\tilde{\theta}_p} H(s)\mathrm{d}s}{\tilde{\theta}_p}} \tag{B19}$$

At the same time, $\tilde{\theta}_p$ must also satisfy (B12). Equation (B12) and (B19) determines the threshold $\tilde{\theta}_p$. Since $(\beta^{\mathrm{se}})'(w^1(\tilde{\theta}_p)) < \eta_\alpha \tilde{\theta}_p$, (B19) implies that

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

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 $\theta \in [\tilde{\theta}_p, \overline{\theta}]$.

[33] 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.