

Creating Experts: Optimal Delegation vs Information Design ^{*}

Péter Eső[†], Sophie Kreutzkamp[‡], and Dezső Szalay[§]

August 24, 2026

Abstract

We analyse a principal–expert model with moral hazard in flexible information acquisition. Preferences are aligned except for the expert’s cost of experimenting. We compare two incentive schemes widely used in practice: principal-led information design followed by cheap talk, and optimal delegation followed by endogenous information acquisition. Extremes-refining, potentially garbled binary experiments (a class within bi-pooling policies) are chosen in both cases. Incentives for information acquisition are provided by effectively excluding actions near the players’ shared ex-ante ideal point. Information design outperforms optimal delegation when the expert’s marginal cost of information-refinement is high; the opposite holds vice versa.

JEL: D82, D83

Keywords: strategic information transmission; cheap talk; moral hazard; flexible information acquisition; optimal delegation; delegated expertise

^{*}Sophie Kreutzkamp gratefully acknowledges funding by the European Union (ERC, IMD-101040122) and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2047/2 – 390685813. Dezső Szalay acknowledges financial support from the German Research Foundation (DFG) through the Cluster of Excellence Hausdorff Center for Mathematics (HCM). We thank conference and seminar participants at the 2026 Theros Conference (Spetses), Oxford, KTI Budapest, as well as thoughtful comments by Anton Kolotilin, Doron Ravid, Péter Vida, and Andy Zapechelnyuk. All errors are our own.

[†]Department of Economics, Oxford University; peter.eso@economics.ox.ac.uk.

[‡]Department of Economics, University of Manchester; sophie.kreutzkamp@manchester.ac.uk.

[§]Department of Economics, University of Bonn, and CEPR; szalay@uni-bonn.de.

1 Introduction

Corporations and organizations, such as university departments, routinely rely on expert committees for policy decisions. Their advice is valuable but effort-intensive. Rules governing how these committees' findings are channelled into the decision-making process vary widely. Consider two contrasting examples from the same master's program at one of the authors' institutions: an admissions committee and a program-reform committee. Admissions decisions rest on pre-specified criteria, e.g. coursework and grades, and require careful file-by-file review by academics to verify compliance. Reviewer(s) communicate assessments to a convenor who makes the final decision. By contrast, the reform committee is asked to identify changes to the status quo, based on any information it deems appropriate to collect. Its recommendations are implemented more-or-less directly as accreditation procedures create time pressure.

The core incentive problem is that faculty members find it costly, in terms of time away from research, to engage in committee work. Monetary incentives are absent; why are the institutional solutions otherwise so different? In admissions, the committee takes on an advisory role and faces restrictions on the information it may acquire; in program reform, the committee has more executive power, with constrained decision-making authority but full discretion over the information it wants to factor in.

A key distinction between the two decision problems lies in the costs and benefits of obtaining more precise information. In admissions, what matters is the average quality of the incoming cohort; the marginal decision about any single applicant has limited impact. Even a largely random selection process would often perform reasonably well. By contrast, the costs of a poorly designed reform of the master's program are substantial. In other words, incentives for the experts to obtain more precise information at the margin seem much better aligned in the latter case.

Our stylized model formalizes these tradeoffs and characterizes the conditions under which each arrangement dominates the other. We consider a setting in which a principal and an expert share identical preferences over actions conditional on a state of nature that is drawn from a continuum. Information about the state can be acquired in arbitrarily flexible ways, at a variable private cost borne by the expert. This creates a problem of moral hazard in information acquisition by the expert.

The principal considers two modes of decision-making that may mitigate moral hazard. One option involves the principal specifying an experiment — a signal-generating process — that the expert is asked to perform privately and report the outcome of. Here,

the expert can shirk by not acquiring information and sending an arbitrary message. This arrangement is termed principal-led information design followed by cheap talk, which corresponds to how the admissions committee operates: they check applicant files against pre-specified criteria and recommend accept/reject decisions. The other option for the principal is optimal delegation, where she defines the set of admissible actions while allowing the expert to decide what information to acquire before making a choice within those constraints. Here the principal cannot observe whether or what type of information is acquired by the expert. This mirrors the workings of the program reform committee.

Key features of the environment where we compare these incentive schemes include a continuum of states and actions with a quadratic mismatch loss, identical for principal and expert, and fully flexible information acquisition whose cost depends on the expectation of the squared posterior mean generated by the expert’s signal. The fact that the cost of a signal depends only on posterior means renders the information design problem tractable; the quadratic specification makes benefits and costs comparable.

Our main result is that optimal delegation outperforms principal-led information design when the marginal benefit of improved information exceeds its marginal cost for the expert; information design trumps delegation when the reverse is true. In the latter case, the expert acquires at most a binary signal, which itself may be garbled via bi-pooling. When the marginal cost of obtaining granular information is high, the principal can prevent garbling by specifying precisely what information must be examined, thereby achieving better outcomes than under delegation.

Our characterization of the optimal delegation set for the case of high marginal cost of signal refinement is of independent interest. We prove that the optimal minimal delegation set — after eliminating redundant choices — contains only two actions. This is so because the expert never acquires more than two signals, no matter what delegation set the principal offers. This “either-or” delegation is common in practice; our explanation in a setting with a continuum of states and actions and involving moral hazard but otherwise perfectly aligned preferences is novel to the delegation literature.

Both incentive schemes, information design and optimal delegation, ‘work’ via the same principle: the principal makes the expert acquire information by preventing him from inducing actions near the status quo. Differences arise because under optimal delegation, moral hazard is mitigated by the distortion of the contractible action, but the expert acquires as much information as he deems optimal. In contrast, in principal-led information design followed by communication, actions remain efficient conditional on the signal the expert is allowed to acquire.

Our analysis is motivated by stylized real-life examples, but it is fundamentally theoretical in nature. It concerns the mitigation of moral hazard in information acquisition by a like-minded expert via realistic incentive schemes. We study the tradeoff between committing to information in cheap-talk communication but obviously not actions (by the very nature of cheap talk communication), and committing to actions but not to information. The comparison fits our leading practical application. From a theoretical point of view, there is a non-trivial tradeoff only when the expert can choose the experiment freely under the optimal delegation of actions, and the principal chooses the information structure under cheap talk. When the information structure is chosen by the same actor, cheap talk is always outperformed by optimal delegation.¹

We find that some form of commitment to an expert’s advice making ‘impact’ improves incentives for information acquisition. Be it by way of not allowing the advisor to suggest or implement incremental changes, or by not allowing him to learn about subtle differences between marginal changes, these institutions involve a commitment to depart from the status quo. The preferred institutional solution depends on the expert’s marginal cost of refining information: optimal delegation is best when this cost is relatively low, whereas principal-led information design followed by communication dominates when the marginal cost of granular signals is high.

Our distinct contribution is to bring information design considerations into the analysis of comparative institutions of decision making in organizations. This allows us to shed new light on the problem of delegated expertise pioneered by Demski and Sappington (1987). Vice versa, our analysis extends the Bayesian persuasion literature into hitherto unexplored domains. As a benchmark, we solve a problem of information design subject to a moral hazard constraint; it involves unobserved, flexible experiment choice followed by cheap talk communication. The alternative institution that we consider is the optimal delegation of actions to incentivize the agent’s costly and flexible information acquisition. Our modeling assumptions, the comparison of real-life institutions, and the explanations in terms of marginal benefits and costs of granular information, are new to the literature. We explain the connections in more depth in section 6.

The paper is structured as follows. In the next section we introduce the general environment and the two modes of communication and decision making that we compare.

¹On the one hand, starting from the optimal information structure in the cheap-talk problem, if the principal could commit to an information structure as well under delegation, she could replicate the cheap-talk outcome, hence optimal delegation would obviously be an improvement. On the other hand, if the expert could freely choose the experiment in the cheap-talk problem, the principal could again do at least as well by a commitment to actions through optimal delegation.

Then, in Section 3, we present our main result concerning when principal-led information design followed by cheap talk dominates optimal delegation and vice versa. The rest of the paper is devoted essentially to the proof. We fully solve the principal's information design problem in Section 4. Then, in Section 5, we establish that optimal delegation improves upon it when the marginal cost of information refinement is below a certain threshold, and that the opposite holds when the marginal cost is above the same threshold. Section 6 discusses related literature and some extensions worthy of future research.

2 Environment and model

There are two players: a principal (P , she) and an expert (E , he). There is a state of nature: random variable $\tilde{\omega}$ with a symmetric distribution F on support $[-\bar{\omega}, \bar{\omega}] \subseteq \mathbb{R}$. The mean of $\tilde{\omega}$ is zero by symmetry; assume its variance is $\sigma^2 < \infty$ and that F has density f . The distribution of $|\tilde{\omega}|$ is denoted by $\bar{F}$, with density $\bar{f}(\omega) = 2f(\omega)$ for all $\omega \in [0, \bar{\omega}]$.

The expert may covertly observe a costly random signal $\tilde{\mu}$ about the state. A typical realization of $\tilde{\mu}$ is $\mu = \mathbb{E}[\tilde{\omega} \mid \tilde{\mu} = \mu]$, i.e., we identify, without loss of generality in our environment, the signal directly with the posterior mean of $\tilde{\omega}$ that it generates. In the delegation model E can choose any $\tilde{\mu}$, in principal-led information design P specifies it directly. The distribution of $\tilde{\mu}$ is π , a mean-preserving contraction of the prior. Let Π denote the set of all mean-preserving contractions of the prior.

We call $\pi \in \Pi$ an experiment and denote its cost, borne by the expert, by $c(\pi)$. The cost of the null experiment, $\pi_\emptyset$ (with mass point at zero, generating $\tilde{\mu} \equiv 0$) is zero. For all $\pi \neq \pi_\emptyset$, the cost depends on the expectation of a convex function of the posterior mean. For tractability, we impose a quadratic specification,

$$c(\pi) = k + \lambda \int \mu^2 d\pi(\mu).$$

For instance, the cost of the experiment that reveals the state is $c(F) = k + \lambda\sigma^2$.

The principal cannot directly observe the realized value of the signal, nor can she observe whether or which signal the expert has chosen. One of the players picks a real-valued action, y . Under cheap talk following information design it is the principal, while under delegation it is the expert. Ex ante, the induced action (rule) is a random variable measurable with respect to π .

The players want the action to match the state: each gets a gross payoff of $-(y - \omega)^2$ from action y taken in state ω . Conditional on posterior mean realization μ , each player's

interim expected gross payoff is

$$-\mathbb{E} \left[(y - \tilde{\omega})^2 \mid \mu \right] = -(y - \mu)^2 - \text{Var} [\tilde{\omega} \mid \mu].$$

Taking expectation with respect to the posterior mean distribution, each player's ex-ante expected gross payoff can be rewritten as

$$\underbrace{- \int (y - \mu)^2 d\pi(\mu)}_{:=U(y,\pi)} + \underbrace{\int \mu^2 d\pi(\mu)}_{:=V(\pi)} - \sigma^2,$$

where we used the variance decomposition identity $\sigma^2 = \int \mu^2 + \text{Var} [\tilde{\omega} \mid \mu] d\pi(\mu)$.

The first term in this formula, $U(y, \pi) = - \int (y - \mu)^2 d\pi(\mu)$, is each player's expected loss from the action mismatching the state. The second term, $V(\pi) = \int \mu^2 d\pi(\mu)$, is the value of information acquired by E ; it is the expected variability of the posterior mean. Both U and V depend on the distribution of the posterior mean, $\pi \in \Pi$; mismatch loss U also depends on the action rule, y , which is measurable with respect to π . The third term is the negative of the baseline variance of the prior.

No information, $\pi = \pi_\emptyset$, coupled with action $y \equiv 0$ yields an expected gross payoff of $-\sigma^2$ to each player; this is the outcome if the principal does not get meaningful advice from the expert. At the other extreme, under full information, $\pi = F$, with the action matching the state, $y \equiv \tilde{\omega}$, each player gets zero expected gross payoff. In each of these examples the induced action is efficient subject to the information acquired by E .

The expert's expected net payoff from experiment $\pi \in \Pi \setminus \pi_\emptyset$ and π -measurable action rule y is

$$U(y, \pi) + V(\pi) - \sigma^2 - c(\pi) = U(y, \pi) + (1 - \lambda)V(\pi) - k - \sigma^2. \quad (1)$$

For $\pi = F$ and $y \equiv \tilde{\omega}$, full information and efficient action, we have $U = 0$ and $V = \sigma^2$, and the expert's expected net payoff becomes $-k - \lambda\sigma^2$.

The principal offers transfer t to the expert for his participation. Knowing that the expert's outside option is $-\sigma^2$, the principal sets t to make his participation constraint bind:

$$t = -U(y, \pi) - (1 - \lambda)V(\pi) + k. \quad (2)$$

Transfers contingent on specific messages or outcomes are not feasible.

By offering t according to (2), the principal becomes 'residual claimant' to social surplus. Her payoff gain relative to her outside option $-\sigma^2$ is $W = U(y, \pi) + V(\pi) - t$, which

can be rewritten as

$$W = 2U(y, \pi) + (2 - \lambda)V - k. \quad (3)$$

By comparing (1) and (3) we note that refining information about the state (i.e., increasing V , the variability of the posterior mean) is valuable for the expert alone if and only if $\lambda < 1$, whereas it is valuable for the principal (and society) if and only if $\lambda < 2$.

2.1 Tension in the model: the key incentive problem

In our environment, tension arises not from the ex-post misalignment of preferences, but from moral hazard: if the cost of information acquisition is too high, then the expert would rather not learn the realization of the state.

It is socially desirable for the expert to learn the state, but his private gain is less than the cost of perfect information acquisition, if $2\sigma^2 \geq k + \lambda\sigma^2 > \sigma^2$, equivalently

$$\sigma^2 \geq k + (\lambda - 1)\sigma^2 > 0. \quad (4)$$

Throughout the paper we assume that condition (4) holds. We also posit $k, \lambda \geq 0$ (non-negative fixed and marginal costs of information acquisition), and hence $\lambda \leq 2$ by (4).

By assumption (4), the first-best — E learning the state and inducing an action that matches it— is unattainable. In the paper we compare two incentive schemes. One is *principal-led information design*, which amounts to E choosing an experiment designed by P , reporting its outcome, and P taking an action. The other is *optimal delegation*, that is, E choosing any experiment and then an action from a delegation set restricted by P .

A specific type of deviation that needs to be ruled out in either model is E agreeing to participate, but then secretly not acquiring information and then inducing the action that is most attractive given the prior, i.e., the action closest to the prior mean, zero. If the smallest action in absolute value that E can induce with no information is y_L , then he prefers experiment π coupled with action rule y as long as his expected net payoff in (1) weakly exceeds $-y_L^2 - \sigma^2$, that is,

$$U(y, \pi) + (1 - \lambda)V(\pi) + y_L^2 \geq k. \quad (5)$$

Note that given assumption (4), if $y_L = 0$, then it is impossible to induce information acquisition because (5) fails to hold. In each problem that we consider there is an analogue to this incentive or ‘no-shirking’ constraint.

2.2 Principal-led information design

The principal-led information design problem is formulated as the following game.

1. The principal selects a transfer $t \in \mathbb{R}$ and an experiment $\pi \in \Pi$.
2. The expert decides whether to participate; if he does then he receives transfer t , otherwise he exits and each players get a payoff of $-\sigma^2$.
3. If he participates, the expert covertly chooses whether or not he privately learns the outcome of experiment π . If he does then he incurs cost $c(\pi)$ and observes signal realization μ drawn according to π . If he chooses not to learn the signal then he incurs no cost and remains uninformed.
4. The expert sends a costless message corresponding to a feasible realization of the signal. He can costlessly misrepresent the value of μ in the support of π even if he has not carried out experiment π . However, if he reports $\mu \notin \text{supp}(\pi)$ then he forfeits the transfer.
5. Finally, the principal chooses an action, $y \in \mathbb{R}$, and payoffs are obtained accordingly.

This is a sequential game with incomplete information: in the final stage, P does not directly observe whether the posterior mean reported by E is indeed the signal he has observed, or indeed, whether he has undertaken the experiment at all. Therefore, the equilibrium concept we use is perfect Bayesian equilibrium (PBE).

The continuation game from stage 4 onwards is a cheap-talk game with perfectly aligned preferences. Even though that game in isolation has multiple equilibria (including babbling), the only play that is consistent with PBE in the overall game is one in which all information obtained by E is truthfully transmitted to P , moreover, the principal's ex-ante expected gain, W , is maximized subject to E 's binding participation and incentive constraints. This is so because P makes a take-it-or-leave-it transfer offer for E 's participation; anticipating any other continuation would make her offer irrational.²

Having made these preliminary observations, we reformulate the PBE of the game as a constrained maximization problem for P as follows.

If E participates and obediently carries out experiment $\pi \neq \pi_\emptyset$ that P has made available for him, then he reports the realization of the posterior mean truthfully, and P takes

²The transfer is a 'bond' that is forfeited by E if he sends a message inconsistent with all posteriors in π . This rules out a deviation where E does not carry out π and conveys to P that he has no information. As a result, the unique PBE that we derive is in fact neologism-proof.

the matching action. Hence the induced action rule in information design is $y \equiv \mu$, and $U(y, \pi)$ vanishes in all expressions concerning the players' payoffs.

If E participates (by accepting transfer t) but does not carry out experiment π , then he reports μ in the support of π closest to zero and induces $y_L = \min \{|\mu| : \mu \in \text{supp}(\pi)\}$. His incentive constraint, (5), becomes

$$(1 - \lambda)V(\pi) + \min_{\mu \in \text{supp}(\pi)} \mu^2 \geq k. \quad (6)$$

In a PBE of the information-design game with information acquisition, P chooses $\pi \in \Pi \setminus \pi_\emptyset$ to maximize (3) subject to E 's incentive constraint, (6). Therefore, if P chooses a non-null experiment, then her payoff gain is the solution to

$$\begin{aligned} W_{\mathcal{I}} &= \max_{\pi \in \Pi \setminus \pi_\emptyset} \{(2 - \lambda)V(\pi) - k\} \\ \text{s.t. } &(1 - \lambda)V(\pi) + \min_{\mu \in \text{supp}(\pi)} \mu^2 \geq k. \end{aligned} \quad (7)$$

The experiment picked by P for E in equilibrium is denoted by $\pi_{\mathcal{I}}$, where subscript $\mathcal{I}$ refers to 'information design' solution. Assumption (4) implies that (5) is violated for a fully revealing experiment; in particular, $\pi_{\mathcal{I}}$ cannot be fully revealing around the prior mean. We characterize the best among such experiments in Proposition 1 of Section 4.

The objective in (7) is convex in the posterior mean generated by experiment π , therefore a more spread-out π , i.e., a more informative signal, increases its value. The left-hand side of the constraint is similarly convex when $\lambda \leq 1$, but it is strictly concave when $\lambda > 1$. A more informative signal relaxes the constraint in the former case and tightens it in the latter case. We expect to find that the principal-optimal experiment is as spread-out as possible, subject to the constraint binding.

If the constraint set is empty or otherwise $W_{\mathcal{I}} < 0$, then P 'gives up' on information design and simply chooses $y = 0$ without consulting E . Therefore, in this game, P 's equilibrium payoff gain relative to $-\sigma^2$ is $W_{\mathcal{I}}^* = \max\{W_{\mathcal{I}}, 0\}$.

2.3 Optimal delegation and delegated information design

The principal's delegation problem is also formulated as a game. The timing and the payoffs are similar to those in principal-led information design, but the moves are different.

1. The principal selects a transfer $t \in \mathbb{R}$ and a delegation set $D \subseteq \mathbb{R}$, where D is closed.

2. The expert decides whether to participate; if he does then he receives transfer t , otherwise he exits and each player gets a payoff of $-\sigma^2$.
3. If E participates, then he covertly chooses an experiment π . If $\pi \neq \pi_\emptyset$ then he pays $c(\pi)$ and privately observes signal realization μ that π generates. If $\pi = \pi_\emptyset$ then he incurs no cost and remains uninformed.
4. Finally, the expert selects action $y \in D$, and payoffs are obtained accordingly.

The players have symmetric information at the time of P 's delegation decision; even though E may generate a private signal after that, P does not have an action in the continuation, hence her beliefs play no role. The game's perfect Bayesian equilibrium is essentially subgame perfect, and we can find it by backward induction. This, again, can be reformulated as an optimization problem for P , characterized next.

Let $y_D(\mu)$ denote the action picked by E from delegation set D (which is chosen by P) and posterior mean μ that E has observed. It is the projection of μ onto D , i.e., an element of D that is closest to μ ; in case of ties, it is immaterial which one, because E and P have aligned preferences over actions and states.

If E chooses to acquire some information, then he picks

$$\pi_D \in \arg \max_{\pi \in \Pi \setminus \pi_\emptyset} \left\{ U(y_D, \pi) + (1 - \lambda)V(\pi) - k - \sigma^2 + t \right\}, \quad (8)$$

where $U(y_D, \pi) = - \int [y_D(\mu) - \mu]^2 d\pi(\mu)$ and $V(\pi) = \int \mu^2 d\pi(\mu)$ as defined earlier, and transfer t is set to make the expert's participation constraint bind.

The expert's incentive constraint is (5) with $y = y_D$, $\pi = \pi_D$, and $y_L = y_D(0)$; the latter follows as $\mu = 0$ in case E remains uninformed. The principal's objective is (3).

We conclude that in any partially informative equilibrium of the delegation game, P 's ex-ante expected payoff gain in excess of $-\sigma^2$ is the solution to

$$\begin{aligned} W_D &= \max_{D \subseteq \mathbb{R}} \{ 2U(y_D, \pi_D) + (2 - \lambda)V(\pi_D) - k \} \\ \text{s.t. } y_D(\mu) &\in \arg \min_{y \in D} (y - \mu)^2 \text{ for all } \mu \in \text{supp}(\pi_D) \\ \pi_D &\in \arg \max_{\pi \in \Pi \setminus \pi_\emptyset} \{ U(y_D, \pi) + (1 - \lambda)V(\pi) \} \\ U(y_D, \pi_D) &+ (1 - \lambda)V(\pi_D) + y_D(0)^2 \geq k. \end{aligned} \quad (9)$$

The first constraint is that y_D is the projection of μ onto set D . The second one, equivalent to (8) with the constant terms omitted, expresses that E chooses the best non-null experi-

ment in response to delegation set D . The third constraint is that E weakly prefers to carry out π_D rather than staying uninformed: this is his incentive or ‘no-shirking’ constraint.

If $W_D < 0$ in the solution to the above problem, then P ‘gives up’ on delegation and simply picks $y = 0$ without consulting E . Therefore, in the delegation game, P ’s equilibrium payoff gain is $W_D^* = \max\{W_D, 0\}$. Subscript D refers to the solution under optimal delegation; delegation set D (typeset in subtly different font) is part of the solution.

We analyse this problem in more depth in Section 5, Propositions 2 and 3, separately for the cases when $\lambda < 1$ and $\lambda > 1$, respectively. We do that under a mild assumption on the prior, namely, that it satisfies

$$\frac{\partial}{\partial x} \mathbb{E}[\tilde{\omega} | a \leq \tilde{\omega} \leq x] \leq 1 \text{ and } \frac{\partial}{\partial x} \mathbb{E}[\tilde{\omega} | x \leq \tilde{\omega} \leq b] \leq 1 \text{ for all } a < x < b. \quad (10)$$

That is, the expected value of the state truncated to an interval does not increase at a rate greater than 1 in the endpoints of the truncation.³ The purpose of this regularity condition is to ensure that the expert’s optimal choice of an experiment has a unique solution.⁴

By assumption (4), if $y_D(0) = 0$, then no information is acquired by E . In particular, this implies that full (unrestricted) delegation cannot induce any learning. This observation originates from Szalay (2005) and carries over to the present delegation game. However, as the information technology in this paper is substantially different, so is the analysis of the optimal delegation problem.

3 Comparison of principal-led information design and optimal delegation with delegated information design

Under assumption (4), the first best, E learning the state and inducing a matching action, is not attainable in our environment: the expert has insufficient incentives to exert effort and become fully informed. Some type of restriction, either in the form of excluding experiments that reveal the state near the prior mean, or excluding actions near the status quo action (the action that would be optimal under the prior, i.e., also the prior mean), is needed to mitigate the expert’s moral hazard in information acquisition. In this section,

³Bagnoli and Bergstrom (2005) show that both are implied by the logconcavity of the density function. Szalay (2012) proves that the two derivatives summing up to at most 1 is equivalent to logconcavity, hence condition (10) is in fact considerably weaker than logconcavity.

⁴The expert’s signal choice is akin to an optimal quantization problem. If (10) does not hold, then his best reply to a symmetric delegation set may be non-unique and asymmetric. Relatedly, Kieffer (1983) shows that a logconcave distribution admits a unique optimal n -level quantizer for any finite n .

we first describe a class of experiments that will be shown to be part of the solution in each incentive scheme. Then we compare the equilibrium of the principal-led information design game introduced in Section 2.2 with that of principal-optimal delegation seen in Section 2.3. That is the main result of the paper.

3.1 Extremes-refining, potentially garbled binary experiments

Throughout the paper we put no restriction whatsoever on the types of experiments that the expert may carry out: any signal that generates distribution π of posterior means that is a mean-preserving contraction of the prior is allowed. However, in the analysis of both problems, we encounter a specific class of signals, which is worth defining in advance. We call these extremes-refining, potentially garbled binary experiments.

An extremes-refining, potentially garbled binary experiment reveals the state perfectly when the state's realization is beyond a certain threshold, otherwise it generates a symmetric binary signal. It has two parameters: a threshold $a \geq 0$ and a posterior mean $\bar{\mu} \geq 0$; the signal reveals ω when $|\omega| \geq a$, and it generates $\pm\bar{\mu}$ with equal conditional probabilities otherwise. Note that $\bar{\mu} \leq \mu_a := \mathbb{E}[\tilde{\omega} \mid \tilde{\omega} \in [0, a]]$ by Bayesian consistency. If $\bar{\mu} < \mu_a$ then the signal sent when $|\tilde{\omega}| < a$ is garbled. This is a bi-pooling policy; see Kleiner, Moldovanu, and Strack (2021), Arieli, Babichenko, Smorodinsky, and Yamashita (2023). The interval $[-a, a]$ is a bi-pool, outside of which all states are perfectly revealed. If there is no garbling involved, then $[-a, 0]$ and $[0, a]$ each constitute a simple pool.

The figure below illustrates an extremes-refining binary experiment with garbling in the central bi-pooling interval. The vertical lines represent perfectly revealed states; the blue and red regions, with the fill gradient indicating garbling, correspond to equiprobable binary signals that generate posterior means $\pm\bar{\mu}$. These areas could be distributed in infinitely many different ways within the bi-pooling interval $(-a, a)$, as long as the posterior means conditional on the 'red' and 'blue' signals are $\bar{\mu}$ and $-\bar{\mu}$, respectively.




Figure 1: Extremes-refining, garbled binary experiment

Definition. Experiment $\pi_{a,\bar{\mu}}$ is an extremes-refining, potentially garbled binary experiment with parameters $a \geq 0$ and $\bar{\mu} \in [0, \mu_a]$ if it generates the following posterior mean of $\tilde{\omega}$:

$$\tilde{\mu} = \begin{cases} \tilde{\omega} & \text{if } |\tilde{\omega}| \geq a \\ \pm \bar{\mu} & \text{with equal probabilities conditional on } |\tilde{\omega}| < a \end{cases}.$$

Note that the cost of experiment $\pi_{a,\bar{\mu}}$ is

$$c(\pi_{a,\bar{\mu}}) = k + \bar{F}(a)\bar{\mu}^2 + \int_a^{\tilde{\omega}} \omega^2 d\bar{F}(\omega).$$

Recall that $\bar{F}$ is the distribution of $|\tilde{\omega}|$, equivalently the distribution of $\tilde{\omega}$ conditional on $\tilde{\omega}$ having a given sign, therefore $\bar{F}(a) = \Pr[|\tilde{\omega}| < a]$ is the probability that $\tilde{\omega}$ falls in the central bi-pooling interval where $\tilde{\mu} = \pm \bar{\mu}$ is generated with 50-50% conditional probability.

If the binary signal sent when $|\tilde{\omega}| < a$ is not garbled, then $\bar{\mu} = \mu_a$. In this case, the signal that is generated when $|\tilde{\omega}| < a$ equals the true mean of $\tilde{\omega}$ conditional on falling into either $(-a, 0]$ or $[0, a)$, so besides $|\tilde{\omega}| < a$, the sign of $\tilde{\omega}$ is revealed as well.

A special case of this type of experiment is when nothing but the sign of $\tilde{\omega}$ is revealed, i.e., $a = \tilde{\omega}$. Using inconsistent but revealing notation, let $\pi_{\pm}$ denote this experiment, and μ_{+} the mean of $\tilde{\omega}$ conditional on $\tilde{\omega} \geq 0$. That is, $\pi_{\pm} \equiv \pi_{\tilde{\omega}}$ and $\mu_{+} \equiv \mu_{\tilde{\omega}}$. The cost of the binary experiment that reveals nothing but the sign of $\tilde{\omega}$ is $c(\pi_{\pm}) = k + \lambda \mu_{+}^2$.

3.2 Main result

The principal's equilibrium payoff gain measured relative to $-\sigma^2$ in the information-design and delegation games are denoted by $W_{\mathcal{I}}^*$ and $W_{\mathcal{D}}^*$, respectively. P can opt out from using either incentive scheme by choosing action $y = 0$ for zero gain. Therefore $W_{\mathcal{I}}^*$ and $W_{\mathcal{D}}^*$ must be non-negative. The next theorem provides conditions for ranking $W_{\mathcal{I}}^*$ and $W_{\mathcal{D}}^*$.

Theorem 1. *In our environment, under assumptions (4) and (10), the principal ranks principal-led information design and optimal delegation as follows.*

- If $\lambda < 1$ then $W_{\mathcal{D}}^* \geq W_{\mathcal{I}}^*$; moreover, if $W_{\mathcal{I}}^* > 0$ (i.e., in the equilibrium of the information design game the expert acquires some information), then $W_{\mathcal{D}}^* > W_{\mathcal{I}}^*$.
- If $\lambda > 1$ then $W_{\mathcal{I}}^* \geq W_{\mathcal{D}}^*$; moreover, if $W_{\mathcal{D}}^* > 0$ (i.e., in the equilibrium of the delegation game the expert acquires some information), then $W_{\mathcal{I}}^* > W_{\mathcal{D}}^*$.

The bottom line of Theorem 1 is that the principal’s ranking of the two incentive schemes — information design and optimal delegation — depends on whether the expert’s marginal cost of acquiring granular information is below or above 1. When refining his signal is cheap for the expert, $\lambda < 1$, the principal prefers delegation. When the expert’s marginal cost is high, $\lambda > 1$, she prefers information design. The contrast is sharp.

Theorem 1 is the main result of the paper. In the rest of the section we provide some intuition as to why it holds and how it is proved. The proof is developed over the next two sections with technical details relegated to the Appendix.

When $\lambda < 1$, the intuition why the principal prefers delegation over information design is as follows. Suppose that the principal can incentivize the expert to learn something about the state using information design. She can only do so by offering him to carry out an experiment that generates posterior means bounded away from zero (the prior mean), which then induce equilibrium actions bounded away from zero as well. Now imagine that the principal chooses a delegation set that allows the expert to pick any action that is farther from zero than the nearest-to-zero actions induced under information design. Facing these choices in the delegation game, the expert learns more refined information in the extremes of the support than he does under information design. This is so because the expanded action set is not optimal for the signal that the expert would have acquired under information design, and learning is inexpensive for him. By revealed preference, the expert is better off under his choice of an information policy; due to the alignment of preferences and the fact that the principal retains all residual social surplus, she is better off, too. It is also cheaper for her to induce the expert not to shirk. The principal benefits from optimal delegation and delegated information design on both accounts.

In contrast, when $\lambda > 1$, it becomes very difficult for the principal to incentivize any learning by the expert using delegation. We show that she can induce only binary experiments that reveal at most the sign of the state. Moreover, among these experiments, garbling by bi-pooling policies is an issue that the principal would prefer to rule out. As a result, optimal delegation involves a commitment to either one of two actions, spread relatively far apart. In contrast, information design can and does rule out garbling in bi-pooling policies, simply because the principal chooses the experiment that the expert will voluntarily carry out, and in addition, it may even provide richer information structures. Hence the principal strictly prefers to use information design in this case.

The proof of Theorem 1 proceeds as follows. In the next section, we fully solve the principal’s information design problem for all parameter values; this is Proposition 1. Then, in Section 5, we consider the delegation game for $\lambda < 1$ and $\lambda > 1$ separately. For

$\lambda < 1$, we derive the expert's best-response experiment to a delegation set that contains all actions beyond the nearest-to-zero actions induced in the information design solution. It is more informative than the experiment under information design, which implies that optimal delegation is preferred when $\lambda < 1$. This is Proposition 2. Finally, for $\lambda > 1$, we solve the principal's optimal delegation problem. We show that P makes E acquire a binary signal when λ is between 1 and $3/2$ and no information at all when $\lambda > 3/2$. This is Proposition 3. We show that for all $\lambda > 1$, P is better off by using information design.

4 Principal-led information design

In Section 2.2 we showed that in the PBE of the information design game with partial information acquisition, the principal's payoff gain (relative to $-\sigma^2$) equals

$$\begin{aligned} W_{\mathcal{I}} &= \max_{\pi \in \Pi \setminus \pi_{\emptyset}} \{(2 - \lambda)V(\pi) - k\} \\ \text{s.t. } &(1 - \lambda)V(\pi) + \min_{\mu \in \text{supp}(\pi)} \mu^2 \geq k. \end{aligned} \tag{11}$$

We now fully characterize the solution to this problem. Throughout this derivation we do not use the regularity condition (10).

The solution to problem (11) must involve an experiment with posterior means bounded away from zero, for if the support of π contains $\mu = 0$, the constraint in (11) fails,

$$(1 - \lambda)V(\pi) \equiv (1 - \lambda) \int \mu^2 d\pi(\mu) \leq (1 - \lambda)\sigma^2 < k$$

by assumption (4), therefore the expert shirks. Proposition 1 goes further: it claims that the solution involves an extremes-refining binary experiment with no garbling. In such an experiment, the posterior means are indeed bounded away from zero; moreover, as the state is fully revealed beyond a threshold, the variability of the posterior mean outside the two simple pooling intervals near the prior mean is maximized. The threshold is as low as possible without violating the constraint in (11), because greater variability of the posterior mean benefits the principal.

As a preparation for stating Proposition 1, define

$$Q_{\mathcal{I}}(a) = (1 - \lambda) \left[\bar{F}(a)\mu_a^2 + \int_a^{\bar{\omega}} \omega^2 d\bar{F}(\omega) \right] + \mu_a^2, \tag{12}$$

which is the left-hand side of the constraint in (11) for an extremes-refining binary experi-

ment with threshold a and no garbling. The expert's no-shirking constraint, that he prefers this experiment over $\pi_\emptyset$, is equivalent to $Q_I(a) \geq k$; therefore $Q_I(a)$ can be interpreted as the 'strength' of the incentives provided by an extremes-refining binary experiment with threshold a and no garbling. Note that $Q_I(0) = (1 - \lambda)\sigma^2$, which is strictly less than k by assumption (4), expressing the central tension in our environment. We derive additional properties of $Q_I(a)$ in the Appendix. $Q_I(a)$ is not necessarily monotonic in a (when $\lambda < 1$), but it always obtains a maximum on $(0, \bar{\omega}]$. Let

$$\bar{k} = \max_{a \in [0, \bar{\omega}]} Q_I(a),$$

which is the greatest incentive that any extremes-refining, non-garbled binary experiment can provide to the expert to choose it over $\pi_\emptyset$.

Proposition 1. *If $\lambda \in [0, 2)$ and $k \leq \bar{k}$, then in the PBE of the principal-led information design game, P selects an extremes-refining binary experiment with threshold $a^* > 0$ and no garbling ($\bar{\mu} = \mu_{a^*}$), where a^* is the smallest positive solution to $Q_I(a^*) = k$.*

In equilibrium, P 's payoff gain equals

$$W_I^* = (2 - \lambda) \left[\bar{F}(a^*) \mu_{a^*}^2 + \int_{a^*}^{\bar{\omega}} \omega^2 d\bar{F}(\omega) - k \right].$$

If $k > \bar{k}$ then there is no information transmitted in any PBE of the game.

Proposition 1 states that in any equilibrium with information acquisition, P asks E to carry out an extremes-refining binary experiment with no garbling such that his no-shirking constraint is binding at the lowest incentive-compatible threshold, a^* . A partially informative equilibrium exists for parameter values k and λ as long as $Q_I(a) \geq k$ holds for some a , i.e., when the constraint set is non-empty. We show below that the latter requirement is satisfied for a non-trivial set of parameter values when the prior is logconcave.

In the rest of this section we explain the logic and steps of the proof of Proposition 1. Technical details of the derivations are found in the Appendix.

To prove Proposition 1, we first solve (11) while fixing the value of the nearest-to-zero posterior mean at $\gamma > 0$. By Bayesian consistency, $\gamma \in [-\mu_+, \mu_+]$. Therefore, for any fixed $\gamma \in (0, \mu_+]$, we first consider

$$\begin{aligned} W_\gamma &= \max_{\pi \in \Pi_\gamma} \{(2 - \lambda)V(\pi) - k\} \\ \text{s.t. } & (1 - \lambda)V(\pi) + \gamma^2 \geq k, \end{aligned} \tag{13}$$

where Π_γ is the set of experiments (mean-preserving contractions of the prior) whose posterior mean closest to zero is γ in absolute value. We show that the optimum is obtained by an extremes-refining binary experiment. The second step of the proof is to optimize W_γ over $\gamma \in (0, \mu_+]$ subject to the constraint set being non-empty. This yields the overall solution, $W_{\mathcal{I}}^*$.

Proposition 1 is established by Lemmas 1–3, which are proved in the Appendix.

Lemma 1. *For $\gamma \in (0, \mu_+]$, the experiment that minimizes $V = \int \mu^2 d\pi(\mu)$ over $\pi \in \Pi_\gamma$, denoted by $\pi_\gamma^{\min}$, assigns equal probability (i.e., 0.5 each) to $-\gamma$ and γ .*

The experiment that maximizes V over Π_γ is $\pi_\gamma^{\max} = \pi_{a, \bar{\mu}}$ with $\bar{\mu} = \mu_a$ such that $\mu_a = \gamma$.

The first statement of Lemma 1 is quite intuitive. Any policy with posterior means at a minimum distance γ from zero generates a variability in the posterior (i.e., expected μ^2) that weakly exceeds γ^2 , and experiment $\pi_\gamma^{\min}$ inducing $\mu = \pm\gamma$ with 50-50% chance achieves this lower bound. It is less evident and hence carefully proved in the Appendix, that the most informative element of Π_γ , denoted by $\pi_\gamma^{\max}$, is an extremes-refining binary experiment with threshold a and no garbling such that $\mu_a = \gamma$.




Figure 2: Examples of $\pi_\gamma^{\min}$ (left) and $\pi_\gamma^{\max}$ (right)

The informativeness of $\pi_\gamma^{\min}$ and that of $\pi_\gamma^{\max}$ are, respectively,

$$V_\gamma^{\min} = \gamma^2 \quad \text{and} \quad V_\gamma^{\max} = \bar{F}(a)\gamma^2 + \int_a^{\bar{\omega}} \omega^2 d\bar{F}(\omega),$$

where threshold $a > 0$ solves $\gamma = \mu_a \equiv \mathbb{E}[\tilde{\omega} | 0 \leq \tilde{\omega} \leq a]$.

By the linearity of the integral, the variability for any experiment can be written as a convex combination of $V_\gamma^{\min}$ and $V_\gamma^{\max}$. Hence (13) is equivalent to finding the optimal weights on the extremes, that is,

$$\begin{aligned} W_\gamma &= \max_{\alpha \in [0,1]} \left\{ (2 - \lambda) \left[\alpha V_\gamma^{\min} + (1 - \alpha) V_\gamma^{\max} - k \right] \right\} \\ \text{s.t.} \quad & (1 - \lambda) \left[\alpha V_\gamma^{\min} + (1 - \alpha) V_\gamma^{\max} \right] + \gamma^2 \geq k. \end{aligned} \tag{14}$$

As the objective and the constraint are both affine in $\alpha \in [0, 1]$, the constrained optimum, if

it exists, must be obtained either at $\alpha(\gamma) \in \{0, 1\}$ or at $0 < \alpha(\gamma) < 1$ where the constraint is binding. Lemma 2 characterizes the solution together with necessary and sufficient conditions for its existence.

Lemma 2. Fix $\gamma \in (0, \mu_+]$. Let $\alpha(\gamma)$ denote the value of α in the solution of (14) when it exists.

- If $\lambda \leq 1$ then a solution exists if and only if $(1 - \lambda)V_\gamma^{max} + \gamma^2 \geq k$, in which case $\alpha(\gamma) = 0$ for $\gamma < \mu_+$, whereas $\alpha(\mu_+)$ can take any value in $[0, 1]$.
- If $\lambda > 1$ then a solution exists if and only if $(1 - \lambda)V_\gamma^{min} + \gamma^2 \equiv (2 - \lambda)\gamma^2 \geq k$. In this case, for $\gamma < \mu_+$, we have $\alpha(\gamma) = 0$ if $k \leq (1 - \lambda)V_\gamma^{max} + \gamma^2$ otherwise $\alpha(\gamma) = \alpha$ solves $(1 - \lambda) \left[\alpha\gamma^2 + (1 - \alpha)V_\gamma^{max} \right] + \gamma^2 = k$, whereas $\alpha(\mu_+)$ can take any value in $[0, 1]$.

Lemma 2 is illustrated in the next figure. The cases of $\lambda < 1$ and $\lambda > 1$ are depicted in the left and right panels, respectively.

For $\lambda < 1$, the intuition is straightforward. The constraint in (14) is relaxed by putting all weight on the maximally informative experiment, which is indeed what the principal prefers to do. A solution exists as long as $(1 - \lambda)V_\gamma^{max} + \gamma^2 \geq k$, as indicated by the grey region on the horizontal axis. A subtle point to note is that (unlike in the case shown in the figure) the left-hand side of this constraint, the strength of expert's incentive to carry out the experiment instead of shirking, is not necessarily monotonically increasing in γ when $\lambda < 1$. However, it obtains a maximum in $\gamma \in (0, \mu_+]$, and a solution to (14) exists whenever k does not exceed this maximum value, $\bar{k}$.

Consider now the case of $\lambda > 1$, illustrated in the right panel of Figure 3. There is a tension between the principal's objective and the expert's incentive constraint in problem (14). A solution exists if and only if $\gamma \geq \hat{\gamma}$, indicated again by the grey region on the horizontal axis. The principal's preference would be to put all the weight on the maximally informative experiment, which she can do provided $\gamma \geq \gamma^*$. For $\gamma \in [\hat{\gamma}, \gamma^*)$ the constraint can only be met by making the experiment less informative.




Figure 3: Illustration of Lemmas 2 and 3 for $\lambda < 1$ (left) and $\lambda > 1$ (right)

The overall solution to (13) is obtained by optimizing W_γ over γ , that is,

$$W_{\mathcal{I}}^* = \max_{\gamma \in [0, \mu_+]} \left\{ (2 - \lambda) \left[\alpha(\gamma) V_\gamma^{\min} + [1 - \alpha(\gamma)] V_\gamma^{\max} - k \right] \right\}, \quad (15)$$

subject to $\alpha(\gamma)$ being well-defined according to Lemma 2. The solution, γ^* , is given in Lemma 3 and is illustrated in Figure 3 above.

Lemma 3. *Problem (15) admits a solution if and only if, for some $\gamma \in (0, \mu_+]$,*

$$(1 - \lambda) V_\gamma^{\max} + \gamma^2 \geq k. \quad (16)$$

If this condition holds, then (15) has a unique solution γ^* , which is the smallest $\gamma > 0$ for which (16) holds as an equality, and the optimum is reached at $\alpha(\gamma^*) = 0$.

The intuition is straightforward for the case of $\lambda < 1$. The principal chooses the least possible distortion (lowest γ) in the incentive-compatible region (the grey interval in the left panel of Figure 3). The condition of existence is that (16) holds for some γ , which simplifies to $k \leq (2 - \lambda) \mu_+^2$ if the left-hand side is monotonic.

For $\lambda > 1$, as the figure in the right panel illustrates, γ^* is unique provided it exists, i.e., $k \leq (2 - \lambda) \mu_+^2$. Indeed, γ^* is the value of γ where the expert's incentive to acquire information is exactly met for the most informative experiment. To better understand why this is the case, compare two experiments, the minimally informative one at $\gamma = \hat{\gamma}$ and the

maximally informative one at $\gamma = \gamma^*$. Using the binding incentive constraint, we have

$$\begin{aligned} W_{\mathcal{I}}(\gamma^*) &= (2 - \lambda) \left[\frac{k - (\gamma^*)^2}{1 - \lambda} - k \right] \\ &> (2 - \lambda) \left[\frac{k - (\hat{\gamma})^2}{1 - \lambda} - k \right] = W_{\mathcal{I}}(\hat{\gamma}). \end{aligned}$$

Hence the principal prefers the maximally informative policy even though $\gamma^* > \hat{\gamma}$. This argument generalizes to all $\gamma \in (\hat{\gamma}, \gamma^*)$ as well.

Lemmas 1–3 together establish all claims of Proposition 1. Notice that in the solution to the principal's information design problem, the left-hand side of the expert's incentive constraint is

$$Q_{\mathcal{I}}(a^*) = (1 - \lambda)V_{\gamma^*}^{\max} + (\gamma^*)^2$$

where $\gamma^* = \mu_{a^*}$; in the solution the constraint is binding, hence $Q_{\mathcal{I}}(a^*) = k$.

Finally, we note that if the prior f is logconcave, then so is $\bar{f}$ (the prior conditional on $\tilde{\omega} \geq 0$). Then, Theorem 1 of Karlin, Proschan, and Barlow (1961) implies $\sigma^2 < 2\mu_+^2$. Hence

$$Q_{\mathcal{I}}(\bar{\omega}) = (2 - \lambda)\mu_+^2 > \left(1 - \frac{1}{2}\lambda\right)\sigma^2 > (1 - \lambda)\sigma^2,$$

and therefore $\bar{k} > (1 - \lambda)\sigma^2$, so the set of parameter values satisfying condition (4) for which principal-led information design induces information acquisition is non-empty. It can also be shown that logconcavity of the prior density implies that $Q_{\mathcal{I}}$ is unimodal for $\lambda = 0$.

5 Delegation and delegated information design

In Section 2.3 we showed that if the delegation game is to have an equilibrium with information acquisition, that is, $W_{\mathcal{D}} > 0$ in the solution to (9), then $y_{\mathcal{D}}(0) \neq 0$. The following analysis builds on the converse insight that by excluding actions near the prior mean, the principal is able to induce some information acquisition.

Since the nature of the problem (as well as the content of Theorem 1 to be proven) is different for $\lambda < 1$ and $\lambda > 1$, we split the analysis according to these two cases. In the former case ($\lambda < 1$), the incentive problem stems from the fixed cost of information acquisition, k , but at the margin the expert prefers better information. In the latter case ($\lambda > 1$), the expert is averse to even small improvements in information.

Recalling Theorem 1, when $\lambda < 1$ we face the comparatively simpler problem of improving on the benchmark solution of principal-led information design by using a delegation mechanism. We achieve this in Section 5.1 by exposing the expert to a delegation set that excludes all actions closer to the prior mean than $\{-\mu_{a^*}, \mu_{a^*}\}$ and allows all other choices. We derive the expert's best reply to this set and show that both the expert and the principal are better off (Proposition 2).

When $\lambda > 1$, we need to meet a much higher hurdle. In this case, we have to show that there exists no delegation mechanism that does at least as well as the principal-led information design solution. This requires characterizing the optimal delegation set, which we do in Proposition 3. Indeed, we show that in response to any delegation set, the expert chooses to acquire either no information at all, or decides to observe a binary, potentially garbled experiment. For relatively low values of λ (between 1 and $\frac{3}{2}$) the principal delegates to the expert but ends up worse off than under information design. For higher values of λ , the principal gives up on delegation.

Throughout the analysis of the delegation game we impose the regularity condition on the prior state distribution introduced in Section 2.3, expressed in equation (10). This says that the expected value of the state conditional on belonging to an interval increases in the endpoints of the interval at a rate that is not greater than 1. This mild condition (weaker than logconcavity of the prior density) ensures that the expert has a unique best-response experiment given any delegation set with a symmetric central exclusion zone.

5.1 Restricted delegation improving on information design when $\lambda < 1$

The next proposition shows how to improve on the information-design solution via optimal delegation.

Proposition 2. *If $\lambda < 1$, then for all $k \leq \bar{k}$ (where $\bar{k}$ is defined as in Proposition 1), the delegation game has a subgame-perfect equilibrium with partial information acquisition.*

If P chooses delegation set $D = [-\bar{\omega}, \bar{\omega}] \setminus (-\mu_{a^*}, \mu_{a^*})$, where $\{-\mu_{a^*}, \mu_{a^*}\}$ are the choices made under information design that are closest to the prior mean, then E acquires a more informative extremes-refining binary signal than he would under information design.

The principal's payoff as well as the parties' joint surplus is greater under delegation than information design.

The key step in establishing Proposition 2 is showing that in response to a symmetric delegation set, the expert picks an extremes-refining binary experiment with no garbling. This result and the optimal value of the threshold are given in Lemma 4.

Lemma 4. *Under the conditions of Proposition 2 ($\lambda < 1$ and $k \leq \bar{k}$), if the principal chooses delegation set $D = \mathbb{R} \setminus (-d, d)$ with $d > 0$, then the expert's best reply is an extremes-refining, non-garbled binary signal, π_z, μ_z . Threshold z in the optimal experiment satisfies $z \geq d$ and*

$$(1 - \sqrt{\lambda})z + \sqrt{\lambda}\mu_z = d \quad (17)$$

unless $d > (1 - \sqrt{\lambda})\bar{\omega} + \sqrt{\lambda}\mu_+$, in which case $z = \bar{\omega}$ and $\mu_z = \mu_+$.

The proof of the lemma is detailed in the Appendix. Here we give an overview of the steps involved.

First, we show that the set of states in $[d, \bar{\omega}]$ that the expert decides to learn perfectly is convex; the same is true for the states he learns in $[-\bar{\omega}, -d]$ as well. This follows by Jensen's inequality, because E 's payoff is convex in the posterior mean when he has no restriction on matching the action to the state. Then we show that the posterior mean distribution that the expert's optimal experiment generates cannot have masspoints outside $[-d, d]$, and it cannot reveal the state inside $(-d, d)$. These claims are proved via perturbations involving splitting and pooling posterior means.

Then, we show that $[-d, 0]$, and separately $[0, d]$, cannot contain bi-pools; moreover, that there cannot be a bi-pool involving one posterior mean from each of these intervals.

These steps establish that the optimal experiment in response to $D = \mathbb{R} \setminus (-d, d)$ has one negative and one positive threshold beyond which the state is fully revealed, whereas between the two thresholds there are exactly two intervals of pooled states. The final step, which relies on assumption (10), is to show that the pooling intervals are symmetric about zero. Hence the experiment is an extremes-refining binary signal with no garbling.

The expert's best-reply experiment for exclusion interval $(-d, d)$, which is an extremes-refining, non-garbled binary policy with threshold z , is illustrated in Figure 4 below.

In Figure 4, on the horizontal axis we measure the state, against which we plot payoff $(1 - \lambda)\omega^2 - (d - \omega)^2 \mathbf{1}_{\omega \in [-d, d]}$ whose expectation the expert maximizes. The second term in this expression is the loss from mismatch in the exclusion region, where $y_D(\omega) = \pm d$ with the sign of d matching that of ω . Note that for $\omega > 0$, this function changes from being strictly concave to being strictly convex at $\omega = d$. By symmetry, for $\omega < 0$, the graph is the mirror image of that for $\omega > 0$.

in Figure 4, an improvement can be generated by revealing states near (just below) z .




Figure 5: Improving the expert’s experiment when z is too low

Note that the above description of adjusting z to resolve E 's tradeoff relies on the fact that his best reply to a symmetric restriction is a *symmetric* extremes-refining binary policy. In order to establish this fact, in the proof, we invoke assumption (10), the regularity condition on the prior. It is easy to construct examples violating this condition in which the most informative policy for E asymmetric.

Now we are ready to prove the main claim of Proposition 2: if $\lambda < 1$ and (partial) information acquisition can be induced using information design, then the same is true for optimal (restricted) delegation; in fact, by choosing the delegation set $D = \mathbb{R} \setminus (-d, d)$ appropriately, the principal can make the expert acquire a more precise signal (in the Blackwell order) than she would have been able to do using information design, and so P strictly prefers restricted delegation to information design.

The idea of the proof is illustrated in Figure 6 and explained below it.




Figure 6: The idea of the proof of Proposition 2

Start with the best, partially informative equilibrium of the information design game: P asks E to carry out an extremes-refining non-garbled binary experiment with threshold a^* such that a^* is the smallest positive solution to $Q_{\mathcal{I}}(a) = k$. This experiment is illustrated in the upper half of Figure 6. Now consider a delegation set that excludes the interval $(-d, d) = (-\mu_{a^*}, \mu_{a^*})$, the set of actions that are closer to zero than any action chosen in equilibrium in the information design game. This is shown in the lower half of Figure 6.

Apply Lemma 4 to this disclosure set, that is, $d = \mu_{a^*}$. The expert's best reply is an extremes-refining, non-garbled binary experiment with threshold z , such that z satisfies equation (17) where $d = \mu_{a^*}$. This yields

$$(1 - \sqrt{\lambda})z + \sqrt{\lambda}\mu_z = \mu_{a^*}.$$

The equation says that μ_{a^*} is a convex combination of z and $\mu_z < z$, therefore we have $\mu_z < \mu_{a^*} < z < a^*$. That is, the expert learns more than he would have in the equilibrium under principal-led information design.

To see that the principal benefits from delegating this way compared to principal-led information design, first, recall that P 's objective is $W \equiv 2U + (2 - \lambda)V - k$. As we have just argued, the experiment chosen by the expert under delegation is more informative than the one imposed on him in principal-led information design, hence $V(\pi_{\mathcal{D}}) > V(\pi_{\mathcal{I}})$. By revealed preference, E 's overall expected payoff (ignoring transfers) is greater as well, because he could have picked the same experiment and actions as under information design. This implies $U(y_{\mathcal{D}}, \pi_{\mathcal{D}}) + (1 - \lambda)V(\pi_{\mathcal{D}}) \geq U(y_{\mathcal{I}}, \pi_{\mathcal{I}}) + (1 - \lambda)V(\pi_{\mathcal{I}})$. A positive weighted sum of the two inequalities yields $W_{\mathcal{D}} > W_{\mathcal{I}}$.

This concludes the proof of the first part of Theorem 1: optimal delegation outperforms information design for the principal when $\lambda < 1$.

5.2 Optimal delegation when $\lambda > 1$

We turn to the analysis of the delegation game when $\lambda > 1$. In Proposition 3, we fully characterize the unique subgame-perfect equilibrium of the delegation game with $\lambda > 1$. The expert learns at most a binary, potentially garbled signal. For $\lambda \in (1, 3/2)$ the optimal delegation set exhibits a symmetric restriction; if $\lambda > 3/2$ then the principal prefers not to delegate at all.

Proposition 3. *If $1 < \lambda < 3/2$ then for all $k \leq (3 - 2\lambda)\lambda\mu_+^2$, the delegation game has a unique subgame-perfect equilibrium with partial information acquisition. The principal chooses*

delegation set $D = \mathbb{R} \setminus (-d^*, d^*)$ where $d^* = \lambda\mu_+$; the expert learns the sign of the state and picks $y_D^* = \pm\lambda\mu_+$ accordingly. Equivalently, P can pick the minimal delegation set, $\{-d^*, d^*\}$.

If $\lambda > 3/2$ then the delegation game has no equilibrium with information acquisition; the principal chooses $y = 0$ for zero social or individual gain.

The proof of Proposition 3 is established by Lemmas 5, 6, and 7.

Lemma 5. *If $\lambda > 1$ then, in response to any delegation set, the expert chooses either to acquire no information at all, or a binary, potentially garbled signal.*

This is the key step in deriving the solution to optimal delegation when $\lambda > 1$. The result that *any* delegation set induces the acquisition of (at most) a *binary* signal is interesting on its own right and is illustrated in Figure 7. The proof, relying on improvements by garbling and applications of Jensen's inequality, is found in the Appendix.

In Figure 7, we show how the expert's best-reply experiment is determined for a delegation set of the form $D = [d'', d'] \cup [d, \bar{\omega}]$ with $d'' < d' < 0 < d$ and $d > |\bar{\omega}|$. The integrand in the expert's objective is drawn with a solid line that changes between black and grey depending on whether the state is in the delegation set. We also plot $(1 - \lambda)\omega^2$ for reference as a grey, dashed line.

Suppose that by using this delegation set, the principal induces information acquisition by the expert; that is, the signal distribution he chooses contains at least two different posterior means. Indeed, E 's optimal experiment generates posteriors $\mu_L < 0 < \mu_H$, both inside interval (d', d) . This is so because the integrand in his payoff is bounded from above by $(1 - \lambda)\omega^2$, which is decreasing in the state's distance from zero. The tangents to the payoff function at μ_L and μ_H must coincide (the common tangent is drawn in brown); otherwise E could use a garbling (or un-garbling) similar to that shown in Figure 6, in order to improve his expected payoff.




Figure 7: Illustration of Lemma 5

A delegation set that induces information acquisition must exclude an open interval of actions near the prior mean, just like the gap (d', d) in Figure 7. When $\lambda > 1$, the expert is motivated to learn solely because of the mismatch loss that he would incur by ignorantly taking one of the nearest-to-zero actions. In contrast, the principal benefits from allowing as many actions as possible (as long as the expert exerts effort to acquire a signal at all), therefore it is without loss of generality to assume that the optimal delegation set is of the form $\mathbb{R} \setminus (d', d)$ where $d' < 0 < d$. Lemma 6 claims that the optimal restriction for P is symmetric, and it also characterizes E 's best response.

Lemma 6. *If $\lambda > 1$, then it is optimal for the principal to use a symmetric restriction, $d' = -d$. If P induces E to acquire some information by excluding actions $(-d, d)$, then E 's best response is a policy that generates, with 50-50% chance,*

$$\mu = \pm \min \left\{ \frac{d}{\lambda}, \mu_+ \right\}. \quad (18)$$

If $d < \lambda\mu_+$ then the expert obtains a garbled signal.

It is easy to show that the first-best binary experiment is symmetric about zero. The intuition is that the first-best problem is a quantization problem (see e.g. Lloyd, 1982, and Kieffer, 1983). Given that the prior satisfies condition (10), the quantization problem has a unique solution which is symmetric. Lemma 6 establishes that the second-best has a similar, symmetric structure.

To prove Lemma 6, we start from an arbitrary, potentially asymmetric delegation set and derive the experiment that the expert chooses as a best reply to it. The two statistics of the gap in the allowed set that matter for payoffs are the size of the gap and the displacement of the gap relative to zero. Joint expected payoff is increasing in the size of the gap whenever such an increase induces a more informative experiment. In contrast, the displacement of the gap relative to zero reduces the informativeness of the experiment — consistently with the intuition from the quantization exercise. Thus, the objective is highest if a gap of a given size is placed symmetrically around zero.

Next we consider the impact of the displacement of the gap around zero on the constraints that the principal faces. Experiments need to satisfy the law of iterated expectations — they need to be Bayes plausible in modern parlance — and the expert needs to prefer the binary experiment over no information acquisition at all. We show that a symmetric gap allows the principal to implement any binary experiment that she wishes, but the same is not true for asymmetrically placed gaps.

In conclusion, the optimal delegation set is to allow all choices except those in a symmetric, open interval around zero. The optimal minimal delegation set, which obtains after elimination of redundant choices, is a symmetric binary set. Moreover, as Lemma 7 claims, P chooses the size of the gap so as to induce E to learn exactly whether the state is positive or negative:

Lemma 7. *If $\lambda > 1$ and the principal induces the expert to acquire some information via delegation, then she excludes actions $(-d^*, d^*)$ such that $d^* = \lambda\mu_+$. In response, the expert learns the sign of the state and carries out $y = \pm\lambda\mu_+$ accordingly. The optimal minimal delegation set consists of just two points, $\pm d^* = \pm\lambda\mu_+$.*




Figure 8: Optimal (minimal) delegation when $\lambda > 1$

Figure 8 illustrates Lemma 7, when $\lambda > 1$ and the principal induces information acquisition via optimal minimal delegation. Actions beyond $\pm d^*$ could be included in the delegation set without changing the outcome.

The principal's expected payoff gain (relative to $-\sigma^2$) when she uses such a restriction and hence the expert learns the sign of the state is

$$\begin{aligned} W_{\mathcal{D}} &= -2(d^* - \mu_+)^2 + (2 - \lambda)\mu_+^2 - k \\ &= (3 - 2\lambda)\lambda\mu_+^2 - k. \end{aligned} \tag{19}$$

If $\lambda > 3/2$ then $W_{\mathcal{D}} < 0$ for all $k \geq 0$, therefore in that case the principal will forego delegation and set $y = 0$.

This establishes all claims of Proposition 3.

In order to prove the second part of Theorem 1 (that the principal prefers information design over optimal delegation when $\lambda > 1$), we need to demonstrate that whenever it is possible to implement information acquisition under delegation, it is also possible to do so under information design. Then, we must verify that the principal's payoff is indeed greater in the latter case.

By Proposition 3, optimal delegation induces information acquisition (specifically, the expert learning the sign of the state) whenever $k \leq (3 - 2\lambda)\lambda\mu_+^2$. In contrast, by Proposition 1, when $\lambda > 1$, information acquisition is implementable via information design for $k \leq (2 - \lambda)\mu_+^2$. It is easy to see that the former condition implies the latter.

As far as the principal's expected payoff gain is concerned, under principal-led information design it is $W_{\mathcal{I}}^* = (2 - \lambda)V^* - k$, where $(1 - \lambda)V^* + \mu_{a^*}^2 = k$. The latter equation is E 's binding incentive constraint with a^* denoting the threshold in the (extremes-refining, non-garbled binary) experiment picked by P for E . But $V^* \geq \mu_+^2$, therefore

$$W_{\mathcal{I}}^* \geq (2 - \lambda)\mu_+^2 - k > (3 - 2\lambda)\lambda\mu_+^2 - k \equiv W_{\mathcal{D}}^*.$$

This completes the proof of the second half of Theorem 1.

6 Related literature and directions for future research

Our main contribution is to the comparison of communication and delegation in the context of delegated expertise. We contrast two natural incentive schemes: principal-led information design followed by cheap talk, and optimal delegation followed by flexible information acquisition by the expert. Our paper connects with several strands of the

economic theory literature, including cheap talk, delegation with endogenous information acquisition, and more general principal-expert models that incorporate information design. Some connections are in terms of shared conceptual questions asked, economic insights gained, while others are methodological, drawing on related technical tools.

In the classical cheap talk literature pioneered by Crawford and Sobel (1982), the comparison of communication and delegation was initiated by Dessein (2002). In his model featuring an exogenously informed sender, simple (unconstrained) delegation performs better whenever meaningful communication is possible at all.⁵

When information acquisition is endogenous, the comparison changes.⁶ Demski and Sappington (1987) introduce the notion of delegated expertise and propose to evaluate arrangements accounting for how well they incentivize the acquisition of information. Allowing for costs of information acquisition, Argenziano, Severinov and Squintani (2016) show that communication can outperform delegation. Although the sender's bias distorts information acquisition under both regimes, the distortion is less harmful under communication: equilibrium cheap talk induces the expert to acquire information whose use by the principal is more valuable to her, resulting in more precise equilibrium decisions.

Deimen and Szalay (2019) demonstrate that communication can outperform delegation when the state is multidimensional and the sender has discretion over which aspects to investigate. When decision-making is delegated to the expert, the expert acquires information that mostly matters to himself. In contrast, under communication advice based on such information would be discounted by the decision-maker resulting in limited information transmission. To avoid such poor communication outcomes, the expert pays more attention to things that matter to the decision-maker.⁷

Holmström (1977, 1980) provided one of the earliest formal treatments of optimal delegation, highlighting the fundamental trade-off between loss of control and the value of the agent's private information. Alonso and Matouschek (2008) and Amador and Bagwell (2013) characterize the optimal delegation set when the principal cannot use message-contingent transfers and must instead commit to a menu of actions from which the (biased but informed) agent chooses. Szalay (2005) introduces moral hazard in information acquisition into optimal delegation. He shows that — even with aligned ex-post preferences — the principal can strengthen incentives for effort by restricting the agent's choice

⁵Subsequent work, e.g. Goltsman, Hörner, Pavlov, and Squintani (2009), compared various other institutions such as mediation, arbitration, negotiation as well in classical cheap talk with exogenous information.

⁶Indeed, some of the predictions of the Crawford–Sobel cheap talk model are changed in important ways when the sender's information is endogenous and costly; see Pei (2015).

⁷For a recent analysis of organizational design to foster delegated learning, see Benkert, Matyskova, and Starkov (2026).

set, in particular by excluding actions close to the status quo. The present paper expands this line of research to an environment with continuous state and action spaces and fully flexible (but costly) information structures. Remarkably, we find that when the marginal cost of refining information is high, the expert acquires at most binary (potentially garbled) signals, and the optimal minimal delegation set is binary. This insight is new to the delegation literature.

The paper is partly inspired by the recent literature on Bayesian persuasion, initiated by Kamenica and Gentzkow (2011) and Rayo and Segal (2010). We extend the domain of applications of information design to organizational design and delegated expertise. An early contribution on information design in cheap talk settings is Ivanov (2010), who characterizes optimal information structures and compares them to full delegation with perfect information. Our setting is novel in its combination of moral hazard in flexible information acquisition with perfectly aligned ex-post preferences. Features of the solution, such as optimal extremes-refining binary signals under principal-led design, are new.

Kolotilin (2018), Dworczak and Martini (2019), and Dworczak and Kolotilin (2024) develop duality and linear programming techniques for persuasion problems. Our delegated information design problem is related but distinct, with the main challenge lying in the joint characterization of the optimal delegation set and induced information structure. For given delegation sets and induced value functions, our results are consistent with their approaches.

Kleiner, Moldovanu, and Strack (2021) and Arieli, Babichenko, Smorodinsky, and Yamashita (2023) provide solution techniques and characterization results for persuasion problems. We build on their results when deriving the solution to our principal-led information design problem.⁸ The former paper demonstrates that the solution methods for finding optimal information structures apply not only to persuasion problems but also to delegation and other problems.

Kolotilin and Zapechelnyuk (2025) provide general conditions under which delegation and persuasion problems are mathematically equivalent. Consistently with their results, we show that in our environment, delegation outperforms information design when the expert’s marginal cost of obtaining granular information is low and vice versa. Kolotilin, Corrao, and Wolitzky (2025) study a class of one-dimensional Bayesian persuasion problems where the receiver’s utility is single-peaked, whereas the sender’s utility is monotonic. Their approach subsumes a large swath of models in the literature as special cases,

⁸Relatedly, in Kreutzkamp and Lou (2025), the sender (with misaligned preferences) can commit to a costless experiment but not truthful reporting; bi-pooling policies are optimal in that setting as well.

but it does not include our environment.

Lipnowski and Ravid (2020) use persuasion methods to study cheap talk with transparent (that is, state-independent) motives. More closely related to our study, with the distinctive features of costly information and otherwise aligned preferences, Lipnowski, Mathevet, and Wei (2020) examine a persuasion version of a rational inattention model, where restricting the set of information policies provides incentives for information acquisition.⁹ Our key innovations, motivated by our application to organizational economics, are the analysis of the optimal restriction of actions inducing experiments (not restricting experiments directly), and the comparison of institutions. While the incentive problems are different, the common elements are moral hazard in information acquisition, conflicts that stem from the cost of experimentation, and that the solutions relate to Szalay (2005), despite differences in assumptions, notably commitment to and observability of information policies.

We compare two mechanisms that are of relevance in the context of organizational economics, where either the information or the actions available to the agent can be chosen by the principal; however, in the latter case the choice of information has to be left to the agent. Clearly, in other contexts — such as constraining a natural monopoly — a regulator may have joint control over action limitations and the information that may be used in action choice. For analyses of such problems see Vairo (2025) and Lou (2026).

Throughout the paper we have assumed a completely flexible information acquisition technology, and, for the analysis of principal-led information design, that the principal can choose any experiment for the expert. In the remainder of this section we report results on several variations on the theme; details are available from the authors upon request.

Suppose that information acquisition is restricted so that only all-or-nothing experiments are feasible: the expert either learns the state at a given cost or remains uninformed at no cost. Then, it follows from Theorem 1 that optimal delegation beats cheap talk; it can also be shown that optimal delegation takes the form as in Szalay (2005). The cheap talk problem in this environment is studied by Eső and Szalay (2020). Communication must take a partitional form; under certain conditions, less nuanced conversations provide stronger incentives.¹⁰

An interesting extension would be to characterize optimal delegation for the case of flexible information acquisition and $\lambda < 1$. We did not need to solve this problem fully for

⁹In recent work, Bilotta, Carnehl, and Preusser (2026) analyse a model of delegated information acquisition, without restrictions on actions, where conflict arises from ex-post preferences over state-action pairs.

¹⁰The reason for fewer messages used in communication could be an exogenous or endogenous restriction, or due to a cost of complexity, as in Sobel (2012).

our purposes (i.e., to show in Theorem 1 that in this case, optimal delegation outperforms competing incentive schemes), but clearly, this is an interesting avenue for future research.

We have placed restrictions on the benefits and costs of information, largely to foster clarity of insight. Quadratic mismatch losses are commonly assumed in the literature on delegation; likewise, much of the literature on information design assumes objectives that depend on the posterior means for tractability. Extending our results to different specifications would be possible. The main tradeoffs would remain, but it goes without saying that the precise dividing point between the two regimes (where principal-led information design outperforms optimal delegation and vice versa) depends very clearly on the exact specification of the model.

7 Appendix: detailed proofs of lemmas

7.1 Proof of Lemma 1

The construction of the experiment that minimizes $V = \int \mu^2 d\pi(\mu)$ for $\pi \in \Pi_\gamma$ is given in the text. As $|\mu| \geq \gamma$ for all $\mu \in \text{supp}(\pi)$, we have $V \geq \gamma^2$. The binary policy that generates $\mu = \pm\gamma$ with equal probabilities, $\pi_\gamma^{\min}$, achieves this lower bound.

Next, we derive the policy that maximizes $V = \int \mu^2 d\pi(\mu)$ on Π_γ for any given $\gamma \in (0, \mu_+]$. The first step of the derivation is to show that the optimum is reached by a bi-pooling policy (cf. Kleiner, Moldovanu, and Strack, 2021, and Arieli, Babichenko, Smorodinsky, and Yamashita, 2023).

An experiment is a bi-pooling policy if there exists a countable collection of pairwise disjoint intervals $(B_n)_n$, each occurring with a positive prior probability, such that the posterior beliefs induced on each such interval have support of size at most two, whereas outside these intervals the state is fully revealed. If two distinct posteriors are induced at states in B_n then the interval is called a bi-pool, if one, then it is called a simple pool.

One such bi-pooling interval is illustrated in Figure 9 below. Associated with every state in interval B , there are two signal realizations in the support of π , denoted by $\mu_{B,1}$ and $\mu_{B,2} \geq \mu_{B,1}$. Both posteriors must lie in B , i.e., $\mu_{B,1}, \mu_{B,2} \in B$; the corresponding probabilities are denoted by $p_{B,1}$ and $p_{B,2}$ and sum up to $\Pr[\tilde{\omega} \in B]$. In the figure, the orange area represents the mass of states assigned to $\mu_{B,1}$ and the yellow area represents the mass of states assigned to $\mu_{B,2}$. Note that by construction, $p_{B,i} > 0$ for each i , as otherwise, $\mu_{B,i}$ would not be in the support of π (a support set is by definition the smallest compact set that has probability 1). Note that if $\mu_{B,2} = \mu_{B,1}$ then B is a simple pooling

interval. The union of two adjacent simple pooling intervals is not a bi-pool either.




Figure 9: Example of a bi-pool

Note that any realization μ in the support of a bi-pooling policy is attained with strictly positive probability if and only if it belongs to a bi-pool.

The derivation of the policy that maximizes V on Π_γ proceeds in five steps.

First, we show the maximum exists and is obtained by a bi-pooling policy. Second, we show that it is in fact obtained by a generalized (possibly asymmetric and garbled) extremes-refining binary experiment. In the third and fourth steps, we show that garbling is suboptimal. Finally, we show that the policy maximizing V on Π_γ is symmetric.

Step 1: There exists an experiment in Π_γ maximizing $V = \int \mu^2 d\pi(\mu)$ and it is a bi-pooling policy.

Proof of Step 1: The set of all Blackwell experiments, denoted by Π , is compact (cf. Strasser, 1985). The set $\Pi_\gamma = \{\pi \in \Pi \mid \text{supp}(\pi) \cap (-\gamma, \gamma) = \emptyset\} \subseteq \Pi$ is a closed subset of the compact set Π , so Π_γ is compact. The objective function $V = \int \mu^2 d\pi(\mu)$ is linear, hence continuous in π . By the extreme value theorem, a maximum exists. By Bauer's maximum principle, the maximum is attained at an extreme point of Π_γ .

The set of extreme points of Π_γ is a subset of the set of extreme points of Π . To see this, assume towards contradiction that there is an extreme point of Π_γ , denoted by π , that is not an extreme point of Π . Hence, there exist some $\pi', \pi'' \in \Pi$ such that $\pi = \alpha\pi' + (1 - \alpha)\pi''$ for some $\alpha \in (0, 1)$ and $\{\pi', \pi''\} \not\subseteq \Pi_\gamma$, that is, $\text{supp}(\pi') \not\subseteq \mathbb{R} \setminus (-\gamma, \gamma)$ or $\text{supp}(\pi'') \not\subseteq \mathbb{R} \setminus (-\gamma, \gamma)$. Since π is a convex combination of π' and π'' , the support of π is the union of the support of π' and π'' . This implies that $\text{supp}(\pi) = \text{supp}(\pi') \cup \text{supp}(\pi'') \not\subseteq \mathbb{R} \cap (-\gamma, \gamma)$, that is, $\pi \notin \Pi_\gamma$, which is a contradiction.

The set of extreme points of Π is the set of bi-pooling policies (cf. Kleiner, Moldovanu, and Strack, 2021) and Arieli, Babichenko, Smorodinsky, and Yamashita, 2023). Hence, every extreme point in Π_γ is a bi-pooling policy. We conclude that the maximum of V on Π_γ exists and is obtained by a bi-pooling policy. $\square$

Next we show that a bi-pooling policy that maximizes V on Π_γ belongs to the class of *generalized* (i.e., potentially asymmetric) extremes-refining binary experiments, charac-

terized by thresholds $a_- < 0 < a_+$, such that outside the bi-pooling interval $[a_-, a_+]$ the state is revealed, whereas inside $[a_-, a_+]$, the induced posterior means are $\mu_-^{\max} \equiv \max\{\mu \in \pi | \mu \leq 0\}$ and $\mu_+^{\min} \equiv \min\{\mu \in \pi | \mu \geq 0\}$.

Step 2: A bi-pooling policy that maximizes V on Π_γ belongs to the class of generalized extremes-refining binary experiments.

Proof of Step 2: Let $\mu_+^{\min} \equiv \min\{\mu \in \pi | \mu \geq 0\}$ be the smallest positive signal realization of π and $\mu_-^{\max} \equiv \max\{\mu \in \pi | \mu \leq 0\}$ be its largest negative signal realization.

First, we show that at least one of $\mu_+^{\min}$ and $\mu_-^{\max}$ must realize with strictly positive probability under π : Since π is a bi-pooling policy, if $\mu_+^{\min}$ and $\mu_-^{\max}$ realized with zero probability, this would imply that the states $\mu_+^{\min}$ and $\mu_-^{\max}$ were fully revealed and that all states in the interval in between, namely in the interval $(\mu_-^{\max}, \mu_+^{\min})$ would have to be assigned to some signal value in the interval $(\mu_-^{\max}, \mu_+^{\min})$. Since $\mu_-^{\max} \leq -\gamma$ and $\mu_+^{\min} \geq \gamma$ as $\pi \in \Pi(\gamma)$, the interval $(\mu_-^{\max}, \mu_+^{\min})$ is non-empty, so there is a positive mass of states in the interval, hence there would have to be a positive mass of signals in the interval $(\mu_-^{\max}, \mu_+^{\min})$, contradicting the definition of $\mu_+^{\min}$ and $\mu_-^{\max}$.

Hence, by definition of a bi-pooling policy, at least one of $\mu_+^{\min}$ and $\mu_-^{\max}$ (any that realizes with strictly positive probability under π) must be part of a bi-pool.

Second, we show that $\mu_+^{\min}$ cannot form a bi-pool with another positive signal realization, and analogously, $\mu_-^{\max}$ cannot form a bi-pool with another negative signal realization. This implies that if they are part of a bi-pool, they can only form a bi-pool with each other.

Suppose $\mu_+^{\min}$ forms a bi-pool with another positive signal realization in the support of π , denoted by μ' . Note that $\mu' \geq \mu_+^{\min}$. Let $p_+^{\min}$ and p' denote the probabilities of these two signal realizations under π . Construct a new experiment π' by separating the mass of states assigned to μ' into two submasses whose posterior means are μ'' and $\mu''' \geq \mu''$. Their corresponding probabilities p'' and p''' sum up to p' . Since $\mu' = (p''\mu'' + p'''\mu''')/p' \geq \mu_+^{\min} \geq \gamma$, it is possible to choose μ'' and μ''' in a way such that $\mu''' \geq \mu'' \geq \gamma$. Note that π' is a strict mean-preserving spread of π iff $\mu''' > \mu''$. In fact, it is possible to choose μ'' and μ''' in a way such that $\mu''' > \mu'' \geq \gamma$ as long as $\mu' > \gamma$. Hence, a bi-pool between $\mu_+^{\min}$ and μ' can only be optimal if $\mu' = \gamma$, implying that $\mu' = \mu_+^{\min} = \gamma$. Hence, if $\mu_+^{\min}$ formed a bi-pool with the next higher signal realization in the support of π , it was essentially a pooling interval.

Analogously, one can show that $\mu_-^{\max}$ cannot form a bi-pool with the next smaller signal realization in the support of π .

We conclude that there are two possible mutually exclusive cases:

Case 1: Both $\mu_+^{\min}$ and $\mu_-^{\max}$ are realized with strictly positive probability. In this case, they either form a strict bi-pool (bi-pool with some garbling), or two separate pooling intervals.

Case 2: Only one of $\mu_+^{\min}$ and $\mu_-^{\max}$ is realized with strictly positive probability, the other realizes with probability zero. In this case, the signal value that realizes with strictly positive probability forms a pooling interval, and the state corresponding to the other signal value is fully revealed.

Note that both the bi-pool or two pooling intervals in case 1 as well as the pooling interval in case 2 are a superset of $(\mu_-^{\max}, \mu_+^{\min})$, i.e., they are intervals around zero.

Third, we show that outside the bi-pool or pooling interval around zero, the state is fully revealed: Suppose π was not fully revealing outside of the bi-pool or pooling interval around zero. Any such experiment is dominated by one where the state is fully revealed beyond the bi-pool or pooling interval around zero as it is a strict mean-preserving spread.

This completes the proof that there exists a maximally informative experiment in Π_γ that is a generalized extremes-refining binary experiment. $\square$

Step 3: A generalized extremes-refining binary policy that maximizes V on Π_γ satisfies $\mu_+^{\min} = \gamma$ and $\mu_-^{\max} = -\gamma$.

Proof of Step 3: Suppose not, i.e., $\mu_+^{\min} > \gamma$ or $\mu_-^{\max} < -\gamma$. Assume $\mu_+^{\min} > \gamma$.

If $\mu_+^{\min}$ and $\mu_-^{\max}$ form a bi-pool (case 1), then we can construct a strict mean-preserving spread by fully revealing the highest states assigned to $\mu_+^{\min}$ such that the posterior mean of the remaining states assigned to $\mu_+^{\min}$ becomes $\mu_+^{\min'} \geq \gamma$.

If $\mu_-^{\max}$ is fully revealed and $\mu_+^{\min}$ forms a pooling interval (case 2), then we can construct a strict mean-preserving spread by fully revealing the highest states of the pooling interval such that the posterior mean of the remaining states in the pooling interval becomes $\mu_+^{\min'} \geq \gamma$. The same is true if $\mu_+^{\min}$ is fully revealed and $\mu_-^{\max}$ forms a pooling interval. Since the highest states of the pooling interval are removed, the posterior mean of the remaining states in the pooling interval becomes $\mu_-^{\max'} < \mu_-^{\max} \leq \gamma$.

The proof for $\mu_-^{\max} < -\gamma$ works analogously.

This completes the proof that $\mu_+^{\min} = \gamma$ and $\mu_-^{\max} = -\gamma$. $\square$

By steps 1–3, there exists a policy maximizing V on Π_γ that is a generalized extremes-refining experiment with posterior means $-\gamma$ and γ induced in the bi-pooling interval $[a_-, a_+]$, which either form a bi-pool or two separate pooling intervals (case 1), or one of them forms a pooling interval and the other one is fully revealed (case 2).

Next, we show that if we have a maximally informative policy of the kind in case 1, then we can rule out a bi-pool with strict garbling.

Step 4: A generalized extremes-refining binary experiment that maximizes V on Π_γ and under which both $-\gamma$ and γ are realized with strictly positive probability is such that both $-\gamma$ and γ form separate pooling intervals.

Proof of Step 4: Let $p_-^{\max}$ be the probability of $\mu_-^{\max}$ and $p_+^{\min}$ be the probability of $\mu_+^{\min}$. Suppose $\mu_-^{\max}$ and $\mu_+^{\min}$ form a bi-pool. Construct a new experiment π' by separating the bi-pool into two separate pools with the same probabilities $p_-^{\max}$ and $p_+^{\min}$.

The corresponding posterior means are $\mu_-^{\max'} < \mu_-^{\max}$ and $\mu_+^{\min'} > \mu_+^{\min}$, therefore $\int \mu^2 d\pi'(\mu) > \int \mu^2 d\pi(\mu)$. $\square$

By step 4, the cases 1 and 2 can be summarized as follows: The bi-pooling interval $[a_-, a_+]$ consists of two pooling intervals, one assigned to $\mu_-^{\max} = -\gamma$ and one assigned to $\mu_+^{\min} = \gamma$. At least one of the intervals has strictly positive mass, i.e., $p_-^{\max} + p_+^{\min} > 0$. Moreover, the state is fully revealed outside $[a_-, a_+]$.

Step 5: A generalized extremes-refining binary policy that maximizes V on Π_γ is symmetric.

Proof of Step 5: Let $[a, c]$ be the set of states assigned to $-\gamma$ and $[c, b]$ be the set of states assigned to γ . (That is, we denote $a_- = a$ and $a_+ = b$.)

We want to show that the maximally informative experiment is symmetric, that is, it satisfies $c = 0$ and $a = -b$.

If $p_-^{\max} = 0$, then $a = -\gamma = c$. If $p_+^{\min} = 0$, then $c = \gamma = b$. If $p_-^{\max} = p_+^{\min}$, then $c = 0$ and $a = -b$ by symmetry of F . Moreover, $a \leq -\gamma < 0 < \gamma \leq b$ and $a \leq -\gamma \leq c \leq \gamma \leq b$.

Think of a and b as functions of c such that $\int_{a(c)}^c \omega dF(\omega) = -\gamma$ and $b = b(c)$ such that $\int_c^{b(c)} \omega dF(\omega) = \gamma$. Differentiating each with respect to c yields

$$cf(c) - a(c)f(a)a'(c) = 0 \quad \Leftrightarrow \quad f(a)a'(c) = \frac{cf(c)}{a(c)} \quad (20)$$

and

$$b(c)f(b)b'(c) - cf(c) = 0 \quad \Leftrightarrow \quad f(b)b'(c) = \frac{cf(c)}{b(c)} \quad (21)$$

Note that

$$V(c) = \int \mu^2 d\pi(\mu) = \int_{\bar{\omega}}^{a(c)} \omega^2 dF(\omega) + \left(\int_{a(c)}^c dF(\omega) + \int_c^{b(c)} dF(\omega) \right) \gamma^2 + \int_{b(c)}^{\bar{\omega}} \omega^2 dF(\omega).$$

We want to show that this expression is maximized at $c = 0$.

Differentiating $V(c)$ yields

$$V'(c) = a(c)^2 f(a) a'(c) + [-f(a) a'(c) + f(b) b'(c)] \gamma^2 - b(c)^2 f(b) b'(c).$$

Plugging in (20) and (21), and rearranging we get

$$\begin{aligned} V'(c) &= a(c)^2 \cdot \frac{cf(c)}{a(c)} + \left[-\frac{cf(c)}{a(c)} + \frac{cf(c)}{b(c)} \right] \gamma^2 - b(c)^2 \cdot \frac{cf(c)}{b(c)} \\ &= \frac{cf(c)}{ab} \cdot (a - b)(ab + \gamma^2), \end{aligned}$$

where on the second line we reverted to $a(c) = a$ and $b(c) = c$ for cleaner notation.

Since $a \leq -\gamma$ and $b \geq \gamma$ and at least one of these inequalities must hold with strict inequality by $p^{\max} p_+^{\min} > 0$, we have that $ab < -\gamma^2$, and therefore $ab + \gamma^2 < 0$ holds everywhere. Moreover, $a - b < 0$ everywhere. Therefore, $V'(c) > 0$ for $c < 0$ and $V'(c) < 0$ for $c > 0$, and so the unique global maximum of $V(c)$ is reached at $c = 0$. $\square$

7.2 Proof of Lemma 2

First, consider $\gamma < \mu_+$. In this case $V_{\mu_+}^{\min} = \mu_+^2 < V_{\mu_+}^{\max}$.

If $\lambda \leq 1$, then both the objective and the left-hand side of the constraint in (11) are decreasing in α . The objective increases and the constraint is relaxed by decreasing α . Therefore, a solution exists if and only if $\alpha = 0$ is feasible, i.e., $(1 - \lambda)V_{\gamma}^{\max} + \gamma^2 \geq k$. Moreover, if a solution exists, then the optimum is reached at $\alpha(\gamma) = 0$.

If $\lambda > 1$, then the objective in (11) is decreasing in α , but the constraint tightens as α goes up. A solution exists if and only if $\alpha = 1$ is feasible: $(1 - \lambda)V_{\gamma}^{\min} + \gamma^2 \geq k$. If this holds, then the optimum is reached at the lowest feasible α . If $(1 - \lambda)V_{\gamma}^{\max} + \gamma^2 \geq k$ then $\alpha(\gamma) = 0$; if $(1 - \lambda)V_{\gamma}^{\min} + \gamma^2 \geq k > (1 - \lambda)V_{\gamma}^{\max} + \gamma^2$ then $\alpha(\gamma)$ is the (unique) solution in α to $(1 - \lambda) [\alpha V_{\gamma}^{\min} + (1 - \alpha)V_{\gamma}^{\max}] + \gamma^2 = k$, that is,

$$\alpha(\gamma) = \frac{V_{\gamma}^{\max} - \frac{k - \gamma^2}{\lambda - 1}}{V_{\gamma}^{\max} - V_{\gamma}^{\min}} \in [0, 1].$$

Finally, if $\gamma = \mu_+$, then $V_{\mu_+}^{\min} = \mu_+^2 = V_{\mu_+}^{\max}$. The objective and the constraint in (11) are independent of α . The constraint becomes $(1 - \lambda)\mu_+^2 + \mu_+^2 \geq k$, and a solution exists iff $(1 - \lambda)V_{\mu_+}^{\min} + \mu_+^2 \geq k$. If this holds then any $\alpha(\mu_+) \in [0, 1]$ solves (11). $\square$

7.3 Proof of Lemma 3

If $\lambda \leq 1$, then we know from Lemma 2 that $\alpha(\gamma) = 0$ is a solution for all γ such that $(1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq k$. The objective, $(2 - \lambda)V_\gamma^{\max} - k$, is strictly decreasing in γ .

Therefore, when $\lambda \leq 1$, a solution to the optimization problem exists iff there exists $\gamma \in [0, \mu_+]$ such that $(1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq k$, and if this holds, then it is optimal for the principal to choose the smallest feasible γ in $[0, \mu_+]$, i.e., the smallest solution in γ to the equation $(1 - \lambda)V_\gamma^{\max} - k = -\gamma^2$. In the solution, $\alpha(\gamma) = 0$.

If $\lambda > 1$, then we know from Lemma 2 that there exists a solution in which $\alpha(\gamma) = 0$ for all γ such that $(1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq k$, and $\alpha(\gamma) = \left(V_\gamma^{\max} - \frac{k - \gamma^2}{\lambda - 1}\right) / \left(V_\gamma^{\max} - V_\gamma^{\min}\right)$ for all γ such that $(1 - \lambda)V_\gamma^{\min} + \gamma^2 \geq k > (1 - \lambda)V_\gamma^{\max} + \gamma^2$.

Therefore, when $\lambda > 1$, the objective function is

$$\begin{cases} (2 - \lambda)V_\gamma^{\max} - k & , \text{ if } (1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq k \\ (2 - \lambda)\frac{k - \gamma^2}{1 - \lambda} - k & , \text{ if } (1 - \lambda)V_\gamma^{\min} + \gamma^2 \geq k > (1 - \lambda)V_\gamma^{\max} + \gamma^2 \end{cases}.$$

The objective is strictly decreasing in γ for all γ such that $(1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq k$, and it is strictly increasing in γ for all γ such that $(1 - \lambda)V_\gamma^{\min} + \gamma^2 \geq k > (1 - \lambda)V_\gamma^{\max} + \gamma^2$.

The constraint set is non-empty iff there is $\gamma \in [0, \mu_+]$ such that $(1 - \lambda)V_\gamma^{\min} + \gamma^2 \geq k$. Since $V_\gamma^{\min} = \gamma^2$ and $\lambda \leq 2$, we find

$$\max_{\gamma \in [0, \mu_+]} (2 - \lambda)\gamma^2 = (2 - \lambda)\mu_+^2.$$

Hence, a solution exists iff $k \leq (2 - \lambda)\mu_+^2$.

If a solution exists, it is one of the following two: either

- (i) the smallest γ in $[0, \mu_+]$ such that $(1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq k$, i.e., the smallest solution in γ to the equation $(1 - \lambda)V_\gamma^{\max} + \gamma^2 = k$, denoted by γ^* , or
- (ii) the largest γ in $[0, \mu_+]$ such that $(1 - \lambda)V_\gamma^{\min} + \gamma^2 \geq k > (1 - \lambda)V_\gamma^{\max} + \gamma^2$, i.e., the largest solution in γ to the equation $(1 - \lambda) \left[\alpha(\gamma)V_\gamma^{\min} + (1 - \alpha(\gamma))V_\gamma^{\max} \right] + \gamma^2 = k$, denoted by γ' .

We will now show that if a solution exists, it is γ^* .

First, note that if a solution exists, that is, if $k \leq (2 - \lambda)\mu_+^2$, then γ^* is well-defined: it exists iff there is $\gamma \in [0, \mu_+]$ such that $(1 - \lambda)V_\gamma^{\max} + \gamma^2 = k$. Since

$$\max_{\gamma \in [0, \mu_+]} (1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq (1 - \lambda)V_{\mu_+}^{\max} + \mu_+^2 = (2 - \lambda)\mu_+^2,$$

$k \leq (2 - \lambda)\mu_+^2$ implies $k \leq \max_{\gamma \in [0, \mu_+]} (1 - \lambda)V_\gamma^{\max} + \gamma^2$. So existence of γ^* is implied by the existence of a solution.

Next, suppose γ' exists, too. We want to show that in that case the objective function attains a higher value at γ^* than at γ' .

First, consider the case of $\gamma^* \leq \gamma'$. The objective equals $(2 - \lambda)V_{\gamma_s}^{\max} - k$ at γ^* since $\alpha(\gamma^*) = 0$, and it is $(2 - \lambda) [\alpha(\gamma')V_{\gamma'}^{\min} + (1 - \alpha(\gamma'))V_{\gamma'}^{\max}] - k$ at γ' . If $\gamma^* \leq \gamma'$ then $\alpha(\gamma')V_{\gamma'}^{\min} + (1 - \alpha(\gamma'))V_{\gamma'}^{\max} \leq V_{\gamma'}^{\max} \leq V_{\gamma^*}^{\max}$, which implies that $(2 - \lambda)V_{\gamma^*}^{\max} - k \geq (2 - \lambda) [\alpha(\gamma')V_{\gamma'}^{\min} + (1 - \alpha(\gamma'))V_{\gamma'}^{\max}] - k$. Therefore, the objective function attains a higher value at γ^* than at γ' .

Second, consider the case of $\gamma^* > \gamma'$. Then, the objective function becomes

$$(2 - \lambda) \frac{k - \gamma^2}{1 - \lambda} - k, \text{ for each } \gamma \in \{\gamma^*, \gamma'\}.$$

If $\gamma^* > \gamma'$, then $(2 - \lambda) \frac{k - (\gamma^*)^2}{1 - \lambda} - k \geq (2 - \lambda) \frac{k - (\gamma')^2}{1 - \lambda} - k$. Therefore, the objective function attains a higher value at γ^* than at γ' .

This completes the proof that if a solution exists, it is the smallest solution in γ to the equation $(1 - \lambda)V_\gamma^{\max} + \gamma^2 = k$.

Finally, recall that if $\lambda > 1$, a solution exists if and only if $k \leq (2 - \lambda)\mu_+^2$. Note that

$$\max_{\gamma \in [0, \mu_+]} (1 - \lambda)V_\gamma^{\max} + \gamma^2 = (1 - \lambda)V_{\mu_+}^{\max} + \mu_+^2 = (2 - \lambda)\mu_+^2$$

as $(1 - \lambda)V_\gamma^{\max}$ is increasing in γ .

So to conclude, no matter whether $\lambda \leq 1$ or $\lambda > 1$, a solution to (15) exists if and only if there is $\gamma \in [0, \mu_+]$ such that $(1 - \lambda)V_\gamma^{\max} + \gamma^2 \geq k$, and it is obtained at γ^* that is the smallest solution in γ to $(1 - \lambda)V_\gamma^{\max} + \gamma^2 = k$. By Lemma 2, we have $\alpha(\gamma^*) = 0$. $\square$

7.4 Proof of Lemma 4

This lemma is proved via a series of claims.

Assume $\lambda < 1$. The expert chooses policy π (a mean-preserving contraction of the prior) to maximize

$$\int -[y_D(\mu) - \mu]^2 d\pi(\mu) + \int (1 - \lambda)\mu^2 d\pi(\mu),$$

where $y_D(\mu)$ is the closest point in $D \equiv \mathbb{R} \setminus (-d, d)$ to posterior mean μ . We have $y_D(\mu) =$

μ for $\mu \in D$ and $y_D(\mu) = d$ or $-d$ (matching the sign of μ) for $\mu \notin D$.

Throughout the proof we refer to $-[y_D(\mu) - \mu]^2 + (1 - \lambda)\mu^2$ as the integrand. For $\mu \in D$, the first term in the integrand is zero, and the function is convex by $\lambda < 1$.

For $\mu \notin D$ the integrand becomes $-(d - \mu)^2 + (1 - \lambda)\mu^2 = -\lambda\mu^2 + 2d\mu - d^2$, which is concave in μ by $\lambda \geq 0$.

Claim 1: The (sub)-set of states $\omega \in [d, \bar{\omega}]$ that is fully revealed is convex.

Proof: suppose not, i.e., the subset of states that is fully revealed in $[d, \bar{\omega}]$ has a gap. Since the integrand is convex in this region, revealing the states increases expected utility by Jensen's inequality.

Claim 2: if π is an optimal policy, then it cannot have masspoints in $\omega \in (d, \bar{\omega}]$.

Proof: Suppose π has a masspoint in $(d, \bar{\omega}]$, say μ' that is realized with probability p' . Since $\mu' > d$, the masspoint must have been generated by a signal that was sent for some states $\omega > d$. Create an improvement as follows: reveal the highest states perfectly until either all mass is used up or the remaining masspoint μ'' is equal to d and occurs with $p'' < p'$. By convexity of the integrand over this range and Jensen's inequality, expected payoff is increased.

Claim 3: if π is an optimal policy, then the set $(-d, d)$ cannot contain any subset within which states are fully revealed.

Proof: suppose contrary to the claim that there is such a set. If it is entirely in the positive domain, garble all states into a single mass point. By concavity of the integrand on $(-d, d)$, this increases expected utility.

If zero is an element of the considered set, then split the states into the positive ones and the negative ones, then garble into one mass point on each side. By concavity this increases the expected payoff.

Claim 4: if π is an optimal policy, then the set $[0, d]$ cannot contain bi-pools.

Proof: Suppose μ_1 and μ_2 form a bi-pool and $\mu_1, \mu_2 \in [0, d]$. Any convex-combination of μ_1 and μ_2 is an element of $[0, d]$ too, in particular the mass point that arises from garbling μ_1 and μ_2 into one signal is in that set as well. Since the integrand is concave for $\mu \in [0, d]$, garbling the signal improves expected utility, contradicting that π is an optimal policy.

Claim 5: if π is an optimal policy, then there cannot be any bi-pool formed by $\mu_{-1} \in [-d, 0)$ and $\mu_1 \in (0, d]$.

Proof: By the previous claims, the set of states that are bunched into the bi-pools must be a superset of the set $[-d, d]$. If this was not the case, then there would have to be more than one masspoint, either for the positive states or for the negative states. Moreover, all states that are not bunched into the bi-pool are revealed by claim 1.

Note that the sets are taken as half open. If any of the masspoints is at zero, then claim 3 shows this is not optimal.

Suppose, contrary to the claim, that there were mass points at $\mu_{-1} \in [-d, 0)$ occurring with p_{-1} and at $\mu_1 \in (0, d]$ occurring with p_1 .

If the bi-pool is symmetric, then we can create a more informative signal by simple pools, formed by all negative states in the bi-pool and all positive states in the bi-pool. Moreover, by symmetry, the probability distribution has equal probabilities on the positive and negative pools. The simple pools generate higher absolute values of the means and hence generate higher expected utility.

Suppose now the bi-pool is asymmetric. Suppose, for any $\varepsilon > 0$, there are positive states of mass ε that are pooled into μ_{-1} and negative states of mass ε that are pooled into μ_1 . Then, create an improvement by pooling the (smallest) negative states into $\mu'_{-1} < \mu_{-1}$ and the (largest) positive states into $\mu'_1 > \mu_1$. Continue as long as this is possible.

When no longer possible, the sets must have the following structure. Without loss, we describe one out of two (isomorphic) situations:

- i) there is a threshold such that all states below the threshold are mapped into μ_{-1} ;
- ii) in addition, a (possibly nonempty) set of states above the threshold (of possibly nonnegative mass) is also mapped into μ_{-1} ;
- iii) all other states are mapped into μ_1 .

Create an improvement as follows:

Pick a mass δ of the highest states above the threshold that were mapped into μ_{-1} and map them into μ_1 . Pick δ such that there exists an interval of states at the low end of the support of μ_1 that are all mapped into μ_1 and map them into μ_{-1} . The probability distribution remains the same, but μ_{-1} decreases and μ_1 increases. This is appreciated.

Claims 1–5 have established that given delegation set $D = \mathbb{R} \setminus (-d, d)$, the expert-optimal policy is characterized by three thresholds: one negative, one positive, and one between the two. Below the low threshold and above the high threshold the state is revealed; the rest of the states form two simple pools divided by the median threshold.

Claim 6: The symmetric solution is characterized by

$$(1 - \sqrt{\lambda})z + \sqrt{\lambda}\mu_z = d. \quad (22)$$

Proof: Suppose S chooses an extremes-refining binary policy with threshold z . If z is optimal then it must weakly exceed d (he does not gain from a learning $\tilde{\omega}$ that he cannot match y to) and hence in the pooling regions he picks $y_D(\mu_z) = d$ and $\mu_D(\mu_{-z}) = -d$.

His expected payoff, $-\int (y - \mu)^2 d\pi(\mu) + (1 - \lambda) \int \mu^2 d\pi(\mu) - k - \sigma^2 + t$, becomes

$$-\bar{F}(z) (d - \mu_z)^2 + (1 - \lambda) \left[\bar{F}(z) \mu_z^2 + \int_z^{\tilde{\omega}} \omega^2 d\bar{F}(\omega) \right] - k - \sigma^2 + t.$$

This expression is strictly concave in z , and the FOC of maximization is

$$2\bar{F}(z) (d - \lambda\mu_z) \frac{d\mu_z}{dz} - \bar{f}(z) (d^2 - 2d\mu_z + \lambda\mu_z^2) - (1 - \lambda)z^2 \bar{f}(z) = 0.$$

Note that

$$\frac{d\mu_z}{dz} = \frac{d}{dz} \left(\frac{\int_0^z \omega d\bar{F}(\omega)}{\bar{F}(z)} \right) = \frac{\bar{f}(z)(z - \mu_z)}{\bar{F}(z)}.$$

Substituting this back into the FOC and canceling terms, we get

$$2(d - \lambda\mu_z)(z - \mu_z) - (d^2 - 2d\mu_z + \lambda\mu_z^2) - (1 - \lambda)z^2 = 0. \quad (23)$$

Completing the square in the middle term and rearranging,

$$2(d - \lambda\mu_z) = \frac{(1 - \lambda)(z^2 - \mu_z^2) + (d - \mu_z)^2}{z - \mu_z}.$$

Notice that the right-hand side is the slope of the chord through the graph of the payoff function at μ_z and z , whereas the left-hand side is the slope of the payoff function, $-(d - \omega)^2 + (1 - \lambda)\omega^2$, evaluated at $\omega = \mu_z$, because

$$\left(-(d - \omega)^2 + (1 - \lambda)\omega^2 \right)' = 2(d - \omega) + 2(1 - \lambda)\omega = 2(d - \lambda\omega).$$

This is the tangency condition illustrated in Figure 3.

Rearrange (23) in a different way, by expanding the brackets and cancelling terms:

$$2dz - 2d\mu_z - 2\lambda z\mu_z + 2\lambda\mu_z^2 - d^2 + 2d\mu_z - \lambda\mu_z^2 - z^2 + \lambda z^2 = 0$$

$$\Leftrightarrow -2\lambda z\mu_z + \lambda\mu_z^2 + \lambda z^2 = z^2 - 2dz + d^2.$$

$$\Leftrightarrow \lambda(z - \mu_z)^2 = (z - d)^2,$$

which is equivalent to the claim because $z \geq d$ and $z \geq \mu_z$.

Claim 7: if the prior satisfies $\frac{\partial}{\partial x}\mathbb{E}[\omega | \omega \in [a, x]] \leq 1$ and $\frac{\partial}{\partial x}\mathbb{E}[\omega | \omega \in [x, b]] \leq 1$ for all $a < x < b$, then the expert-optimal policy given delegation set $\mathbb{R} \setminus (-d, d)$ is to learn whether the state belongs to $[-z, 0]$ or $[0, z]$, or its exact value otherwise.

Proof: Suppose that the three thresholds characterizing the expert-optimal policy are $z_{-1} < 0 < z_1$ and $x \in [z_{-1}, z_1]$. Let μ_{-1} and μ_1 denote the posterior means when the state is in $[z_{-1}, x]$ and $[x, z_1]$, respectively. Write the expert's expected payoff, with constant terms omitted, as a function of the intermediate threshold x :

$$\Gamma(x) = \max_{z_{-1}, z_1} \left\{ - \int_{z_{-1}}^x (-d - \mu_{-1})^2 dF(\omega) - \int_x^{z_1} (y - \mu_1)^2 dF(\omega) \right. \\ \left. + (1 - \lambda) \left[(F(x) - F(z_{-1})) \mu_{-1}^2 + (F(z_1) - F(x)) \mu_1^2 + \int_{\underline{\omega}}^{z_{-1}} \omega^2 dF(\omega) + \int_{z_1}^{\bar{\omega}} \omega^2 dF(\omega) \right] \right\}.$$

Maximize this in x taking into account that z_{-1} and z_1 are set optimally.

By the Envelope Theorem,

$$\Gamma'(x) = -(-d - \mu_{-1})^2 f(\omega) + (d - \mu_1)^2 f(\omega) \\ + 2 \int_{z_{-1}}^x (-d - \mu_{-1}) \frac{\partial \mu_{-1}}{\partial x} dF(\omega) + 2 \int_x^{z_1} (d - \mu_1) \frac{\partial \mu_1}{\partial x} dF(\omega) \\ + (1 - \lambda) \left\{ 2[F(x) - F(z_{-1})] \mu_{-1} \frac{\partial \mu_{-1}}{\partial x} + 2[F(z_1) - F(x)] \mu_1 \frac{\partial \mu_1}{\partial x} + (\mu_{-1}^2 - \mu_1^2) f(\omega) \right\}.$$

Using

$$\frac{\partial}{\partial x} \int_{z_{-1}}^x \frac{\omega dF(\omega)}{F(x) - F(z_{-1})} = \frac{(x - \mu_{-1})f(x)}{F(x) - F(z_{-1})}$$

and

$$\frac{\partial}{\partial x} \int_x^{z_1} \frac{\omega dF(\omega)}{F(z_1) - F(x)} = \frac{(\mu_1 - x)f(x)}{F(z_1) - F(x)}$$

we can write

$$\begin{aligned}\Gamma'(x) &= -(-d - \mu_{-1})^2 f(x) + (d - \mu_1)^2 f(x) \\ &\quad + 2(-d - \mu_{-1})(x - \mu_{-1}) f(x) + 2(d - \mu_1)(\mu_1 - x) f(x) \\ &\quad + (1 - \lambda) \left[2\mu_{-1}(x - \mu_{-1}) f(x) + 2\mu_1(\mu_1 - x) f(x) + f(x)\mu_{-1}^2 - f(x)\mu_1^2 \right]\end{aligned}$$

which simplifies to

$$\begin{aligned}\Gamma'(x) &= -(-d - \mu_{-1} - (x - \mu_{-1}))^2 f(x) + (x - \mu_{-1})^2 f(x) \\ &\quad + (d - x)^2 f(x) - (x - \mu_1)^2 f(x) + (1 - \lambda) f(x) \left[(x - \mu_1)^2 - (x - \mu_{-1})^2 \right]\end{aligned}$$

and further to

$$\begin{aligned}\Gamma'(x) &= -(-d - x)^2 f(x) + (d - x)^2 f(x) - \lambda f(x) \left[(x - \mu_1)^2 - (x - \mu_{-1})^2 \right] \\ &= f(x) [-4dx + \lambda(\mu_1 - \mu_{-1})(2x - \mu_1 - \mu_{-1})]\end{aligned}$$

Clearly, $x = 0$ is a stationary point. We need to investigate whether there are any other.

Suppose wlog that $x > 0$ and rewrite the condition as

$$[-4d + 2\lambda(\mu_1 - \mu_{-1})]x = \lambda(\mu_{-1}^2 - \mu_1^2).$$

The solution has the feature that $\mu_1 < d$ and $\mu_{-1} > -d$ so that $-\mu_{-1} < d$. Hence $\mu_1 - \mu_{-1} < 2d$ and

$$-4d + 2\lambda(\mu_1 - \mu_{-1}) < -4d + 4\lambda d = -4(1 - \lambda)d < 0$$

Hence, if there is a solution with $x > 0$, then it needs to be the case that $\mu_{-1}^2 < \mu_1^2$, and moreover

$$\begin{aligned}x &= \frac{\lambda(\mu_{-1}^2 - \mu_1^2)}{-4d + 2\lambda(\mu_1 - \mu_{-1})} \\ &= (\mu_1 + \mu_{-1}) \frac{\lambda(\mu_{-1} - \mu_1)}{-4d + 2\lambda(\mu_1 - \mu_{-1})}.\end{aligned}$$

We note that the term $\frac{-\lambda(\mu_1 - \mu_{-1})}{-4d + 2\lambda(\mu_1 - \mu_{-1})}$ is positive and increasing in $(\mu_1 - \mu_{-1})$.

Moreover, by $-4d + 4\lambda (\mu_1 - \mu_{-1}) < 0$, we have

$$\frac{\lambda (\mu_{-1} - \mu_1)}{-4d + 2\lambda (\mu_1 - \mu_{-1})} < \frac{1}{2}$$

and hence the right side is smaller than $\frac{\mu_1 + \mu_{-1}}{2}$.

From the first-order conditions for the learning bounds,

$$z_1^* = \frac{d - \mu_1 \sqrt{\lambda}}{1 - \sqrt{\lambda}} \quad \text{and} \quad z_{-1}^* = \frac{-d - \mu_{-1} \sqrt{\lambda}}{1 - \sqrt{\lambda}}$$

Therefore

$$\frac{dz_{-1}^*}{dx} = \frac{-\sqrt{\lambda} \frac{\partial}{\partial x} \mu_{-1}^*}{\left(1 - \sqrt{\lambda} + \sqrt{\lambda} \frac{\partial}{\partial z_{-1}^*} \mu_{-1}^*\right)} < 0 \quad \text{and} \quad \frac{dz_1^*}{dx} = \frac{-\frac{\partial}{\partial x} \mu_1^* \sqrt{\lambda}}{\left(1 - \sqrt{\lambda}\right) + \sqrt{\lambda} \frac{\partial}{\partial z_1^*} \mu_1^*} < 0$$

Compute

$$\frac{d\mu_1}{dx} = \frac{\left(1 - \sqrt{\lambda}\right) \frac{\partial}{\partial x} \mu_1^*}{\left(1 - \sqrt{\lambda}\right) + \sqrt{\lambda} \frac{\partial}{\partial z_1^*} \mu_1^*} \quad \text{and} \quad \frac{d\mu_{-1}}{dx} = \frac{\left(1 - \sqrt{\lambda}\right) \frac{\partial}{\partial x} \mu_{-1}^*}{\left(1 - \sqrt{\lambda} + \sqrt{\lambda} \frac{\partial}{\partial z_{-1}^*} \mu_{-1}^*\right)}$$

Observe that we get $\frac{d\mu_{-1}}{dx} \leq 1$ if

$$0 \leq \left(1 - \sqrt{\lambda}\right) \left(1 - \frac{\partial}{\partial x} \mu_{-1}^*\right) + \sqrt{\lambda} \frac{\partial}{\partial z_{-1}^*} \mu_{-1}^*$$

and likewise we get $\frac{d\mu_1}{dx} \leq 1$ if

$$0 \leq \left(1 - \sqrt{\lambda}\right) \left(1 - \frac{\partial}{\partial x} \mu_1^*\right) + \sqrt{\lambda} \frac{\partial}{\partial z_1^*} \mu_1^*.$$

Both conditions hold by assumption (10), hence there cannot be any positive solution to

$$x = (\mu_1 + \mu_{-1}) \frac{\lambda (\mu_{-1} - \mu_1)}{-4y + 2\lambda (\mu_1 - \mu_{-1})}.$$

By symmetry, no solution with $x < 0$ can arise either.

7.5 Proof of Lemma 5

Assume $\lambda > 1$ and denote the delegation set by D . The expert chooses π (a mean-preserving contraction of the prior) in order to maximize

$$- \int (y_D^*(\mu) - \mu)^2 d\pi(\mu) - (\lambda - 1) \int \mu^2 d\pi(\mu). \quad (24)$$

The first term is the loss from expected mismatch, the second term is the net cost of information.

If the delegation set includes zero, then E acquires no information. In what follows assume this is not the case. Let (d_{-1}, d_1) denote the maximal interval around zero excluded by D . Suppose D excludes actions in the possibly empty sets (d_{-3}, d_{-2}) and (d_2, d_3) as well, implying that the actions $[d_{-2}, d_{-1}]$ and $[d_1, d_2]$ are feasible.

We prove that any optimal information structure takes one of three forms: there are two posteriors $\mu_{-1} \leq 0 \leq \mu_1$ such that either

- i) $d_{-1} \leq \mu_{-1} \leq 0 \leq \mu_1 \leq d_1$; or
- ii) $\mu_{-1} < d_{-1} < 0 < \mu_1 \leq d_1$; or
- iii) $d_{-1} \leq \mu_{-1} < 0 < d_1 < \mu_1$,

and d_{-2} is either equal to the lower bound of the support or equal to the upper bound of some gap (d_{-3}, d_{-2}) ; likewise, d_2 is either equal to the upper bound of the support or equal to the lower bound of some gap (d_2, d_3) in the set of available actions.

Let $M_H = \{\mu \in \text{supp}(\pi) : \mu > d_1\}$ and $p_H = \Pr_\pi[\mu \in M_H]$. Likewise, let $M_L = \{\mu \in \text{supp}(\pi) : \mu < d_{-1}\}$ and $p_L := \Pr_\pi[\mu \in M_L]$.

Claim 1: At least one of the sets M_H and M_L is empty, possibly both.

Proof of Claim 1: Start with an arbitrary experiment, π . Let $\mu_H = \mathbb{E}_\pi[\mu | \mu \in M_H]$ and $\mu_L = \mathbb{E}_\pi[\mu | \mu \in M_L]$. Suppose towards contradiction that $p_H, p_L > 0$, and so

$$\mu^* = \frac{p_L}{p_L + p_H} \mu_L + \frac{p_H}{p_L + p_H} \mu_H$$

is well defined.

We will show that there exists a garbled experiment $\hat{\pi}$ such that $\{\text{supp}(\hat{\pi}) : \mu \geq d_1\} = d_1$ and $\{\text{supp}(\hat{\pi}) : \mu \leq d_{-1}\} = d_{-1}$ that the expert prefers over π .

Starting from distribution π , bunch all the posteriors in M_H together and all the posteriors in M_L together, and garble the bunched posteriors with a binary signal $s \in \{s_L, s_H\}$.

By Bayes rule,

$$\Pr[\mu_L | s = s_L] = \frac{p_L \Pr[s_L | \mu = \mu_L]}{p_L \Pr[s_L | \mu = \mu_L] + p_H \Pr[s_L | \mu = \mu_H]},$$

and likewise

$$\Pr[\mu_L | s = s_H] = \frac{p_L \Pr[s_H | \mu = \mu_L]}{p_L \Pr[s_H | \mu = \mu_L] + p_H \Pr[s_H | \mu = \mu_H]}.$$

Note that $\mu_L < d_{-1}$ and $\mu_H > d_1$.

There exist conditional distributions $\Pr[\mu_L | s = s_L] = \alpha \in (0, 1)$ and $\Pr[\mu_L | s = s_H] = \beta \in (0, 1)$ such that

$$d_{-1} = \alpha\mu_L + (1 - \alpha)\mu_H \text{ and } d_1 = \beta\mu_L + (1 - \beta)\mu_H,$$

and a marginal distribution $\Pr[s = s_L] = \gamma_L$ and $\Pr[s = s_H] = \gamma_H$ such that

$$\gamma_L + \gamma_H = p_L + p_H$$

and

$$\frac{\gamma_L}{\gamma_L + \gamma_H} d_{-1} + \frac{\gamma_H}{\gamma_L + \gamma_H} d_1 = \mu^*$$

if, and only if, $\mu^* \in (d_{-1}, d_1)$. Hence, if $\mu^* \in (d_{-1}, d_1)$, then the proof is complete.

So, suppose next that $\mu^* \notin (d_{-1}, d_1)$, and without loss, $\mu^* \geq d_1$.

If $\mu^* = d_1$, then bunching the posteriors in M_H and M_L together into a single signal eliminates all mismatch and reduces the variation in the posterior means, improving (24) on both fronts.

Consider thus the case $\mu^* > d_1$. In this case, an improvement can be found as follows. Take a binary signal σ with values σ_L and σ_H such that σ_L is sent with probability $1 - \delta$ if $\mu \in M_L$ and with $1 - \varepsilon$ if $\mu \in M_H$, whereas σ_H is sent with probability δ if $\mu \in M_L$ and with ε if $\mu \in M_H$.

Hence, letting $\Delta\mu = \mu_H - \mu_L$, we have

$$\begin{aligned} \mathbb{E}[\mu | \sigma = \sigma_L] &= \frac{p_L(1 - \delta)}{p_L(1 - \delta) + p_H(1 - \varepsilon)} \mu_L + \frac{p_H(1 - \varepsilon)}{p_L(1 - \delta) + p_H(1 - \varepsilon)} \mu_H \\ &= \mu_L + \frac{p_H(1 - \varepsilon)}{p_L(1 - \delta) + p_H(1 - \varepsilon)} \Delta\mu \end{aligned}$$

and

$$\begin{aligned}\mathbb{E}[\mu | \sigma = \sigma_H] &= \frac{p_L \delta}{p_L \delta + p_H \varepsilon} \mu_L + \frac{p_H \varepsilon}{p_L \delta + p_H \varepsilon} \mu_H \\ &= \mu_L + \frac{p_H \varepsilon}{p_L \delta + p_H \varepsilon} \Delta \mu.\end{aligned}$$

The prior expectation taken with respect to this signal distribution is

$$\frac{(p_L (1 - \delta) + p_H (1 - \varepsilon))}{p_L + p_H} \mathbb{E}[\mu | \sigma = \sigma_L] + \frac{p_L \delta + p_H \varepsilon}{p_L + p_H} \mathbb{E}[\mu | \sigma = \sigma_H] = \mu^*.$$

As $\mu_H > d_1$, we can find ε and δ such that

$$\mu_L + \frac{p_H \varepsilon}{p_L \delta + p_H \varepsilon} \Delta \mu = d_1.$$

All the solutions exhibit $\delta/\varepsilon = c$, a constant, defined as

$$c = \frac{p_H}{p_L} (\mu_H - d_1).$$

Likewise, as $\mu_L < d_{-1}$, we can find ε and δ such that

$$\mu_L + \frac{p_H (1 - \varepsilon)}{p_L (1 - \delta) + p_H (1 - \varepsilon)} \Delta \mu = d_{-1},$$

such that for all solutions, $(1 - \delta)/(1 - \varepsilon)$ is constant. Substitute in $\delta/\varepsilon = c$, so that

$$\frac{1 - \delta}{1 - \varepsilon} = \frac{1 - \frac{\delta}{\varepsilon} \varepsilon}{1 - \varepsilon} = \frac{1 - c\varepsilon}{1 - \varepsilon}.$$

The range of this expression depends on the level of c . If $c = 1$, then it is identically equal to 1 for all ε . If $c < 1$, then it increases in ε from 1 to ∞ . If $c > 1$, then it decreases in ε from 1 to 0 and is well defined for all $\varepsilon \leq 1/c$.

As $\mu^* > d_1$, the latter case, $c = \delta/\varepsilon > 1$, is relevant (because if $\delta < \varepsilon$ then we would have $\mathbb{E}[\mu | \sigma = \sigma_H] > \mu^*$, increasing the variation). We need to demonstrate that

$$\frac{p_H}{p_L \left(\frac{1 - c\varepsilon}{1 - \varepsilon} \right) + p_H} = \frac{d_{-1} - \mu_L}{\Delta \mu}$$

has a solution. For $\frac{\delta}{\varepsilon} = c$ and $\varepsilon \rightarrow 0$

$$\lim_{\varepsilon \rightarrow 0} \left(\mu_L + \frac{p_H}{p_L \left(\frac{1-c\varepsilon}{1-\varepsilon} \right) + p_H} \Delta\mu \right) = \mu^* > d_1 > d_{-1},$$

while

$$\lim_{\varepsilon \rightarrow 1} \left(\mu_L + \frac{p_H}{p_L \left(\frac{1-c\varepsilon}{1-\varepsilon} \right) + p_H} \Delta\mu \right) = \mu_L < d_{-1}.$$

Hence, there exist ε^* and δ^* such that $\mathbb{E}[\mu | \sigma = \sigma_L] = d_{-1}$ and $\mathbb{E}[\mu | \sigma = \sigma_H] = d_1$ as claimed.

It follows that at most one of the sets M_H and M_L is nonempty and Claim 1 is proven.

□

Claim 2: Any two positive (negative) posteriors cannot have a convex set of included actions between them.

Proof of Claim 2: Suppose not, and there is a convex set of included actions between two positive posteriors. Then, we can find a garbling (using the ε, δ garbling technology seen in the proof of Claim 1) that results in posteriors that are closer together and at the bounds of included sets. This garbling reduces variation and does not increase mismatch. □

It follows from Claim 2 that, if there are more than two positive posteriors, then they are both elements of the same excluded set.

Let $M := \{\mu \in \text{supp}(\pi) : \mu \in [d_{-1}, d_1]\}$.

Claim 3: the set M has at most two elements.

Proof of Claim 3: Suppose to the contrary that there are more than two posteriors in M , so that the same action is taken for two posteriors. Since both mismatch and variation are convex in the posterior, bunching those two posteriors together reduces variation and the expected cost of (squared) mismatch.

By Claims 1-3, there can be at most three posteriors. However:

Claim 4: In an optimal information structure, there is at most one positive and one negative posterior.

Proof of Claim 4: By Claims 1-3 there can be at most three posteriors. By Claim 2, if there are two positive posteriors, then they must lie in a restricted set. Without loss of generality, suppose there are three posteriors, μ_{-1}, μ_2 , and μ_3 , where $d_{-1} \leq \mu_{-1} < 0$ (by Claim 1) and

$0 < d_1 < d_2 < \mu_2 < \mu_3 < d_3$, where (d_2, d_3) is a non-empty set of excluded actions. Let p_{-1} , p_2 and p_3 denote the probabilities with which the corresponding posteriors are realized. The expert's expected payoff (ignoring constants $-\sigma^2 - k + t$) is

$$-p_{-1} (y_{-1} - \mu_{-1})^2 - p_2 (y_2 - \mu_2)^2 - p_3 (y_3 - \mu_3)^2 - (\lambda - 1) (p_{-1} \mu_{-1}^2 + p_2 \mu_2^2 + p_3 \mu_3^2).$$

This cannot be an optimal information structure. The expert could get a higher utility with a garbled version of this information structure, where the garbling is done as follows.

There are two signals, $s \in \{s_L, s_H\}$. In the states in which the posterior mean μ_{-1} is induced under π , the expert receives signal $s = s_L$ (revealing that $\mu = \mu_{-1}$) with probability δ . In these states, signal $s = s_H$ is sent with probability $1 - \delta$, but s_H is also sent all the time in the states in which either μ_2 or μ_3 was induced under π . Let

$$\hat{\mu} = \frac{p_2}{p_2 + p_3} \mu_2 + \frac{p_3}{p_2 + p_3} \mu_3,$$

$$\mu_1(\delta) = \frac{p_{-1}(1 - \delta)}{p_{-1}(1 - \delta) + p_2 + p_3} \mu_{-1} + \frac{p_2 + p_3}{p_{-1}(1 - \delta) + p_2 + p_3} \hat{\mu}$$

Let $p_1 = p_2 + p_3$. Note that for $\delta = 0$, $\mu_1(\delta) = 0$, while for $\delta = 1$, $\mu_1(\delta) = \hat{\mu}$. Moreover, $\mu_1(\delta)$ is increasing and convex in δ with derivatives

$$\frac{\partial \mu_1(\delta)}{\partial \delta} = \frac{p_1 p_{-1}}{(p_{-1}(1 - \delta) + p_1)^2} (\hat{\mu} - \mu_{-1}) > 0$$

$$\frac{\partial^2 \mu_1(\delta)}{\partial \delta^2} = 2 \frac{p_1 p_{-1}^2}{(p_{-1}(1 - \delta) + p_1)^3} (\hat{\mu} - \mu_{-1}) > 0$$

Hence, it is possible to induce $\mu_1(\delta) \in [d_1, d_2]$. If that is the case, then, under the new information structure, the expert's expected payoff (ignoring constants $-\sigma^2 - k + t$) is

$$-p_{-1} \delta (y_{-1} - \mu_{-1})^2 - (\lambda - 1) (p_{-1} \delta \mu_{-1}^2 + (p_{-1}(1 - \delta) + p_1) \mu_1^2).$$

Clearly, the mismatch is reduced. We need to show that the variation is reduced as well.

We first show that the variation is an increasing function of δ . By differentiation,

$$\frac{\partial}{\partial \delta} (p_{-1} \delta \mu_{-1}^2 + (p_{-1}(1 - \delta) + p_1) \mu_1^2) = p_{-1} (\mu_{-1}^2 - \mu_1^2) + 2\mu_1 \frac{p_1 p_{-1}}{p_{-1}(1 - \delta) + p_1} (\hat{\mu} - \mu_{-1})$$

and the variation goes to zero for $\delta = 0$, hence the variation is increasing in δ for δ small.

For $\delta = 1$, and thus $\mu_1 = \hat{\mu}$ and $\frac{p_1}{p_{-1}(1-\delta)+p_1} = 1$,

$$\begin{aligned}\frac{\partial}{\partial \delta} \left(p_{-1} \delta \mu_{-1}^2 + (p_{-1}(1-\delta) + p_1) \mu_1^2 \right) \Big|_{\delta=1} &= p_{-1} (\mu_{-1}^2 - \hat{\mu}_1^2) + 2\hat{\mu} p_{-1} (\hat{\mu} - \mu_{-1}) \\ &= p_{-1} (\mu_{-1} - \hat{\mu})^2 \\ &> 0\end{aligned}$$

Hence, the variation is increasing for δ close to one.

The variation is a convex function of δ . To see this, differentiate once more,

$$\begin{aligned}\frac{\partial^2}{\partial \delta^2} \left(p_{-1} \delta \mu_{-1}^2 + (p_{-1}(1-\delta) + p_1) \mu_1^2 \right) &= -2p_{-1} \mu_1 \frac{p_1 p_{-1}}{(p_{-1}(1-\delta) + p_1)^2} (\hat{\mu} - \mu_{-1}) \\ &\quad + 2 \frac{(p_1 p_{-1})^2}{(p_{-1}(1-\delta) + p_1)^3} (\hat{\mu} - \mu_{-1})^2 \\ &\quad + 2\mu_1 \frac{p_1 p_{-1}^2}{(p_{-1}(1-\delta) + p_1)^2} (\hat{\mu} - \mu_{-1}) \\ &= 2 \frac{(p_1 p_{-1})^2}{(p_{-1}(1-\delta) + p_1)^3} (\hat{\mu} - \mu_{-1})^2 \\ &> 0\end{aligned}$$

so the variation is convex in δ , and since it is nondecreasing in δ for $\delta = 0$, it is increasing in δ for all $\delta \in (0, 1]$.

Finally, for $\delta = 1$, the variation takes the value

$$p_{-1} \mu_{-1}^2 + p_1 \mu_1^2 = p_{-1} \mu_{-1}^2 + p_1 \hat{\mu}_1^2 < p_{-1} \mu_{-1}^2 + p_1 \left(\frac{p_2}{p_2 + p_3} \mu_2^2 + \frac{p_3}{p_2 + p_3} \mu_3^2 \right)$$

where the inequality is due to Jensen's inequality.

It follows that we can choose δ so that $\mu_1(\delta) \in [d_1, d_2]$ which eliminates mismatch and reduces variation, contradicting that π was an optimal information structure. $\square$

7.6 Proof of Lemma 6

Continue to assume $\lambda > 1$. By Lemma 5, the expert's best response to any delegation set is an experiment with at most two posteriors. (If delegation induces some information acquisition, there are exactly two.) We now show that the optimal delegation set for the principal has a symmetric gap around zero; $D = \mathbb{R} \setminus (-d, d)$, and the minimal optimal

delegation set is simply $D = \{-d, d\}$.

Consider a delegation set with an arbitrary gap around zero, (d_{-1}, d_1) , and denote the two posteriors in the expert's best-response experiment by μ_{-1} and μ_1 , with probability p placed on the latter. By Bayesian consistency, $(1 - p)\mu_{-1} + p\mu_1 = 0$. Without loss of generality assume $d_{-1} + d_1 \geq 0$.

If $\mu_1 < \frac{d_{-1} + d_1}{2}$, then there is no information acquisition, since in this case, action d_{-1} (which is weakly closer to zero than action d_1) is optimal for both posteriors. But then, acquiring no information saves costs.

Moreover, posterior μ_1 cannot reveal more than the statistically optimal information. Taken together, feasible posteriors satisfy

$$\frac{d_{-1} + d_1}{2} \leq \mu_1 \leq \mathbb{E} \left[\omega \mid \omega \geq \frac{d_{-1} + d_1}{2} \right].$$

Likewise, for any feasible posterior,

$$p \in \left(0, 1 - \bar{F} \left(\frac{d_{-1} + d_1}{2} \right) \right],$$

because a feasible posterior must be an element of the set of states.

The expert's problem is to choose p and μ_1 to maximize

$$-(1 - p) (\min \{d_{-1}, \mu_{-1}\} - \mu_{-1})^2 - p (\max \{d_1, \mu_1\} - \mu_1)^2 - (\lambda - 1) \left((1 - p) \mu_{-1}^2 + p \mu_1^2 \right)$$

subject to

$$\mu_{-1} = -\frac{p}{1 - p} \mu_1$$

as well as

$$p \in \left(0, 1 - \bar{F} \left(\frac{d_{-1} + d_1}{2} \right) \right], \mu_1 \in \left[\frac{d_{-1} + d_1}{2}, \mathbb{E} \left[\omega \mid \omega \geq \frac{d_{-1} + d_1}{2} \right] \right].$$

Substitute in for μ_{-1} and consider the relaxed problem of maximizing

$$-(1 - p) \left(\min \left\{ d_{-1}, -\frac{p}{1 - p} \mu_1 \right\} + \frac{p}{1 - p} \mu_1 \right)^2 - p (\max \{d_1, \mu_1\} - \mu_1)^2 - (\lambda - 1) \frac{p}{1 - p} \mu_1^2$$

in p and μ_1 while dropping the constraints on both of these variables.

Simplify the problem even further by conjecturing that $\max \{d_1, \mu_1\} = d_1$ and that

$$\min \{d_{-1}, \mu_{-1}\} = d_{-1}.$$

We solve this problem by sequentially maximizing over μ_1 given p , and then optimizing over p .

By straightforward calculus, the solution is

$$\mu_1^* = (1 - p) \frac{(d_1 - d_{-1})}{\lambda}$$

and

$$p^* = \frac{\lambda (d_1 + d_{-1})}{2 (d_{-1} - d_1)} + \frac{1}{2}$$

which implies that

$$\mu_1^* = \left(\frac{1}{2} + \frac{\lambda (d_1 + d_{-1})}{2 (d_{-1} - d_1)} \right) \frac{(d_1 - d_{-1})}{\lambda} = \frac{d_1 + d_{-1}}{2} + \frac{(d_1 - d_{-1})}{2\lambda}$$

and

$$\mu_{-1}^* = -\frac{p^*}{1 - p^*} \mu_1^* = \frac{d_{-1} + d_1}{2} - \frac{(d_1 - d_{-1})}{2\lambda}$$

Note that $\mu_1^* < d_1$ and $\mu_{-1}^* > d_{-1}$ as conjectured.

Next we show that for a given size of the gap, that is, for a fixed $d_1 - d_{-1}$, the expert's and the principal's payoffs are maximized by a symmetrically placed gap. Recall that expected payoffs depend only on the expected squared mismatch and the expected variation.

The expected squared mismatch is

$$(1 - p^*) \left(d_{-1} + \frac{p^*}{1 - p^*} \mu_1^* \right)^2 + p^* (d_1 - \mu_1^*)^2 = \frac{1}{4\lambda^2} (\lambda - 1)^2 (d_{-1} - d_1)^2,$$

while the expected variation is

$$(1 - p^*) (\mu_{-1}^*)^2 + p^* (\mu_1^*)^2 = \frac{1}{4\lambda^2} (d_{-1} - d_1)^2 - \left(\frac{d_{-1} + d_1}{2} \right)^2.$$

Note that only the variation depends on the placement of the gap, and the expected variation is smaller for an asymmetrically placed gap.

Also note that for a symmetric gap around zero, the constraint

$$\frac{(d_1 - d_{-1})}{2\lambda} + \frac{d_{-1} + d_1}{2} \leq \mathbb{E} \left[\omega \mid \omega \geq \frac{d_{-1} + d_1}{2} \right]$$

is met whenever $d_1 = -d_{-1}$ satisfies

$$\frac{d_1}{\lambda} \leq \mu_+.$$

Moreover, for a symmetric gap around zero, the other ignored constraint,

$$p^* = \frac{\lambda (d_{-1} + d_1)}{2(d_{-1} - d_1)} + \frac{1}{2} \leq 1 - \bar{F} \left(\frac{d_{-1} + d_1}{2} \right)$$

is automatically satisfied as well.

The two players' joint expected surplus is

$$-2 \frac{1}{4\lambda^2} (\lambda - 1)^2 (d_{-1} - d_1)^2 + (2 - \lambda) \left(\frac{1}{4\lambda^2} (d_{-1} - d_1)^2 - \left(\frac{d_{-1} + d_1}{2} \right)^2 \right),$$

which simplifies to

$$\lambda (3 - 2\lambda) \frac{1}{4\lambda^2} (d_{-1} - d_1)^2 - (2 - \lambda) \left(\frac{d_{-1} + d_1}{2} \right)^2.$$

For $\lambda \in (1, \frac{3}{2})$, the joint expected surplus is largest for a gap of a given size $d_1 - d_{-1}$, if the gap is placed symmetrically around zero.

The principal maximizes this expression subject to the expert's no-shirking constraint,

$$-\frac{1}{4\lambda^2} (\lambda - 1)^2 (d_{-1} - d_1)^2 - (\lambda - 1) \left(\frac{1}{4\lambda^2} (d_{-1} - d_1)^2 - \left(\frac{d_{-1} + d_1}{2} \right)^2 \right) \geq -(d_{-1})^2,$$

which simplifies to

$$-\lambda (\lambda - 1) \frac{1}{4\lambda^2} (d_{-1} - d_1)^2 + (\lambda - 1) \left(\frac{d_{-1} + d_1}{2} \right)^2 \geq -(d_{-1})^2.$$

For a symmetric gap, the constraint further simplifies to

$$-\frac{\lambda (\lambda - 1)}{\lambda^2} d_1^2 \geq -d_1^2,$$

which always holds as $\lambda^2 - \lambda^2 + \lambda \geq 0$.

It is easy to see that in the most informative symmetric experiment we must have

$$\frac{d_1}{\lambda} = \mu_+,$$

and any $\mu \leq \mu_+$ is implementable with a symmetric delegation set where $d_1 \leq \lambda\mu_+$.

It follows that the principal cannot gain by making the set asymmetric.

7.7 Proof of Lemma 7

Continue to assume $\lambda > 1$. By Lemmas 5 and 6, if the expert decides to acquire some information in the delegation game, then the signal he chooses puts 50-50% on posterior means $\mu = \pm \min\{d/\lambda, \mu_+\}$, where d is the threshold of the (symmetric) restriction picked by the principal. His net payoff from participating and best-responding is

$$-(d - \mu)^2 - (\lambda - 1)\mu^2 - \sigma^2 - k + t = -d^2 + 2\mu d - \lambda\mu^2 - \sigma^2 - k + t.$$

If $d \leq \lambda\mu_+$ and hence $\mu = d/\lambda$ then this net payoff is

$$-d^2 + \frac{d^2}{\lambda} - \sigma^2 - k + t$$

and E 's incentive constraint (to not deviate to $\pi_\emptyset$) becomes

$$-d^2 + \frac{d^2}{\lambda} - \sigma^2 - k + t \geq -d^2 - \sigma^2 + t,$$

equivalently $d \geq \sqrt{\lambda k}$. Together with $d \leq \lambda\mu_+$ this yields the condition

$$\sqrt{\lambda k} \leq d \leq \lambda\mu_+,$$

which holds for some d as long as $k \leq \lambda\mu_+^2$.

If $d \geq \lambda\mu_+$ then the expert's net payoff from participating and best-responding is

$$-d^2 + 2\mu_+d - \lambda\mu_+^2 - \sigma^2 - k + t,$$

and his incentive constraint becomes

$$-d^2 + 2\mu_+d - \lambda\mu_+^2 - \sigma^2 - k + t \geq -d^2 - \sigma^2 + t,$$

equivalently $2\mu_+d - \lambda\mu_+^2 \geq k$. By $d \geq \lambda\mu_+$, the left-hand side weakly exceeds $\lambda\mu_+^2$, therefore $k \leq \lambda\mu_+^2$ is sufficient in this case as well.

The principal's payoff gain (relative to $-\sigma^2$) is twice the expert's net payoff ignoring the terms $-\sigma^2 + t$ with the cost added back in once. If $k \leq \lambda\mu_+^2$ then this is

$$W_D = -2d^2 + 4\mu d - \lambda\mu^2 - k,$$

where $\mu = \pm \min\{d/\lambda, \mu_+\}$.

If $d < \lambda\mu_+$ then, by differentiation, we find that P 's payoff is strictly increasing in d if and only if $\lambda < \frac{3}{2}$.

If $d \geq \lambda\mu_+$ then, again by differentiation, we find that P 's payoff is strictly decreasing in d .

Therefore, if $\lambda \leq \frac{3}{2}$ then P sets $d^* = \lambda\mu_+$ provided E weakly prefers the non-garbled binary (sign-only) experiment to the null experiment; this is indeed the case when $k \leq \lambda\mu_+^2$.

However, if $\lambda > \frac{3}{2}$ then $d^* = 0$ is optimal for P ; that is, she chooses unrestricted delegation, equivalent to no delegation at all.

References

- Alonso, R., and Matouschek, N., (2008): "Optimal Delegation," *Review of Economic Studies*, 75, 259–293.
- Amador, M., and Bagwell, K., (2013): "The Theory of Optimal Delegation with An Application to Tariff Caps," *Econometrica*, 81(4), 1541–1599.
- Argenziano, R., Severinov, S., and Squintani, F., (2016): "Strategic Information Acquisition and Transmission," *American Economic Journal: Microeconomics* 8(3), 119–155.
- Arieli, I., Babichenko, Y., Smorodinsky, R., and Yamashita, T., (2023): "Optimal Persuasion," *Theoretical Economics* 18(1), 15–36.
- Bagnoli, M., and Bergstrom, T., (2005): "Logconcave Probability and its Applications", *Economic Theory*, 26(2), 445–469.
- Benkert, J.-M., Matyskova, L., and Starkov, E., (2025): "Strategic Attribute Learning," *mimeo*, <https://arxiv.org/pdf/2412.10024v2>
- Bilotta, F., Carnehl, C., and Preusser, J. (2026): "Delegated Information Provision," *mimeo*, <https://arxiv.org/pdf/2603.10867>

- Crawford, V., and Sobel, J., (1982): “Strategic Information Transmission”, *Econometrica*, 50, 1431–1451.
- Deimen, I., and Szalay, D., (2019): “Delegated Expertise, Authority, and Communication”, *American Economic Review* 109(4), 1349–74.
- Demski, J., and Sappington, D., (1987): “Delegated expertise,” *Journal of Accounting Research*, 25, 68–89.
- Dworczak, P., and Kolotilin, A. (2024): “The Persuasion Duality,” *Theoretical Economics*, 19, 1701–1755.
- Dworczak, P., and Martini, G., (2019): “The Simple Economics of Optimal Persuasion,” *Journal of Political Economy* 127(5), <https://doi.org/10.1086/701813>
- Eső, P., and Szalay, D., (2020): “Brevity,” *mimeo*.
- Goltsman, M., Hörner, J., Pavlov, G., and Squintani, F., (2009): “Mediation, arbitration and negotiation,” *Journal of Economic Theory* 144, 1397–1420.
- Holmström, B., (1977): “On Incentives and Control in Organizations,” *Ph.D. dissertation*, Stanford University
- Holmström, B., (1980): “On the Theory of Delegation,” Discussion Paper No. 438, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- Ivanov, M., (2010): “Informational Control and Organizational Design,” *Journal of Economic Theory* 145, 721–751.
- Kamenica, E., and Gentzkow, M., (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- Karlin, S., Proschan, F., and Barlow, R. E., (1961): “Moment inequalities of Pólya frequency functions.” *Pacific Journal of Mathematics*, 11(4), 1023–1033.
- Kleiner, A., Moldovanu, B., and Strack, P., (2021): “Extreme Points and Majorization: Economic Applications”, *Econometrica*, 89(4), 1557–1593.
- Kieffer, J., (1983): “Uniqueness of locally optimal quantizer for log-concave density and convex error weighting function”, *IEEE Transactions on Information Theory*, 29, 42–47.
- Kolotilin, A., (2018): “Optimal information disclosure: a linear programming approach,” *Theoretical Economics* 13, 607–636.
- Kolotilin, A., Corrao, R., and Wolitzky, A., (2024): “Persuasion and Matching: Optimal

- Productive Transport," *Journal of Political Economy* 133(4), 1334–1381.
- Kolotilin, A., and Zapechelnyuk, A., (2025): "Persuasion meets delegation," *Econometrica* 93(1), 195–228.
- Kreutzkamp, S., and Lou, Y., (2025): "Persuasion without ex-post commitment," *Journal of Economic Theory* 228, 1006058, <https://doi.org/10.1016/j.jet.2025.106058>
- Lipnowski, E., Ravid, D., (2020): "Cheap talk with transparent motives," *Econometrica* 88, 1631–1660.
- Lipnowski, E., Mathevet, L., and Wei, D., (2020): "Attention Management", *American Economic Review: Insights*, 2(1), 17–32.
- Lou, Y., (2026): "Optimal Delegation: Coarsification through Information Design," *mimeo*.
- Pei, H. D., (2015): "Communication with Endogenous Information Acquisition," *Journal of Economic Theory*, 160, 132–149.
- Rayo, L., and Segal, I., (2010): "Optimal Information Disclosure," *Journal of Political Economy* 118(5), 949–87.
- Sobel, J., (2012): "Complexity versus Conflict in Communication," in *Proceedings of 46th Annual CISS (Conference on Information)*, Princeton, NJ, 2012, 1–6.
- Szalay, D., (2005): "The Economics of Extreme Options and Clear Advice," *Review of Economic Studies*, 72, 1173–1198.
- Szalay, D., (2012): "Strategic information transmission and stochastic orders," *mimeo*.
- Vairo, M., (2025): "Information vs. Flexibility", *mimeo*.