Benefiting from Bias

Ball, Ian; Gao, Xin

DOI:
10.1016/j.jet.2024.105816

J. Econ. Theory 217 (2024) 105816

Contents lists available at ScienceDirect

Journal of Economic Theory

journal homepage: www.elsevier.com/locate/jet

Benefiting from bias: Delegating to encourage information acquisition ☆

Ian Ball [a], Xin Gao [b,*]

[a] Department of Economics, MIT, United States of America
[b] Department of Economics, University of Birmingham, UK

ARTICLE INFO

JEL classification:
D82
D86

Keywords:
Delegation
Information acquisition
Benefiting from bias

ABSTRACT

A principal delegates decisions to a biased agent. Payoffs depend on a state that the principal cannot observe. Initially, the agent does not observe the state, but he can acquire information about it at a cost. We characterize the principal's optimal delegation set. This set features a cap on high decisions and a gap around the agent's ex ante favorite decision. The set may even induce ex-post Pareto-dominated decisions. Under certain conditions on the cost of information acquisition, we show that the principal prefers delegating to an agent with a small bias than to an unbiased agent.

1. Introduction

Delegation is ubiquitous. The leader of an organization cannot make every decision, so she must delegate: CEOs delegate to middle managers and politicians delegate to policy advisors. In the classical model of delegation (Holmström, 1977), a principal delegates decisions to an agent with better information but biased preferences. The principal faces a tradeoff. To utilize the agent's private information, the principal must give the agent discretion, but discretion allows the agent to bias the decision in his own favor.

In many applications, the agent does not initially have better information, but he can acquire information at a cost. For example, managers gather information about employee performance, and advisors conduct policy research. In these settings, offering the agent discretion motivates him to acquire information. Conversely, if the agent has little control over the decision, he has little to gain from additional information.

In our model, an uninformed principal (she) delegates a single-dimensional decision to an agent (he) who is biased toward higher decisions. The agent is *initially* uninformed of the state. After the agent observes the delegation set, he chooses a level of costly private effort on experimentation. This effort determines the probability with which the experiment succeeds. If the experiment succeeds, the agent learns the state. Otherwise, the experiment fails and the agent learns nothing. After the experiment's outcome is realized, the agent chooses a decision from the delegation set.

We first characterize the form of optimal delegation. Theorem 1 says that any optimal delegation set must take one of three forms: *hollow*, *interval*, and *high-point*. A hollow set has a gap around the agent's ex ante favorite decision and a cap that prohibits

☆ This paper was previously circulated under the titles "Biased and Uninformed: Delegating to Encourage Information Acquisition" and "Benefiting from Bias." The paper was completed while Ian Ball was visiting Northwestern University. For insightful comments and discussions, we are grateful to Kyle Bagwell, Dirk Bergemann, Tilman Börgers, Hans Peter Grüner, Marina Halac, Johannes Hörner, Navin Kartik, Eric Maskin, Volker Nocke, Martin Peitz, Larry Samuelson, Nicolas Schutz, and Thomas Tröger, as well as audiences at Cambridge, Mannheim, and Yale.

* Corresponding author.
E-mail addresses: ianball@mit.edu (I. Ball), x.gao.1@bham.ac.uk (X. Gao).

https://doi.org/10.1016/j.jet.2024.105816
Received 26 July 2021; final version received 20 November 2023; Accepted 25 February 2024

Available online 12 March 2024

high decisions. The gap lowers the agent's payoff from remaining uninformed, thus motivating the agent to acquire information. The cap restricts the agent's bias if he learns that the state is high. If the optimal cap is low enough to exclude the agent's ex ante favorite decision, then there is no gap, and the optimal delegation set is an interval. Finally, a high-point set consists of an interval together with a singleton decision (the high point) that the principal and agent agree is too high in every state. The agent chooses the high point only if he learns that the state is very high. Including the high point induces the agent to acquire more information, and this benefit outweighs the direct loss from the agent's higher decisions.

In Theorem 2, we provide conditions under which the delegation set takes the different forms. Hollow delegation is optimal if the bias is sufficiently small. Interval delegation is optimal if the bias is sufficiently large and the agent's optimal effort choice is sufficiently concave as a function of the return from learning the state. We show by example that high-point delegation can be optimal if information acquisition is very costly.

Finally, we analyze the principal's preferences over the agent's level of bias. In the classical delegation problem with an informed agent, the principal wants the agent's bias to be as small as possible. In our setting, the agent's bias can benefit the principal by creating a wedge between the ex ante favorite decisions of the principal and the agent. With this wedge, the principal can select a delegation set that punishes the agent if his experiment fails, at a lower cost to the principal. Theorem 3 shows that for a range of cost functions, the principal prefers delegating to an agent with a small bias than to an unbiased agent.

The rest of the paper is organized as follows. Section 1.1 discusses related literature. Section 2 presents the model. Section 3 analyzes the players' optimization problems. In Section 4, we characterize the three forms that optimal delegation can take. We then give conditions under which each form is optimal. Section 5 studies the principal's preferences over the agent's level of bias. Section 6 is the conclusion. Proofs are in Appendix A.

1.1. Related literature

Our paper connects the classical delegation literature with more recent work on information acquisition. Holmström (1977) introduces the delegation problem with an informed agent. He proves that interval delegation is optimal in the UQC setting—*uniform* state distribution, *quadratic* losses, and *constant* (state-independent) bias. Subsequent work shows that interval delegation remains optimal under successive relaxations of the UQC assumptions. In a uniform–quadratic setting with affine bias, Melumad and Shibano (1991) characterize the incentive-compatible decision rules. They show that interval delegation is optimal as long as the bias is not very sensitive to the state. If the bias is sufficiently sensitive to the state, then a two-point delegation set is optimal. In a quadratic setting with constant bias, Martimort and Semenov (2006) give a sufficient condition on the state distribution for interval delegation to be optimal. In a quadratic setting with arbitrary bias, Alonso and Matouschek (2008) characterize whether delegation is valuable. They provide a general characterization of the optimal delegation set, and they give a condition that is necessary and sufficient for interval delegation to be optimal. In a more general setting that allows for non-quadratic preferences, Amador and Bagwell (2013) give separate necessary and sufficient conditions for interval delegation to be optimal.

We work in the tractable setting of Krähmer and Kováč (2016); we discuss the contribution of their paper below. This setting allows for a non-uniform state, non-quadratic losses, and non-constant bias.[1] In this setting, if the agent were informed, then the optimal delegation set would be an interval.

In our model, the agent is biased and information acquisition is costly. Both features are necessary to identify the principal's benefit from bias. Earlier work studies information acquisition and bias separately. Szalay (2005) studies delegation to an unbiased agent who can pay a cost to privately learn about the state.[2] Under certain conditions, the optimal delegation set has a gap around the players' ex ante favorite decision.[3] This gap encourages the agent to acquire information. Semenov (2018) studies delegation to a biased agent who, prior to contracting, privately learns the state with some exogenous probability. The agent cannot choose to acquire additional information.[4] The optimal delegation set has a gap around the agent's favorite uninformed decision. If the agent has not learned the state, then this gap induces him to select a less biased decision.

Krähmer and Kováč (2016) study a delegation problem in which the agent initially observes a private, binary signal of the state. Then, after contracting, the agent learns the state perfectly. Their paper focuses on whether sequential screening—by offering a menu of delegation sets—outperforms the optimal static delegation set. The optimal static delegation set is an interval, but if sequential screening is strictly optimal, the optimal menu contains non-interval delegation sets. In their model, static interval delegation is optimal if the agent's bias is sufficiently small. In our model, hollow delegation is optimal if the agent's bias is sufficiently small. Finally, Ivanov (2010) compares delegation with cheap-talk in a setting in which the principal can design the agent's private information.[5]

Beyond the setting of delegation, misaligned preferences can bring benefits for different reasons.[6] In communication games with endogenous information acquisition, the sender may acquire more information if he is more biased (Argenziano et al., 2016) or

[1] Kleiner (2022) uses a similar decision setting to study multi-dimensional delegation.

[2] Similarly, in Demski and Sappington (1987), the agent can acquire information at a cost. In their model, the principal cannot restrict the delegation set, and transfers are allowed.

[3] Aghion and Tirole (1997) observe that giving the agent discretion can encourage information acquisition, but they do not study a constrained delegation problem.

[4] As an extension, Semenov (2018) allows the agent to acquire information at a cost. The paper does not analyze the design of the delegation set in that setting. The paper suggests that the characterization of optimal delegation continues to hold, but this overlooks the effect of the delegation set on the agent's choice of effort.

[5] In a fixed UQC setting, Goltsman et al. (2009) compare delegation with two other protocols: mediation and negotiation (multi-round cheap talk).

[6] Li (2001) and Gerardi and Yariv (2008) show that an undesirable default option can encourage information acquisition. In a model of delegating regulations, Bubb and Warren (2014) study a different form of preference bias that directly increases the agent's marginal benefit from learning the state. In a macroeconomic model, Rogoff (1985) shows that appointing a central banker who is biased towards inflation stabilization can mitigate the problem of time-inconsistency.

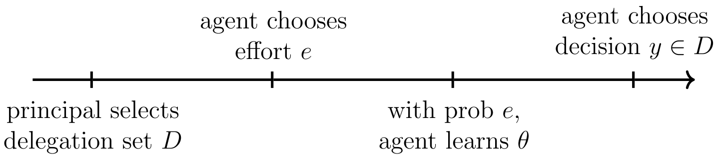

Fig. 1. Timing.

has more divergent beliefs (Che and Kartik, 2009). Our analysis is quite different because the principal (receiver) has commitment power.

2. Model

2.1. Delegation setting

There are two players: a principal (she) and an agent (he). The principal controls a decision $y \in \mathbf{R}$. Payoffs from the decision depend on a state $\theta \in \mathbf{R}$, drawn from a commonly known distribution function F with continuously differentiable, strictly positive density f on its support $[\underline{\theta}, \bar{\theta}]$. The principal and agent have state-dependent utilities $u_P(y, \theta)$ and $u_A(y, \theta)$, which will be specified in Section 2.2.

The principal does not observe the state. Initially, the agent does not observe the state either, but he can privately experiment at a cost. As in Szalay (2005), the agent chooses experimentation effort $e \in [0, 1)$. This effort level determines the probability with which the experiment succeeds. With probability e, independent of the state, the experiment succeeds and the agent learns the state. Otherwise, the experiment fails and the agent does not learn the state. The principal does not observe the agent's effort choice or the realization of the experiment. The rest of the setting is common knowledge.

The agent has a thrice continuously differentiable effort cost function

$$c : [0, 1) \to \mathbf{R},$$

satisfying $c(0) = c'(0) = 0$; $c''(e) > 0$ for all e in $(0, 1)$; and $\lim_{e \uparrow 1} c'(e) = \infty$.

Fig. 1 shows the timing. First, the principal selects a compact delegation set D.[7] The agent observes the set D and chooses effort e. Next, the outcome of the experiment is realized. The agent observes whether the experiment succeeds. He observes the state if and only if the experiment succeeds. Then he selects a decision y from D, and payoffs are realized.

2.2. Decision preferences

Following Krähmer and Kováč (2016),[8] utilities for the principal and agent are given by

$$u_P(y, \theta) = \theta y + a(y),$$

$$u_A(y, \theta) = (\theta + b(\theta))y + a(y),$$

where $a : \mathbf{R} \to \mathbf{R}$ is strictly concave and twice continuously differentiable, and $b : [\underline{\theta}, \bar{\theta}] \to \mathbf{R}$ is twice continuously differentiable and satisfies $b'(\theta) > -1$ for all θ. Thus, $\theta + b(\theta)$ is strictly increasing in θ. This setting nests quadratic-loss utilities.[9]

To ensure that the principal and agent have utility-maximizing decisions in each state, we assume

$$\lim_{y \to -\infty} -a'(y) < \underline{\theta} \wedge (\underline{\theta} + b(\underline{\theta})) \quad \text{and} \quad \lim_{y \to \infty} -a'(y) > \bar{\theta} \vee (\bar{\theta} + b(\bar{\theta})),$$

where $\wedge$ and $\vee$ denote the minimum and maximum operators, respectively.

In each state θ, the principal's and the agent's favorite decisions, denoted $y_P(\theta)$ and $y_A(\theta)$, are defined by the first-order conditions

$$-a'(y_P(\theta)) = \theta, \qquad -a'(y_A(\theta)) = \theta + b(\theta).$$

Their ex ante favorite decisions, denoted $y_{P,0}$ and $y_{A,0}$, are given by

$$-a'(y_{P,0}) = \mathbf{E}[\theta], \qquad -a'(y_{A,0}) = \mathbf{E}[\theta + b(\theta)].$$

---

[7] Offering a delegation set is equivalent to committing to a deterministic decision rule on an abstract message space. Kováč and Mylovanov (2009) give conditions under which stochastic mechanisms cannot improve upon deterministic mechanisms; see Footnote 11. For an analysis of delegation with money-burning, see Amador and Bagwell (2013) and Ambrus and Egorov (2017).

[8] Krähmer and Kováč (2016) put bias $-b(\theta)$ in the principal's utility; we put bias $b(\theta)$ in the agent's utility. That is, we normalize the state relative to the *principal's* utility. Our normalization makes some formulas more complicated, but it allows us to vary the agent's bias, without changing the state distribution.

[9] Take $a(y) = -(1/2)y^2$. Add the decision-irrelevant terms $-(1/2)\theta^2$ to u_P and $-(1/2)(\theta + b(\theta))^2$ to u_A. Then scale these utilities by 2.

The functions y_A and y_P are strictly increasing. Since $y_A(\theta) = y_P(\theta + b(\theta))$, we interpret $b(\theta)$ as the agent's bias in state θ. In particular, we have $y_A(\theta) \geq y_P(\theta)$ if and only if $b(\theta) \geq 0$.

We maintain from Krähmer and Kováč (2016) the following joint assumptions on the bias function and state distribution.[10] Let $B(\theta) = b(\theta)/(1 + b'(\theta))$. Recall that $1 + b'(\theta) > 0$, so $B(\theta)$ is well-defined and has the same sign as $b(\theta)$.

A1. $f(\theta) + B'(\theta)f(\theta) + B(\theta)f'(\theta) > 0$ for all θ.
A2. $b(\underline{\theta}) \geq 0$; $b(\bar{\theta}) > 0$; and $\mathbf{E}[b(\theta)] > 0$.
A3. $\underline{\theta} + b(\underline{\theta}) < \mathbf{E}[\theta]$.

Assumption A1 ensures that in the informed-agent delegation problem, interval delegation is optimal. For each state θ, the inequality guarantees that if the induced decision rule has a jump at state θ, then the principal can strictly increase her payoff by adding the decision $y_A(\theta)$ to the delegation set. In the special case of constant bias $b(\theta) = \beta > 0$, Assumption A1 reduces to the inequality $(\log f(\theta))' = f'(\theta)/f(\theta) > -1/\beta$. That is, the density f does not decay (multiplicatively) too quickly. The smaller the agent's bias, the more permissive is this constraint.

Assumptions A2 and A3 are less substantive. Under Assumption A1, the optimal informed-agent delegation set is an interval. Assumption A2 ensures that this interval excludes high decisions, but not low decisions. Assumption A3 ensures that this interval is not a singleton. With constant bias $b(\theta) = \beta$, Assumptions A2 and A3 together are equivalent to the inequality $0 < \beta < \mathbf{E}[\theta] - \underline{\theta}$.

We distinguish the *decision setting*, parameterized by (F, a, b), from the *information technology*, parameterized by c. Much of the delegation literature focuses on the uniform–quadratic–constant (UQC) decision setting, where the state is uniformly distributed on the unit interval $[0, 1]$, and utilities are quadratic with constant bias β. In this case, Assumptions A1–A3 hold if and only if $0 < \beta < 1/2$. The UQC decision setting will serve as a running example.

3. The players' optimization problems

Once the principal selects a delegation set, the agent faces a dynamic decision problem. We analyze the agent's problem and then the principal's.

3.1. Agent's problem

Suppose that the principal has selected a delegation set D. We analyze the agent's problem backwards, starting with his choice after each outcome of the experiment.

After a failure If the experiment fails, the agent does not learn the state, so he solves

$$\underset{y \in D}{\text{maximize}} \quad \mathbf{E}[u_A(y, \theta)].$$

The maximum value of this problem, denoted $u_{A,0}(D)$, is called the agent's *uninformed payoff* from D. Assume that ties are broken in the principal's favor, and denote the maximizer by $y_{A,0}(D)$, called the agent's *uninformed decision* from D.

After a success If the experiment succeeds, the agent observes the realized state θ, and he solves

$$\underset{y \in D}{\text{maximize}} \quad u_A(y, \theta).$$

Denote the maximum value of this problem by $u_A(D, \theta)$. Assume that ties are broken in the principal's favor, and denote the maximizer by $y_A(D, \theta)$, called the agent's *informed decision* from D in state θ. The agent's *informed payoff* from D is given by

$$u_{A,1}(D) = \mathbf{E}[u_A(D, \theta)].$$

Effort choice The agent's expected utility gain from observing the state is

$$\Delta_A(D) = u_{A,1}(D) - u_{A,0}(D).$$

Learning the state can only help the agent, so $\Delta_A(D) \geq 0$ for every delegation set D. With this notation, the agent's effort choice problem is

$$\underset{e \in [0,1)}{\text{maximize}} \quad u_{A,0}(D) + e \cdot \Delta_A(D) - c(e).$$

The assumptions on the cost function c ensure that there is a unique maximizer, denoted $\hat{e}(D)$, which is given by the first-order condition

[10] Our expressions look different because of our different normalization; see Footnote 8.

$$c'(\hat{e}(D)) = \Delta_A(D).$$

3.2. Principal's problem

For each delegation set D, the principal's expected utility, conditional on each outcome of the experiment, is determined by the agent's subsequent decisions. The principal's uninformed-agent utility and informed-agent utility are defined as

$$u_{P,0}(D) = \mathbf{E}[u_P(y_{A,0}(D), \theta)], \qquad u_{P,1}(D) = \mathbf{E}[u_P(y_A(D,\theta), \theta)].$$

Putting all this together, the principal's delegation problem is

$$\underset{D}{\text{maximize}} \quad U_P(D) = (1 - \hat{e}(D))u_{P,0}(D) + \hat{e}(D)u_{P,1}(D), \tag{1}$$

where the maximization is over all compact subsets D of $\mathbf{R}$. The principal's payoff from a delegation set D depends on three things: her uninformed-agent payoff $u_{P,0}(D)$, her informed-agent payoff $u_{P,1}(D)$, and the agent's expected return $\Delta_A(D)$ from learning the state, which pins down the agent's effort choice $\hat{e}(D)$.

The principal's payoff does not change if she adds to the delegation set superfluous decisions that the agent will never choose. We focus on delegation sets that are *minimal* in the sense that each decision in the delegation set is the agent's uniquely optimal choice in some state.

Lemma 1 *(Existence). The principal's delegation problem* (1) *has a solution. Moreover, there exists an optimal delegation set that is minimal.*

The principal maximizes over all compact subsets of $\mathbf{R}$. There is no loss in restricting to compact subsets of a fixed, sufficiently large interval. With the Hausdorff metric, this restricted domain is compact. By our tie-breaking assumption, the principal's objective is upper semicontinuous. Thus, a solution exists.

Hereafter, we restrict attention to delegation sets that are minimal, without explicit reference. In particular, we call a minimal delegation set *the* optimal delegation set if every other *minimal* delegation set is strictly worse.

4. Optimal delegation

4.1. Informed-agent benchmark

As a benchmark, consider the classical delegation problem in which the agent knows the state. In our notation, this problem is

$$\underset{D}{\text{maximize}} \quad u_{P,1}(D), \tag{2}$$

where the maximization is over all compact subsets D of $\mathbf{R}$. In this problem, the principal faces a tradeoff between utilizing the agent's private information and restricting the agent's expression of bias.

For any delegation set D, the principal can guarantee the uninformed-agent payoff $u_{P,0}(D)$ in the informed-agent problem by offering the singleton delegation set $\{y_{A,0}(D)\}$. Therefore, the principal's value from delegating to an informed agent (2) is an upper bound on the value from delegating to an initially uninformed agent (1).

Following Kováč and Mylovanov (2009) and Krähmer and Kováč (2016), we use the envelope theorem to express the principal's objective in (2) as a weighted average of the *agent's* payoff in each state.

Lemma 2 *(Utility representation). For any delegation set D, we have*

$$u_{P,1}(D) = u_A(D,\underline{\theta})B(\underline{\theta})f(\underline{\theta}) - u_A(D,\bar{\theta})B(\bar{\theta})f(\bar{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} u_A(D,\theta)[f(\theta) + B'(\theta)f(\theta) + B(\theta)f'(\theta)]\,\mathrm{d}\theta. \tag{3}$$

By Assumption A1, the coefficient on $u_A(D,\theta)$ in the integral is strictly positive for each state θ. By Assumption A2, the coefficient on $u_A(D,\underline{\theta})$ is nonnegative and the coefficient on $u_A(D,\bar{\theta})$ is strictly negative. Call $u_A(D,\theta)$ the payoff of type θ. As the principal considers enlarging the delegation set, she must balance the gain from increasing the payoff of each type θ in $[\underline{\theta},\bar{\theta})$ against the loss from increasing the payoff of the highest type $\bar{\theta}$. Therefore, it is optimal to give the agent discretion to choose any decision up to a cap.

The next lemma is equivalent to Krähmer and Kováč (2016, Lemma 1, p. 856).[11]

[11] In a quadratic setting with an informed agent, Kováč and Mylovanov (2009) show that, under their Assumption 1, interval delegation is optimal in the larger class of stochastic mechanisms. Their Assumption 1 is slightly weaker than our assumption A1, which is from Krähmer and Kováč (2016). Their Assumption 1 requires the inequality in A1 to hold only in states θ for which $y_A(\theta)$ lies between the endpoints of the optimal (informed-agent) delegation set. In our main model with information acquisition, we need the stronger assumption A1 because optimal delegation may allow more extreme decisions in order to encourage information acquisition.

Lemma 3 *(Informed-agent delegation). With an informed agent, the optimal delegation set is the interval $[y_A(\underline{\theta}), y_A(\hat{\theta})]$, where $\hat{\theta}$ is the unique solution of*

$$\hat{\theta} + b(\hat{\theta}) = \mathbf{E}[\theta | \theta \geq \hat{\theta}]. \tag{4}$$

Consider the informed agent's choice from the interval $[y_A(\underline{\theta}), y_A(\hat{\theta})]$. If $\theta \leq \hat{\theta}$, then the agent chooses his favorite decision $y_A(\theta)$. If $\theta > \hat{\theta}$, then the agent chooses the endpoint $y_A(\hat{\theta})$. Condition (4) says that the decision $y_A(\hat{\theta})$ is the principal's favorite, conditional on the state θ being at least $\hat{\theta}$. Perturbing the cap away from $y_A(\hat{\theta})$ perturbs the agent's decision in each state θ with $\theta \geq \hat{\theta}$. Thus, (4) is the first-order condition for the optimal cap. Assumptions A1–A3 ensure that (4) has a unique solution. It can be shown that if the bias function b strictly increases pointwise, then the cap $y_A(\hat{\theta})$ strictly decreases. Therefore, it is optimal to give an informed agent less discretion if he is more biased.

4.2. Structure of optimal delegation

We return to the main delegation problem (1) in which the agent is initially uninformed. By Lemma 3, the interval delegation set $[y_A(\underline{\theta}), y_A(\hat{\theta})]$ maximizes the principal's informed-agent payoff $u_{P,1}$. As the principal modifies this delegation set, she must trade off the loss from reducing her informed-agent payoff against the potential benefits: increasing her uninformed-agent payoff and inducing greater information acquisition (by decreasing the agent's uninformed payoff or increasing the agent's informed payoff).

We first characterize the possible "gaps" in an optimal delegation set. Formally, for $d_1 < d_2$, a delegation set D *has a gap* (d_1, d_2) if $[d_1, d_2] \cap D = \{d_1, d_2\}$.[12]

Lemma 4 *(Gaps). If an optimal delegation set has a gap (d_1, d_2), then $d_1 < y_{A,0} < d_2$. Moreover, either*

(i) $\mathbf{E}[u_A(d_1, \theta)] = \mathbf{E}[u_A(d_2, \theta)]$*; or*
(ii) $\mathbf{E}[u_A(d_1, \theta)] > \mathbf{E}[u_A(d_2, \theta)]$ *and* $d_2 > y_A(\bar{\theta})$.

In words, if an optimal delegation set has a gap, then the gap must contain the agent's ex ante favorite decision $y_{A,0}$. Moreover, the agent must be ex ante indifferent between the endpoints of the gap, unless the right endpoint d_2 is strictly above $y_A(\bar{\theta})$, in which case the agent may strictly prefer the left endpoint.

To prove Lemma 4, we use the utility representation in Lemma 2 to show that the principal strictly prefers to fill (at least partially) any gap that violates (i) and (ii). The intuition for Lemma 4 is as follows. Consider a delegation set D satisfying $D \subset [y_A(\underline{\theta}), y_A(\bar{\theta})]$. Gaps in D clearly reduce the agent's informed payoff. By Assumption A1, gaps in D also reduce the principal's informed-agent payoff. Therefore, any benefit from a gap must come from its effect on the agent's uninformed decision. If a delegation set D is optimal, then it must have the minimal gap necessary to induce the uninformed decision $y_{A,0}(D)$. Such a gap contains exactly those decisions that the agent strictly prefers to $y_{A,0}(D)$. Hence, (i) follows. This argument does not apply to gaps (d_1, d_2) with $d_2 > y_A(\bar{\theta})$. Filling in decisions near d_2 can reduce the principal's informed-agent payoff. In some states, the agent switches his decision from d_1 to a higher, newly available decision. The agent's decision decreases only in those extremely high states in which d_2 was chosen. We will give an example of an optimal delegation set with a gap that satisfies (ii).

Using Lemma 4, we now characterize the three forms that optimal delegation can take. Recall the definition of the cutoff type $\hat{\theta}$ from the optimal informed-agent delegation set (Lemma 3).

Theorem 1 *(Optimal delegation—three forms). If a delegation set D^* is optimal, then one of the following holds.*

I. *Hollow: $D^* = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup [y_1, y_2]$ for some y_0, y_1, and y_2 satisfying $y_0 < y_{A,0} < y_1 \leq y_2$, where $\mathbf{E}[u_A(y_0, \theta)] = \mathbf{E}[u_A(y_1, \theta)]$ and $y_2 > y_A(\hat{\theta})$.*
II. *Interval: $D^* = [y_A(\underline{\theta}), y_0]$ for some y_0 satisfying $y_{P,0} < y_0 < y_{A,0}$.*
III. *High-point: $D^* = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{\bar{y}\}$ for some y_0 and $\bar{y}$ satisfying $y_0 < y_{A,0} < y_A(\bar{\theta}) < \bar{y}$ and $\mathbf{E}[u_A(y_0, \theta)] > \mathbf{E}[u_A(\bar{y}, \theta)]$.*

Fig. 2 shows an example of each of the three delegation forms. In every case, the agent's ex ante favorite decision $y_{A,0}$ is excluded. This is because the only first-order effect of hollowing out a small gap around $y_{A,0}$ (satisfying (i) from Lemma 4) is to increase the principal's uninformed-agent payoff. Starting from the optimal informed-agent delegation set $[y_A(\underline{\theta}), y_A(\hat{\theta})]$,[13] a *hollow* delegation set is formed as follows. First, the cap is increased from $y_A(\hat{\theta})$ to y_2. Raising the cap now has the additional benefit of increasing information acquisition. Second, the principal hollows out a gap (y_0, y_1) around the agent's ex ante favorite decision. If the experiment fails, the agent is indifferent between y_0 and y_1, but he chooses the lower decision y_0 by the tie-breaking assumption. The gap reduces

[12] According to this definition, a proper subset of a gap is not a gap.

[13] The optimal informed-agent delegation set can be quite irregular if the state distribution violates our assumption A1; see Alonso and Matouschek (2008, Proposition 2, p. 273). If the density decreases too quickly, then the optimal delegation set may have gaps in order to deter deviations by lower types. In our model, there is a positive probability that the agent's posterior mean equals the prior mean. Therefore, the density effectively decreases quickly to the right of the prior mean. The new feature in our setting is the endogeneity of the agent's type distribution.

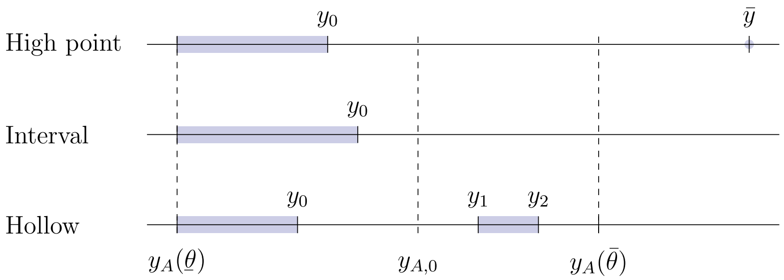

Fig. 2. Three delegation forms.

both the agent's informed and uninformed payoffs, but it reduces the agent's uninformed payoff by more. Therefore, enlarging the gap increases information acquisition, but it also decreases the principal's informed-agent payoff.[14]

If the agent's bias is large enough, then allowing decisions above the agent's ex ante favorite decision $y_{A,0}$ entails a direct loss that outweighs the information acquisition benefit. In this case, an *interval* delegation set is optimal. Each interval delegation set includes the principal's ex ante favorite decision but excludes the agent's ex ante favorite decision.

A *high-point* delegation set consists of an interval together with an isolated point $\bar{y}$ that is strictly higher than even the agent's favorite decision in the highest state. If the experiment fails, the agent strictly prefers y_0 to the high point $\bar{y}$. This strict preference distinguishes high-point delegation from the special case of hollow delegation with $y_1 = y_2$. The gap $(y_0, \bar{y})$ is consistent with Lemma 4 because $\bar{y} > y_A(\bar{\theta})$. The optimal value of $\bar{y}$ balances two competing forces. Increasing the high point discourages information acquisition by decreasing the agent's informed payoff, while leaving the agent's uninformed payoff unchanged. On the other hand, perturbing the high point below its optimal value reduces the principal's informed-agent payoff by inducing the agent to change his decision from y_0 to the new, lower high point in certain states.[15]

We claim that high-point delegation is never optimal in the related models of Szalay (2005) and Semenov (2018). Consider an arbitrary high-point delegation set D. If the agent is unbiased, as in Szalay (2005),[16] then the principal strictly prefers to perturb the high point downward, thus increasing the common informed-agent payoff $u_{P,1} = u_{A,1}$ and inducing greater information acquisition. If the agent is biased but exogenously informed of the state with a fixed probability, as in Semenov (2018), then the principal's payoff is a fixed convex combination of her uninformed-agent and informed-agent payoffs. Perturbing the high point affects only the principal's informed-agent payoff $u_{P,1}$. As a function of the high point, this payoff $u_{P,1}$ is strictly quasiconvex (see Lemma 5) and hence cannot have an interior maximizer.

4.3. Optimal delegation form

The characterization in Theorem 1 reduces the class of candidate optimal delegation sets to a small parametric family. In any example, we can numerically optimize over the parameters to find an optimal delegation set. We now give general conditions under which the different delegation forms are optimal.

Recall that the agent's optimal effort level $\hat{e}(D)$ depends only on the agent's gain $\Delta_A(D)$ from learning the state. Let $\hat{e}(x)$ denote the agent's optimal effort level if $\Delta_A(D) = x$. Formally, $\hat{e}(x)$ is defined by $c'(\hat{e}(x)) = x$ for any $x > 0$. Below, the context should make clear whether the argument of $\hat{e}$ is a delegation set or a number. For the next result, say that a delegation set D *lies below* y if $\max D \le y$.

Theorem 2 *(Optimal delegation form). Fix a decision setting* (F, a, b).

1. *If* $\mathbf{E}[\theta + b(\theta)] \le \hat{\theta} + b(\hat{\theta})$, *then for every cost function, every optimal delegation set is hollow.*
2. *If* $\mathbf{E}[\theta + b(\theta)] > \hat{\theta} + b(\hat{\theta})$, *then there exists a positive threshold* $K = K(F, a, b)$ *such that if* $-\hat{e}''(x)/\hat{e}'(x) \ge K$ *for all positive* x, *then every optimal delegation set is an interval.*
3. *If* $u_{P,1}(\{y_{P,0}\}) \ge u_{P,1}([y_A(\underline{\theta}), y_{A,0}])$, *then for every cost function, every optimal delegation set lying below* $y_A(\bar{\theta})$ *is an interval.*

The inequality $\mathbf{E}[\theta + b(\theta)] \le \hat{\theta} + b(\hat{\theta})$ is a low-bias condition.[17] It holds if and only if the agent's ex ante favorite decision $y_{A,0}$ lies inside the optimal informed-agent delegation set $[y_A(\underline{\theta}), y_A(\hat{\theta})]$ from Lemma 3. In this case, given any interval (or high-point) delegation set, the principal strictly prefers to add some decision above $y_{A,0}$ that does not change the agent's uninformed decision. This modification increases the principal's informed-agent payoff and encourages information acquisition. Thus, hollow delegation is

[14] Enlarging the gap also affects the principal's uninformed-agent payoff. The sign of this effect depends on the relative values of y_0 and $y_{P,0}$.

[15] The principal's informed-agent payoff is a strictly quasiconvex function of the high-point (Lemma 5). At the optimal high-point, this function must be locally strictly increasing; otherwise the principal would profit from reducing the high point.

[16] Szalay (2005) restricts the principal to delegation sets within $[y_A(\underline{\theta}), y_A(\bar{\theta})]$, so high-point delegation is not feasible. Our argument above shows that even if this restriction is dropped, high-point delegation cannot be optimal.

[17] Formally, if this inequality holds for some bias function b, then it can be shown to hold for any pointwise smaller bias function.

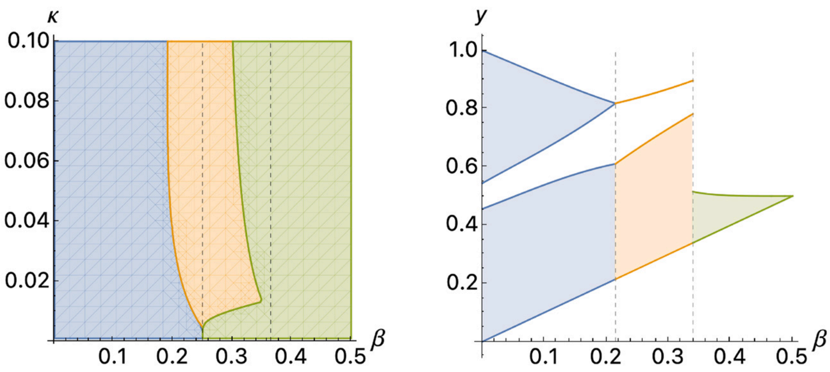

Fig. 3. Optimal delegation forms. (For interpretation of the colors in the figure, the reader is referred to the web version of this article.)

optimal (part 1). If instead $y_{A,0}$ lies outside $[y_A(\underline{\theta}), y_A(\hat{\theta})]$, then adding such decisions above $y_{A,0}$ necessarily reduces the principal's informed-agent payoff. If $\hat{e}$ is sufficiently concave, then this loss cannot be outweighed by the information acquisition benefit. Hence, interval delegation is optimal (part 2). Finally, for very high bias, hollow delegation sets lying below $y_A(\bar{\theta})$ cannot be optimal, no matter the cost function (part 3). The principal prefers the singleton delegation set $\{y_{P,0}\}$ to any hollow delegation set lying below $y_A(\bar{\theta})$.

We illustrate Theorem 2 in the UQC setting with induced effort function $\hat{e}(x) = 1 - e^{-x/\kappa}$, where $\kappa > 0$. This effort function is convenient because the Arrow–Pratt coefficient $-\hat{e}''(x)/\hat{e}'(x)$ equals $1/\kappa$ for all x. The function $\hat{e}$ is induced by the cost function

$$c(e) = \kappa[(1-e)\log(1-e) + e], \tag{5}$$

which is the leading example in Szalay (2005). This UQC–exponential setting is parameterized by the bias parameter β and the cost parameter κ.

The left panel of Fig. 3 indicates the optimal form of delegation as a function of the parameter vector (β, κ). For this parameter range, numerical optimization suggests that high-point delegation is not optimal. Moving from left to right, the three regions (blue, orange, green) respectively indicate where the following delegation forms are optimal: hollow with non-degenerate upper interval $[y_1, y_2]$, hollow with singleton upper interval, and interval. In particular, the interval delegation region is not convex. For some values of β, hollow delegation is optimal only for intermediate values of κ. The upper interval of the hollow delegation set encourages information acquisition. For low κ, information acquisition is high without the upper interval. For high κ, information acquisition is low even with the upper interval.

The conditions in Theorem 2 have simple expressions in terms of (β, κ), which we illustrate in the left panel of Fig. 3. By part 1, hollow delegation is optimal if $\beta \le 1/4$, i.e., left of the first dotted line. By part 2, for $\beta > 1/4$, interval delegation is optimal if $\kappa \le 1/K(\beta)$, where K is a positive function. In the UQC setting, part 3 implies that interval or high-point delegation is optimal if $\beta \ge (\sqrt{33} + 3)/24 \approx 0.36$, i.e., right of the second dotted line.

The right panel of Fig. 3 plots the optimal delegation set, as a function of the bias β, for fixed $\kappa = 0.02$. These parameters correspond to the horizontal line $\kappa = 0.02$ in the left panel. As the bias increases from $\beta = 0$, the top interval of the optimal delegation set shrinks until it becomes a point. Thereafter, the upper interval remains degenerate until the optimal delegation form switches from hollow to interval. At this transition, information acquisition jumps down, and the principal's uninformed-agent payoff jumps up.

It is difficult to give general sufficient conditions for high-point delegation to be optimal. Under a certain condition on the decision setting, we construct a cost function for which high-point delegation is optimal; see Appendix A.9. In the UQC setting, our condition is satisfied if $\beta > 0.29$. In our construction, the effort function $\hat{e}$ is a differentiable, strictly increasing approximation of a step function: $\hat{e}(x) \approx \varepsilon[x \ge x_0]$, where x_0 is slightly larger than $\Delta_A([y_A(\underline{\theta}), y_{P,0}])$, and ε is a small positive parameter. For any delegation set D, the agent's experiment fails with high probability. Therefore, a necessary condition for a delegation set to be optimal is that it induces an uninformed decision sufficiently near $y_{P,0}$. Among delegation sets satisfying this necessary condition, intervals do not induce sufficient information acquisition, and hollow sets result in high decisions that bring down the principal's informed-agent payoff. The principal strictly prefers a high-point delegation set. The high point encourages information acquisition (by pushing the agent's return from learning the state above x_0), while preserving the principal's informed-agent payoff (by ensuring that decisions above $y_{A,0}$ are selected with low probability).

5. Principal's preferences over the agent

Suppose that the principal initially selects an agent from a pool of candidates with different information acquisition technologies and different biases. Whom should the principal select? Of course, the principal's first choice would be an unbiased agent who can perfectly learn the state for free, but such an agent may not be available. It is intuitive that the principal would prefer cheaper information acquisition and lower bias. Cheaper information is indeed better for the principal, but we show that reducing the agent's bias can hurt the principal.

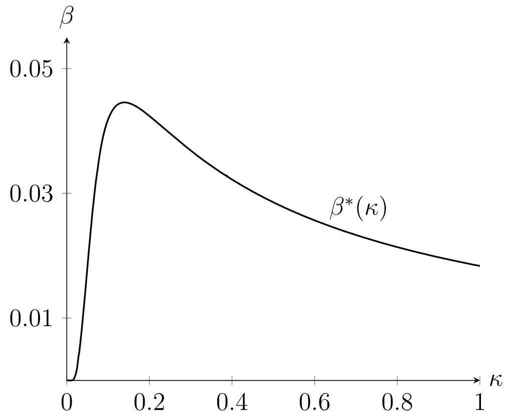

Fig. 4. Principal's favorite bias.

Consider an optimal delegation set D^* for a fixed cost function c. If the marginal cost of information acquisition, c', decreases pointwise, then the agent's induced effort choice function $\hat{e}$ increases pointwise. In particular, $\hat{e}(D^*)$ increases, so the principal's payoff from D^* increases.[18] Hence, the principal's value from delegation must increase.

We turn to the comparative statics in the agent's bias. To build intuition, consider the UQC setting. Fix an arbitrary delegation set D. As the agent's bias β decreases, the principal's informed-agent payoff $u_{P,1}(D)$ and uninformed-agent payoff $u_{P,0}(D)$ both increase. The effect on the agent's return $\Delta_A(D)$ from learning the state depends on the form of D. If D is an interval delegation set, then $\Delta_A(D)$ also increases, and therefore the principal strictly benefits. For hollow delegation sets, however, the agent's return from learning the state can decrease. The resulting loss from reduced information acquisition can outweigh the benefit of less biased decisions after each outcome of the experiment.

Theorem 3 *(Benefiting from bias). Assume that the players have quadratic-loss utilities and the state distribution is symmetric about its mean. Let e_0 denote the level of effort chosen by an unbiased agent faced with an unrestricted delegation set. If*

$$\frac{c'(e_0)}{c''(e_0)} > 1 - e_0, \tag{6}$$

then there exists $\bar{\beta} > 0$ such that for all β in $(0, \bar{\beta})$, the principal strictly prefers an agent with constant bias β to an unbiased agent.

In the application of a politician delegating to a policy advisor who conducts costly research, Theorem 3 indicates that the politician may be better off selecting an advisor whose political preferences differ from her own. The intuition for Theorem 3 is as follows. If the agent is unbiased, then the principal and agent agree on the ex ante optimal decision. Removing decisions near this ex ante optimum encourages information acquisition by punishing the agent if his experiment fails. But this punishment for the agent also punishes the principal. If the agent is biased, then the principal and agent prefer different decisions ex ante. The principal can therefore select a delegation set that excludes the agent's ex ante favorite decision but includes her own. This intuition relies on the property of the information acquisition technology that, with positive probability, the agent learns little about the state. On the other hand, we believe that the potential benefit from bias is robust to relaxing the assumption that the agent perfectly learns the state with positive probability.

By Theorem 2 (part 1), optimal delegation is hollow if the agent's bias is sufficiently small (and nonzero). In the limit as the bias tends to zero, the hollow delegation gap may or may not vanish, depending on the cost of information acquisition. By Szalay (2005, Proposition 3, p. 1181), if condition (6) holds, then with an unbiased agent, any optimal delegation set D^* must have a gap around the ex ante favorite decision $y_{P,0} = y_{A,0}$. Hence, the agent's uninformed decision is strictly below $y_{P,0}$. We show that for any sufficiently small bias β, the principal's payoff from offering the shifted delegation set $\beta + D^*$ to an agent with bias β is strictly higher than the principal's payoff from offering the delegation set D^* to an unbiased agent. The first-order benefit from inducing a less biased uninformed decision outweighs the second-order loss from inducing more biased decisions when the agent learns the state. As the bias β varies, shifting the delegation set in this way keeps constant the agent's return from learning the state, and hence the induced effort level.

To illustrate Theorem 3, we return to the UQC setting with the cost function from (5). We have $c'(e_0)/c''(e_0) = -(1-e_0)\log(1-e_0)$. The effort level e_0 depends on κ. It can be checked that (6) holds if and only if $\kappa \leq 1/12$. Fig. 4 plots the principal's favorite level of agent bias, $\beta^*(\kappa)$, as a function of the cost parameter κ. The principal's favorite bias achieves a maximum value of approximately 0.045. With an informed agent, it is optimal for the principal to delegate if and only if $\beta < 0.5$. So over the range of biases for which delegation is worthwhile, the bottom 9% may be more aligned than is optimal. The principal's gain from an optimally biased agent

[18] By Assumption A3, a singleton delegation set is suboptimal, so we must have $\Delta_A(D^*) > 0$ and $u_{P,1}(D^*) > u_{P,0}(D^*)$.

relative to an unbiased agent peaks between $\kappa = 0.07$ and $\kappa = 0.08$. With $\kappa = 0.08$, the principal's gain from facing an agent with the optimal bias $\beta^*(0.08) \approx 0.036$ (rather than an unbiased agent) is equivalent to her loss, in the informed-agent problem, from facing an agent with bias $\beta \approx 0.027$ (rather than an unbiased agent).

6. Conclusion

In the classical delegation literature, the agent's bias is all that prevents the principal from achieving her first-best decision rule. We consider an additional friction—costly information acquisition for the agent. We characterize the principal's optimal delegation set, and we show that the principal may prefer delegating to an agent with a small bias than to an unbiased agent. Preference misalignment allows the principal to punish the agent at a lower cost to herself. These punishments occur on-path because the agent cannot learn the state with certainty.

We work with a flexible decision setting, but the information acquisition technology takes a simple parametric form. This structure allows us to characterize the optimal delegation set. Allowing flexible information acquisition presents a few challenges. In our model, the agent's optimal effort choice depends on a simple statistic of the delegation set, and the resulting distribution of the agent's posterior beliefs lies in a one-parameter family. With flexible information acquisition, the agent's optimal choice of information structure depends on the delegation set in a complicated way, and there are no exogenous restrictions on the resulting distribution of the agent's posterior beliefs. Incorporating flexible information acquisition into models of delegation is a promising direction for future work.

CRediT authorship contribution statement

Ian Ball: Writing – original draft, Writing – review & editing. **Xin Gao:** Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Appendix A. Proofs

In the proofs, we use the following terminology. Given a state θ and a decision y with $y \neq y_A(\theta)$, the *θ-conjugate* of y is the unique decision y' distinct from y that satisfies $u_A(y', \theta) = u_A(y, \theta)$. Decision $y_A(\theta)$ is defined to be the θ-conjugate of itself. Similarly, given a decision y with $y \neq y_{A,0}$, the *ex ante conjugate* of y is the unique decision y' distinct from y that satisfies $\mathbf{E}[u_A(y', \theta)] = \mathbf{E}[u_A(y, \theta)]$. Decision $y_{A,0}$ is defined to be the ex ante conjugate of itself.

A.1. Proof of Lemma 1

Existence Select a sufficiently large compact set $Y \subset \mathbf{R}$ so that restricting to compact subsets of Y does not change the supremum. Denote by $\mathcal{K}_Y$ the space of nonempty compact subsets of Y, endowed with the Hausdorff metric. This space $\mathcal{K}_Y$ is compact (Aliprantis and Border, 2006, Theorem 3.85). To prove existence, it suffices to check that the map

$$D \mapsto (1 - \hat{e}(D))u_{P,0}(D) + \hat{e}(D)u_{P,1}(D)$$

is upper semicontinuous. We prove that $u_{P,0}$ and $u_{P,1}$ are upper semicontinuous and $\hat{e}$ is continuous.

Define correspondences $Y_0^* : \mathcal{K}_Y \twoheadrightarrow Y$ and $Y_1^* : \mathcal{K}_Y \times [\underline{\theta}, \bar{\theta}] \twoheadrightarrow Y$ by

$$Y_0^*(D) = \operatorname*{argmax}_{y \in D} \mathbf{E}[u_A(y, \theta)], \qquad Y_1^*(D, \theta) = \operatorname*{argmax}_{y \in D} u_A(y, \theta).$$

Endow $\mathcal{K}_Y \times [\underline{\theta}, \bar{\theta}]$ with the product topology. By Berge's theorem (Aliprantis and Border, 2006, Theorem 17.31), these correspondences are upper hemicontinuous.[19] By our tie-breaking assumption, the utility functions $u_{P,0}$ and $u_{P,1}$ can be expressed as

$$u_{P,0}(D) = \max_{y \in Y_0^*(D)} \mathbf{E}[u_P(y, \theta)], \qquad u_{P,1}(D) = \mathbf{E}\left[\max_{y \in Y_1^*(D,\theta)} u_P(y, \theta)\right].$$

[19] The function $y \mapsto \mathbf{E}[u_A(y, \theta)]$ is continuous by dominated convergence. For Y_0^*, we apply Berge's theorem to the correspondence φ_0 from $\mathcal{K}_Y$ into Y defined by $\varphi_0(D) = D$. The identity function on $\mathcal{K}_Y$ is clearly continuous, so φ_0 is continuous by Theorem 17.15 in Aliprantis and Border (2006). For Y_1^*, we apply Berge's theorem to the correspondence φ_1 from $\mathcal{K}_Y \times [\underline{\theta}, \bar{\theta}]$ into Y defined by $\varphi_1(D, \theta) = D$. This correspondence is continuous because it is the composition of φ_0 with the projection map $(D, \theta) \mapsto D$.

By a variant of Berge's theorem (Aliprantis and Border, 2006, Lemma 17.30), it follows that $u_{P,0}$ and $u_{P,1}$ are upper semicontinuous.[20]

Now we check that $\hat{e}$ is continuous. Recall that $\Delta_A(D) = u_{A,1}(D) - u_{A,0}(D)$. It follows, as above, from Berge's theorem (Aliprantis and Border, 2006, Theorem 17.31) that Δ_A is a continuous function on $\mathcal{K}_Y$. The agent's effort-choice first-order condition gives $\hat{e}(D) = (c')^{-1}(\Delta_A(D))$. The inverse function $(c')^{-1}$ is well-defined and continuous because c' is strictly increasing, with $c'(0) = 0$ and $\lim_{e \uparrow 1} c'(e) = \infty$. Therefore, $\hat{e}$ is the composition of continuous functions, and hence is continuous.

Minimality Formally, a delegation set D is *minimal* if for each decision d in D there exists some state θ in $[\underline{\theta}, \bar{\theta}]$ such that $\{d\} = \operatorname{argmax}_{y \in D} u_A(y, \theta)$.

Now we turn to the proof. Given an arbitrary compact delegation set D, let

$$\underline{y} = \max \left[\operatorname*{argmax}_{y \in D} u_A(y, \underline{\theta}) \right], \qquad \bar{y} = \min \left[\operatorname*{argmax}_{y \in D} u_A(y, \bar{\theta}) \right].$$

Let $D' = [\underline{y}, \bar{y}] \cap D$. By construction, D' is minimal.[21] The agent makes the same choice from D' as from D, unless he learns that the state is in $\{\underline{\theta}, \bar{\theta}\}$, which occurs with probability 0. Therefore, $U_P(D') = U_P(D)$.

A.2. Proof of Lemma 2

Fix a delegation set D. By the envelope theorem (Milgrom and Segal, 2002), we have

$$u_A(D, \theta) - u_A(D, \underline{\theta}) = \int_{\underline{\theta}}^{\theta} (1 + b'(\xi)) y_A(D, \xi) \, \mathrm{d}\xi, \tag{7}$$

for each state θ. Since

$$u_P(y, \theta) - u_A(y, \theta) = -b(\theta) y = -(1 + b'(\theta)) B(\theta) y,$$

we have

$$u_{P,1}(D) - u_{A,1}(D) = -\int_{\underline{\theta}}^{\bar{\theta}} (1 + b'(\theta)) y_A(D, \theta) B(\theta) f(\theta) \, \mathrm{d}\theta.$$

Now integrate by parts, using $u_A(D, \theta)$ as an antiderivative, from (7), to get

$$u_{P,1}(D) - u_{A,1}(D) = -u_A(D, \theta) B(\theta) f(\theta) \Big|_{\underline{\theta}}^{\bar{\theta}} + \int_{\underline{\theta}}^{\bar{\theta}} u_A(D, \theta) [B'(\theta) f(\theta) + B(\theta) f'(\theta)] \, \mathrm{d}\theta.$$

Bringing $u_{A,1}(D)$ to the right side and expressing it as integral gives (3).

A.3. Proof of Lemma 3

Let D be a maximizer of $u_{P,1}$ that is minimal, which exists by the proof of Lemma 1 (Appendix A.1). By A3, we know $y_{P,0} > y_A(\underline{\theta})$, so $\max D > y_A(\underline{\theta})$. We claim that $\max D \le y_A(\bar{\theta})$. Otherwise, letting $\bar{d}$ be the $\bar{\theta}$-conjugate of $\max D$,[22] the principal would strictly prefer $[y_A(\underline{\theta}) \wedge \bar{d}, \bar{d}]$ by the utility representation (Lemma 2).

We have shown that $y_A(\underline{\theta}) < \max D \le y_A(\bar{\theta})$, so $\max D = y_A(\hat{\theta})$ for some $\hat{\theta}$ in $(\underline{\theta}, \bar{\theta}]$. Hence, $D = [y_A(\underline{\theta}), y_A(\hat{\theta})]$, for otherwise the principal would strictly prefer $[y_A(\underline{\theta}), y_A(\hat{\theta})]$ by the utility representation (Lemma 2). It remains to maximize over the value of $\hat{\theta}$. We have

$$u_{P,1}([y_A(\underline{\theta}), y_A(\hat{\theta})]) = \int_{\underline{\theta}}^{\hat{\theta}} u_P(y_A(\theta), \theta) f(\theta) \, \mathrm{d}\theta + \int_{\hat{\theta}}^{\bar{\theta}} u_P(y_A(\hat{\theta}), \theta) f(\theta) \, \mathrm{d}\theta.$$

Differentiate with respect to $\hat{\theta}$, noting that $a'(y_A(\hat{\theta})) = -\hat{\theta} - b(\hat{\theta})$, to get

[20] For $u_{P,0}$, conclude by dominated convergence that the function $y \mapsto \mathbf{E}[u_P(y, \theta)]$ is continuous, hence upper semicontinuous. For $u_{P,1}$, Berge's theorem implies that the integrand is upper semicontinuous in (D, θ), so $u_{P,1}$ is upper semicontinuous by dominated convergence.

[21] Fix d in D'. We claim that $\{d\} = \operatorname{argmax}_{y \in D} u_A(y, \theta)$ for some θ. If d is in $[y_A(\underline{\theta}), y_A(\bar{\theta})]$, take $\theta = y_A^{-1}(d)$. If $d < y_A(\underline{\theta})$, then $d = \underline{y}$, so take $\theta = \underline{\theta}$. If $d > y_A(\bar{\theta})$, then $d = \bar{y}$, so take $\theta = \bar{\theta}$.

[22] See the beginning of Appendix A for the definition.

$$y_A'(\hat{\theta}) \int_{\hat{\theta}}^{\bar{\theta}} [\theta - \hat{\theta} - b(\hat{\theta})] f(\theta)\, \mathrm{d}\theta. \tag{8}$$

Since y_A' is strictly positive, setting (8) to zero gives condition (4). We analyze the sign of the integral in (8), as a function of $\hat{\theta}$. It is positive at $\hat{\theta} = \underline{\theta}$ since $\mathbf{E}[\theta] - \underline{\theta} - b(\underline{\theta}) > 0$ (by A3), and it is negative near $\bar{\theta}$ since $-b(\bar{\theta}) < 0$ (by A2). To prove that this integral in (8) is strictly single-crossing from above, it suffices to show that it is strictly quasiconcave. Differentiating gives

$$b(\hat{\theta}) f(\hat{\theta}) - (1 + b'(\hat{\theta}))(1 - F(\hat{\theta})) = -(1 + b'(\hat{\theta}))(1 - F(\hat{\theta}) - B(\hat{\theta}) f(\hat{\theta})).$$

This expression is strictly single-crossing from above since $1 + b'(\hat{\theta}) > 0$ for all $\hat{\theta}$ and $1 - F(\hat{\theta}) - B(\hat{\theta}) f(\hat{\theta})$ is strictly decreasing in $\hat{\theta}$ (by A1).

A.4. Proof of Lemma 4

Let D^* be an optimal delegation set that is minimal. We must have $\Delta_A(D^*) > 0$ and $u_{P,1}(D^*) > u_{P,0}(D^*)$; otherwise, since $y_{P,0} > y_A(\underline{\theta})$ by A3, the principal would strictly prefer the delegation set $[y_A(\underline{\theta}), y_{P,0}]$ by the utility representation (Lemma 2).

Suppose, for a contradiction, that D^* has a gap (d_1, d_2) that violates (i) and (ii). There are four cases. In each case, we construct a delegation set D' that the principal strictly prefers.

1. $y_{A,0} \notin (d_1, d_2)$ and $d_2 \le y_A(\bar{\theta})$. Take $D' = D^* \cup (d_1, d_2)$.
2. $y_{A,0} \notin (d_1, d_2)$ and $d_2 > y_A(\bar{\theta})$. Since D^* is minimal, $\{d_2\} = \operatorname{argmax}_{y \in D^*} u_A(y, \bar{\theta})$. In particular, $u_A(d_2, \bar{\theta}) > u_A(d_1, \bar{\theta})$. Let d_1' be the $\bar{\theta}$-conjugate of d_2. Take $D' = D^* \cup (d_1, d_1']$.
3. $y_{A,0} \in (d_1, d_2)$ and $\mathbf{E}[u_A(d_1, \theta)] < \mathbf{E}[u_A(d_2, \theta)]$. Let d_1' be the ex ante conjugate of d_2. Take $D' = D^* \cup (d_1, d_1']$.
4. $y_{A,0} \in (d_1, d_2)$ and $\mathbf{E}[u_A(d_1, \theta)] > \mathbf{E}[u_A(d_2, \theta)]$ and $d_2 \le y_A(\bar{\theta})$. Let d_2' be the ex ante conjugate of d_1. Take $D' = D^* \cup [d_2', d_2)$.

In every case, we have $u_{P,0}(D') \ge u_{P,0}(D^*)$; $\Delta_A(D') > \Delta_A(D^*)$; and $u_{P,1}(D') > u_{P,1}(D^*)$ by the utility representation (Lemma 2). Therefore, $U_P(D') > U_P(D^*)$.

A.5. Proof of Theorem 1

The proof uses the following lemma.

Lemma 5 *(Effects of delegation parameters). The following hold for any minimal delegation set within each regime from Theorem 1.*[23]

I. *Hollow: $u_{P,1}$ is strictly decreasing in the gap $y_1 - y_0$ and single-peaked about $y_2 = y_A(\hat{\theta})$; Δ_A is strictly increasing in the gap $y_1 - y_0$ and in the cap y_2.*
II. *Interval: $u_{P,1}$ is single-peaked about $y_0 = y_A(\hat{\theta})$; Δ_A is strictly increasing in y_0.*
III. *High-point: $u_{P,1}$ is strictly increasing in y_0 and strictly quasiconvex in $\bar{y}$; Δ_A is strictly convex in y_0 and strictly decreasing in $\bar{y}$.*

Let D^* be an optimal delegation set that is minimal. As shown in the proof of Lemma 4 (Appendix A.4), D^* is not a singleton, so $\max D^* > y_A(\underline{\theta})$. By the utility representation (Lemma 2), $\min D^* \le y_A(\underline{\theta})$. We separate into three cases according to the value of $\max D^*$.

I. $y_{A,0} < \max D^* \le y_A(\bar{\theta})$. We show that D^* is hollow. Let $y_2 = \max D^*$. By Lemma 4, either (a) $D^* = [y_A(\underline{\theta}), y_2]$ or (b) $D^* = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup [y_1, y_2]$ for some y_0 and y_1 satisfying $y_0 < y_{A,0} < y_1 \le y_2$ and $\mathbf{E}[u_A(y_0, \theta)] = \mathbf{E}[u_A(y_1, \theta)]$. We first rule out case (a). For positive r, let $y_0(r) = y_{A,0} - r$ and let $y_1(r)$ be the ex ante conjugate of $y_0(r)$. By the implicit function theorem, $y_1(r)$ is differentiable. For r in $(0, y_0^{-1}(y_A(\underline{\theta})) \wedge y_1^{-1}(y_2))$, let $D(r) = [y_A(\underline{\theta}), y_0(r)] \cup [y_1(r), y_2]$. It is straightforward to check that $\frac{\mathrm{d}}{\mathrm{d}r} U_P(D(r))$ is strictly positive in a neighborhood of $r = 0$ because as r tends to 0, the first-order effects of r on Δ_A and $u_{P,1}$ vanish, and the first-order effect of r on $u_{P,0}$ is positive (because $y_{A,0} > y_{P,0}$). Therefore, $[y_A(\underline{\theta}), y_2]$ cannot be optimal, and we must be in case (b). We claim that $y_2 > y_A(\hat{\theta})$. As a function of the cap y_2, we know from Lemma 5 that $u_{P,1}$ is single-peaked about $y_A(\hat{\theta})$; Δ_A is strictly increasing; and $u_{P,0}$ is constant.
II. $y_A(\underline{\theta}) < \max D^* \le y_{A,0}$. Let $y_0 = \max D^*$. We show that D^* is an interval delegation set. By Lemma 4, we must have $D^* = [y_A(\underline{\theta}), y_0]$. We claim that $y_0 > y_{P,0}$. As a function of the cap y_0, it follows from Lemma 5 that $u_{P,1}$ is single-peaked about $y_A(\hat{\theta})$; Δ_A is strictly increasing; and $u_{P,0}$ is single-peaked about $y_{P,0}$, which is strictly below $y_A(\hat{\theta})$. It remains to check that $y_0 \ne y_{A,0}$. As a function of the cap y_0, both Δ_A and $u_{P,1}$ are differentiable at $y_0 = y_{A,0}$, but $u_{P,0}$ has only one-sided derivatives. The jump

[23] These comparative statics still hold if some of the optimality conditions in Theorem 1 are dropped. In particular, for the hollow regime, we can drop the requirement that $y_2 > y_A(\hat{\theta})$. For the interval regime, we can relax the condition $y_0 > y_{P,0}$ to $y_0 > y_A(\underline{\theta})$. In some of the proofs below, we apply Lemma 5 over these extended domains.

from the left derivative to the right derivative is strictly positive since $y_{A,0} > y_{P,0}$.[24] Thus, U_P cannot achieve a maximum at $y_0 = y_{A,0}$.

III. $\max D^* > y_A(\bar{\theta})$. Let $\bar{d} = \max D^*$. We show that D^* is hollow or high-point. Since D^* is minimal, $\bar{d}$ must be an isolated point of D^*. Let $y_0 = \max(D^* \setminus \{\bar{d}\})$. By Lemma 4, we have $D^* = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{\bar{d}\}$ for some y_0 satisfying $y_0 < y_{A,0}$ and $\mathbf{E}[u_A(y_0,\theta)] \geq \mathbf{E}[u_A(\bar{d},\theta)]$. There are two cases. If $\mathbf{E}[u_A(y_0,\theta)] = \mathbf{E}[u_A(\bar{d},\theta)]$, then D^* is hollow, with $y_1 = y_2 = \bar{d}$. If $\mathbf{E}[u_A(y_0,\theta)] > \mathbf{E}[u_A(\bar{d},\theta)]$, then D^* is high-point, with $\bar{y} = \bar{d}$. In this case, we establish one more property of D^*. Since D^* is minimal, $u_A(\bar{y},\bar{\theta}) > u_A(y_0,\bar{\theta})$. Let $\tilde{y}$ be the $\bar{\theta}$-conjugate of $\bar{y}$. We have $y_0 < \tilde{y} < y_A(\bar{\theta})$. We now compare the delegation set $\tilde{D} = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{\tilde{y}\}$ with D^*. Clearly, $u_{A,1}(\tilde{D}) > u_{A,1}(D^*)$. By the utility representation (Lemma 2), $u_{P,1}(\tilde{D}) > u_{P,1}(D^*)$. Since D^* is optimal, we must have $y_{A,0}(\tilde{D}) \neq y_0$, hence $\mathbf{E}[u_A(\tilde{y},\theta)] > \mathbf{E}[u_A(y_0,\theta)]$.

A.6. Proof of Lemma 5

I. Hollow The comparative statics for $u_{P,1}$ follow from the utility representation (Lemma 2) and the proof of Lemma 3 (Appendix A.3). Clearly, Δ_A is strictly increasing in y_2. We show that Δ_A is strictly increasing in the gap $y_1 - y_0$.

It is convenient to parameterize the gap by the left-radius r. That is, let $y_0(r) = y_{A,0} - r$, and let $y_1(r)$ be the ex ante conjugate of $y_0(r)$. We view other quantities as functions of r as well, over the domain $(0, y_1^{-1}(y_2)]$. Define θ_A by $\theta_A + b(\theta_A) = \mathbf{E}[\theta + b(\theta)]$. We have

$$-u'_{A,0}(r) = \mathbf{E}[\theta + b(\theta)] + a'(y_0(r)) = \theta_A + b(\theta_A) + a'(y_0(r)).$$

Apply the implicit function theorem to the ex ante conjugacy condition to get

$$y_1'(r) = -\frac{\theta_A + b(\theta_A) + a'(y_0(r))}{\theta_A + b(\theta_A) + a'(y_1(r))}. \tag{9}$$

For $i = 0, 1$, let $\theta_i(r) = y_A^{-1}(y_i(r))$, with the convention that $y_A^{-1}(y)$ equals $\underline{\theta}$ if $y < y_A(\underline{\theta})$ and equals $\bar{\theta}$ if $y > y_A(\bar{\theta})$. We have

$$\begin{aligned} u'_{A,1}(r) &= -\int_{\theta_0(r)}^{\theta_A} \left[\theta + b(\theta) + a'(y_0(r))\right] f(\theta)\,\mathrm{d}\theta + y_1'(r) \int_{\theta_A}^{\theta_1(r)} \left[\theta + b(\theta) + a'(y_1(r))\right] f(\theta)\,\mathrm{d}\theta \\ &> -\int_{\theta_0(r)}^{\theta_A} \left[\theta_A + b(\theta_A) + a'(y_0(r))\right] f(\theta)\,\mathrm{d}\theta + y_1'(r) \int_{\theta_A}^{\theta_1(r)} \left[\theta_A + b(\theta_A) + a'(y_1(r))\right] f(\theta)\,\mathrm{d}\theta. \end{aligned}$$

Substituting in (9) and simplifying gives

$$\begin{aligned} u'_{A,1}(r) &> -[\theta_A + b(\theta_A) + a'(y_0(r))][F(\theta_1(r)) - F(\theta_0(r))] \\ &\geq -[\theta_A + b(\theta_A) + a'(y_0(r))] \\ &= u'_{A,0}(r), \end{aligned}$$

where the second inequality holds because $\theta_A + b(\theta_A) + a'(y_0(r)) \geq 0$. Thus, $\Delta'_A(r) > 0$.

II. Interval By the proof of Lemma 3 (Appendix A.3), $u_{P,1}$ is single-peaked about $y_0 = y_A(\hat{\theta})$. We check that Δ_A is strictly increasing in y_0. Let $\theta_0(y_0) = y_A^{-1}(y_0)$. We view other quantities as functions of y_0 as well. We have

$$u'_{A,1}(y_0) = \int_{\theta_0(y_0)}^{\bar{\theta}} [\theta + b(\theta) + a'(y_0)] f(\theta)\,\mathrm{d}\theta > \int_{\underline{\theta}}^{\bar{\theta}} [\theta + b(\theta) + a'(y_0)] f(\theta)\,\mathrm{d}\theta = u'_{A,0}(y_0),$$

where the strict inequality holds because $\theta + b(\theta) + a'(y_0) \geq 0$ if and only if $\theta \geq \theta_0(y_0)$.

III. High-point From the utility representation (Lemma 2), $u_{P,1}$ is strictly increasing in y_0. We check that $u_{P,1}$ is strictly quasiconvex in $\bar{y}$. With y_0 fixed, let $\theta^*(\bar{y})$ denote the type that is indifferent between y_0 and $\bar{y}$. By the implicit function theorem, θ^* is differentiable. From the utility representation (Lemma 2), we have

$$u'_{P,1}(\bar{y}) = \int_{\theta^*(\bar{y})}^{\bar{\theta}} [\theta + b(\theta) + a'(\bar{y})][f(\theta) + B'(\theta) f(\theta) + B(\theta) f'(\theta)]\,\mathrm{d}\theta - [\bar{\theta} + b(\bar{\theta}) + a'(\bar{y})] B(\bar{\theta}) f(\bar{\theta}).$$

[24] The derivative jump is $-(\mathbf{E}[\theta] + a'(y_{A,0})) = \mathbf{E}[b(\theta)]$, which is strictly positive by A2.

Therefore,

$$u''_{P,1}(\bar{y}) = a''(\bar{y})\left[\int_{\theta^*(\bar{y})}^{\bar{\theta}} [f(\theta) + B'(\theta)f(\theta) + B(\theta)f'(\theta)]\,\mathrm{d}\theta - B(\bar{\theta})f(\bar{\theta})\right] \\ - \theta^{*\prime}(\bar{y})\left[\theta^*(\bar{y}) + b(\theta^*(\bar{y})) + a'(\bar{y})\right]\left[f(\theta^*(\bar{y})) + B'(\theta^*(\bar{y}))f(\theta^*(\bar{y})) + B(\theta^*(\bar{y}))f'(\theta^*(\bar{y}))\right]. \tag{10}$$

The expression on the second line of (10) is strictly positive by A1 since $-\theta^{*\prime}(\bar{y}) < 0$ and $\theta^*(\bar{y}) + b(\theta^*(\bar{y})) + a'(\bar{y}) < 0$. We claim that if $u'_{P,1}(\bar{y}) \geq 0$, then the first line of (10) is nonnegative. We know $a''(\bar{y}) \leq 0$. For any state θ,

$$\theta + b(\theta) + a'(\bar{y}) \leq \bar{\theta} + b(\bar{\theta}) + a'(\bar{y}) < 0.$$

Using A1, it can be shown that the term in brackets on the first line of (10) is bounded above by $u'_{P,1}(\bar{y})/(\bar{\theta} + b(\bar{\theta}) + a'(\bar{y}))$, which is nonpositive if $u'_{P,1}(\bar{y}) \geq 0$. We conclude that $u''_{P,1}(\bar{y}) > 0$ whenever $u'_{P,1}(\bar{y}) \geq 0$. Hence, $u_{P,1}$ is strictly quasiconvex.

We turn to the comparative statics for Δ_A. Clearly, Δ_A is strictly decreasing in $\bar{y}$. We check that Δ_A is strictly convex in y_0. With $\bar{y}$ fixed, let $\theta^*(y_0)$ denote the type that is indifferent between y_0 and $\bar{y}$. By the implicit function theorem, θ^* is differentiable. Let $\theta_0(y_0) = y_A^{-1}(y_0)$, with the convention that $y_A^{-1}(y)$ equals $\underline{\theta}$ if $y < y_A(\underline{\theta})$. We have

$$u'_{A,0}(y_0) = \int_{\underline{\theta}}^{\bar{\theta}} [\theta + b(\theta) + a'(y_0)]f(\theta)\,\mathrm{d}\theta, \qquad u'_{A,1}(y_0) = \int_{\theta_0(y_0)}^{\theta^*(y_0)} [\theta + b(\theta) + a'(y_0)]f(\theta)\,\mathrm{d}\theta.$$

Differentiate again, noting that $\theta_0(y_0) + b(\theta_0(y_0)) + a'(y_0) = 0$ if $y_0 \geq y_A(\underline{\theta})$, to get

$$u''_{A,0}(y_0) = a''(y_0), \qquad u''_{A,1}(y_0) = \theta^{*\prime}(y_0)[\theta^*(y_0) + b(\theta^*(y_0)) + a'(y_0)]f(\theta^*(y_0)) + a''(y_0)[F(\theta^*(y_0)) - F(\theta_0(y_0))].$$

In the expression for $u''_{A,1}(y_0)$, the first term is a product of three strictly positive factors. Since $a''(y_0) \leq 0$, it follows that

$$\Delta''_A(y_0) = u''_{A,1}(y_0) - u''_{A,0}(y_0) > 0.$$

A.7. Proof of Theorem 2

Lemma 6 *(Adding a point). Fix decisions y_0 and y_1, with $y_0 < y_1$. If $\mathbf{E}[u_A(y_0,\theta)] = \mathbf{E}[u_A(y_1,\theta)]$, then we have*

$$u_{P,1}([y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{y_1\}) \geq u_{P,1}([y_A(\underline{\theta}) \wedge y_0, y_0]) \iff \mathbf{E}[\theta + b(\theta)] \leq \hat{\theta} + b(\hat{\theta}).$$

We use Lemma 6 to complete the proof.

Part 1 Suppose $\mathbf{E}[\theta + b(\theta)] \leq \hat{\theta} + b(\hat{\theta})$. For a contradiction, suppose that there is an optimal minimal delegation set that takes the interval form $D = [y_A(\underline{\theta}), y_0]$. In particular, $u_{P,1}(D) > u_{P,0}(D)$. Let y_1 be the ex ante conjugate of y_0, and set $D' = [y_A(\underline{\theta}), y_0] \cup \{y_1\}$. By construction, $u_{P,0}(D') = u_{P,0}(D)$ and $\Delta_A(D') > \Delta_A(D)$. By Lemma 6, $u_{P,1}(D') \geq u_{P,1}(D)$. Therefore, the principal strictly prefers D' to D.

Next, suppose for a contradiction that there is an optimal minimal delegation set that takes the high-point form $D = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{\bar{y}\}$. In particular, $u_{P,1}(D) > u_{P,0}(D)$. Let y_1 be the ex ante conjugate of y_0. Define the sets

$$D' = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{y_1\}, \qquad D'' = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{y_1 \vee y_A(\bar{\theta})\}.$$

By construction, $u_{P,0}(D') = u_{P,0}(D'') = u_{P,0}(D)$. Since D is optimal, we must have $u_A(y_1,\bar{\theta}) > u_A(\bar{y},\bar{\theta})$ (which implies that $y_1 < \bar{y}$; see the end of the proof of Theorem 1 (Appendix A.5). Hence, $\Delta_A(D') > \Delta_A(D)$ and $\Delta_A(D'') > \Delta_A(D)$. To prove that the principal strictly prefers either D' or D'' to D, we show that $u_{P,1}(D) \leq \max\{u_{P,1}(D'), u_{P,1}(D'')\}$. Let $\hat{y}$ be the $\bar{\theta}$-conjugate of y_0. On the domain $[y_1 \vee y_A(\bar{\theta}), \hat{y}]$, consider the map $g(y) = u_{P,1}([y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{y\})$. By Lemma 5, g is quasiconvex, so $g(\bar{y}) \leq \max\{g(y_1 \vee y_A(\bar{\theta})), g(\hat{y})\}$. By Lemma 6, $g(\hat{y}) = u_{P,1}([y_A(\underline{\theta}), y_0]) \leq u_{P,1}([y_A(\underline{\theta}), y_0] \cup \{y_1\})$. Therefore, $u_{P,1}(D) \leq \max\{u_{P,1}(D''), u_{P,1}(D')\}$, as desired.

Part 2 Suppose $\mathbf{E}[\theta + b(\theta)] > \hat{\theta} + b(\hat{\theta})$. Define $\mathcal{D}$ to be the collection of all hollow or high-point minimal delegation sets D with the property that $u_{P,1}(D) > u_{P,1}(\{y_{P,0}\})$. (Delegation sets violating this inequality cannot be optimal, no matter the cost function.) The union $\cup\mathcal{D}$ is bounded. We claim that there exists a positive constant $L = L(F,a,b)$ such that for every delegation set D in $\mathcal{D}$, the interval delegation set $\tilde{D} = [y_A(\underline{\theta}), y_{A,0}(D) \vee y_{P,0}]$ satisfies

$$u_{P,1}(\tilde{D}) > u_{P,1}(D) \quad \text{and} \quad \Delta_A(D) - \Delta_A(\tilde{D}) \leq L[u_{P,1}(\tilde{D}) - u_{P,1}(D)]. \tag{11}$$

We complete the proof, taking this claim as given.

First, we establish a bound on the induced effort function. Define a positive constant $L' = L'(F,a,b)$ by $L' = \sup_{D\in\mathcal{D}}[u_{P,1}(D) - u_{P,0}(D)]$. The supremum is finite because $\cup\mathcal{D}$ is bounded. Choose a positive constant $K = K(F,a,b)$ such that

$$K^{-1}\exp\left(K\Delta_A([y_A(\underline{\theta}),y_{P,0}])\right)-K^{-1}>LL'. \tag{12}$$

Assume $-\hat{e}''(x)/\hat{e}'(x)\geq K$ for $x>0$. Let $h(x)=\hat{e}(x)/\hat{e}'(x)$, for $x>0$. We have

$$h'(x)=1-\frac{\hat{e}(x)\hat{e}''(x)}{(\hat{e}'(x))^2}=1+h(x)\left(-\frac{\hat{e}''(x)}{\hat{e}'(x)}\right)\geq 1+Kh(x).$$

Apply Grönwall's inequality to the function $\tilde{h}=h+K^{-1}$ at initial points converging downward to 0. We conclude that for all $x>0$,

$$\frac{\hat{e}(x)}{\hat{e}'(x)}=\tilde{h}(x)-K^{-1}\geq K^{-1}e^{Kx}-K^{-1}.$$

For $x\geq\Delta_A([y_A(\underline{\theta}),y_{P,0}])$, it follows from (12) that

$$\frac{\hat{e}(x)}{\hat{e}'(x)}>LL'. \tag{13}$$

We show that the principal can strictly improve upon any delegation set in $\mathcal{D}$. Fix D in $\mathcal{D}$. The definition of $\mathcal{D}$ implies that $u_{P,1}(D)>u_{P,0}(D)$. Let $\tilde{D}=[y_A(\underline{\theta}),y_{A,0}(D)\vee y_{P,0}]$. We prove that $U_P(\tilde{D})>U_P(D)$. By construction, $u_{P,0}(\tilde{D})\geq u_{P,0}(D)$. By (11), $u_{P,1}(\tilde{D})>u_{P,1}(D)$. If $\Delta_A(\tilde{D})\geq\Delta_A(D)$, then we immediately get $U_P(\tilde{D})>U_P(D)$. So we hereafter assume $\Delta_A(\tilde{D})<\Delta_A(D)$. To show that $U_P(\tilde{D})>U_P(D)$, it suffices to show that

$$\hat{e}(\Delta_A(\tilde{D}))[u_{P,1}(\tilde{D})-u_{P,1}(D)]>[\hat{e}(\Delta_A(D))-\hat{e}(\Delta_A(\tilde{D}))][u_{P,1}(D)-u_{P,0}(D)].$$

Since $\hat{e}$ is concave,

$$\hat{e}(\Delta_A(D))-\hat{e}(\Delta_A(\tilde{D}))\leq\hat{e}'(\Delta_A(\tilde{D}))[\Delta_A(D)-\Delta_A(\tilde{D})].$$

Therefore, it suffices to show that

$$\frac{\hat{e}(\Delta_A(\tilde{D}))}{\hat{e}'(\Delta_A(\tilde{D}))}>\frac{[\Delta_A(D)-\Delta_A(\tilde{D})][u_{P,1}(D)-u_{P,0}(D)]}{u_{P,1}(\tilde{D})-u_{P,1}(D)}. \tag{14}$$

By Lemma 5, we have $\Delta_A(\tilde{D})\geq\Delta_A([y_A(\underline{\theta}),y_{P,0}])$, so by (13), the left side of (14) is strictly greater than LL'. By (11) and the definition of L', the right side of (14) is at most LL'.

Part 2—Proof of claim It suffices to prove a modified claim: There exists a positive constant $L=L(F,a,b)$ such that for every delegation set D in $\mathcal{D}$, the delegation set $D_0=[y_A(\underline{\theta})\wedge y_{A,0}(D),y_{A,0}(D)]$ satisfies

$$u_{P,1}(D_0)>u_{P,1}(D)\quad\text{and}\quad u_{A,1}(D)-u_{A,1}(D_0)\leq L[u_{P,1}(D_0)-u_{P,1}(D)]. \tag{15}$$

The original claim then follows from Lemma 5 since for every delegation set D in $\mathcal{D}$, the associated sets $\tilde{D}=[y_A(\underline{\theta}),y_{A,0}(D)\vee y_{P,0}]$ and $D_0=[y_A(\underline{\theta})\wedge y_{A,0}(D),y_{A,0}(D)]$ satisfy

$$u_{P,1}(\tilde{D})\geq u_{P,1}(D_0)\quad\text{and}\quad\Delta_A(D)-\Delta_A(\tilde{D})\leq\Delta_A(D)-\Delta_A(D_0)=u_{A,1}(D)-u_{A,1}(D_0).$$

Before proving the modified claim, we record a calculation that will be useful throughout the rest of the proof. If some type θ^* is indifferent between decisions y and y', then

$$a(y)-a(y')=-(\theta^*+b(\theta^*))(y-y').$$

Therefore, for any state θ, we have

$$\begin{aligned}u_{A,1}(y,\theta)-u_{A,1}(y',\theta)&=\left[\theta+b(\theta)-\theta^*-b(\theta^*)\right](y-y'),\\u_{P,1}(y,\theta)-u_{P,1}(y',\theta)&=\left[\theta-\theta^*-b(\theta^*)\right](y-y').\end{aligned} \tag{16}$$

To prove the modified claim, we first define suitable constants. Define θ_A by $\theta_A+b(\theta_A)=\mathbf{E}[\theta+b(\theta)]$. By assumption, $\theta_A+b(\theta_A)>\hat{\theta}+b(\hat{\theta})$, so $\theta_A>\hat{\theta}$. Define the function $g:[y_A(\underline{\theta}),y_A(\bar{\theta})]\to\mathbf{R}$ by $g(y)=u_{P,1}([y_A(\underline{\theta}),y])$. From the proof of Lemma 3 (Appendix A.3), the function g is continuously differentiable and single-peaked about $y=y_A(\hat{\theta})$. Define the constant

$$M_1=\min\left\{|g'(y)|:y_{A,0}\leq y\leq y_A(\bar{\theta})\right\}.$$

Since $y_{A,0}>y_A(\hat{\theta})$, we know that M_1 is strictly positive. Next, define the positive constants

$$M_2=\int_{\theta_A}^{\bar{\theta}}[\theta_A+b(\theta_A)-\theta]f(\theta)\,\mathrm{d}\theta,\qquad M_3=\int_{\theta_A}^{\bar{\theta}}[\theta+b(\theta)-\theta_A-b(\theta_A)]f(\theta)\,\mathrm{d}\theta.$$

In particular, M_2 is positive by the proof of Lemma 3 (Appendix A.3) because $\theta_A>\hat{\theta}$.

With these constants defined, we can complete the proof. There are two cases.

First we consider hollow delegation sets. Let $D = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup [y_1, y_2]$ be a hollow delegation set in $\mathcal{D}$. Let $D_0 = [y_A(\underline{\theta}) \wedge y_0, y_0]$. We have

$$\begin{aligned} u_{P,1}([y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{y_1\}) - u_{P,1}(D) &= u_{P,1}([y_A(\underline{\theta}), y_1]) - u_{P,1}([y_A(\underline{\theta}), y_2]) \\ &\geq M_1(y_2 - y_1). \end{aligned} \tag{17}$$

From the calculations in (16), we have

$$u_{P,1}(D_0) - u_{P,1}([y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{y_1\}) = M_2(y_1 - y_0). \tag{18}$$

Let $M = \min\{M_1, M_2\}$. Combining (17) and (18) gives

$$u_{P,1}(D_0) - u_{P,1}(D) \geq M(y_2 - y_0). \tag{19}$$

From the calculations in (16), we have

$$\begin{aligned} u_{A,1}(D) - u_{A,1}(D_0) &= \int_{\theta_A}^{\bar{\theta}} [\theta + b(\theta) - \theta_A - b(\theta_A)](y_A(D,\theta) - y_0) f(\theta)\, \mathrm{d}\theta \\ &\leq M_3(y_2 - y_0). \end{aligned} \tag{20}$$

So in the hollow case, by (19) and (20), the desired bound (15) holds with $L = M_3/M$.

Second, we consider high-point delegation sets. Let $D = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{\bar{y}\}$ be a high-point delegation set in $\mathcal{D}$, and let $D_0 = [y_A(\underline{\theta}) \wedge y_0, y_0]$. Let θ^* be the type that is indifferent between y_0 and $\bar{y}$. Thus, $\theta^* > \theta_A > \hat{\theta}$. Apply (16) to get

$$u_{P,1}(D_0) - u_{P,1}(D) = (\bar{y} - y_0) \int_{\theta^*}^{\bar{\theta}} [\theta^* + b(\theta^*) - \theta] f(\theta)\, \mathrm{d}\theta > M_2(\bar{y} - y_0), \tag{21}$$

where the inequality follows from the proof of Lemma 3 (Appendix A.3). By (16), we have

$$u_{A,1}(D) - u_{A,1}(D_0) = (\bar{y} - y_0) \int_{\theta^*}^{\bar{\theta}} [\theta + b(\theta) - \theta^* - b(\theta^*)] f(\theta)\, \mathrm{d}\theta < M_3(\bar{y} - y_0). \tag{22}$$

So in the high-point case, by (21) and (22), the desired bound (15) holds with $L = M_3/M_2$.

To cover both cases, set $L = \max\{M_3/M, M_3/M_2\} = M_3/M$.

Part 3 Suppose $u_{P,1}(\{y_{P,0}\}) \geq u_{P,1}([y_A(\underline{\theta}), y_{A,0}])$. In particular, we have $y_{A,0} > y_A(\hat{\theta})$. Let $\bar{y}_A$ be the $\bar{\theta}$-conjugate of $y_{A,0}$. Note that $\bar{y}_A > y_A(\bar{\theta})$. We prove a stronger version of the theorem statement, with $\bar{y}_A$ in place of $y_A(\bar{\theta})$. Let D be a hollow or high-point minimal delegation set lying below $\bar{y}_A$. If $\max D \leq y_A(\bar{\theta})$, let $\bar{d} = \max D$. If $\max D > y_A(\bar{\theta})$, let $\bar{d}$ be the $\bar{\theta}$-conjugate of $\max D$. In both cases, we have $y_{A,0} \leq \bar{d} \leq y_A(\bar{\theta})$. Therefore,

$$u_{P,1}(D) < u_{P,1}([y_A(\underline{\theta}), \bar{d}]) \leq u_{P,1}([y_A(\underline{\theta}), y_{A,0}]) \leq u_{P,1}(\{y_{P,0}\}),$$

where the first inequality follows from the utility representation (Lemma 2); the second inequality follows from Lemma 5 since $\bar{d} \geq y_{A,0} > y_A(\hat{\theta})$; and the last inequality is by assumption. Since $u_{P,0}(D) \leq u_{P,1}(\{y_{P,0}\})$, we conclude that $U_P(D) < u_{P,1}(\{y_{P,0}\}) = U_P(\{y_{P,0}\})$. Therefore, the principal strictly prefers the singleton $\{y_{P,0}\}$ to every hollow delegation set that lies below $\bar{y}_A$.

A.8. Proof of Lemma 6

Define θ_A by $\theta_A + b(\theta_A) = \mathbf{E}[\theta + b(\theta)]$. Fix decisions y_0 and y_1 with $y_0 < y_1$. If $\mathbf{E}[u_A(y_0, \theta)] = \mathbf{E}[u_A(y_1, \theta)]$, then $u_A(y_0, \theta_A) = u_A(y_1, \theta_A)$. Applying the calculation in (16), we have

$$u_{P,1}([y_A(\underline{\theta}) \wedge y_0, y_0] \cup \{y_1\}) - u_{P,1}([y_A(\underline{\theta}) \wedge y_0, y_0]) = (y_1 - y_0) \int_{\theta_A}^{\bar{\theta}} [\theta - \theta_A - b(\theta_A)] f(\theta)\, \mathrm{d}\theta.$$

As shown in the proof of Lemma 3 (Appendix A.3), the integral on the right is nonnegative if and only if $\theta_A \leq \hat{\theta}$, or equivalently, $\mathbf{E}[\theta + b(\theta)] = \theta_A + b(\theta_A) \leq \hat{\theta} + b(\hat{\theta})$.

A.9. Optimality of high-point delegation

We first outline the argument. Given a bias function b satisfying a condition described below, we select a suitable high-point delegation set $D_{\bar{y}} = [y_A(\underline{\theta}), y_{P,0}] \cup \{\bar{y}\}$. We then construct a family of cost functions indexed by ε. If ε is sufficiently small, so that information acquisition is sufficiently costly, then a necessary condition for a delegation set D to be optimal is that $u_{P,0}(D)$ is sufficiently large. We check that for all ε sufficiently small, the principal strictly prefers the delegation set $D_{\bar{y}}$ to (1) any delegation set D violating the necessary condition (because $u_{P,0}(D)$ is too small); (2) any *interval* delegation set D satisfying the necessary condition (because $\hat{e}(D)$ is too small); and (3) any *hollow* delegation set satisfying the necessary condition (because $u_{P,1}(D)$ is too small).

The construction requires one condition on the decision setting. Let $\hat{y}_P$ denote the ex ante conjugate of $y_{P,0}$. Assume

$$u_{P,1}(\{y_{P,0}\}) > u_{P,1}([y_A(\underline{\theta}), y_{P,0}] \cup \{\hat{y}_P\}). \tag{23}$$

By the utility representation (Lemma 2), condition (23) implies that $u_{P,1}([y_A(\underline{\theta}), y_{P,0}]) > u_{P,1}([y_A(\underline{\theta}), y_{P,0}] \cup \{\hat{y}_P\})$. By Lemma 6, it follows that $\mathbf{E}[\theta + b(\theta)] > \hat{\theta} + b(\hat{\theta})$. In the UQC setting, (23) reduces to the inequality $-1/12 > u_{P,1}([\beta, 1/2] \cup \{1/2 + 2\beta\})$. This condition is cubic in β, and it holds if $\beta > 0.29$.

We turn to the proof. For $\bar{y} > y_A(\bar{\theta})$, let $D_{\bar{y}} = [y_A(\underline{\theta}), y_{P,0}] \cup \{\bar{y}\}$. First, we select suitable thresholds. Let $\bar{y}_P$ denote the $\bar{\theta}$-conjugate of $y_{P,0}$.

Lemma 7 *(Thresholds). If (23) holds, then there exist ex ante conjugates y_0^* and y_1^* with $y_{P,0} < y_0^* < y_{A,0} < y_1^*$ and a high point $\bar{y}$ in $(y_A(\bar{\theta}), \bar{y}_P)$ such that the following are satisfied:*

$$\Delta_A(D_{\bar{y}}) > \Delta_A([y_A(\underline{\theta}), y_0^*]), \tag{24}$$

$$u_{P,1}(D_{\bar{y}}) > u_{P,1}(\{y_{P,0}\}) > u_{P,1}([y_A(\underline{\theta}), y_0^*] \cup \{y_1^*\}). \tag{25}$$

Next, we define the cost function. Define $\mathcal{D}$ to be the collection of all hollow or interval minimal delegation sets D with the property that $u_{P,1}(D) > u_{P,1}(\{y_{P,0}\})$. (Delegation sets violating this inequality cannot be optimal, no matter the cost function.) The union $\cup\mathcal{D}$ is bounded, so we can choose a positive constant H such that $H \geq \Delta_P(D)$ for all D in $\mathcal{D}$. By (24), for each ε in $(0, 1/2)$ there exists a cost function c_ε that induces an effort function $\hat{e} = \hat{e}_\varepsilon$ satisfying

$$\hat{e}([y_A(\underline{\theta}), y_0^*]) \leq \varepsilon^2 < \varepsilon(1-\varepsilon) \leq \hat{e}(D_{\bar{y}}) \leq \sup_{D \in \mathcal{D}} \hat{e}(D) \leq \varepsilon.$$

Hereafter, assume that the cost function is c_ε. To simplify notation, let $u_0 = u_{P,1}(\{y_{P,0}\})$ and $u_1 = \mathbf{E}[u_P(y_P(\theta), \theta)]$. Since $\hat{e}(D_{\bar{y}}) \geq \varepsilon(1-\varepsilon)$, it follows from (25) that

$$\begin{aligned} U_P(D_{\bar{y}}) &= u_0 + \hat{e}(D_{\bar{y}})(u_{P,1}(D_{\bar{y}}) - u_0) \\ &\geq u_0 + \varepsilon(1-\varepsilon)(u_{P,1}(D_{\bar{y}}) - u_0). \end{aligned} \tag{26}$$

Now we bound the principal's payoff over all delegation sets in $\mathcal{D}$. Fix D in $\mathcal{D}$. We separate into cases.

1. Suppose $u_{P,0}(D) \leq u_{P,0}([y_A(\underline{\theta}), y_0^*])$. Since $\hat{e}(D) \leq \varepsilon$, we have

$$\begin{aligned} U_P(D) &= (1 - \hat{e}(D))u_{P,0}(D) + \hat{e}(D)u_{P,1}(D) \\ &\leq (1-\varepsilon)u_{P,0}([y_A(\underline{\theta}), y_0^*]) + \varepsilon u_1. \end{aligned}$$

2. Suppose D is an interval delegation set satisfying $u_{P,0}(D) > u_{P,0}([y_A(\underline{\theta}), y_0^*])$. Thus, $D = [y_A(\underline{\theta}), y_0]$ for some y_0 in $(y_{P,0}, y_0^*)$. By Lemma 5, $\Delta_A(D) < \Delta_A([y_A(\underline{\theta}), y_0^*])$, so $\hat{e}(D) < \hat{e}([y_A(\underline{\theta}), y_0^*]) \leq \varepsilon^2$. Recalling the definition of H above, we have

$$\begin{aligned} U_P(D) &= u_{P,0}(D) + \hat{e}(D)\Delta_P(D) \\ &< u_0 + \varepsilon^2 H. \end{aligned}$$

3. Suppose D is a hollow delegation set satisfying $u_{P,0}(D) > u_{P,0}([y_A(\underline{\theta}), y_0^*])$. Thus, $D = [y_A(\underline{\theta}) \wedge y_0, y_0] \cup [y_1, y_2]$ for some ex ante conjugates y_0 and y_1 and some cap y_2 satisfying $y_0 < y_{A,0} < y_1 \leq y_2$. Since $u_{P,0}(D) > u_{P,0}([y_A(\underline{\theta}), y_0^*])$, we know $y_0 < y_0^*$. By Lemma 5, we have

$$\begin{aligned} u_{P,1}(D) &< u_{P,1}([y_A(\underline{\theta}), y_0^*] \cup [y_1^*, y_2]) \\ &< u_{P,1}([y_A(\underline{\theta}), y_0^*] \cup \{y_1^*\}), \end{aligned} \tag{27}$$

where the last inequality holds because we have $y_2 \geq y_1 > y_1^* > y_{A,0} > y_A(\hat{\theta})$ by the assumption that $\mathbf{E}[\theta + b(\theta)] > \hat{\theta} + b(\hat{\theta})$. Using (25), we conclude from (27) that $u_{P,1}(D) < u_0$. Clearly, $u_{P,0}(D) \leq u_0$, so $U_P(D) < u_0$.

Given cost function c_ε, we conclude from these cases that for every delegation set D in $\mathcal{D}$,

$$U_P(D) \leq \max\left\{(1-\varepsilon)u_{P,0}([y_A(\underline{\theta}), y_0^*]) + \varepsilon u_1, u_0 + \varepsilon^2 H\right\}. \tag{28}$$

Since $u_{P,0}([y_A(\underline{\theta}), y_0^*]) < u_0$, it follows that for all ε sufficiently small, the right side of (28) is strictly smaller than the lower bound on $U_P(D_{\bar{y}})$ in (26), and hence $\sup_{D \in \mathcal{D}} U_P(D) < U_P(D_{\bar{y}})$.

A.10. Proof of Lemma 7

Assume (23) holds. First we choose $\bar{y}$. Recall that $\bar{y}_P$ denotes the $\bar{\theta}$-conjugate of $y_{P,0}$. We have

$$\lim_{\bar{y}'\uparrow\bar{y}_P} u_{P,1}(D_{\bar{y}'}) = u_{P,1}([y_A(\underline{\theta}), y_{P,0}]) > u_{P,1}(\{y_{P,0}\}),$$

where the inequality follows from the utility representation (Lemma 2). Thus, we may select $\bar{y}$ in $(y_A(\bar{\theta}), \bar{y}_P)$ such that $u_{P,1}(D_{\bar{y}}) > u_{P,1}(\{y_{P,0}\})$. Next, we choose y_0^*. Clearly, $\Delta_A(D_{\bar{y}}) > \Delta_A([y_A(\underline{\theta}), y_{P,0}])$. By (23), we have $u_{P,1}(\{y_{P,0}\}) > u_{P,1}([y_A(\underline{\theta}), y_{P,0}] \cup \{\hat{y}_P\})$, where $\hat{y}_P$ is the ex ante conjugate of $y_{P,0}$. The payoffs Δ_A and $u_{P,1}$ are continuous in the delegation set parameters, and ex ante conjugation is a continuous function. Therefore, we may select y_0^* in $(y_{P,0}, y_{A,0})$ such that (24) and (25) are satisfied (with y_1^* equal to the ex ante conjugate of y_0^*).

A.11. Proof of Theorem 3

Following the proof of Szalay (2005, Proposition 3, p. 1181), it can be shown that if (5) holds, then with an unbiased agent, the delegation set

$$D^* = \left[\underline{\theta} \wedge (\mathbf{E}[\theta] - r), \mathbf{E}[\theta] - r\right] \cup \left[\mathbf{E}[\theta] + r, \bar{\theta} \vee (\mathbf{E}[\theta] + r)\right]$$

is optimal for some radius $r > 0$.[25] Let e^* denote the unbiased agent's induced effort level $\hat{e}(D^*)$. For all $\beta \geq 0$, let $D(\beta) = \beta + D^*$. For all $\beta \geq 0$, if the principal offers the delegation set $D(\beta)$ to an agent with constant bias β, then the agent's induced effort level is e^*. Let $V(\beta)$ denote the value of the principal's delegation problem when facing an agent with constant bias β. Write $U_P(\cdot;\beta)$, $u_{P,0}(\cdot;\beta)$, $u_{P,1}(\cdot;\beta)$, and $y_A(\cdot;\beta)$ to indicate the dependence of these functions on the agent's bias β. We have

$$\begin{aligned}
V(\beta) - V(0) &\geq U_P(D(\beta);\beta) - U_P(D^*;0) \\
&= (1-e^*)\left[u_{P,0}(D(\beta);\beta) - u_{P,0}(D^*;0)\right] + e^*\left[u_{P,1}(D(\beta);\beta) - u_{P,1}(D^*;0)\right] \\
&= (1-e^*)\left[-(\beta-r)^2 + r^2\right] + e^*\,\mathbf{E}\left[-\left(y_A(D^*,\theta;0) + \beta - \theta\right)^2 + \left(y_A(D^*,\theta;0) - \theta\right)^2\right] \\
&= (1-e^*)\left[-(\beta-r)^2 + r^2\right] - e^*\beta^2 - 2e^*\beta\,\mathbf{E}\left[y_A(D^*,\theta;0) - \theta\right].
\end{aligned}$$

Since the state distribution is symmetric about its mean, the expectation in the last line vanishes, so we have

$$\begin{aligned}
V(\beta) - V(0) &\geq (1-e^*)\left[-(\beta-r)^2 + r^2\right] - e^*\beta^2 \\
&= \beta\left[2r(1-e^*) - \beta\right].
\end{aligned}$$

Let $\bar{\beta} = 2r(1-e^*)$. We conclude that $V(\beta) > V(0)$ for all β in $(0, \bar{\beta})$.

References

Aghion, Philippe, Tirole, Jean, 1997. Formal and real authority in organizations. J. Polit. Econ. 105 (1), 1–29.
Aliprantis, Charalambos D., Border, Kim C., 2006. Infinite Dimensional Analysis: A Hitchhiker's Guide, 3rd ed. Springer.
Alonso, Ricardo, Matouschek, Niko, 2008. Optimal delegation. Rev. Econ. Stud. 75 (1), 259–293.
Amador, Manuel, Bagwell, Kyle, 2013. The theory of optimal delegation with an application to tariff caps. Econometrica 81 (4), 1541–1599.
Ambrus, Attila, Egorov, Georgy, 2017. Delegation and nonmonetary incentives. J. Econ. Theory 171, 101–135.
Argenziano, Rossella, Severinov, Sergei, Squintani, Francesco, 2016. Strategic information acquisition and transmission. Am. Econ. J. Microecon. 8 (3), 119–155.
Bubb, Ryan, Warren, Patrick L., 2014. Optimal agency bias and regulatory review. J. Leg. Stud. 43, 95–135.
Che, Yeon-Koo, Kartik, Navin, 2009. Opinions as incentives. J. Polit. Econ. 117 (5), 815–860.
Demski, Joel S., Sappington, David E.M., 1987. Delegated expertise. J. Account. Res. 25 (1), 68–89.
Gerardi, Dino, Yariv, Leeat, 2008. Costly expertise. Am. Econ. Rev. Pap. Proc. 98 (2), 187–193.
Goltsman, Maria, Hörner, Johannes, Pavlov, Gregory, Squintani, Francesco, 2009. Mediation, arbitration and negotiation. J. Econ. Theory 144 (4), 1397–1420.
Holmström, Bengt, 1977. On Incentives and Control in Organizations. Ph.D. dissertation. Stanford University.
Ivanov, Maxim, 2010. Informational control and organizational design. J. Econ. Theory 145 (2), 721–751.
Kleiner, Andreas, 2022. Optimal delegation in a multidimensional world. arXiv:2208.11835.
Kováč, Eugen, Mylovanov, Tymofiy, 2009. Stochastic mechanisms in settings without monetary transfers: the regular case. J. Econ. Theory 144 (4), 1373–1395.
Krähmer, Daniel, Kováč, Eugen, 2016. Optimal sequential delegation. J. Econ. Theory 163, 849–888.
Li, Hao, 2001. A theory of conservatism. J. Polit. Econ. 109 (3), 617–636.
Martimort, David, Semenov, Aggey, 2006. Continuity in mechanism design without transfers. Econ. Lett. 93 (2), 182–189.
Melumad, Nahum D., Shibano, Toshiyuki, 1991. Communication in settings with no transfers. Rand J. Econ. 22 (2), 173–198.
Milgrom, Paul, Segal, Ilya, 2002. Envelope theorems for arbitrary choice sets. Econometrica 70 (2), 583–601.
Rogoff, Kenneth, 1985. The optimal degree of commitment to an intermediate monetary target. Q. J. Econ. 100 (4), 1169–1189.
Semenov, Aggey, 2018. Delegation to a possibly ignorant agent. Can. J. Econ. 51 (1), 64–93.
Szalay, Dezsö, 2005. The economics of clear advice and extreme options. Rev. Econ. Stud. 72 (4), 1173–1198.

[25] Our conclusion is slightly weaker than that in Szalay (2005, Proposition 3, p. 1181) because we do not restrict to delegation sets within $[y_A(\underline{\theta}), y_A(\bar{\theta})]$. If the state is uniformly distributed on $[0,1]$, then $u_{P,1}(\{0,1\}) = u_{P,1}(\{1/2\})$, so the optimal radius must be strictly smaller than $1/2$.