Optimally Biased Expertise*

Pavel Ilinov[1], Andrei Matveenko[2], Maxim Senkov[3], and Egor Starkov[†4]

[1]Ecole Polytechnique Fédérale de Lausanne, School of Architecture, Civil and Environmental Engineering
[2]University of Mannheim, Department of Economics
[3]European Research University, ERUNI Open Research Institute
[4]University of Copenhagen, Department of Economics

February 23, 2024

Report TRANSP-OR 240223
Transport and Mobility Laboratory
School of Architecture, Civil and Environmental Engineering
Ecole Polytechnique Fédérale de Lausanne
transp-or.epfl.ch

*We are grateful to Jean-Michel Benkert, Gregorio Curello, Mogens Fosgerau, Alessandro Ispano, Ole Jann, Eugen Kovac, Matthias Lang, Stephan Lauermann, Filip Matějka, Jeffrey Mensch, Meg Meyer, Fabio Michelucci, Volker Nocke, Nicolas Schutz, Ludvig Sinander, Peter Norman Sørensen, Allen Vong, Jan Zápal, seminar participants at Bern, Copenhagen, EPFL, CERGE-EI, NRU HSE, and THEMA Cergy, as well as the participants of 12th Conference on Economic Design, Oligo Workshop 2022, SAET 2022 and YEM 2022 conferences, CBS Microeconomics workshop, and Nordic Theory Group seminar for helpful comments and valuable feedback. This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 740369) and from the Grant Agency of Charles University (GAUK project No. 323221). Funding by the German Research Foundation (DFG) through CRC TR 224 (project B03) and by the Czech Academy of Sciences through the Lumina Quaeruntur fellowship (LQ300852101) is gratefully acknowledged.

†Ilinov: pavel.ilinov@epfl.ch; Matveenko: matveenko@uni-mannheim.de; Senkov: maxim.senkov@eruni.org; Starkov: egor.starkov@econ.ku.dk.

Abstract

This paper shows that in delegation problems, a biased principal can strictly benefit from hiring an agent with misaligned preferences or beliefs. We consider a "delegated expertise" problem in which the agent has an advantage in acquiring information relative to the principal. We show that it is optimal for a principal who is ex ante biased towards one action to select an agent who is less biased. Such an agent is more uncertain ex ante about what the best course of action is and would acquire more information. The benefit to the principal from a more informed decision always outweighs the cost of a small misalignment. We show that selecting an optimally misaligned agent is a valuable tool, which performs on par with optimal contracting (while imposing no additional cost on the principal) and outperforms restricted delegation. All results continue to hold when delegation is replaced by communication.

JEL-codes: D82, D83, D91, M51.

Keywords: delegation, rational inattention, heterogeneous beliefs, discrete choice.

1 Introduction

Presidents, CEOs, and other leaders are often touted as visionaries, paving the way to a brighter tomorrow. Yet they are not doing this alone. In their reforms they rely on advice and expertise of others – yet they may hire advisors and experts who do not necessarily share the same vision. For example, Lyndon Johnson was passionate about his economic reform, “the War on Poverty”: “That’s my kind of program. I’ll find money for it one way or another. If I have to, I’ll take away money from things to get money for people. ... Give it the highest priority. Push ahead full tilt” (Bailey and Duquette, 2014). Chairing Johnson’s Council of Economic Advisers was Walter Heller, who, while being one of the original authors of the program, was no stranger to pushing against it, advocating for fiscal responsibility and frugality, especially later in the 1960s.[1] Similarly, Ronald Reagan’s radical “Reaganomics” reforms clashed since early on with a more restrained position of the Federal Reserve and its then-chairman Paul Volcker,[2] but that did not stop Reagan from renominating Volcker for a second term in 1983.

Why can it be beneficial for a partisan principal to hire an agent with a misaligned vision? At first sight such a decision looks counterintuitive – e.g., Holmström (1980) suggests that misalignment between a principal and an agent leads to a conflict of interest, since from the principal’s point of view, the agent then makes suboptimal decisions. A similar conclusion could be drawn from the political economy literature, which suggests that political leaders trade off competence versus loyalty when selecting appointees (Lewis, 2011) – one would think that misalignment depresses loyalty, while not necessarily benefitting the competence. Yet, we show in this paper that even conditional on competence, misalignment between a principal and an agent can lead to better decisions or recommendations, and thus benefit a partisan principal.

To show this, we consider a delegation model in which a principal (she) and an agent (he) have common payoffs from different actions given an unobserved state of the world, but have misaligned prior beliefs about the state of the world.[3] The agent does not have any preexisting knowledge about the case he is asked to consider, but can use his expertise to acquire additional information to make the best decision. The agent’s cost of learning is not internalized by the principal, and her own cost of learning is prohibitively high. This setting was labeled by Demski and Sappington (1987) as the “delegated expertise” problem.

We show that when the principal is ex ante *biased* towards some action (in the sense of having a non-uniform prior belief over which of the actions is optimal), it is optimal

[1] Crichton, K. 1987. “Walter Heller: Presidential Persuader.” *The New York Times*, June 21. `https://www.nytimes.com/1987/06/21/business/walter-heller-presidential-persuader.html`

[2] Atkinson, C. 1982. “Reagan, Volcker Meet to Discuss Policy Rift.” *Washington Post*, February 17. `https://www.washingtonpost.com/archive/business/1982/02/17/reagan-volcker-meet-to-discuss-policy-rift/f0e448ae-a08d-46b9-aefd-5f4d5449b04f/`

[3] In the context of our motivating examples, this can be interpreted as both the President and the Economist having the same objective (strong economy, low inflation, high employment, etc.), but different views on which monetary and fiscal tools should be used in order to achieve these objectives in a given situation. For instance, the President could believe tax cuts are more likely to benefit the economy, whereas the Economist could advocate for a more restricted fiscal policy for sake of maintaining control over monetary policy.

for her to hire a *misaligned* agent. In particular, she benefits the most from delegating to an agent who is ex ante less biased and hence more uncertain than she is about what the best course of action is (Propositions 1 and 9, Theorem 1). This is because the more uncertain the agent is, the more he learns about the state, and the better his action fits the state – which benefits the principal.[4] This, however, has to be balanced against the *tilt*: any kind of misalignment between the principal and the agent leads to a tilt in the agent's decisions relative to what the principal would prefer. In the end, the principal prefers to hire an agent who is more uncertain than she is, and who thus conducts a more thorough investigation than an aligned agent would, – but who still favors the same action ex ante. This result holds regardless of who has the final decision rights: the optimal delegation strategy is the same whether the principal delegates the decision rights to the agent or merely expects a recommendation on the optimal course of action (Proposition 10).

This conclusion has implications in various settings. One relates to bilateral relationships – e.g., when an authority in a public organization wants to find the best expert to delegate a decision to. Our findings provide conditions on the set of experts for delegation to be beneficial, as well as an upper bound for the expected gains in such a delegation. Moreover, if it is relatively straightforward for the authority to order available experts in regards to their attitudes, we provide a useful *directional* behavioral tool: the authority should look for an expert who shares similar views but is more uncertain or moderate. Our second interpretation covers large organizations. We take heterogeneous priors as different views of the people in organizations such as research teams, firms, and political parties. Our results speak in favor of diversity of views in such organizations. We characterize a useful diversity strategy for the leader: she benefits from having workers with slightly more *moderate* views. Although the optimal agent is unique in our model, our problem is static and one-shot. For other decision problems, the leader may have different opinions and, therefore, benefits from having workers with different views in the organization.[5]

Importantly, the optimal degree of misalignment is non-monotone in the strength of principal's own bias (Corollary 1). If the principal is unbiased, then she would prefer an aligned agent, since delegating to such an agent both maximizes learning and minimizes tilt. The same, however, is true if the principal is extremely biased – in this case she is almost certain that one action is better than all others, and might either take this action by herself, or find an agent who is aligned with her and just as biased. Hiring a misaligned agent is most optimal for the somewhat-biased principal, who ex ante prefers one action over another, but still values the information that would be collected by the agent.

We further show that the principal can equivalently benefit from leveraging misalignment in preferences rather than misalignment in beliefs. Namely, Theorem 2 states that the best delegation outcome can be implemented by hiring an agent with either optimally misaligned beliefs, or optimally misaligned preferences (or, equivalently, offering action-contingent payments). This result has a mirror implication for the empirical literature

[4]"Learning more" here is meant in the sense of the expected uncertainty reduction from the principal's prior to posterior belief (as measured by mutual information between the two). This measure extends but is not equivalent to the Blackwell order.

[5]Banerjee and Somanathan (2001) and Li and Suen (2004) present a counterargument to the benefits of diversity, arguing that if a decision needs to be made by a collective, diverse collectives may have a harder time agreeing on a decision and may produce worse outcomes. For a recent review on diversity in organizations, see Shore et al. (2009).

estimating discrete choice models: Theorem 2 implies that the agent's observed action choice probabilities alone do not allow an external observer to jointly identify the decision maker's beliefs and preferences in our setting.

The main conclusion of our paper is that delegation to an agent with misaligned beliefs is an instrument that is available – and valuable – to the principal. Not only that, but in our setting this instrument can not be outperformed by contracts with action-contingent payments (Theorem 2) or outcome-contingent payments (Proposition 4). Conversely, misalignment in beliefs can itself outperform contracting if we take the cost of contracts for the principal into account. Further, misalignment is typically better than restricting the agent's choice set (Proposition 5). This benefit of misalignment challenges the opinion that disagreement between the principal and the agent inevitably leads to a conflict, so the principal should seek to hire an agent who is most aligned with her preferences and beliefs (see Holmström, 1980; Crawford and Sobel, 1982; Prendergast, 1993; Alonso and Matouschek, 2008; Egorov and Sonin, 2011; Che et al., 2013 for some examples of such a message).

The existence of the principal's trade-off between the amount of information acquired by an agent and the tilt in his resulting decisions relies on the flexibility of the agent's learning technology. We capture this flexibility by using the Shannon model of discrete rational inattention, which allows the agent to acquire arbitrary signals and parametrizes the cost of such a signal through the expected entropy reduction (see Maćkowiak et al., 2023 for a recent survey of the literature on rational inattention).[6] The choice of a signal in this model depends on the agent's prior belief: an agent whose prior is skewed towards some state of the world chooses a signal which is relatively more informative regarding that state and thus allows him to make a better decision in that state. This dimension of flexibility is what enables the tilt in the misaligned agent's decisions.

Our results, however, are *not* specific to the Shannon entropy model and continue to hold with other information cost specifications that allow for flexible learning, as shown in Section 6.2 for the channel capacity cost (Woodford, 2012), the log-likelihood ratio cost (Pomatto et al., 2023), and the class of uniformly posterior-separable information cost functions that generalizes the Shannon entropy cost (Caplin et al., 2022). These robustness exercises also confirm that our results are driven by the fact that agents with different prior beliefs seek out different information, and are not purely driven by the artifact that the cost of information is dependent on the agent's prior belief (and so different agents have different concepts of what the *cheaper* information is), which is inherent to the posterior-separable cost functions.

Our paper mainly connects to the literature on delegation. Most papers on delegation follow Holmström (1980) in assuming that the agent has preexisting private information relevant to the decision. We adopt instead the "delegated expertise" setting of Demski and Sappington (1987), where the agent has no information advantage over the principal ex ante, but rather has to collect information, and the expertise grants him a *learning* advan-

[6] The entropy parametrization has been rationalized in both information theory as a cost function arising from the optimal encoding problem (see Cover and Thomas, 2012) and decision theory as arising naturally from Wald's sequential sampling model (see Hébert and Woodford, 2019), and has been shown to work as a microfoundation of the logit choice rule commonly used in choice estimation (Matějka and McKay, 2015). We mainly explore the model of finding the best alternative, studied in Caplin et al. (2019); the Shannon model with other preference specifications, including a binary-quadratic model closer to classic delegation framework, is studied in Section 6.1.

tage over the principal.[7] Demski and Sappington (1987) explore a contracting problem in a setting in which the agent chooses between a finite number of signal structures. Lindbeck and Weibull (2020) extend this analysis to a rationally inattentive agent (who can acquire any information subject to entropy costs).

Deimen and Szalay (2019) study a version of the delegated expertise problem, in which the agent's bias is also unknown ex ante, and show that communication outperforms delegation in such a problem. Our results imply that this conclusion depends on the model specifics, as in the setting we consider, communication and delegation yield identical results. Szalay (2005) shows that restricting the agent's action set could be a useful tool in such a setting, since banning an ex ante optimal "safe" action can nudge the agent to acquire more information about which of the risky actions is the best. Our grand message is similar: the principal is willing to sacrifice something in exchange for the agent acquiring more information, but we present a different channel through which the principal can achieve this.

The closest to our paper is contemporary work by Ball and Gao (2021). They consider a model of delegated expertise and demonstrate a result similar to that of Szalay (2005): that banning the ex ante safe actions can lead to more information acquisition by the agent, which benefits the principal. However, where Szalay (2005) looks at the scenario in which the principal's and the agent's preferences coincide ex post (i.e., net of information costs), Ball and Gao (2021) explore a model with misaligned preferences and show that the principal may benefit from some misalignment between her preferences and those of the agent. In their setting, this is due to divergence between the principal's and the agent's ex ante optimal actions (due to preference misalignment), which makes banning the ex ante agent-preferred action less costly for the principal. Our paper suggests a different channel through which misalignment may incentivize the agent's information acquisition: using a flexible information acquisition framework, we show that misalignment can lead to more information acquisition by the mere virtue of the agent being more uncertain than the principal about what the optimal action is.

The effects of misalignment in prior beliefs are also studied by Che and Kartik (2009). They analyze a delegated expertise game in which the principal retains the decision rights, and the agent, after acquiring the relevant information, chooses whether to disclose it to the principal. They show that the need to communicate may incentivize a misaligned agent to acquire more information than an aligned one, in order to more effectively persuade the principal about which action needs to be taken, as well as to avoid the punishment for concealing evidence. As we show in Section 6.3, both the persuasion and the prejudice avoidance channels are absent from our model, even if we consider communication (as opposed to delegation, as in the baseline model). Our explanation of the desirability of misalignment is thus completely separate from that of Che and Kartik (2009). We argue instead that agents are heterogeneous in their ex ante uncertainty regarding the optimal action, and this heterogeneity can be exploited by hiring a more uncertain agent, who would put more effort into learning the state – even if such an agent would be misaligned relative to the principal. Our setup further allows us to obtain novel comparative

[7] Graham et al. (2015) show that delegation tends to be used when the decision-making demands more evidence that the delegatee can provide. Alternatively, the choice to delegate a decision is often associated with a volatile environment that a delegator faces (Foss and Laursen, 2005; Ekinci and Theodoropoulos, 2021), where any knowledge quickly becomes obsolete.

statics results and show that the optimal misalignment is non-monotone in the principal's bias.

Finally, there exists a literature that argues in favor of misaligned delegation in strategic settings, as a way to commit to a certain strategy. Examples include Rogoff (1985); Segendorff (1998); Kockesen and Ok (2004); Stepanov (2020), and Ispano and Vida (2022). We differ from that literature in focusing on delegation of *non-strategic* decisions, showing how misalignment may be beneficial even in the absence of a strategic counterparty. Additionally, a separate literature argues that people may have intrinsic preferences for potentially biased information, see Masatlioglu et al. (2023) for a recent example. We abstract from such considerations and focus on a setting in which information has purely instrumental value.

The remainder of the paper is organized as follows: Section 2 formulates the main model, which is then analyzed in Section 3 for the special case of binary states and actions, while Section 4 analyzes the general problem. Section 5 compares misaligned beliefs as a delegation tool to other tools, such as misaligned preferences, payments, and restricting the action set. Section 6 explores a number of extensions of the baseline model, and Section 7 concludes.

2 Model

2.1 The Story

We begin by explaining verbally the outline of the model and justifying some of the assumptions made therein; the formal setup follows in Section 2.2.

Consider a principal (she) who would like to implement an optimal decision that depends on the unknown state of the world. To choose the best course of action, the principal delegates the decision to an expert (an agent, he), who has a learning advantage in acquiring information about the state and the optimal decision. For simplicity, we assume that the agent's learning costs (defined further) are finite and the principal's are infinite, but the results extend naturally to the case when the principal's learning costs are finite but larger than the agent's. Further, Section 6.3 demonstrates that communication is equivalent to delegation in our setting (barring the equilibrium multiplicity), so it is not important for our results whether the principal or the agent makes the final decision.

There are many experts available to the principal, and all experts have a common interest with the principal, but differ in their opinions on the issue (Section 5.1 demonstrates the connection of our results to the case of common beliefs but misaligned preferences). These prior beliefs of different agents are observable by the principal – e.g., due to the agents' reputation concerns (i.e., agents needing to publicly establish a particular stance on a broad policy question for sake of earning, and subsequently capitalizing on, a specific reputation). Experts with different initial opinions would acquire different information, and thus possibly make different final decisions. The principal is thus concerned with choosing the best agent for the job. Alternatively, our results can be interpreted as comparative statics for a game between a principal and a given agent with some fixed misalignment, w.r.t. the degree of misalignment. That said, we believe that a literal interpretation of selecting one agent from a population with heterogeneous beliefs is valid as well. Kahneman et al. (2021) survey a large body of evidence suggesting that similar experts and decision-makers in similar conditions make extremely different judgements and predictions, with a large share of these differences attributable to the interpersonal heterogeneity (and a smaller share being due to intra-personal noise in decision-making). We argue that this heterogeneity can be leveraged by the principal through selecting an agent whose bias fits a given problem the most.[8]

2.2 The Setup

The above can be modeled as a game played between a principal and a population of agents. Let $\mathcal{A}$ denote the set of actions with a typical element a, and Ω denote the set of states with a typical element ω. We assume that both $\mathcal{A}$ and Ω are finite. The principal has a prior belief $\mu_p \in \Delta(\Omega)$, where $\Delta(\Omega)$ denotes the set of all probability distributions on Ω. Every agent in the population has some prior belief $\mu \in \Delta(\Omega)$, which is observable by the principal.[9] In what follows, we refer to an agent according to his prior belief. Let $\mathcal{M} \subseteq \Delta(\Omega)$ denote the set of prior beliefs of all agents in the population.[10]

The terminal payoff that both the principal and the agent selected by the principal receive when action a is chosen in state ω is given by $u(a, \omega)$. Prior to making the decision, the selected agent can acquire additional information about the realized state. We assume that the agent can choose any signal structure defined by the respective conditional probability system $\phi : \Omega \to \Delta(\mathcal{S})$, which prescribes a distribution over signals $s \in \mathcal{S}$ for all states $\omega \in \Omega$, where $\mathcal{S}$ is arbitrarily rich. The information is costly: when choosing a signal structure ϕ, the agent must incur cost $c(\phi, \mu)$ that depends on the informativeness of the signal ϕ and the agent's prior belief μ.[11]

The cost function we consider is the Shannon entropy cost function used in rational inattention models (Matějka and McKay, 2015). In this specification, the cost is proportional to the expected reduction in entropy of the agent's belief resulting from receiving the signal (we consider other cost functions in Section 6.2 to show that our results do not depend on this particular specification). Namely, let $\eta : \mathcal{S} \to \Delta(\Omega)$ denote the agent's

[8]Note that the evidence presented by Kahneman et al. (2021) implies that the population of principals would also be heterogeneous in their judgements of a given decision problem. This, together with the inherent heterogeneity of problems, would create demand for a wide variety of experts – and hence mitigate the agents' desire to conceal or misrepresent their biases.

[9]To clarify, we work with a model of non-common prior beliefs about ω, and we take this assumption at face value. Such settings are not uncommon in economic theory: see Morris (1995); Alonso and Câmara (2016); Che and Kartik (2009) for some examples and discussion. It is well known (see Aumann, 1976; Bonanno and Nehring, 1997) that agents starting with a common prior can not commonly know that they hold differing beliefs. We allow the agents to have heterogeneous prior beliefs, and thus to "agree to disagree". While it may be possible to replicate our results in a common-prior model with asymmetric information, where an agent's ex ante belief is affected by some private information not observed by the principal, such a model would feature signaling concerns (e.g., an agent learning something about the principal's information about the state from the fact that he was chosen for the job, and the principal then exploiting this inference channel). We prefer to abstract from such signaling and simply assume non-common priors from the start.

[10]For many of the results we assume that the population of agents is rich enough to represent the whole spectrum of viewpoints: $\mathcal{M} = \Delta(\Omega)$.

[11]Similar to, e.g., Alonso and Câmara (2016), we assume that the agent and the principal share the understanding of the signal structure. Combined with them having different (subjective) prior beliefs over states, this implies they would also have different (subjective) posterior beliefs if both could observe the signal realization.

posterior belief system, obtained from μ and ϕ using the Bayes' rule. The cost is then defined as

$$c(\phi,\mu) \equiv \lambda\Bigg(-\sum_{\omega\in\Omega} \mu(\omega)\ln\mu(\omega) + \\ + \sum_{\omega\in\Omega}\sum_{s\in\mathcal{S}} \left(\sum_{\omega'\in\Omega} \mu(\omega')\phi(s|\omega')\right) \eta(\omega|s)\ln\eta(\omega|s)\Bigg), \tag{1}$$

where $\lambda \in \mathbb{R}_{++}$ is a cost parameter.[12] We assume that the principal does not internalize the cost of learning, and the agent fully bears this cost. The main interpretation (shared by, e.g., Lipnowski et al., 2020) of this assumption is that the cost reflects the cognitive process of the agent. Information acquisition costs thus lead to moral hazard, with the agent potentially not willing to acquire the amount of information desired by the principal. This is the main conflict between the two parties in our model.

In line with the delegation literature, we assume that the principal cannot use monetary or other kinds of transfers to manage the agent's incentives. This is primarily because learning is non-contractible in most settings – indeed, it is difficult to think of a setting, in which a learning-based contract could be enforceable, i.e., either the principal or the agent could demonstrate beyond reasonable doubt exactly how much effort the agent has put into learning the relevant information, and what kind of conclusions he has arrived at. A simpler justification of the no-transfer assumption could be that such transfers are institutionally prohibited in some settings (see Laffont and Tirole, 1990; Armstrong and Sappington, 2007; Alonso and Matouschek, 2008 for some examples and a discussion of such settings). Section 5 shows, however, that even when contracting is feasible, it does not improve upon hiring an agent with a misaligned belief, and neither can restricting the set of actions that the agent is allowed to choose from.

The game proceeds as follows. In the first stage, the principal selects an agent from the population based on the agent's prior belief μ. In the second stage, the selected agent chooses signal structure ϕ and pays cost $c(\phi,\mu)$. In the third stage, the agent receives signal s according to the chosen signal structure ϕ and selects action a given s. Payoffs $u(a,\omega)$ are then realized for the principal and the agent.

The following subsections describe the respective optimization problems faced by the principal and her selected agent, and introduce the equilibrium concept.

2.3 The Agent's Problem

The agent selected by the principal chooses a signal structure $\phi : \Omega \to \Delta(\mathcal{S})$ and a choice rule $\sigma : \mathcal{S} \to \mathcal{A}$ to maximize his expected payoff net of the information costs. The agent's objective function is

$$\mathbb{E}[u(a,\omega)|\mu] - c(\phi,\mu) = \sum_{\omega\in\Omega}\mu(\omega)\sum_{s\in\mathcal{S}}\phi(s|\omega)u(\sigma(s),\omega) - c(\phi,\mu).$$

The agent's problem can then be written down as

$$\max_{\phi,\sigma}\left\{\sum_{\omega\in\Omega}\mu(\omega)\sum_{s\in\mathcal{S}}\phi(s|\omega)u(\sigma(s),\omega) - c(\phi,\mu)\right\}. \tag{2}$$

[12]We follow the standard convention and let $0\ln 0 = 0$.

Lemma 1 in Matějka and McKay (2015) shows that problem (2) with entropy cost function can be reframed as a problem of selecting a collection of conditional choice probabilities. This reformulation is presented in Section 2.6.

2.4 The Principal's Problem

The principal's problem is to choose an agent to delegate to. The choice is based on the agent's prior belief $\mu \in \mathcal{M}$ in order to maximize her expected utility from the action eventually chosen by the chosen agent. Her objective function is

$$\mathbb{E}[u(a,\omega)|\mu_p] = \sum_{\omega\in\Omega} \mu_p(\omega) \sum_{s\in\mathcal{S}} \phi(s|\omega)u(\sigma(s),\omega),$$

so her optimization problem can be written down as

$$\begin{aligned} &\max_{\mu} \left\{ \sum_{\omega\in\Omega} \mu_p(\omega) \sum_{s\in\mathcal{S}} \phi_\mu(s|\omega)u(\sigma_\mu(s),\omega) \right\}, \\ &\text{s.t. } (\phi_\mu, \sigma_\mu) \text{ solves (2) given } \mu, \end{aligned} \tag{3}$$

where the choice of agent μ affects the signal structure ϕ_μ and the choice rule σ_μ chosen by the agent. Therefore, the principal's problem is effectively that of choosing a pair (ϕ, σ) from a menu given by the agents' equilibrium strategies.

2.5 Equilibrium Definition

We now present the equilibrium notion used throughout the paper; the discussion follows.

Definition (Equilibrium)**.** *An equilibrium of the game is given by* $(\mu^*, \{\phi^*_\mu, \sigma^*_\mu\}_{\mu\in\mathcal{M}})$*: the principal's choice* $\mu^* \in \mathcal{M}$ *of the agent who the task is delegated to and a collection of the agents' information acquisition strategies* $\phi^*_\mu : \Omega \to \Delta(\mathcal{S})$ *and choice rules* $\sigma^*_\mu : \mathcal{S} \to \mathcal{A}$ *for all* $\mu \in \mathcal{M}$*, such that:*

1. ϕ^*_μ *and* σ^*_μ *constitute a solution to* (2) *for every* $\mu \in \mathcal{M}$*;*

2. μ^* *is a solution to* (3) *given* $(\phi^*_\mu, \sigma^*_\mu)$.

Note that the above effectively defines a Subgame-Perfect Nash Equilibrium. Ravid (2020) and Cusumano et al. (2024) show that solving games with rationally inattentive players becomes tricky when such players' beliefs are endogenous, and care is required in defining an equilibrium. While our game features incomplete information (about the state of the world chosen by Nature), and the players' beliefs play a central role in the analysis, problem formulations (2) and (3) allow us to treat these beliefs as exogenous parameters or regular maximization variables entering the terminal payoff functions. This is because one player's actions do not affect another player's beliefs in this game, hence a belief consistency requirement is not needed (however, we do require internal consistency in that the agent's posterior belief η is obtained by updating his prior belief μ via Bayes' rule given his requested signal structure ϕ).

2.6 Preliminary Analysis

Matějka and McKay (2015) show that with entropy costs, the agent's problem of choosing the information structure and choice rule can be reduced to the problem of choosing the conditional action probabilities. Namely, the maximization problem of the agent can be rewritten as that of choosing a decision rule $\pi : \Omega \to \Delta(\mathcal{A})$ (which is a single state-contingent action distribution, as opposed to the combination of a signal strategy $\phi : \Omega \to \Delta(\mathcal{S})$ and a choice rule $\sigma : \mathcal{S} \to \mathcal{A}$):

$$\max_{\pi} \left\{ \sum_{\omega \in \Omega} \mu(\omega) \left(\sum_{a \in \mathcal{A}} \pi(a|\omega) u(a, \omega) \right) - c(\pi, \mu) \right\}, \tag{4}$$

where $c(\pi, \mu)$ denotes, with abuse of notation, the information cost induced by the action distribution π.[13] Lemma 2 in the online appendix of Matějka and McKay (2015) implies that with the utility functions we adopt in the remainder of the paper, the agent's problem has a unique solution in either formulation (up to signal labels). Let $\beta(a_i)$ denote the respective unconditional probability of choosing alternative a_i (calculated using the agent's own prior belief μ):

$$\beta(a) \equiv \sum_{\omega \in \Omega} \mu(\omega) \pi(a|\omega).^{14} \tag{5}$$

The principal's problem can then be rewritten as choosing $\mu \in \mathcal{M}$ that solves

$$\begin{aligned} &\max_{\mu} \left\{ \sum_{\omega \in \Omega} \mu_p(\omega) \left(\sum_{a \in \mathcal{A}} \pi_\mu(a|\omega) u(a, \omega) \right) \right\}, \\ &\text{s.t. } \pi_\mu \text{ solves (4) given } \mu. \end{aligned} \tag{6}$$

In what follows, we refer to problem (6) as the principal's **full problem**. Our main interest in what follows lies in the properties of the solution μ^* of the full problem and the chosen agent's optimal strategy π_{μ^*}.

We now proceed to analyze the model described above.

3 Binary Case

We start by looking at the binary-state, binary-action version of the model, since the results can be presented more clearly in such a setting than in either the general model with N states (explored in Section 4), or a quadratic-loss model more frequently encountered in delegation literature (explored in Section 6.1).[15] We show that the principal has to balance

[13] Cost $c(\pi, \mu)$ is calculated as the cost $c(\phi, \mu)$ of the cheapest strategy (ϕ, σ) that generates decision rule π. The choice rule in such a strategy is deterministic, and the signal strategy prescribes at most one signal per action (Matějka and McKay, 2015, Lemma 1).

[14] Given this notation, we can express $c(\pi, \mu)$ as

$$c(\pi, \mu) = \lambda \left(\sum_{\omega \in \Omega} \mu(\omega) \left(\sum_{a \in \mathcal{A}} \pi(a|\omega) \ln \pi(a|\omega) \right) - \sum_{a \in \mathcal{A}} \beta(a) \ln \beta(a) \right).$$

[15] Such a double-binary model is common in the delegation literature, see e.g. Li and Suen (2004) with a slightly different informal story.

off the amount of information acquired against the *nature* of information acquired – since agents with different prior beliefs tilt their learning towards different states. This makes the principal favor agents who are somewhat more uncertain than her regarding the state, but who do not necessarily have a uniform prior belief (Proposition 1).

Assume that the state space is $\Omega = \{l, r\}$ and with abuse of notation let us represent beliefs μ by the probability they assign to state r, so $\mu \in [0, 1]$. Assume further that the action set is $\mathcal{A} = \{L, R\}$, and the common utility net of information costs that the principal and the agent get from the decision is given by $u(L|l) = u(R|r) = 1$ and $u(L|r) = u(R|l) = 0$. We proceed by backward induction, looking at the agent's problem first, and then using the agent's optimal behavior to solve the principal's problem of choosing the best agent.

The agent is allowed to choose any informational structure (Blackwell experiment) he wants, paying the cost which is proportional to the expected reduction of the Shannon entropy of his belief. Using the result presented in Section 2.6, the agent's problem can be reformulated as the problem of choosing a stochastic decision rule $\pi : \Omega \to \Delta(\mathcal{A})$, which solves

$$\max_{\pi} \left\{ \mu\pi(R|r) + (1-\mu)\pi(L|l) - c(\pi, \mu) \right\}. \tag{7}$$

The solution to this problem can be summarized by the two precisions $\{\pi(R|r), \pi(L|l)\}$ or, alternatively, the two unconditional probabilities $\{\beta(R), \beta(L)\}$. Using Theorem 1 in Matějka and McKay (2015), we get that

$$\pi(L|l) = \frac{\beta(L)e^{\frac{1}{\lambda}}}{\beta(L)e^{\frac{1}{\lambda}} + \beta(R)}, \qquad \pi(R|r) = \frac{\beta(R)e^{\frac{1}{\lambda}}}{\beta(L) + \beta(R)e^{\frac{1}{\lambda}}}, \tag{8}$$

and their Corollary 2 implies that

$$\beta(R) = \frac{\mu e^{\frac{1}{\lambda}} - (1-\mu)}{e^{\frac{1}{\lambda}} - 1}, \qquad \beta(L) = \frac{e^{\frac{1}{\lambda}}(1-\mu) - \mu}{e^{\frac{1}{\lambda}} - 1}, \tag{9}$$

cropped to $[0, 1]$.[16] Combining (8) and (9), we get that the solution to problem (7) is given by[17]

$$\begin{aligned} \pi(R|r) &= \frac{\left(\mu e^{\frac{1}{\lambda}} - (1-\mu)\right) e^{\frac{1}{\lambda}}}{\left(e^{\frac{2}{\lambda}} - 1\right)\mu}, \\ \pi(L|l) &= \frac{\left((1-\mu)e^{\frac{1}{\lambda}} - \mu\right) e^{\frac{1}{\lambda}}}{\left(e^{\frac{2}{\lambda}} - 1\right)(1-\mu)}, \end{aligned} \tag{10}$$

cropped to $[0, 1]$. Figure 1 demonstrates how the agent's action precisions choice depends on his prior belief.

[16]Hereinafter, "$f = g$ cropped to $[0, 1]$" means "$f = \min\left\{1, \max\left\{0, g\right\}\right\}$".

[17]This solution takes the form of the so-called rational inattention (RI) logit. In comparison to the standard logit behavior, under RI-logit the decision-maker (the agent in our case) has a stronger tendency to select the ex ante optimal alternatives more frequently.

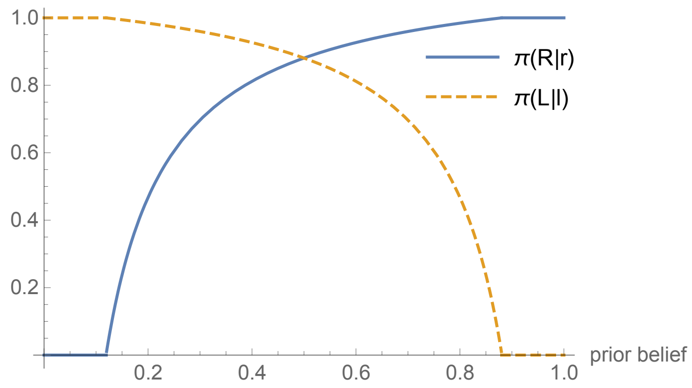

Figure 1: Solution of the agent's problem (7) for different prior beliefs μ.

In turn, the principal's problem is

$$\begin{aligned} &\max_{\mu} \{\mu_p \pi_\mu(R|r) + (1-\mu_p)\pi_\mu(L|l)\}, \\ &\text{s.t. } \pi_\mu \text{ solves problem (7) given } \mu. \end{aligned} \tag{11}$$

It is easy to see by comparing the payoffs in (7) and (11) that the principal benefits from higher precisions $\pi(R|r)$ and $\pi(L|l)$, the same as the agent. However, the relative weights the principal and the agent assign to these precisions depend on their respective priors μ_p and μ, and are hence different. Hence, in order to understand the trade-offs that the principal faces in hiring agents with different priors, we need to explore how the agent's optimal strategy (10) depends on his prior belief μ.

When solving problem (7), the agent faces a trade-off between increasing the precision of his decisions, $\pi(R|r)$ and $\pi(L|l)$, and the cost of information. Further, he prefers to learn more about the more probable event: the higher is the probability that the agent's prior belief assigns to $\omega = r$, the more important is precision $\pi(R|r)$ for his payoff, compared to $\pi(L|l)$. Therefore, two agents with different beliefs would acquire *different* information, leading to different precisions $\pi(R|r)$ and $\pi(L|l)$.[18] At the same time, the closer is the prior belief μ to the extremes ($\mu = 0$ or $\mu = 1$), the more confident is the agent about what the state is, and the less relevant is the precision in the other state for him, leading to such an agent acquiring less information in total.

To summarize, the agent's belief μ affects his optimal decision precisions in two ways: a more uncertain agent acquires more information (and hence makes a better decision on average) than an agent who believes one state is more likely. However, the latter is more concerned with choosing the correct action in the ex ante more likely state, while neglecting the other state.

The principal prefers, ceteris paribus, to hire an agent who acquires more information and hence makes better choices – i.e., a more uncertain agent (μ close to 0.5). However, if she believes that, e.g., state r is ex ante more likely ($\mu_p > 0.5$), then she, for all the same reasons as the agent, cares more about the agent choosing the optimal action in state

[18]This feature of the flexible information acquisition model was analyzed in the application to belief polarization by Nimark and Sundaresan (2019), as well as in the marketing literature (see Jerath and Ren, 2021).

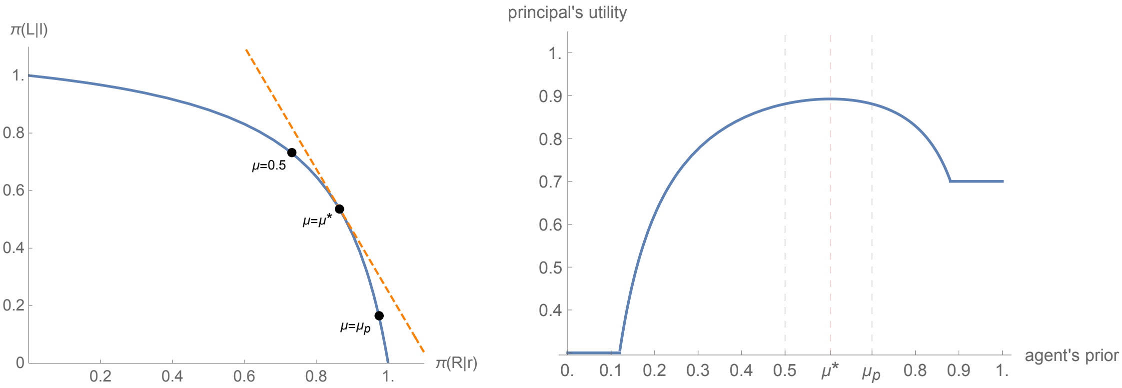

(a) Indifference curve and "budget set". (b) Principal's expected utility.

Figure 2: The principal's problem ($\mu_p = 0.7$).

r than in state l. The latter leads her to prefer an agent who is not completely uncertain ($\mu \neq 0.5$), favoring those who agree with her in terms of which state is more likely ($\mu > 0.5$). Balancing the two issues leads to the principal optimally hiring an agent who has a belief different from hers: $\mu \neq \mu_p$, yet who agrees with her ex ante on the optimal action: $\mu \geq 0.5 \iff \mu_p \geq 0.5$.

Figure 2 presents the principal's problem visually when $\mu_p = 0.7$. Panel (a) plots the principal's "budget set" in the space of two precisions $\pi(R|r)$ and $\pi(L|l)$ (solid blue line), given by the precision pairs optimally chosen by an agent with some prior belief $\mu \in [0, 1]$. Also plotted is the highest indifference curve the principal can achieve (yellow dashed line). The indifference curves are linear, with the slope given by the ratio of probabilities the principal assigns to the two states: $\frac{\mu_p}{1-\mu_p}$. Panel (b) plots the principal's expected utility from hiring an agent as a function of the agent's belief μ.

We can see from the two panels of Figure 2 that the principal with a prior belief $\mu_p = 0.7$ would prefer to hire an agent with a prior belief $\mu \approx 0.6$. Note that the graph in panel (b) is flat for very high and very low μ, which corresponds to the agents who do not learn anything, and simply always choose the ex ante optimal action. Further, agents with low μ (e.g., $\mu \approx 0.15$) acquire non-trivial information, but hiring them is worse for the principal than simply taking the ex ante optimal action for sure (equivalent to hiring an agent with $\mu = 1$). In other words, if an agent is too biased, the information he acquires does not benefit the principal due to the tilt in the agent's actions relative to what the principal would have chosen.

It should be noted that throughout this paper, when we say that a more unbiased agent acquires "more information", the latter is meant in terms of mutual information of the optimal signal structure ϕ^*_μ and the principal's belief μ_p, given by $\frac{c(\phi^*_\mu, \mu_p)}{\lambda}$. Intuitively, one can think of this informativeness measure as the number of bits that the principal needs to encode the information contained in ϕ^*_μ. This informativeness measure imposes an order on experiments ϕ that completes (and therefore respects) the Blackwell order. That said, our statements should not be understood in the sense of Blackwell order: Figure 1 shows that the experiments chosen by any two different agents are (either the same, or) Blackwell-incomparable, since the Blackwell order would require for both precisions $\pi(R|r), \pi(L|l)$ to be higher in one experiment than in another.

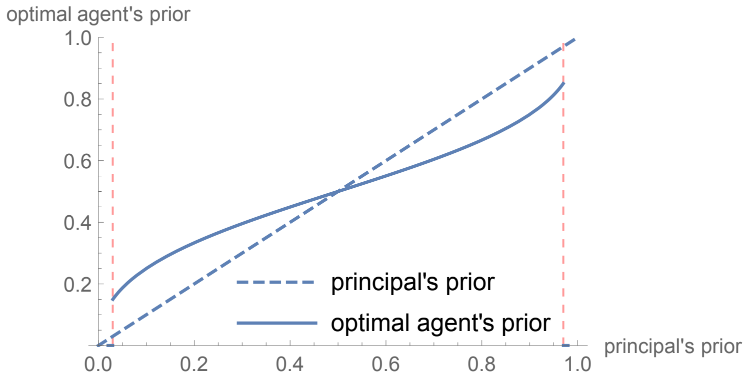

Figure 3: The optimal delegation strategy μ^* as a function of the principal's prior belief μ_p.

Proposition 1 below formalizes the intuition above and provides a closed-form solution for the optimal delegation strategy given the principal's prior belief μ_p. Figure 3 visualizes the optimal delegation strategy as a function of μ_p.

Proposition 1. *If $\mathcal{M} = [0,1]$, then the principal's optimal delegation strategy is given by*

$$\mu^*(\mu_p) = \frac{\sqrt{\mu_p}}{\sqrt{\mu_p} + \sqrt{1-\mu_p}}. \tag{12}$$

Therefore, if $\mu_p \in \left(\frac{1}{2}, 1\right)$, the principal optimally delegates to an agent with belief $\mu \in \left(\frac{1}{2}, \mu_p\right)$.

Remark. *The result in Proposition 1, including the functional form* (12)*, is not exclusive to the Shannon model and is robust to other information cost specifications (see Section 6.2) and other preference specifications (see Section 6.1).*

From Proposition 1 and Figure 3 we can immediately see that misalignment is most beneficial to a moderately-biased principal, while if μ_p is close to either 0.5 or 1, then it is best to hire an (almost-)aligned agent. This is summarized by Corollary 1 below.

Corollary 1. *The optimal misalignment $|\mu_p - \mu^*(\mu_p)|$ is single-peaked in $\mu_p \in \left(\frac{1}{2}, 1\right)$.*

One thing to note about Proposition 1 is that the optimal delegation strategy (12) does not depend on the agent's information cost factor, λ. While it is immediate that the higher is λ, the less information the agent with any given prior μ collects, Proposition 1 serves to show that the trade-off between the quantity of information and the tilt in the decisions does not depend on the absolute quantity of information the agent acquires. Section 6.2 does, however, suggest that this specific conclusion is likely an artifact produced by the entropy information cost function.

Figure 4 demonstrates the difference in action precisions $\pi(R|r), \pi(L|l)$ between delegating to a perfectly aligned agent ($\mu = \mu_p$) and the optimally misaligned agent as given by (12). Optimal delegation leads to the agent consuming more information, lowers the probability of correctly matching the ex ante more likely (according to the principal's belief μ_p) state, $\pi(R|r)$, and increases $\pi(L|l)$, thereby bringing the two closer together.

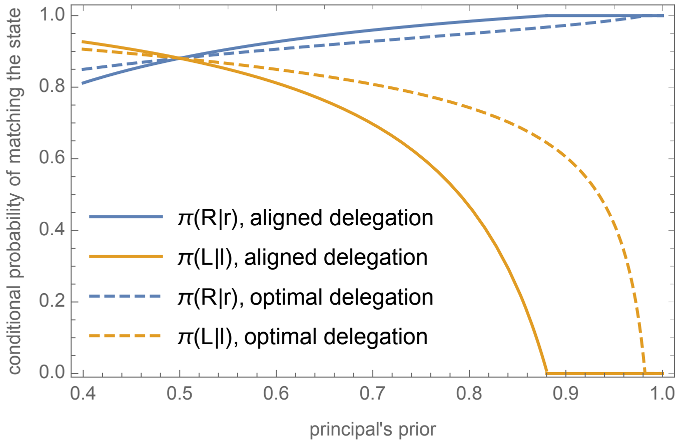

Figure 4: Action precisions under optimal delegation and delegating to the aligned agent.

Overall, under the optimal delegation, the ex ante less attractive option (as seen by both the principal and the agent) is implemented relatively more frequently as compared to the case of the aligned delegation. The principal's benefit from a higher $\pi(L|l)$ under optimal delegation outweighs her loss from a lower $\pi(R|r)$ than under aligned delegation.

Here, an interesting connection can be made to prospect theory (see Barberis, 2013 for a review). In particular, Tversky and Kahneman (1992) suggest that in problems of choice under risk, individual decision-makers tend to succumb to cognitive biases such as overweighing small probabilities and underweighing large probabilities. They propose a probability weighting function that decision-makers unconsciously use, which is reminiscent of our optimal delegation strategy (12), with μ_p being the objective probability and μ^* being the decision-maker's perceived probability. Our result can thus be interpreted as one possible evolutionary explanation of the probability weighting functions. Namely, suppose that "Nature" (evolutionary pressure) is the principal and "Human" is the agent. They both have common utility function $u(a, \omega)$ representing the survival probability of the individual/population, but natural selection is indifferent towards the human's cognitive costs $c(\phi, \mu)$ involved in the decision-making process. In this setting, natural selection would lead humans to develop probabilistic misperceptions according to (12), since these maximize the survival probability.[19]

In the next section, we generalize the binary model, assuming more available alternatives, while keeping the structure of the payoffs the same.

4 General Case

In this section, we extend the analysis to a general problem of finding the best alternative, allowing for $N > 2$ actions and states. We show that the principal's optimal delegation strategy is qualitatively the same as in the binary case, i.e., it is optimal to hire a "more uncertain" agent who investigates more actions in search of the best one than a fully aligned agent. Further, we characterize the whole set of decision rules that can be achieved

[19]Steiner and Stewart (2016) suggest an alternative explanation of probabilistic misperceptions using a similar nature-as-a-principal approach, but a different source of conflict between Nature and Human.

by selecting the agent's prior belief and show that it coincides with what can be achieved by selecting action-contingent subsidies for the agent.

We are now looking at the model with $\mathcal{A} \equiv \{a_1, ..., a_N\}$ and $\Omega \equiv \{\omega_1, ..., \omega_N\}$ for some N, and the preferences are given by $u(a_i, \omega_i) = 1$ and $u(a_i, \omega_j) = 0$ if $i \neq j$. Without loss of generality, we assume that the principal's belief $\mu_p \in \Delta(\Omega)$ is such that $\mu_p(\omega_1) \geq \mu_p(\omega_2) \geq \ldots \geq \mu_p(\omega_N)$ (otherwise states and actions can be relabeled as necessary; here $\mu_p(\omega)$ denotes the probability that belief μ_p assigns to state ω). As before, results from Section 2.6 apply, meaning that the agent's problem is equivalent to choosing the action distribution $\pi : \Omega \to \Delta(\mathcal{A})$ to maximize (4), and the principal selects an agent according to his prior $\mu \in \mathcal{M}$ to maximize (6). We do not restrict the choice of agents and let $\mathcal{M} = \Delta(\Omega)$ (i.e., for any probability distribution $\mu \in \Delta(\Omega)$, the principal can find and hire an agent with prior belief μ).

4.1 Agent's Problem

Proceeding by backward induction, we start by looking at the problem of an agent with some prior belief μ. Invoking Theorem 1 from Matějka and McKay (2015), as we did in the binary case, we obtain that the agent's optimal decision rule satisfies:

$$\pi(a_i|\omega_j) = \frac{\beta(a_i)e^{\frac{u(a_i,\omega_j)}{\lambda}}}{\sum_{k=1}^{N} \beta(a_k)e^{\frac{u(a_k,\omega_j)}{\lambda}}}, \tag{13}$$

where $\beta(a_i)$, defined in (5), is the unconditional choice probability according to the agent's prior belief μ, and itself depends on $\{\pi(a_i|\omega_j)\}_{j=1}^{N}$. While (13) does not provide a closed-form solution for the decision rule $\pi(a_i|\omega_j)$, it implies that the conditional choice probabilities π are uniquely determined given the unconditional choice probabilities β, and this mapping depends solely on the agent's payoffs and not on his prior belief. In what follows, we use the implication that a collection of the unconditional choice probabilities β pins down the whole decision rule π and use β to summarize the agent's chosen decision rule.

The above is not to say that closed-form expressions cannot be obtained. Caplin et al. (2019) show (see their Theorem 1) that in the setting of Matějka and McKay (2015), an agent with a prior belief μ optimally chooses a decision rule that generates unconditional choice probabilities

$$\beta(a_i) = \max\left\{0, \frac{1}{\delta}\left(\frac{(K(\beta)+\delta)\mu(\omega_i)}{\sum\limits_{j\in C(\beta)} \mu(\omega_j)} - 1\right)\right\}, \tag{14}$$

where $\delta \equiv e^{\frac{1}{\lambda}}-1$; $C(\beta) \equiv \{i \in \{1, ..., N\} : \beta(a_i) > 0\}$ denotes the **consideration set**, i.e., the set of actions that are chosen with strictly positive probabilities, and $K(\beta) \equiv |C(\beta)|$ denotes the power (number of actions in) this set.

4.2 Principal's Relaxed Problem

As mentioned previously, (13) implies that a collection of the unconditional choice probabilities β pins down the whole decision rule π. Let us then consider a **relaxed problem** for the principal, in which instead of choosing the agent's prior μ, she is free to select the unconditional choice probabilities $\beta \in \Delta(\mathcal{A})$ directly:

$$\max_{\beta} \left\{ \sum_{j=1}^{N} \mu_p(\omega_j) \left(\sum_{i=1}^{N} \frac{\beta(a_i) e^{\frac{u(a_i,\omega_j)}{\lambda}}}{\sum_{k=1}^{N} \beta(a_k) e^{\frac{u(a_k,\omega_j)}{\lambda}}} u(a_i, \omega_j) \right) \right\}. \tag{15}$$

In the above, we used (13) to represent the conditional probabilities $\pi(a_i|\omega_j)$ in (6) in terms of the unconditional probabilities $\beta(a_i)$. In Section 4.3 we show that the solution to this relaxed problem is implementable in the full problem – i.e., that there exists an agent's belief μ that generates the principal-optimal choice probabilities β.

Note that $\beta(a_i)$ in the above represents the probability with which *an agent* expects to select action a_i. The principal's expected probability of a_i being selected, $\sum_{j=1}^{N} \mu_p(\omega_j)\pi(a_i|\omega_j)$, would generically be different, since her prior belief μ_p is different. Despite the potential confusion this enables, analyzing the principal's problem through the prism of choosing β is the most convenient approach due to the RI-logit structure of the solution to the agent's problem.

Given the state-matching preferences $u(a_j, \omega_j) = 1$, $u(a_i, \omega_j) = 0$ if $i \neq j$, we can simplify (15) to

$$\max_{\beta} \left\{ \sum_{j=1}^{N} \mu_p(\omega_j) \frac{\beta(a_j) e^{\frac{1}{\lambda}}}{1 + \delta\beta(a_j)} \right\}. \tag{16}$$

We can now state the solution to the principal's problem as follows.

Lemma 1. *The solution to the principal's relaxed problem* (16) *is given by*

$$\beta^*(a_i) = \max \left\{ 0, \frac{1}{\delta} \left(\frac{(K(\beta^*) + \delta)\sqrt{\mu_p(\omega_i)}}{\sum\limits_{j \in C(\beta^*)} \sqrt{\mu_p(\omega_j)}} - 1 \right) \right\},$$

where $\delta \equiv e^{\frac{1}{\lambda}} - 1$.

Lemma 1 describes the solution in terms of the action choice probabilities, which do not necessarily give the reader a good idea of its features and the intuition behind this solution. We explore these in more detail in Section 4.4. Before that, however, we need to ensure that this solution is attainable in the principal's full problem, which is done in the following section.

4.3 Principal's Full Problem

The question this section explores is: can the principal generate choice probabilities β^* by appropriately choosing the agent's prior belief μ? In the binary case, the answer was trivially "yes": due to continuity of the agent's strategy, by varying the agent's belief $\mu(r)$ between 0 and 1, the principal could induce any unconditional probability $\beta(R)$. In the multidimensional case, this is not immediate. However, the following result shows that the result still holds with N actions and states under state-matching preferences.

Lemma 2. *In the principal's full problem* (6)*, any vector* $\beta \in \Delta(\mathcal{A})$ *of unconditional choice probabilities is* ***implementable****: there exists a prior belief* $\mu \in \Delta(\Omega)$ *such that* $\beta(a_i) = \sum_{j=1}^{N} \mu(\omega_j)\pi^*_\mu(a_i|\omega_j)$*, where* π^*_μ *solves the agent's problem* (4) *given* μ*.*

The lemma states that if $\mathcal{M} = \Delta(\Omega)$, then the principal can generate any vector of unconditional action probabilities. Note that this does not imply that she is able to select any decision rule $\pi(a_i|\omega_j)$ – if this were the case, under the state-matching preferences she would simply choose to have $\pi(a_i|\omega_i) = 1$ for all i. However, Lemma 2 does imply that the choice probabilities described in Lemma 1 – those that solve the principal's relaxed problem, – are implementable and thus also solve her full problem. The result does, however, rely on the state-matching preferences: we show in Section 5.1 that it does not hold for arbitrary payoff functions.

4.4 Properties of the Optimal Delegation Strategy

While Lemma 1 presents the solution of the principal's problem in terms of the unconditional choice probabilities, this representation is not the most visual. We now demonstrate some implications of this solution in terms of other variables. Namely, Theorem 1 extends Proposition 1 and shows how the chosen agent's prior belief relates to that of the principal. Proposition 2 then compares actions taken under optimal delegation vs aligned delegation.

We begin by looking at the optimal agent choice in terms of the agent's belief μ^*.

Theorem 1. *The principal's equilibrium delegation strategy* μ^* *is such that for all* $i, j \in \{1, ..., N\}$*:*

$$\frac{\mu^*(\omega_i)}{\mu^*(\omega_j)} = \frac{\sqrt{\mu_p(\omega_i)}}{\sqrt{\mu_p(\omega_j)}}.$$

This characterization directly implies the following:

1. $\mu^*(\omega_1) \geq ... \geq \mu^*(\omega_N)$*;*
2. $\mu^*(\omega_1) \leq \mu_p(\omega_1)$ *and* $\mu^*(\omega_N) \geq \mu_p(\omega_N)$*, with equalities if and only if* $\mu_p(\omega_1) = ... = \mu_p(\omega_j)$*;*
3. *distribution* μ^* *is majorized by* μ_p*;*[20]
4. μ^* *has higher entropy than* μ_p*.*

The intuition behind the result above is the same as that behind Proposition 1: the optimally chosen agent is more uncertain than the principal between any given pair of states. To see this, note that if $\mu_p(\omega_i) > \mu_p(\omega_j)$ then $1 < \frac{\mu^*(\omega_i)}{\mu^*(\omega_j)} < \frac{\mu_p(\omega_i)}{\mu_p(\omega_j)}$ – i.e., the agent believes state ω_i is ex ante more likely than ω_j, as the principal does, but he assigns

[20]A discrete probability distribution μ_p is said to majorize another distribution μ^* if both probability vectors are sorted in a descending order and for all $k \in \{1, ..., N\}$:

$$\sum_{i=1}^{k} \mu^*(\omega_i) \leq \sum_{i=1}^{k} \mu_p(\omega_i).$$

relatively less weight to ω_i. This applies to any pair of states. Thus, the implication is that the optimal agent must assign a lower ex ante probability to ω_1, the most likely state according to the principal, than she does, and vice versa for ω_N. Note further that the result in Theorem 1 is again independent of λ, implying that the optimal delegation strategy is determined by the relative trade-off between the quantity of information acquired and the tilt introduced in actions by the misalignment in beliefs, but the absolute quantity of information acquired is irrelevant. In particular, hiring an agent with μ^* is optimal even when he acquires no information, and another agent μ is available, who would be willing to invest effort in learning ω (since such a μ-agent would be too misaligned relative to the principal).

We now switch to comparing the choices made under optimal delegation to those that would arise under **aligned delegation** – i.e., if the principal selected an agent with $\mu = \mu_p$. Let $\bar{\beta}$ denote the choice probabilities that would be generated under aligned delegation. Caplin et al. (2019) show that these probabilities $\bar{\beta}$, as a function of the agent's prior μ, are given by (see their Theorem 1)

$$\bar{\beta}(a_i) = \max\left\{0, \frac{1}{\delta}\left(\frac{(K(\bar{\beta})+\delta)\mu(\omega_i)}{\sum_{j\in C(\bar{\beta})}\mu(\omega_j)} - 1\right)\right\}. \tag{17}$$

Since $\mu_p(\omega_1) > ... > \mu_p(\omega_N)$, the consideration set in the aligned problem is then simply $C(\bar{\beta}) = \{1, ..., \bar{K}\}$, and its size $\bar{K} \equiv K(\bar{\beta})$ is the unique solution of

$$\mu_p(\omega_{\bar{K}}) > \frac{1}{\bar{K}+\delta}\sum_{j=1}^{\bar{K}}\mu_p(\omega_j) \geq \mu_p(\omega_{\bar{K}+1}). \tag{18}$$

In turn, we can see from Lemma 1 that under optimal delegation, size $K^* = K(\beta^*)$ of the consideration set under optimal choice is

$$\sqrt{\mu_p(\omega_{K^*})} > \frac{1}{K^*+\delta}\sum_{j=1}^{K^*}\sqrt{\mu_p(\omega_j)} \geq \sqrt{\mu_p(\omega_{K^*+1})}. \tag{19}$$

These two conditions allow us to compare K^* and $\bar{K}$ directly, which is done by the following proposition.

Proposition 2. *Optimal delegation weakly expands the consideration set relative to aligned delegation:*

$$K(\beta^*) \geq K(\bar{\beta}).$$

In other words, delegating to an optimally misaligned agent leads to a wider variety of actions played in equilibrium. This is a direct consequence of delegation to a more uncertain agent – since he is less sure than the principal of what the optimal action is ex ante, he considers more actions worth investigating. Every action has some positive probability of actually being optimal, and thus a more uncertain agent plays a wider range of different actions ex post. We could already see this effect at play in the binary case, where if μ_p is extreme, then an aligned agent takes the ex ante optimal action without acquiring any additional information, whereas the optimally chosen agent could investigate both actions.

5 Misaligned Beliefs Versus Other Instruments

The preceding analysis set the foundation for using misalignment in beliefs as an instrument in delegation. This section studies how this instrument compares to the other instruments, such as contracting or restricting the delegation set. We keep the overall structure of the problem the same as in Sections 3 and 4, but modify the problem to allow for different tools at the principal's disposal, and compare the outcomes in these modified problems to those in the baseline problem of choosing an agent with the optimal beliefs.

5.1 Contracting on Actions/Misaligned Preferences

The most basic delegation tool is contracting: if the principal could offer the agent a contract that specifies contingent payments, this would be the most direct way to provide incentives (see Laffont and Martimort (2009) for many examples). We begin by looking at *action*-contingent contracts $\tau : \mathcal{A} \to \mathbb{R}$, which allow the principal to incentivize the agent by offering payments that depend on the action that the agent selects. This assumes that actions are contractible (i.e., observable and verifiable) and the principal has the institutional power to make such contracts – either of which may or may not hold in any given setting. We assume that all agents and the principal have a common prior belief μ_p, all players' preferences are quasilinear in payments, and the principal's marginal utility of money is ρ, and the agent's marginal utility of money is 1.[21]

Note that instead of contracting, we can interpret this setup as a problem of selecting an agent with misaligned preferences by setting $\rho = 0$. Schedule τ then represents not payments, but rather an agent's "biases", i.e., inherent preferences towards certain actions on top of the "unbiased" utility $u(a, \omega)$. Such a problem of selecting an agent with optimally misaligned preferences is a natural counterpart to our baseline problem of selecting an agent with optimally misaligned beliefs.

The agent's problem (again using the equivalence presented in 2.6) is then given by

$$\max_{\pi} \left\{ \sum_{j=1}^{N} \mu_p(\omega_j) \sum_{i=1}^{N} \pi(a_i|\omega_j)\big(u(a_i, \omega_j) + \tau(a_i)\big) - c(\phi, \mu_p) \right\}, \tag{20}$$

given τ, and the principal's **contracting problem** is

$$\begin{aligned} &\max_{\tau} \left\{ \sum_{j=1}^{N} \mu_p(\omega_j) \sum_{i=1}^{N} \pi(a_i|\omega_j)\big(u(a_i, \omega_j) - \rho\tau(a_i)\big) \right\} \\ &\text{s.t. } \pi \text{ solves (20) given } \tau. \end{aligned} \tag{21}$$

Instead of providing a closed-form solution to this problem, we appeal to Lemma 2 to argue that regardless of ρ, the principal cannot obtain higher expected utility than in the baseline problem of choosing an agent with a misaligned belief μ. In particular, Lemma 2

[21] In line with the baseline problem, we do not impose any explicit participation constraints on the agent that would impose a lower bound on the transfers. The implicit assumption here is that the agent is being paid some non-negotiable unconditional salary if he is hired, which is sufficient to ensure participation. Payments $\{\tau(a_i)\}$ should then be treated as premia, with the limited liability assumption implying they must be non-negative.

implies that any unconditional choice probabilities $\beta \in \Delta(\mathcal{A})$ generated by an agent, who is incentivized by payments or misaligned preferences, can also be obtained by selecting an agent with appropriately misaligned beliefs. Moreover, by using Proposition 3 of Matveenko and Mikhalishchev (2021) we can also show the converse – that any decision rule achievable with misaligned beliefs can be replicated with payments τ, and the two tools are thus equivalent in what they allow to achieve.[22] These results are formalized by the following theorem.

Theorem 2. *The principal's problem of contracting on actions* (21) *is equivalent to her full (delegation) problem* (6)*:*

1. *For any vector* $\tau : \mathcal{A} \to \mathbb{R}$ *of payments/biases and a corresponding* $\beta : \Omega \to \Delta(\mathcal{A})$ *that solves* (20) *given* τ*, there exists a prior belief* $\mu \in \Delta(\Omega)$ *such that* β *also solves* (4) *given* μ.
2. *For any* $\mu \in \Delta(\Omega)$ *and the corresponding* $\beta : \Omega \to \Delta(\mathcal{A})$ *that solves* (4) *given* μ*, there exist payments* $\tau : \mathcal{A} \to \mathbb{R}$ *such that* β *also solves* (20) *given* τ.

The theorem above directly implies that neither of the two instruments (contracting on actions or searching for an agent with stronger/weaker preferences for specific actions) can yield strictly better results than hiring an agent with an optimally misaligned belief. On the other hand, neither can misaligned beliefs yield better outcomes than action-contingent contracts. However, if the principal's contract choice is subject to the limited liability constraint ($\tau(a_i) \geq 0$ for all i), then it is immediate that contracting on actions is strictly worse, since it cannot yield a better decision rule, but requires payments from the principal – payments which are avoidable if she instead hires an agent who is intrinsically motivated by his beliefs over states or preferences towards specific actions.

Further, our Lemma 2 and the results of Matveenko and Mikhalishchev (2021) also imply that no combination of misaligned beliefs, misaligned preferences, and payments for actions can perform better than any individual instrument. Moreover, they also imply that suboptimal misalignment along any single dimension can be amended using other instruments. That is, if a given agent holds a non-optimal prior belief (that does not coincide with the principal's either), the optimal behavior might be induced via action-contingent transfers. Conversely, if an agent has biased preference towards certain actions, this misalignment can be compensated for by selecting an agent with an approprite prior belief. Proposition 3 presents an example of such equivalence in the context of a model with $N = 2$.

Proposition 3. *Consider the binary setting of Section 3. Consider the principal's problem of contracting on actions* (21)*, where* $\rho = 0$ *and the agent holds prior belief* $\mu \neq \mu_p$. *Then:*

1. *for any* μ*, there exist payments/biases* $\{\tau^*(L), \tau^*(R)\}$ *that implement the optimal conditional choice probabilities from Section 3;*

[22] Lemma 1 of Matveenko and Mikhalishchev (2021) implies that a third equivalent tool is setting the quotas, i.e., imposing specific unconditional choice probabilities for a different action.

2. *these payments/biases are such that*[23]

$$\tau^*(R) \geq \tau^*(L) \iff \mu \leq \mu^* = \frac{\sqrt{\mu_p}}{\sqrt{\mu_p} + \sqrt{1-\mu_p}}.$$

It is easy to see the intuition behind the proposition: if the agent's prior belief μ assigns lower probability to state $\omega = r$ compared to the principal-optimal prior μ^* given in Proposition 1, such an agent is ex ante too biased towards action $a = L$ for the principal's taste, even though he potentially acquires more information than an agent with belief μ^*. Therefore, the principal can nudge the agent towards action $a = R$ by offering higher payment if he selects R (or find an agent whose preference bias towards R offsets his belief bias towards state l).[24] This discussion also emphasizes that what matters for our results is not the agent's uncertainty about the state per se, but the agent's uncertainty about *what the optimal action is*. E.g., an agent who assigns very high probability to state $\omega = l$ can be optimal for the principal, as long as the agent's preferences are sufficiently biased in favor of action $a = R$ – so the agent is actually uncertain about which action to take and chooses to acquire additional information to break the indifference.

5.2 Contracting on Outcomes

We now turn to exploring *outcome*-contingent contracts. An outcome in our model can be measured by whether a correct action was chosen ($a = a_j$ when $\omega = \omega_j$) or not. We thus let the principal select payments $\bar{\tau}, \underline{\tau}$ that the agent receives, so that $\tau(a_i, \omega_i) = \bar{\tau}$ and $\tau(a_i, \omega_j) = \underline{\tau}$ if $i \neq j$.[25] We assume limited liability ($\bar{\tau}, \underline{\tau} \geq 0$), quasilinearity of preferences in payments for all agents, and let the agent's marginal utility of money to be 1, and the principal's marginal utility of money to be ρ.

The agent's problem is then choosing $\pi : \Omega \to \Delta(\mathcal{A})$ that solves[26]

$$\max_{\pi} \left\{ \sum_{j=1}^{N} \mu(\omega_j) \sum_{i=1}^{N} \pi(a_i|\omega_j)\big(u(a_i, \omega_j) + \tau(a_i, \omega_j)\big) - c(\phi, \mu) \right\}, \tag{22}$$

given τ, and the principal's **contracting problem** is

$$\begin{aligned} \max_{\bar{\tau}, \underline{\tau}} & \left\{ \sum_{j=1}^{N} \mu_p(\omega_j) \sum_{i=1}^{N} \pi(a_i|\omega_j)\big(u(a_i, \omega_j) - \rho\tau(a_i, \omega_j)\big) \right\}, \qquad (23) \\ \text{s.t. } & \tau(a_i, \omega_i) = \bar{\tau} \text{ for all } i, \\ & \tau(a_i, \omega_j) = \underline{\tau} \text{ for all } i, j \neq i, \\ & \pi \text{ solves (22) given } \bar{\tau}, \underline{\tau}. \end{aligned}$$

[23] The closed-form expressions are available in the proof in the Appendix.

[24] This is broadly related to the findings of Espitia (2023), who shows that the bias in the agent's preferences can be counteracted by the bias in beliefs (although, the belief biases in his paper are limited to over- and underconfidence).

[25] If the principal could contract on both actions and outcomes, she would have the freedom to select any payment schedule $\{\tau(a_i, \omega_j)\}$. Lindbeck and Weibull (2020) study such a problem with N states and two actions.

[26] While it is more common in the literature to consider an agent who yields no intrinsic utility from actions and is motivated exclusively via payments, for sake of consistency, we maintain the assumption that the agent enjoys the same intrinsic utility $u(a, \omega)$ as the principal, albeit possibly to a different magnitude.

It is trivially optimal for the principal to set $\underline{\tau} = 0$, since her objective is to provide incentives for the agent to *match* the state. Then the agent's (ex post) payoff net of information cost becomes $u(a_i, \omega_j) + \tau(a_i, \omega_j) = (1 + \bar{\tau})u(a_i, \omega_j)$, and the principal's payoff is $u(a_i, \omega_j) - \tau(a_i, \omega_j) = (1 - \rho\bar{\tau})u(a_i, \omega_j)$. In other words, by increasing the incentive payment $\bar{\tau}$, the principal effectively lowers the relative cost of information for the agent, at the cost of decreasing her own payoff. It then appears like an instrument that could be universally useful for the principal – even when she chooses an agent with the optimal prior belief, she could still benefit from reducing the agent's information cost, which would result in him acquiring more information. The following proposition shows, however, that this is not the case: while contracting on outcomes may be a useful instrument, it cannot improve on delegating to the optimally misaligned agent when payments are costly to the principal.

Proposition 4. *Consider the principal's contracting problem* (23) *in the binary setting of Section 3 and suppose* $\mu_p \geq 1/2$. *Then for any* $\rho \geq \min\left\{1, \frac{1}{2\lambda}\right\}$ *there exist* $\bar{\mu}_L, \bar{\mu}_R$ *and* $\hat{\mu}_L, \hat{\mu}_R$ *such that:*

1. $\hat{\mu}_L \leq \bar{\mu}_L < \mu^* < \mu_p < \bar{\mu}_R \leq \hat{\mu}_R$, *where* μ^* *is given by* (12)*;*
2. *the principal's problem* (23) *is solved by* $\bar{\tau} > 0$ *if* $\mu \in (\hat{\mu}_L, \bar{\mu}_L) \cup (\bar{\mu}_R, \hat{\mu}_R)$*;*
3. *the principal's problem* (23) *is solved by* $\bar{\tau} = 0$ *otherwise.*

The proposition states that the principal uses the incentive payments, $\bar{\tau} > 0$, when she has an intermediate degree of misalignment in opinions with the agent. This may happen if the agent is moderately more biased than the principal, $\mu > \mu_p$, and acquires too little information for the principal's taste (which is the case when $\mu_p < \bar{\mu}_R < \mu < \hat{\mu}_R$). An additional reward for matching the state then incentivizes the agent to acquire more information and improves the principal's payoff, despite her giving a part of it to the agent. If the agent is too biased, however ($\mu > \hat{\mu}_R$), then the incentives become too costly for the principal to provide, and she chooses $\bar{\tau} = 0$. The logic is analogous if the agent is sufficiently biased in the opposite direction ($\mu \ll 0.5$). Finally, if the agent is sufficiently aligned with the principal, $\mu \in [\bar{\mu}_L, \bar{\mu}_R]$, then providing bonus payments does not provide enough of an additional incentive to the agent to justify the cost for the principal. This latter case includes both the aligned agent ($\mu = \mu_p$) and the optimally biased agent ($\mu = \mu^*(\mu_p)$). Therefore, the principal's ability to offer incentive payments is not beneficial to her when she has access to a broadly-aligned agent.

5.3 Restricting the Delegation Set

Another instrument commonly explored in the delegation literature is restricting the delegation set – i.e., the set of actions that the agent may take (see, e.g., Holmström, 1980). In particular, in the context of "delegated expertise" problems, Szalay (2005) and Ball and Gao (2021) show that it may be optimal to rule out an ex ante optimal action in order to force the agent to exert effort and learn which of the ex post optimal (but ex ante risky) actions is best. Lipnowski et al. (2020) show a similar result in a Bayesian Persuasion setting in which the receiver is rationally inattentive to the sender's message.

In our setting, however, due to the state-matching utility assumption, there are no "safe" actions that the principal could rule out, as Propostion 2 suggests. Assuming that the principal and the agent hold the same prior belief μ_p, and $\mu_p(\omega_1) > ... > \mu_p(\omega_N)$, action a_1 is the "safest" in the sense of being the most likely ex ante to be optimal. However, it could not be optimal for the principal to ban a_1 – since, indeed, this is the action that is ex ante most likely to be ex post optimal! In other words, while excluding a_1 from the delegation set would lead the agent to acquire more information, it would also lead to larger ex post losses due to the agent being unable to select action a_1 in cases in which it is optimal to do so. Thus while the general idea of the principal being willing to nudge the agent to acquire more information/information about ex ante suboptimal actions holds true in our setting, restricting the delegation set is not an instrument that lends any value to the principal.

Proposition 5 below summarizes this logic. Consider the agent's problem as given by

$$\max_{\pi} \left\{ \sum_{j=1}^{N} \mu_p(\omega_j) \sum_{i=1}^{N} \pi(a_i|\omega_j) u(a_i, \omega_j) - c(\phi, \mu_p) \right\}, \tag{24}$$

given $A^* \subseteq \mathcal{A}$ (and the maximization is w.r.t. a mapping $\pi : \Omega \to \Delta(A^*)$), and the principal's **restriction problem**

$$\begin{aligned} &\max_{A^*} \left\{ \sum_{j=1}^{N} \mu_p(\omega_j) \sum_{i=1}^{N} \pi(a_i|\omega_j) u(a_i, \omega_j) \right\}, \\ &\text{s.t. } \pi : \Omega \to \Delta(A^*) \text{ solves (24) given } A^*. \end{aligned} \tag{25}$$

Then we can state the result as follows.

Proposition 5. *The unrestricted delegation set $A^* = \mathcal{A}$ is always a solution to the principal's restriction problem* (25).

6 Extensions

6.1 Alternative Preference Specifications

The analysis in Sections 4 and 5.1 is heavily reliant on state-matching preferences that we assume are shared by both the principal and the agent(s). It is reasonable to ask whether our conclusions hold under other preference specifications. We first consider a "binary-quadratic" model, which is closely related to the "uniform-quadratic" model commonly used in delegation literature starting with Holmström, 1980 and Crawford and Sobel, 1982. In such a uniform-quadratic model, the state is $\omega \sim U[0, 1]$, the action set is $\mathcal{A} = [0, 1]$, and the players' preferences are described by $u_i(a, \omega) = -(a-(\omega+b_i))^2$ for $i \in \{P, A\}$ and some "biases" b_A, b_P. Such a uniform-quadratic model is not tractable with entropy information costs. Instead, we consider a related model with a binary state and quadratic loss and show that the characterization from Proposition 1 continues to hold in this setting. We then proceed to explore general preferences with finite states and actions, and whether the results of Section 5.1 on the equivalence of misalignment in beliefs and misalignment in preferences carry over to such settings.

Binary-quadratic model. Consider a *binary-quadratic problem* that looks as follows: the state space is $\Omega = \{0, 1\}$; all beliefs μ are represented in terms of probability $\mu \in [0, 1]$ assigned to state $\omega = 1$. The action set is $\mathcal{A} = [0, 1]$; the principal's and the agent's common utility function is $u(a, \omega) = -(a - \omega)^2$; the agent's information cost $c(\phi, \mu)$ is given by the Shannon entropy (1). Everything else is analogous to the binary model presented in Section 3.

We show in the following proposition that our characterization (12) of the optimal delegation strategy continues to hold in such a model, implying that our result is not specific to the problem of finding the best alternative considered in the baseline model.

Proposition 6. *In the binary-quadratic problem, there exists* $\hat{\mu} \geq 0.5$ *such that the principal's optimal delegation strategy is given by:*

$$\mu^{**}(\mu_p) = \begin{cases} \mu^*(\mu_p) & \textit{as given by } (12) \textit{ if } \mu_p \in (1 - \hat{\mu}, \hat{\mu}), \\ \mu_p & \textit{otherwise}. \end{cases}$$

While the result above may suggest that it must then be the entropy cost assumption that is driving the result, Section 6.2 below demonstrates that this is not the case, and the same characterization can be obtained in the binary setting of Section 3 with general posterior-separable cost functions.

Alternative principal's preferences. We now turn to the question of how general the results of Section 5.1 are. We first generalize the principal's utility function $u_p(a, \omega)$ while maintaining the agent's intrinsic preference for matching the state: $u_A(a_i, \omega_i) = 1$, $u_A(a_i, \omega_j) = 0$ if $i \neq j$. Naturally, the specific functional forms of the optimal delegation strategies (such as those presented in Proposition 1, Theorem 1, and Lemma 1) depend on the specific form of the principal's utility function. However, Lemma 2 only depends on the agent's utility function, meaning that Theorem 2 still holds: any outcome that can be achieved by contracting on actions or hiring an agent with misaligned intrinsic preferences, can also be achieved by hiring an agent with misaligned beliefs (and vice versa). Meaning that regardless of the principal's objective function, hiring an agent with state-matching preferences and a suitable belief is as good as hiring an agent with aligned prior belief, state-matching preferences, and either some additional preference over actions, or action-contingent payments on top of that.

Alternative aligned preferences. The argument above does, however, hinge on the agent having state-matching preferences as a baseline. If we allow arbitrary preferences for the agent – even if they align with the principal's preferences net of the information cost – the equivalence stated in Theorem 2 breaks down. In such a general case, finding an agent with optimally misaligned preferences may yield strictly better results for the principal than hiring an agent with an optimally misaligned belief, and hence contracting on actions may, in principle, yield better results too. This is due to the equivalence presented in Section 4.2 breaking down with general preferences, as stated by the following proposition.

Proposition 7. *There exists a utility function* $u(a, \omega)$ *such that the solution to the principal's relaxed problem* (16) *cannot be attained as a solution to the full problem* (6).

Corollary 2. *There exists a non-state-matching utility function $u(a, \omega)$ such that the conclusions of Theorem 2 do not hold.*

The proposition above states that with general preferences, the principal is no longer able to implement any vector of unconditional choice probabilities β via an appropriate choice of the agent's prior μ – which is still possible through the choice of action-contingent contracts as in Section 5.1 (see Proposition 3 in Matveenko and Mikhalishchev, 2021). Therefore, hiring an agent with misaligned preferences can be strictly better for the principal than hiring an agent with misaligned beliefs for general utility functions $u(a, \omega)$. Equivalently, performance of an agent with optimally misaligned beliefs may, in general, be improved upon by providing further incentives through action-contingent contracts.

6.2 Alternative cost functions

Our analysis used Shannon entropy cost function (1) to model the agent's cost of acquiring information. It has an undesirable property that the cost of a given signal structure/Blackwell experiment depends on the agent's prior belief (see Mensch, 2018; Denti et al., 2022). To demonstrate that our main result does not hinge on this or any other specific properties of the entropy cost parametrization, this section explores a number of alternative specifications of the cost function in the binary setting of Section 3. We show that in all cases, the principal's optimal delegation strategy looks similar to what is obtained in Proposition 1: if μ_p is not too extreme, it is optimal for the principal to delegate to an ex ante more unbiased agent: $\mu^*(\mu_p) \in [0.5, \mu_p)$.

To remind, in the binary setting of Section 3, we assumed $\Omega = \{l, r\}$, $\mathcal{A} = \{L, R\}$, and the common utility function net of information costs given by $u(L|l) = u(R|r) = 1$ and $u(L|r) = u(R|l) = 0$. In all settings below, we will be looking for an equilibrium as defined in Section 2.5, with the $c(\phi, \mu)$ in the agent's problem (4) replaced by one of the respective cost functions defined below.

Channel capacity cost function. The first cost function we consider is the channel capacity cost proposed by Woodford (2012). We follow the analysis by Nimark and Sundaresan (2019), hereinafter referred to as NS. The channel capacity cost of a given signal structure $\phi : \Omega \to \Delta(\mathcal{S})$ is given by

$$c_C(\phi) \equiv \max_{\mu \in \Delta(\Omega)} c(\phi, \mu),$$

where $c(\phi, \mu)$ is the entropy cost (1). Intuitively, the channel capacity measures the maximum amount of information that can be extracted from signal ϕ by any agent. By definition, cost $c_C(\phi)$ of a given signal structure does not depend on the selected agent's prior μ, unlike the Shannon entropy cost function.

NS show that the argument from Section 2.6 continues to hold with the channel capacity cost: the agent optimally selects a "recommender" signal structure, where each signal realization is associated with a unique action. Thus we can reduce the agent's problem to that of choosing a decision rule $\pi : \Omega \to \Delta(\mathcal{A})$ which solves

$$\max_{\pi} \Big\{ \mu\pi(R|r) + (1 - \mu)\pi(L|l) - c_C(\pi, \mu) \Big\}, \tag{26}$$

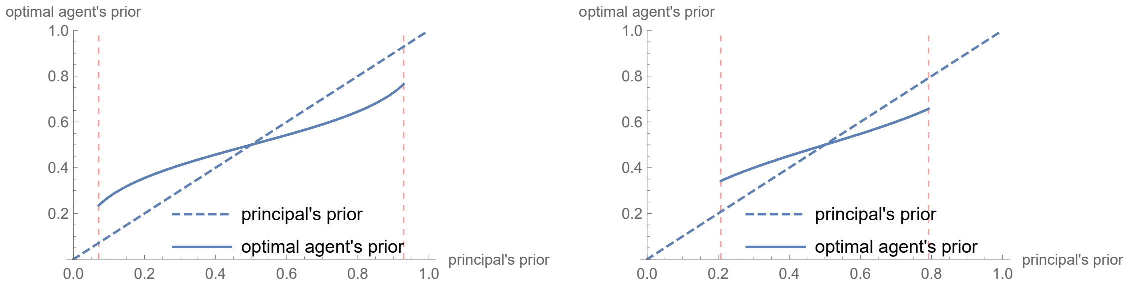

(a) $\lambda = 0.5$ (b) $\lambda = 1$

Figure 5: The optimal delegation strategy $\mu^*(\mu_p)$ with channel capacity cost function.

where $c_C(\pi, \mu)$ denotes the information cost induced by π (which, in this formulation, does depend on μ).

NS show that the agent's problem (26) is well-defined and the solution always exists. It shares the same broad features as the solution with entropy costs: an agent with $\mu > 0.5$ chooses an experiment such that $\pi(R|r) > \pi(L|l)$ and vice versa. More broadly, $\pi(R|r)$ is continuous and increasing in μ, while the opposite is true for $\pi(L|l)$; a more uncertain agent also acquires more information in total.

The continuity of π w.r.t. μ implies that the principal's problem (11) always has a solution. The fact that the agent's behavior is qualitatively the same as a function of μ as it was with entropy costs implies that the principal's trade-off also remains fundamentally the same: more information vs less tilt. While the principal's problem proved to be analytically intractable, Figure 5 presents numerical solutions for two values of λ and all μ_p.

Observation 1. *Figure 5 suggests that the solution of the principal's problem with channel capacity costs* $c_C(\phi)$ *is qualitatively similar to the solution described in Proposition 1 and Corollary 1.*

We see that both plots in Figure 5 demonstrate a delegation strategy that is qualitatively the same as in Figure 3, which plotted it for entropy costs: if the principal's belief μ_p is not too extreme, she chooses an agent with belief μ between μ_p and 0.5. For extreme μ_p, same as before, she selects (an agent who acquires no information and chooses) the ex ante optimal action. However, Figure 5 also highlights a difference relative to the Shannon model, in that the principal's solution now depends on the cost parameter λ: higher λ leads to less delegation under channel capacity costs. This suggests that the independence of the principal's strategy from λ is a special feature of the Shannon model (or, as we see later, of the posterior-separable cost functions more broadly).

Log-likelihood ratio cost function. We now move on to consider the log-likelihood ratio (LLR) cost function proposed by Pomatto et al. (2023), hereinafter referred to as PTS. PTS derive the LLR cost function axiomatically as the cost of *acquiring* information (as opposed to the entropy cost being that of *processing* information, according to their argument) from a set of intuitive cost linearity axioms. The LLR cost of a given signal

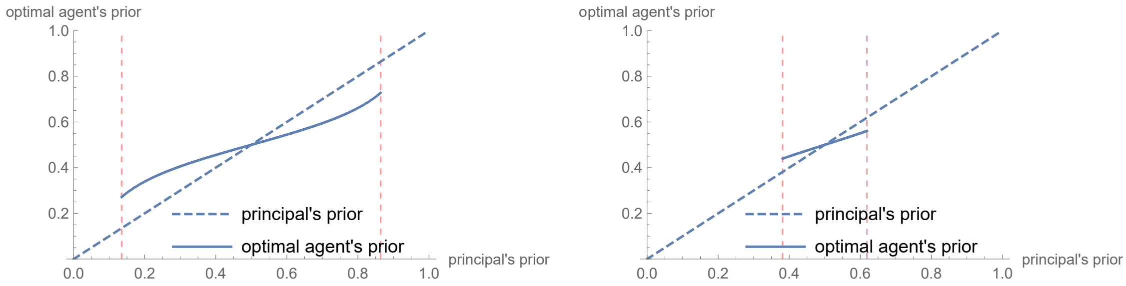

(a) $\lambda = 0.1$ (b) $\lambda = 0.5$

Figure 6: The optimal delegation strategy $\mu^*(\mu_p)$ with LLR cost function.

structure $\phi : \Omega \to \Delta(\mathcal{S})$ is defined as

$$c_L(\phi) \equiv \sum_{\omega_i, \omega_j \in \Omega} \lambda_{ij} \int_{\mathcal{S}} \ln \left(\frac{\phi(s|\omega_i)}{\phi(s|\omega_j)} \right) d\phi(s|\omega_i),$$

where λ_{ij} are the parameters encoding the "closeness" of states ω_i and ω_j (how difficult it is to distinguish them). In our binary setting, we assume $\lambda_{LR} = \lambda_{RL} = \lambda$. As in the case of channel capacity costs, PTS' main representation theorem shows that the LLR cost of a given experiment ϕ does not depend on the prior belief μ.

In the binary setting, the agent optimally chooses no more than two signal realizations, because LLR cost is monotone with respect to the Blackwell order. Therefore, we can again invoke the logic of Section 2.6 and reduce the agent's problem to that of choosing a decision rule $\pi : \Omega \to \Delta(\mathcal{A})$ subject to cost $c_L(\pi)$.[27]

PTS explore a binary problem in their Sections 6.1 and 6.6 but only demonstrate an analytical solution to the agent's problem for the case $\mu = 0.5$. We have found the agent's problem to be analytically intractable for $\mu \neq 0.5$, and therefore solve both the agent's and the principal's problems numerically. Figure 6 demonstrates our findings.

Observation 2. *Figure 6 suggests that the solution of the principal's problem with log-likelihood ratio costs $c_L(\phi)$ is qualitatively similar to the solution described in Proposition 1 and Corollary 1.*

One can see that the principal's optimal delegation strategy plotted therein looks qualitatively the same as for entropy and channel capacity costs (Figures 3 and 5, respectively): if the principal's belief μ_p is not too extreme, the principal chooses an agent with belief μ between μ_p and 0.5. Further, similarly to the setting with channel capacity costs in Figure 5, the principal delegates less for higher values of the information cost parameter λ.

Uniformly posterior-separable cost functions. Another class of cost functions we consider is a family of uniformly posterior-separable cost functions (Caplin et al., 2022) that generalizes the Shannon entropy cost function (1). As shown below, the representation from Proposition 1 continues to hold for this wide class of information cost functions. In particular, not only is the main insight ("hiring a less biased agent is beneficial for the

[27] PTS provide a representation for $c_L(\pi)$ not presented here.

principal; more so for intermediate values of the principal's ex ante bias") valid for such cost functions, but even the exact expression for the optimally misaligned agent applies.

Consider again the binary setting of Section 3. A UPS cost function over signal structures $\phi : \Omega \to \Delta(\mathcal{S})$ conditional on the agent's prior belief μ is defined as[28]

$$c_{\text{UPS}}(\phi, \mu) \equiv \lambda \Big[\mathbb{E}[\hat{c}(\eta)|\phi] - \hat{c}(\mu) \Big]$$

where $\eta \in \Delta(\Omega)$ are the posterior beliefs induced by the signal structure ϕ, the expectation is taken over all posteriors η w.r.t. the distribution induced by ϕ conditional on μ, and function $\hat{c}(\eta)$ satisfies the following conditions:

- $\hat{c}(\eta)$ is of Legendre type (strictly convex, differentiable and satisfies the Inada conditions: $\lim_{\eta \to 0+} \hat{c}'(\eta) = -\infty$, $\lim_{\eta \to 1-} \hat{c}'(\eta) = +\infty$; Rockafellar, 1970);
- $\hat{c}(\eta)$ is twice continuously differentiable;
- $\hat{c}(\eta)$ is symmetric around $\eta = 0.5$: if $\eta' + \eta'' = 1$, then $\hat{c}(\eta') = \hat{c}(\eta'')$.

An example of such a function is negative entropy, $\hat{c}(\eta) = -H(\eta)$, which generates the Shannon cost (1). We can now show that for any other function $\hat{c}(\eta)$ that satisfies the conditions above, Proposition 1 continues to hold.

Proposition 8. *In the binary setting of Section 3 with UPS information cost function, the principal's optimal delegation strategy is given by* (12).

This result suggests that the entropy specification is not crucial for our conclusions, which continue to hold for a wide class of posterior-separable cost functions. A question arises then of how far we can push our result – in other words, can it be obtained with any information cost? As we show next, it is the flexibility of the information acquisition technology that plays a significant role in our results.

Symmetric cost functions. Finally, we explore a family of simple "symmetric" cost functions, which restrict the agent to symmetric signals. This analysis highlights the importance of flexibility of the agent's learning technology for the trade-off we identify. In particular, suppose that instead of being able to choose an arbitrary signal structure $\phi : \Omega \to \Delta(\mathcal{S})$, the agent is restricted to a binary signal space $\mathcal{S} \equiv \{l, r\}$ and can only choose signal precision that we denote, abusing notation, by $\phi \equiv \phi(r|r) = \phi(l|l) \in [1/2, 1]$. The cost of information is then given by some function $c_S(\phi)$ that is strictly increasing, convex, differentiable in $\phi \in [1/2, 1]$, and satisfies $c_S(1/2) = 0$ and the Inada conditions $c_S'(1/2) = 0$, $\lim_{\phi \to 1} c_S'(\phi) = +\infty$.

Let ϕ^*_μ denote the precision optimally chosen by an agent with prior belief μ. The agent only acquires information ($\phi^*_\mu > 1/2$) if he intends to follow the signal ($\sigma(r) = R$ and $\sigma(l) = L$), since this trivially dominates doing the converse, and conditional on

[28] Under the conditions on $\hat{c}$ imposed below, this is also known as Bregman information cost (Banerjee et al., 2005; Fosgerau et al., 2023).

ignoring the signal, acquiring an uninformative signal $\phi = 1/2$ is strictly cheaper. Hence if $\phi^*_\mu > 1/2$, then

$$\phi^*_\mu = \arg\max_{\phi} \{\mu\phi + (1-\mu)\phi - c_S(\phi)\}. \tag{27}$$

Let ϕ^{**} denote the candidate solution given by the FOC of (27): $c_S'(\phi^{**}) = 1$. Note that ϕ^{**} does not depend on the agent's belief μ. The agent's expected utility from acquiring no information ($\phi = 1/2$) and taking the ex ante optimal action is given by $\max\{\mu, 1-\mu\}$. The expected utility from choosing $\phi = \phi^{**}$ is given by $\phi^{**} - c_S(\phi^{**}) \in [1/2, 1]$.[29] Then denoting the agents who are indifferent between the two options by $\bar{\mu}_R \equiv \phi^{**} - c_S(\phi^{**})$ and $\bar{\mu}_L \equiv 1 - (\phi^{**} - c_S(\phi^{**}))$, we can describe the agent's optimal choice of precision by

$$\phi^*_\mu = \begin{cases} \phi^{**} & \text{if } \mu \in [\bar{\mu}_L, \bar{\mu}_R], \\ 1/2 & \text{otherwise.} \end{cases}$$

Moving on to the principal's problem (and maintaining the assumption that $\mu_p > 1/2$), the principal's payoff is given by ϕ^{**} if $\phi^*_\mu = \phi^{**}$, by μ_p if $\phi^*_\mu = 1/2$ and $\mu > 1/2$, and by $1 - \mu_p$ if $\phi^*_\mu = 1/2$ and $\mu < 1/2$. Therefore, the principal's preferred agent is

$$\mu^*(\mu_p) = \begin{cases} \mu \in [\bar{\mu}_L, \bar{\mu}_R] & \text{if } \mu_p \le \phi^{**}, \\ \mu \in (\bar{\mu}_R, 1] & \text{if } \mu_p > \phi^{**}. \end{cases} \tag{28}$$

Notably, if $\mu_p \in (\bar{\mu}_R, \phi^{**})$, then the principal strictly prefers a misaligned agent, whose prior belief is more uncertain than the principal's. Further, there exists a selection from (28) that supports the following proposition (which is proved by the preceding argument):

Proposition 9. *Given a symmetric information cost function $c_S(\phi)$, there exists an equilibrium in which the principal always delegates to a more uncertain agent: for any $\mu_p > 1/2$, $\mu^*(\mu_p) \in [1/2, \mu_p)$.*

It is evident, however, that the statement of Proposition 9 is quite weak. Symmetric information cost leads to the principal actually being indifferent between all agents $\mu \in [\bar{\mu}_L, \bar{\mu}_R]$, as well as between all agents $\mu \in (\bar{\mu}_R, 1]$, since within these intervals, μ affects neither the tilt in the agent's decisions, nor the amount of information acquired. Consequently, if $\mu_p \notin (\bar{\mu}_R, \phi^{**})$, then hiring an aligned agent or, possibly, even a more certain agent, is just as optimal for the principal as hiring a more uncertain agent. Conversely, if the principal strictly prefers to hire a learning agent, then she might as well hire an agent with $\mu = 1/2$, whose decisions would not be any more tilted than those of an agent with $\mu = \bar{\mu}_R$.

It is straightforward that the result above continues to hold for any weakly increasing $c_S(\phi)$, whereas all the other assumptions on $c_S(\phi)$ are not strictly necessary and were adopted to simplify the argument. For example, we could also consider a "Pandora box" cost function, under which the agent can either learn the state perfectly at a cost, or learn nothing. This can be captured as $c_S(\phi) = \lambda \cdot \mathbb{I}\{\phi > 1/2\}$, where $\mathbb{I}\{\cdot\}$ is the indicator

[29] Function $\phi - c_S(\phi)$ is strictly concave in ϕ due to the assumptions made, evaluates to $1/2$ when $\phi = 1/2$ (hence $\phi^{**} - c_S(\phi^{**}) \ge 1/2$), and from $c_S(\phi) \ge 0$ we have that $\phi - c_S(\phi) \le 1$ for all $\phi \le 1$.

function. Under such a cost function, the agent would learn the state perfectly if he is sufficiently uncertain, and stick to his prior belief otherwise; and it is thus always weakly optimal for the principal to choose the most uncertain agent: $\mu^*(\mu_p) = 1/2$. Another learning technology that is also symmetric across states and signals, but does not fall under the parametrization above, is the one used by Szalay (2005) and Ball and Gao (2021). In their respective models, the agent selects an effort $e \in [0, 1]$ subject to cost $c_F(e)$, which allows him to perfectly learn the state with probability e, and with the complementary probability $1 - e$ the agent observes no signal. It is easy to see that under this technology, the learning effort e^*_μ is higher when μ is closer to $1/2$ (more uncertain agents learn more). However, same as with symmetric cost functions, there is no disadvantage to hiring a misaligned agent – the principal would strictly prefer to hire the most uncertain agent, $\mu = 1/2$.

The goal of this exercise is to demonstrate that to fully capture the trade-off that the principal faces – that between the *quantity of information* acquired by a misaligned agent and the *decision tilt* that such a misalignment introduces, – a flexible learning technology is necessary. Inflexible technologies, such as those described by the symmetric cost functions, lack the detail to capture fully the trade-off that the principal faces. At the same time, the robustness checks presented above that use the channel capacity and the log-likelihood ratio cost functions demonstrate that our results are not specific to the entropy cost – that it is indeed the flexibility of the agent's learning technology and not the specific features of the cost function that drive our results.

6.3 Communication

In this section, we consider the importance of decision rights in our model with misaligned beliefs. In particular, we juxtapose the *delegation* scheme explored so far, under which the agent has the power to make the final decision, to *communication*, where an agent must instead communicate his findings to the principal, who then chooses the action. A large literature in organizational economics is devoted to comparing delegation and communication in various settings (see Dessein, 2002; Alonso et al., 2008; Rantakari, 2008 for some examples) Deimen and Szalay 2019 in particular explore this question in the context of a delegated expertise problem similar to ours and show that communication outperforms delegation in the class of problems they consider. In contrast, we show that in our setting, communication performs exactly as well as delegation – i.e., the principal will always find it optimal to follow the agent's recommendation.

Although the principal and the agent have the same preferences, it is generally unclear whether it is optimal for the principal to follow a recommendation of an agent due to the misalignment in their beliefs. Namely, since the principal and the agent start from different prior beliefs, the same is true for posteriors: if the principal could observe the information that the agent obtained, her posterior belief would be different from that of the agent. This implies that ex interim, the principal could prefer an action different from the agent's choice, and could benefit from overruling the agent's decision if she had the power to do so. However, this would mean that the agent's incentives to acquire information are different from the baseline model, and the principal could have some influence over the agent's learning strategy via her final choice rule.[30] We show below that, in the end,

[30]Argenziano et al. (2016) provide one example of how the principal can manipulate the agent's infor-

none of these effects come into play, and there exists a communication equilibrium that replicates the delegation equilibrium.

The setup follows the baseline model from Section 2 with state-matching preferences, with the exception that the final stage ("agent selecting action $a \in \mathcal{A}$") is replaced by two. First, after observing signal $s \in \mathcal{S}$ generated by his signal structure ϕ, the agent selects a recommendation (message) $\tilde{a} \in \mathcal{A}$ to the principal. After that, the principal observes the recommendation $\tilde{a}$, uses it to update her belief $\mu_p(\omega|\tilde{a})$ about the state of the world, and then selects an action $a \in \mathcal{A}$ that determines both parties' payoffs.[31] The equilibrium of the communication game is then defined as follows.

Definition (Communication Equilibrium)**.** *An equilibrium of the cheap talk game is characterized by* $(\mu^*, \{\phi^*_\mu, \tilde{\sigma}^*_\mu\}_{\mu\in\mathcal{M}}, \sigma^*, \mu_p)$*, which consists of the following:*

1. *the principal's posterior beliefs* $\mu_p : \mathcal{A} \to \Delta(\Omega)$ *that are consistent with* $(\phi^*_\mu, \sigma^*_\mu)$ *(i.e., satisfy Bayes' rule on the equilibrium path);*

2. *the principal's choice rule* $\sigma^* : \mathcal{A} \to \mathcal{A}$*, which solves the following for every* $\tilde{a} \in \mathcal{A}$*, given the posterior* μ_p*:*

$$\max_{\sigma(\tilde{a})} \left\{ \sum_{\omega\in\Omega} \mu_p(\omega|\tilde{a}) u(\sigma(\tilde{a}), \omega) \right\};$$

3. *a collection of the agents' information acquisition strategies* $\phi^*_\mu : \Omega \to \Delta(\mathcal{S})$ *and communication strategies* $\tilde{\sigma}^*_\mu : \mathcal{S} \to \mathcal{A}$ *that solve the following given* σ *for every* $\mu \in \mathcal{M}$*:*

$$\max_{\phi,\tilde{\sigma}} \left\{ \sum_{\omega\in\Omega} \mu(\omega) \sum_{s\in\mathcal{S}} \phi(s|\omega) u(\sigma(\tilde{\sigma}(s)), \omega) - c(\phi, \mu) \right\};$$

4. *the principal's choice* $\mu^* \in \mathcal{M}$ *of the agent to whom the task is delegated, which solves the following given* $(\phi^*_\mu, \sigma^*_\mu)$*,* σ^**, and* μ_p*:*

$$\max_{\mu} \left\{ \sum_{\omega\in\Omega} \mu_p(\omega) \sum_{s\in\mathcal{S}} \phi(s|\omega) u(\sigma(\tilde{\sigma}(s)), \omega) \right\}.$$

We can then state the result as follows.

Proposition 10. *There exists a communication equilibrium* $(\mu^*, \{\phi^*_\mu, \tilde{\sigma}^*_\mu\}_{\mu\in\mathcal{M}}, \sigma^*, \mu_p)$ *that is outcome-equivalent to the equilibrium* $(\mu^*, \{\phi^*_\mu, \sigma^*_\mu\}_{\mu\in\mathcal{M}})$ *of the original game, in the sense that* μ^* *and* $\phi^*_{\mu^*}$ *coincide across the two equilibria,* $\tilde{\sigma}^*_{\mu^*} = \sigma^*_{\mu^*}$*, and* σ^* *is the identity mapping.*

mation acquisition incentives under cheap talk communication.

[31] Given that message labels are arbitrary, we focus w.l.o.g. on "direct" equilibria, in which the agent's message corresponds to an action recommendation. Further, for simplicity we assume that the principal only observes the recommendation made by the agent, and not the signal he received or the signal structure he requested.

The result is, perhaps, unsuprising, since Holmström (1980) showed that communication is equivalent to restricting the agent's action set, and this latter instrument was shown in Section 5.3 to be irrelevant in our setting, as long as the principal can select an agent with the prior belief she prefers. The result in Proposition 10, however, is subject to a few caveats. First, cheap talk models are plagued by equilibrium multiplicity: for any informative equilibrium, there exist equilibria with less informative communication, up to completely uninformative (babbling) equilibria. In our setting, this means that, in addition to the equilibrium outlined in Proposition 10 above, there also exists a babbling equilibrium in which the agent acquires no information and makes a random recommendation, and the principal always ignores it and selects the ex ante optimal action.[32] There would also likely exist multiple equilibria of intermediate informativeness – e.g., equilibria with a limited vocabulary, where only some actions $\tilde{\mathcal{A}} \subset \mathcal{A}$ are recommended on the equilibrium path. In practice, this means that, under communication, there is a risk of miscoordination on uninformative equilibria, whereas under delegation the equilibrium is unique. The same force may also work the other way, and there may be equilibria that are preferred by the principal to the delegation equilibrium, that can only be sustained under cheap talk (see Argenziano et al., 2016 for an example of how such equilibria may arise). However, the question of whether such equilibria exist is beyond the scope of this paper.

The second caveat lies in the fact that Proposition 10 relies on state-matching preferences. In our setting (with the exception of Section 6.1), any action is either "right" or "wrong", without any degrees of correctness. The misalignment of beliefs across the principal and the agent is thus small enough to not warrant the principal overriding the agent's suggested action. In contrast, in a uniform-quadratic framework of Argenziano et al. (2016) or a normal-quadratic framework of Che and Kartik (2009), both states and actions lie in a continuum, and the principal's loss is proportional to the distance between the realized state and the chosen action. In such a setting, any misalignment (be it in preferences or beliefs) between the principal and the agent would lead to the principal being willing to override the agent's recommendation, leading to the delegation equilibrium being no longer directly sustainable under communication. This ability to exploit interim misalignment is also what drives the persuasion and prejudice avoidance channels that underlie the result of Che and Kartik (2009). By shutting these channels down we provide a novel explanation for the optimality of bias in delegation.

On a separate note, it is immediate from Proposition 10 that the same equilibrium would survive in a setting with *verifiable communication* a là Che and Kartik (2009), where an agent chooses between disclosing a signal that he received and disclosing nothing – as opposed to cheap talk communication assumed above, where the agent can send any message. Since in the cheap talk equilibrium described in Proposition 10, the principal always finds it optimal to follow the (optimally chosen) agent's recommendation and take the agent's most preferred action even in the absence of evidence, the same is true when evidence can be presented. In other words, the agent would never have an incentive to conceal evidence from the principal.

[32] If an agent makes uninformed recommendations, it is optimal for the principal to ignore it. If the principal ignores the recommendation, it is optimal for the agent to not acquire any information. Neither agent in this situation can unilaterally deviate to informative communication.

7 Conclusion

We show that hiring an agent with beliefs that are misaligned with those of the principal can be beneficial for the principal, especially when the principal is ex ante biased. We show this in the context of a model where the agent can acquire costly information before making a decision. More specifically, a biased principal prefers to delegate to an agent who is ex ante more uncertain about what the best action is but is somewhat biased towards the same action as her. This is mainly due to a more uncertain agent being willing to acquire more information about the state, which enables more efficient actions to be taken. As we show, exploiting belief misalignment can be a valid instrument that the principal can use in delegation, which in our setting performs on par with or better than contingent transfers or restriction of the action set from which the agent can choose. The value of this instrument is highest to a moderately-biased principal, whereas both an unbiased and an extremely biased principals would optimally select an aligned agent.

In the analysis, we use the workhorse rational inattention model for discrete choice, the Shannon entropy model. It allows us to provide a richer demonstration of the consequences of delegation to a misaligned agent by allowing the agent to acquire information flexibly, which tilts the decisions of an agent with misaligned beliefs relative to an aligned agent. We show that misaligned delegation is optimal *despite* the tilt introduced by this flexibility. While adopting a particular cost function specification simplifies the analysis, we do show that our results are not specific to the entropy information cost, and continue to hold for a much more general class of posterior-separable information cost functions, as well as other cost functions recently proposed in the literature.

Due to the inherent complexity of rational inattention models, we mostly confine our exploration to a discrete state-matching model, which strays away from the continuous models more commonly used in delegation problems. In a model with a continuum action space, the scope for an agent to manifest his tilt is much larger, and hence the trade-off between the agent's information acquisition and tilted decision-making would supposedly be different. We show, however, that our results extend verbatim to a quadratic-loss framework common in literature on delegation. Due to tractability concerns, we confine the analysis to a binary-quadratic model, as opposed to the more commonly encountered uniform-quadratic (Holmström, 1980) or normal-quadratic (Che and Kartik, 2009) models, but there are few reasons to believe that this dimension of richness would produce results that are substantially different.

An assumption that may feel excessively strong in our analysis is the common knowledge of all agents' and the principal's prior beliefs. It may be more reasonable to assume that agents are strategic in presenting their viewpoints to the principal, as well as making inferences from the fact that they were chosen for the job. Such signaling concerns could yield an economically meaningful effect, yet we abstract from them completely in our paper. A more careful investigation is in order.

References

Aliprantis, C. D. and Border, K. C. (2006). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, third edition.

Alonso, R. and Câmara, O. (2016). Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706.

Alonso, R., Dessein, W., and Matouschek, N. (2008). Centralization Versus Decentralization: An Application to Price Setting by a Multi-Market Firm. *Journal of the European Economic Association*, 6(2-3):457–467.

Alonso, R. and Matouschek, N. (2008). Optimal delegation. *The Review of Economic Studies*, 75(1):259–293.

Argenziano, R., Severinov, S., and Squintani, F. (2016). Strategic Information Acquisition and Transmission. *American Economic Journal: Microeconomics*, 8(3):119–155.

Armstrong, M. and Sappington, D. E. (2007). Recent Developments in the Theory of Regulation. In *Handbook of Industrial Organization*, volume 3, pages 1557–1700. Elsevier.

Aumann, R. J. (1976). Agreeing to Disagree. *The Annals of Statistics*, 4(6):1236–1239.

Bailey, M. J. and Duquette, N. J. (2014). How johnson fought the war on poverty: The economics and politics of funding at the office of economic opportunity. *The journal of economic history*, 74(2):351–388.

Ball, I. and Gao, X. (2021). Benefiting from bias. Working paper.

Banerjee, A., Merugu, S., Dhillon, I. S., Ghosh, J., and Lafferty, J. (2005). Clustering with bregman divergences. *Journal of machine learning research*, 6(10).

Banerjee, A. and Somanathan, R. (2001). A simple model of voice. *The Quarterly Journal of Economics*, 116(1):189–227.

Barberis, N. C. (2013). Thirty Years of Prospect Theory in Economics: A Review and Assessment. *Journal of Economic Perspectives*, 27(1):173–196.

Bonanno, G. and Nehring, K. (1997). Agreeing to disagree: a survey. Invited lecture at the Workshop on Bounded Rationality and Economic Modeling, Firenze, July 1 and 2, 1997 (Interuniversity Centre for Game Theory and Applications and Centre for Ethics, Law and Economics).

Caplin, A. and Dean, M. (2013). The behavioral implications of rational inattention with shannon entropy.

Caplin, A., Dean, M., and Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, 86(3):1061–1094.

Caplin, A., Dean, M., and Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6):1676–1715.

Che, Y.-K., Dessein, W., and Kartik, N. (2013). Pandering to Persuade. *The American Economic Review*, 103(1):47–79.

Che, Y.-K. and Kartik, N. (2009). Opinions as incentives. *Journal of Political Economy*, 117(5):815–860.

Cover, T. M. and Thomas, J. A. (2012). *Elements of information theory*. John Wiley & Sons.

Crawford, V. P. and Sobel, J. (1982). Strategic Information Transmission. *Econometrica*, 50(6):1431–1451.

Cusumano, C. M., Fabbri, F., and Pieroth, F. (2024). Competing to commit: Markets with rational inattention. *American Economic Review*, 114(1):285–306.

Deimen, I. and Szalay, D. (2019). Delegated expertise, authority, and communication. *American Economic Review*, 109(4):1349–1374.

Demski, J. S. and Sappington, D. E. (1987). Delegated expertise. *Journal of Accounting Research*, pages 68–89.

Denti, T., Marinacci, M., and Rustichini, A. (2022). Experimental cost of information. *American Economic Review*, 112(9):3106–23.

Dessein, W. (2002). Authority and communication in organizations. *The Review of Economic Studies*, 69(4):811–838.

Egorov, G. and Sonin, K. (2011). Dictators and their viziers: Endogenizing the loyalty–competence trade-off. *Journal of the European Economic Association*, 9(5):903–930.

Ekinci, E. and Theodoropoulos, N. (2021). Determinants of delegation: Evidence from british establishment data. *Bogazici Journal: Review of Social, Economic & Administrative Studies*, 35(1).

Espitia, A. (2023). Confidence and organizations. Working paper.

Fosgerau, M., Sethi, R., and Weibull, J. (2023). Equilibrium screening and categorical inequality. *American Economic Journal: Microeconomics*, 15(3):201–242.

Foss, N. J. and Laursen, K. (2005). Performance pay, delegation and multitasking under uncertainty and innovativeness: An empirical investigation. *Journal of Economic Behavior & Organization*, 58(2):246–276.

Graham, J. R., Harvey, C. R., and Puri, M. (2015). Capital allocation and delegation of decision-making authority within firms. *Journal of financial economics*, 115(3):449–470.

Hébert, B. and Woodford, M. (2019). Rational Inattention when Decisions Take Time. Technical Report w26415, National Bureau of Economic Research, Cambridge, MA.

Holmström, B. (1980). On the theory of delegation. Discussion Papers 438, Northwestern University, Center for Mathematical Studies in Economics and Management Science.

Ispano, A. and Vida, P. (2022). Good cop-bad cop: delegating interrogations. Technical report, THEMA (Théorie Economique, Modélisation et Applications), Université de Cergy-Pontoise.

Jain, V. and Whitemeyer, M. (2021). Search and competition with endogenous investigations. Technical report, Working Paper.

Jerath, K. and Ren, Q. (2021). Consumer rational (in) attention to favorable and unfavorable product information, and firm information design. *Journal of Marketing Research*, 58(2):343–362.

Kahneman, D., Sibony, O., and Sunstein, C. R. (2021). *Noise: a flaw in human judgment*. Hachette UK.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kockesen, L. and Ok, E. A. (2004). Strategic delegation by unobservable incentive contracts. *The Review of Economic Studies*, 71(2):397–424.

Laffont, J.-J. and Martimort, D. (2009). *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

Laffont, J.-J. and Tirole, J. (1990). The Politics of Government Decision Making: Regulatory Institutions. *Journal of Law, Economics, and Organization*, 6(1):1–32.

Lewis, D. E. (2011). Presidential Appointments and Personnel. *Annual Review of Political Science*, 14(1):47–66.

Li, H. and Suen, W. (2004). Delegating decisions to experts. *Journal of Political Economy*, 112(S1):S311–S335.

Lindbeck, A. and Weibull, J. (2020). Delegation of investment decisions, and optimal remuneration of agents. *European Economic Review*, 129:103559.

Lipnowski, E., Mathevet, L., and Wei, D. (2020). Attention management. *American Economic Review: Insights*, 2(1):17–32.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273.

Marshall, A. W., Olkin, I., and Arnold, B. C. (1979). Inequalities: theory of majorization and its applications.

Masatlioglu, Y., Orhun, Y., and Raymond, C. (2023). Intrinsic information preferences and skewness. *American Economic Review*, 113(10):2615–44.

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.

Matveenko, A. and Mikhalishchev, S. (2021). Attentional role of quota implementation. *Journal of Economic Theory*, 198:105356.

Matyskova, L. and Montes, A. (2023). Bayesian persuasion with costly information acquisition. *Journal of Economic Theory*, page 105678.

Mensch, J. (2018). Cardinal representations of information. *Available at SSRN 3148954*.

Morris, S. (1995). The common prior assumption in economic theory. *Economics & Philosophy*, 11(2):227–253.

Nimark, K. P. and Sundaresan, S. (2019). Inattention and belief polarization. *Journal of Economic Theory*, 180:203–228.

Pomatto, L., Strack, P., and Tamuz, O. (2023). The cost of information: The case of constant marginal costs. *American Economic Review*.

Prendergast, C. (1993). A theory of "Yes Men". *The American Economic Review*, 83(4):757–770.

Rantakari, H. (2008). Governing Adaptation. *The Review of Economic Studies*, 75(4):1257–1285.

Ravid, D. (2020). Ultimatum bargaining with rational inattention. *American Economic Review*, 110(9):2948–63.

Rockafellar, R. T. (1970). *Convex analysis*. Number 28. Princeton university press.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *The quarterly journal of economics*, 100(4):1169–1189.

Segendorff, B. (1998). Delegation and Threat in Bargaining. *Games and Economic Behavior*, 23(2):266–283.

Shore, L. M., Chung-Herrera, B. G., Dean, M. A., Ehrhart, K. H., Jung, D. I., Randel, A. E., and Singh, G. (2009). Diversity in organizations: Where are we now and where are we going? *Human resource management review*, 19(2):117–133.

Steiner, J. and Stewart, C. (2016). Perceiving Prospects Properly. *American Economic Review*, 106(7):1601–1631.

Stepanov, S. (2020). Biased performance evaluation in a model of career concerns: incentives versus ex-post optimality. *Journal of Economic Behavior & Organization*, 179:589–607.

Szalay, D. (2005). The economics of clear advice and extreme options. *The Review of Economic Studies*, 72(4):1173–1198.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323.

Woodford, M. (2012). Inattentive valuation and reference-dependent choice. Working paper.

A Main Proofs

A.1 Proof of Proposition 1

Throughout this proof, we will refer to the delegation rule under consideration,

$$\mu^* = \frac{\sqrt{\mu_p}}{\sqrt{\mu_p} + \sqrt{1-\mu_p}},$$

as the **candidate rule**. It is straightforward that under the candidate rule, if $\mu_p > \frac{1}{2}$ then $\mu^* \in \left(\frac{1}{2}, \mu_p\right)$, since

$$\frac{\mu^*}{1-\mu^*} = \frac{\sqrt{\mu_p}}{\sqrt{1-\mu_p}} < \frac{\mu_p}{1-\mu_p}$$

when $\mu_p > \frac{1}{2}$, so $\mu^* < \mu_p$, and also $\sqrt{\mu_p} > \sqrt{1-\mu_p}$ in that case, so $\mu^* > \frac{1}{2}$. It thus remains to show that the candidate rule is indeed optimal for the principal. While a shorter proof exists that invokes Lemma 1 that derives an optimal strategy for the case of N states and actions, we choose to present a more direct, albeit a somewhat longer, proof.

Plugging the solution to the agent's problem (10) (assuming this solution is interior for now) into the principal's problem (11), we get that the principal's payoff looks as follows:

$$\begin{aligned}
\mu_p\pi(R|r) + (1-\mu_p)\pi(L|l) &= \mu_p \frac{\left(\mu e^{\frac{1}{\lambda}} - (1-\mu)\right) e^{\frac{1}{\lambda}}}{\left(e^{\frac{2}{\lambda}} - 1\right)\mu} + (1-\mu_p)\frac{\left((1-\mu)e^{\frac{1}{\lambda}} - \mu\right) e^{\frac{1}{\lambda}}}{\left(e^{\frac{2}{\lambda}} - 1\right)(1-\mu)} \\
&= \frac{e^{\frac{1}{\lambda}}}{e^{\frac{2}{\lambda}} - 1}\left[\mu_p\left(e^{\frac{1}{\lambda}} - \frac{1-\mu}{\mu}\right) + (1-\mu_p)\left(e^{\frac{1}{\lambda}} - \frac{\mu}{1-\mu}\right)\right] \\
&\propto e^{\frac{1}{\lambda}} - \mu_p\frac{1-\mu}{\mu} - (1-\mu_p)\frac{\mu}{1-\mu}.
\end{aligned}$$

The FOC for the principal's maximization problem above w.r.t. μ is

$$\begin{aligned}
&\frac{\mu_p}{\mu^2} - \frac{1-\mu_p}{(1-\mu_p)^2} = 0 \\
\iff &\frac{\mu}{1-\mu} = \frac{\sqrt{\mu_p}}{\sqrt{1-\mu_p}}.
\end{aligned} \tag{29}$$

It is trivial to verify that the second-order condition holds as well, hence as long as (29) yields an interior solution (i.e., the probabilities in (10) are in $[0,1]$), the candidate solution is indeed optimal among all such interior solutions.

We now check for which μ the solution (10) is interior. Using the expressions (10), one can easily verify that $\pi(R|r) \geq 0 \iff \frac{\mu}{1-\mu} \geq e^{-\frac{1}{\lambda}}$ and $\pi(R|r) \leq 1 \iff \frac{\mu}{1-\mu} \leq e^{\frac{1}{\lambda}}$, and the conditions $\pi(L|l) \in [0,1]$ yield the same two interiority conditions. This implies that if $\frac{\mu}{1-\mu} \in \left[e^{-\frac{1}{\lambda}}, e^{\frac{1}{\lambda}}\right]$, then the agent acquires some information and selects both actions with positive probabilities, and otherwise $(\pi(R|r), \pi(L|l)) \in \{(1,0),(0,1)\}$, meaning that the agent simply chooses the ex ante optimal action for sure without acquiring any information about the state.

The candidate rule then suggests that the principal delegates to a learning agent iff $\frac{\mu_p}{1-\mu_p} \in \left[e^{-\frac{2}{\lambda}}, e^{\frac{2}{\lambda}}\right]$, and otherwise delegates to an agent who plays the ex ante optimal action. We have shown that the candidate rule selects the optimal among the learning agents; it is left to verify that such a criterion for choosing between learning and non-learning agents is optimal for the principal.

Consider $\mu_p \geq \frac{1}{2}$; then among the non-learning agents, the principal would obviously choose the one who plays $a = R$ (rather than $a = L$), and such a choice yields the principal expected payoff $\mu_p \cdot 1 + (1-\mu_p) \cdot 0 = \mu_p$. Optimal delegation to a learning agent yields (by plugging the candidate rule into the principal's payoff obtained above)

$$\frac{e^{\frac{1}{\lambda}}}{e^{\frac{2}{\lambda}}-1}\left[e^{\frac{1}{\lambda}} - \mu_p \frac{1-\mu^*}{\mu^*} - (1-\mu_p)\frac{\mu^*}{1-\mu^*}\right] = \frac{e^{\frac{1}{\lambda}}}{e^{\frac{2}{\lambda}}-1}\left[e^{\frac{1}{\lambda}} - 2\sqrt{\mu_p(1-\mu_p)}\right]. \quad (30)$$

Taking the difference between (30) and μ_p, the payoff from delegating to a non-learning agent, let us find belief μ_p of a principal who would be indifferent between the two:

$$\begin{aligned} &\frac{e^{\frac{1}{\lambda}}}{e^{\frac{2}{\lambda}}-1}\left[e^{\frac{1}{\lambda}} - 2\sqrt{\mu_p(1-\mu_p)}\right] - \mu_p = 0 \\ \iff\ & e^{\frac{2}{\lambda}} - 2e^{\frac{1}{\lambda}}\sqrt{\mu_p(1-\mu_p)} = \mu_p e^{\frac{2}{\lambda}} - \mu_p \\ \iff\ & \left(e^{\frac{1}{\lambda}}\sqrt{1-\mu_p} - \sqrt{\mu_p}\right)^2 = 0 \\ \iff\ & \frac{\sqrt{\mu_p}}{\sqrt{1-\mu_p}} = e^{\frac{1}{\lambda}}. \end{aligned}$$

Hence, the principal prefers a learning agent when $\frac{\sqrt{\mu_p}}{\sqrt{1-\mu_p}} < e^{\frac{1}{\lambda}}$ and a non-learning agent when $\frac{\sqrt{\mu_p}}{\sqrt{1-\mu_p}} > e^{\frac{1}{\lambda}}$. Therefore, the candidate rule is indeed optimal for $\mu_p \geq \frac{1}{2}$. A mirror argument can be used to establish optimality for $\mu_p \leq \frac{1}{2}$. This concludes the proof of Proposition 1.

A.2 Proof of Corollary 1

Proof of Proposition 1 shows that $\mu^*(\mu_p) < \mu_p$ for all $\mu_p \in (0.5, 1)$, hence we can ignore the absolute value operator. Using expression (12) we then obtain

$$\frac{d}{d\mu_p}\left(\mu_p - \mu^*(\mu_p)\right) = \frac{4\mu_p(1-\mu_p) + 2\sqrt{\mu_p(1-\mu_p)} - 1}{2\sqrt{\mu_p(1-\mu_p)} \cdot \left(\sqrt{\mu_p} + \sqrt{1-\mu_p}\right)^2},$$

where the denominator is weakly positive for all $\mu_p \in (0.5, 1)$, and the numerator is positive if and only if $\sqrt{\mu_p(1-\mu_p)} \notin \left(\frac{-1-\sqrt{5}}{4}, \frac{-1+\sqrt{5}}{4}\right)$, which is equivalent to $\mu_p \leq \frac{1}{2} + \frac{1}{2} \cdot \sqrt{\frac{\sqrt{5}-1}{2}} \approx 0.893$. Then $|\mu_p - \mu^*(\mu_p)|$ is increasing for these values of μ_p and decreasing otherwise, meaning it satisfies single-peakedness.

A.3 Proof of Lemma 1

The goal is to find the optimal choice probabilities $\beta^* \in \Delta(\mathcal{A})$ which maximize the principal's expected utility (16). First, let us rewrite expression (16) using $\delta \equiv e^{\frac{1}{\lambda}} - 1$:

$$\sum_{j=1}^{N} \mu_p(\omega_j) \frac{\beta(a_j) e^{\frac{1}{\lambda}}}{1+\delta\beta(a_j)} = \sum_{j\in C(\beta)} e^{\frac{1}{\lambda}} \frac{\frac{\mu_p(\omega_j)}{\delta}(1+\delta\beta(a_j)) - \frac{\mu_p(\omega_j)}{\delta}}{1+\delta\beta(a_j)}$$
$$= \sum_{j\in C(\beta)} e^{\frac{1}{\lambda}} \left(\frac{\mu_p(\omega_j)}{\delta} - \frac{\mu_p(\omega_j)}{\delta\,(1+\delta\beta(a_j))} \right).$$

The first term in the brackets above is independent of β, so the principal's maximization problem is equivalent to

$$\min_{\beta} \sum_{j\in C(\beta)} \frac{\mu_p(\omega_j)}{1+\delta\beta(a_j)}. \tag{31}$$

Let ξ denote the Lagrange multiplier corresponding to the constraint $\sum_{j=1}^{N} \beta(a_j) = 1$. Then the first-order condition for $\beta(a_i)$ with $i \in C(\beta)$ is

$$(1+\delta\beta(a_i))^2 = -\frac{\mu_p(\omega_i)}{\xi}. \tag{32}$$

Summing up these equalities over all $j \in C(\beta)$, we get that

$$\sum_{j\in C(\beta)} (1+\delta\beta(a_j))^2 = -\frac{\sum_{j\in C(\beta)} \mu_p(\omega_j)}{\xi} \tag{33}$$

Combining (32) and (33):

$$1+\delta\beta(a_i) = \frac{\sqrt{\mu_p(\omega_i)}}{\sqrt{\sum_{j\in C(\beta)} \mu_p(\omega_j)}} \sqrt{\sum_{j\in C(\beta)} (1+\delta\beta(a_j))^2} \tag{34}$$

Once again summing up these equalities over all $j \in C(\beta)$, we get that

$$K(\beta) + \delta = \frac{\sum_{j\in C(\beta)} \sqrt{\mu_p(\omega_j)}}{\sqrt{\sum_{j\in C(\beta)} \mu_p(\omega_j)}} \sqrt{\sum_{j\in C(\beta)} (1+\delta\beta(a_j))^2}.$$

Expressing $\sqrt{\sum_{j\in C(\beta)} (1+\delta\beta(a_j))^2}$ from this expression and plugging it into (34) allows us to express $\beta(a_i)$ (for $i \in C(\beta)$) in closed form as

$$\beta(a_i) = \frac{1}{\delta} \left(\frac{(K(\beta)+\delta)\sqrt{\mu_p(\omega_i)}}{\sum_{j\in C(\beta)} \sqrt{\mu_p(\omega_j)}} - 1 \right). \tag{35}$$

The necessary condition for option i to be in a consideration set ($i \in C(\beta)$) is $\beta(a_i) \geq 0$ or, equivalently,

$$\sqrt{\mu_p(\omega_i)} > \frac{1}{K(\beta)+\delta} \sum_{j\in C(\beta)} \sqrt{\mu_p(\omega_j)}.$$

Now let ξ_k denote the Lagrange multiplier for the constraint $\beta(a_k) \geq 0$. Then the first-order condition for an alternative $k \notin C(\beta)$ that is not chosen is

$$\mu_p(\omega_k) = -\xi - \xi_k \qquad \Rightarrow \qquad \mu_p(\omega_k) \leq -\xi.$$

Plugging in ξ from (32) into the inequality above yields

$$\mu_p(\omega_k) \leq \frac{\sum_{j\in C(\beta)} \mu_p(\omega_j)}{\sum_{j\in C(\beta)}(1+\delta\beta(a_j))^2} \quad \Leftrightarrow \quad \sqrt{\mu_p(\omega_k)} \leq \frac{1}{K(\beta)+\delta} \sum_{j\in C(\beta)} \sqrt{\mu_p(\omega_j)}$$

for all $k \notin C(\beta)$.

Since the minimization problem has a convex objective function and linear constraints, the Kuhn-Tucker conditions are necessary and sufficient. Thus the necessary and sufficient conditions that the solution β^* must satisfy are given by:

$$\begin{cases} \sqrt{\mu_p(\omega_i)} > \frac{1}{K(\beta^*)+\delta} \sum\limits_{j\in C(\beta^*)} \sqrt{\mu_p(\omega_j)} & \text{for all } i \in C(\beta^*), \\ \sqrt{\mu_p(\omega_k)} \leq \frac{1}{K(\beta^*)+\delta} \sum\limits_{j\in C(\beta^*)} \sqrt{\mu_p(\omega_j)} & \text{for all } k \notin C(\beta^*). \end{cases}$$

Recall that we assumed, without loss of generality, that $\mu_p(\omega_1) \geq \mu_p(\omega_2) \geq \ldots \geq \mu_p(\omega_N)$. Suppose that the solution β^* is such that $K(\beta^*) = K'$. Clearly then, in the optimum, the consideration set $C(\beta^*)$ will consist of the first K' alternatives.

Denote $\Delta_L \equiv (L+\delta)\sqrt{\mu_p(\omega_L)} - \sum\limits_{j=1}^{L} \sqrt{\mu_p(\omega_j)}$. Notice that for all $L > 1$:

$$\begin{aligned} \Delta_L \equiv& (L+\delta)\sqrt{\mu_p(\omega_L)} - \sum_{j=1}^{L} \sqrt{\mu_p(\omega_j)} \\ =& (L-1+\delta)\sqrt{\mu_p(\omega_{L-1})} - \sum_{j=1}^{L-1} \sqrt{\mu_p(\omega_j)} - \sqrt{\mu_p(\omega_L)} \\ & - (L-1+\delta)\sqrt{\mu_p(\omega_{L-1})} + (L+\delta)\sqrt{\mu_p(\omega_L)} \\ =& \Delta_{L-1} - (L-1+\delta)\left(\sqrt{\mu_p(\omega_{L-1})} - \sqrt{\mu_p(\omega_L)}\right). \end{aligned}$$

Therefore, Δ_L decreases in L. Since $\Delta_1 > 0$, there either exists unique K' such that $\Delta_{K'} > 0$ and $\Delta_{K'+1} \leq 0$, or $\Delta_L > 0$ for all L. In the former case, $K(\beta^*) = K'$, and in the latter case, $K(\beta^*) = N$.

In the end, the solution to the principal's problem is given by $\beta^*(a_i)$ as in (35) if $i \in C(\beta^*)$, $\beta^*(a_i) = 0$ if $i \notin C(\beta^*)$, and $C(\beta^*) = 1, ..., K(\beta^*)$, where $K(\beta^*)$ is as described above.

A.4 Proof of Lemma 2

Corollary 2 from Matějka and McKay (2015) shows that a vector of the unconditional choice probabilities $\beta \in \Delta(\mathcal{A})$ solves (13) only if it solves the system of equations given by

$$\sum_{j=1}^{N} \mu(\omega_j) \frac{e^{\frac{u(a_i,\omega_j)}{\lambda}}}{\sum_{k=1}^{N} \beta(a_k) e^{\frac{u(a_k,\omega_j)}{\lambda}}} = 1, \tag{36}$$

for every $i \in \{1, ..., N\}$ such that $\beta(a_i) > 0$.

The question then is: given a vector $\beta \in \Delta(\mathcal{A})$ of unconditional choice probabilities, can we find $\mu \in \mathbb{R}_+^N$ that solves the following system:

$$\begin{cases} \mu(\omega_1) + \mu(\omega_2) + \ldots + \mu(\omega_N) = 1, \\ \sum\limits_{j=1}^{N} \mu(\omega_j) \frac{e^{\frac{u(a_i,\omega_j)}{\lambda}}}{\sum\limits_{k=1}^{N} \beta(a_k) e^{\frac{u(a_k,\omega_j)}{\lambda}}} = 1 \qquad \forall i \in C(\beta). \end{cases} \tag{37}$$

The system above is a linear system of $K(\beta) + 1$ equations with N unknowns. To prove the solution exists, we use the Farkas' lemma (Aliprantis and Border, 2006, Corollary 5.85). It states that given some matrix $A \in \mathbb{R}^{m \times n}$ and a vector $b \in \mathbb{R}^m$, the linear system $Ax = b$ has a non-negative root $x \in \mathbb{R}_+^n$ if and only if there exists no vector $y \in \mathbb{R}^m$ such that $A'y \geq 0$ with $b'y < 0$. The two latter inequalities applied to our case form the following system:

$$\begin{cases} y_0 \left(\sum\limits_{k=1}^{N} \beta(a_k) e^{\frac{u(a_k,\omega_j)}{\lambda}} \right) + \left(\sum\limits_{i \in C(\beta)} y_i e^{\frac{u(a_i,\omega_j)}{\lambda}} \right) \geq 0 \quad \forall j \in \{1, ..., N\}, \\ y_0 + \sum\limits_{i \in C(\beta)} y_i < 0. \end{cases} \tag{38}$$

We need to show there exists no $y \in \mathbb{R}^{K(\beta)+1}$ that solves the system above. Let us define $z_i \equiv y_i + y_0 \beta(a_i)$ for $i \in C(\beta)$. Then, recalling that $e^{\frac{u(a_i,\omega_i)}{\lambda}} = e^{\frac{1}{\lambda}}$ and $e^{\frac{u(a_i,\omega_j)}{\lambda}} = 1$ for $i \neq j$, system (38) transforms to

$$\begin{cases} z_j e^{\frac{1}{\lambda}} + \sum\limits_{i \in C(\beta) \setminus \{j\}} z_i \geq 0 \quad \forall j \in C(\beta), \\ \sum\limits_{i \in C(\beta)} z_i \geq 0 \qquad\qquad\qquad \forall j \in \{1, ..., N\} \setminus C(\beta), \\ \sum\limits_{i \in C(\beta)} z_i < 0. \end{cases} \tag{39}$$

System (39) above does not have a solution. Indeed, if $C(\beta) \subsetneq \{1, ..., N\}$, then the middle set of inequalities directly contradicts the latter inequality. If $C(\beta) = \{1, ..., N\}$, then subtracting the latter inequality from the former, for a given $j \in C(\beta)$, yields $z_j \delta \geq 0 \iff z_j \geq 0$. Since this must hold for all $j \in C(\beta)$, we obtain a contradiction with the latter inequality, $\sum\limits_{i \in C(\beta)} z_i < 0$.

By the Farkas' lemma, we then conclude that for any vector $\beta \in \Delta(\mathcal{A})$ there exists a belief $\mu \in \Delta(\Omega)$ that solves system (37). This concludes the proof.

A.5 Proof of Theorem 1

This proof proceeds in two parts. First, we show that the delegation strategy introduced in the proposition (hereinafter referred to as "the candidate strategy") is optimal for the principal. Then we establish that it does indeed possess the stated properties.

Consider an agent with a prior belief

$$\mu(\omega_i) = \frac{\sqrt{\mu_p(\omega_i)}}{\sum_{j=1}^{N} \sqrt{\mu_p(\omega_j)}}. \tag{40}$$

It is trivial to verify that prior belief μ defined this way satisfies the candidate strategy in the statement of the proposition, and hence represents the candidate strategy. Consider an agent hired in accordance with the candidate rule. Substituting (40) into (14) yields

$$\beta(a_i) = \max\left\{0, \frac{1}{\delta}\left(\frac{(K(\beta^*)+\delta)\sqrt{\mu_p(\omega_i)}}{\sum_{j\in C(\beta^*)} \sqrt{\mu_p(\omega_j)}} - 1\right)\right\}, \tag{41}$$

which are exactly the probabilities stated in Lemma 1. Therefore, an agent hired according to the candidate strategy makes decisions in such a way that generates the principal-optimal unconditional choice probabilities. Therefore, delegation according to the candidate strategy is indeed optimal for the principal.

Now we show that the candidate strategy satisfies the properties stated in the proposition. **The first property** follows clearly from (40): $\mu^*(\omega_1) \geq \mu^*(\omega_2) \geq \ldots \geq \mu^*(\omega_N)$.

For **the second property**, we want to show that $\mu^*(\omega_1) \leq \mu_p(\omega_1)$ and $\mu^*(\omega_N) \geq \mu_p(\omega_N)$. The former inequality can be shown as follows:

$$\begin{aligned}
&\mu^*(\omega_1) \leq \mu_p(\omega_1)\\
\iff &\frac{\sqrt{\mu_p(\omega_i)}}{\sum_{j=1}^{N} \sqrt{\mu_p(\omega_j)}} \leq \mu_p(\omega_1)\\
\iff &1 \leq \sqrt{\mu_p(\omega_1)} \cdot \left(\sum_{j=1}^{N} \sqrt{\mu_p(\omega_j)}\right)\\
\iff &1 \leq \mu_p(\omega_1) + \sqrt{\mu_p(\omega_1)\mu_p(\omega_2)} + \ldots + \sqrt{\mu_p(\omega_1)\mu_p(\omega_N)},
\end{aligned}$$

and the latter inequality holds because $\mu_p(\omega_1)+\ldots+\mu_p(\omega_N) = 1$ and $\sqrt{\mu_p(\omega_1)\mu_p(\omega_j)} \geq \mu_p(\omega_j)$ for all $j \in \{1, \ldots, N\}$, since $\mu_p(\omega_1) \geq \mu_p(\omega_j)$. Note that $\mu^*(\omega_1) = \mu_p(\omega_1)$ only if $\mu_p(\omega_1) = \ldots = \mu_p(\omega_N)$.

Similarly, the inequality $\mu^*(\omega_N) \geq \mu_p(\omega_N)$ is equivalent to

$$1 \geq \sqrt{\mu_p(\omega_1)\mu_p(\omega_N)} + \ldots + \sqrt{\mu_p(\omega_{N-1})\mu_p(\omega_N)} + \mu_p(\omega_N),$$

which holds because $\sqrt{\mu_p(\omega_j)\mu_p(\omega_N)} \leq \mu_p(\omega_j)$ for all $j \in \{1, \ldots, N\}$, with equalities only if $\mu_p(\omega_1) = \ldots = \mu_p(\omega_N)$.

We now proceed to show **the third property** (majorization). We want to show that for all $k \in \{1, ..., N\}$,

$$\sum_{i=1}^{k} \mu^*(\omega_i) \leq \sum_{i=1}^{k} \mu_p(\omega_i).$$

Using (40), this inequality can be rewritten as

$$\frac{\sum_{i=1}^{k} \sqrt{\mu_p(\omega_i)}}{\sum_{j=1}^{N} \sqrt{\mu_p(\omega_j)}} \leq \sum_{l=1}^{k} \mu_p(\omega_l),$$

$$\iff 1 \leq \frac{\sum_{l=1}^{k} \mu_p(\omega_l) \sum_{j=1}^{N} \sqrt{\mu_p(\omega_j)}}{\sum_{i=1}^{k} \sqrt{\mu_p(\omega_i)}}$$

$$\iff 1 \leq \sum_{l=1}^{k} \mu_p(\omega_i) + \sum_{j=k+1}^{N} \sqrt{\mu_p(\omega_j)} \cdot \rho, \tag{42}$$

where $\rho \equiv \frac{\sum_{l=1}^{k} \mu_p(\omega_l)}{\sum_{i=1}^{k} \sqrt{\mu_p(\omega_i)}}$. Since $\sum_{l=1}^{N} \mu_p(\omega_i) = 1$, to verify that (42) holds, it is enough to show that $\rho \geq \sqrt{\mu_p(\omega_{k+1})}$ (since $\sqrt{\mu_p(\omega_{k+1})} \geq ... \geq \sqrt{\mu_p(\omega_N)}$), with the former inequality being equivalent to

$$\frac{\sum_{l=1}^{k} \mu_p(\omega_l)}{\sum_{i=1}^{k} \sqrt{\mu_p(\omega_i)}} \geq \sqrt{\mu_p(\omega_{k+1})}$$

$$\iff \mu_p(\omega_1) + ... + \mu_p(\omega_k) \geq \sqrt{\mu_p(\omega_1)\mu_p(\omega_{k+1})} + ... + \sqrt{\mu_p(\omega_k)\mu_p(\omega_{k+1})},$$

which is true because $\mu_p(\omega_{k+1}) \leq \mu_p(\omega_k) \leq ... \leq \mu_p(\omega_1)$. We conclude that μ_p majorizes μ^*.

To show **the fourth property**, observe that the Shannon entropy for some probability vector μ is given by

$$H(p) \equiv -\sum_{i=1}^{N} \mu(\omega_i) \cdot \log \mu(\omega_i).$$

Marshall et al. (1979, p.101) show that $H(p)$ is a Schur-concave function (i.e., it is decreasing w.r.t. the majorization order). It therefore holds that if μ_p majorizes μ^*, then $H(\mu_p) \leq H(\mu^*)$. This concludes the proof of Theorem 1.

A.6 Proof of Proposition 2

It follows from (18) that the size of the consideration set in the aligned problem, $\bar{K}$, is such that

$$\sum_{j=1}^{\bar{K}} \frac{\mu_p(\omega_j)}{\mu_p(\omega_{\bar{K}})} < \bar{K} + \delta \leq \sum_{j=1}^{\bar{K}} \frac{\mu_p(\omega_j)}{\mu_p(\omega_{\bar{K}+1})}$$

Since $\frac{\mu_p(\omega_i)}{\mu_p(\omega_{\bar{K}})} > 1$ for all $i < \bar{K}$, we have that $\frac{\mu_p(\omega_i)}{\mu_p(\omega_{\bar{K}})} > \frac{\sqrt{\mu_p(\omega_i)}}{\sqrt{\mu_p(\omega_{\bar{K}})}} > 1$ holds for all $i < K$. Therefore,

$$\sum_{j=1}^{\bar{K}} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{\bar{K}})}} < \bar{K} + \delta. \tag{43}$$

From (19), K^* is the unique solution of

$$\sum_{j=1}^{K^*} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{K^*})}} < K^* + \delta \leq \sum_{j=1}^{K^*} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{K^*+1})}}. \tag{44}$$

Two cases are possible, depending on whether

$$\bar{K} + \delta \gtreqless \sum_{j=1}^{\bar{K}} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{\bar{K}+1})}}. \tag{45}$$

If $\bar{K} + \delta \leq$ RHS in (45) (where RHS refers to the right-hand side), then together with (43) this implies that $\bar{K}$ solves (44), and thus $\bar{K} = K^*$, which satisfies that statement of the proposition.

If, however, $\bar{K} + \delta >$ RHS in (45), then $\bar{K}$ does not solve (44). In this case, note that going from K by $K + 1$ increases the LHS of (45) by 1 and increases the RHS by the amount strictly greater than 1, since a new term $\frac{\sqrt{\mu_p(\omega_{K+1})}}{\sqrt{\mu_p(\omega_{K+2})}} > 1$ is added to the sum, and all existing terms increase because $\mu_p(\omega_{K+1}) < \mu_p(\omega_K)$. This holds for all K, meaning that if $\bar{K} + \delta >$ RHS in (45), then $K + \delta > \sum_{j=1}^{K} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{K+1})}}$ for all $K < \bar{K}$. Therefore, the unique solution K^* of (44) must be such that $K_M > \bar{K}$. This concludes the proof.

A.7 Proof of Theorem 2

Part 2 of the statement follows immediately from Proposition 3 of Matveenko and Mikhalishchev (2021).

To show part 1, we invoke Theorem 1 from Matějka and McKay (2015) stated in (13), which claims that in the contracting problem, the $\beta : \Omega \to \Delta(\mathcal{A})$ that solves the agent's problem (20) is given by

$$\begin{aligned}\pi(a_i|\omega_j) &= \frac{\beta(a_i)e^{\frac{u(a_i,\omega_j)+\tau(a_i)}{\lambda}}}{\sum_{k=1}^N \beta(a_k)e^{\frac{u(a_k,\omega_j)+\tau(a_k)}{\lambda}}} \\ &= \frac{\beta'(a_i)e^{\frac{u(a_i,\omega_j)}{\lambda}}}{\sum_{k=1}^N \beta'(a_k)e^{\frac{u(a_k,\omega_j)}{\lambda}}}, \\ \text{where } \beta(a_i) &= \sum_{j=1}^N \mu(\omega_j)\pi(a_i|\omega_j). \\ \text{and } \beta'(a_i) &\equiv \frac{\beta(a_i)e^{\frac{\tau(a_i)}{\lambda}}}{\sum_{k=1}^N \beta(a_k)e^{\frac{\tau(a_k)}{\lambda}}}.\end{aligned} \tag{46}$$

Since β' is a valid probability distribution on $\mathcal{A}$, representation (46) together with (13) imply that such a collection of conditional probabilities π is a valid solution to the agent's problem (4) when the agent's preferences net of information costs are given by $u(a_i, \omega_j)$. That is, the principal can implement the desired conditional choice probabilities π by choosing an agent with unbiased preferences and some belief μ, such that the unconditional choice probabilities selected by this agent are given by β'. Lemma 2 implies that such a belief $\mu \in \Delta(\Omega)$ does indeed exist.

A.8 Proof of Proposition 3

Plugging (12) into (10) yields the principal-optimal conditional choice probabilities for the binary model, given by

$$\begin{aligned}\pi^*(R|r) &= \left(e^{\frac{2}{\lambda}} - 1\right)^{-1} e^{\frac{1}{\lambda}} \left(e^{\frac{1}{\lambda}} - \sqrt{\frac{1-\mu_p}{\mu_p}}\right), \\ \pi^*(L|l) &= \left(e^{\frac{2}{\lambda}} - 1\right)^{-1} e^{\frac{1}{\lambda}} \left(e^{\frac{1}{\lambda}} - \sqrt{\frac{\mu_p}{1-\mu_p}}\right),\end{aligned} \tag{47}$$

cropped to $[0, 1]$.

The agent's preferences only depend on the difference $\tau(R) - \tau(L)$. Assuming all $\tau(R) \in \mathbb{R}$ are available to the principal (no limited liability), it is without loss to set $\tau(L) = 0$. The agent's problem is given by (20). Solving it given $\tau = (\tau(R), 0)$ yields

$$\begin{aligned}\pi(R|r) &= 1 - \frac{e^{\frac{2}{\lambda}}(1-\mu) - e^{\frac{1+\tau(R)}{\lambda}} + \mu}{\left(e^{\frac{2}{\lambda}} - 1\right)\left(e^{\frac{1+\tau(R)}{\lambda}} - 1\right)\mu}, \\ \pi(L|l) &= \frac{e^{\frac{1}{\lambda}}\left(e^{\frac{2}{\lambda}}(1-\mu) - e^{\frac{1+\tau(R)}{\lambda}} + \mu\right)}{\left(e^{\frac{2}{\lambda}} - 1\right)\left(e^{\frac{1}{\lambda}} - e^{\frac{\tau(R)}{\lambda}}\right)(1-\mu)},\end{aligned} \tag{48}$$

cropped to $[0, 1]$.

The principal's contracting problem (21) in the binary setting with $\rho = 0$ is similar to (11):

$$\begin{aligned} &\max_{\tau(R)} \{\mu_p \pi(R|r) + (1-\mu_p)\, \pi(L|l)\} \\ &\text{s.t. } \pi(R|r), \pi(L|l) \text{ are given by (48).} \end{aligned} \tag{49}$$

Assuming the probabilities in (48) are interior, the F.O.C. for (49) yields the candidate solution $\tau(R)$ given by

$$\tau^*(R) = \lambda \ln \left[\frac{\frac{1-\mu}{\mu} e^{\frac{1}{\lambda}} + \sqrt{\frac{1-\mu_p}{\mu_p}}}{\frac{1-\mu}{\mu} + e^{\frac{1}{\lambda}} \sqrt{\frac{1-\mu_p}{\mu_p}}} \right], \tag{50}$$

where the expression under the $\ln(\cdot)$ is non-negative for any λ, μ_p, μ, and thus the candidate $\tau(R)$ exists for any μ that yields interior probabilities (48).

Plugging (50) into (48) yields, after some routine manipulations, the conditional choice probabilities that coincide with (47) (hence, probabilities (48) are interior given μ and $\tau^*(R)$ if and only if probabilities (47) are interior). Thus, condition (50) is not only necessary, but also sufficient. Hence, for any μ_p for which (47) are interior, $\tau^*(R)$ as given by (50) solves the principal's problem (49), and this solution exists for any μ.

If λ and μ_p are such that probabilities (47) are not interior, then the principal would like the agent to take the ex ante (principal-)preferred action (it can be verified that the expressions in (47) are such that $\pi^*(R|r) \geq 1 \iff \pi^*(L|l) \leq 0$ and $\pi^*(R|r) \leq 0 \iff \pi^*(L|l) \geq 1$). The candidate transfers (50) yield exactly such non-interior probabilities (when plugged into (48)), and hence they still solve the principal's problem (49) for any respective μ.[33] This concludes the proof of part 1 of the proposition.

To show part 2, consider (50) as a function of μ. It is strictly decreasing in μ on $[0, 1]$, and the equation $\tau^*(R)(\mu) = 0$ has a unique root in $[0, 1]$ equal to

$$\mu^* = \frac{\sqrt{\mu_p}}{\sqrt{\mu_p} + \sqrt{1-\mu_p}},$$

meaning that $\tau(R) \geq 0 = \tau(L) \iff \mu \leq \mu^*$.

A.9 Proof of Proposition 4

As argued in the text, it is immediate that $\underline{\tau} = 0$. Proceeding analogously to Section 3, we obtain that the agent's problem (22) given the incentive payment $\bar{\tau} \geq 0$ is solved by

$$\pi(R|r) = \min\{1, \max\{0, \pi_u(R|r)\}\} \text{ and } \pi(L|l) = \min\{1, \max\{0, \pi_u(L|l)\}\},$$

$$\begin{aligned} \text{where } \pi_u(R|r) &= \frac{e^{\frac{1+\bar{\tau}}{\lambda}} \left(e^{\frac{1+\bar{\tau}}{\lambda}} \mu - (1-\mu) \right)}{\left(e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1 \right) \mu} = \frac{e^{\frac{1+\bar{\tau}}{\lambda}}}{e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1} \left(e^{\frac{1+\bar{\tau}}{\lambda}} - \frac{1-\mu}{\mu} \right), \\ \pi_u(L|l) &= \frac{e^{\frac{1+\bar{\tau}}{\lambda}} \left(e^{\frac{1+\bar{\tau}}{\lambda}} (1-\mu) - \mu \right)}{\left(e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1 \right) (1-\mu)} = \frac{e^{\frac{1+\bar{\tau}}{\lambda}}}{e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1} \left(e^{\frac{1+\bar{\tau}}{\lambda}} - \frac{\mu}{1-\mu} \right). \end{aligned} \tag{51}$$

[33] Note that $\tau^*(R)$ is not the unique solution in this case. If $\pi^*(R|r) = 1, \pi^*(L|l) = 0$, then any $\tau(R) \geq \lambda \ln \left(\mu + (1-\mu) e^{\frac{2}{\lambda}} \right) - 1$ yields the optimal choice probabilities, and if $\pi^*(R|r) = 0, \pi^*(L|l) = 1$, then any $\tau(R) \leq 1 - \lambda \ln \left(\mu e^{\frac{2}{\lambda}} + (1-\mu) \right)$ solves the principal's problem.

The principal's full contracting problem (23) can be rewritten as

$$\begin{aligned}&\max_{\bar{\tau}\in\mathbb{R}_+}\left\{(1-\rho\bar{\tau})\left(\mu_p\pi(R|r)+(1-\mu_p)\pi(L|l)\right)\right\},\\&\text{ s.t. } \pi(R|r),\pi(L|l) \text{ are given by (51).}\end{aligned}\tag{52}$$

We use τ^* to denote the solution to this problem.

To begin with, note that $\tau^* \geq 0$ (due to limited liability) and $\tau^* < 1/\rho$ (otherwise the principal's payoff is zero or negative, hence such τ^* are dominated by $\bar{\tau} = 0$). Further, if $\tau^* > 0$, then $\pi = \pi_u$, since otherwise the principal could reduce $\bar{\tau}$ without affecting the agent's choice.

Let us define the principal's **relaxed contracting problem** as

$$\begin{aligned}&\max_{\bar{\tau}\in\mathbb{R}}\left\{(1-\rho\bar{\tau})\left(\mu_p\pi_u(R|r)+(1-\mu_p)\pi_u(L|l)\right)\right\},\\&\text{s.t. } \pi_u(R|r),\pi_u(L|l) \text{ are given by (51).}\end{aligned}\tag{53}$$

It differs from the full problem (52) in that it ignores the constraints $\bar{\tau} = 0$ and $\pi(R|r), \pi(L|l) \in [0,1]$. We use τ^{**} to denote the interior solution of this relaxed problem, whenever it exists. So far, we can conclude that the principal's problem (52) is solved by $\tau^* \in \{0, \tau^{**}\}$. The local maximizer τ^{**} is optimal if it satisfies all of the following three properties (and $\tau^* = 0$ otherwise):[34]

Feasibility: τ^{**} exists and $\tau^{**} \in [0, 1/\rho]$.[35]

Effectiveness: τ^{**} generates $\pi = \pi_u$.

Preferability: τ^{**} is preferred to $\bar{\tau} = 0$.

The FOC of problem (53) (that must be solved by τ^{**}) is given by

$$\mu_p\frac{1-\mu}{\mu}+(1-\mu_p)\frac{\mu}{1-\mu}=e^{\frac{1+\bar{\tau}}{\lambda}}\cdot\frac{\lambda\rho\left(e^{2\frac{1+\bar{\tau}}{\lambda}}-1\right)+2(1-\rho\bar{\tau})}{\lambda\rho\left(e^{2\frac{1+\bar{\tau}}{\lambda}}-1\right)+\left(e^{2\frac{1+\bar{\tau}}{\lambda}}+1\right)(1-\rho\bar{\tau})}.\tag{54}$$

Let $\gamma(\mu, \mu_p)$ denote the LHS and $\chi(\bar{\tau})$ the RHS of (54), respectively. Note that $\chi(\bar{\tau})$ is continuous in $\bar{\tau}$ and only depends on $\bar{\tau}$, λ, and ρ, but not on μ or μ_p. Further, Lemma 3 below shows that if $\rho \geq \min\left\{1, \frac{1}{2\lambda}\right\}$ then for all λ, $\chi(\bar{\tau})$ is increasing in $\bar{\tau} \in [0, 1/\rho]$ (recall that $\tau^{**} > 1/\rho$ obviously violates preferability, hence we drop this case). We maintain this restriction on ρ throughout the rest of the proof. Monotonicity implies that a feasible τ^{**} exists for given $\mu, \mu_p, \lambda, \rho$ if and only if $\chi(0) \leq \gamma(\mu, \mu_p) \leq \chi(1/\rho)$, where the "if" part follows from the intermediate value theorem, and the "only if" part follows from $\tau \geq 1/\rho$ never being optimal. The strict monotonicity of $\chi(\bar{\tau})$ also means that the objective function in (52) is strictly concave in $\bar{\tau}$, so if τ^{**} exists, then it is unique and it is a local maximizer of (52).

[34]Feasibility and preferability should be self-explanatory. Effectiveness means that the incentive payment is effective at inducing the agent to acquire a non-trivial amount of information. Note that $\bar{\tau} = 0$ is effective when the agent acquires information in the absence of a transfer.

[35]Note that $\tau^{**} \leq 1/\rho$ is not an exogenous restriction, but is rather implied by preferability, as established previously. It is, however, convenient to include this an explicit restriction.

Lemma 3. *Function* $\chi(\bar{\tau})$ *is continuous and increasing in* $\bar{\tau} \in [0, 1/\rho)$ *for all* λ *and all* $\rho \geq \min\{1, 1/2\lambda\}$.

Proof. Denote $\xi = \xi(\tau, \lambda) \equiv e^{\frac{1+\tau}{\lambda}}$. For sake of brevity we drop the arguments of $\xi(\tau, \lambda)$ and the bar from $\bar{\tau}$ throughout the proof of this lemma. Then we can rewrite

$$\chi(\tau) = \xi \frac{\lambda\rho(\xi^2 - 1) + 2(1 - \rho\tau)}{\lambda\rho(\xi^2 - 1) + (\xi^2 + 1)(1 - \rho\tau)}.$$

This function is trivially continuous and differentiable in $\tau \in [0, 1/\rho]$. Hence it suffices to show that $\frac{d\chi(\tau)}{d\tau} > 0$:

$$\begin{aligned}
\frac{d\chi(\tau)}{d\tau} =& \frac{\left(\lambda\rho(3\xi^2 - 1) + 2(1 - \rho\tau)\right) \frac{\partial\xi}{\partial\tau} - 2\rho\xi}{\lambda\rho(\xi^2 - 1) + (\xi^2 + 1)(1 - \rho\tau)} \\
& - \left[2\xi(\lambda\rho + 1 - \rho\tau)\frac{\partial\xi}{\partial\tau} - \rho(\xi^2 + 1)\right] \cdot \frac{\xi\left[\lambda\rho(\xi^2 - 1) + 2(1 - \rho\tau)\right]}{[\lambda\rho(\xi^2 - 1) + (\xi^2 + 1)(1 - \rho\tau)]^2} \\
=& \frac{\xi\left(\xi^2 - 1\right)}{\lambda} \cdot \frac{2\left(\xi^2 - 1 + 2\frac{1-\rho\tau}{\lambda\rho}\right) + \frac{1-\rho\tau}{\lambda\rho} \cdot \left(\xi^2 - 1 - 2\frac{1-\rho\tau}{\lambda\rho}\right)}{\left[(\xi^2 - 1) + (\xi^2 + 1)\frac{1-\rho\tau}{\lambda\rho}\right]^2}.
\end{aligned}$$

The latter expression is strictly positive for $\tau \in [0, 1/\rho)$ if and only if

$$2\left(\xi^2 - 1 + 2\frac{1 - \rho\tau}{\lambda\rho}\right) + \frac{1 - \rho\tau}{\lambda\rho} \cdot \left(\xi^2 - 1 - 2\frac{1 - \rho\tau}{\lambda\rho}\right) > 0. \tag{55}$$

The first term is strictly positive (since $\xi > 1$ and $\tau \leq 1/\rho$). The second term is nonnegative for the given range of τ if $\xi^2 - 1 \geq 2\frac{1-\rho\tau}{\lambda\rho}$. Note that $\xi^2 - 1 \geq 2\frac{1+\tau}{\lambda}$, hence (55) holds if $\frac{1+\tau}{\lambda} \geq \frac{1-\rho\tau}{\lambda\rho}$ for all $\tau \in [0, 1/\rho)$, which holds if $\rho \geq 1$.

Alternatively, we can rewrite (55) as

$$(\xi^2 - 1)\left(2 + \frac{1 - \rho\tau}{\lambda\rho}\right) + 2\frac{1 - \rho\tau}{\lambda\rho}\left(2 - \frac{1 - \rho\tau}{\lambda\rho}\right) \geq 0.$$

In the above expression, the first term is again always strictly positive; the second term is nonnegative if $\lambda\rho \geq 1/2$ (since $\tau \leq 1/\rho$).

We thus conclude that if either $\rho \geq 1$, or $\rho \geq 1/2\lambda$, then $\frac{d\chi(\tau)}{d\tau} > 0$ for $\tau \in [0, 1/\rho)$, so $\chi(\tau)$ is indeed increasing in τ on that interval. □

Let us define the following cutoffs on μ that will prove helpful in establishing the properties of interest of τ^{**} (feasibility, effectiveness, and preferability):

1. Let $\mu_{L1} \in (0, \mu^*)$ and $\mu_{R1} \in (\mu_p, 1)$ be such that $\gamma(\mu_{L1}, \mu_p) = \gamma(\mu_{R1}, \mu_p) = \chi(0)$. Lemma 4 below establishes that these cutoffs exist.

2. Let $\mu_{L2} \equiv \max\left\{\mu : \pi^0_\mu(L|l) = 1\right\}$, $\mu_{R2} \equiv \min\left\{\mu : \pi^0_\mu(R|r) = 1\right\}$, where $\pi^0_\mu(L|l)$ and $\pi^0_\mu(R|r)$ stand for the respective probabilities (51) given μ and $\bar{\tau} = 0$. In words, μ_{L2} and μ_{R2} are the most extreme beliefs μ for which the agent voluntarily acquires information in the absence of incentive payment. Closed-form expressions can be obtained from (51), with $\mu_{L2} = \frac{1}{1+e^{\frac{1}{\lambda}}}$ and $\mu_{R2} = \frac{e^{\frac{1}{\lambda}}}{1+e^{\frac{1}{\lambda}}}$.

3. Let $\mu_{L3} \equiv \inf\{\mu : \tau^* > 0\}$, $\mu_{R3} \equiv \sup\{\mu : \tau^* > 0\}$. In words, these denote the most extreme beliefs μ up to which the principal is willing to offer incentive contracts. Lemma 6 below shows that μ_{L3} and μ_{R3} are always well-defined (i.e., that the set $\{\mu : \tau^* > 0\}$ is nonempty).

4. Let $\mu_{L4} \equiv \max\left\{\mu : \pi^*_\mu(L|l) = 1\right\}$, $\mu_{R4} \equiv \min\left\{\mu : \pi^*_\mu(R|r) = 1\right\}$, where $\pi^*_\mu(L|l)$ and $\pi^*_\mu(R|r)$ stand for the respective probabilities (51) given μ and $\bar{\tau} = \tau^{**}$. In words, μ_{L4} and μ_{R4} are the most extreme beliefs μ for which the agent acquires information given the candidate-optimal incentive τ^{**}. Closed-form expressions can be obtained, with $\mu_{L4} = \frac{1}{1+e^{\frac{1+\tau^{**}}{\lambda}}}$ and $\mu_{R4} = \frac{e^{\frac{1+\tau^{**}}{\lambda}}}{1+e^{\frac{1+\tau^{**}}{\lambda}}}$.

5. Let $\mu_{L5} \in (0, \mu^*)$ and $\mu_{R5} \in (\mu^*, 1)$ be such that $\gamma(\mu_{L5}, \mu_p) = \gamma(\mu_{R5}, \mu_p) = \chi(1/\rho)$. These cutoffs exists due to the properties of $\gamma(\mu, \mu_p)$ established for μ_{L1}, as well as the fact that $\chi(1/\rho) = e^{\frac{1+1/\rho}{\lambda}} > 1$. Closed-form expressions can be obtained, with

$$\mu_{L5}, \mu_{R5} = \frac{e^{\frac{1+1/\rho}{\lambda}} + 2\mu_p \mp \sqrt{e^{2\frac{1+1/\rho}{\lambda}} + 4\mu_p(\mu_p - 1)}}{2\left(1 + e^{\frac{1+1/\rho}{\lambda}}\right)}.$$

   As $\chi(0) < \chi(1/\rho)$ for all λ (see Lemma 3), it follows that $\mu_{L5} < \mu_{L1}$ and $\mu_{R5} > \mu_{R1}$.

We now proceed to establishing the conditions on μ for which the three properties of τ^{**} (feasibility, effectiveness, preferability) do or do not hold. To begin with, as was previously claimed, a *feasible* τ^{**} exists if and only if $\chi(0) \leq \gamma(\mu, \mu_p) \leq \chi(1/\rho)$, which, due to the monotonicity of $\chi(\bar{\tau})$ in $\bar{\tau}$, is equivalent to

$$\mu \in [\mu_{L5}, \mu_{L1}] \cup [\mu_{R1}, \mu_{R5}]. \tag{56}$$

As shown by construction above, μ_{L5}, μ_{R5} are always well-defined and are located to the outside of μ_{L1} and μ_{R1}, respectively. It thus remains to verify that μ_{L1} and μ_{R1} are also well-defined, which is done by the following lemma.

Lemma 4. *For all* λ*, if* $\rho \geq \min\{1, 1/2\lambda\}$*, then* μ_{L1} *and* μ_{R1} *exist.*

Proof. Function $\chi(\bar{\tau})$ is independent of μ. Function $\gamma(\mu, \mu_p)$ is single-dipped in μ, with $\min_\mu \gamma(\mu, \mu_p) = 2\sqrt{\mu_p(1-\mu_p)} < 1$ achieved at $\mu = \mu^*(\mu_p)$ as given by (12), and $\sup_\mu \gamma(\mu, \mu_p) = +\infty$ achieved by $\mu \to \{0, 1\}$. Hence a sufficient condition for the cutoffs of interest to exist is

$$\chi(0) \geq \gamma(\mu^*(\mu_p), \mu_p). \tag{57}$$

In the inequality above, only $\chi(0)$ depends on ρ. Note further that $\frac{d\chi(0)}{d\rho} > 0$. Therefore, if (57) holds for some $\tilde{\rho}$, then it also holds – and, consequently, μ_{L1} and μ_{R1} exist – for all $\rho \geq \tilde{\rho}$.

Observe that $\chi(0) \geq 1$ for $\rho = \frac{1}{2\lambda}$: denoting $\xi = \xi(\lambda) \equiv e^{\frac{1}{\lambda}}$, we have

$$\begin{aligned} \chi(0) = \xi\frac{\lambda\rho(\xi^2-1)+2}{\lambda\rho(\xi^2-1)+(\xi^2+1)} &\geq 1 \\ \iff (\xi-1)^2\left(\lambda\rho(\xi+1)-1\right) &\geq 0 \end{aligned} \tag{58}$$

Since $\xi = e^{\frac{1}{\lambda}} \geq 1$, a sufficient condition for (58) is given by

$$\rho \geq \frac{1}{\lambda(\xi + 1)}, \tag{59}$$

which obviously holds if $\rho \geq \frac{1}{2\lambda}$. Further, $e^{\frac{1}{\lambda}} \geq 1 + \frac{1}{\lambda} \iff \lambda(\xi + 1) \geq 1$, hence the RHS of (59) is weakly smaller than 1, so the inequality also holds for all $\rho \geq 1$.

We conclude that if $\rho \geq \min\left\{1, \frac{1}{2\lambda}\right\}$, then $\chi(0) \geq 1$, and hence (57) is satisfied and the relevant cutoffs exist. □

Moving on to *effectiveness*, it should be immediate from the analysis in Section 3 that for a feasible τ^{**} to yield interior choice probabilities (51), it must be that $\mu \in [\mu_{L4}, \mu_{R4}]$. The following lemma establishes the location of μ_{L4}, μ_{R4} relative to other cutoffs.

Lemma 5. *Cutoffs μ_{L4} and μ_{R4} are such that $\mu_{L4} \in [\mu_{L5}, \mu_{L2}]$ and $\mu_{R4} \in [\mu_{R2}, \mu_{R5}]$.*

Proof. Denoting $\mu_L(\bar{\tau}) \equiv \max\{\mu : \pi_\mu(L|l, \bar{\tau}) = 1\} = \frac{1}{1+e^{\frac{1+\bar{\tau}}{\lambda}}}$ and observing that it is strictly decreasing in $\bar{\tau}$, we get $\mu_L(0) \geq \mu_L(\tau^{**}) \geq \mu_L(1/\rho)$, which is equivalent to $\mu_{L2} \geq \mu_{L4} \geq \mu_L(1/\rho)$. Routine calculations using the closed-form expression for μ_{L5} then demonstrate that $\mu_L(1/\rho) \geq \mu_{L5}$, implying that in the end, $\mu_{L4} \in [\mu_{L5}, \mu_{L2}]$. The result for μ_{R4} is shown analogously. □

Finally, we need to establish when the principal *prefers* a feasible $\bar{\tau} = \tau^{**}$ to $\bar{\tau} = 0$, which is done by the following lemma.

Lemma 6. *The principal weakly prefers a feasible $\bar{\tau} = \tau^{**}$ to $\bar{\tau} = 0$ if and only if $\mu \in [\mu_{L3}, \mu_{R3}]$. Further, these cutoffs satisfy $\mu_{L3} \in [\mu_{L4}, \mu_{L2}]$ and $\mu_{R3} \in [\mu_{R2}, \mu_{R4}]$.*

Proof. For $\mu \in [\mu_{L4}, \mu_{L2}]$, the principal compares his payoff from choosing $\bar{\tau} = 0$, given by $1 - \mu_p$, and his payoff from choosing $\bar{\tau} = \tau^{**}$. Thus, μ_{L3} satisfies the following indifference condition

$$\mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}] \equiv \big(1 - \rho\tau^{**}(\mu)\big)\big(\mu_p \pi^*_\mu(R|r) + (1 - \mu_p)\pi^*_\mu(L|l)\big) = 1 - \mu_p. \tag{60}$$

The LHS of (60) is single-peaked in μ:

$$\begin{aligned}
\frac{d\mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}]}{d\mu} &= \frac{\partial \mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}]}{\partial \mu} \\
&= (1 - \rho\tau^{**}(\mu)) \left(\mu_p \frac{\partial \pi^*_\mu(R|r)}{\partial \mu} + (1 - \mu_p)\frac{\partial \pi^*_\mu(L|l)}{\partial \mu}\right) \\
&= (1 - \rho\tau^{**}(\mu)) \frac{e^{2\frac{1+\bar{\tau}}{\lambda}}}{e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1} \left(\frac{\mu_p}{\mu^2} - \frac{1 - \mu_p}{(1 - \mu)^2}\right)
\end{aligned}$$

where the first equality follows from the envelope theorem. The final expression is strictly positive for $\mu < \mu^*(\mu_p)$ and strictly negative for $\mu > \mu^*(\mu_p)$, hence the single-peakedness follows.

Thus, $\frac{d}{d\mu}\mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}] > 0$ for $\mu \in [\mu_{L4}, \mu_{L2}]$ (since also $\mu_{L2} < 1/2 \leq \mu^*$). Hence we can show that $\mu_{L3} \in [\mu_{L4}, \mu_{L2}]$ by establishing that

$$\mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}; \mu = \mu_{L4}] \leq 1 - \mu_p \leq \mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}; \mu = \mu_{L2}]$$

and applying the intermediate value theorem. The first inequality follows from the fact that at μ_{L4}, τ^{**} is such that the agent does not acquire information, yet the principal pays a positive transfer to him (which is trivially dominated by $\bar{\tau} = 0$). The second inequality follows from the fact that given $\mu = \mu_{L2}$, (54) holds for all $\bar{\tau} \in [0, 1/\rho]$, so if τ^{**} exists, it is preferred to $\bar{\tau} = 0$. We conclude that $\mu_{L3} \in [\mu_{L4}, \mu_{L2}]$, and the mirror argument can establish that $\mu_{R3} \in [\mu_{R2}, \mu_{R4}]$.

Finally, the single-peakedness of $\mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}]$ in μ implies that $\bar{\tau} = \tau^{**}$ is preferred to $\bar{\tau} = 0$ for all $\mu \in [\mu_{L3}, \mu_{R3}]$. $\square$

To summarize, the principal's problem (23) is solved by $\tau^* \in \{0, \tau^{**}\}$, with τ^{**} being the solution if and only if it is feasible, effective, and preferable. It is feasible if and only if (56) holds; effective if and only if $\mu \in [\mu_{L4}, \mu_{R4}]$, and preferable if and only if $\mu \in [\mu_{L3}, \mu_{R3}]$. Further, we have established that $\mu_{L5} \le \mu_{L4} \le \mu_{L3} \le \mu_{L2}$ (and the converse holds for the other set of cutoffs), as well as $\mu_{L1} < \mu^* \le \mu_p < \mu_{R1}$. Therefore, $\tau^* = \tau^{**}$ if and only if $\mu \in [\mu_{L3}, \mu_{L1}] \cup [\mu_{R1}, \mu_{R3}]$, whenever these intervals are non-empty. After denoting $\hat{\mu}_L \equiv \mu_{L3}$, $\hat{\mu}_R \equiv \mu_{R3}$, $\bar{\mu}_L \equiv \max\{\mu_{L3}, \mu_{L1}\}$, $\bar{\mu}_R \equiv \min\{\mu_{R1}, \mu_{R3}\}$ and excluding the endpoints, at which $\tau^{**} = 0$, we obtain the statement of the Proposition.

A.10 Proof of Proposition 5

Using Theorem 1 of Caplin et al. (2019), the agent's problem (24) given some restriction set A^* is solved by π such that the corresponding $\beta \in \Delta(A^*)$ satisfies (17) for all $a_i \in A^*$. Further, recall from Section 4 that π and β are connected in the optimum by relation (13) (where we set $\pi(a_i|\omega_j) \equiv \beta(a_i) \equiv 0$ for all $a_i \notin A^*$ and all $\omega_j \in \Omega$). Then by plugging (13) and the state-matching utility into the principal's expected payoff, it can be rewritten as in (16):

$$\sum_{i=1}^{N} \mu_p(\omega_i) \frac{\beta(a_i) e^{\frac{1}{\lambda}}}{1 + \delta\beta(a_i)} = \sum_{i \in C(\beta)} \mu_p(\omega_i) \frac{(1+\delta)\beta(a_i)}{1 + \delta\beta(a_i)}.$$

Plugging in (17) for β in the expression above transforms it to

$$\sum_{i \in C(\beta)} \frac{\frac{1+\delta}{\delta} \mu_p(\omega_i) \left[(K(\beta) + \delta)\, \mu_p(\omega_i) - \sum_{j \in C(\beta)} \mu(\omega_j) \right]}{(K(\beta) + \delta) \mu_p(\omega_i)}$$
$$= \frac{1+\delta}{\delta} \left[\sum_{i \in C(\beta)} \mu_p(\omega_i) - \sum_{i \in C(\beta)} \frac{\sum_{j \in C(\beta)} \mu(\omega_j)}{(K(\beta) + \delta)} \right]$$
$$= \frac{1+\delta}{K(\beta) + \delta} \sum_{i \in C(\beta)} \mu_p(\omega_i). \quad (61)$$

To prove the proposition statement, we proceed by induction. Consider some arbitrary action set $A_- \subset \mathcal{A}$ such that $a_k \notin A_-$ for some $k \in \{1, ..., N\}$ and another action set $A_+ \equiv A_- \cup \{a_k\}$. Let β_+ denote the unconditional choice probabilities corresponding to the solution of (24) given A_+, let $C_+ \equiv C(\beta_+)$ and $K_+ \equiv K(\beta_+)$, and define β_-, C_-, K_- analogously given A_-.

Our goal is to show that that selecting $A^* = A_+$ is weakly better for the principal than $A^* = A_-$. If $a_k \notin C_+$, then the payoffs in the two cases are equal, and the statement is trivially true. Otherwise, using (61) for the principal's expected payoff, the statement amounts to:

$$\begin{aligned} 0 \leq & \left(\frac{1+\delta}{K_+ + \delta} \sum_{i \in C_+} \mu_p(\omega_i) \right) - \left(\frac{1+\delta}{K_- + \delta} \sum_{i \in C_-} \mu_p(\omega_i) \right) \\ \iff 0 \leq & \left((K_- + \delta) \sum_{i \in C_+} \mu_p(\omega_i) \right) - \left((K_+ + \delta) \sum_{i \in C_-} \mu_p(\omega_i) \right) \\ \iff 0 \leq & (K_- + \delta)\mu_p(\omega_k) - \left(\sum_{i \in C_-} \mu_p(\omega_i) \right). \end{aligned} \tag{62}$$

Since $a_k \in C_+$ by assumption, $\beta_+(a_k) > 0$, which, from (17), implies that

$$\begin{aligned} 0 < & \frac{(K(\bar{\beta}) + \delta)\mu(\omega_i)}{\sum_{j \in C(\bar{\beta})} \mu(\omega_j)} - 1 \\ \iff 0 < & (K_+ + \delta)\mu_p(\omega_k) - \left(\sum_{i \in C_+} \mu_p(\omega_i) \right) \\ \iff 0 < & (K_- + 1 + \delta)\mu_p(\omega_k) - \left(\sum_{i \in C_-} \mu_p(\omega_i) + \mu_p(\omega_k) \right) \\ \iff 0 < & (K_- + \delta)\mu_p(\omega_k) - \left(\sum_{i \in C_-} \mu_p(\omega_i) \right), \end{aligned}$$

which immediately implies that (62) holds. Therefore, it is indeed better for the principal to choose A_+ over A_-. Since A_- was arbitrary, this proves by induction that allowing a larger action set is always weakly better for the principal, and hence proves the original proposition.

A.11 Proof of Proposition 6

The proof follows in two main parts: first we solve the agent's problem in the binary-quadratic setting, and then we proceed to analyze the principal's problem.

The agent's problem. We solve the agent's problem in the binary-quadratic problem using the so-called posterior approach, where instead of a signal structure ϕ, the agent maximizes over a distribution of posterior beliefs.[36] This reformulation is as follows.

Consider an agent who chooses a signal strategy ϕ with signal realizations in $S \in \mathcal{S}$. Each signal realization $s \in S$ is associated with a corresponding posterior belief $\eta \in [0, 1]$

[36]This approach is popular in the Bayesian Persuasion literature (Kamenica and Gentzkow, 2011). For some examples of this approach being used in problems with rationally inattentive agents, see, e.g., Jain and Whitemeyer (2021) and Matyskova and Montes (2023), with Caplin et al. (2022) presenting a general treatment.

(as with the priors, we represent all posterior beliefs $\eta \in \Delta(\Omega)$ in terms of the probability they assign to state $\omega = 1$). Hence an experiment ϕ induces a distribution over posterior beliefs. It is commonly known from the literature on rational inattention (see, e.g., Caplin and Dean (2013)) that instead of considering the set of all Blackwell experiments ϕ, one can consider the set of distributions of posterior beliefs that average out to the prior.

In order to avoid introducing new notation for such distributions, without loss of generality we restrict attention to "direct" signal structures, in which signals are labeled according to the posterior belief they induce. Formally, for this proof, let $\mathcal{S} \equiv \Delta(\Omega)$ and restrict attention to the *set of feasible direct experiments* ϕ:

$$\Phi_\mu \equiv \{\phi \in \Delta(\Delta(\Omega)) : \ \forall s \in \Delta(\Omega) : \eta(s) = s; \ \mathbb{E}[s|\phi] = \mu\}.$$

In other words, any experiment in Φ_μ prescribes some distribution over signals $s \in \Delta(\Omega)$ such that any signal induces a posterior belief $\eta(s) = s$, and the signals average out to the prior belief μ. An agent with prior belief μ is limited to choosing a signal structure $\phi \in \Phi_\mu$. In what follows, we suppress the signal notation and refer to all signals s according to the posterior beliefs $\eta(s)$ that they induce.

It can be easily seen from (1) that the cost of a direct experiment ϕ can be expressed as

$$c(\phi, \mu) = \lambda\Big[-\mathbb{E}[H(\eta)|\phi] + H(\mu)\Big],$$

where $H(\mu) \equiv \mu \ln \mu + (1-\mu)\ln(1-\mu)$ is the entropy of belief μ.

Under the quadratic preferences $u(a, \omega) = -(a-\omega)^2$, given some posterior belief η, the agent's optimal choice rule is $\sigma^*(\eta) = \eta$. His interim expected payoff given posterior η, thus, equals $\mathbb{E}[u(\eta, \omega)|\eta] = -\eta(1-\eta)$. Therefore, the agent's problem is given by

$$\max_{\phi \in \Phi_\mu} \mathbb{E}\Big[-\eta(1-\eta) + \lambda H(\eta) - \lambda H(\mu) \mid \phi\Big]. \tag{63}$$

Since the last term $-\lambda H(\mu)$ does not affect the maximization, it can be safely ignored. Then we can define the agent's *net utility* as $\nu(\eta) \equiv -\eta(1-\eta) + \lambda H(\eta)$. Problem (63) is then equivalent to $\max_{\phi\in\Phi_\mu} \mathbb{E}[\nu(\eta)|\phi]$, which can be solved using a concavification approach (Kamenica and Gentzkow, 2011; Caplin et al., 2022).

In particular, let $\hat{\nu}(\eta)$ denote the *concavified net utility function* (also known as the concave closure of $\nu(\eta)$; Rockafellar, 1970), which is defined as the minimal concave function that majorizes all net utilities. We characterize $\hat{\nu}(\eta)$ by establishing some properties of $\nu(\eta)$ and its derivatives ν', ν''. Note first that $\nu(\eta)$ is symmetric around 0.5: for any η_1, η_2 s.t. $\eta_1 + \eta_2 = 1$, it is true that $\nu(\eta_1) = \nu(\eta_2)$, $\nu'(\eta_1) = -\nu'(\eta_2)$, $\nu''(\eta_1) = \nu''(\eta_2)$. Further, simple algebra shows that $\nu''(\eta) = 2 - \lambda\frac{1}{\eta(1-\eta)}$. Two cases are then possible, as depicted in Figure 7.

Case 1: $\lambda \geq 0.5$. Then $\nu''(\eta) \leq 0$ and function $\nu(\eta)$ is concave on $[0, 1]$, meaning $\hat{\nu}(\mu) = \nu(\mu)$ for all μ. The results of Caplin et al. (2022) then imply that the agent chooses degenerate distribution ϕ that puts all mass on $\eta = \mu$ and does not acquire any information, regardless of μ. This concludes the analysis of the agent's problem in case 1, and the remainder of this section is devoted to the analysis of the case when $\lambda < 0.5$.

Case 2: $\lambda < 0.5$. Then equation $\nu''(\eta) = 0$ has two roots, denote them as η_l, η_r, with $\eta_l < \eta_r$. Function $\nu(\eta)$ is concave on $[0, \eta_l) \cup (\eta_r, 1]$ and convex on (η_l, η_r), and so its derivative $\nu'(\eta)$ is decreasing on $[0, \eta_l) \cup (\eta_r, 1]$ and increasing on (η_l, η_r). Simple algebra

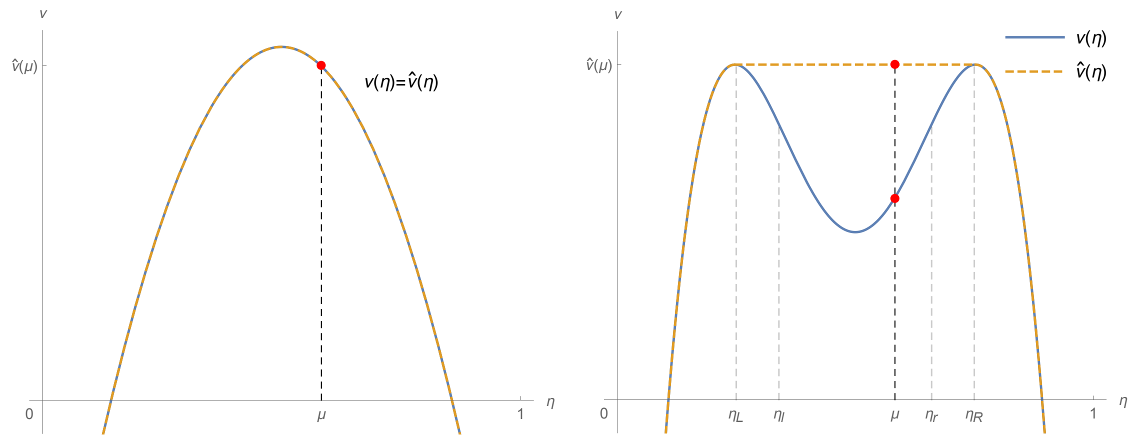

(a) $\lambda \geq \frac{1}{2}$ (b) $\lambda < \frac{1}{2}$

Figure 7: The net utility function $\nu(\eta)$ and the concavified net utility function $\hat{\nu}(\eta)$.

shows that $\nu'(0.5) = 0$. Together with the fact that $\lim_{\eta \to 0+} \nu(\eta) = -\infty = \lim_{\eta \to 1-} \nu(\eta)$, it implies that there exist two other roots of $\nu'(\eta) = 0$ and they belong to the intervals $[0, \eta_l]$ and $[\eta_r, 1]$. We denote these roots as η_L, η_R, respectively, and the symmetry of $\nu(\eta)$ implies that $\eta_L = 1 - \eta_R$.

The behavior of $\nu'(\eta)$ suggests that η_L, η_R are local and global maxima of $\nu(\eta)$. Therefore, $\hat{\nu}(\eta) = \nu(\eta)$ if $\eta \in [0, \eta_L] \cup [\eta_R, 1]$, and if $\eta \in (\eta_L, \eta_R)$ then $\hat{\nu}(\eta)$ lies on the straight line connecting points $(\eta_L, \nu(\eta_L))$ and $(\eta_R, \nu(\eta_R))$. Therefore, if $\mu \in [0, \eta_L] \cup [\eta_R, 1]$, tangent lines to $\nu(\eta)$ and $\hat{\nu}(\eta)$ coincide, and the agent does not acquire information; if $\mu \in (\eta_L, \eta_R)$, the agent chooses ϕ with two posterior beliefs in the support: η_L and η_R.

We proceed to explicitly calculate the optimal signal structure ϕ^* in the latter case (i.e., when the agent acquires information). First, by the law of total probability:

$$\phi^*(\eta_R)\eta_R + \phi^*(\eta_L)\eta_L = \mu \qquad \Rightarrow \qquad \phi^*(\eta_R) = \frac{\mu - \eta_L}{\eta_R - \eta_L}. \tag{64}$$

Further, from the Bayes' rule, we have that $\eta_R = \frac{\mu \cdot \phi^*(\eta_R|1)}{\mu \cdot \phi^*(\eta_R|1) + (1-\mu) \cdot \phi^*(\eta_R|0)}$. From the law of total probability, $\mu \cdot \phi^*(\eta_R|1) + (1-\mu) \cdot \phi^*(\eta_R|0) = \phi^*(\eta_R)$. Combining the two, we get

$$\begin{aligned} \phi^*(\eta_R|1) &= \frac{\eta_R}{\mu} \cdot \frac{\mu - \eta_L}{\eta_R - \eta_L} = 1 - \phi^*(\eta_L|1), \\ \phi^*(\eta_R|0) &= \frac{1 - \eta_R}{1 - \mu} \cdot \frac{\mu - \eta_L}{\eta_R - \eta_L} = 1 - \phi^*(\eta_L|0), \end{aligned} \tag{65}$$

which completes the characterization of ϕ^*. However, it will also prove useful to note that while not reflected in the notation, ϕ^* does depend on the prior, μ, and this dependence is given by

$$\begin{aligned} \frac{\partial \phi^*(\eta_R|1)}{\partial \mu} &= \frac{\eta_R \cdot \eta_L}{\mu^2 \cdot (\eta_R - \eta_L)} = -\frac{\partial \phi^*(\eta_L|1)}{\partial \mu}, \\ \frac{\partial \phi^*(\eta_R|0)}{\partial \mu} &= \frac{(1 - \eta_R) \cdot (1 - \eta_L)}{(1-\mu)^2 \cdot (\eta_R - \eta_L)} = -\frac{\partial \phi^*(\eta_L|0)}{\partial \mu}. \end{aligned} \tag{66}$$

The principal's problem. Begin by assuming $\lambda < 0.5$. The principal's expected utility from hiring a learning agent with prior belief $\mu \in (\eta_L, \eta_R)$ is given by

$$\mu_p\Big(\phi^*(\eta_R|1)\bar{u} + \phi^*(\eta_L|1)\underline{u}\Big) + (1-\mu_p)\Big(\phi^*(\eta_R|0)\underline{u} + \phi^*(\eta_L|0)\bar{u}\Big), \tag{67}$$

where $\bar{u} \equiv -\eta_L^2 = -(1-\eta_R)^2$ and $\underline{u} \equiv -\eta_R^2 = -(1-\eta_L)^2 < \bar{u}$. If it is optimal to hire a learning agent, then the optimal agent's prior μ^* must maximize (67), hence μ^* must solve the FOC given by

$$\mu_p \frac{\eta_R \cdot \eta_L}{(\mu^*)^2 \cdot (\eta_R - \eta_L)}(\bar{u} - \underline{u}) = (1-\mu_p)\frac{(1-\eta_R)\cdot(1-\eta_L)}{(1-\mu^*)^2\cdot(\eta_R-\eta_L)}(\bar{u}-\underline{u})$$
$$\iff \frac{\mu^*}{1-\mu^*} = \sqrt{\frac{\mu_p}{1-\mu_p}},$$

where the first line uses (66), and then the second line follows from $\eta_L = 1-\eta_R$. It is trivial to verify that the SOC also hold. Representation (12) therefore applies conditional on the principal hiring a learning agent. To fully solve the principal's problem it is then left to characterize her choice between a learning and a non-learning agent.

Conditional on hiring a non-learning agent, it is trivially optimal for the principal to hire an aligned agent $\mu = \mu_p$ (since such an agent chooses $a \in \mathcal{A}$ to maximize $\mathbb{E}[u(a,\omega)|\mu_p]$, as opposed to some other function induced by an expectation w.r.t. another belief).

Suppose w.l.o.g. $\mu_p \geq 0.5$ (the logic for $\mu_p < 0.5$ is analogous). We first characterize the principal's optimal strategy for $\mu_p \in \big[\eta_R, (\mu^*)^{-1}(\eta_R)\big)$ – i.e., such μ_p that an agent with $\mu = \mu_p$ is not learning and an agent with $\mu^*(\mu_p)$ is learning.[37] We then show that when one of these is violated (either $\mu = \mu^*(\mu_p)$ does not learn, or $\mu = \mu_p$ does), an extrapolation of such a strategy is optimal.

If $\mu_p \in \big[\eta_R, (\mu^*)^{-1}(\eta_R)\big)$, hiring an aligned agent $\mu = \mu_p$ yields the principal an expected payoff of $-\mu_p(1-\mu_p)$, while hiring agent $\mu^*(\mu_p)$ yields

$$-\mu_p\Big(\phi^*(\eta_R|1)(1-\eta_R)^2 + \phi^*(\eta_L|1)(1-\eta_L)^2\Big) - (1-\mu_p)\Big(\phi^*(\eta_R|0)\eta_R^2 + \phi^*(\eta_L|0)\eta_L^2\Big) =$$
$$= -2\eta_R(1-\eta_R)\sqrt{\mu_p(1-\mu_p)},$$

where the equality uses (65) and the fact that $\eta_R = 1-\eta_L$. The latter (hiring $\mu^*(\mu_p)$) is preferred if and only if

$$-2\eta_R(1-\eta_R)\sqrt{\mu_p(1-\mu_p)} \geq -\mu_p(1-\mu_p)$$
$$\iff 2\eta_R(1-\eta_R) \leq \sqrt{\mu_p(1-\mu_p)}. \tag{68}$$

Note that the RHS of the inequality above is strictly decreasing in $\mu_p \geq 0.5$, while the LHS does not depend on μ_p. Furthermore,

- if $\mu_p = \eta_R$, then (68) holds: $2\eta_R(1-\eta_R) \leq \sqrt{\eta_R(1-\eta_R)} \iff \sqrt{\eta_R(1-\eta_R)} \leq \frac{1}{2}$, which holds because $\eta_R(1-\eta_R) \leq \frac{1}{4}$ for any $\eta_R \in [0.5, 1]$;

[37] Here $(\mu^*)^{-1}(\mu)$ is the inverse of function $\mu^*(\mu_p)$ given by (12).

- if $\mu_p = (\mu^*)^{-1}(\eta_R) = \frac{\eta_R^2}{\eta_R^2+(1-\eta_R)^2}$, then (68) is equivalent to $\eta_R^2+(1-\eta_R)^2 \le \frac{1}{2} \iff 2\left(\eta_R + \frac{1}{2}\right)^2 \le 0$, which fails to hold for any $\eta_R \in [0.5, 1]$.

By the intermediate value theorem, we conclude that there exists $\hat{\mu} \in \left[\eta_R, (\mu^*)^{-1}(\eta_R)\right)$ such that (68) is satisfied – and delegation to a non-learning agent is preferred – if and only if $\mu_p \ge \hat{\mu}$.

If $\mu_p \ge (\mu^*)^{-1}(\eta_R)$, then the optimal non-learning agent $\mu = \mu_p$ is still available, but the optimal learning agent $\mu = \mu^*(\mu_p)$ is not, hence (if the principal wants a learning agent, she has to deviate away from $\mu^*(\mu_p)$, and so) the principal has an even stronger preference towards hiring a non-learning agent than in the case above. The mirror logic also holds: if $\mu_p < \eta_R$, then the optimal learning agent is available, but the optimal non-learning agent is not, hence the principal has a stronger preference towards hiring a learning agent than in the case above.

We conclude that there exists $\hat{\mu} \in \left[\eta_R, (\mu^*)^{-1}(\eta_R)\right)$ that solves (68) w.r.t. μ_p such that the principal's optimal strategy is given by:

$$\mu^{**}(\mu_p) = \begin{cases} \mu^*(\mu_p) & \text{as given by (12) if } \mu_p < \hat{\mu}, \\ \mu_p & \text{if } \mu_p \ge \hat{\mu}. \end{cases}$$

This concludes the proof for the case $\lambda < 0.5$. In case $\lambda \ge 0.5$, no agent acquires any information, hence by the logic above, hiring an agent $\mu = \mu_p$ is optimal. Setting $\hat{\mu} = 0.5$, we get the result in this case as well, which completes the proof.

A.12 Proof of Proposition 7

We provide an example for $N = 3$. We use the same version of the Farkas' Lemma as in the proof of Lemma 2. To show that there is no prior belief that solves the system of the first-order conditions for the problem, it is sufficient to show that there is a solution to the following dual inequality system

$$\begin{cases} z_1 e^{\frac{u(a_1,\omega_1)}{\lambda}} + z_2 e^{\frac{u(a_1,\omega_2)}{\lambda}} + z_3 e^{\frac{u(a_1,\omega_3)}{\lambda}} \ge 0, \\ z_1 e^{\frac{u(a_2,\omega_1)}{\lambda}} + z_2 e^{\frac{u(a_2,\omega_2)}{\lambda}} + z_3 e^{\frac{u(a_2,\omega_3)}{\lambda}} \ge 0, \\ z_1 e^{\frac{u(a_3,\omega_1)}{\lambda}} + z_2 e^{\frac{u(a_3,\omega_2)}{\lambda}} + z_3 e^{\frac{u(a_3,\omega_3)}{\lambda}} \ge 0, \\ z_1 + z_2 + z_3 < 0. \end{cases} \tag{69}$$

Let us normalize $\lambda = 1$ and consider payoffs given by the following matrix:

$$\begin{pmatrix} u(a_1,\omega_1) & u(a_2,\omega_1) & u(a_3,\omega_1) \\ u(a_1,\omega_2) & u(a_2,\omega_2) & u(a_3,\omega_2) \\ u(a_1,\omega_3) & u(a_2,\omega_3) & u(a_3,\omega_3) \end{pmatrix} = \begin{pmatrix} \ln 3 & 0 & \ln(2+\varepsilon) \\ 0 & \ln 3 & \ln(2+\varepsilon) \\ 0 & 0 & \ln(2+\varepsilon) \end{pmatrix}$$

Notice that vector $(z_1, z_2, z_3) = (-1-\delta, -1-\delta, 2)$ for small enough $\delta, \varepsilon \ge 0$ solves system (69): the two latter inequalities hold trivially for all such z, and the two former inequalities hold if $\varepsilon \ge 3^{\frac{1+\delta}{2}} - 2$. Therefore, there exists no μ that solves system (37) given $\beta \in \Delta(\Theta)$.

A.13 Proof of Proposition 8

The proof is largely analogous to that of Proposition 6 and proceeds in two main parts: we first solve the agent's problem in the binary setting of Section 3 with uniformly posterior-separable information costs, and then we proceed to analyze the principal's problem.

The agent's problem. We solve the agent's problem using the so-called posterior approach, where instead of a signal structure ϕ, the agent maximizes over a distribution of posterior beliefs (see Footnote 36). In particular, consider an agent that chooses a signal strategy ϕ with signal realizations in $S \in \mathcal{S}$. Each signal realization $s \in S$ is associated with a corresponding posterior belief $\eta \in [0, 1]$ (as with the priors, we represent all posterior beliefs $\eta \in \Delta(\Omega)$ in terms of the probability they assign to state $\omega = r$). Hence an experiment ϕ induces a distribution over posterior beliefs. It is commonly known from the literature on rational inattention (see, e.g., Caplin and Dean (2013)) that instead of considering the set of all Blackwell experiments ϕ, one can consider the set of distributions of posterior beliefs that average out to the prior.

In order to avoid introducing new notation for such distributions, we instead consider "direct" signal structures, in which signals are labeled according to the posterior belief they induce. Formally, for this proof, let $\mathcal{S} \equiv \Delta(\Omega)$ and restrict attention to the *set of feasible direct experiments* ϕ:

$$\Phi_\mu \equiv \{\phi \in \Delta(\Delta(\Omega)) : \forall s \in \Delta(\Omega) : \eta(s) = s;\ \mathbb{E}[s|\phi] = \mu\}.$$

In other words, any experiment in Φ_μ prescribes some distribution over signals $s \in \Delta(\Omega)$ such that any signal induces a posterior belief $\eta(s) = s$, and the signals average out to the prior belief μ. An agent with prior belief μ is limited to choosing a signal structure $\phi \in \Phi_\mu$. In what follows, we suppress the signal notation and refer to all signals s according to the posterior beliefs $\eta(s)$ that they induce.

With state-matching preferences, the agent's optimal choice rule as a function of his posterior belief η is given by (up to a tie-breaking rule)

$$\sigma^*(\eta) = \begin{cases} R & \text{if } \eta \geq 0.5, \\ L & \text{if } \eta < 0.5. \end{cases}$$

The agent's expected payoff is hence $\mathbb{E}[u(\sigma^*(\eta), \omega)|\eta] = \max\{\eta, 1-\eta\}$. Since the information cost (of a given direct experiment ϕ) is given by

$$c_{\mathsf{UPS}}(\phi, \mu) \equiv \lambda\Big[\mathbb{E}[\hat{c}(\eta)|\phi] - \hat{c}(\mu)\Big],$$

the agent's problem then amounts to

$$\max_{\phi \in \Phi_\mu} \mathbb{E}\Big[\max\{\eta, 1-\eta\} - \lambda\hat{c}(\eta) + \lambda\hat{c}(\mu) \mid \phi\Big]. \tag{70}$$

Since the last term $\lambda\hat{c}(\mu)$ does not affect the maximization, it can be safely ignored. Then we can define the agent's *net utility* as $\nu(\eta) \equiv \max\{\eta, 1-\eta\} - \lambda\hat{c}(\eta)$. Problem (70) is then equivalent to $\max_{\phi\in\Phi_\mu} \mathbb{E}[\nu(\eta)|\phi]$, which can be solved using a concavification approach (Kamenica and Gentzkow, 2011; Caplin et al., 2022). In particular, let $\hat{\nu}(\eta)$ denote the *concavified net utility function* (also known as the concave closure of $\nu(\eta)$; Rockafellar, 1970), which is defined as the minimal concave function that majorizes all net utilities. We characterize $\hat{\nu}(\eta)$ by establishing some properties of $\nu(\eta)$ and its derivatives ν', ν''.

Note first that since $\hat{c}(\eta)$ is assumed to be symmetric around 0.5, $\nu(\eta)$ is also symmetric around 0.5: for any η_1, η_2 s.t. $\eta_1 + \eta_2 = 1$, it is true that $\nu(\eta_1) = \nu(\eta_2)$, $\nu'(\eta_1) = -\nu'(\eta_2)$, $\nu''(\eta_1) = \nu''(\eta_2)$. We thus limit our analysis to the interval $\eta \in [0, 0.5)$, and the analysis for $\eta \in (0.5, 1]$ is analogous. Clearly, if $\eta \in [0, 0.5)$ then $\nu(\eta) = 1 - \eta - \lambda\hat{c}(\eta)$. Function $\hat{c}(\eta)$ is assumed to be convex, thus $\nu''(\eta) = -\lambda\hat{c}''(\eta) < 0$ for $\eta \in [0, 0.5)$. Therefore, the derivative $\nu'(\eta)$ is decreasing for $\eta \in [0, 0.5)$. Convexity and symmetry of $\hat{c}(\eta)$ also imply that it attains minimum at $\eta = 0.5$, so $\hat{c}'(0.5) = 0$ and $\lim_{\eta \to 0.5-} \nu'(\eta) < 0$. On the other hand, from the Inada properties of $\hat{c}(\eta)$, we also have that $\lim_{\eta \to 0+} \nu'(\eta) = +\infty$. Since derivative $\nu'(\eta)$ is continuous and decreasing for all $\eta \in [0, 0.5)$, there exists a unique root of the equation $\nu'(\eta) = 0$. Denote this root as η_L and note that it is then a maximizer of $\nu(\eta)$ on the interval $[0, 0.5)$. Symmetry of $\nu(\eta)$ implies that there exists a symmetric maximizer $\eta_R \in (0.5, 1]$ of $\nu(\eta)$ on $(0.5, 1]$. Moreover, $\eta_L + \eta_R = 1$ and $\nu(\eta_L) = \nu(\eta_R)$.

It further follows from the monotonicity of $\nu'(\eta)$ on each interval that $\nu(\eta)$ is concave on $\eta \in [0, \eta_L] \cup [\eta_R, 1]$, hence $\hat{\nu}(\eta) = \nu(\eta)$ for such η. Conversely, if $\eta \in (\eta_L, \eta_R)$, then $\hat{\nu}(\eta)$ is the straight line connecting points $(\eta_L, f(\eta_L))$ and $(\eta_R, f(\eta_R))$. Thus, if the agent's prior is $\mu \in [0, \eta_L] \cup [\eta_R, 1]$, then tangent lines to $\nu(\eta)$ and $\hat{\nu}(\eta)$ coincide, and the agent does not acquire any information; if $\mu \in (\eta_L, \eta_R)$, then the agent optimally chooses experiment ϕ with two posterior beliefs in the support, η_L and η_R, which are such that $\eta_L + \eta_R = 1$.

In the latter case, we can derive the optimal experiment ϕ^* in closed form. Conditional on η_L and η_R, expressions (64)–(66) apply in this context as well after replacing $\phi^*(\eta|1)$ and $\phi^*(\eta|0)$ with $\phi^*(\eta|r)$ and $\phi^*(\eta|l)$, respectively.

The principal's problem. The principal's expected utility from hiring a learning agent with prior belief $\mu \in (\eta_L, \eta_R)$ is given by

$$\mathbb{E}[u(\sigma^*(\eta), \omega)|\mu_p, \phi^*] = \mu_p \phi^*(\eta_R|r) + (1 - \mu_p)\phi^*(\eta_L|l). \tag{71}$$

If it is optimal to hire a learning agent, then the optimal agent's prior μ^* must maximize (71), hence μ^* must solve the FOC given by

$$\mu_p \frac{\eta_R \cdot \eta_L}{(\mu^*)^2 \cdot (\eta_R - \eta_L)} = (1 - \mu_p)\frac{(1 - \eta_R) \cdot (1 - \eta_L)}{(1 - \mu^*)^2 \cdot (\eta_R - \eta_L)}$$
$$\iff \frac{\mu^*}{1 - \mu^*} = \sqrt{\frac{\mu_p}{1 - \mu_p}},$$

where the first line uses (66), and then the second line follows from $\eta_L = 1 - \eta_R$. It is trivial to verify that the SOC also hold at this point. Representation (12) therefore applies conditional on the principal hiring a learning agent. To fully solve the principal's problem it is then left to characterize her choice between a learning and a non-learning agent.

Suppose w.l.o.g. $\mu_p \geq 0.5$ (the logic for $\mu_p < 0.5$ is analogous). We first characterize the principal's optimal strategy for $\mu_p \in \left[0.5, (\mu^*)^{-1}(\eta_R)\right)$ – i.e., such μ_p that an agent with $\mu^*(\mu_p)$ is learning. Here $(\mu^*)^{-1}(\mu)$ is the inverse of function $\mu^*(\mu_p)$ given by (12), so $(\mu^*)^{-1}(\eta_R) = \frac{\eta_R^2}{\eta_R^2+(1-\eta_R)^2}$. Conditional on hiring a non-learning agent, the principal is indifferent between all agents $\mu \in [\eta_R, 1]$, since any such non-learning agent chooses $a = R$ and yields the principal an expected payoff equal to μ_p. If $\mu_p \in \left[0.5, (\mu^*)^{-1}(\eta_R)\right)$, hiring an optimal learning agent $\mu^*(\mu_p)$ yields

$$\mu_p \phi^*(\eta_R|r) + (1-\mu_p)\phi^*(\eta_L|l) = \frac{\eta_R}{\eta_R - \eta_L}\left(\eta_R - 2\eta_L\sqrt{\mu_p(1-\mu_p)}\right). \tag{72}$$

Routine algebraic manipulations help establish that (72) is convex in μ_p, and it is tangent to μ_p at $\mu_p = (\mu^*)^{-1}(\eta_R)$. Therefore, hiring agent $\mu^*(\mu_p)$ is better than hiring a non-learning agent for all $\mu_p \in \left[0.5, (\mu^*)^{-1}(\eta_R)\right)$.

In turn, for $\mu_p \geq (\mu^*)^{-1}(\eta_R)$, maximizing (71) yields a corner solution $\mu = \eta_R$, hence hiring a learning agent is not optimal, and the principal prefers instead to hire any non-learning agent $\mu \in [\eta_R, 1]$ – hence representation (12) prescribes one optimal strategy for such μ_p. This concludes the proof.

A.14 Proof of Proposition 10

We first show that there exists an equilibrium in the communication game that replicates the deletation equilibrium: the optimal agent acquires the same information, makes a truthful action recommendation, and the principal follows the recommendation.

Suppose that under delegation, the optimally chosen agent follows a decision rule β^* that yields a consideration set $C(\beta^*) = \{1, ..., K^*\}$. By Lemma 1, we have that

$$\begin{aligned}
\sqrt{\mu(\omega_{K^*})} &\geq \frac{1}{K^*+\delta}\sum_{i=1}^{K^*}\sqrt{\mu(\omega_i)} \\
\iff \delta\sqrt{\mu(\omega_{K^*})} &\geq \sum_{i=1}^{K^*-1}\left(\sqrt{\mu(\omega_i)} - \sqrt{\mu(\omega_{K^*})}\right)
\end{aligned} \tag{73}$$

Suppose the agent reports truthfully. Given the state-matching payoffs, for the principal to follow recommendation $\tilde{a} = \tilde{a}_{K^*}$ whenever it is issued, it must hold that

$$\mu_p(\omega_{K^*}|\tilde{a}_{K^*}) = \max_i \mu_p(\omega_i|\tilde{a}_{K^*}), \tag{74}$$

where $\mu_p(\omega|\tilde{a})$ is the probability that the principal's posterior belief assigns to state ω after hearing recommendation $\tilde{a}$ from the agent. In equilibrium, the principal's posterior $\mu_p(\omega_{K^*}|\tilde{a}_{K^*})$ must satisfy Bayes' rule:

$$\begin{aligned}
\mu_p(\omega_{K^*}|\tilde{a}_{K^*}) &= \frac{\pi(a_{K^*}|\omega_{K^*})\mu_p(\omega_{K^*})}{\sum_{i=1}^N \mu_p(\omega_i)\pi(a_{K^*}|\omega_i)} \\
&= \frac{\beta(a_{K^*})e^{\frac{1}{\lambda}}}{\beta(a_1)+...+\beta(a_{K^*-1})+\beta(a_{K^*})e^{\frac{1}{\lambda}}}\cdot\frac{\mu_p(\omega_{K^*})}{\sum_{i=1}^N \mu_p(\omega_i)\pi(a_{K^*}|\omega_i)} \\
&= \frac{\beta(a_{K^*})e^{\frac{1}{\lambda}}}{1+\delta\beta(a_{K^*})}\cdot\frac{\mu_p(\omega_{K^*})}{\sum_{i=1}^N \mu_p(\omega_i)\pi(a_{K^*}|\omega_i)} \\
&= \frac{\sum_{i=1}^{K^*}\sqrt{\mu_p(\omega_i)}}{K^*+\delta}\cdot\beta(a_{K^*})e^{\frac{1}{\lambda}}\cdot\frac{\sqrt{\mu(\omega_{K^*})}}{\sum_{i=1}^N \mu_p(\omega_i)\pi(a_{K^*}|\omega_i)},
\end{aligned}$$

Where the last line is obtained by plugging the expression for $\beta(a_{K^*})$ from Lemma 1 in the denominator of the preceding line. Similarly, we can calculate the probability that the principal's posterior assigns to any other state ω_j:

$$\mu_p(\omega_j|\tilde{a}_{K^*}) = \begin{cases} \frac{\sum_{i=1}^{K^*}\sqrt{\mu_p(\omega_i)}}{K^*+\delta} \cdot \beta(a_{K^*})e^{\frac{1}{\lambda}} \cdot \frac{\sqrt{\mu(\omega_j)}}{\sum_{i=1}^{N}\mu_p(\omega_i)\pi(a_{K^*}|\omega_i)} & \text{if } j < K^*, \\ 0 & \text{if } j > K^*. \end{cases}$$

For condition (74) to hold, it is then enough for

$$e^{\frac{1}{\lambda}}\sqrt{\mu(\omega_{K^*})} \geq \sqrt{\mu(\omega_1)} \quad \iff \quad \delta\sqrt{\mu(\omega_{K^*})} \geq \sqrt{\mu(\omega_1)} - \sqrt{\mu(\omega_{K^*})}, \tag{75}$$

to be satisfied. Note, however, that it is strictly weaker than (73), since

$$\sqrt{\mu(\omega_1)} - \sqrt{\mu(\omega_{K^*})} < \sum_{i=1}^{K^*-1}\left(\sqrt{\mu(\omega_i)} - \sqrt{\mu(\omega_{K^*})}\right).$$

Therefore, we conclude that (75) holds, and thus it is optimal for the principal to choose action a_{K^*} when the agent with prior belief μ^* recommends it.

Following the same argument, we can show the same for any other recommendation $\tilde{a}_i$ for $i \in C(\beta^*)$: the necessary and sufficient condition for the principal to find it optimal to follow the recommendation would be

$$e^{\frac{1}{\lambda}}\sqrt{\mu(\omega_i)} \geq \sqrt{\mu(\omega_1)},$$

which is implied by (74), since $\mu(\omega_i) \geq \mu(\omega_{K^*})$ for $i \in C(\beta^*)$. This concludes the proof.