

CHARLES UNIVERSITY

FACULTY OF SOCIAL SCIENCES

Center for Economic Research and Graduate Education

Dissertation Thesis

2023

Maxim Senkov

CHARLES UNIVERSITY

FACULTY OF SOCIAL SCIENCES

Center for Economic Research and Graduate Education

Maxim Senkov

Essays in Strategic Information Provision

Dissertation Thesis

Prague 2023

Author: Maxim Senkov

Supervisor: PhDr. Jan Zápala, Ph.D.

Year of the defense: 2023

References:

SENKOV, Maxim. Essays in Strategic Information Provision. Praha, 2023. 124 pages. Dissertation thesis (PhD.). Charles University, Faculty of Social Sciences, Center for Economic Research and Graduate Education. Supervisor PhDr. Jan Zápál, Ph.D.

Abstract:

In the first chapter, we show that 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 of a better-informed decision always outweighs the cost of a small misalignment.

In the second chapter, I study a game between an agent and a principal in a dynamic information design framework. A principal funds a multistage project and retains the right to cut the funding if it stagnates at some point. An agent wants to convince the principal to fund the project as long as possible, and can design the flow of information about the progress of the project in order to persuade the principal. If the project is sufficiently promising ex ante, then the agent commits to providing only the good news that the project is accomplished. If the project is not promising enough ex ante, the agent persuades the principal to start the funding by committing to provide not only good news but also the bad news that a project milestone has not been reached by an interim deadline.

In the third chapter, we study an information design model in which the state space is finite, the sender and the receiver have state-dependent quadratic loss functions, and their disagreement regarding the preferred action is of arbitrary form. This framework enables us to focus on the understudied sender’s trade-off between the informativeness of the signal and the concealment of the state-dependent disagreement about the preferred action. In particular, we study which states are pooled together in the supports of posteriors of the optimal signal.

Abstrakt:

V první kapitole ukážeme, že zaujatý principál může mít prospěch z najmutí agenta s neshodnými preferencemi nebo přesvědčením. Zkoumáme problém delegované expertízy“, ve kterém má agent oproti principálovi výhodu při získávání informací. Ukazujeme, že pro principála, který je ex ante zaujatý vůči jedné akci, je optimální vybrat agenta, který je méně zaujatý. Takový agent je ex ante více nejistý ohledně toho, jaký postup je nejlepší, a získá tak více informací. Přínos z lépe informovaného rozhodnutí pro principála vždy převáží náklady na malou neshodu v preferencích.

Ve druhé kapitole studuji hru mezi agentem a principálem v rámci dynamického informačního designu. Principál financuje vícefázový projekt a ponechává si právo omezit financování, pokud v určitém okamžiku začne projekt stagno-

vat. Agent chce přemluvit principála, aby financoval projekt tak dlouho, jak je to možné, a může navrhnout tok informací o průběhu projektu, aby principála přesvědčil. Pokud je projekt dostatečně slibný ex ante, pak se agent zavazuje poskytovat pouze dobré zprávy, že projekt je dokončen. Pokud projekt není dostatečně slibný ex ante, agent přesvědčí principála, aby zahájil financování tím, že se zaváže poskytovat nejen dobré, ale i špatné zprávy ohledně nedosažení milníku projektu v prozatímním termínu.

Ve třetí kapitole, se zabýváme modelem Bayesiánského přesvědčování s konečně mnoha stavy a kvadratickými ztrátovými funkcemi odesílatele a příjemce závisujícími na stavu. Nesouhlas mezi odesílatelem a příjemcem ohledně optimální akce může mít libovolný tvar. Tento model umožňuje zaměřit se na relativně neprozkoumaný kompromis mezi informativností signálu a utajením nesouhlasu ohledně optimální akce. Konkrétně se zaměřujeme na to, jak odesílatel sdružuje stavy v posteriorech optimálního signálu.

Keywords:

information design, dynamic information design, preference misalignment, informational incentives, rational inattention, discrete choice.

Klíčová slova:

informační design, dynamický informační design, tvar nesouladu preferencí, informační pobídky, racionální nepozornost, diskrétní volba.

Length of the work:

319 370.

Declaration

1. I hereby declare that I have compiled this thesis using the listed literature and resources only.
2. I hereby declare that my thesis has not been used to gain any other academic title.
3. I fully agree to my work being used for study and scientific purposes.

In Prague on
29 May

Maxim Senkov

Contents

Introduction	3
1 Optimally Biased Expertise	5
1.1 Introduction	5
1.2 Model	9
1.2.1 The Story	9
1.2.2 The Setup	10
1.2.3 The Agent's Problem	11
1.2.4 The Principal's Problem	12
1.2.5 Equilibrium Definition	12
1.2.6 Preliminary Analysis	13
1.3 Binary Case	13
1.4 General Case	18
1.4.1 Agent's Problem	18
1.4.2 Principal's Relaxed Problem	19
1.4.3 Principal's Full Problem	20
1.4.4 Properties of the Optimal Delegation Strategy	20
1.5 Misaligned Beliefs Versus Other Instruments	22
1.5.1 Contracting on Actions/Misaligned Preferences	22
1.5.2 Contracting on Outcomes	24
1.5.3 Restricting the Delegation Set	26
1.6 Extensions	26
1.6.1 Alternative Preference Specifications	26
1.6.2 Alternative cost functions	27
1.6.3 Communication	31
1.7 Conclusion	34
1.A Main Proofs	36
1.A.1 Proof of Proposition 1.3	36
1.A.2 Proof of Corollary 1.3	37
1.A.3 Proof of Lemma 1	38
1.A.4 Proof of Lemma 2	39
1.A.5 Proof of Theorem 1.4.4	40
1.A.6 Proof of Proposition 1.4.4	41
1.A.7 Proof of Theorem 1.5.1	42
1.A.8 Proof of Proposition 1.5.1	43
1.A.9 Proof of Proposition 1.5.2	44
1.A.10 Proof of Proposition 1.5.3	49
1.A.11 Proof of Proposition 1.6.1	50
1.A.12 Proof of Proposition 1.6.3	51
2 Setting Interim Deadlines to Persuade	53
2.1 Introduction	53
2.2 Related literature	55
2.3 The model	57
2.3.1 Setup	57

2.3.2	Discussion of assumptions	58
2.4	No-information and full-information benchmarks	59
2.4.1	No-information benchmark	59
2.4.2	Full-information benchmark	61
2.5	Agent’s choice of information policy	64
2.5.1	The structure of optimal information disclosure	64
2.5.2	Postponed disclosure of project completion	65
2.5.3	Immediate disclosure of completion and an interim deadline	72
2.6	General preferences	75
2.7	Conclusion	77
2.A	Notational conventions	78
2.B	The principal’s continuation value	78
2.C	Proofs	79
2.D	Disclosure of project completion with a deterministic delay	93
2.E	The case of no project completion deadline	95
3	Form of Preference Misalignment Linked to State-Pooling Structure in Bayesian Persuasion	96
3.1	Introduction	96
3.2	Related literature	98
3.3	Model	100
3.4	General results about the optimal signal	101
3.4.1	Characterization of non-disclosure	102
3.4.2	Full disclosure	102
3.4.3	“Extremization” – non-existence of an interior posterior	103
3.5	State-pooling structure of the optimal signal	103
3.5.1	Definitions	103
3.5.2	Procedure for discovery of the state-pooling structure of the optimal signal	105
3.5.3	Discussion of the procedure	106
3.6	Characterization of the state-pooling structure for $n = 3$	107
3.7	Conclusion	109
3.A	Technical details and proofs	110
3.A.1	The structure of the sender’s problem	110
3.A.2	Proofs	111
3.B	Comment on Assumption 3.1	115
3.C	Demonstration of the procedure for discovery of the state-pooling structure of the optimal signal	116
	Bibliography	119

Introduction

Information is valuable for efficient decision-making, and thus, in many situations, the owners of information can strategically use it to their advantage. Retail goods sellers tend to conceal certain attributes of a product while being transparent about others, biased prosecutors strategize when inviting a witness to a court hearing, and experts obfuscate specific pieces of evidence when advising a policymaker - all these are examples of strategic information provision, which is the overarching theme of this dissertation. The first chapter examines a setting in which a boss and an expert agree on an issue *ex post* but disagree *ex ante* due to differences in opinions on the issue, and studies the boss's choice of expert based on these differences. The second and third chapters consider the expert and the decision-maker with misaligned preferences, focusing on the dynamic and static aspects of the expert's choice of information provision to the decision-maker, respectively.

In the first chapter, we consider a principal and an agent who share the same preference for choosing an action that matches the unobservable state of the world, but who have different prior beliefs regarding the state. In contrast to the principal, the agent can access costly information about the state of the world. We study the principal's choice in the "delegated expertise" problem: the principal selects an agent based on the agent's prior belief, and then the agent acquires information and chooses the action. Surprisingly, we show that it is not optimal for the principal to delegate the task to an agent whose prior belief is perfectly aligned with the principal's prior. Instead, we demonstrate that it is optimal for the principal to delegate the task to an agent who agrees with the principal on the best action choice *ex ante* but who is more uncertain than the principal (i.e., the optimal agent's prior belief is relatively closer to the uniform prior belief than the principal's). We demonstrate that, instead of delegating to an agent with a misaligned prior, the principal can achieve the optimal delegation outcome by delegating to an agent with misaligned preferences over the action choices. Finally, we show that delegation to a misaligned agent as an incentive tool performs at least as well as action-contingent or outcome-contingent transfers to the agent. The results serve as a useful directional behavioral tool for delegation in public organizations and provide support for diversity in large organizations.

In the second chapter, I examine a game between a principal who funds an innovative multi-stage project and decides when to cut funding, and an agent who controls the information on the project's progress toward completion. I assume that the project has two stages, the agent prefers the principal to postpone cutting the funding, and the agent can commit to a dynamic information policy specifying which pieces of information will be provided to the principal and when. I study the agent's choice of information policy and show that if the project is sufficiently attractive to the principal *ex ante*, then the agent promises to disclose only the completion of the second stage of the project, doing so with a postponement. However, if the project is not sufficiently attractive *ex ante*, the agent promises to provide information regarding the completion of both the second and first stages of the project. This particular structure of optimal disclosure is preserved under a more general preference specification. Intriguingly, the optimal form of disclosure

for the first stage completion is a deterministic interim deadline: at the outset of the game, the agent announces a date at which she will disclose if the first stage is already completed or not, and upon receiving the bad news, the principal cuts funding for the project precisely at the interim deadline. The results shed light on the information disclosure in venture finance and bureaucracies.

In the third chapter, we explore the structure of optimal signals in the Bayesian persuasion model with a continuous action space and discrete state space. We assume that both the sender and the receiver share quadratic loss preferences, which implies that they have state-dependent preferred actions. The main twist is that we assume the difference between the sender's and the receiver's state-dependent preferred actions can have an arbitrary form, and we show how this form determines the sender's choice of optimal signal. In particular, we define the state-pooling structure of a signal. Given a signal, this structure specifies which states of the world are pooled together in the supports of posterior beliefs constituting the signal. Unexpectedly, we demonstrate that the state-pooling structure of an optimal signal can be explored using a simple condition on the alignment of the sender's and receiver's preferred actions at pairs of states. Using this condition, we provide a graph procedure that takes the set of states of the world and the form of preference misalignment as input, and delivers the optimal and candidate-optimal pools of states as the output.

1. Optimally Biased Expertise

Co-authored with Pavel Ilinov (CERGE-EI), Andrei Matveenko (University of Mannheim, Department of Economics), and Egor Starkov (University of Copenhagen, Department of Economics).

1.1 Introduction

Presidents, CEOs, and other leaders are often touted as visionaries, paving the way to a brighter tomorrow. However, they cannot do this alone. They regularly rely on the advice and expertise of others, and 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, p. 354]. Chairing Johnson’s Council of Economic Advisers was Walter Heller, who, while being one of the original authors of the War on Poverty, 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 from their early days 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 to 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, because, 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 for loyalty when selecting appointees [Lewis, 2011] – one would think that misalignment depresses loyalty, while not necessarily benefitting the competence. Nevertheless, in this paper we show 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

¹Crichton, 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>

²Atkinson, 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/>

³In 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

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 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 1.3 and 1.6.2, Theorem 1.4.4). 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. This, however, must 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 1.6.3).

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 offer conditions on the set of experts for delegation to be beneficial, as well as an upper bound for the expected gains from such a delegation. Moreover, if it is relatively straightforward for the authority to rank 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 who 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.⁴

Importantly, the optimal degree of misalignment is non-monotone in the strength of principal’s own bias (Corollary 1.3). If the principal is unbiased, then she would prefer an unbiased agent, who would hence be aligned with her vision. The same would apply if a principal is extremely biased – in this case she is almost certain that one action is better than all others, and may either

likely to benefit the economy, whereas the Economist could advocate for a more restricted fiscal policy for sake of maintaining control over monetary policy.

⁴Banerjee 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 can have a harder time agreeing on a decision and may produce worse outcomes. For a recent review on diversity in organizations, see Shore, Chung-Herrera, Dean, Ehrhart, Jung, Randel, and Singh [2009].

take this action on her own, or find an agent who is equally as biased. Hiring a misaligned agent is hence most optimal for the somewhat-biased principal, who has some *ex ante* preference for one action over another, but who 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. Our Theorem 1.5.1 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 estimating discrete choice models: Theorem 1.5.1 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. Further, in our setting, this instrument can perform equally as well as contracts with action-contingent payments (Theorem 1.5.1) and outcome-contingent payments (Proposition 1.5.2), – and even better if we take the principal’s contract costs into account. Further, misalignment is typically better than restricting the agent’s choice set (Proposition 1.5.3). This benefit of misalignment challenges the opinion that disagreement between the principal and the agent inevitably leads to conflict, and thus the principal should seek to hire an agent who is most aligned with her preferences and beliefs (see Holmström, 1980, Crawford and Sobel, 1982b, 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 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).⁵ 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 relative tilt in the misaligned agent’s decisions. We demonstrate (see Section 1.6.2) that our results are not specific to the entropy parametrization and continue to hold with other information cost specifications that allow for flexible learning, such as the channel capacity cost [Woodford, 2012] and the log-likelihood ratio cost [Pomatto, Strack, and Tamuz, 2023].

Our paper is mainly connected to the literature on delegation. Most papers on delegation follow Holmström [1980] in assuming that the agent has preexist-

⁵The 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 general preferences is studied in Section 6.1.

ing 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* advantage over the principal.⁶ 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). Szalay [2005] shows that restricting the agent’s action set can be a useful tool in such a setting, because 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 overarching message is similar: the principal is willing to sacrifice something in exchange for the agent’s acquiring more information, but we present a different channel through which the principal can achieve this.

The closest study 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 *ex ante* safe actions can lead to more information acquisition by the agent, which benefits the principal. However, where Szalay [2005] looks at a 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 have also been 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 punishment for concealing evidence. As we show in Section 1.6.3, 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

⁶Graham, Harvey, and Puri [2015] show that delegation tends to be applied when the decision-making demands more evidence than the delegatee can provide. Alternatively, the choice to delegate a decision is often exercised when a delegator faces a volatile environment [Foss and Laursen, 2005, Ekinci and Theodoropoulos, 2021], in which any knowledge quickly becomes obsolete.

by hiring a more uncertain agent, who will 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 statics results and to show that the optimal misalignment is non-monotone in the principal’s bias.

Finally, a literature exists 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 by focusing on delegation of *non-strategic* decisions, showing how misalignment may be beneficial even in the absence of a strategic counterparty.

The remainder of the paper is organized as follows: Section 1.2 formulates the main model, which is analyzed in Section 1.3 for the special case of binary states and actions, while Section 1.4 analyzes the general problem. Section 1.5 compares misaligned beliefs as a delegation tool to other tools, such as misaligned preferences, payments, and restricting the action set. Section 1.6 explores a number of extensions of the baseline model, and Section 1.7 concludes.

1.2 Model

1.2.1 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 1.2.2.

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 1.6.3 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 1.5.1 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, Sibony, and Sunstein [2021] survey a large body of evidence suggesting that similar experts and decision-makers in similar

conditions make extremely different judgement 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.⁷

1.2.2 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 . 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.⁸ 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.⁹

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 $\phi : \Omega \rightarrow \Delta(\mathcal{S})$, 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 .¹⁰

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

⁷Note that the evidence presented by Kahneman, Sibony, and Sunstein [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.

⁸To 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.

⁹For many of the results we assume that the population of agents is rich enough to represent the whole spectrum of viewpoints: $\mathcal{M} = \Delta(\Omega)$.

¹⁰Similar 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 observed the signal realization.

1.6.2 to show that our results do not depend on this particular specification). Namely, let $\eta : \mathcal{S} \rightarrow \Delta(\Omega)$ denote the agent's posterior belief system, obtained from μ and ϕ using the Bayes' rule. The cost is then defined as

$$c(\phi, \mu) \equiv \lambda \left(- \sum_{\omega \in \Omega} \mu(\omega) \ln \mu(\omega) + \sum_{\omega \in \Omega} \sum_{s \in \mathcal{S}} \left(\sum_{\omega' \in \Omega} \mu(\omega') \phi(s|\omega') \right) \eta(\omega|s) \ln \eta(\omega|s) \right), \quad (1.1)$$

where $\lambda \in \mathbb{R}_{++}$ is a cost parameter.¹¹ 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, Mathevet, and Wei, 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 Triole, 1990, Armstrong and Sappington, 2007, Alonso and Matouschek, 2008 for some examples and a discussion of such settings). Section 1.5 shows 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.

1.2.3 The Agent's Problem

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

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

¹¹We follow the standard convention and let $0 \ln 0 = 0$.

The agent's problem can then be written down as

$$\max_{\phi, \sigma} \left\{ \sum_{\omega \in \Omega} \mu(\omega) \sum_{s \in \mathcal{S}} \phi(s|\omega) u(\sigma(s), \omega) - c(\phi, \mu) \right\}. \quad (1.2)$$

Lemma 1 in Matějka and McKay [2015] shows that problem (1.2) with entropy cost function can be reframed as a problem of selecting a collection of conditional choice probabilities. This reformulation is presented in Section 1.2.6.

1.2.4 The Principal's Problem

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

$$\begin{aligned} \max_{\mu} \left\{ \sum_{\omega \in \Omega} \mu_p(\omega) \sum_{s \in \mathcal{S}} \phi_{\mu}(s|\omega) u(\sigma_{\mu}(s), \omega) \right\}, \\ \text{s.t. } (\phi_{\mu}, \sigma_{\mu}) \text{ solves (1.2) given } \mu, \end{aligned} \quad (1.3)$$

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.

1.2.5 Equilibrium Definition

We now present the equilibrium notion used throughout the paper; the discussion follows. [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 $\phi_{\mu}^* : \Omega \rightarrow \Delta(\mathcal{S})$ and choice rules $\sigma_{\mu}^* : \mathcal{S} \rightarrow \mathcal{A}$ for all $\mu \in \mathcal{M}$, such that:

1. ϕ_{μ}^* and σ_{μ}^* constitute a solution to (1.2) for every $\mu \in \mathcal{M}$;
2. μ^* is a solution to (1.3) given $(\phi_{\mu}^*, \sigma_{\mu}^*)$.

Note that the above effectively defines a Subgame-Perfect Nash 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 (1.2) and (1.3) allow us to treat these beliefs as just some exogenous functions entering the terminal payoff functions. This is primarily 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).

1.2.6 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 $\pi : \Omega \rightarrow \Delta(\mathcal{A})$ (which is a single state-contingent action distribution, as opposed to the combination of a signal strategy $\phi : \Omega \rightarrow \Delta(\mathcal{S})$ and a choice rule $\sigma : \mathcal{S} \rightarrow \mathcal{A}$):

$$\max_{\pi} \left\{ \sum_{\omega \in \Omega} \mu(\omega) \left(\sum_{a \in \mathcal{A}} \pi(a|\omega) u(a, \omega) \right) - c(\pi, \mu) \right\}, \quad (1.4)$$

where $c(\pi, \mu)$ denotes, with abuse of notation, the information cost induced by the action distribution π .¹² Lemma 2 in the online appendix of Matějka and McKay [2015] implies in our setting that 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):

$$\beta(a) \equiv \sum_{\omega \in \Omega} \mu(\omega) \pi(a|\omega). \quad (1.5)$$

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

$$\begin{aligned} \max_{\mu} \left\{ \sum_{\omega \in \Omega} \mu_p(\omega) \left(\sum_{a \in \mathcal{A}} \pi_{\mu}(a|\omega) u(a, \omega) \right) \right\} \\ \text{s.t. } \pi_{\mu} \text{ solves (1.4) given } \mu. \end{aligned} \quad (1.6)$$

In what follows, we refer to problem (1.6) 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.

1.3 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 the general model.¹³ We show that the principal has to balance 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 1.3).

¹²Cost $c(\pi, \mu)$ is calculated as the cost $c(\phi, \mu)$ of the cheapest strategy (ϕ, σ) that generates 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]. Given this, we have that

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

¹³Such a binary model is common in the delegation literature, see e.g. Li and Suen [2004] with a slightly different informal story.

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 the 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 1.2.6, the agent's problem can be reformulated as the problem of choosing a stochastic decision rule $\pi : \Omega \rightarrow \Delta(\mathcal{A})$, which solves

$$\max_{\pi} \left\{ \mu \pi(R|r) + (1 - \mu) \pi(L|l) - c(\pi, \mu) \right\}. \quad (1.7)$$

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

$$\pi(L|l) = \frac{\beta(L)e^{\frac{1}{\lambda}}}{\beta(L)e^{\frac{1}{\lambda}} + \beta(R)}, \quad \pi(R|r) = \frac{\beta(R)e^{\frac{1}{\lambda}}}{\beta(L) + \beta(R)e^{\frac{1}{\lambda}}}, \quad (1.8)$$

and their Corollary 2 implies that

$$\beta(R) = \frac{\mu e^{\frac{1}{\lambda}} - (1 - \mu)}{e^{\frac{1}{\lambda}} - 1}, \quad \beta(L) = \frac{e^{\frac{1}{\lambda}}(1 - \mu) - \mu}{e^{\frac{1}{\lambda}} - 1}, \quad (1.9)$$

cropped to $[0, 1]$. Combining (1.8) and (1.9), we get that the solution to problem (1.7) is given by¹⁴

$$\begin{aligned} \pi(R|r) &= \frac{(\mu e^{\frac{1}{\lambda}} - (1 - \mu)) e^{\frac{1}{\lambda}}}{(e^{\frac{2}{\lambda}} - 1) \mu}, \\ \pi(L|l) &= \frac{((1 - \mu) e^{\frac{1}{\lambda}} - \mu) e^{\frac{1}{\lambda}}}{(e^{\frac{2}{\lambda}} - 1) (1 - \mu)}, \end{aligned} \quad (1.10)$$

cropped to $[0, 1]$. Figure 1.1 demonstrates how the agent's action precisions choice depends on his prior belief.

In turn, the principal's problem is

$$\begin{aligned} \max_{\mu} \{ &\mu_p \pi_{\mu}(R|r) + (1 - \mu_p) \pi_{\mu}(L|l) \} \\ \text{s.t. } &\pi_{\mu} \text{ solves problem (1.7) given } \mu. \end{aligned} \quad (1.11)$$

It is easy to see by comparing the payoffs in (1.7) and (1.11) 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,

¹⁴This 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 1.1: Solution of problem (1.7) for different prior beliefs μ .

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 (1.10) depends on his prior belief μ .

When solving problem (1.7), 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 $\pi(L|l)$.¹⁵ 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 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 1.2 plots the principal's expected utility from hiring an agent as a function of the agent's belief μ when $\mu_p = 0.7$. We can see 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 is flat for very high and very low μ , which corresponds to the

¹⁵This 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].




Figure 1.2: Expected utility of principal with prior belief $\mu_p = 0.7$ as a function of the agent's prior belief μ .




Figure 1.3: The optimal delegation strategy μ^* as a function of the principal's prior belief μ_p .

agents who do not learn anything, and simply always choose the ex ante optimal action. Further, agents with low $\mu \sim (0.15, 0.2)$ acquire non-trivial information, but hiring them is worse for the principal than taking the ex ante optimal action (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.

Proposition 1.3 below formalizes this intuition and provides a closed-form solution for the optimal delegation strategy given the principal's prior belief μ_p . Figure 1.3 visualizes the optimal delegation strategy as a function of μ_p .

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

$$\mu^*(\mu_p) = \frac{\sqrt{\mu_p}}{\sqrt{\mu_p} + \sqrt{1 - \mu_p}}. \quad (1.12)$$

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

From Proposition 1.3 and Figure 1.3 we can immediately see that misalignment is the 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 1.3 below.

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

One thing to note about Proposition 1.3 is that the optimal delegation strategy




Figure 1.4: Action precisions under optimal delegation and delegating to the aligned agent.

(1.12) 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 1.3 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 1.6.2 does, however, suggest that this specific conclusion is likely an artifact produced by the entropy information cost function.

Figure 1.4 demonstrates the difference in the action precisions between delegating to a perfectly aligned agent ($\mu = \mu_p$) and the optimally misaligned agent as given by (1.12). 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. 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 overweighting small probabilities and underweighting large probabilities. They propose a probability weighting function that decision-makers unconsciously use, which is reminiscent of our optimal delegation strategy (1.12), 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 misper-

ceptions according to (1.12), since these maximize the survival probability.¹⁶

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

1.4 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 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 $\mathcal{A} \equiv \{a_1, \dots, a_N\}$ and $\Omega \equiv \{\omega_1, \dots, \omega_N\}$ 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 μ_p is such that $\mu_p(\omega_1) \geq \mu_p(\omega_2) \geq \dots \geq \mu_p(\omega_N)$ (otherwise states and actions can be relabeled as necessary). As before, results from Section 1.2.6 apply, meaning that the agent’s problem is equivalent to choosing the action distribution $\pi : \Omega \rightarrow \Delta(\mathcal{A})$ to maximize (1.4), and the principal selects an agent according to his prior $\mu \in \mathcal{M}$ to maximize (1.6). 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).

1.4.1 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:

$$\pi(a_i|\omega_j) = \frac{\beta(a_i)e^{\frac{u(a_i, \omega_j)}{\lambda}}}{\sum_{k=1}^N \beta(a_k)e^{\frac{u(a_k, \omega_j)}{\lambda}}}, \quad (1.13)$$

where $\beta(a_i)$, defined in (1.5), 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 (1.13) 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 an agent with a prior belief μ

¹⁶Steiner 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.

optimally chooses a decision rule that generates unconditional choice probabilities

$$\beta(a_i) = \max \left\{ 0, \frac{1}{\delta} \left(\frac{(K(\beta) + \delta)\mu(\omega_i)}{\sum_{j \in C(\beta)} \mu(\omega_j)} - 1 \right) \right\}, \quad (1.14)$$

where $\delta \equiv e^{\frac{1}{\lambda}} - 1$; $C(\beta) \equiv \{i \in \{1, \dots, N\} : \beta(a_i) > 0\}$ 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.

1.4.2 Principal's Relaxed Problem

As mentioned previously, (1.13) 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:

$$\max_{\beta} \left\{ \sum_{j=1}^N \mu_p(\omega_j) \left(\sum_{i=1}^N \frac{\beta(a_i) e^{\frac{u(a_i, \omega_j)}{\lambda}}}{\sum_{k=1}^N \beta(a_k) e^{\frac{u(a_k, \omega_j)}{\lambda}}} u(a_i, \omega_j) \right) \right\}. \quad (1.15)$$

In the above, we used (1.13) to represent the conditional probabilities $\pi(a_i|\omega_j)$ in (1.6) in terms of the unconditional probabilities $\beta(a_i)$. In Section 1.4.3 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 (1.15) to

$$\max_{\beta} \left\{ \sum_{j=1}^N \mu_p(\omega_j) \frac{\beta(a_j) e^{\frac{1}{\lambda}}}{1 + \delta \beta(a_j)} \right\}. \quad (1.16)$$

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

Lemma 1. *The solution to the principal's relaxed problem (1.16) is given by*

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

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 1.4.4. 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.

1.4.3 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 (1.6), 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 (1.4) 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 1.5.1 that it does not hold for arbitrary payoff functions.

1.4.4 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 1.4.4 extends Proposition 1.3 and shows how the chosen agent’s prior belief relates to that of the principal. Proposition 1.4.4 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^* .

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

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

In particular, $\mu^*(\omega_1) \geq \dots \geq \mu^*(\omega_N)$. Further, $\mu^*(\omega_1) \leq \mu_p(\omega_1)$ and $\mu^*(\omega_N) \geq \mu_p(\omega_N)$, with equalities if and only if $\mu_p(\omega_1) = \dots = \mu_p(\omega_j)$.

The intuition behind the result above is the same as that behind Proposition 1.3: 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 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 1.4.4 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)

$$\bar{\beta}(a_i) = \max \left\{ 0, \frac{1}{\delta} \left(\frac{(K(\bar{\beta}) + \delta)\mu(\omega_i)}{\sum_{j \in C(\bar{\beta})} \mu(\omega_j)} - 1 \right) \right\}. \quad (1.17)$$

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

$$\mu_p(\omega_{\bar{K}}) > \frac{1}{\bar{K} + \delta} \sum_{j=1}^{\bar{K}} \mu_p(\omega_j) \geq \mu_p(\omega_{\bar{K}+1}). \quad (1.18)$$

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

$$\sqrt{\mu_p(\omega_{K^*})} > \frac{1}{K^* + \delta} \sum_{j=1}^{K^*} \sqrt{\mu_p(\omega_j)} \geq \sqrt{\mu_p(\omega_{K^*+1})}. \quad (1.19)$$

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

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.

1.5 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 1.3 and 1.4, 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.

1.5.1 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 $\tau : \mathcal{A} \rightarrow \mathbb{R}$, 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 1.¹⁷

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 1.2.6) is then given by

$$\max_{\pi} \left\{ \sum_{j=1}^N \mu_p(\omega_j) \sum_{i=1}^N \pi(a_i | \omega_j) \left(u(a_i, \omega_j) + \tau(a_i) \right) - c(\phi, \mu_p) \right\}, \quad (1.20)$$

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

$$\begin{aligned} \max_{\tau} \left\{ \sum_{j=1}^N \mu_p(\omega_j) \sum_{i=1}^N \pi(a_i | \omega_j) \left(u(a_i, \omega_j) - \rho \tau(a_i) \right) \right\} \\ \text{s.t. } \pi \text{ solves (1.20) given } \tau. \end{aligned} \quad (1.21)$$

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

¹⁷In 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.

belief μ . In particular, Lemma 2 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 Mikhailishchev [2021] we can also show the converse – that any decision rule achievable with misaligned beliefs can be replicated with payments τ (or by setting the quotas, i.e., imposing specific unconditional choice probabilities for a different action). These results are formalized by the following theorem.¹⁸

The principal’s problem of contracting on actions (1.21) is equivalent to her full (delegation) problem (1.6):

1. For any vector $\tau : \mathcal{A} \rightarrow \mathbb{R}$ of payments/biases and a corresponding $\beta : \Omega \rightarrow \Delta(\mathcal{A})$ that solves (1.20) given τ , there exists a prior belief $\mu \in \Delta(\Omega)$ such that β also solves (1.4) given μ .
2. For any $\mu \in \Delta(\Omega)$ and the corresponding $\beta : \Omega \rightarrow \Delta(\mathcal{A})$ that solves (1.4) given μ , there exist payments $\tau : \mathcal{A} \rightarrow \mathbb{R}$ such that β also solves (1.20) 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. Further, 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 Mikhailishchev [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 appropriate prior belief. The following proposition presents one example of such equivalence, in the context of a model with $N = 2$.

Consider the binary setting of Section 1.3. Consider the principal’s problem of contracting on actions (1.21), 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 1.3;

¹⁸The result regarding quotas is not included in the theorem, yet it follows immediately from Lemma 1 of Matveenko and Mikhailishchev [2021].

2. these payments/biases are such that¹⁹

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

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 1.3, 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).²⁰ 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.

1.5.2 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 $i \neq j$.²¹ 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 $\pi : \Omega \rightarrow \Delta(\mathcal{A})$ that solves²²

$$\max_{\pi} \left\{ \sum_{j=1}^N \mu(\omega_j) \sum_{i=1}^N \pi(a_i | \omega_j) \left(u(a_i, \omega_j) + \tau(a_i, \omega_j) \right) - c(\phi, \mu) \right\}, \quad (1.22)$$

¹⁹The closed-form expressions are available in the proof in the Appendix.

²⁰This 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).

²¹If 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.

²²While 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.

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

$$\begin{aligned} & \max_{\bar{\tau}, \underline{\tau}} \left\{ \sum_{j=1}^N \mu_p(\omega_j) \sum_{i=1}^N \pi(a_i | \omega_j) (u(a_i, \omega_j) - \rho \tau(a_i, \omega_j)) \right\}, & (1.23) \\ & \text{s.t. } \tau(a_i, \omega_i) = \bar{\tau} \text{ for all } i, \\ & \quad \tau(a_i, \omega_j) = \underline{\tau} \text{ for all } i, j \neq i, \\ & \quad \pi \text{ solves (1.22) given } \bar{\tau}, \underline{\tau}. \end{aligned}$$

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.

Consider the principal's contracting problem (1.23) in the binary setting of Section 1.3 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 (1.12);
2. the principal's problem (1.23) 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 (1.23) 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.

1.5.3 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, there are no “safe” actions that the principal could rule out, as Proposition 1.4.4 suggests. Assuming that the principal and the agent hold the same prior belief μ_p , and $\mu_p(\omega_1) > \dots > \mu_p(\omega_N)$, action a_1 is the “safest” in the sense of being the most likely ex ante to be optimal. However, it would be trivially suboptimal 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 1.5.3 below summarizes this logic. Consider the agent’s problem as given by

$$\max_{\pi} \left\{ \sum_{j=1}^N \mu_p(\omega_j) \sum_{i=1}^N \pi(a_i|\omega_j) u(a_i, \omega_j) - c(\phi, \mu_p) \right\}, \quad (1.24)$$

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

$$\begin{aligned} \max_{A^*} \left\{ \sum_{j=1}^N \mu_p(\omega_j) \sum_{i=1}^N \pi(a_i|\omega_j) u(a_i, \omega_j) \right\}, \\ \text{s.t. } \pi : \Omega \rightarrow \Delta(A^*) \text{ solves (1.24) given } A^*. \end{aligned} \quad (1.25)$$

Then we can state the result as follows.

The unrestricted delegation set $A^* = \mathcal{A}$ is always a solution to the principal’s restriction problem (1.25).

1.6 Extensions

1.6.1 Alternative Preference Specifications

The analysis in Sections 1.4 and 1.5.1 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. Since the utility function $u(a, \omega)$ is shared by both the principal and the agent, it is reasonable to generalize one at a time.

We begin by generalizing 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 1.3, Theorem 1.4.4, 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 1.5.1 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.

The 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 1.5.1 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 1.4.2 breaking down with general preferences, as stated by the following proposition.

There exists a utility function $u(a, \omega)$ such that the solution to the principal’s relaxed problem (1.16) cannot be attained as a solution to the full problem (1.6). There exists a non-state-matching utility function $u(a, \omega)$ such that the conclusions of Theorem 1.5.1 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 1.5.1 (see Proposition 3 in Matveenko and Mikhailishchev, 2021).

1.6.2 Alternative cost functions

Our analysis uses the Shannon entropy cost function (1.1) 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]. To demonstrate that our main result does not hinge on this or any other specific properties of the entropy cost parametrization, this section explores three alternative specifications of the cost function in the binary setting of Section 1.3. We show that, in all cases, the principal’s optimal delegation strategy looks similar to what we obtain in Proposition 1.3: unless μ_p is too extreme, it is optimal for the principal to delegate to an ex ante more unbiased agent: $\mu^*(\mu_p) \in [0.5, \mu_p)$.

As noted, in the binary setting of Section 1.3, 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 look for an equilibrium as defined in Section 1.2.5, with the $c(\phi, \mu)$ in the agent’s problem (1.4)




Figure 1.5: The optimal delegation strategy $\mu^*(\mu_p)$ with channel capacity cost function.

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 Nirmark and Sundaresan [2019], hereinafter referred to as NS. The channel capacity cost of a given signal structure $\phi : \Omega \rightarrow \Delta(\mathcal{S})$ is given by

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

where $c(\phi, \mu)$ is the entropy cost (1.1). Intuitively, the channel capacity measures the maximum amount of information that can be extracted from signal ϕ by any agent. By definition, the 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 1.2.6 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 $\pi : \Omega \rightarrow \Delta(\mathcal{A})$ which solves

$$\max_{\pi} \left\{ \mu \pi(R|r) + (1 - \mu) \pi(L|l) - c_C(\pi, \mu) \right\}, \quad (1.26)$$

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 (1.26) is well-defined and a 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 (1.11) 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 1.5 presents numerical solutions for two values of λ and all μ_p .

Both plots in Figure 1.5 demonstrate a delegation strategy that is qualitatively the same as in Figure 1.3, which plotted the same strategy 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 1.5 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.

Log-likelihood ratio cost function. We now move on to consider the log-likelihood ratio (LLR) cost function proposed by Pomatto, Strack, and Tamuz [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 structure $\phi : \Omega \rightarrow \Delta(\mathcal{S})$ 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 1.2.6 and reduce the agent’s problem to that of choosing a decision rule $\pi : \Omega \rightarrow \Delta(\mathcal{A})$ subject to cost $c_L(\pi)$.²³

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 1.6 demonstrates our findings. The principal’s optimal delegation strategy plotted therein looks qualitatively the same as for entropy and channel capacity costs (Figures 1.3 and 1.5, 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 1.5, the principal delegates less for higher values of the information cost parameter λ .

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 in 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 $\phi : \Omega \rightarrow \Delta(\mathcal{S})$, the agent is restricted to a binary signal space $\mathcal{S} \equiv \{l, r\}$ and can only choose the 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

²³PTS provide a representation for $c_L(\pi)$ not presented here.




Figure 1.6: The optimal delegation strategy $\mu^*(\mu_p)$ with LLR cost function.

in $\phi \in [1/2, 1]$, and satisfies $c_S(1/2) = 0$ and the Inada conditions $c'_S(1/2) = 0$, $\lim_{\phi \rightarrow 1} c'_S(\phi) = +\infty$.

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 $\sigma(l) = L$), because this trivially dominates doing the converse, and conditional on ignoring the signal, acquiring an uninformative signal $\phi = 1/2$ is strictly cheaper. Hence if $\phi_\mu^* > 1/2$, then

$$\phi_\mu^* = \arg \max_{\phi} \{ \mu\phi + (1 - \mu)\phi - c_S(\phi) \}. \quad (1.27)$$

Let ϕ^{**} denote the candidate solution given by the FOC of (1.27): $c'_S(\phi^{**}) = 1$. 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 $\phi^{**} - c_S(\phi^{**}) \in [1/2, 1]$.²⁴ 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

$$\mu^*(\mu_p) = \begin{cases} \mu \in [\bar{\mu}_L, \bar{\mu}_R] & \text{if } \mu_p \leq \phi^{**}, \\ \mu \in (\bar{\mu}_R, 1] & \text{if } \mu_p > \phi^{**}. \end{cases} \quad (1.28)$$

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 (1.28) that supports the following proposition (which is proved by the preceding argument): Given a symmetric information cost function $c_S(\phi)$,

²⁴Function $\phi - c_S(\phi)$ is strictly concave in ϕ due to the assumptions made, evaluates to $1/2$ when $\phi = 1/2$ (hence $\phi^{**} - c_S(\phi^{**}) \geq 1/2$), and because $c_S(\phi) \geq 0$, we have that $\phi - c_S(\phi) \leq 1$ for all $\phi \leq 1$.

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 1.6.2 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 $\mu \in (\bar{\mu}_R, 1]$, because, 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 equally 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 will 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’s 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 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 fully capture 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.

1.6.3 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, Dessein, and Matouschek, 2008, Rantakari, 2008 for some examples). We show that in our setting, communication performs 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.²⁵ We show below that, in the end, 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 1.2 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.²⁶ The equilibrium of the communication game is then defined as follows. [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 $\mu_p : \mathcal{A} \rightarrow \Delta(\Omega)$ that are consistent with $(\phi_\mu^*, \sigma_\mu^*)$ (i.e., satisfy Bayes’ rule on the equilibrium path);
2. the principal’s choice rule $\sigma^* : \mathcal{A} \rightarrow \mathcal{A}$, 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 $\phi_\mu^* : \Omega \rightarrow \Delta(\mathcal{S})$ and communication strategies $\tilde{\sigma}_\mu^* : \mathcal{S} \rightarrow \mathcal{A}$ that solve the following given σ for every $\mu \in \mathcal{M}$:

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

²⁵Argenziano, Severinov, and Squintani [2016] provide one example of how the principal can manipulate the agent’s information acquisition incentives under cheap talk communication.

²⁶Given 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.

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.

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.

The result is, perhaps, unsurprising, since Holmström [1980] showed that communication is equivalent to restricting the agent’s action set, and this latter instrument was shown in Section 1.5.3 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 1.6.3, 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 1.6.3 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.²⁷ 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 1.6.3 relies on state-matching preferences. In our setting (with the exception of Section 1.6.1), 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

²⁷If 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.

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 1.6.3 that the same equilibrium would survive in a setting with *verifiable communication* à la 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 1.6.3, 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.

1.7 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. A biased principal prefers to delegate to an agent who is *ex ante* more uncertain about what the best action is, but who is somewhat biased towards the same action as the principal. 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 a 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 greatest 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 the exact trade-offs obviously do depend on the particular cost function specification, we do show that, qualitatively, our results are not specific to the entropy information cost.

Due to the added complexity of entropy models, we confine our exploration to a discrete state-matching model, which strays from the continuous models more commonly used in delegation problems. In a model with continuous actions, 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 again be different. Exploration of the effects of misalignment in a continuous model of delegated expertise could be an interesting direction for further research.

Yet another 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 and that they make inferences from the fact that they were chosen for the job. Such signaling concerns could yield an economically meaningful effect, but we abstract from them completely in our paper. A more careful investigation is in order.

1.A Main Proofs

1.A.1 Proof of Proposition 1.3

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 $\mu^* \in (\frac{1}{2}, \mu_p)$, 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 (1.10) (assuming this solution is interior for now) into the principal's problem (1.11), we get that the principal's payoff looks as follows:

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

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

$$\begin{aligned} \frac{\mu_p}{\mu^2} - \frac{1 - \mu_p}{(1 - \mu_p)^2} &= 0 \\ \iff \frac{\mu}{1 - \mu} &= \frac{\sqrt{\mu_p}}{\sqrt{1 - \mu_p}}. \end{aligned} \tag{1.29}$$

It is trivial to verify that the second-order condition holds as well, hence as long as (1.29) yields an interior solution (i.e., the probabilities in (1.10) are in $[0, 1]$), the candidate solution is indeed optimal among all such interior solutions.

We now check for which μ the solution (1.10) is interior. Using the expressions (1.10), 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 $\frac{\mu}{1 - \mu} \in [e^{-\frac{1}{\lambda}}, e^{\frac{1}{\lambda}}]$, 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 $\frac{\mu_p}{1-\mu_p} \in [e^{-\frac{2}{\lambda}}, e^{\frac{2}{\lambda}}]$, 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)

$$\frac{e^{\frac{1}{\lambda}}}{e^{\frac{2}{\lambda}} - 1} \left[e^{\frac{1}{\lambda}} - \mu_p \frac{1 - \mu^*}{\mu^*} - (1 - \mu_p) \frac{\mu^*}{1 - \mu^*} \right] = \frac{e^{\frac{1}{\lambda}}}{e^{\frac{2}{\lambda}} - 1} \left[e^{\frac{1}{\lambda}} - 2\sqrt{\mu_p(1 - \mu_p)} \right]. \quad (1.30)$$

Taking the difference between (1.30) 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 1.3.

1.A.2 Proof of Corollary 1.3

Proof of Proposition 1.3 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 (1.12) we then obtain

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

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.

1.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 (1.16). First, let us rewrite expression (1.16) using $\delta \equiv e^{\frac{1}{\lambda}} - 1$:

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

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

$$\min_{\beta} \sum_{j \in C(\beta)} \frac{\mu_p(\omega_j)}{1 + \delta \beta(a_j)}. \quad (1.31)$$

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

$$(1 + \delta \beta(a_i))^2 = -\frac{\mu_p(\omega_i)}{\xi}. \quad (1.32)$$

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

$$\sum_{j \in C(\beta)} (1 + \delta \beta(a_j))^2 = -\frac{\sum_{j \in C(\beta)} \mu_p(\omega_j)}{\xi} \quad (1.33)$$

Combining (1.32) and (1.33):

$$1 + \delta \beta(a_i) = \frac{\sqrt{\mu_p(\omega_i)}}{\sqrt{\sum_{j \in C(\beta)} \mu_p(\omega_j)}} \sqrt{\sum_{j \in C(\beta)} (1 + \delta \beta(a_j))^2} \quad (1.34)$$

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 (1.34) allows us to express $\beta(a_i)$ (for $i \in C(\beta)$) in closed form as

$$\beta(a_i) = \frac{1}{\delta} \left(\frac{(K(\beta) + \delta) \sqrt{\mu_p(\omega_i)}}{\sum_{j \in C(\beta)} \sqrt{\mu_p(\omega_j)}} - 1 \right). \quad (1.35)$$

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

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

Plugging in ξ from (1.32) 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:

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

Recall that we assumed, without loss of generality, that $\mu_p(\omega_1) \geq \mu_p(\omega_2) \geq \dots \geq \mu_p(\omega_N)$. 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 $\Delta_L \equiv (L + \delta)\sqrt{\mu_p(\omega_L)} - \sum_{j=1}^L \sqrt{\mu_p(\omega_j)}$. Notice that for all $L > 1$:

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

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 (1.35) if $i \in C(\beta^*)$, $\beta^*(a_i) = 0$ if $i \notin C(\beta^*)$, and $C(\beta^*) = 1, \dots, K(\beta^*)$, where $K(\beta^*)$ is as described above.

1.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 (1.13) only if it solves the system of equations given by

$$\sum_{j=1}^N \mu(\omega_j) \frac{e^{\frac{u(a_j, \omega_j)}{\lambda}}}{\sum_{k=1}^N \beta(a_k) e^{\frac{u(a_k, \omega_j)}{\lambda}}} = 1, \quad (1.36)$$

for every $i \in \{1, \dots, N\}$ 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:

$$\begin{cases} \mu(\omega_1) + \mu(\omega_2) + \dots + \mu(\omega_N) = 1, \\ \sum_{j=1}^N \mu(\omega_j) \frac{e^{\frac{u(a_i, \omega_j)}{\lambda}}}{\sum_{k=1}^N \beta(a_k) e^{\frac{u(a_k, \omega_j)}{\lambda}}} = 1 \quad \forall i \in C(\beta). \end{cases} \quad (1.37)$$

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:

$$\begin{cases} y_0 \left(\sum_{k=1}^N \beta(a_k) e^{\frac{u(a_k, \omega_j)}{\lambda}} \right) + \left(\sum_{i \in C(\beta)} y_i e^{\frac{u(a_i, \omega_j)}{\lambda}} \right) \geq 0 \quad \forall j \in \{1, \dots, N\}, \\ y_0 + \sum_{i \in C(\beta)} y_i < 0. \end{cases} \quad (1.38)$$

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 (1.38) transforms to

$$\begin{cases} z_j e^{\frac{1}{\lambda}} + \sum_{i \in C(\beta) \setminus \{j\}} z_i \geq 0 \quad \forall j \in C(\beta), \\ \sum_{i \in C(\beta)} z_i \geq 0 \quad \forall j \in \{1, \dots, N\} \setminus C(\beta), \\ \sum_{i \in C(\beta)} z_i < 0. \end{cases} \quad (1.39)$$

System (1.39) above does not have a solution. Indeed, if $C(\beta) = \{1, \dots, N\}$, then the middle set of inequalities directly contradicts the latter inequality. If $C(\beta) \neq \{1, \dots, N\}$, 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, $\sum_{i \in C(\beta)} z_i < 0$.

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 (1.37). This concludes the proof.

1.A.5 Proof of Theorem 1.4.4

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

$$\mu(\omega_i) = \frac{\sqrt{\mu_p(\omega_i)}}{\sum_{j=1}^N \sqrt{\mu_p(\omega_j)}}. \quad (1.40)$$

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 (1.40) into (1.14) yields

$$\beta(a_i) = \max \left\{ 0, \frac{1}{\delta} \left(\frac{(K(\beta^*) + \delta) \sqrt{\mu_p(\omega_i)}}{\sum_{j \in C(\beta^*)} \sqrt{\mu_p(\omega_j)}} - 1 \right) \right\}, \quad (1.41)$$

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. Firstly, it follows clearly from (1.40) that $\mu^*(\omega_1) \geq \mu^*(\omega_2) \geq \dots \geq \mu^*(\omega_N)$. It remains 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:

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

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

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

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

which holds because $\sqrt{\mu_p(\omega_j)\mu_p(\omega_N)} \leq \mu_p(\omega_j)$ for all $j \in \{1, \dots, N\}$, with equalities only if $\mu_p(\omega_1) = \dots = \mu_p(\omega_N)$. This concludes the proof of Theorem 1.4.4.

1.A.6 Proof of Proposition 1.4.4

It follows from (1.18) 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,

$$\sum_{j=1}^{\bar{K}} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{\bar{K}})}} < \bar{K} + \delta. \quad (1.42)$$

From (1.19), K^* is the unique solution of

$$\sum_{j=1}^{K^*} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{K^*})}} < K^* + \delta \leq \sum_{j=1}^{K^*} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{K^*+1})}}. \quad (1.43)$$

Two cases are possible, depending on whether

$$\bar{K} + \delta \sum_{j=1}^{\bar{K}} \frac{\sqrt{\mu_p(\omega_j)}}{\sqrt{\mu_p(\omega_{\bar{K}+1})}}. \quad (1.44)$$

If $\bar{K} + \delta \leq RHS$ in (1.44) (where RHS refers to the right-hand side), then together with (1.42) this implies that \bar{K} solves (1.43), and thus $\bar{K} = K^*$, which satisfies that statement of the proposition.

If, however, $\bar{K} + \delta > RHS$ in (1.44), then \bar{K} does not solve (1.43). In this case, note that going from K by $K + 1$ increases the LHS of (1.44) 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 (1.44), 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 (1.43) must be such that $K_M > \bar{K}$. This concludes the proof.

1.A.7 Proof of Theorem 1.5.1

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 (1.13), which claims that in the contracting problem, the $\beta : \Omega \rightarrow \Delta(\mathcal{A})$ that solves the agent's problem (1.20) is given by

$$\begin{aligned} \pi(a_i|\omega_j) &= \frac{\beta(a_i) e^{\frac{u(a_i, \omega_j) + \tau(a_i)}{\lambda}}}{\sum_{k=1}^N \beta(a_k) e^{\frac{u(a_k, \omega_j) + \tau(a_k)}{\lambda}}} \\ &= \frac{\beta'(a_i) e^{\frac{u(a_i, \omega_j)}{\lambda}}}{\sum_{k=1}^N \beta'(a_k) e^{\frac{u(a_k, \omega_j)}{\lambda}}}, \end{aligned} \quad (1.45)$$

$$\text{where } \beta(a_i) = \sum_{j=1}^N \mu(\omega_j) \pi(a_i|\omega_j).$$

$$\text{and } \beta'(a_i) \equiv \frac{\beta(a_i) e^{\frac{\tau(a_i)}{\lambda}}}{\sum_{k=1}^N \beta(a_k) e^{\frac{\tau(a_k)}{\lambda}}}.$$

Since β' is a valid probability distribution on \mathcal{A} , representation (1.45) together with (1.13) imply that such a collection of conditional probabilities π is a valid solution to the agent's problem (1.4) 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.

1.A.8 Proof of Proposition 1.5.1

Plugging (1.12) in (1.10) yields the optimal conditional choice probabilities for the binary model, given by

$$\begin{aligned}\pi^*(R|r) &= \left(e^{\frac{2}{\lambda}} - 1\right)^{-1} e^{\frac{1}{\lambda}} \left(e^{\frac{1}{\lambda}} - \sqrt{\frac{1-\mu_p}{\mu_p}}\right), \\ \pi^*(L|l) &= \left(e^{\frac{2}{\lambda}} - 1\right)^{-1} e^{\frac{1}{\lambda}} \left(e^{\frac{1}{\lambda}} - \sqrt{\frac{\mu_p}{1-\mu_p}}\right),\end{aligned}\tag{1.46}$$

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 (1.20). Solving it given $\tau = (\tau(R), 0)$ yields

$$\begin{aligned}\pi(R|r) &= 1 - \frac{e^{\frac{2}{\lambda}}(1-\mu) - e^{\frac{1+\tau(R)}{\lambda}} + \mu}{\left(e^{\frac{2}{\lambda}} - 1\right) \left(e^{\frac{1+\tau(R)}{\lambda}} - 1\right) \mu}, \\ \pi(L|l) &= \frac{e^{\frac{1}{\lambda}} \left(e^{\frac{2}{\lambda}}(1-\mu) - e^{\frac{1+\tau(R)}{\lambda}} + \mu\right)}{\left(e^{\frac{2}{\lambda}} - 1\right) \left(e^{\frac{1}{\lambda}} - e^{\frac{\tau(R)}{\lambda}}\right) (1-\mu)},\end{aligned}\tag{1.47}$$

cropped to $[0, 1]$.

The principal's contracting problem (1.21) in the binary setting with $\rho = 0$ is similar to (1.11):

$$\begin{aligned}\max_{\tau(R)} & \{\mu_p \pi(R|r) + (1-\mu_p) \pi(L|l)\} \\ \text{s.t.} & \pi(R|r), \pi(L|l) \text{ are given by (1.47)}.\end{aligned}\tag{1.48}$$

Assuming the probabilities in (1.47) are interior, the F.O.C. for (1.48) yields the candidate solution $\tau(R)$ given by

$$\tau^*(R) = \lambda \ln \left[\frac{\frac{1-\mu}{\mu} e^{\frac{1}{\lambda}} + \sqrt{\frac{1-\mu_p}{\mu_p}}}{\frac{1-\mu}{\mu} + e^{\frac{1}{\lambda}} \sqrt{\frac{1-\mu_p}{\mu_p}}} \right],\tag{1.49}$$

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 (1.47).

Plugging (1.49) into (1.47) yields, after some routine manipulations, the conditional choice probabilities that coincide with (1.46) (hence, the probabilities (1.47) are interior given μ and $\tau^*(R)$ if and only if the probabilities (1.46) are interior). Thus, the condition (1.49) is not only necessary, but also sufficient. Hence, for any μ_p for which (1.46) are interior, $\tau^*(R)$ as given by (1.49) solves the principal's problem (1.48), and this solution exists for any μ .

If λ and μ_p are such that probabilities (1.46) 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 (1.46) are such that $\pi^*(R|r) \geq 1 \iff \pi^*(L|l) \leq 0$ and vice versa). The candidate transfers (1.49) yield exactly such non-interior probabilities (when plugged into (1.47)), and hence they still solve the principal's problem (1.48) for any respective μ .²⁸ This concludes the proof of part 1 of the proposition.

To show part 2, consider (1.49) 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 $\tau(R) \geq 0 \iff \mu \leq \mu^*$.

1.A.9 Proof of Proposition 1.5.2

As argued in the text, it is immediate that $\underline{\tau} = 0$. Proceeding analogously to Section 1.3, we obtain that the agent's problem (1.22) given the incentive payment $\bar{\tau} \geq 0$ is solved by

$$\begin{aligned} \pi(R|r) &= \min \{1, \max \{0, \pi_u(R|r)\}\} \quad \text{and} \quad \pi(L|l) = \min \{1, \max \{0, \pi_u(L|l)\}\}, \\ \text{where } \pi_u(R|r) &= \frac{e^{\frac{1+\bar{\tau}}{\lambda}} \left(e^{\frac{1+\bar{\tau}}{\lambda}} \mu - (1 - \mu) \right)}{\left(e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1 \right) \mu} = \frac{e^{\frac{1+\bar{\tau}}{\lambda}}}{e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1} \left(e^{\frac{1+\bar{\tau}}{\lambda}} - \frac{1 - \mu}{\mu} \right), \\ \pi_u(L|l) &= \frac{e^{\frac{1+\bar{\tau}}{\lambda}} \left(e^{\frac{1+\bar{\tau}}{\lambda}} (1 - \mu) - \mu \right)}{\left(e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1 \right) (1 - \mu)} = \frac{e^{\frac{1+\bar{\tau}}{\lambda}}}{e^{\frac{2(1+\bar{\tau})}{\lambda}} - 1} \left(e^{\frac{1+\bar{\tau}}{\lambda}} - \frac{\mu}{1 - \mu} \right). \end{aligned} \tag{1.50}$$

The principal's full contracting problem (1.23) can be rewritten as

$$\begin{aligned} \max_{\bar{\tau} \in \mathbb{R}_+} & \left\{ (1 - \rho \bar{\tau}) \left(\mu_p \pi(R|r) + (1 - \mu_p) \pi(L|l) \right) \right\}, \\ \text{s.t. } & \pi(R|r), \pi(L|l) \text{ are given by (1.50).} \end{aligned} \tag{1.51}$$

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.

²⁸Note 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.

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

$$\begin{aligned} \max_{\bar{\tau} \in \mathbb{R}} & \left\{ (1 - \rho\bar{\tau}) \left(\mu_p \pi_u(R|r) + (1 - \mu_p) \pi_u(L|l) \right) \right\}, \\ \text{s.t. } & \pi_u(R|r), \pi_u(L|l) \text{ are given by (1.50)}. \end{aligned} \quad (1.52)$$

It differs from the full problem (1.51) 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 (1.51) 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):²⁹

Feasibility: τ^{**} exists and $\tau^{**} \in [0, 1/\rho]$.³⁰

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

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

The FOC of problem (1.52) (that must be solved by τ^{**}) is given by

$$\mu_p \frac{1 - \mu}{\mu} + (1 - \mu_p) \frac{\mu}{1 - \mu} = e^{\frac{1+\bar{\tau}}{\lambda}} \cdot \frac{\lambda \rho \left(e^{2\frac{1+\bar{\tau}}{\lambda}} - 1 \right) + 2(1 - \rho\bar{\tau})}{\lambda \rho \left(e^{2\frac{1+\bar{\tau}}{\lambda}} - 1 \right) + \left(e^{2\frac{1+\bar{\tau}}{\lambda}} + 1 \right) (1 - \rho\bar{\tau})}. \quad (1.53)$$

Let $\gamma(\mu, \mu_p)$ denote the LHS and $\chi(\bar{\tau})$ the RHS of (1.53), 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 (1.51) is strictly concave in $\bar{\tau}$, so if τ^{**} exists, then it is unique and it is a local maximizer of (1.51).

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

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

²⁹Feasibility 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.

³⁰Note 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.

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

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

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

$$2\left(\xi^2 - 1 + 2\frac{1 - \rho\tau}{\lambda\rho}\right) + \frac{1 - \rho\tau}{\lambda\rho} \cdot \left(\xi^2 - 1 - 2\frac{1 - \rho\tau}{\lambda\rho}\right) > 0. \quad (1.54)$$

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 (1.54) 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 (1.54) 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. \square

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 $\mu_{L2} \equiv \max\{\mu : \pi_\mu^0(L|l) = 1\}$, $\mu_{R2} \equiv \min\{\mu : \pi_\mu^0(R|r) = 1\}$, where $\pi_\mu^0(L|l)$ and $\pi_\mu^0(R|r)$ stand for the respective probabilities (1.50) 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 (1.50), 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 $\mu_{L4} \equiv \max\{\mu : \pi_\mu^*(L|l) = 1\}$, $\mu_{R4} \equiv \min\{\mu : \pi_\mu^*(R|r) = 1\}$, where $\pi_\mu^*(L|l)$ and $\pi_\mu^*(R|r)$ stand for the respective probabilities (1.50) 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 exist 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

$$\mu \in [\mu_{L5}, \mu_{L1}] \cup [\mu_{R1}, \mu_{R5}]. \quad (1.55)$$

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 (1.12), and $\sup_{\mu} \gamma(\mu, \mu_p) = +\infty$ achieved by $\mu \rightarrow \{0, 1\}$. Hence a sufficient condition for the cutoffs of interest to exist is

$$\chi(0) \geq \gamma(\mu^*(\mu_p), \mu_p). \quad (1.56)$$

In the inequality above, only $\chi(0)$ depends on ρ . Note further that $\frac{d\chi(0)}{d\rho} > 0$. Therefore, if (1.56) 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

$$\begin{aligned} \chi(0) &= \xi \frac{\lambda\rho(\xi^2 - 1) + 2}{\lambda\rho(\xi^2 - 1) + (\xi^2 + 1)} \geq 1 \\ &\iff (\xi - 1)^2 (\lambda\rho(\xi + 1) - 1) \geq 0 \end{aligned} \quad (1.57)$$

Since $\xi = e^{\frac{1}{\lambda}} \geq 1$, a sufficient condition for (1.57) is given by

$$\rho \geq \frac{1}{\lambda(\xi + 1)}, \quad (1.58)$$

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 (1.58) 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 (1.56) is satisfied and the relevant cutoffs exist. \square

Moving on to *effectiveness*, it should be immediate from the analysis in Section 1.3 that for a feasible τ^{**} to yield interior choice probabilities (1.50), 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. \square

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

$$\mathbb{E}[u(a, \omega) \mid \mu_p, \tau^{**}] \equiv (1 - \rho\tau^{**}(\mu)) \left(\mu_p \pi_\mu^*(R|r) + (1 - \mu_p) \pi_\mu^*(L|l) \right) = 1 - \mu_p. \quad (1.59)$$

The LHS of (1.59) 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 $\mu = \mu_{L2}$, (1.53) 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 (1.23) 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 (1.55) 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 $\mu_{L5} \leq \mu_{L4} \leq \mu_{L3} \leq \mu_{L2}$ (and the converse holds for the other set of cutoffs), as well as $\mu_{L1} < \mu^* \leq \mu_p < \mu_{R1}$. 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.

1.A.10 Proof of Proposition 1.5.3

Using Theorem 1 of Caplin et al. [2019], the agent's problem (1.24) given some restriction set A^* is solved by π such that the corresponding $\beta \in \Delta(A^*)$ satisfies (1.17) for all $a_i \in A^*$. Further, recall from Section 1.4 that π and β are connected in the optimum by relation (1.13) (where we set $\pi(a_i \mid \omega_j) \equiv \beta(a_i) \equiv 0$ for all $a_i \notin A^*$ and all $\omega_j \in \Omega$). Then by plugging (1.13) and the state-matching utility into the principal's expected payoff, it can be rewritten as in (1.16):

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

Plugging in (1.17) for β in the expression above transforms it to

$$\begin{aligned} \sum_{i \in C(\beta)} \frac{\frac{1+\delta}{\delta} \mu_p(\omega_i) \left[(K(\beta) + \delta) \mu_p(\omega_i) - \sum_{j \in C(\beta)} \mu(\omega_j) \right]}{(K(\beta) + \delta) \mu_p(\omega_i)} \\ = \frac{1 + \delta}{\delta} \left[\sum_{i \in C(\beta)} \mu_p(\omega_i) - \sum_{i \in C(\beta)} \frac{\sum_{j \in C(\beta)} \mu(\omega_j)}{(K(\beta) + \delta)} \right] \\ = \frac{1 + \delta}{K(\beta) + \delta} \sum_{i \in C(\beta)} \mu_p(\omega_i). \quad (1.60) \end{aligned}$$

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 $k \in \{1, \dots, N\}$ and another action set $A_+ \equiv A_- \cup \{a_k\}$. Let β_+ denote the unconditional choice probabilities corresponding to the solution of (1.24) 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 (1.60) for the principal's expected payoff, the statement amounts to:

$$\begin{aligned}
0 &\leq \left(\frac{1+\delta}{K_+ + \delta} \sum_{i \in C_+} \mu_p(\omega_i) \right) - \left(\frac{1+\delta}{K_- + \delta} \sum_{i \in C_-} \mu_p(\omega_i) \right) \\
\iff 0 &\leq \left((K_- + \delta) \sum_{i \in C_+} \mu_p(\omega_i) \right) - \left((K_+ + \delta) \sum_{i \in C_-} \mu_p(\omega_i) \right) \\
\iff 0 &\leq (K_- + \delta) \mu_p(\omega_k) - \left(\sum_{i \in C_-} \mu_p(\omega_i) \right). \tag{1.61}
\end{aligned}$$

Since $a_k \in C_+$ by assumption, $\beta_+(a_k) > 0$, which, from (1.17), 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 (1.61) 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.

1.A.11 Proof of Proposition 1.6.1

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

$$\begin{cases}
z_1 e^{\frac{u(a_1, \omega_1)}{\lambda}} + z_2 e^{\frac{u(a_1, \omega_2)}{\lambda}} + z_3 e^{\frac{u(a_1, \omega_3)}{\lambda}} \geq 0, \\
z_1 e^{\frac{u(a_2, \omega_1)}{\lambda}} + z_2 e^{\frac{u(a_2, \omega_2)}{\lambda}} + z_3 e^{\frac{u(a_2, \omega_3)}{\lambda}} \geq 0, \\
z_1 e^{\frac{u(a_3, \omega_1)}{\lambda}} + z_2 e^{\frac{u(a_3, \omega_2)}{\lambda}} + z_3 e^{\frac{u(a_3, \omega_3)}{\lambda}} \geq 0, \\
z_1 + z_2 + z_3 < 0.
\end{cases} \tag{1.62}$$

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 $\delta, \varepsilon \geq 0$ solves system (1.62): the two latter inequalities hold trivially for all such z , and the two former inequalities hold if $\varepsilon \geq 3^{\frac{1+\delta}{2}} - 2$. Therefore, there exists no μ that solves system (1.37) given $\beta \in \Delta(\Theta)$.

1.A.12 Proof of Proposition 1.6.3

We first show that there exists an equilibrium in the communication game that replicates the deletion 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 $C(\beta^*) = \{1, \dots, K^*\}$. By Lemma 1, we have that

$$\begin{aligned} \sqrt{\mu(\omega_{K^*})} &\geq \frac{1}{K^* + \delta} \sum_{i=1}^{K^*} \sqrt{\mu(\omega_i)} \\ \iff \delta \sqrt{\mu(\omega_{K^*})} &\geq \sum_{i=1}^{K^*-1} \left(\sqrt{\mu(\omega_i)} - \sqrt{\mu(\omega_{K^*})} \right) \end{aligned} \quad (1.63)$$

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

$$\mu_p(\omega_{K^*} | \tilde{a}_{K^*}) = \max_i \mu_p(\omega_i | \tilde{a}_{K^*}), \quad (1.64)$$

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:

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

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 (1.64) to hold, it is then enough for

$$e^{\frac{1}{\lambda}} \sqrt{\mu(\omega_{K^*})} \geq \sqrt{\mu(\omega_1)} \iff \delta \sqrt{\mu(\omega_{K^*})} \geq \sqrt{\mu(\omega_1)} - \sqrt{\mu(\omega_{K^*})}, \quad (1.65)$$

to be satisfied. Note, however, that it is strictly weaker than (1.63), 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 (1.65) 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 (1.64), since $\mu(\omega_i) \geq \mu(\omega_{K^*})$ for $i \in C(\beta^*)$. This concludes the proof.

2. Setting Interim Deadlines to Persuade

2.1 Introduction

The development of any innovation requires investment of both time and capital, while the outcome of this investment is inherently stochastic. Usually, the investor, being the principal, retains the option to stop funding the innovative project if at some point it proves unsuccessful. It is widely documented that the agent running the project tends to prefer the principal to postpone the stopping of the funding to enjoy either the extra funds or an additional chance to turn her research idea into a success story.¹ In such an agent-principal relationship, the agent's technological expertise and the quality of her innovative idea often allow her to manipulate the principal by designing how and when the outcomes of the research and development process are announced.

Recently, venture capital firms have started to pour billions into startups focused on the development of quantum computers, which are known for their technological complexity and difficulty of construction. The economic viability of quantum computing is questioned by a number of experts; however, the startups promise the investors a completed product in the foreseeable future.² For instance, a quantum startup PsiQuantum announced to potential investors that it hopes to develop a commercially-viable quantum computer within five years and managed to raise more than \$200 million in 2019.³

This paper studies the implications of the agent's control of information during the progress of a research and development project when the agent and the principal disagree about the timing of when to abandon the research idea. I ask: What is the degree of transparency to which an agent should commit before starting to work on an innovative project? In particular, which terms for self-reporting on the progress of the project should a startup propose while discussing the term sheet with a venture capitalist? As I show, depending on the properties of the project, the startup would strategically choose both the timing for the disclosure of updates on the progress of the project and the type of news it discloses - either good or bad.

I study a game between a startup and an investor. The startup controls the information on the progress of the project and has the power to propose the terms for self-reporting on it to the venture capitalist.⁴ The startup has an intertemporal commitment power and commits to a dynamic information policy, which can be interpreted as designing the terms of the contract specifying how the information on the progress of the project is disclosed over time as the project unfolds. In return, the investor continuously provides funds and chooses when to

¹Agency conflict in which the agent prefers the principal to postpone abandoning the project that the agent is working on is studied in Admati and Pfleiderer [1994], Gompers [1995], Bergemann and Hege [1998, 2005], Cornelli and Yosha [2003].

²"The Quantum Computing Bubble." *Financial Times*, August 25, 2022.

³"Bristol Professor's Secretive Quantum Computing Start-Up Raises £179m." *The Telegraph*, November 16, 2019.

⁴I discuss the reasoning behind this assumption in Section 2.3.2.

stop funding the project.

The project has two stages and evolves stochastically over time toward completion, conditional on continuous investment in it. The completion of each of the stages of project occurs according to a Poisson process. The completion of the first stage serves as a milestone, such as the development of a prototype, while completion of the second stage achieves the project. The investor gets a lump-sum project completion profit if and only if he stops investing after the project is completed and before an exogenous project completion deadline, and the startup prefers the principal to postpone stopping the funding.⁵

As the investor receives the reward only after a prolonged period of investment, he initially invests without being able to see if the investment is worthwhile. Hence, it is individually rational for the investor to start investing only if he is sufficiently optimistic regarding the future of the project. An important feature of the setting that I consider is that the *information is symmetric at the outset*: not only the investor, but also the startup is unable to find out if the project will bring profit to the investor or not - this can be inferred only as time goes on and some evidence is accumulated. The only tool that the startup has for persuading the investor to start investing is the promise of future reports on the progress of the project.

Clearly, the good news about the completion of the project is valuable to the investor as it helps him to stop investing in a timely manner. Further, as evidence regarding the project accumulates over time, failure to pass the milestone in a reasonable time makes the project unlikely to be accomplished in time - and the investor prefers to stop investing after observing such bad news. When designing the information policy, the startup chooses optimally between the provision of these two types of evidence in order to postpone the investor's stopping decision for as long as possible.

I show that the startup's choice of information policy depends on the ex ante attractiveness of the project for the investor. The attractiveness is captured by the *flow cost-benefit ratio of the project*. Thus, a project is relatively more attractive ex ante to the investor when its flow investment cost is lower, its project completion profit is higher, or the Poisson rate, at which completion of one stage of the project occurs, is higher.

When the project is sufficiently attractive ex ante to the investor, promises to provide information only on the completion of the project serve as a sufficiently strong incentive device to motivate the investor to start the funding at the outset. Further, the future news on the completion of the project does not harm the total expected surplus generated by the interaction of the startup and investor, while the future news on the project being poor decreases the surplus that the startup can potentially extract from the investor. Accordingly, the startup commits to providing only the good news, but not the bad news on the project in the future: it *discloses the completion of the project and postpones the disclosure* in order to ensure the extraction of as much surplus as possible from the investor. In the context of quantum computing, the startup optimally chooses and announces to the venture capitalist the date by which it plans to have a fully developed quantum computer. When the date comes, the startup reports completion if the quantum computer has been completed; if not, the startup reports the completion

⁵I discuss the reasons for the presence of the project completion deadline in Section 2.3.1.

as soon as it occurs.

The situation changes when the project does not look promising to the investor *ex ante*. In that case, if the startup commits to disclosing only the completion of the project, the investor will not have the sufficient motivation to start investing in it. Thus, the startup extends the information policy to encompass not only the good news but also the bad. As in the case of the promising project, the startup discloses the project’s completion and does so without any postponement, thereby fully exploiting its preferred incentive tool. In addition, the startup sets a date at which the bad news is released if the milestone of the project has not yet been reached - this date is *the interim reporting deadline*.

Setting the interim deadline, the startup chooses a deterministic date, which it optimally postpones. As the startup prefers the investor to postpone stopping the funding, it prefers the interim deadline to be at a later expected date. Further, the completion of the stages of project according to a Poisson process makes both the startup and the investor risk-averse with respect to the date of the interim deadline. Thus, the startup prefers to set the interim deadline at a deterministic date and to postpone it as late in time as possible in order to extract all the surplus from the investor. In the context of quantum computing, the startup optimally chooses and announces a fixed date by which it plans to have a prototype of the quantum computer. When the date comes, reporting having the prototype at hand convinces the investor to continue the funding, and reporting not having the prototype leads to termination of the project.

Finally, I demonstrate that the outlined structure of the optimal information disclosure holds for a broad class of preferences of the startup and the investor. I allow for profit-sharing between the startup and the investor, varying degrees of the startup’s benefit from the flow of funding, and exponential discounting, and show that the startup prefers not to set any interim deadlines whenever the project is sufficiently promising to the investor. The future disclosure of the completion of the project promises investor profit in exchange for a prolonged investment, while the disclosure of the stagnation of the project at the interim deadline promises investor only saved costs, as further investment stops. Thus, when the project is attractive, the startup can make the funding and the beneficial experimentation relatively longer by setting no interim deadlines.

2.2 Related literature

My paper is related to the *literature on dynamic information design*. The closest paper in this strand of literature is by Ely and Szydlowski [2020]. Similarly to my paper, they study the optimal persuasion of a receiver facing a lump-sum payoff to incur costly effort for a longer time. In my model, as in theirs, the sender is concerned to satisfy the beginning-of-the-game individual rationality constraint of the receiver when choosing the information policy. Further, the “leading on” information policy in Ely and Szydlowski [2020] has a similar flavor to the “postponed disclosure of completion” information policy in my paper: promises of news on completion of the project serve as an incentive device sufficient to satisfy the receiver’s individual rationality constraint.

However, there are several substantial differences between Ely and Szydlowski [2020] and my paper. While in their model the state of the world is static and

drawn at the beginning of the game, in my model it evolves endogenously over time, given the receiver’s investment. As a result, the initial disclosure used in the “moving goalposts” policy in Ely and Szydlowski [2020] cannot provide additional incentives for the receiver in my model. The sender in my model uses another incentive device to incentivize the receiver to opt in at the initial period: she commits to an interim deadline at which she discloses that the first stage of the project is not completed.

Another closely related paper is by Orlov et al. [2020]. The main similarity to my paper lies in the sender’s incentive to postpone the receiver’s irreversible stopping decision. The sender in their paper prefers to backload the information provision, which is in line with the properties of the optimal information policy in my paper. However, there are a number of substantial differences between our papers. In Orlov et al. [2020], the sender does not have the intertemporal commitment power; further, the receiver potentially obtains a non-negative payoff at each moment of time, and thus the sender does not need to persuade the receiver to opt in at the beginning of the game.

Ely [2017], Renault et al. [2017], Ball [2019] also analyze dynamic information design models. However, their papers focus on persuading a receiver who chooses an action and receives a payoff at each moment of time, whereas in my paper the receiver takes an irreversible action and receives a lump-sum project completion payoff. Henry and Ottaviani [2019] consider a dynamic Bayesian persuasion model in which, similarly to my model, the receiver needs to take an irreversible decision. However, the incentives of the sender and receiver differ from my model: the receiver wants to match the static state of the world and the sender is concerned with both the receiver’s action choice and with the timing of that choice. Basak and Zhou [2020] study dynamic information design in a regime change game. The optimal information policy in their model resembles the interim deadline policy in my model: at a fixed date, the principal sends the recommendation to attack if the regime is substantially weak by that time.

My paper is also related to the *literature on the dynamic provision of incentives for experimentation* [Bergemann and Hege, 1998, 2005, Curello and Sinander, 2020, Madsen, 2022]. The closest papers in this strand of literature are by Green and Taylor [2016] and Wolf [2017]. Similarly to my model, both papers consider design of a contract regarding a two-stage project, in which the completion of stages arrives at a Poisson rate. In Green and Taylor [2016], there is no project completion deadline and the quality of the project is known to be good, while in Wolf [2017] the quality of the project is uncertain. In contrast to my paper, both papers focus on a canonical moral-hazard problem and give the power to design the terms of the contract to the investor (principal) rather than the startup (agent). In particular, the contract in Green and Taylor [2016] specifies the terms for the agent’s reporting on the completion of the first stage of the project. Similarly to my model, the optimal reporting takes the form of a deterministic interim deadline: at a principal-chosen date, the agent truthfully reports if she has already completed the first stage, which determines the further funding of the project.⁶

⁶In a broad sense, my paper also relates to the small strand of theoretical literature on dynamic startup-investor and startup-worker relations under information asymmetry [Kaya, 2020, Ekmekci et al., 2020]. However, while these papers focus on the signaling of the type of

2.3 The model

2.3.1 Setup

I consider a game between an agent (she, sender) and a principal (he, receiver). Time is continuous and there is a publicly observable deadline T , $t \in [0, T]$.⁷ For each t , *the principal* chooses sequentially to invest in the project ($a_t = 1$) or not ($a_t = 0$). The flow cost of the investment is constant and given by c . The choice of $a_t = 0$ at some t is irreversible and induces the end of the game.⁸

The assumption that the project needs to be completed in finite time is natural in many economic settings. The main interpretation for T is an expected change in market conditions that renders the project unprofitable. In the context of a research and development project, T could stand for the date at which the competitor's innovative product is expected to enter the market, or the date at which the competitor is expected to get a patent on the competing innovation.

The *state of the world* at time t is captured by the number of stages of the project completed by t , x_t , and the project has two stages, $x_t \in \{0, 1, 2\}$. The state process begins at the state $x_0 = 0$ and, conditional on the continuation of the investment by the principal, it increases at a Poisson rate $\lambda > 0$. Information on the number of stages completed is controlled by *the agent*. Thus, when making investment decisions, the principal relies on the information provided by the agent.

The project brings the profit v if and only if the *second stage of the project* has been completed by the time of stopping, and a payoff of 0, otherwise. I assume that all of the profit goes to the principal. This assumption simplifies the exposition without affecting the main results of the paper. I relax this assumption and consider the profit-sharing between the agent and the principal in Section 2.6.

There is a *conflict of interest* between the agent and the principal as the agent benefits from using the funds provided by the principal for running the project, possibly diverting them for her benefit. Thus, the agent faces the flow payoff of c and wants the principal to postpone his irreversible decision to stop as long as possible.

I study the agent's choice of information provision to the principal. The agent chooses an information policy to maximize her expected long-run payoff. I assume that the agent has the power to announce and commit to a policy. An *information policy* σ is a rule that for each date t and for each past history $h(t)$ specifies a probability distribution on the set of messages M . The history includes all past and current realizations of the process and all past message draws and principal's action choices.

When choosing an information policy, the agent faces a rich strategy space. First, she can choose if the information on the completion of the first, or second, stage of the project will be disclosed by the policy. Second, she can choose

startup, I study the provision of information by the startup on the progress of the project.

⁷The results for the setting without a deadline are easily obtained by considering $T \rightarrow \infty$. They are presented in Appendix 2.E.

⁸The absence of the principal's commitment to an investment policy and the irreversibility of the stopping decision capture the venture capitalist's option to abandon the project, e.g., in the case of its negative net present value.

how to disclose the completion of a stage of the project: for instance, to do so immediately or to postpone the disclosure.

The timing of the game is as follows. First, at $t = 0$, the agent publicly commits to an information policy σ . Second, at each t the principal observes the message generated by the information policy and makes her investment decision. The game ends when the principal chooses to stop investing or at T , if he keeps investing until T . I assume that whenever indifferent about investing or not, the principal chooses to invest, and whenever indifferent about disclosing information or not, the agent chooses not to disclose.

Throughout the paper, I use the following intuitive notational convention: for any two dates at which the principal stops investing, S and τ ,

$$\begin{aligned} S \wedge \tau & \min \{S, \tau\}, \\ S \vee \tau & \max \{S, \tau\}. \end{aligned}$$

2.3.2 Discussion of assumptions

The main interpretation of the considered dynamic information design problem is the contracting between the agent (startup) and the principal (investor) on the terms of reporting on the completion of stages of the project that are not publicly observed. The terms could take the form of a proposed formal reporting schedule or a schedule of meetings with the investor. Non-observability of the stage completions stems from the fact that, while the technology is being developed in the lab, the principal either does not have sufficient expertise to assess the progress or the full access to the lab.

I assume that the principal does not have the power to propose the terms for reporting to the agent and, e.g., make her fully disclose the progress achieved in the lab. The most natural interpretation of such an asymmetry in the bargaining power is the asymmetry in the market for private equity: there are sufficiently many investors willing to invest in a particular technology or sufficiently few startups working on the technology.⁹ For instance, investors' interest in quantum computing has grown markedly in recent years, while there are reports of a shortage of human capital in this industry.^{10,11} Another example is the communication software industry, which has recently experienced increased investment activity.¹²

As the agent enjoys the power of full control over the information on the progress of the project, she is completely free to offer what is disclosed and when. In particular, the contract between the agent and the principal can specify that the completion of the second stage of the project is disclosed with a delay rather than immediately. The agent who has an advantage in expertise over the principal

⁹In the alternative interpretation of the model, contracting concerns internal corporate research and development and takes place between the leading researcher and the headquarters of a company. The leading researcher's bargaining power in proposing the terms for disclosure again stems from the market asymmetry: the specialists having the desired level of expertise might be in a short supply.

¹⁰"The Quantum Computing Bubble." *Financial Times*, August 25, 2022.

¹¹"Quantum Computing Funding Remains Strong, but Talent Gap Raises Concern", a report by McKinsey Digital, <https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/quantum-computing-funding-remains-strong-but-talent-gap-raises-concern/>.

¹²"This Is Insanity: Start-Ups End Year in a Deal Frenzy." *Best Daily Times*, December 07, 2020.

can rationalize such a condition by saying that before the success is reported to the principal, it is worth re-checking the data, which takes time.

Even though the principal can not dictate to the agent which information and how she should disclose, the principal can potentially hire an external monitor who would visit the lab and prepare an additional report on the progress of the project. In that case, the contract signed between the agent and the principal will account for both free information that the agent promised to provide and additional costly information which the principal obtains with the help of a monitor. In the baseline version of the model, I assume that the principal can not use the help of a monitor. This can be rationalized by the shortage of experts in the field, which makes hiring a monitor prohibitively costly. Alternative interpretation is that the agent restricts the principal's access to additional information on the progress of the project by stating that a potential information leak would put the technology being developed at risk.¹³

The information policy relies upon the agent's commitment power, which holds not only within each date but also between the dates. The agent's commitment within each date follows from prohibitively high legal costs of cooking up the evidence. The agent's intertemporal commitment stems from the rigidity of terms and form of reporting fixed in the contract that the agent and the principal sign at the outset of the game.

2.4 No-information and full-information benchmarks

2.4.1 No-information benchmark

First, I consider the simple case when the information policy is given by σ^{NI} : the same message m is sent for all histories $h(t)$ and all dates t . Thus, the agent provides no information regarding the progress of the project. As I demonstrate, in this case the principal starts investing in the project if and only if it is sufficiently promising for the principal from the ex ante perspective and invests until a deterministic interior date.

As no news arrives, the principal bases his decision about when to stop investing on his unconditional belief regarding the completion of the second stage of the project. I denote the unconditional belief that n stages of the project were completed by t , by $p_n(t)$, i.e., $p_n(t)P(x_t = n)$. The state of the world is fully determined by $p(t)$ given by

$$\begin{aligned} p_0(t) &= e^{-\lambda t}, \\ p_1(t) &= \lambda t e^{-\lambda t}, \\ p_2(t) &= 1 - e^{-\lambda t} - \lambda t e^{-\lambda t}. \end{aligned}$$

The principal's sequential choice of $a_t \in \{0, 1\}$ can be restated equivalently as

¹³In particular, this rationale was used to restrict the investors' access to information on the progress of the project in the case of Theranos, see "What Red Flags? Elizabeth Holmes Trial Exposes Investors' Carelessness." *The New York Times*, November 04, 2021.

the choice of deterministic stopping time $S^{NI} \in [0, T]$ chosen at $t = 0$.¹⁴ Given the principal's continuous investment, the probability of completion of the second stage of the project, $p_2(t)$, increases monotonously over time, making obtaining the payoff v more likely. However, postponing the stopping is costly.

To decide on S^{NI} , the principal trades off the flow benefits and flow costs of postponing the stopping decision, while keeping the individual rationality constraint in mind. The flow cost of postponing the stopping for Δt is given by $c \cdot \Delta t$ and the flow benefit is given by $v \cdot p_1(t) \lambda \Delta t$.¹⁵ Thus, the necessary condition for the principal's stopping at some interior moment of time ($0 < S < T$) is given by

$$v \cdot p_1(S) \lambda = c. \quad (2.1)$$

Let

$$\kappa \frac{c}{v\lambda},$$

the ratio of the flow cost of investment c to the gross project payoff v normalized using λ , the rate at which a project stage is completed in expectation. The intuitive interpretation of κ is the *flow cost-benefit ratio of the project*. κ is an inverse measure of how ex ante promising the project is for the principal. (2.1) is equivalently given by¹⁶

$$\underbrace{p_1(S)}_{\text{flow benefit of waiting}} = \underbrace{\kappa}_{\text{flow cost of waiting}} \quad (2.2)$$

and presented graphically in Figure 2.1. As the state process transitions monotonously from 0 to 1 and then to 2, $p_1(t)$ first increases and after some time starts to decrease. Thus, the principal has two candidate interior stopping times satisfying (2.2), \bar{S} and \bar{S}^{NI} . The principal prefers to postpone stopping to \bar{S}^{NI} , as during (\bar{S}, \bar{S}^{NI}) the flow benefits are higher than the flow costs.

The forward-looking principal can guarantee himself a payoff of 0 if he does not start investing at $t = 0$. Thus, he will choose to start investing at $t = 0$ only if his flow gains accumulated up to $T \wedge \bar{S}^{NI}$ are larger than his flow losses, and his expected payoff is given by

$$V^{NI} \max \left\{ 0, \int_0^{T \wedge \bar{S}^{NI}} (v \cdot p_1(s) \lambda - c) ds \right\}. \quad (2.3)$$

Geometrically, the integral in (2.3) represents the difference between the shaded areas in Figure 2.2 that correspond to the accumulated gains and losses. The principal starts investing at $t = 0$ if, given T and λ , the normalized cost-benefit ratio κ is low enough, so that the shaded area of the accumulated gains is at least as large as that of the accumulated losses. I denote such a cutoff value of κ by $\kappa^{NI}(T, \lambda)$ and summarize the principal's choice without information in Lemma 7.

¹⁴Note that the dynamic belief system that he faces is deterministic in a sense of being fully specified from $t = 0$ perspective.

¹⁵To observe this, note that the probability of the completing both the first and second stages within a very short time Δt is negligibly small; thus, during some Δt , the principal receives the project completion payoff v iff the first stage has already been completed.

¹⁶Here I WLOG express the flow benefits and flow costs of investing for the principal in different units of measurement.




Figure 2.1: Principal's choice under no information:

left plot: postponing stopping increases the chance of getting a project payoff v ;
right plot: principal trades off costs and benefits and optimally stops at \bar{S}^{NI} .

Lemma 7. *Assume no information regarding the progress of the project arrives over time. Denote the time at which the principal stops investing by S^{NI} . If $\kappa > \kappa^{NI}(T, \lambda)$, then the principal does not start investing in the project, i.e., $S^{NI} = 0$. If $\kappa \leq \kappa^{NI}(T, \lambda)$, then the principal's choice of stopping time is given by*

$$S^{NI} = \begin{cases} \bar{S}^{NI}, & \text{if } \frac{1}{\lambda} \leq T \text{ and } \kappa \geq e^{-\lambda T} \lambda T \\ T, & \text{otherwise,} \end{cases} \quad (2.4)$$

the closed-form expressions for \bar{S}^{NI} and $\kappa^{NI}(T, \lambda)$ are presented in the proof in Appendix 3.A.

2.4.2 Full-information benchmark

Here, I consider the case in which the information policy is given by σ^{FI} : $M = \{m_0, m_1, m_2\}$ and the message m_n is sent for all t such that $x_t = n$, $n \in \{0, 1, 2\}$. Thus, the principal fully observes the progress of the project at each t .¹⁷ I characterize the cutoff level of the cost-benefit ratio below which the principal is willing to start investing. Further, I show that the principal chooses to stop when no stages of the project are completed and the project completion deadline T is sufficiently close.

At each t , the principal uses the information on the number of stages completed to decide either to stop investing or to postpone the stopping. The news on completion of the second stage of the project makes the principal stop immediately, as this way he immediately receives the project payoff v and stops incurring the costs of further investment. If only the first stage of the project is completed, the principal faces the following trade-off. The instantaneous probability that the second stage will be completed during Δt is given by $\lambda \Delta t$, which is constant over the time. Thus, the expected benefit of postponing the stopping for Δt is given

¹⁷This benchmark corresponds to equilibrium in the setting, where the principal has the full power to propose the terms of self-reporting to the agent.




Figure 2.2: Principal's choice to start investing at $t = 0$ or not:
left plot: $T > \bar{S}^{NI}$; the project deadline is distant and decision-irrelevant;
right plot: $T \leq \bar{S}^{NI}$; the project deadline is close, which leads to lower expected benefits of investing.

In both plots the expected accumulated gains are higher than the losses, so the principal starts to invest at $t = 0$.

by $v \cdot \lambda \Delta t$. Meanwhile, the expected cost of postponing the stopping is given by $c \cdot \Delta t$. As a result, if $\kappa \leq 1$, then the principal who knows that the first stage of the project has already been completed invests until either the completion of the second stage or until the project deadline T is reached.

Consider now the case in which the principal knows that the first stage has not yet been completed. The principal's trade-off with respect to the stopping decision is now more involved. Postponing the stopping for Δt leads to the completion of the first stage of the project with the instantaneous probability $\lambda \Delta t$. Completion of the first stage of the project at some t implies that the principal receives the continuation value of the game, conditional on having the first stage completed. I denote the continuation value of the principal at time t under full information and conditional on the completion of first stage of the project by $V_{t|1}^{FI}$. This is given by¹⁸

$$V_{t|1}^{FI} = \left(v - \frac{c}{\lambda} \right) \left(1 - e^{-\lambda(T-t)} \right). \quad (2.5)$$

The principal's expected benefit from postponing the stopping for Δt is given by $V_{t|1}^{FI} \cdot \lambda \Delta t$ and the cost of postponing the stopping is, as before, given by $c \cdot \Delta t$. The continuation value, $V_{t|1}^{FI}$, shrinks over time and approaches 0 as the project deadline T approaches. This is because the shorter the time left before the project deadline, the less likely it is that the second stage of the project will be completed before T . If at some t , and given that no stages are completed yet, the expected net benefit of investing reaches 0, it is optimal for the principal to stop at that t .¹⁹ I denote this date by S_0^F and plot it in Figure 2.3.

¹⁸See the derivation in the proof of Lemma 8 in the Appendix.

¹⁹If at t the expected benefit of investing becomes lower than the cost, then, after t , the net expected benefit remains negative. Thus, it is optimal for the principal to stop investing precisely at t .




Figure 2.3: The principal optimally sets an interim deadline $t = S_0^F$ under full information: given that the first stage of the project has not been completed by S_0^F , it is optimal to stop investing at S_0^F .

As the principal has an incentive to stop at S_0^F only if he knows that the first stage or the milestone of the project has not been reached, the economic interpretation of S_0^F is that it is *the interim deadline* that the principal sets for the project. If the milestone has not been reached by the interim deadline, then it is sufficiently unlikely that the project will be completed before the project deadline T . Thus, it is optimal for the principal to “give up” on the project and stop investing at $t = S_0^F$. If the milestone is reached by the interim deadline, then the principal has an incentive not to stop investing until either the second stage is completed or T is hit.

Finally, given the plan to stop either at the interim deadline, or at the completion of the second stage of the project, it is individually rational to start investing only if the principal’s expected payoff from the $t = 0$ perspective is non-negative. I denote the upper bound for the normalized cost-benefit ratio such that the principal starts investing at $t = 0$ by $\kappa^{FI}(T, \lambda)$. Intuitively, $\kappa^{FI}(T, \lambda) > \kappa^{NI}(T, \lambda)$: whenever the principal is willing to start investing under no information, he is also willing to start under the full information. I summarize the principal’s choice under full information in Lemma 8.

Lemma 8. *Assume that the progress of the project is fully observable at each moment in time. If $\kappa > \kappa^{FI}(T, \lambda)$, where $\kappa^{FI}(T, \lambda) > \kappa^{NI}(T, \lambda)$, then the principal does not start investing in the project. If $\kappa \leq \kappa^{FI}(T, \lambda)$, the principal invests either until the random date at which the second stage of the project is completed, $t = \tau_2$, or until the interim deadline, $t = S_0^F$, at which he stops if the first stage has not yet been completed. Formally, the time at which the principal stops investing is a random variable τ given by:*

$$\tau = \begin{cases} \tau_2 \wedge T, & \text{if } x_{S_0^F} \neq 0 \\ S_0^F, & \text{otherwise,} \end{cases}$$

where $S_0^F = T + \frac{1}{\lambda} \log\left(\frac{1-2\kappa}{1-\kappa}\right)$ and the expression for $\kappa^{FI}(T, \lambda)$ is presented in the proof in Appendix 3.A.

Assume now that the agent chooses which information to provide to the principal. As for $\kappa > \kappa^{FI}(T, \lambda)$ the principal is not willing to start investing even

under full information, there is no way in which the agent can strategically conceal the information to her benefit. In Section 2.5, I assume $\kappa \leq \kappa^{FI}(T, \lambda)$ and analyze how the agent can strategically provide information on the progress of the project and extract the principal's surplus.

2.5 Agent's choice of information policy

In this Section, I present how the agent's choice of information policy changes with the ex ante attractiveness of the project, which is captured by the cost-benefit ratio κ . In Section 2.5.1, I start with Proposition 2.5.1 which summarizes the comparative statics result. In Sections 2.5.2-2.5.3, I proceed with the detailed discussion of the economic mechanisms that determine the outlined structure of the optimal information policy. Throughout Section 2.5, I maintain the following technical assumption:

Assumption 2.1. $e^{\lambda T} > \lambda T(\lambda T + 1) + 1$.

For the sake of a clearer exposition, this assumption rules out the case in which T is so low that whenever the principal is willing to start investing in the no-information benchmark, he invests until T . Relaxing this assumption does not change the comparative statics result in Proposition 2.5.1 qualitatively.²⁰

2.5.1 The structure of optimal information disclosure

There exist cost-benefit ratio cutoffs $\kappa^{ND}(T, \lambda)$, $\kappa^{ND}(T, \lambda) < \kappa^{NI}(T, \lambda)$, and $\tilde{\kappa}(T, \lambda)$, $\kappa^{NI}(T, \lambda) < \tilde{\kappa}(T, \lambda) < \kappa^{FI}(T, \lambda)$, such that, depending on the cost-benefit ratio of the project, the optimal information policy has the following form:

1. when $\kappa \leq \kappa^{ND}(T, \lambda)$, the agent provides no information and the principal invests until T ;
2. when $\kappa^{ND}(T, \lambda) < \kappa \leq \tilde{\kappa}(T, \lambda)$, the agent discloses only the completion of the second stage of the project and does that with the postponement;
3. when $\tilde{\kappa}(T, \lambda) < \kappa < \kappa^{FI}(T, \lambda)$, the agent immediately discloses the completion of the second stage of the project whenever it occurs and specifies a deterministic interim deadline, at which it discloses if the first stage is already completed;
4. when $\kappa \geq \kappa^{FI}(T, \lambda)$, the agent provides no information as the principal's long-run payoff is non-positive even under full information.

Figure 2.4 illustrates Proposition 2.5.1 and presents the partition of the cost-benefit ratio space based on the corresponding forms of the optimal information policy.

The structure of optimal disclosure presented in Proposition 2.5.1 follows the simple and intuitive pattern. The lower is the value of cost-benefit ratio, the higher is ex ante attractiveness of the project to the principal. First, for

²⁰I discuss the implications of relaxing this assumption in the proof of Proposition 2.5.1.




Figure 2.4: Comparative statics of the form of optimal information policy with respect to the cost-benefit ratio of the project, $\kappa(T, \lambda)$.

$\kappa \leq \kappa^{ND}(T, \lambda)$, the project is so attractive that the principal is willing to keep investing until the project deadline T even in the no-information benchmark. Thus, there is no need to disclose any information. For the higher values of κ , there emerges a room for strategic disclosure, and the higher is the value of κ (i.e., the lower is the ex ante attractiveness of the project), the more information the agent has to disclose to incentivize the principal. For $\kappa \geq \kappa^{FI}(T, \lambda)$, the project gets so unattractive that the principal can not strictly benefit from investing even in the full-information benchmark. In this extreme case, the agent chooses not to disclose any information.

The most important part of the result in Proposition 2.5.1 demonstrates which additional pieces of information the agent chooses to disclose and when she chooses to disclose them as κ gets higher and higher. When $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$, the agent discloses only the completion of the second stage of the project and does not promise any information on the completion of the first stage of the project. Further, as κ increases from $\kappa^{ND}(T, \lambda)$ to $\tilde{\kappa}(T, \lambda)$, the agent adjusts the timing of the disclosure: she postpones the disclosure of the second stage completion less and less and discloses immediately for $\tilde{\kappa}(T, \lambda)$. For $\kappa \in (\tilde{\kappa}(T, \lambda), \kappa^{FI}(T, \lambda))$, the agent not only discloses the completion of the second stage of the project immediately, but also provides information on the completion of the first stage at the interim deadline that she optimally chooses.

In the subsequent Sections, I provide details on the mechanisms that shape the aforementioned comparative statics results. I omit the trivial case of non-disclosure under $\kappa \leq \kappa^{ND}(T, \lambda)$ and start the discussion from the optimal information policy under $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$.

2.5.2 Postponed disclosure of project completion

In this Section, I restrict attention to $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$ and explain why the optimal information policy has the particular form presented in the Proposition 2.5.1: the agent discloses only the completion of the project and does this with the postponement.

Agent's problem

To characterize the agent's choice of information policy, I consider an equivalent problem, in which the agent directly chooses the stochastic history-contingent length of investment subject to the principal's individual rationality constraints that ensure optimality of such action process for the principal. An *investment schedule* is a random variable $\tau : \Omega \rightarrow [0, T]$ defined on the probability space

$(\Omega, \mathcal{F}, \mathbb{P})$ and adapted to the filtration $F = (\mathcal{F}_t)_{t \geq 0}$ generated by the stochastic process x_t . As I demonstrate in Section 2.5.2, restricting attention to random variables adapted to the *natural* filtration of x_t is without loss of generality for the agent's equilibrium expected payoff when $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$.²¹

Informally, an investment schedule τ is a random variable with support $[0, T]$ specified by a rule that suggests when to stop investing depending on the history of previous realizations of the number of completed stages x_t .²² The agent chooses this rule at $t = 0$. $P(x_\tau = 2)$ captures the belief about two stages of the project completed by τ , the random time of stopping in the future, and $E[\tau]$ captures the expected length of investment from $t = 0$ perspective.

Given an investment schedule τ , the long-run payoff of the agent and the principal are given, respectively, by

$$\begin{aligned} W(\tau) & E[\tau] c, \\ V(\tau) & P(x_\tau = 2) v - E[\tau] c. \end{aligned}$$

As an investment schedule τ is an action recommendation rule, the action recommendations generated by this rule have to be obedient for the principal. In other words, at each date and for each possible history the principal's actions suggested by τ have to be optimal for the principal. An object useful for characterizing if an investment schedule τ generates obedient action recommendations is given by the principal's continuation value at some interim date t promised by the investment schedule τ . This continuation value depends on the beliefs of the principal.

As the principal does not commit to a policy at $t = 0$, he rationally updates his belief given an investment schedule τ and assesses the costs and benefits of either further following the investment schedule τ provided by the agent or deviating from it. The absence of stopping by some date t and, thus, the fact that the game continues at t serves as a source of inference for the principal. First, he forms a belief regarding the number of completed stages of the project by t , conditional on the game still continuing at t , $P(x_t = n | t < \tau)$. Second, he forms a belief regarding the number of completed stages of the project at the random date of stopping in the future, τ , $P(x_\tau = n | t < \tau)$.

Given the absence of stopping by t , the principal's expected payoff promised by the schedule is given by $P(x_\tau = 2 | t < \tau) v - E[\tau - t | t < \tau] c$. The principal's expected payoff from stopping at t is given by $P(x_t = 2 | t < \tau) v$. The principal's *continuation value at t* given the investment schedule τ is the difference between these two expected payoffs, I denote it by $V_t(\tau)$:

$$V_t(\tau) [P(x_\tau = 2 | t < \tau) - P(x_t = 2 | t < \tau)] v - E[\tau - t | t < \tau] c. \quad (2.6)$$

This way of formulating the continuation value is intuitive: if the continuation value $V_t(\tau)$ gets negative then it is not valuable to continue investing for the

²¹In other words, there is no need for external randomization devices to optimally incentivize the principal when $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$.

²²The stopping rules from the no-information and full-information benchmarks are given in Lemmas 7 and 8, respectively. Further examples of such rules include "stop 1 minute after the second stage of the project is completed" and "stop at $t = S$ if only the first stage of the project is completed by $t = S$ ".

principal, and he is better-off stopping immediately rather than following the schedule. The following Lemma shows the necessary and sufficient conditions for an investment schedule τ to generate obedient action recommendations for the principal.

Lemma 9. *An investment schedule τ is the principal's best response to at least one information policy σ if and only if*

$$V_t(\tau) \geq 0, \forall t \geq 0 \text{ and } V_\tau^{NI} < 0, \quad (2.7)$$

where V_t^{NI} is the principal's optimal continuation value in the absence of any additional information from the agent starting from the date t .

$V_t(\tau) \geq 0, \forall t \geq 0$ ensures that the principal does not want to stop before the date of stopping suggested by the investment schedule is reached, and $V_\tau^{NI} < 0$ ensures that the principal does not want to continue conditional on reaching the date of stopping suggested by the investment schedule. Conditions from Lemma 9 constitute the system of constraints for the agent's problem.

As the agent chooses an investment schedule τ to maximize her long-run payoff, the constraint $V_\tau^{NI} < 0$ is inactive at optimum.²³ Thus, without loss of generality, I omit this constraint from the agent's problem, and the problem that the agent solves at $t = 0$ is given by

$$\begin{aligned} & \max_{\tau \in \mathcal{T}} \{c \cdot E[\tau]\} \\ & \text{s.t. } V_t(\tau) \geq 0, \forall t \geq 0, \end{aligned} \quad (2.8)$$

where \mathcal{T} is the set of stopping times with respect to the natural filtration of x_t . As the principal's choice to postpone the stopping of funding is costly, it is natural to interpret the system of constraints in (2.8) as the *system of principal's individual rationality constraints*.

The agent's problem is complex, and thus I solve it in steps. First, I characterize the investment schedule, which solves the *relaxed* version of (2.8) with the principal's individual rationality constraints only for some initial periods. Second, I demonstrate that there exists an investment schedule solving the *relaxed* agent's problem and satisfying the *full* system of the principal's individual rationality constraints (2.7). This investment schedule pins down optimal information policy.

Solution to the agent's relaxed problem

In this Section, I consider the agent's relaxed problem and discuss its solution. This sheds light on the technical intuition behind the key properties of the optimal information policy. The agent's relaxed problem for the parametric case of $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$ is given by (2.8) with the *principal's individual rationality constraint only for $t \in [0, \bar{S}^{NI}]$* . The agent's relaxed problem for the parametric case of $\kappa \in (\kappa^{NI}(T, \lambda), \tilde{\kappa}(T, \lambda)]$ is given by (2.8) with the *principal's individual rationality constraint only for $t = 0$* .

²³Otherwise, the agent can prolong the expected funding by choosing a different τ .

Consider the agent's long-run payoff given an investment schedule, $W(\tau)$. This can be restated equivalently as follows:

$$\begin{aligned} W(\tau) &= [W(\tau) + V(\tau)] - V(\tau) \\ &= \underbrace{P(x_\tau = 2)v}_{\text{total surplus}} - \underbrace{[P(x_\tau = 2)v - E[\tau]c]}_{\text{principal's surplus}}. \end{aligned} \quad (2.9)$$

The solution to the agent's relaxed problem for both considered parametric cases follows a simple idea: the optimal investment schedule *ensures that the total surplus is maximal and that the principal's surplus is minimal*. Consider a schedule τ such that the stopping occurs after the completion of the second stage of the project, unless the project deadline T was hit, i.e., the schedule satisfies the condition $\tau \geq \tau_2 \wedge T$. Such a schedule leads to

$$P(x_\tau = 2) = P(x_T = 2). \quad (2.10)$$

Given a schedule τ satisfying (2.10), if τ is individually rational for the principal at date $t = 0$ then the total surplus generated achieves its upper bound and is given by $P(x_T = 2)v$, which depends on the exogenously given project deadline T and the profit v . However, the stopping only after the second stage completion is not individually rational for the principal at $t = 0$ when the cost of funding is sufficiently high, the profit is sufficiently low, or the expected time until a project stage completion is sufficiently high.

Lemma 10 elaborates on the cost-benefit ratio cutoff value $\tilde{\kappa}(T, \lambda)$: it distinguishes the case in which stopping only after the second stage completion is individually rational at $t = 0$ from the case in which it is not. Based on this partition, when $\kappa \in (\kappa^{ND}, \tilde{\kappa}(T, \lambda)]$, I call the project *ex ante promising* for the principal.

Lemma 10. *For each (T, λ) there exists $\tilde{\kappa}(T, \lambda)$, $\kappa^{NI}(T, \lambda) < \tilde{\kappa}(T, \lambda) < \kappa^{FI}(T, \lambda)$, such that if $\kappa \leq \tilde{\kappa}(T, \lambda)$ ($\kappa > \tilde{\kappa}(T, \lambda)$) then an investment schedule τ in which stopping after $\tau_2 \wedge T$ happens with probability one is individually rational at $t = 0$ (not individually rational at $t = 0$) for the principal.*

For $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$, the schedule $\tau \geq \tau_2 \wedge T$ is individually rational for the principal at $t = 0$, and it maximizes the total surplus. In addition to choosing $\tau \geq \tau_2 \wedge T$, it is optimal for the agent to choose the investment schedule with a higher expected date of stopping the funding to extract all the principal's surplus subject to his individual rationality constraints. For $\kappa \in (\kappa^{NI}(T, \lambda), \tilde{\kappa}(T, \lambda)]$, the agent chooses such τ that the principal's individual rationality constraint at $t = 0$ is binding. As a result, $V(\tau) = V^{NI}$, i.e., the principal gets his no-information benchmark payoff given by 0.

For $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$, as in the no-information benchmark the principal invests until \bar{S}^{NI} with certainty, the agent chooses the investment schedule as to postpone the start of information provision at least until \bar{S}^{NI} . Further, the agent chooses τ with a higher expected date of stopping so that the principal's individual rationality constraint at $t = \bar{S}^{NI}$ is binding. The absence of stopping until at least \bar{S}^{NI} and the fact that individual rationality constraint binds at $t = \bar{S}^{NI}$ taken together imply that $V(\tau) = V^{NI}$, i.e., from $t = 0$ perspective, the

principal gets her no-information benchmark payoff, which is non-negative and given by (2.3).

The next Lemma summarizes the *necessary* conditions for an investment schedule to solve the agent's relaxed problem when the project is promising. These conditions are shared both by the relaxed problem formulated for the case of $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$ and the relaxed problem formulated for the case of $\kappa \in (\kappa^{NI}(T, \lambda), \tilde{\kappa}(T, \lambda)]$. The conditions that are *both necessary and sufficient* for an investment schedule to solve the agent's relaxed problem are presented in the Proof of Lemma 11.

Lemma 11. *Assume $\kappa \in (\kappa^{ND}, \tilde{\kappa}(T, \lambda)]$. If an investment schedule τ solves agent's relaxed problem, then*

1. *with probability one, stopping occurs after $\tau_2 \wedge T$;*
2. *$V(\tau) = V^{NI}$, where V^{NI} is the principal's expected payoff in the no-information benchmark, given by (2.3).*

Optimal information policy

In this Section, I show that there exists an information policy that both solves the agent's relaxed problem and satisfies the full system of the individual rationality constraints. Given this, as Lemma 11 describes the solution to the relaxed problem, it also sheds light on the properties of the optimal information policy for the case of a promising project. These properties have a clear-cut economic interpretation as an investment schedule τ can be easily interpreted in terms of action recommendations for the principal.

An investment schedule τ can be without loss of generality implemented using a *direct recommendation mechanism* - a simple policy which has $M = \{0, 1\}$, where $m = 1$ received at date t is a recommendation to continue investing at t for the principal and $m = 0$ received at date t is a recommendation to stop investing at t .²⁴ Keeping this in mind, it is clear from Lemma 11 that the *optimal information policy* has to satisfy the following conditions. *First*, whenever the agent recommends the principal to stop, the second stage of the project is already completed. *Second*, the recommendation to stop is postponed so that the principal's individual rationality constraint is binding, which manifests in $V(\tau) = V^{NI}$. The first condition presents the key feature of the optimal information policy for the case of promising project: the agent discloses the completion of the second stage of the project, but *stays silent regarding the completion of the first stage of the project*. The intuition behind the agent's choice is simple: a recommendation to stop when no stages of the project are completed and the project deadline T is close does indeed incentivize the principal; however, it also reduces the total surplus generated that can be extracted via the agent's control of information. Meanwhile, the recommendation to stop when the two stages of the project are completed incentivizes the principal without reducing the total surplus generated. When $\kappa \leq \tilde{\kappa}(T, \lambda)$, a partially informative policy that discloses only the

²⁴The connection between an investment schedule τ and a direct recommendation mechanism implementing the schedule τ is simple: whenever, based on the evolution of the state process, τ suggests stopping the funding, the direct recommendation mechanism sends the message $m = 0$.

completion of the second stage provides sufficient incentives to the principal, and thus the agent uses it.²⁵

I proceed with obtaining an investment schedule that not only satisfies the conditions in Lemma 11 and solves the relaxed problem, but also satisfies the full system of the principal's individual rationality constraints in Lemma 9. Ensuring both is non-trivial. For instance, consider a mechanism that implements an investment schedule solving the agent's relaxed problem and assume it recommends to continue for $t \in [0, S^*)$, then at S^* recommends stopping if the second stage is already completed, but recommends to continue at all the subsequent dates $t \in (S^*, T]$. A no stopping recommendation drawn at S^* reveals that the state is either 0 or 1. Clearly, after sufficient time passes after S^* , the principal would attach a high probability to the second stage already being completed and would potentially be tempted to deviate from the recommendation to continue.²⁶ However, a *direct recommendation mechanism that implements an optimal investment schedule exists*. I present it in Proposition 2.5.2.

Assume $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$. The optimal mechanism does not provide a recommendation to stop during $t \in [0, S^*)$. At $t = S^*$, if the second stage of the project is already completed, then the mechanism recommends the principal to stop. If the second stage of the project is not yet completed, then the mechanism recommends the principal to stop at the moment of its completion $t = \tau_2$. Formally,

$$\tau = S^* \vee (\tau_2 \wedge T),$$

where S^* is chosen such that $V(\tau) = V^{NI}$, i.e., the respective constraint in the system of principal's individual rationality constraints is binding.

The recommendation mechanism starting from S^* generates recommendations to stop if the second stage is completed. As the recommendation to stop comes immediately at the completion of the second stage for all $t > S^*$, hearing no recommendation to stop reveals that the state is either 0 or 1. Further, as time goes on, the principal attaches a higher and higher probability to the state being 1, which ensures obedience to the recommendation to continue at each date. Further, the start of information provision S^* is sufficiently postponed to ensure that the principal's individual rationality constraint is binding either at $t = \bar{S}^{NI}$ or at $t = 0$.

The choice of S^* is driven by extraction of the principal's surplus and depends on κ in an intuitive way. First, consider the case $\kappa \in (\kappa^{ND}, \kappa^{NI}(T, \lambda)]$, the principal is willing to start investing and invests until $t = \bar{S}^{NI}$ in the no-information benchmark. The agent's optimal choice is to set $S^* > \bar{S}^{NI}$. Given such an information policy, the principal does not stop at \bar{S}^{NI} , the date of stopping in the no-information benchmark, and with probability one continues to invest during $t \in [\bar{S}^{NI}, S^*)$ even though the mechanism provides absolutely no information for all $t < S^*$. This is driven by the fact that the expected benefit from stopping at some future date $t \in [S^*, T]$ and obtaining the project payoff v with certainty

²⁵The "leading on" information policy in Ely and Szydlowski [2020] is similar: the only information that the policy provides is that the task is already completed and, thus, it is time to stop investing.

²⁶In other words, $V_t(\tau)$ drifts down over time and can get negative at some date.

compensates the flow losses of investing during $t \in [\bar{S}^{NI}, S^*]$.²⁷ Further, the agent sufficiently postpones S^* to ensure that she extracts the principal’s surplus and the principal gets precisely $V^{NI} \geq 0$.

In the case $\kappa \in (\kappa^{NI}(T, \lambda), \tilde{\kappa}(T, \lambda)]$, the principal is not willing to start in the no-information benchmark as his expected payoff from investing is negative. Thus, the agent chooses S^* to guarantee that the principal gets his reservation value $V^{NI} = 0$ and is thus willing to start investing at $t = 0$. The value of S^* is relatively lower as compared to the previous case: as the project is less attractive, to provide the principal sufficient incentives, the agent needs to start the information provision regarding the completion of the project earlier.

Finally, there exist many information policies that both solve the agent’s relaxed problem and satisfy the full system of constraints (2.7). This constitutes an important advantage for the agent: she can choose an optimal policy that is easier to implement from the real-world perspective, depending on the particular environment. In the optimal mechanism from Proposition 2.5.2, the recommendation to stop at some date t depends only on the current state of the world x_t . In an alternative delayed disclosure mechanism, the recommendation to stop arrives with a pre-specified delay after the second stage was completed. Thus, the recommendation depends only on the past history and not on the current state of the world. In an optimal delayed disclosure mechanism, the delay becomes smaller as more time passes. I characterize such a mechanism in Appendix 2.D.²⁸

Recall that, as Lemma 11 suggests, the key idea of the optimal information policy is that the agent postpones the disclosure of the completion of the project to extract more surplus, which makes the principal’s individual rationality constraint bind. The higher the cost-benefit ratio of the project κ becomes, the higher additional value the agent’s information policy needs to provide to the principal to ensure that his active individual rationality constraint is satisfied. The implication of this for the optimal information policy is presented in Lemma 12.

Lemma 12. *Assume $\kappa \in (\kappa^{ND}(T, \lambda), \tilde{\kappa}(T, \lambda)]$. Given the recommendation mechanism implementing an optimal investment schedule τ , for a fixed Poisson rate λ , the expected length of investment $E[\tau]$ decreases in the cost-benefit ratio κ .*

The intuition is that the higher the cost-benefit ratio of the project becomes, the sooner after the second stage of the project is completed the agent recommends the principal to stop. For the cost-benefit ratio as high as $\tilde{\kappa}(T, \lambda)$, the agent provides the recommendation to stop immediately at the date of completion of the second stage of the project. Further, for $\kappa > \tilde{\kappa}(T, \lambda)$, the optimal information policy satisfying the conditions in Lemma 11 ceases to be individually rational for the principal. As I show in the next Section, for $\kappa > \tilde{\kappa}(T, \lambda)$,

²⁷Similarly to the “leading on” information policy in Ely and Szydlowski [2020], the promises of future disclosure of the completion of the project are used as a “carrot” to make the receiver continue investing beyond the point at which he stops in the no-information benchmark.

²⁸The rich set of optimal direct recommendation mechanisms in my model encompasses both mechanisms in which the information provision depends only on the current state, similarly to the optimal mechanism in Ely and Szydlowski [2020], and the mechanisms that use delay, similarly to the delayed beep from Ely [2017].

in addition to immediate disclosure of the project completion, the agent provides the information regarding the completion of the first stage of the project.

2.5.3 Immediate disclosure of completion and an interim deadline

When $\kappa > \tilde{\kappa}(T, \lambda)$, the project is not promising for the principal and any investment schedule in which stopping occurs after $\tau_2 \wedge T$ with probability one violates the principal's individual rationality constraint. In other words, from the ex ante perspective the future reports disclosing only the completion of the project do not motivate the principal to start investing. Thus, an investment schedule that provides an individually rational expected payoff to the principal should assign a positive probability not only to stopping after the completion of the project, but also to stopping in either state 0, when no stages of the project are completed, or state 1, when only the first stage of the project is completed. I present the necessary conditions for an investment schedule to be optimal when the project is not promising in Lemma 13.

Lemma 13. *Assume $\kappa \in (\tilde{\kappa}(T, \lambda), \kappa^{FI}(T, \lambda))$. If an investment schedule τ solves agent's problem, then it satisfies the conditions*

1. *conditional on no completed stages of the project, stopping of the funding happens with a positive probability;*
2. *conditional on one completed stage of the project, stopping of the funding happens with probability zero;*
3. *conditional on two completed stages of the project, stopping of the funding happens immediately (at $t = \tau_2$) and with probability one.*

Stopping when only the first stage of the project is already completed is clearly inefficient. In state 1, the principal prefers to continue investing until the completion of the second stage and this principal's incentive to wait is aligned with the agent's incentive to postpone the stopping. Further, stopping in state 1 does not help overcome the problem of the violated individual rationality constraint under $\kappa > \tilde{\kappa}(T, \lambda)$. Meanwhile, assigning a positive probability to stopping when no stages are completed helps to overcome the problem of violated individual rationality constraint, as the principal benefits from stopping at some date t when the first stage of the project is not yet completed and the project deadline T is sufficiently close. Further, the agent clearly prefers the stopping of funding after the completion of the second stage rather than in state 0 as the former does not harm the total surplus generated. Thus, a schedule that is optimal assigns probability 1 to immediate stopping when the second stage is completed.

Lemma 13 implies that in an investment schedule, optimal for the agent, stopping after the completion of the second stage of the project happens immediately and stopping also happens given that no stages of the project are completed - i.e., at the *interim deadline chosen by the agent, which I denote by S_0^A* . Thus, Lemma 13 drastically simplifies the strategy space in the agent's design problem: it is only left to characterize the optimal interim deadline. At the outset of the

game, the agent designs a device that privately randomizes over the interim deadlines S_0^A . That is, the agent publicly chooses a distribution $F_{S_0^A}$, then an interim deadline is drawn according to it and privately observed by the agent. Next, the information starts to flow. The action recommendation to stop the funding satisfies the following investment schedule

$$\tau = \begin{cases} S_0^A, & \text{if } x_{S_0^A} = 0 \\ \tau_2 \wedge T, & \text{otherwise,} \end{cases} \quad (2.11)$$

where the principal knows only the distribution $F_{S_0^A}$, but not the draw of S_0^A .

Given that the completion of the second stage of the project is disclosed immediately, stopping at the interim deadline in state 0 leads to a loss of expected further investment flow for the agent, and a potential savings from abandoning a “stagnating” project for the principal. The agent’s payoff can be without loss of generality restated as the expected loss in future investment due to stopping at the interim deadline S_0^A in state 0 (rather than at $\tau_2 \wedge T$). Given this, the agent’s problem can be expressed as

$$\min_{F_{S_0^A}} E_{F_{S_0^A}} \left[\underbrace{P(x_{S_0^A} = 0) E[\tau_2 \wedge T - S_0^A | x_{S_0^A} = 0]}_{\text{expected loss in future investment given } S_0^A} \right], \quad (2.12)$$

subject to the system of the principal’s individual rationality constraints, which also have a natural interpretation as the expectation of principal’s savings on the future investment, which discontinues at S_0^A in state 0, minus the loss in the project completion profit due to stopping the funding at S_0^A in state 0.²⁹

Inspecting the agent’s expected loss in future investment in (2.12) reveals that if the agent postpones the interim deadline S_0^A , then two effects arise. First, the probability that stopping at the interim deadline will happen decreases. Second, the expected loss in total surplus due to stopping at the interim deadline rather than at $\tau_2 \wedge T$ decreases. Thus, the agent’s expected loss in future investment is decreasing in the date of interim deadline and the agent prefers an interim deadline with a later expected date.

Agent’s choice of the interim deadline distribution $F_{S_0^A}$ is affected by the two factors. First, as the expected loss in future investment in (2.12) is decreasing and convex in the date of the interim deadline, and thus the agent is risk-averse with respect to random interim deadlines. Thus, given some random interim deadline, the agent *directly* benefits from inducing a mean-preserving contraction. Second, the agent benefits from inducing a mean-preserving contraction *indirectly*. Inspecting the principal’s long-run payoff for some fixed S_0^A reveals that the principal is also risk-averse with respect to random interim deadlines. Thus, inducing a mean-preserving contraction makes the principal better-off and relaxes the individual rationality constraint at $t = 0$, hence, allowing the agent to postpone the expected interim deadline further. As a result the optimal for the agent interim deadline takes the form of a *deterministic date*. In other words, it is optimal for the agent to *publicly announce the interim deadline S_0^A at the outset*, so that the principal knows it.

²⁹The principal’s individual rationality constraint is presented in (2.39).

The agent has an incentive to postpone the interim deadline and uses her control of the information environment to postpone the deadline as much as possible so that the principal's individual rationality constraint at $t = 0$ binds. Figure 2.5 demonstrates the principal's long-run payoff as a function of the interim deadline, which I denote by S_0 . It is maximized at the principal-preferred interim deadline S_0^P , which was characterized in Lemma 8. The agent-preferred interim deadline S_0^A yields the principal's expected payoff of 0.




Figure 2.5: Principal's long-run payoff, V , as a function of an interim reporting deadline chosen by the agent, S_0 .

The next Proposition summarizes the optimal investment schedule, which can be without loss of generality implemented using a direct recommendation mechanism:

Assume $\kappa \in (\tilde{\kappa}(T, \lambda), \kappa^{FI}(T, \lambda))$. The optimal information policy is given by a direct recommendation mechanism that generates

- (a) the recommendation to stop at the moment of completion of the second stage of the project, $t = \tau_2$, and
- (b) a conditional recommendation to stop at the publicly announced interim deadline $t = S_0^A$. At S_0^A , stopping is recommended with certainty if the first stage of the project has not yet been completed.

Formally,

$$\tau = \begin{cases} S_0^A, & \text{if } x_{S_0^A} = 0 \\ \tau_2 \wedge T, & \text{otherwise,} \end{cases}$$

where S_0^A is chosen so that the principal's individual rationality constraint at $t = 0$ is binding, i.e., $V(\tau) = 0$.

A stopping recommendation at any date other than the interim deadline $t = S_0^A$ fully reveals that project is accomplished. Further, observing a recommendation to stop at the interim deadline, the principal learns that the milestone of the project has not yet been reached and becomes sufficiently pessimistic that the project will be completed by T .

A notable feature of the optimal information policy when the project is ex ante unattractive is its uniqueness. The only optimal instrument through which the agent fine tunes the incentive provision to the principal is the choice of interim

deadline, and there is a unique optimal way to set the deadline to make the principal's individual rationality constraint bind.

I proceed by considering the comparative statics of the interim deadline. Both the agent-preferred and the principal-preferred interim deadline, S_0^A and S_0^P , respectively, increase in the exogenous deadline T . This is because less time pressure relaxes the principal's individual rationality constraint and allows the agent to postpone the deadline further in order to extract the principal's surplus.

As the cost-benefit ratio increases up to κ^{FI} , the agent-preferred deadline converges to the principal-preferred deadline. An increase in the cost-benefit ratio of the project makes the principal's individual rationality constraint tighter.³⁰ As a result, for a higher κ , in the absence of completion of the first stage, the principal is willing to wait for a shorter time before stopping. Thus, both the interim deadline preferred by the principal S_0^P and the interim deadline chosen by the agent S_0^A are lower for a higher κ . Further, for a higher κ the agent has to choose an information policy relatively closer to the full-information benchmark to ensure that the individual rationality constraint at $t = 0$ is satisfied. Hence, the agent-chosen interim deadline S_0^A approaches S_0^P , the interim deadline preferred by the principal. The comparative statics of S_0^P and S_0^A with respect to the cost-benefit ratio of the project κ are presented in Figure 2.6.




Figure 2.6: Interim deadline chosen by the agent S_0^A (**dashed**) and preferred by the principal S_0^P (**thick**), as functions of the cost-benefit ratio of the project κ .

2.6 General preferences

In this Section, I allow for profit-sharing between the agent and the principal, varying degree of the agent's benefit from the flow of funds, and exponential discounting, and demonstrate that the optimal information policy still has the same properties as in the baseline model.

First, I assume that the agent and the principal share the project completion profit v : the principal gets $\alpha \cdot v$, while the agent gets $(1 - \alpha) \cdot v$, $\alpha \in (0, 1]$. Thus,

³⁰This is because the increase in κ makes the principal's instantaneous benefit from waiting decrease, and the normalized instantaneous cost of waiting becomes higher.

now the agent benefits not only from the flow of funds provided by the principal for running the project but also from the share in the profit. The assumption that the agent gets a share in the project completion profit is natural in many situations. In particular, the research documents that the entrepreneurs in innovative startups are up to some extent driven by giving vent to their entrepreneurial mindset and bringing their innovative ideas to life [Gundolf et al., 2017]. In such a context, a positive profit share of the agent captures that the agent is motivated by the success of the project.

Second, I assume that given a flow cost of c incurred by the principal, the agent obtains a flow benefit βc , $\beta \geq 0$. β can be interpreted as the agent's marginal benefit from using the funds provided by the principal for funding the project. Alternatively, for $\beta \in [0, 1]$ the loss of $1 - \beta$ of the amount of the transfer at each date can be interpreted as the transaction costs. Finally, setting $\beta = 0$ for some $\alpha < 1$ allows for abstracting from the agent's motives for diverting the funds and considering the agent motivated only by the success of the project.

Third, I allow for exponential discounting at a rate $r > 0$. Thus, the present value of a profit obtained at a date t is given by ve^{-rt} and the present value of a stream of funding up to date t is given by $\frac{1}{r}(1 - e^{-rt})c$. The following Proposition demonstrates that given the more general preference specification, the structure of the optimal disclosure, present in the baseline model, preserves.

- (a) When the cost-benefit ratio of the project is low, $\kappa \leq \tilde{\kappa}(T, \lambda, r, \alpha)$, the optimal investment schedule τ satisfies $\tau \geq \tau_2 \wedge T$, i.e., the agent recommends the principal to stop only after the completion of the second stage of the project.
- (b) When $\kappa > \tilde{\kappa}(T, \lambda, r, \alpha)$, the optimal investment schedule τ assigns positive probability both to the stopping in state 2 and state 0, i.e., the agent not only discloses the completion of the second stage of the project, but also specifies an interim deadline for the completion of the first stage.

Similarly to the baseline model, allowing the principal to stop after the project completion brings profit to the principal and thus leads to a relatively higher total surplus, which the agent can extract. Meanwhile, allowing the principal to stop at the interim deadline does not increase total surplus and serves solely as an expected payoff transfer from the agent to the principal. To see that, note that stopping when the first stage of the project is still incomplete allows the principal to save on the further costs of funding the project when over time the project proves to be “unsuccessful”. This can not be beneficial for the agent as she does not internalize the costs of running the project. Further, stopping at the interim deadline is strictly detrimental for the agent as she strictly prefers the principal to postpone the stopping of funding when no stages of the project are completed.³¹

When the project is sufficiently ex-ante attractive, the agent can motivate the principal to start funding the project without promising to stop the stagnant project at the interim deadline, and this is strictly beneficial for the agent. Thus, when the project is promising, the agent sets no interim deadlines, which in expectation gives her more funds and more experimentation for free.

³¹The probability of project success and stock of obtained funds are non-decreasing in the date of stopping irrespective of the number of the completed stages of the project.

2.7 Conclusion

A transparent flow of information is crucial for the successful management of any innovative project. However, the researcher, who controls the information on the progress of the project, often tends to have different motives than the investor. This leads to the question of how a researcher chooses the transparency of the flow of information about the progress of a project in order to manipulate the investor's funding decisions. I address this question in a dynamic information design model in which the agent commits to providing information to the principal with an incentive to postpone the principal's irreversible stopping of the funding.

I contribute to the dynamic information design literature by studying the problem of the dynamic provision of information regarding the progress of a multistage project, which evolves endogenously over time and needs to be completed before a deadline. I show that the agent's choice of which pieces of information to provide and when depends on the project being either *ex ante* attractive for the principal or not. In the case of a promising project, the agent provides only the good news that the project is completed and postpones the reports. In the case of an unattractive project, to motivate the principal to start funding the project the agent not only reports the completion of the project, but also helps the principal to find out when the project stagnates. To achieve this, the agent announces an interim deadline for the project – a certain date at which she recommends the principal to cut the funding of the project if the milestone of the project has not been reached.

2.A Notational conventions

The state process is given by $x_t, \forall t \in \mathbb{R}_+$, defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $t \in \mathbb{R}_+$. Its natural filtration is denoted by $F = (\mathcal{F}_t)_{t \geq 0}$. Throughout Appendices 2.B and 3.A, the following notational conventions are used:

1. I denote the random time at which the n th stage of the project is completed by τ_n . Formally, $\tau_n \in \mathbb{R}_+$ is a continuously distributed random variable that represents the first hitting time of $x_t = n$.

2. Consider some stopping time τ . Whenever “ τ ” stands as a term in an inequality, it stands for a realization of the stopping time τ and it should be read as “for each $\omega \in \Omega$ and corresponding $\tau(\omega)$ ”.

Example 1. “ $\tau_2 \wedge T \geq \tau$ ” should be read as “ $\tau_2(\omega) \wedge T \geq \tau(\omega)$, for all $\omega \in \Omega$ ”.

Example 2. “for all $t \in [S, \tau)$ ” should be read as “for all $t \in [S, \tau(\omega))$, for all $\omega \in \Omega$ ”.

3. The continuation value of the agent at time t , given τ , and conditional on $t < \tau$:

$$W_t(\tau) \mathbb{E}[\tau - t | t < \tau] c.$$

4. The total (continuation) surplus at time t , given τ , and conditional on $t < \tau$:

$$SV_t(\tau) W_t(\tau) + V_t(\tau).$$

5. Shorthand for posterior beliefs:

$$\begin{aligned} q_n(t) \mathbb{P}(x_t = n | t < \tau), \\ r_n(t) \mathbb{P}(x_\tau = n | t < \tau). \end{aligned}$$

2.B The principal’s continuation value

Here I present the alternative formulation of the principal’s continuation value (2.6). This helps me to study some of its properties for further use in Appendix 3.A. The continuation value of the principal at time t and given the investment schedule τ is given by (2.6). Postponing the stopping for Δt brings a benefit in the form of project completion payoff v iff the second stage of the project is completed within Δt . As x_t follows the Poisson process, the probability of two increments in a very short time Δt is negligibly small. Thus, during some Δt , the principal gets the project completion payoff v iff the first stage of the project has already been completed at some earlier time. Thus, postponing the stopping for Δt brings the principal v with probability $\lambda q_1(t) \Delta t$. The second stage is not completed within Δt with the complementary probability of $1 - \lambda q_1(t) \Delta t$. The principal’s continuation value is thus given by

$$\begin{aligned} V_t(\tau) &= (v \lambda q_1(t) - c) \Delta t + (1 - \lambda q_1(t) \Delta t) V_{t+\Delta t}(\tau) \\ &= v \lambda (q_1(t) - \kappa) \Delta t + (1 - \lambda q_1(t) \Delta t) V_{t+\Delta t}(\tau). \end{aligned}$$

Differentiating both sides w.r.t. Δt and considering $\lim_{\Delta t \rightarrow 0} (\cdot)$ yields

$$0 = v \lambda (q_1(t) - \kappa) - \lambda q_1(t) V_t(\tau) + \dot{V}_t(\tau),$$

which, after rearranging becomes

$$\dot{V}_t(\tau) = \lambda q_1(t) V_t(\tau) + v \lambda (\kappa - q_1(t)). \quad (2.13)$$

2.C Proofs

Proof of Lemma 7. The beliefs regarding the number of stages of the project completed by time t , x_t , evolve according to the Poisson process. The principal's unconditional beliefs are given by $p_0(0) = 1$ and for any t such that the game still continues,

$$\begin{aligned}\dot{p}_0(t) &= -\lambda p_0(t), \\ \dot{p}_1(t) &= \lambda(p_0(t) - p_1(t)), \\ \dot{p}_2(t) &= \lambda p_1(t),\end{aligned}\tag{2.14}$$

where $p_0(t) = e^{-\lambda t}$ and $p_1(t) = \lambda t e^{-\lambda t}$, $p_2(t) = 1 - p_0(t) - p_1(t)$. The principal's problem is given by

$$\max_{S \in [0, T]} \{v \cdot p_2(S) - c \cdot S\}.\tag{2.15}$$

I start with analyzing the choice of S for the interior solution case, $S \in (0, T)$. F.O.C. for (2.15) is given by

$$v \cdot \dot{p}_2(S) = c,\tag{2.16}$$

or, equivalently, $p_1(S) = \kappa$. There are two values satisfying (2.16): \bar{S} and \bar{S}^{NI} , $\bar{S} < \bar{S}^{NI}$. At each $t \in (\bar{S}, \bar{S}^{NI})$ the principal receives a net positive payoff flow. Thus, stopping at \bar{S} is not optimal and the only candidate for optimal stopping is \bar{S}^{NI} .³² Further, one can obtain the closed form expression for the interior stopping time \bar{S}^{NI} from (2.16):

$$\bar{S}^{NI} = -\frac{1}{\lambda} \mathcal{W}_{-1}(-\kappa),\tag{2.17}$$

where $\mathcal{W}_{-1}(x)$ denotes the negative branch of the Lambert W function. \bar{S}^{NI} is well-defined for any $\kappa < e^{-1}$.

Thus, the solution to (2.15) could potentially be 0, \bar{S}^{NI} , or T . I proceed with a useful lemma.

Lemma 14. *The following is true regarding the principal's continuation value in the no-information benchmark, \bar{V}_t^{NI} : if $\bar{V}_t^{NI} \geq 0$, for some $t \in [0, \bar{S}^{NI} \wedge T]$, then $V^{NI}(s) \geq 0$, for all $s \in [t, \bar{S}^{NI} \wedge T]$.*

Proof. The principal's continuation value in the no-information benchmark is given by

$$\bar{V}_t^{NI} = [p_2(T \wedge \bar{S}^{NI}) - p_2(t)]v - (T \wedge \bar{S}^{NI} - t)c.\tag{2.18}$$

Further,

$$\dot{V}^{NI}(t) = v\lambda(\kappa - e^{-\lambda t}\lambda t) = v\lambda(\kappa - p_1(t)).$$

$p_1(t) \leq \kappa$ for all $t \in [0, \bar{S}]$ and $p_1(t) \geq \kappa$ for all $t \in [\bar{S}, \bar{S}^{NI} \wedge T]$. Thus, \bar{V}_t^{NI} increases for $t \in [0, \bar{S}]$, decreases for $t \in [\bar{S}, T \wedge \bar{S}^{NI}]$, and $V^{NI}(T \wedge \bar{S}^{NI}) = 0$, which establishes the result. \square

³² \bar{S} is a local minimum of the objective.

Lemma 14 implies that if $V^{NI}(0) \geq 0$ and the principal chooses to opt in at $t = 0$, then $\bar{V}_t^{NI} \geq 0$, $t \in [0, \bar{S}^{NI} \wedge T]$, i.e., he invests until $t = T \wedge \bar{S}^{NI}$. This implies that the solution to (2.15) is either $T \wedge \bar{S}^{NI}$ or 0.

Finally, at $t = 0$ the principal chooses to start investing or not. The condition for the principal to start investing at $t = 0$ is given by

$$V^{NI} \geq 0. \quad (2.19)$$

To specify the set of parameters for which (2.19) is satisfied, I obtain the cutoff value of κ given T and λ . Such a parameterization is intuitive: κ above the cutoff level corresponds to a project with sufficiently high normalized cost-benefit ratio and implies that the principal does not opt in. I denote this cutoff by $\kappa^{NI}(T, \lambda)$. This solves (2.19) holding with equality. Two cases are possible.

Case 1: $T \leq \bar{S}^{NI} \iff T \leq -\frac{1}{\lambda} \mathcal{W}_{-1}(-\kappa)$. This inequality is satisfied when either $\frac{1}{\lambda} > T$ or $\begin{cases} \frac{1}{\lambda} \leq T \\ \kappa \leq e^{-\lambda T} \lambda T. \end{cases}$ Given $T \leq \bar{S}^{NI}$, (2.19) holding with equality becomes

$$p_2(T)v - Tc = 0.$$

Solving it for κ yields $\kappa = e^{-\lambda T} \left(\frac{e^{\lambda T} - 1}{\lambda T} - 1 \right)$.

Case 2: $T > \bar{S}^{NI}$. This inequality is satisfied when $\frac{1}{\lambda} \leq T$ and $\kappa > e^{-\lambda T} \lambda T$. Given $T > \bar{S}^{NI}$, (2.19) holding with equality becomes

$$vp_2(\bar{S}^{NI}) - c\bar{S}^{NI} = 0 \iff v(1 - p_0(\bar{S}^{NI}) - p_1(\bar{S}^{NI})) = c\bar{S}^{NI},$$

where (recall that $\dot{p}_2(\bar{S}^{NI}) = \frac{c}{v}$)

$$p_0(\bar{S}^{NI}) = \frac{1}{\lambda^2 \bar{S}^{NI}} \dot{p}_2(\bar{S}^{NI}) = \frac{c}{\lambda^2 \bar{S}^{NI} v} = \frac{\kappa}{\lambda \bar{S}^{NI}}$$

and

$$p_1(\bar{S}^{NI}) = \frac{1}{\lambda} \dot{p}_2(\bar{S}^{NI}) = \frac{c}{\lambda v} = \kappa.$$

Consequently,

$$vp_2(\bar{S}^{NI}) - c\bar{S}^{NI} = v - v \cdot \kappa \left(1 + \lambda \bar{S}^{NI} + \frac{1}{\lambda \bar{S}^{NI}} \right).$$

Let $y\lambda\bar{S}^{NI}$. Note that, by definition, $y > 1$. Then $\kappa = ye^{-y}$, and so

$$(vp_2(\bar{S}^{NI}) - c\bar{S}^{NI})/v = 1 - e^{-y}(1 + y + y^2).$$

It follows that $V^{NI}(0)$ is nonnegative whenever $\lambda\bar{S}^{NI} \geq y_0 \doteq 1.79328$, which is equivalent to

$$\kappa \leq \kappa_0 \doteq 0.298426.$$

Finally, putting the two cases together yields

$$\kappa^{NI}(T, \lambda) = \begin{cases} \kappa_0 \doteq 0.298426, & \text{if } \frac{1}{\lambda} \leq T \text{ and } \kappa \geq e^{-\lambda T} \lambda T \\ e^{-\lambda T} \left(\frac{e^{\lambda T} - 1}{\lambda T} - 1 \right), & \text{otherwise.} \end{cases} \quad (2.20)$$

□

Proof of Lemma 8. The principal chooses $a_t \in \{0, 1\}$ sequentially given the observed realizations of $x_t \in \{0, 1, 2\}$. Whenever the principal observes $t = \tau_2$, he immediately chooses $a_t = 0$ and gets v .

Consider the case $x_t = 1, t < T$, i.e., the first stage of the project has already been completed. As x_t follows a Poisson process, in expectation it would take $\frac{1}{\lambda}$ units of time for the second stage to be completed and its completion brings the principal the value of v . Thus, the necessary and sufficient condition for the principal to invest at t when $x_t = 1, t < T$ is given by

$$v - c \cdot \frac{1}{\lambda} \geq 0 \iff \kappa \leq 1$$

Assume that $\kappa \leq 1$ holds and $x_t = 1$; thus, the principal chooses to invest at t . In that case, the principal invests until $\tau_2 \wedge T$. To see this, recall that the only news that the principal can receive given $x_t = 1, t < T$ is the completion of the second stage of the project, τ_2 , which leads to immediate stopping. At each $t < \tau_2 \wedge T$ such that $x_t = 1$, choosing $a_t = 0$ yields an instantaneous expected payoff of 0, while choosing $a_t = 1$ yields an instantaneous expected payoff of $\lambda v \Delta t - c \Delta t$. Thus, $\kappa \leq 1$ suffices for the principal to invest until $\tau_2 \wedge T$.

Consider now the case of $x_t = 0, t < T$, i.e., no stages of the project have yet been completed. Postponing the stopping for Δt brings the instantaneous expected payoff of $V_{t|1}^{FI} \lambda \Delta t - c \Delta t$, where $V_{t|1}^{FI}$ is the principal's continuation value at time t under full information, conditional on the completion of the first stage of the project. I proceed with obtaining the expression for $V_{t|1}^{FI}$. By definition, the principal gets v whenever the second stage is completed not later than T . The principal invests until $\tau_2 \wedge T$, and knows that at t the first stage of the project is already completed; thus, $V_{t|1}^{FI}$ is given by

$$V_{t|1}^{FI} = v \text{P}(\tau_2 \leq T | x_t = 1) - c \text{E}[\tau_2 \wedge T - t | x_t = 1].$$

$\tau_2 | x_t = 1$ corresponds to the time between two consecutive Poisson arrivals, and thus has exponential distribution. First, consider $\text{P}(\tau_2 \leq T | x_t = 1)$:

$$\text{P}(\tau_2 \leq T | x_t = 1) = 1 - e^{-\lambda(T-t)}.$$

Next, consider $\text{E}[\tau_2 \wedge T - t | x_t = 1]$:

$$\begin{aligned} & \text{E}[\tau_2 \wedge T | x_t = 1] - t \\ &= \text{P}(\tau_2 \leq T | x_t = 1) \int_t^T z \cdot \frac{\lambda e^{-\lambda(z-t)}}{\text{P}(\tau_2 \leq T | x_t = 1)} dz + \text{P}(\tau_2 > T | x_t = 1) T - t \\ &= \frac{1}{\lambda} \left(1 - e^{-\lambda(T-t)} \right) + t - e^{-\lambda(T-t)} T + \text{P}(\tau_2 > T | x_t = 1) T - t \\ &= \frac{1}{\lambda} \left(1 - e^{-\lambda(T-t)} \right). \end{aligned}$$

(2.21)

Thus,

$$\begin{aligned} V_{t|1}^{FI} &= v \left(1 - e^{-\lambda(T-t)}\right) - c \frac{1}{\lambda} \left(1 - e^{-\lambda(T-t)}\right) \\ &= \left(v - \frac{c}{\lambda}\right) \left(1 - e^{-\lambda(T-t)}\right). \end{aligned} \quad (2.22)$$

From (2.22) one observes that $V_{t|1}^{FI}$ decreases in t . If the net instantaneous benefit given by $V_{t|1}^{FI} \lambda \Delta t - c \Delta t$ gets as low as 0 at some t , then the principal chooses to stop investing at this t . I denote the time at which the net instantaneous benefit reaches 0 by S_0^P . S_0^P can be obtained from $(\lambda V_1^{FI}(S_0^P) - c) \Delta t = 0$. Thus,

$$S_0^P = T + \frac{1}{\lambda} \log \left(\frac{1 - 2\kappa}{1 - \kappa} \right). \quad (2.23)$$

The principal is willing to start investing iff at $t = 0$ the expected payoff from investing at $t = 0$ covers the costs of investing, i.e. $(\lambda V_1^{FI}(0) - c) \Delta t \geq 0$. From (2.23), this corresponds to $S_0^P \geq 0$. I denote the upper bound on the cost-benefit ratio κ such that the principal chooses to start investing in $t = 0$ under full information by $\kappa^{FI}(T, \lambda)$, I solve $S_0^P = 0$ for κ and obtain

$$\kappa^{FI}(T, \lambda) = \frac{1 - e^{-\lambda T}}{2 - e^{-\lambda T}}. \quad (2.24)$$

In summary, under full information, if $\kappa \leq \kappa^{FI}(T, \lambda)$, then the principal starts investing at $t = 0$. Further, he stops at S_0^P if the first stage of the project has not been completed by that time. Otherwise, he proceeds to invest until $\tau_2 \wedge T$. \square

Proof of Proposition 2.5.1. I provide the proof for each of the four parametric cases below.

1. *The case of $\kappa \leq \kappa^{ND}(T, \lambda)$.*

$\kappa^{ND}(T, \lambda)$ is defined as follows: for any $\kappa \leq \kappa^{ND}(T, \lambda)$, the principal invests until T in the no-information benchmark. From Lemma 7, if the principal is willing to start investing, i.e., $\kappa \leq \kappa^{NI}(T, \lambda)$, then

$$S^{NI} = \bar{S}^{NI} \wedge T.$$

For the sake of instruction, below I consider relaxing the Assumption 2.1 and demonstrate how the relation between $\kappa^{ND}(T, \lambda)$ and $\kappa^{NI}(T, \lambda)$ changes between Case a (relaxed Assumption 2.1) and Case b (Assumption 2.1 holds).

Case a. $e^{\lambda T} \leq \lambda T (\lambda T + 1) + 1$. In this case, whenever the principal is willing to start investing in the no-information benchmark, she invests until T , i.e., $\kappa^{ND}(T, \lambda) = \kappa^{NI}(T, \lambda)$, where $\kappa^{NI}(T, \lambda)$ is given by (2.20). To see that, first, consider the extreme sub-case in which $T < \frac{1}{\lambda}$. As $-\lambda \bar{S}^{NI}$ must belong to -1 axis of Lambert W function, it has a lower bound corresponding to $\frac{1}{\lambda}$. Thus, $T < \bar{S}^{NI}$ for any $\kappa(T, \lambda)$. Second, consider $\lambda T \in [1, \tilde{\lambda} T]$, where $\tilde{\lambda} T$ solves $e^{\lambda T} = \lambda T (\lambda T + 1) + 1$. In this case, from (2.17), if $\kappa(T, \lambda) \leq e^{-\lambda T} \lambda T$ ($\kappa(T, \lambda) \geq e^{-\lambda T} \lambda T$, respectively), then $T \leq \bar{S}^{NI}$ ($T \geq \bar{S}^{NI}$, respectively). However, $\kappa^{NI}(T, \lambda) \leq e^{-\lambda T} \lambda T$. Thus, $\kappa^{ND}(T, \lambda) = \kappa^{NI}(T, \lambda)$.

Case b. $e^{\lambda T} > \lambda T (\lambda T + 1) + 1$. As before, it holds that if $\kappa(T, \lambda) \leq e^{-\lambda T} \lambda T$ ($\kappa(T, \lambda) \geq e^{-\lambda T} \lambda T$), then $T \leq \bar{S}^{NI}$ ($T \geq \bar{S}^{NI}$, respectively). Denote

$$\kappa^{ND}(T, \lambda) e^{-\lambda T} \lambda T.$$

As $\kappa^{NI}(T, \lambda) > \kappa^{ND}(T, \lambda)$, two cases emerge. If $0 < \kappa \leq \kappa^{ND}(T, \lambda)$, then $T \leq \bar{S}^{NI}$, and from $\kappa \leq \kappa^{NI}(T, \lambda)$, it holds that $S^{NI} = T$ and as the agent does not strictly benefit from disclosing any information, she chooses non-disclosure. If $\kappa > \kappa^{ND}(T, \lambda)$, then $T > \bar{S}^{NI}$ and the agent can potentially benefit from information disclosure.

2. *The case of $\kappa^{ND}(T, \lambda) < \kappa \leq \tilde{\kappa}(T, \lambda)$.*

The result is established in Proposition 2.5.2.

3. *The case of $\tilde{\kappa}(T, \lambda) < \kappa < \kappa^{FI}(T, \lambda)$.*

The result is established in Proposition 2.5.3.

4. *The case of $\kappa \geq \kappa^{FI}(T, \lambda)$.*

The principal's long-run payoff in the full-information benchmark non-positive. Thus, the agent can not strictly benefit from information disclosure and chooses non-disclosure. □

Proof of Lemma 9. Necessity. Assume $V_t(\tau) < 0$ for some t . In that case, it is optimal for the principal to deviate to stopping at $t < \tau$. Thus, there is no information policy σ , for which this τ is the principal's best reply. Assume $V_\tau^{NI} \geq 0$. Thus, the principal deviates to stopping at $t > \tau$, and there is no σ , for which this τ is the best reply.

Sufficiency. Assume (2.7) holds. $V_t(\tau) \geq 0$ for all $t < \tau$ implies that the principal prefers to continue rather than to stop the funding for all $t < \tau$. Thus, it can not be that case that the principal stops before τ . Further, $V_\tau^{NI} < 0$ implies that, conditional on reaching the date of stopping τ , it is better for the principal to stop immediately rather than to stop at $t > \tau$.

Consider a direct recommendation mechanism σ with $M = \{0, 1\}$ such that whenever, based on the evolution of the state process, the considered investment schedule τ suggests stopping the funding, the direct recommendation mechanism sends the message $m = 0$ to the principal. As it is optimal for the principal to stop at τ , τ is the principal's best reply to σ . □

Proof of Lemma 10. Consider the recommendation mechanism immediately disclosing the completion of the second stage of the project; it is given by $\tau = \tau_2 \wedge T$. There exists such $\tilde{\kappa}(T, \lambda)$ that solves the principal's binding $t = 0$ individual rationality constraint when $\tau = \tau_2 \wedge T$:

$$V(\tau_2) = 0, \tag{2.25}$$

where

$$\begin{aligned} V(\tau_2) &= p_2(T) v - E[\tau_2 \wedge T] c \\ &= v \left(1 - e^{-\lambda T} - \lambda T e^{-\lambda T} \right) - c \frac{1}{\lambda} \left(2 - 2e^{-\lambda T} - \lambda T e^{-\lambda T} \right). \end{aligned} \tag{2.26}$$

The solution to equation (2.25) is given by

$$\tilde{\kappa}(T, \lambda) = \frac{1 - e^{\lambda T} + \lambda T}{2 - 2e^{\lambda T} + \lambda T}. \tag{2.27}$$

Further, $\kappa > \tilde{\kappa}(T, \lambda) \Rightarrow V(\tau_2) < 0$ and $\kappa \leq \tilde{\kappa}(T, \lambda) \Rightarrow V(\tau_2) \geq 0$. \square

Proof of Lemma 11. Consider the case of $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$. The agent's relaxed problem for this case has the individual rationality constraints only for $t \in [0, \bar{S}^{NI}]$, and it is given by

$$\begin{aligned} & \max_{\tau \in \mathcal{T}} \{c \cdot \mathbb{E}[\tau]\} \\ & \text{s.t. } V_t(\tau) \geq 0, \forall t \in [0, \bar{S}^{NI}], \end{aligned} \quad (2.28)$$

where $V_t(\tau)$ is given by (2.6) and \mathcal{T} is the set of stopping times with respect to the natural filtration of x_t .

Consider the candidate investment schedule τ such that $\tau \geq \bar{S}^{NI} \vee (\tau_2 \wedge T)$ and $V(\tau) = V^{NI}$, where V^{NI} is given by (2.3). I start with arguing that the candidate τ satisfies the system of individual rationality constraints. From Lemma 7, given candidate τ , the principal invests until \bar{S}^{NI} with certainty and the constraints in (2.28) are satisfied for all $t \in [0, \bar{S}^{NI}]$. Further, τ implies that $V_{\bar{S}^{NI}}(\tau) = 0$, i.e., the individual rationality constraint at $t = \bar{S}^{NI}$ is binding.

I proceed with arguing that the candidate τ maximizes the agent's objective function in (2.28). The agent's objective can be WLOG written out as:

$$W(\tau) = \underbrace{\mathbb{P}(x_\tau = 2)v}_{\text{total surplus}} - \underbrace{V(\tau)}_{\text{principal's surplus}}. \quad (2.29)$$

By Lemma 10, an investment schedule τ that assigns probability one to $\tau \geq \tau_2 \wedge T$ satisfies the individual rationality constraint at $t = 0$ in (2.28). Note that, given $\tau \geq \tau_2 \wedge T$, the total surplus in (2.29) is given by $\mathbb{P}(x_T = 2)v$, i.e., total surplus achieves its upper bound determined by the exogenously given project deadline T . The principal's surplus in (2.29) is given by $V(\tau) = V^{NI}$, i.e., principal's surplus achieves its lower bound specified by (2.3). This can be seen from the principal's decision problem, in which he best replies to an information policy σ . As σ allows the principal to condition his actions on the information regarding the evolution of the state process, the principal's equilibrium payoff can not be lower than V^{NI} , his equilibrium payoff when he is restricted to choosing actions without conditioning them on the information about the state process. Thus, τ solves the relaxed problem (2.28).

Consider the case of $\kappa \in (\kappa^{NI}(T, \lambda), \tilde{\kappa}(T, \lambda)]$. The agent's relaxed problem for this case has the individual rationality constraint only for the initial period, and it is given by

$$\begin{aligned} & \max_{\tau \in \mathcal{T}} \{c \cdot \mathbb{E}[\tau]\} \\ & \text{s.t. } V(\tau) \geq 0, \end{aligned} \quad (2.30)$$

where $V(\tau) = \mathbb{P}(x_\tau = 2)v - \mathbb{E}[\tau]c$.

Consider candidate investment schedule τ such that $\tau \geq \tau_2 \wedge T$ and $V(\tau) = V^{NI}$. For such τ , agent's expected payoff (2.29) is given by $\mathbb{P}(x_T = 2)v - V^{NI}$. As discussed for the parametric case $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$, the first term is at its upper bound. To see that the second term is at its lower bound, note that, from Lemma 7, $V^{NI} = 0$, and thus the individual rationality constraint in (2.30) is binding. Hence, τ solves the relaxed problem (2.30). \square

Proof of Proposition 2.5.2. The proof covers the case $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$ and the case $\kappa \in (\kappa^{NI}(T, \lambda), \tilde{\kappa}(T, \lambda)]$ separately.

1. *The case of $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$.*

I start with proving the existence of S^* such that $V(\tau) = V^{NI}$. Assume that $S^* > \bar{S}^{NI}$. For all $t \in [\bar{S}^{NI}, S^*)$, stopping never occurs, at $t = S^*$ it occurs if $x_{S^*} = 2$, and for all $t \in (S^*, \tau)$ it occurs at $t = \tau_2 \wedge T$. For $t \in [S^*, \tau)$, the absence of stopping induces posteriors $q_n(t)$. Further, for $t \in [S^*, \tau)$ the principal discounts future benefits from postponing stopping using the probability of the state being 2. Hence, the continuation value at $t = \bar{S}^{NI}$ can be written as

$$V_{\bar{S}^{NI}}(\tau) = v\lambda \left(\int_{\bar{S}^{NI}}^{S^*} p_1(z) - \kappa dz + \int_{S^*}^T (q_1(z) - \kappa)(1 - P(x_z = 2)) dz \right). \quad (2.31)$$

The principal's long-run payoff is given by

$$V(\tau) = \int_0^{\bar{S}^{NI}} (v \cdot p_1(s) \lambda - c) ds + V_{\bar{S}^{NI}}(\tau),$$

where $\int_0^{\bar{S}^{NI}} (v \cdot p_1(s) \lambda - c) ds = V^{NI}$. Thus, to ensure that S^* makes the individual rationality constraint bind at $t = \bar{S}^{NI}$, i.e., $V(\tau) = V^{NI}$, it is necessary and sufficient that $V_{\bar{S}^{NI}}(\tau) = 0$. Using (2.31), this equation can be written as

$$\int_{\bar{S}^{NI}}^{S^*} \kappa - p_1(z) dz = \int_{S^*}^T (q_1(z) - \kappa)(1 - P(x_z = 2)) dz.$$

Let $g(S) = \int_{\bar{S}^{NI}}^S \kappa - p_1(z) dz$ and $k(S) = \int_S^T (q_1(z) - \kappa)(1 - P(x_z = 2)) dz$, $S \in [\bar{S}^{NI}, \tau)$. $q_1(t) \geq \kappa$, for all $t \in [S^*, T)$. Thus, $g(\bar{S}^{NI}) = 0, k(\bar{S}^{NI}) > 0$. Further, $p_1(t) < \kappa$, for all $t \in (\bar{S}^{NI}, T]$. Hence, $g(T) > 0, k(T) = 0$. Finally, $p_1(t) \leq \kappa$, for all $t \in [\bar{S}^{NI}, T]$ implies that $g'(S) \geq 0$, for all $s \in [\bar{S}^{NI}, T]$, and $q_1(t) \geq \kappa$, for all $t \in [S^*, T]$ implies that $k'(S) \leq 0$, for all $s \in [S^*, T]$. Thus, by the intermediate value theorem, there exists S^* solving $V_{\bar{S}^{NI}}(\tau) = 0$. Thus, there exists $S^* > \bar{S}^{NI}$ such that principal's individual rationality constraint is binding at $t = \bar{S}^{NI}$.

I proceed with proving that the investment schedule τ satisfies the conditions in Lemma 9 and thus it is obedient.

First, consider $t \leq \bar{S}^{NI}$. The principal's continuation value for all $t \in [0, \bar{S}^{NI}]$ can be written as

$$V_t(\tau) = \int_t^{\bar{S}^{NI}} v\lambda(p_1(s) - \kappa) ds + V_{\bar{S}^{NI}}(\tau).$$

Given the binding individual rationality constraint, it becomes

$$V_t(\tau) = \int_t^{\bar{S}^{NI}} v\lambda(p_1(s) - \kappa) ds, \text{ for all } t \in [0, \bar{S}^{NI}).$$

Finally, note that $V_t(\tau)$ above is equivalent to V_t^{NI} given by (2.18). Lemma 7 implies that given $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$, $V^{NI}(0) = V(\tau) \geq 0$. Further, Lemma 14 implies that $V(\tau) \geq 0 \Rightarrow V_t(\tau) \geq 0, \forall t \in [0, \bar{S}^{NI})$.

Second, consider $t \in [\bar{S}^{NI}, S^*]$. Given $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$, $p_1(t) \leq \kappa, \forall t \in [\bar{S}^{NI}, S^*]$. Thus, $V_t^{NI} = 0, \forall t \in [\bar{S}^{NI}, S^*]$. The principal's continuation value is given by

$$V_t(\tau) = \int_t^{S^*} v\lambda(p_1(s) - \kappa) ds + V_{S^*}(\tau). \quad (2.32)$$

As $p_1(t) \leq \kappa, \forall t \in [\bar{S}^{NI}, S^*]$, $\int_t^{S^*} v\lambda(p_1(s) - \kappa) ds \leq 0$ and it is increasing in t . As $V_{\bar{S}^{NI}}(\tau) = 0$, where $V_{\bar{S}^{NI}}(\tau)$ is given by (2.31), it follows that $V_t(\tau) \geq 0, \forall t \in [\bar{S}^{NI}, S^*]$.

Third, consider $t \in [S^*, \tau)$. The absence of stopping at $t \geq S^*$ reveals that $x_t \neq 2$. Thus, $q_1(t) = \frac{p_1(t)}{p_0(t) + p_1(t)} = \frac{\lambda t}{1 + \lambda t}$, $\forall t \in [S^*, \tau)$, and, thus, $\dot{q}_1(t) > 0$. Further, $q_1(S^*) > \kappa$. The continuation value $\forall t \in [S^*, \tau)$ is given by

$$V_t(\tau) = E[\int_t^\tau v\lambda(q_1(z) - \kappa) dz | t < \tau].$$

Thus, $V_t(\tau) \geq 0, \forall t \in [S^*, \tau)$.

2. The case of $\kappa^{NI}(T, \lambda) < \kappa \leq \tilde{\kappa}(T, \lambda)$.

I start with proving the existence of S^* such that $V(\tau) = 0$. For all $t \in [0, S^*)$, stopping never occurs, at $t = S^*$ it occurs if $x_{S^*} = 2$, and for all $t \in (S^*, T]$ it occurs at $t = \tau_2 \wedge T$. The principal's long-run payoff can be written as

$$V(\tau) = v\lambda \left(\int_0^{S^*} p_1(z) - \kappa dz + \int_{S^*}^T (q_1(z) - \kappa)(1 - P(x_z = 2)) dz \right). \quad (2.33)$$

To ensure that S^* makes the individual rationality constraint bind at $t = 0$, it is necessary and sufficient that $V(\tau) = 0$. The next step of the proof consist of inspecting (2.33) to establish that there exists S^* ensuring that $V(\tau) = 0$. It follows the respective part from the proof for the parametric case $\kappa^{ND}(T, \lambda) < \kappa \leq \kappa^{NI}(T, \lambda)$, imposing $\bar{S}^{NI} = 0$ in it everywhere; thus, I omit it for the sake of space.

I proceed with proving that the investment schedule τ satisfies the conditions in Lemma 9 and thus it is obedient. The principal's continuation value is given by (2.32). As $\kappa \in (\kappa^{NI}(T, \lambda), \tilde{\kappa}(T, \lambda)]$, it follows from Lemma 7 that $V_t^{NI} = 0, \forall t \in [0, S^*]$. First, assume $S^* \leq \bar{S}^{NI}$. From the proof of Lemma 7, it follows that $p_1(t) \leq \kappa, \forall t \in [0, \bar{S}]$, and $p_1(t) \geq \kappa, \forall t \in [\bar{S}, \bar{S}^{NI}]$. Thus,

$$\int_t^{\bar{S}^{NI}} v\lambda(p_1(s) - \kappa) ds \geq \int_0^{\bar{S}^{NI}} v\lambda(p_1(s) - \kappa) ds, \forall t \in [0, \bar{S}^{NI}]. \quad (2.34)$$

As $V_t(\tau)$ is given by (2.32), $V(\tau) = 0$ and (2.34) imply that $V_t(\tau) \geq 0, \forall t \in [0, S^*]$. Second, assume $S^* \geq \bar{S}^{NI}$. As $V(\tau) = 0$ and $\int_0^{\bar{S}^{NI}} v\lambda(p_1(s) - \kappa) ds < 0$, it must be that $V(\bar{S}^{NI}) > 0$. Further, $\int_t^{S^*} v\lambda(p_1(s) - \kappa) ds$ increases in t for $t \in [\bar{S}^{NI}, S^*]$. Thus, $V_t(\tau) \geq 0, \forall t \in [0, S^*]$.

Finally, the proof that $V_t(\tau) \geq 0, \forall t \in [S^*, \tau)$ follows the the respective part of the proof for the parametric case $\kappa \in (\kappa^{ND}(T, \lambda), \kappa^{NI}(T, \lambda)]$; thus, I omit it for the sake of space.

□

Proof of Lemma 12. I provide the proof for the parametric cases $\kappa^{ND}(T, \lambda) < \kappa \leq \kappa^{NI}(T, \lambda)$ and $\kappa^{NI}(T, \lambda) < \kappa \leq \tilde{\kappa}(T, \lambda)$ separately.

1. *The case of $\kappa^{ND}(T, \lambda) < \kappa \leq \kappa^{NI}(T, \lambda)$.*

Under any obedient optimal policy, the principal's individual rationality constraint is binding, thus, $V(\tau) = V^{NI}$, or equivalently $p_2(T)v - E[\tau]c = p_2(\bar{S}^{NI})v - \bar{S}^{NI}c$. Thus,

$$E[\tau] = \frac{1}{\lambda\kappa} \left(p_2(T) - p_2(\bar{S}^{NI}) \right) + \bar{S}^{NI}.$$

Differentiating both sides with respect to κ yields

$$\frac{\partial E[\tau]}{\partial \kappa} = \frac{e^{-T\lambda}(1 + T\lambda) - e^{-\bar{S}^{NI}\lambda} - \kappa}{\kappa^2\lambda}.$$

The equation

$$e^{-T\lambda}(1 + T\lambda) - e^{-\bar{S}^{NI}\lambda} - \kappa = 0$$

can be equivalently rewritten as

$$e^{-T\lambda} - e^{-\bar{S}^{NI}\lambda} = \kappa - e^{-T\lambda}T\lambda.$$

It has a unique solution corresponding to $\kappa = \kappa^{ND}(T, \lambda)e^{-T\lambda}T\lambda$. As $\kappa > \kappa^{ND}(T, \lambda)$, it holds that $\partial E[\tau]/\partial \kappa < 0$.

2. *The case of $\kappa^{NI}(T, \lambda) < \kappa \leq \tilde{\kappa}(T, \lambda)$.*

The principal's long-run payoff under any obedient optimal policy is given by

$$E[\tau]c = p_2(T)v.$$

Rewriting it equivalently, $E[\tau] = \frac{1}{\lambda} \frac{1}{\kappa} p_2(T) \Rightarrow \partial E[\tau]/\partial \kappa < 0$. □

Proof of Lemma 13. Lemma 10 implies that if a schedule τ assigns zero probability to stopping in states 0 and 1 then $V(\tau) < 0$ and the individual rationality constraint is violated. Thus, the necessary condition for a schedule τ to be individually rational under $\kappa \in (\tilde{\kappa}(T, \lambda), \kappa^{FI}(T, \lambda))$ is that it assigns a positive probability to stopping not only in state 2, but also to stopping in either state 0 or state 1. Consider a schedule τ that assigns a positive probability to stopping in state 1. Consider an alternative schedule τ' which is induced by reallocating the probability mass of stopping in state 1 to stopping at $\tau_2 \wedge T$. Lemma 8 suggests that in state 1 the principal strictly benefits from postponing the stopping until the second stage of the project is completed. Thus, $V(\tau') > V(\tau)$. Further, under τ' the principal invests strictly longer, in expectation. Thus, $W(\tau') > W(\tau)$. Thus, for a schedule to be optimal it should not assign a positive probability to stopping in state 1.

Next, consider a schedule τ which assigns a positive probability to stopping in states 0 and 2. Assume that the stopping in state 0 happens at date S , which can be either deterministic or stochastic: if $x_S = 0$ then $\tau = S$, otherwise, $\tau \geq \tau_2 \wedge T$ and there exists $\omega \in \Omega$ such that $\tau(\omega) > \tau_2(\omega)$, i.e., with a positive probability, stopping in state 2 happens strictly after the date of transition to state 2. Assume that $V(\tau) = 0$. Consider the following investment schedule $\tilde{\tau}$: if $x_{\tilde{S}} = 0$ then

$\tilde{\tau} = \tilde{S}$, $E[\tilde{S}] > E[S]$, otherwise, $\tilde{\tau} = \tau_2 \wedge T$, and $V(\tilde{\tau}) = 0$. Further, from (2.9), the agent's objective is given by

$$\begin{aligned} W(\tilde{\tau}) - W(\tau) &= (SV(\tilde{\tau}) - V(\tilde{\tau})) - (SV(\tau) - V(\tau)) \\ &= SV(\tilde{\tau}) - SV(\tau). \end{aligned}$$

The change from $\tau \geq \tau_2 \wedge T$ to $\tau = \tau_2 \wedge T$ induces no loss in total surplus as the measure of $\omega \in \Omega$ satisfying the event $\{\tau_2 \leq T\}$ is equal for both schedules. Further, the change from conditional stopping at S to conditional stopping at \tilde{S} induces an increase in total surplus as $P(x_{\tilde{S}} = 0) < P(x_S = 0)$ and thus, in the latter case, conditional stopping happens less frequently. Hence, $SV(\tilde{\tau}) \geq SV(\tau)$. Thus, for a schedule that assigns positive probability to stopping in states 0 and 2 to be optimal, it is necessary that stopping in state 2 happens at τ_2 with probability 1. \square

Proof of Proposition 2.5.3. Given Lemma 13, the space of candidate optimal investment schedules under $\kappa \in (\tilde{\kappa}(T, \lambda), \kappa^{FI}(T, \lambda)]$ simplifies to schedules such that stopping in state 2 happens at τ_2 , and also stopping in state 0 happens with positive probability. Thus, to characterize the optimal schedule under $\kappa \in (\tilde{\kappa}(T, \lambda), \kappa^{FI}(T, \lambda)]$, I need to characterize the assignment of the probability mass of stopping in state 0 that is optimal for the agent given the principal's individual rationality constraints. To do this, I consider the agent's optimal design of a *device that randomizes over the dates of stopping in state 0*.

At $t = 0$, the agent chooses a distribution F_ρ on $[0, T]$, observable to both the agent and the principal. ρ stands for the random date at which the stopping occurs if the state is 0 by that date. ρ is drawn at $t = 0$ according to F_ρ , which is independent from the state process x_t , and the draw privately observed by the agent.

I proceed with a useful lemma.

Lemma 15. *Given an investment schedule*

$$\tau = \begin{cases} \rho, & \text{if } x_\rho = 0 \\ \tau_2 \wedge T, & \text{otherwise,} \end{cases} \quad (2.35)$$

where $\rho \in [0, T]$, it holds that

$$P(x_\tau = 2) = P(x_T = 2) - P(x_\rho = 0)P(x_T = 2|x_\rho = 0)$$

and

$$E[\tau] = E[\tau_2 \wedge T] - P(x_\rho = 0)E[\tau_2 \wedge T - \rho|x_\rho = 0].$$

Proof. $P(x_\tau = 2)$ stands for the mass of events such that the principal gets v . Given (2.35), the principal gets v either if the second stage is completed not later than ρ or if the first stage is completed not later than ρ and the second stage is completed not later than T . Thus,

$$P(x_\tau = 2) = P(\{x_\rho = 1\} \cap \{\tau_2 \leq T\}) + P(x_\rho = 2).$$

Further,

$$P(\{x_\rho = 1\} \cap \{\tau_2 \leq T\}) = P(x_\rho = 1) P(\tau_2 \leq T | x_\rho = 1).$$

Thus,

$$P(x_\tau = 2) = P(x_\rho = 1) P(\tau_2 \leq T | x_\rho = 1) + P(x_\rho = 2). \quad (2.36)$$

Further, from the full probability formula,

$$\begin{aligned} P(x_\rho = 1) P(\tau_2 \leq T | x_\rho = 1) &= \\ &= P(x_T = 2) \\ &\quad - P(x_\rho = 0) P(\tau_2 \leq T | x_\rho = 0) \\ &\quad - P(x_\rho = 2) P(\tau_2 \leq T | x_\rho = 2). \end{aligned}$$

Plugging this into (2.36) yields

$$P(x_\tau = 2) = P(x_T = 2) - P(x_\rho = 0) P(\tau_2 \leq T | x_\rho = 0).$$

I proceed with proving the second result of Lemma 15. Given (2.35), it holds that

$$\begin{aligned} E[\tau] &= P(x_\rho = 0) E[\tau | x_\rho = 0] + P(x_\rho > 0) E[\tau | x_\rho > 0] \\ &= P(x_\rho = 0) \rho + P(x_\rho > 0) E[\tau_2 \wedge T | x_\rho > 0]. \end{aligned} \quad (2.37)$$

Further, from the full probability formula,

$$\begin{aligned} P(x_\rho > 0) E[\tau_2 \wedge T | x_\rho > 0] &= E[\tau_2 \wedge T] \\ &\quad - P(x_\rho = 0) E[\tau_2 \wedge T | x_\rho = 0]. \end{aligned}$$

Plugging this into (2.37) yields

$$E[\tau] = E[\tau_2 \wedge T] - P(x_\rho = 0) E[\tau_2 \wedge T - \rho | x_\rho = 0].$$

□

I proceed to solving the agent's problem:

$$\begin{aligned} &\max_{F_\rho} \left\{ E_{F_\rho} [cE[\tau]] \right\} \\ &\text{s.t. } E_{F_\rho} [V_t(\tau) | t < \tau] \geq 0, \forall t \geq 0, \end{aligned} \quad (2.38)$$

where τ is given by (2.35).

I proceed in two steps: first, I formulate and *solve the relaxed version of* (2.38) with individual rationality constraint only for $t = 0$; second, I demonstrate that the solution to the relaxed problem *satisfies the full system of constraints in* (2.38).

The individual rationality constraint in the relaxed problem is given by

$$P(x_\tau = 2) v - E[\tau] c \geq 0.$$

Using Lemma 15, the agent's relaxed problem can be written out as:

$$\begin{aligned} &\min_{F_\rho} \left\{ E_{F_\rho} [P(x_\rho = 0) E[\tau_2 \wedge T - \rho | x_\rho = 0]] \right\} \\ &\text{s.t. } E_{F_\rho} [P(x_\rho = 0) (cE[\tau_2 \wedge T - \rho | x_\rho = 0] - v P(\tau_2 \leq T | x_\rho = 0))] \geq -V(\tau_2). \end{aligned} \quad (2.39)$$

The Lagrangian function for the problem is

$$\begin{aligned} \mathcal{L} = & \mathbb{E}_{F_\rho} [\mathbb{P}(x_\rho = 0) \mathbb{E}[\tau_2 \wedge T - \rho | x_\rho = 0]] \\ & - \mu \left(\mathbb{E}_{F_\rho} [\mathbb{P}(x_\rho = 0) (c \mathbb{E}[\tau_2 \wedge T - \rho | x_\rho = 0] - v \mathbb{P}(\tau_2 \leq T | x_\rho = 0))] + V(\tau_2) \right), \end{aligned}$$

where $\mathbb{P}(x_\rho = 0) = e^{-\lambda\rho}$,

$$\begin{aligned} & \mathbb{E}[\tau_2 \wedge T - \rho | x_\rho = 0] \\ = & \mathbb{P}(\tau_2 \leq T | x_\rho = 0) \int_\rho^T z \cdot \frac{\lambda^2 (z - \rho) e^{-\lambda(z-\rho)}}{\mathbb{P}(\tau_2 \leq T | x_\rho = 0)} dz + \mathbb{P}(\tau_2 > T | x_\rho = 0) T - \rho \\ = & \frac{2}{\lambda} - \frac{2}{\lambda} e^{-\lambda(T-\rho)} - e^{-\lambda(T-\rho)} (T - \rho) \end{aligned} \tag{2.40}$$

and

$$\mathbb{P}(\tau_2 \leq T | x_\rho = 0) = 1 - e^{-\lambda(T-\rho)} - \lambda(T-\rho) e^{-\lambda(T-\rho)}. \tag{2.41}$$

I obtain the F.O.C., which needs to hold for each value of ρ that has a positive probability in F_ρ :

$$e^{-\lambda T} \left(c \left(2e^{-\lambda(T-\rho)} - 1 \right) (\mu - 1) - \mu \lambda v \left(e^{-\lambda(T-\rho)} - 1 \right) \right) = 0. \tag{2.42}$$

The derivative of the left-hand side of (2.42) w.r.t. ρ is given by $e^{-\lambda\rho} \lambda (2c + \mu(\lambda v - 2c))$. As $\kappa^{FI}(T, \lambda) < \frac{1}{2}$, the derivative is positive. Thus, there exists at most one ρ that satisfies the FOC (2.42). Thus, the optimal F_ρ is degenerate. I denote it with S_0^A , the interim deadline.

I proceed with characterizing the optimal S_0^A :

$$\begin{aligned} & \min_{S \in [0, T]} \{ \mathbb{P}(x_S = 0) \mathbb{E}[\tau_2 \wedge T - S | x_S = 0] \} \\ & \text{s.t. } \mathbb{P}(x_S = 0) \left(c \mathbb{E}[\tau_2 \wedge T - S | x_S = 0] - S V_{S|0}(\tau_2) \right) \geq -V(\tau_2). \end{aligned} \tag{2.43}$$

The system of F.O.C. is given by

$$\begin{cases} e^{-\lambda T} c \left(2e^{-\lambda(T-S)} - 1 \right) (\mu - 1) & \geq 0 \text{ if } S = 0 \\ - e^{-\lambda T} \mu \lambda v \left(e^{-\lambda(T-S)} - 1 \right) & = 0 \text{ if } S \in (0, T) \\ & \leq 0 \text{ if } S = T \\ \frac{c}{\lambda} e^{-\lambda T} \left(2 \left(e^{-\lambda(T-S)} - 1 \right) - \lambda(T-S) \right) & \\ - v e^{-\lambda T} \left(\left(e^{-\lambda(T-S)} - 1 \right) - \lambda(T-S) \right) + V(\tau_2) \geq 0 & = 0 \text{ if } \mu > 0. \end{cases}$$

Assume $\mu = 0$. In this case, the first F.O.C. wrt S yields $-c e^{-\lambda T} \left(2e^{-\lambda(T-S)} - 1 \right)$. The expression is negative for all $S \in (0, T)$. Thus, $\mu > 0$, and optimal S solves the binding constraint. Thus, I proceed with inspecting the corresponding equation given by

$$\begin{aligned} & \frac{c}{\lambda} e^{-\lambda T} \left(2 \left(e^{-\lambda(T-S)} - 1 \right) - \lambda(T-S) \right) \\ & - v e^{-\lambda T} \left(\left(e^{-\lambda(T-S)} - 1 \right) - \lambda(T-S) \right) \\ = & -V(\tau_2), \end{aligned} \tag{2.44}$$

where $V(\tau_2)$ is given by (2.26).

The solution to (2.44) is given by

$$S = \frac{1}{\lambda} [\gamma + \mathcal{W}(-\gamma e^{-\gamma})], \quad (2.45)$$

where $\gamma = e^{\lambda T \frac{1-2\kappa}{1-\kappa}}$ and $\mathcal{W}(\cdot)$ denotes the Lambert W function.

Denote the 0 and -1 branches of the Lambert W function by $\mathcal{W}_0(\cdot)$ and $\mathcal{W}_{-1}(\cdot)$. $\kappa \in (0, \frac{1}{2})$, thus, $\gamma > 0$. (2.45) depends on γ and for each $\gamma \neq 1$ corresponds to two points as the Lambert W function has two branches. The values of (2.45) as a function of γ are presented in Figure 2.7. They are given by

$$S = \begin{cases} \left(\frac{1}{\lambda} [\gamma + \mathcal{W}_{-1}(-\gamma e^{-\gamma})], 0 \right), & \text{if } \gamma < 1 \\ \left(0, \frac{1}{\lambda} [\gamma + \mathcal{W}_0(-\gamma e^{-\gamma})] \right), & \text{if } \gamma > 1 \\ 0, & \text{if } \gamma = 1. \end{cases}$$




Figure 2.7: Roots of equation (2.44) as a function of the parameter γ :
 root corresponding to branch 0 of the Lambert W function - **thick**;
 root corresponding to branch -1 of the Lambert W function - **dashed**.

γ is decreasing in κ , and $\gamma|_{\kappa=\kappa^{FI}} = 1$. As $\kappa \leq \kappa^{FI}$, which corresponds to $\gamma \geq 1$, the solution to (2.44) is given by

$$S_A = 0, \quad S_B = \frac{1}{\lambda} [\gamma + \mathcal{W}_0(-\gamma e^{-\gamma})].$$

As the objective of (2.43) is decreasing in S and $S_B > S_A$, the solution to (2.43) is given by

$$S_0^A = \frac{1}{\lambda} [\gamma + \mathcal{W}_0(-\gamma e^{-\gamma})], \quad \gamma = e^{\lambda T \frac{1-2\kappa}{1-\kappa}}. \quad (2.46)$$

Finally, I can describe the solution to (2.39): τ is the stopping time such that stopping occurs either at the moment of completion of the second stage of the project or at S_0^A , conditional on the absence of the completion of the first stage of the project, i.e.

$$\tau = \begin{cases} S_0^A, & \text{if } x_{S_0^A} = 0 \\ \tau_2 \wedge T, & \text{otherwise,} \end{cases} \quad (2.47)$$

where S_0^A is given by (2.46).

I proceed with the second part of the proof: I demonstrate that (2.47) satisfies the full system of constraints in (2.38), and thus solves (2.38). To do this, I need to demonstrate that $V_t(\tau) \geq 0$, for all $t \in [0, \tau]$. If the recommendation mechanism τ is given by (2.47), then, for $t < S_0^A$ the absence of stopping at some t reveals that $x_t \neq 2$. Thus,

$$q_1(t) = \frac{p_1(t)}{p_1(t) + p_0(t)} = \frac{\lambda t}{1 + \lambda t}, \forall t < S_0^A.$$

Hence, $\dot{q}_1(t) > 0$, for all $t < S_0^A$. Further, for $t \geq S_0^A$, the absence of stopping reveals that $x_t = 1$. Thus, $q_1(t) = 1$, for all $t \geq S_0^A$.

Writing out $\dot{V}_t(\tau)$ based on (2.13) yields

$$\dot{V}_t(\tau) = \lambda q_1(t) V_t(\tau) + v\lambda(\kappa - q_1(t)). \quad (2.48)$$

$q_1(0) = 0$ and $\dot{q}_1(t) > 0$, for all $t < S_0^A$. I define \tilde{t} as the solution of $\frac{\lambda t}{1 + \lambda t} = \kappa$. $q_1(t) < \kappa$, for all $t \in [0, \tilde{t} \wedge S_0^A]$.

I argue that $V(\tau) \geq 0 \Rightarrow V_t(\tau) \geq 0$, for all $t \in (0, \tilde{t} \wedge S_0^A)$. Assume that this is not true, then $\exists \hat{t}$ such that $\hat{t} \inf \{t \in (0, \tilde{t} \wedge S_0^A) : V_t(\tau) < 0\}$. As $V_t(\tau)$ is continuous in t , it follows that $V_{\hat{t}}(\tau) = 0$, and by the mean value theorem there must be $\bar{t} \in (0, \hat{t})$ such that $\dot{V}_{\bar{t}}(\tau) \leq 0$. But this is in contradiction with the fact that $V_{\bar{t}}(\tau) \geq 0$ and 2.48.

Consider now $t \in [\tilde{t} \wedge S_0^A, \tau)$. The continuation value can be written as

$$V_t(\tau) = \mathbb{E} \left[\int_t^\tau v\lambda(q_1(z) - \kappa) dz \mid t < \tau \right]. \quad (2.49)$$

As $\kappa < \frac{1}{2}$ and $q_1(t) = 1$, for all $t \in [S_0^A, \tau)$, it holds that $q_1(t) \geq \kappa$, $\forall t \in [\tilde{t} \wedge S_0^A, \tau)$. Thus, it can be seen from (2.49) that $V_t(\tau) \geq 0$, $\forall t \in [\tilde{t} \wedge S_0^A, \tau)$. \square

Proof of Proposition 2.6. I assume it is not the case that $\alpha = 1$ and $\beta = 0$ as, otherwise, agent is indifferent and discloses no information. I start with proving existence of $\tilde{\kappa}$ and then proceed to proving that when the project is promising, an investment schedule, in which stopping never occurs in state 0, is optimal. Proving existence of $\tilde{\kappa}$ follows the steps of the proof of Lemma 10. The principal's expected payoff is given by

$$V(\tau) = \alpha \mathbb{P}(x_\tau = 2) v \mathbb{E} \left[e^{-r\tau} \mid \tau_2 \leq \tau \right] - \mathbb{E} \left[\int_0^\tau e^{-rs} ds \right] c.$$

$\tilde{\kappa}$ solves $V(\tau_2) = 0$, or, equivalently

$$\alpha \mathbb{P}(x_{\tau_2 \wedge T} = 2) v \mathbb{E} \left[e^{-r \cdot \tau_2 \wedge T} \mid \tau_2 \leq T \right] = \mathbb{E} \left[\int_0^{\tau_2 \wedge T} e^{-rs} ds \right] c, \quad (2.50)$$

where $\mathbb{P}(x_{\tau_2 \wedge T} = 2) = p_2(T)$. Solving (2.50) for κ yields

$$\tilde{\kappa}(T, \lambda, r, \alpha) = \frac{1}{\lambda \alpha} \frac{\mathbb{P}(x_{\tau_2 \wedge T} = 2) \mathbb{E} \left[e^{-r \cdot \tau_2 \wedge T} \mid \tau_2 \leq T \right]}{\mathbb{E} \left[\int_0^{\tau_2 \wedge T} e^{-rs} ds \right]}.$$

Finally, $V(\tau)$ decreases in κ . Thus, if $\kappa < \tilde{\kappa}(T, \lambda, r, \alpha)$, then an investment schedule $\tau = \tau_2 \wedge T$ satisfies the principal's individual rationality constraint.

Consider now the agent's expected payoff $W(\tau)$ given by

$$W(\tau) = (1 - \alpha) P(x_\tau = 2) v E[e^{-r\tau} | \tau_2 \leq \tau] + E\left[\int_0^\tau e^{-rs} ds\right] \beta c.$$

Consider the case $\kappa \leq \tilde{\kappa}(T, \lambda, r, \alpha)$. Consider an investment schedule τ given by (2.35), i.e., such that stopping happens either immediately at the moment of the second stage completion, or in state 0 at a possibly random interim deadline. Further, consider an alternative investment schedule $\hat{\tau} = \tau_2 \wedge T$. Given the two investment schedules, $P(x_{\hat{\tau}} = 2) > P(x_\tau = 2)$. Further, $E[e^{-r\hat{\tau}} | \tau_2 \leq \hat{\tau}] = E[e^{-r\tau} | \tau_2 \leq \tau]$ and $E\left[\int_0^{\hat{\tau}} e^{-rs} ds\right] > E\left[\int_0^\tau e^{-rs} ds\right]$. As $W(\hat{\tau}) > W(\tau)$ and $\kappa < \tilde{\kappa}(T, \lambda, r, \alpha)$, the agent prefers to implement an investment schedule $\hat{\tau}$ rather than τ .

Consider now the case $\kappa > \tilde{\kappa}(T, \lambda, r, \alpha)$. The application of the arguments from the proof of Lemma 13 establishes the result. \square

2.D Disclosure of project completion with a deterministic delay

Assume $\kappa \in (0, \kappa^{NI}(T, \lambda)]$ and $T > \bar{S}^{NI}$. The optimal mechanism provides no information until $t = \bar{S}^{NI}$. At each $t \geq \bar{S}^{NI}$, it generates a recommendation to stop iff the second stage of the project was completed at date $\pi(t)$ in the past, where

$$\pi(t) = -\frac{1}{\lambda} \left(1 + \frac{1}{\lambda} \mathcal{W}_{-1}\left(-\frac{1}{\kappa} e^{-1-\lambda t}\right) \right),$$

where $\mathcal{W}_{-1}(\cdot)$ denotes the -1 branch of Lambert W function.

The mechanism from Proposition 2.D does not recommend stopping until the second stage of the project is completed, and thus maximizes the total surplus. The mechanism makes the principal's individual rationality constraint bind, $V_{\bar{S}^{NI}}(\tau) = 0$. The absence of a stopping recommendation after $t = \bar{S}^{NI}$ induces posterior beliefs $q_1(t) = \kappa, \forall t \geq \bar{S}^{NI}$. Note that the principal's expected instantaneous payoff within Δt is given by

$$v \cdot q_1(t) \lambda \Delta t - c \cdot \Delta t = v \lambda \Delta t (q_1(t) - \kappa).$$

No information is provided until \bar{S}^{NI} and after \bar{S}^{NI} the mechanism keeps the principal's expected instantaneous payoff precisely at 0, $\forall t \geq \bar{S}^{NI}$. As a result, the principal's continuation value is kept at 0 for all $t \in [\bar{S}^{NI}, \tau)$.

The delay is given by $t - \pi(t)$. At the beginning of the disclosure, $t = \bar{S}^{NI}$, the delay is \bar{S}^{NI} . To keep the belief regarding state 1 constant, the delay decreases for all $t \in (\bar{S}^{NI}, \tau)$.

Proof of Proposition 2.D. Posterior beliefs at date π induced by the disclosure of the absence of second stage completion are given by

$$\begin{aligned} q_0(\pi) &= \frac{p_0(\pi)}{p_0(\pi) + p_1(\pi)}, \\ q_1(\pi) &= \frac{p_1(\pi)}{p_0(\pi) + p_1(\pi)}. \end{aligned}$$

As no other evidence is provided during $(\pi, t]$, the beliefs evolve according to

$$\begin{aligned} q_0(s) &= \frac{e^{-\lambda s}}{1 + \lambda\pi}, \\ q_1(s) &= \frac{e^{-\lambda s} \lambda (s + \pi)}{1 + \lambda\pi}, \end{aligned}$$

where $s \geq \pi$.

The belief regarding state 1 at current date t is given by

$$q_1(t) = \frac{e^{-\lambda(t-\pi)} \lambda t}{1 + \lambda\pi}. \quad (2.51)$$

The dynamic of the state is the same as in the no-information benchmark until $t = \bar{S}^{NI}$. Therefore,

$$q_0(\bar{S}^{NI}) = p_0(\bar{S}^{NI}) = \frac{\kappa}{\lambda \bar{S}^{NI}} \quad \text{and} \quad q_1(\bar{S}^{NI}) = p_1(\bar{S}^{NI}) = \kappa.$$

The dynamics for $t \geq \bar{S}^{NI}$ then is $q_1(t) = \kappa$, $\dot{q}_1(t) = 0$. Solving from (2.51),

$$\pi = -\frac{1}{\lambda} \left(1 + \frac{1}{\lambda} \mathcal{W}_{-1} \left(-\frac{1}{\kappa} e^{-1-\lambda t} \lambda t \right) \right).$$

The recommendation mechanism τ is obedient. $\tau \geq \tau_2 \wedge T$ implies that the recommendation to stop comes only if the second stage of the project has already been completed, and thus immediate stopping is clearly optimal for the principal. The recommendation not to stop is also obedient. $V_t(\tau) \geq 0, \forall t \in [0, \bar{S}^{NI}]$ is formally demonstrated in the proof of obedience for Proposition 2.5.2. I proceed by showing that $V_t(\tau) = 0, \forall t \in [\bar{S}^{NI}, \tau)$. Writing out $V_t(\tau)$ in the recursive form yields

$$\begin{aligned} V_t(\tau) &= (v\lambda q_1(t) - c) \Delta t + (1 - \lambda q_1(t) \Delta t) V_{t+\Delta t}(\tau) \\ &= v\lambda (q_1(t) - \kappa) \Delta t + (1 - \lambda q_1(t) \Delta t) V_{t+\Delta t}(\tau). \end{aligned}$$

As $q_1(t) = \kappa, \forall t \in [\bar{S}^{NI}, \tau)$, it becomes

$$V_t(\tau) = (1 - \lambda q_1(t) \Delta t) V_{t+\Delta t}(\tau), \forall t \in [\bar{S}^{NI}, \tau).$$

Differentiating both sides w.r.t. Δt yields

$$0 = -\lambda q_1(t) V_{t+\Delta t}(\tau) + \dot{V}_{t+\Delta t}(\tau).$$

This differential equation together with the boundary condition $V_T(\tau) = 0$ has a unique solution $V_t(\tau) = 0$ for all $t \in [\bar{S}^{NI}, T]$. □

2.E The case of no project completion deadline

Importantly, the presence of a hard project deadline T serves as one of the necessary and sufficient conditions for the agent to commit to an interim reporting deadline. Without a hard deadline T , the principal's incentives under full information are different. Recall from Lemma 8 the principal's incentive to continue investing decreases in the length of absence of the first stage completion. In the case $T \rightarrow \infty$, the continuation value $V_{t|1}^{FI}$ is constant and given by $v(1 - \kappa)$. As a result, the principal's incentive to continue investing given the absence of stage completion does not change over time. Thus, if the principal opts in, he never chooses to stop investing before the completion of the second stage occurs. As a result, setting an interim deadline stops serving as an agent's tool to incentivize the principal's investment. The agent's information policy in the case of no project deadline is given in Lemma 16.

Lemma 16. *Assume that $T \rightarrow \infty$. In that case, if $\kappa < \frac{1}{2}$, then the agent uses the information policy presented in Proposition 2.5.1, Case 2.*

Proof of Lemma 16. Under full information and the absence of an exogenous deadline, the principal assigns value v_x to each state $x \in \{0, 1, 2\}$. Clearly, $v_2 = v$ as the principal stops immediately and gets v . In state 1, at each t the principal gets $v\Delta t$ with probability $\lambda\Delta t$ and pays $c\Delta t$. As a result, the principal's continuation value is constant. Assume that $\kappa < 1$, as otherwise $c \geq \lambda v$ and the principal chooses not to invest in state 1. As the principal's continuation value in state 1 does not change over time,

$$0 = \lambda \cdot (v_2 - v_1) - c,$$

and so

$$v_1 = v - \frac{c}{\lambda} = v(1 - \kappa).$$

Thus, the principal wants to invest in state 0 if $c \leq \lambda v_1$, i.e., $\kappa \leq \frac{1}{2}$.

Finally, as the information regarding τ_1 is not decision-relevant for the principal, for $\kappa < \frac{1}{2}$, the agent chooses the information policy that discloses only the completion of the second stage of the project and optimally postpones the disclosure to make the principal's individual rationality constraint bind. □

3. Form of Preference Misalignment Linked to State-Pooling Structure in Bayesian Persuasion

Co-authored with Rastislav Rehak (CERGE-EI).

3.1 Introduction

Bayesian persuasion, pioneered by Kamenica and Gentzkow [2011], studies strategic disclosure of information when the sender controls the information environment (called *signal*) and the receiver controls the choice of action to be taken. As a review by Kamenica [2019] suggests, this literature has provided many extensions of the original model of Kamenica and Gentzkow [2011] with interesting qualitative insights. However, full characterization of the optimal signal is generally difficult even in the original model. There has been little progress on this front, and it has been limited to a small number of special cases.¹

We contribute to this literature by studying a special case of the original model that has received little attention – a Bayesian persuasion model in which both the sender and the receiver have *state-dependent preferred actions*. We characterize a qualitative property of the optimal signal called *state-pooling structure*, which describes pools of states that cannot be discerned from one another by the optimal signal. Specifically, we ask how the structure of state-dependent preference misalignment affects the state-pooling structure of the optimal signal.

To illustrate the main point of this paper, we present an example of a politician (receiver, he) and his advisor (sender, she). They both wish to implement some level of government spending $a \in \mathbb{R}$ that is adapted to the current economic situation captured by GDP per capita y , which takes one of three possible values: 1, 2, or 3. However, they each have a different vision of optimal spending as a function of GDP per capita. The advisor's payoff is $u_A(a, y) = -(a - \omega_A(y))^2$ and the politician's payoff is $u_P(a, y) = -(a - \omega_P(y))^2$, where $\omega_A(y)$ and $\omega_P(y)$ represent the preferred spending of the advisor and the politician in state y , respectively. The advisor designs an investigation (a signal) that can inform the politician about the realization of GDP per capita. She does that strategically to influence the spending choice of the politician. We are interested in how the structure of this signal depends on the form of misalignment between the advisor's and politician's preferences captured by ω_A and ω_P , respectively.

Figure 3.1 illustrates how the form of disagreement between the advisor's and politician's preferred spending influences the structure of the optimal signal.² In the case presented *in the left plot*, the advisor's optimal signal fully reveals whether the state of the economy is low or not, i.e., one of the two outcomes of

¹We return to this point in the discussion of related literature in Section 3.2.

²The structures of the optimal signals for the two cases considered in Figure 3.1 are derived using results from Section 3.6.

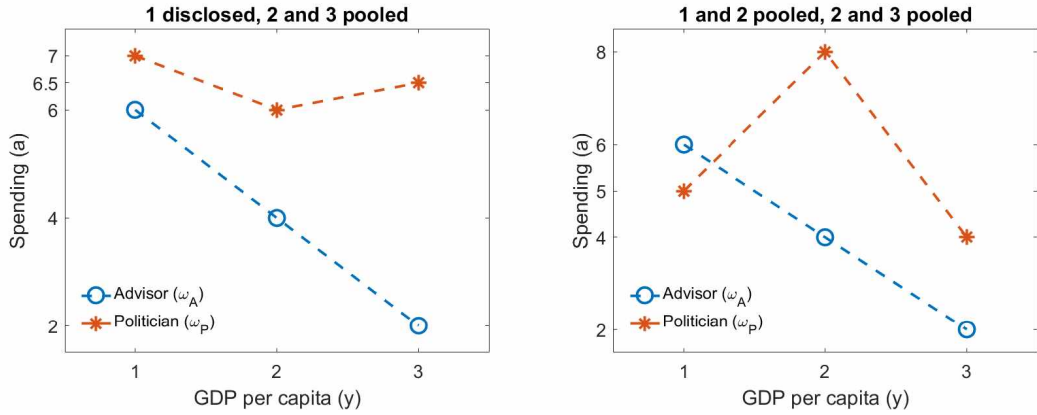


Figure 3.1: The form of disagreement between the advisor (ω_A) and the politician (ω_P) matters for the structure of the optimal signal:

left plot: the advisor fully reveals state 1 and pools states 2 and 3 together;

right plot: the advisor pools states 1 and 2 together and states 2 and 3 together

(note: we consider only three levels of GDP per capita; the lines are drawn only for clarity of the picture)

her investigation fully reveals the low state and the other leaves the politician uncertain about the high and middle states – we say they are pooled together. Intuitively, both the advisor and the politician want the highest spending in the low state, so their goals are aligned in this state and the advisor wants to reveal it perfectly. However, they disagree about whether the spending should be higher in the middle or high state, so the advisor wants to attenuate this disagreement by pooling these two states together. In the case presented *in the right plot*, the advisor’s optimal signal reveals whether the economy is above or below average, i.e., one of the two outcomes of her investigation pools the low and middle states, while the other pools the middle and high states. Intuitively, the advisor and the politician disagree about whether the spending should be higher in the low or middle state, so the advisor wants to attenuate this disagreement by pooling these two states together. However, they both agree that the spending should be higher in the middle state than in the high state, but the politician prefers a greater spending difference between these two states than the advisor. Therefore, the advisor wants to moderate the politician’s actions by pooling these two states together.

In Section 3.3, we describe our model. We use the Bayesian persuasion framework of Kamenica and Gentzkow [2011] with one-dimensional finite state space – the sender’s preferred action. Both the sender and the receiver have quadratic loss functions with bliss points depending on the state of the world. The structure of misalignment is captured by function ρ mapping the state of the world (the sender’s preferred action) to the receiver’s preferred action. The case of linear ρ with slope 1 corresponds to the benchmark of perfect alignment.³ We do not impose any requirements on this function and we analyze the role of its shape for the qualitative structure of the optimal signal in terms of state pooling.

³A state-independent intercept does not affect the choice of the signal because it is a “sunk cost” for the sender.

In Section 3.4, we present general results on the pooling structure of the optimal signal. The patterns of pooling are driven by the sender’s trade-off between (i) the informativeness of the signal, which leads to better adaptation of the action to the state of the world in states of alignment, and (ii) the revelation of the realized mismatch of the sender’s and receiver’s preferred actions, which drives the action of the receiver away from the sender’s preferred action. First, we show that the sender generically benefits from revealing some information. The only cases in which non-disclosure is optimal are when ρ is linear with a slope sufficiently different from 1. Second, we demonstrate that the optimal signal does not induce an interior belief (except in cases of non-disclosure).

In Section 3.5, we propose a simple graph procedure to characterize the optimal structure of state pooling for a given ρ . This procedure consists of an analysis of ρ on pairs of states and a test of pooling of more than two states. The crucial element of this procedure is the slope of ρ between pairs of states, which plays the role of an index of misalignment – if it is too high (disagreement about magnitude) or lower than zero (disagreement about order), then it indicates space for pooling; otherwise, it indicates space for separation.

In Section 3.6, we provide a full characterization of the state-pooling structure in the case of three states of the world. The state-pooling structure is completely pinned down by the shape of ρ except for the case in which ρ has a slope sufficiently different from 1 for each of the three pairs of states. In that case, the choice of a particular state-pooling structure depends both on the shape of ρ and the prior.

3.2 Related literature

First, we relate our work to the *Bayesian persuasion* literature. The most relevant results from the seminal paper by Kamenica and Gentzkow [2011] are (i) conditions for full disclosure or non-disclosure in the general form and (ii) comparative statics of more aligned preferences. Regarding point (i), we go beyond these two “corner” cases for the optimal signal, similarly as in the recent studies of Arieli et al. [2020] and Kolotilin and Wolitzky [2020]. We discuss the connection of our work to Kolotilin and Wolitzky [2020] in more detail later in this section. Regarding point (ii), we perform a different exercise with preference misalignment: we fix the preferences and analyze how the structure of preference misalignment is related to the structure of state pooling of the optimal signal.

The methodological progress in Bayesian persuasion on the front of providing a general characterization of the structure of the optimal signal has been scarce. First, with two or three states of the world, concavification provides an insightful graphical method of solving the sender’s problem [Kamenica and Gentzkow, 2011]. Second, when the sender’s utility depends only on the expected state, the “Rothschild-Stiglitz approach” [Gentzkow and Kamenica, 2016] and linear programming methods [Kolotilin, 2018, Dworzak and Martini, 2019] have been used to solve these problems. However, we are interested in situations with the *sender’s state-dependent preferred* action and the role of the structure of preference misalignment, where these methods do not deliver immediate answers. We propose a new concavification-based approach of characterizing the state-pooling structure of the optimal signal.

The closest paper to ours is Kolotilin and Wolitzky [2020]. However, we differ along several directions, and our paper can be viewed as complementary to theirs. First, their sender prefers higher actions independently of the state, but experiences state-dependent loss from mismatching the preferred action. In contrast, our sender has state-dependent preferred actions, but her loss from mismatching the preferred action is state-independent. Second, their receiver prefers higher actions in higher states; we do not impose this assumption. Third, they provide sufficient (and “almost necessary”) conditions for special patterns of “assortative” disclosure. However, they do not provide a procedure for finding the pooling structure of the optimal signal explicitly, and they avoid characterization of more complicated patterns. In contrast, we work in a more specialized quadratic setting and do not restrict ourselves to characterization of specific (pairwise) pooling structures. Instead, we propose a general procedure for finding the pooling structure. Finally, the mechanisms driving the results in the two papers are different: in Kolotilin and Wolitzky [2020], the information does not have value for the sender alone, so state pooling emerges from pure persuasion concerns, while state pooling in our model is driven by the interplay of the sender’s incentives to disclose the state and to hide misalignment.

Two other related papers in Bayesian persuasion literature are Alonso and Câmara [2016] and Galperti [2019]. Similar to our paper, both rely on the concavification technique to obtain insights regarding the optimal signal. Alonso and Câmara [2016] consider the standard Bayesian persuasion model, but assume that the sender and the receiver have heterogeneous prior beliefs. While the sender in Alonso and Câmara [2016] uses the variation of the difference between the sender’s and receiver’s prior beliefs across the states of the world to design the optimal disclosure, our sender uses the variation in the misalignment of the sender’s and receiver’s bliss points across the states of the world.⁴ Galperti [2019] considers the standard Bayesian persuasion model in which the sender and the receiver have a special type of heterogeneous prior beliefs: the receiver attaches zero probability to some states that are perceived with positive probability by the sender. While we restrict attention to a sender with state-dependent bliss actions and study the general patterns of state pooling, Galperti [2019] makes weaker assumptions about preferences and focuses on patterns of pooling of the states that have a priori zero probability for the receiver.

Second, the results of our study are connected to the *literature on persuasion games*, in which the sender chooses how to disclose her private verifiable information regarding the state of the world. Milgrom [1981] and Milgrom and Roberts [1986] analyze the conventional model of a persuasion game and establish the result on “unraveling” of the sender’s private information leading to full disclosure. Dye [1985] and Shin [1994] study state pooling in a similar game but with (second-order) uncertainty of the receiver about whether the sender actually has some private information or not. Seidmann and Winter [1997] analyze a persuasion game in which the sender has state-dependent preferred actions, and they demonstrate that the “unraveling” result still holds. The combination of these two features – second-order uncertainty and state-dependent preferred actions –

⁴They demonstrate that, under some mild conditions on the sender’s and receiver’s preferences, the sender generically chooses at least partial disclosure over non-disclosure. Similarly, in our model, the non-disclosure conditions are stringent.

has been studied in a small number of recent papers. The closest paper to ours is Hummel et al. [2018], in which unraveling does not occur due to the presence of the receiver’s second-order uncertainty. In the Bayesian persuasion model that we study, the sender’s disclosure mechanism serves a similar role to the one in Hummel et al. [2018]: the sender moderates the receiver’s actions via pooling of the states for which the sender’s bliss-point line is sufficiently flat relative to that of the receiver.

Finally, Miura [2018] studies how pooling equilibria can be characterized based on a procedure that uses a *masquerade graph* introduced in Hagenbach et al. [2014]. In his procedure, a pool of states is formed by the types of the sender who are mutually interested in masquerading, i.e., being perceived by the receiver as some other type in the pool. In spirit, this resembles the procedure for discovery of the state-pooling structure we introduce: a masquerade edge between two nodes (types) in Miura’s graph procedure plays a similar role as an edge between two nodes (states) in our graph procedure – it captures a motive for manipulative non-disclosure.

3.3 Model

We consider the standard Bayesian persuasion framework: a sender (S , she) designs and commits to an information structure (a Blackwell experiment) about an unknown state of the world $\omega \in \Omega$ to influence the action $a \in A$ of a receiver (R , he). The state space is finite, $\Omega \subset \mathbb{R}$, $|\Omega| = n$, and the action space is continuous, $A = \mathbb{R}$. The sender and the receiver have a common prior $p_0 \in \Delta(\Omega)$. They have the following preferences:

$$\begin{aligned} u_S &= -(a - \omega)^2, \\ u_R &= -(a - \rho(\omega))^2, \end{aligned}$$

where $\rho : \Omega \rightarrow \mathbb{R}$ is arbitrary. Hence, state ω represents the preferred action of the sender and $\rho(\omega)$ the preferred action of the receiver.⁵

As is standard, the sender can be seen equivalently as choosing a Bayes-plausible distribution over posteriors, which we refer to as *signal*: $\pi \in \Delta(\Delta(\Omega))$ such that

$$\sum_{p \in \text{supp}(\pi)} \pi(p) p(\omega) = p_0(\omega) \quad \forall \omega \in \Omega. \quad (3.1)$$

The timing is as follows: the sender chooses a signal π , a posterior belief p is drawn according to π , and the receiver takes an action a given the belief p . The solution concept is subgame perfect equilibrium. Going backwards, the receiver’s optimal action given a posterior belief p is $a(p) = E_p[\rho(\omega)]$. Hence, the game

⁵This model can be seen as a reduced form of a model in which the state of the world is two-dimensional, $y = (\omega_S, \omega_R)$, and the sender can design the experiment only about the dimension that is relevant for her, ω_S . The receiver then forms expectations about his relevant dimension, ω_R , using a common prior $p_0 \in \Delta(\Omega^2)$, so $\rho(\omega_S) = E_{p_0}[\omega_R | \omega_S]$. This formulation maps better to the example with a politician and his advisor presented in the Introduction.

⁶Kamenica and Gentzkow [2011] show that there exists an optimal π such that $|\text{supp}(\pi)| \leq \min\{|\Omega|, |A|\}$. Hence, we restrict our search for the optimal signal only to signals satisfying $|\text{supp}(\pi)| \leq n$.

reduces to the following problem of the sender:

$$\max_{\pi \in \Delta(\Delta(\Omega))} -E_{\pi} \left[E_p \left[(E_p [\rho(\omega)] - \omega)^2 \right] \right] \text{ s.t. } \sum_{p \in \text{supp}(\pi)} \pi(p)p = p_0, \quad (3.2)$$

where $E_{\pi}[\cdot]$ is the expectation over posteriors with respect to π and $E_p[\cdot]$ is the expectation over states with respect to p .

3.4 General results about the optimal signal

In this section, we present general results about the optimal signal, and combine them in the next section to construct the procedure that allows us to discover which states are “pooled” together in the optimal signal.

To better understand how the sender chooses the signal, we start by inspecting the trade-off she faces. We can rewrite the objective function from her problem (3.2) as

$$\text{var}_{\pi} (E_p [\omega]) - E_{\pi} \left[(E_p [\omega - \rho(\omega)])^2 \right]. \quad (3.3)$$

The first term captures the benefit of a more informative (in the sense of Blackwell) π – ideally, she would like to reveal all states perfectly.⁷ The second term captures the “cost” of revealed misalignment – ideally, she would like to “pool” some states to hide the largest misalignment. Hence, the sender prefers to reveal the most information so that the action is well adapted to the state. However, since she does not control the action directly, she wants to exploit the form of misalignment captured by ρ to manipulate the action of the receiver.

We can notice that the intercept of ρ does not play a role for the optimal signal. Formally, consider any function ρ and take $\rho' = b + \rho$ for some arbitrary constant $b \in \mathbb{R}$. The sender’s objective function

$$\text{var}_{\pi} (E_p [\omega]) - E_{\pi} \left[(E_p [\omega - \rho'(\omega)])^2 \right] \quad (3.4)$$

can be rewritten in the form

$$\text{var}_{\pi} (E_p [\omega]) - E_{\pi} \left[(E_p [\omega - \rho(\omega)])^2 \right] - 2bE_{p_0} [\omega - \rho(\omega)] + b^2. \quad (3.5)$$

The last two terms in (3.5) do not depend on π , so the optimal signals under ρ and ρ' coincide. Hence, a state-independent bias b (no matter how large) does not affect the optimal signal.⁸ Intuitively, the state-independent bias acts as a sunk cost for the sender. She cannot hide it by any manipulation of the signal because it is perfectly known ex ante.

It follows from the irrelevance of the intercept of ρ that what matters for the optimal signal is the overall shape of ρ , not agreement in particular states. In particular, perfect agreement between the sender and the receiver about the preferred action in a state of the world does not suffice for disclosure of that state. For example, consider two states $\omega_1 < \omega_2$, $\rho(\omega_1) = \omega_1$, $\rho(\omega_2) = 2\omega_1 - \omega_2$. Even

⁷To illustrate this point, imagine an interior prior p_0 , a signal π^1 with only interior beliefs, and a signal π^2 similar to π^1 , but with more extreme beliefs: $p_k^2 = p_k^1 + \varepsilon(p_k^1 - p_0) \forall k$, for some small enough $\varepsilon > 0$. Then, $\text{var}_{\pi^2} (E_p [\omega]) = (1 + \varepsilon)^2 \text{var}_{\pi^1} (E_p [\omega]) > \text{var}_{\pi^1} (E_p [\omega])$.

⁸We can contrast this feature with cheap talk [Crawford and Sobel, 1982a] in which the value of b matters for the informativeness of the equilibrium communication.

though the sender and the receiver perfectly agree about the preferred action in ω_1 , they substantially disagree in ω_2 . It will be evident from the results in this section that full disclosure of the “perfect-agreement state” ω_1 is not optimal. Intuitively, due to the Bayesian consistency constraint, full disclosure of ω_1 would limit the opportunity to moderate the substantial disagreement in ω_2 .⁹

3.4.1 Characterization of non-disclosure

In this subsection, we characterize the situation in which the sender does not benefit from revealing any information to the receiver.

The sender never (i.e., for any prior) benefits from providing any information if and only if ρ is linear with the slope from $(-\infty, 0] \cup [2, +\infty)$.

Proof. The proof is in Appendix 3.A. It identifies the conditions for concavity of the expected utility of the sender as a function of the induced posterior by the principal-minor test of the Hessian matrix of this function. \square

Surprisingly, it is relatively easy to introduce some information revelation in our setting: it is sufficient to have a nonlinearity in ρ . The intuition for this generic taste for information revelation is that information has high value for the sender who wants to match the state of the world. The cases of optimal non-disclosure identified in Proposition 3.4.1 are intuitive too: (i) *misalignment in order*, i.e., when the sender and the receiver disagree about the order of the bliss actions (slope of ρ negative) or (ii) *misalignment in magnitude*, i.e., when they agree about the order, but the receiver overreacts relative to the sender (slope of ρ greater than two).

The non-disclosure characterized in Proposition 3.4.1 is never uniquely optimal for $n \geq 3$. To resolve such cases of indifference, we make the following assumption.

Assumption 3.1. Under indifference, the sender chooses not to disclose the states.

This assumption can be justified by the sender’s interest in saving effort on communication when it is not needed. Technically, it greatly simplifies the analysis. Substantively, it leads us to identify the least informative signal in the indifference set of the sender. In Appendix 3.B, we analyze the structure of our problem that gives rise to the cases of indifference, and discuss the role of Assumption 3.1 as opposed to other selection criteria.

3.4.2 Full disclosure

In the next proposition, we provide a sufficient condition for full disclosure of the state of the world.

If ρ is linear with a slope in $[0, 2]$, full revelation of the state is always optimal (i.e., for any prior).

⁹In fact, Proposition 3.4.1 will imply that it is optimal not to disclose anything in this example.

Proof. The proof is in Appendix 3.A. It mostly follows from the proof of Proposition 3.4.1. \square

For general n , Proposition 3.4.2 provides only a sufficient condition for full disclosure, but for $n = 2$ we can provide a full characterization. This special case is a cornerstone of our analysis of the case with general n .

Lemma 17. *For $n = 2$, the sender strictly prefers full revelation if and only if the slope of ρ is in $(0, 2)$. The sender is indifferent between any feasible signals if and only if the slope of ρ is either zero or two. The sender strictly prefers no revelation if and only if the slope of ρ is in $(-\infty, 0) \cup (2, \infty)$.*

Proof. The proof is in Appendix 3.A. \square

3.4.3 “Extremization” – non-existence of an interior posterior

After analyzing the conditions for extreme signals (non-disclosure and full disclosure), we look at more structured signals. The following proposition provides the key result enabling that analysis.

[Extremization] Suppose non-disclosure is not optimal. Then, it is never optimal to induce an interior posterior.

Proof. The proof is in Appendix 3.A. It is constructed by contradiction with the optimality of the signal, based on an improvement by splitting one of its posteriors. We call this result “extremization” because it leads us from the interior of the simplex to its extreme (boundary) subsimplexes. \square

We can apply Proposition 3.4.3 iteratively to eliminate the areas of posteriors that will not appear in the optimal signal. This sharpens the idea about the structure of the optimal signal, which is our main interest, and simplifies the search for it. We use this idea in the next section.

3.5 State-pooling structure of the optimal signal

In this section, we go beyond the extreme cases of full disclosure and non-disclosure and study how preference misalignment, captured by ρ , affects a qualitative property of the optimal signal that we call state pooling. We define the state-pooling structure of a signal and present an illustrative procedure for its discovery that builds on the general results from Section 3.4.

3.5.1 Definitions

We say that states $\omega_{k_1}, \dots, \omega_{k_m}$, for some $k_1, \dots, k_m \in \{1, \dots, n\}$, are *pooled* together (or form a *pool of states*) under signal π if the set $M = \{\omega_{k_1}, \dots, \omega_{k_m}\}$ satisfies

$$\exists p \in \text{supp}(\pi) : \text{supp}(p) = M \ \& \ \forall p' \in \text{supp}(\pi) \text{ s.t. } p' \neq p : M \text{supp}(p'), \quad (3.6)$$

where $\text{supp}(\cdot)$ denotes support.¹⁰ The set of all pools of states that signal π induces is called the *state-pooling structure* of signal π .

The state-pooling structure of a signal can be captured *graphically* by representing each state of the world by a node and each pool by highlighting the corresponding set of nodes; an example is presented in Figure 3.2.




Figure 3.2: Example of a graphical representation of the state-pooling structure when $n = 4$ and the signal induces posteriors supported on $\{\omega_1\}$ and $\{\omega_2, \omega_3, \omega_4\}$

In the next subsection, we propose a procedure that aims to find the state-pooling structure of the optimal signal for a given form of preference misalignment captured by ρ . This procedure can easily be represented graphically; its desired output is a graphical representation of the state-pooling structure of the type depicted in Figure 3.2, i.e., nodes representing states and highlighted pools. However, the proposed procedure may not identify the state-pooling structure of the optimal signal completely in some cases, but may offer only candidates for optimal pools. Nevertheless, we can often identify which of the candidate pools are certainly a part of the optimal state-pooling structure. Hence, we introduce two types of highlighting in the procedure – *dashed* (highlighting candidate pools) and *full* (highlighting pools certainly belonging to the optimal state-pooling structure). Naturally, highlighting in full is superior to highlighting in dashed because it expresses certainty.

An important working component of the graphical procedure is the *edges* between pairs of nodes – they represent a pooling tendency of the corresponding states. We will see that this pooling tendency is driven by the slope of ρ between pairs of corresponding states; we denote the *slope of ρ between states ω_i and ω_j* by

$$s_{ij} = \frac{\rho(\omega_j) - \rho(\omega_i)}{\omega_j - \omega_i}. \quad (3.7)$$

This object represents an index of misalignment between the receiver (the numerator) and the sender (the denominator).¹¹

A subroutine of our procedure relates to the well-known problem from computer science called the *clique problem*. Thus, we borrow a few notions from graph theory. Let $G = (V, E)$ be an undirected graph (with V denoting the set of nodes and E denoting the set of edges). We call a subset of nodes $C \subseteq V$ *clique* if the subgraph of G induced by C is complete (i.e., the nodes in C are fully connected). A *clique* C is called *maximal* if there does not exist another clique strictly above C (in the sense of inclusion). The version of the clique problem

¹⁰In intuitive terms, $\omega_{k_1}, \dots, \omega_{k_m}$ are pooled together under signal π if π reveals whether the event $\{\omega_{k_1}, \dots, \omega_{k_m}\}$ occurred.

¹¹A similar object plays an important role for the pooling structure (of types) in Hummel et al. [2018].

that we are interested in is finding all maximal cliques in an undirected graph. Systematic inspection of all subsets of nodes or the *Bron–Kerbosch algorithm* can be used to solve this problem.

3.5.2 Procedure for discovery of the state-pooling structure of the optimal signal

We present a procedure that inspects the form of misalignment function ρ and reflects its implications for the state-pooling structure of the optimal signal on a graph. The output are pools highlighted in full (which are certainly present in the state-pooling structure of the optimal signal) and candidate pools highlighted in dashed (which may be present in the state-pooling structure of the optimal signal). We present an example of the output of this procedure at the end of this subsection and a step-by-step illustration of the procedure leading to this output in Appendix 3.C.

Procedure for discovery of the state-pooling structure of the optimal signal:

Input: Set of states Ω ($|\Omega| = n$) and preference-misalignment function $\rho : \Omega \rightarrow \mathbb{R}$.

1. Create a fully connected graph on n nodes where node i corresponds to state ω_i .
2. Eliminate all edges ij such that the slope of ρ on $\omega_i < \omega_j$, s_{ij} , is in $(0, 2)$.
3. Highlight in full each isolated node (i.e., a node with no edges leading to any other node) as a singleton pool.
4. Among the remaining (i.e., non-isolated) nodes, list all maximal cliques.
5. For each maximal clique C :
 - for k from $|C|$ to 2:
 - for all subsets $M \subseteq C$ such that $|M| = k$:
 - leftmargin=2cm If M was ever inspected before, do nothing and continue iteration.
 - leftmargin=2cm If M is a subset of a highlighted set of nodes, do nothing and continue iteration.
 - leftmargin=2cm Otherwise, apply the non-disclosure test to the inspected pool M : Is ρ linear with slope in $(-\infty, 0] \cup [2, \infty)$ on the states corresponding to the nodes in M ?
 - If yes, highlight pool M in dashed on output and continue iteration.
 - If no, denote M as inspected and continue iteration.
6. If any node belongs only to one highlighted pool (in dashed), highlight the corresponding pool in full (if not already highlighted in full).

An example of the output produced by this procedure appears in the right panel of Figure 3.3; an example of function ρ leading to this output is depicted

in the left panel.¹² State 1 is isolated because the sender and the receiver agree on its position relative to other states both in order and in magnitude, so there is no reason for the sender to leverage this state for manipulation of beliefs. States 2, 3, and 4 are pooled together (they pass the non-disclosure test) because the sender tries to moderate the action of the receiver, who would overreact in these states (disagreement about magnitude). States 3 and 5 may be pooled together (disagreement about order) and 4 and 5 may also be pooled together (disagreement about order), but states 3, 4, and 5 are not pooled together even though they form a maximal clique (because they do not pass the non-disclosure test) – the sender prefers to exploit some variation in this collection of nodes. Hence, the optimal signal will induce posterior $p_1 = \delta_1$ and posterior p_2 supported on 2, 3, and 4. Moreover, it will induce at least one of the posteriors p_3 or p_4 supported on 3 and 5 or 4 and 5, respectively.




Figure 3.3: Output of the graph procedure (right panel) for function ρ on the left panel: 1 is isolated; 2, 3, and 4 are pooled together (they pass the non-disclosure test); 3 and 5 may be pooled together; 4 and 5 may be pooled together; 3, 4, and 5 are not pooled together (they do not pass the non-disclosure test)

3.5.3 Discussion of the procedure

The idea underlying our proposed procedure is the iterative application of Proposition 3.4.1 and Proposition 3.4.3, which we call a *top-down approach*. Starting from the full $(n - 1)$ -dimensional simplex,¹³ we can check whether non-disclosure is optimal using Proposition 3.4.1. If it is optimal, the sender chooses a completely uninformative signal. If it is not, Proposition 3.4.3 suggests that the optimal signal will induce posteriors on the boundary of the $(n - 1)$ -dimensional simplex. Hence, we focus on each of the $(n - 2)$ -dimensional boundary simplexes and apply the same test. Specifically, by restricting the sender's expected utility (as a function of the posterior) on a particular $(n - 2)$ -dimensional simplex, we use Proposition 3.4.1 to check if non-disclosure is optimal there:

¹²A step-by-step illustration of the procedure leading to this output appears in Appendix 3.C.

¹³We start from the $(n - 1)$ -dimensional simplex because $p_n = 1 - p_1 - \dots - p_{n-1}$.

- If it is optimal, then the sender cannot benefit from splitting the pool of states corresponding to the vertices of the inspected $(n - 2)$ -dimensional simplex. However, the sender might not want to choose this pool of states at all, so this pool of states constitutes only a candidate pool for the optimal signal.¹⁴
- If it is not optimal, then by Proposition 3.4.3 we eliminate all interior points from the inspected $(n - 2)$ -dimensional simplex and restrict our focus to its $(n - 3)$ -dimensional boundary simplexes; for each of them, we repeat the same steps.

Along the path from the full $(n - 1)$ -dimensional simplex to lower-dimensional simplexes due to elimination of “interior” posteriors outlined in the second bullet point, we move closer to the trivial case of 1-dimensional simplexes where we apply Lemma 17.

Our procedure relies on this top-down approach in Step 5. However, compared to the top-down approach, the procedure starts with a simplification of the problem by identifying the only relevant subsets of nodes for this inspection – the maximal cliques (Steps 2 and 4). This step is justified by the fact that the necessary condition for optimality of non-disclosure on a simplex is optimality of non-disclosure on its boundary simplexes, which follows easily from Proposition 3.4.1. Hence, if we have a given collection of nodes with some pair of nodes in it that is not pooled, this whole collection of nodes cannot form a pool.

In Steps 3 and 6 of the procedure, we exploit Bayesian consistency (and the interior prior). In particular, the structure of the graph obtained after Step 2 is informative about the state-pooling structure by itself: any isolated node represents a state that is fully disclosed. In Step 5, we can identify only candidates for optimal pools, but, in Step 6, Bayesian consistency can help us to determine which of them will be certainly a part of the optimal pooling structure.

Note that we have not mentioned the prior in our identification of the optimal pooling structure. This prior-independence of our procedure relies on a feature of the quadratic setting: constant convexity/concavity structure in all points. However, even in the quadratic setting, the pooling structure of the optimal signal itself is not always prior-independent. This feature imposes a limit on how far we can go with our simple prior-independent procedure in identifying the full pooling structure of the optimal signal. In some cases, we also need to incorporate the prior into our analysis at the end of the procedure (see Section 3.6 for examples).

3.6 Characterization of the state-pooling structure for $n = 3$

In this section, we use the above procedure to characterize the state-pooling structure of the optimal signal in the simplest interesting case of three states (the case of two states is trivial and is fully characterized in Lemma 17). We describe the state-pooling structure for all possible cases of the form of ρ , which we capture

¹⁴Here, we also use Assumption 3.1. This simplifies the analysis because we do not need to keep track of all equivalent splits.

through s_{12} , s_{23} , and s_{13} . For clarity of exposition, we divide the cases into five classes (i)-(v) based on the features of the resulting state-pooling structure and the role of the prior. Class (i) corresponds to full disclosure, class (ii) corresponds to signals that fully disclose one of the states, classes (iii) and (iv) correspond to signals that reveal some information without fully revealing any of the states, and class (v) corresponds to non-disclosure. Within a given class, we use letters to distinguish between particular state-pooling structures.

Assume that there are three states of the world, $\Omega = \{\omega_1, \omega_2, \omega_3\}$. Depending on the form of ρ , as pinned down by s_{12} , s_{23} , and s_{13} , the state-pooling structure of the optimal signal is as follows:

	s_{12}	s_{23}	s_{13}	state-pooling structure
i	$\in (0, 2)$	$\in (0, 2)$	$\in (0, 2)$	$\{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}\}$
ii.a	$\in (0, 2)$	$\notin (0, 2)$	$\in (0, 2)$	$\{\{\omega_1\}, \{\omega_2, \omega_3\}\}$
ii.b	$\notin (0, 2)$	$\in (0, 2)$	$\in (0, 2)$	$\{\{\omega_3\}, \{\omega_1, \omega_2\}\}$
iii.a	$\notin (0, 2)$	$\in (0, 2)$	$\notin (0, 2)$	$\{\{\omega_1, \omega_2\}, \{\omega_1, \omega_3\}\}$
iii.b	$\in (0, 2)$	$\notin (0, 2)$	$\notin (0, 2)$	$\{\{\omega_2, \omega_3\}, \{\omega_1, \omega_3\}\}$
iii.c	$\notin (0, 2)$	$\notin (0, 2)$	$\in (0, 2)$	$\{\{\omega_1, \omega_2\}, \{\omega_2, \omega_3\}\}$
iv ¹⁵	$\notin (0, 2)$	$\notin (0, 2)$	$\notin (0, 2)$	depending on s_{12}, s_{23}, s_{13} , and prior, either (iii.a), (iii.b), or (iii.c) pooling
v	$s_{12} = s_{23} = s_{13} = s \notin (0, 2)$			$\{\{\omega_1, \omega_2, \omega_3\}\}$

Proof. The proof is in Appendix 3.A. □

The observed state-pooling structures emerge from the interaction of the two main forces that drive the sender's choice. On the one hand, the sender wants to disclose the states so that the induced receiver's actions vary sufficiently with the state of the world. On the other hand, she wants to pool the states together to dampen that variation if there is a severe misalignment in either order or magnitude in some pairs of states. The slope of ρ for states ω_i and ω_j , s_{ij} ($i, j \in \{1, 2, 3\}$, $i \neq j$), serves as an index that can capture the misalignment in either order or magnitude in that pair of states.

In case (i), there is no severe preference misalignment in either pair of states, so the sender fully discloses each state. In case (ii.a), s_{23} captures a severe preference misalignment in the pair of states ω_2, ω_3 , so the sender pools these states together to conceal the misalignment but reveals state ω_1 to maximize the informativeness of the signal. In case (iii.a), s_{12} and s_{13} capture a severe preference misalignment in two pairs of states, so the sender pools the respective pairs together but still reveals some information: $\{\{\omega_1, \omega_2\}, \{\omega_1, \omega_3\}\}$. In case (iv), there is a misalignment in each of the three pairs of states and the optimal state-pooling structure is sensitive to the prior and to the relation between the slopes of ρ .

¹⁵ $s_{12} = s_{23} = s_{13} = s \notin (0, 2)$ corresponds to non-disclosure, so we exclude this combination from case (iv) and denote it as a separate case (v). See Appendix 3.A for details on the choice from (iii.a), (iii.b), and (iii.c).

A notable feature of the state-pooling structure of the optimal signal under $n = 3$ is that the sender never chooses to fully disclose the middle state of the world ω_2 and pool ω_1 and ω_3 together. For that to be the case, it would need to hold $s_{13} \notin (0, 2)$, $s_{12} \in (0, 2)$, and $s_{23} \in (0, 2)$, which cannot happen.¹⁶ The intuition is that full disclosure of ω_2 and pooling of ω_1 and ω_3 is not in line with the sender's preference for maximizing the variance of the induced posterior beliefs. A potentially better way to leverage state ω_2 is to form two pools $\{\omega_2, \omega_1\}$ and $\{\omega_2, \omega_3\}$ because it can induce relatively more variation in the receiver's actions.

3.7 Conclusion

We consider a Bayesian persuasion model in which both the sender and the receiver have state-dependent preferred actions. We specialize to a quadratic-utility setting to simplify the otherwise nontrivial problem of characterizing the optimal signal. In this framework, we make the trade-off that drives the sender's choice of the signal transparent: on the one hand, the sender wants to reveal information to adapt the action to the state of the world; on the other hand, she wants to hide information to conceal the misalignment between her and the receiver.

We focus on characterization of the state-pooling structure of the optimal signal. In particular, we link the form of misalignment between the sender and the receiver in their preferred (state-dependent) actions to the state-pooling structure of the sender's optimal signal. To achieve this goal, we propose an illustrative graphical procedure for finding the sets of states that are pooled together in the supports of posteriors of the optimal signal.

Our model naturally suits the analysis of influence in political economy. The sender's and receiver's (state-dependent) single-peaked preferences over the continuous action space are consistent with ideology-based preferences over a continuous set of policy alternatives. That set could represent potential allocations of a resource such as the amount of budget spending on a public good. Thus, our framework can capture an arbitrary form of ideological disagreement between a lobbyist and a policymaker regarding the preferred state-dependent policy and yield predictions about the structure of the lobbyist's chosen information disclosure.

Our analysis motivates a number of directions for further research. First, further investigation and economic interpretation of particular state-pooling patterns that emerge when there are more than three states of the world might be of interest. Second, more progress could be made on analyzing state-pooling patterns that may emerge under loss functions of a more general form.

¹⁶Note that $s_{13} = \frac{\rho(\omega_3) - \rho(\omega_1)}{\omega_3 - \omega_1} = \frac{1}{(\omega_3 - \omega_2) + (\omega_2 - \omega_1)} (s_{23}(\omega_3 - \omega_2) + s_{12}(\omega_2 - \omega_1))$ and $(0, 2)$ is a convex set.

3.A Technical details and proofs

3.A.1 The structure of the sender's problem

We are interested in the solution of the sender's problem

$$\max_{\pi \in \Delta(\Delta(\Omega))} -E_{\pi} \left[E_p \left[(E_p [\rho(\omega)] - \omega)^2 \right] \right] \text{ s.t. } \sum_{p \in \text{supp}(\pi)} \pi(p)p = p_0. \quad (3.8)$$

We can rewrite the objective function as

$$-E_{\pi} \left[E_p [\rho(\omega)]^2 - 2E_p [\rho(\omega)] E_p [\omega] + E_p [\omega^2] \right]. \quad (3.9)$$

Using the Bayesian consistency condition $\sum_p \pi(p)p = p_0$, we can see that the last term becomes

$$-E_{p_0} [\omega^2]. \quad (3.10)$$

Therefore, the solution to the problem above is the same as the solution to the problem

$$\max_{\pi \in \Delta(\Delta(\Omega))} E_{\pi} [E_p [\rho(\omega)] (2E_p [\omega] - E_p [\rho(\omega)])] \text{ s.t. } \sum_{p \in \text{supp}(\pi)} \pi(p)p = p_0. \quad (3.11)$$

A general approach to solving this problem is concavification of the function

$$g(p) = E_p [\rho(\omega)] (2E_p [\omega] - E_p [\rho(\omega)]). \quad (3.12)$$

We use the parametrization $g(p) = g(p_1, p_2, \dots, p_{n-1})$, where $p_n = 1 - p_1 - \dots - p_{n-1}$. We collect the free variables in the vector

$$\bar{p} = (p_1, \dots, p_{n-1})'.$$

We also denote

$$\begin{aligned} \bar{\rho} &= (\rho(\omega_1) - \rho(\omega_n), \dots, \rho(\omega_{n-1}) - \rho(\omega_n))', \\ \bar{\omega} &= (\omega_1 - \omega_n, \dots, \omega_{n-1} - \omega_n)'. \end{aligned}$$

With this notation, we can write

$$g(\bar{p}) = \underbrace{\bar{p}' [2\bar{\rho}\bar{\omega}' - \bar{\rho}\bar{\rho}']}_{G} \bar{p} + [2\omega_n \bar{\rho}' - \rho_n \bar{\rho}' + 2\rho_n \bar{\omega}' - \rho_n \bar{\rho}'] \bar{p} + 2\rho_n \omega_n - \rho_n^2. \quad (3.13)$$

Hence, the curvature of g is driven by matrix G because the Hessian matrix is

$$H = G + G'.^{17} \quad (3.15)$$

The ij element ($i, j \in \{1, \dots, n-1\}$) of H is

$$\begin{aligned} H_{ij} &= \frac{\partial^2 g(p)}{\partial p_i \partial p_j} = 2\{[\rho(\omega_i) - \rho(\omega_n)](\omega_j - \omega_n) - [\rho(\omega_i) - \rho(\omega_n)][\rho(\omega_j) - \rho(\omega_n)] \\ &\quad + [\rho(\omega_j) - \rho(\omega_n)](\omega_i - \omega_n)\}. \end{aligned} \quad (3.16)$$

¹⁷We can also rewrite g as a linear-quadratic form

$$g(\bar{p}) = \frac{1}{2} \bar{p}' H \bar{p} + [2\omega_n \bar{\rho}' - \rho_n \bar{\rho}' + 2\rho_n \bar{\omega}' - \rho_n \bar{\rho}'] \bar{p} + 2\rho_n \omega_n - \rho_n^2. \quad (3.15)$$

This special structure of the problem implies that general submatrices of order 3 (for $n \geq 4$) of the Hessian matrix H have zero determinants.¹⁸ Hence, by the Laplace expansion of determinants, all submatrices of order $k \geq 3$ have zero determinants. We can deduce from this observation, using the fact that the determinant rank of a matrix is equal to the column/row rank of the matrix,¹⁹ that H has at most two non-zero eigenvalues. Therefore, there are at least $n - 3$ orthogonal directions (in space $\mathbb{R}^{n-1} \ni \bar{p}$) that span the space along which g is linear, and at most two orthogonal directions that span the space (orthogonal to the space spanned by the linear directions) on which g has a less trivial shape.

3.A.2 Proofs

Proof of Proposition 3.4.1. The sender does not benefit from providing any information if and only if g is concave.²⁰ g is concave if and only if its Hessian matrix is negative semidefinite, which can be checked with the test on its principal minors.

Suppose $n \geq 3$ (the case $n = 2$ is covered separately in Lemma 17). Let Δ_k be a principal minor of order k of the Hessian matrix of g . Since $\Delta_k = 0$ for $k \geq 3$ (see the discussion above), a necessary and sufficient condition for g to be concave is $\Delta_1 \leq 0$ and $\Delta_2 \geq 0$ for all Δ_1, Δ_2 .

Let Δ_1^i be the first-order principal minor obtained from row (column) i :

$$\Delta_1^i = 2(\rho(\omega_i) - \rho(\omega_n))(2(\omega_i - \omega_n) - (\rho(\omega_i) - \rho(\omega_n))). \quad (3.17)$$

Let Δ_2^{ij} be the second-order principal minor obtained from rows (columns) i and j :

$$\Delta_2^{ij} = -4[(\rho(\omega_i) - \rho(\omega_j))(\omega_j - \omega_n) - (\rho(\omega_j) - \rho(\omega_n))(\omega_i - \omega_j)]^2. \quad (3.18)$$

We can see that $\Delta_2^{ij} \leq 0$. Hence, g is concave or convex only if $\Delta_2 = 0$ for all Δ_2 . This condition yields a system of $\frac{(n-1)(n-2)}{2}$ equations

$$\Delta_2^{ij} = 0, \quad i, j \in \{1, \dots, n-1\}, \quad i \neq j. \quad (3.19)$$

Under the natural assumption that $\omega_1 < \dots < \omega_n$ (which is without loss of generality), we obtain from $\Delta_2^{ij} = 0$

$$\frac{\rho(\omega_j) - \rho(\omega_i)}{\omega_j - \omega_i} = \frac{\rho(\omega_n) - \rho(\omega_j)}{\omega_n - \omega_j} \quad (3.20)$$

or, equivalently,

$$\frac{\rho(\omega_j) - \rho(\omega_i)}{\omega_j - \omega_i} = \frac{\rho(\omega_n) - \rho(\omega_i)}{\omega_n - \omega_i}. \quad (3.21)$$

¹⁸Proof is available upon request. It is basically just tedious algebra.

¹⁹The *determinant rank* of H is the size k of the largest $k \times k$ submatrix with a non-zero determinant. The *column/row rank* of H is the dimension of the space spanned by the columns/rows of H . It is straightforward to show that these ranks are equal.

²⁰The “if” part follows directly from the definition of concavity. The “only if” part would also follow directly from the definition of concavity if the sender did not benefit from providing any information for every prior. But if the sender does not benefit from providing any information only in one prior, because g is a linear-quadratic form, this property extends to all priors.

Therefore, the system of equations (3.19) gives rise to $\frac{(n-1)(n-2)}{2}$ slope equality conditions. From (3.20) and (3.21), we have

$$\begin{aligned}
j = n - 1, i = n - 2 : \frac{\rho(\omega_n) - \rho(\omega_{n-1})}{\omega_n - \omega_{n-1}} &= \frac{\rho(\omega_{n-1}) - \rho(\omega_{n-2})}{\omega_{n-1} - \omega_{n-2}} = \frac{\rho(\omega_n) - \rho(\omega_{n-2})}{\omega_n - \omega_{n-2}}, \\
j = n - 2, i = n - 3 : \frac{\rho(\omega_n) - \rho(\omega_{n-2})}{\omega_n - \omega_{n-2}} &= \frac{\rho(\omega_{n-2}) - \rho(\omega_{n-3})}{\omega_{n-2} - \omega_{n-3}} = \frac{\rho(\omega_n) - \rho(\omega_{n-3})}{\omega_n - \omega_{n-3}}, \\
&\vdots \\
j = 2, i = 1 : \frac{\rho(\omega_n) - \rho(\omega_2)}{\omega_n - \omega_2} &= \frac{\rho(\omega_2) - \rho(\omega_1)}{\omega_2 - \omega_1} = \frac{\rho(\omega_n) - \rho(\omega_1)}{\omega_n - \omega_1}.
\end{aligned}$$

Hence, system (3.19) is equivalent to a linearity of ρ :

$$s \frac{\rho(\omega_2) - \rho(\omega_1)}{\omega_2 - \omega_1} = \frac{\rho(\omega_3) - \rho(\omega_2)}{\omega_3 - \omega_2} = \dots = \frac{\rho(\omega_n) - \rho(\omega_{n-1})}{\omega_n - \omega_{n-1}}. \quad (3.22)$$

Finally, given that $\Delta_2 = 0$ for all Δ_2 holds, one can establish whether g is concave or convex based on the sign of Δ_1 . Inspecting the sign of (3.17) yields:

$$\Delta_1^i \geq 0 \iff (\rho(\omega_n) - \rho(\omega_i) \geq 0) \wedge \frac{\rho(\omega_n) - \rho(\omega_i)}{\omega_n - \omega_i} \leq 2 \iff 0 \leq s \leq 2. \quad (3.23)$$

The complement identifies the concavity slopes (including the borderline slopes $s \in \{0, 2\}$). \square

Proof of Proposition 3.4.2. This proposition is basically proven in the proof of Proposition 3.4.1, using the fact that g is convex if and only if $\Delta_1 \geq 0$ and $\Delta_2 \geq 0$ for all Δ_1, Δ_2 . The only difference is that the convexity of g is only sufficient for optimality of full disclosure, but is not necessary (we can provide an example of optimal full disclosure with non-convex g). \square

Proof of Lemma 17. For $n = 2$, g is a quadratic function, so its second derivative completely characterizes its curvature, which completely characterizes the type of optimal signals. In particular, let $\omega_1 < \omega_2$. Then,

$$\frac{\partial^2 g(p_1)}{\partial p_1^2} = 2(\rho(\omega_1) - \rho(\omega_2))(2(\omega_1 - \omega_2) - (\rho(\omega_1) - \rho(\omega_2))), \quad (3.24)$$

which is strictly positive if and only if the slope of ρ is in $(0, 2)$ (strict convexity and full disclosure), strictly negative if and only if the slope of ρ is in $(-\infty, 0) \cup (2, \infty)$ (strict concavity and non-disclosure), and zero if and only if the slope of ρ is either zero or two (linearity and indifference). \square

Proof of Proposition 3.4.3. Non-disclosure is optimal if and only if g is concave. Hence, if non-disclosure is not optimal, g is not concave. Therefore, g has to have a direction along which it is strictly convex.²¹

Suppose (toward contradiction) that it is optimal to induce an interior posterior, i.e., there exists a posterior p in the support of the optimal signal π such that $p(\omega) > 0 \forall \omega$. Then, we can split p along a strictly convex direction to q_1

²¹This is independent of the position because g is a linear-quadratic form.

and q_2 , i.e., there exists some $\lambda \in (0, 1)$ such that $p = \lambda q_1 + (1 - \lambda)q_2$. Then, π' formed from π by replacing p by q_1 with probability $\lambda\pi(p)$ and q_2 with probability $(1 - \lambda)\pi(p)$ is Bayes-plausible and it induces a strict improvement for the sender because, from strict convexity of g along the direction determined by q_1 and q_2 ,

$$E_{\pi'} [g(p)] - E_{\pi} [g(p)] = \pi(p)(\lambda g(q_1) + (1 - \lambda)g(q_2) - g(p)) > 0. \quad (3.25)$$

This is a contradiction with optimality of π . \square

Proof of Proposition 3.6. We derive the state-pooling structure for the form of ρ for each case presented in the table of Proposition 3.6 using the graph procedure presented in Section 3.5.2.

Case (i). Since $s_{12}, s_{23}, s_{13} \in (0, 2)$, Step 2 of the procedure eliminates all edges, so each node is highlighted in full in Step 3. Thus immediately after Step 3, the procedure yields the state-pooling structure of the optimal signal $\{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}\}$.

Case (ii.a). Since $s_{12}, s_{13} \in (0, 2)$ and $s_{23} \notin (0, 2)$, after Step 2 of the procedure, node 1 is isolated (thus, it is highlighted in full in Step 3) and there is an edge left between nodes 2 and 3. Since the pool $\{2, 3\}$ is a maximal clique (Step 4) and ρ is obviously linear with slope from $(-\infty, 0] \cup [2, \infty)$ on states ω_2 and ω_3 , this pool is highlighted in dashed in Step 5. Finally, it is highlighted in full in Step 6 because nodes 2 and 3 belong only to this pool. Therefore, the state-pooling structure of the optimal signal is $\{\{\omega_1\}, \{\omega_2, \omega_3\}\}$.

Case (ii.b). Analogous to case (ii.a).

Case (iii.a). Since $s_{12}, s_{13} \notin (0, 2)$ and $s_{23} \in (0, 2)$, after Step 2 of the procedure, there are two edges left: one between nodes 1 and 2 and one between nodes 1 and 3. Since both pools $\{1, 2\}$ and $\{1, 3\}$ are maximal cliques (Step 4) and ρ is obviously linear with slope from $(-\infty, 0] \cup [2, \infty)$ on states ω_1, ω_2 and ω_1, ω_3 , respectively, these pools are highlighted in dashed in Step 5. Finally, they are highlighted in full in Step 6 because node 2 belongs only to pool $\{1, 2\}$ and node 3 belongs only to pool $\{1, 3\}$. Therefore, the state-pooling structure of the optimal signal is $\{\{\omega_1, \omega_2\}, \{\omega_1, \omega_3\}\}$.

Case (iii.b). Analogous to case (iii.a).

Case (iii.c). Analogous to case (iii.a).

Case (iv). We assume that $s_{12} = s_{23} = s_{13} = s \notin (0, 2)$ does not hold (this case is covered by case (v)). Thus, the graph procedure yields the candidate pools $\{\omega_1, \omega_2\}$, $\{\omega_2, \omega_3\}$, and $\{\omega_1, \omega_3\}$ (corresponding to the pools of nodes highlighted in dashed in the graph). To determine the optimal state-pooling structure given the set of candidate pools is non-trivial.

Denote the n -th directional derivative of a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ along a direction (a, b) by $D_{(a,b)}^n f$. Denote $p_1 \Pr(\omega_1)$ and $p_2 \Pr(\omega_2)$. From the proof of Proposition 3.4.1, the nonlinearity in ρ implies that there exists a direction (a, b) along which $g(p_1, p_2)$ (defined in (3.12)) is strictly convex. The set of all such directions is pinned down by the condition

$$D_{(a,b)}^2 g(p) > 0, \quad (3.26)$$

which rewrites as (assuming $s_{13} \neq 0$ and $s_{23} \neq 0$; see below for the discussion of

these cases)

$$\begin{aligned}
& a^2 (\rho(\omega_1) - \rho(\omega_3)) [2(\omega_1 - \omega_3) - (\rho(\omega_1) - \rho(\omega_3))] + \\
& b^2 (\rho(\omega_2) - \rho(\omega_3)) [2(\omega_2 - \omega_3) - (\rho(\omega_2) - \rho(\omega_3))] + \\
& ab \frac{\rho(\omega_2) - \rho(\omega_3)}{\rho(\omega_1) - \rho(\omega_3)} (\rho(\omega_1) - \rho(\omega_3)) [2(\omega_1 - \omega_3) - (\rho(\omega_1) - \rho(\omega_3))] + \\
& ab \frac{\rho(\omega_1) - \rho(\omega_3)}{\rho(\omega_2) - \rho(\omega_3)} (\rho(\omega_2) - \rho(\omega_3)) [2(\omega_2 - \omega_3) - (\rho(\omega_2) - \rho(\omega_3))] > 0.
\end{aligned} \tag{3.27}$$

Next, $s_{13} \notin (0, 2) \wedge s_{23} \notin (0, 2)$ implies²²

$$\begin{cases} (\rho(\omega_1) - \rho(\omega_3)) [2(\omega_1 - \omega_3) - (\rho(\omega_1) - \rho(\omega_3))] \leq 0, \\ (\rho(\omega_2) - \rho(\omega_3)) [2(\omega_2 - \omega_3) - (\rho(\omega_2) - \rho(\omega_3))] \leq 0. \end{cases} \tag{3.28}$$

We can see from (3.27) and (3.28) that if (a, b) is a direction along which g is strictly convex, both a and b have to be non-zero. Thus, we can normalize the direction (a, b) to $(\frac{a}{b}, 1)$ and denote $x := \frac{a}{b}$. Hence, the set of directions along which g is strictly convex is characterized by

$$\begin{aligned}
& x^2 (\rho(\omega_1) - \rho(\omega_3)) [2(\omega_1 - \omega_3) - (\rho(\omega_1) - \rho(\omega_3))] + \\
& (\rho(\omega_2) - \rho(\omega_3)) [2(\omega_2 - \omega_3) - (\rho(\omega_2) - \rho(\omega_3))] + \\
& x \frac{\rho(\omega_2) - \rho(\omega_3)}{\rho(\omega_1) - \rho(\omega_3)} (\rho(\omega_1) - \rho(\omega_3)) [2(\omega_1 - \omega_3) - (\rho(\omega_1) - \rho(\omega_3))] + \\
& x \frac{\rho(\omega_1) - \rho(\omega_3)}{\rho(\omega_2) - \rho(\omega_3)} (\rho(\omega_2) - \rho(\omega_3)) [2(\omega_2 - \omega_3) - (\rho(\omega_2) - \rho(\omega_3))] > 0.
\end{aligned} \tag{3.29}$$

Inspecting (3.29) given (3.28), one observes that the first two terms in (3.29) are non-positive. Therefore, the sum of the last two terms must necessarily be strictly positive for any direction along which g is strictly convex. Further, if the third term is strictly negative, the fourth term is non-positive and vice versa. So, if either of the last two terms is strictly negative, their sum is also strictly negative. Equivalently, if their sum is non-negative, they both have to be non-negative. Moreover, if their sum is strictly positive, they cannot both be zero. But if any one of the last two terms in (3.29) is strictly positive, then by (3.28)

$$x \frac{\rho(\omega_1) - \rho(\omega_3)}{\rho(\omega_2) - \rho(\omega_3)} < 0. \tag{3.30}$$

To summarize, if $(x, 1)$ is a direction along which g is strictly convex, then

$$\begin{cases} x > 0 & \text{if } \frac{\rho(\omega_1) - \rho(\omega_3)}{\rho(\omega_2) - \rho(\omega_3)} < 0 \quad (\iff \frac{s_{13}}{s_{23}} < 0), \\ x < 0 & \text{if } \frac{\rho(\omega_1) - \rho(\omega_3)}{\rho(\omega_2) - \rho(\omega_3)} > 0 \quad (\iff \frac{s_{13}}{s_{23}} > 0). \end{cases} \tag{3.31}$$

By similar arguments, if $s_{13} = 0$,²³ the necessary condition for $(x, 1)$ being the direction along which g is strictly convex is

$$\begin{cases} x > 0 & \text{if } s_{23} > 0, \\ x < 0 & \text{if } s_{23} < 0 \end{cases} \tag{3.32}$$

²²At least one of these terms is non-zero due to the assumption that $s_{12} = s_{23} = s_{13} = s \notin (0, 2)$ does not hold.

²³Notice that s_{13} and s_{23} cannot be simultaneously zero by assumption, because this would lead to case (v).

and if $s_{23} = 0$, the necessary condition for $(x, 1)$ being the direction along which g is strictly convex is

$$\begin{cases} x > 0 & \text{if } s_{13} > 0, \\ x < 0 & \text{if } s_{13} < 0. \end{cases} \quad (3.33)$$

Given some interior prior, the sender splits it along a direction along which g is strictly convex and induces posteriors that lie on two edges of the simplex. We can distinguish the following cases:

1. If $\frac{s_{13}}{s_{23}} < 0$ or $s_{13} = 0 \wedge s_{23} > 0$ or $s_{23} = 0 \wedge s_{13} > 0$, then $x > 0$. Hence, the optimal split is either of the form $(q_1, 0, 1 - q_1), (1 - q_2, q_2, 0)$ (pooling case (iii.a)) or of the form $(q_1, 1 - q_1, 0), (0, q_2, 1 - q_2)$ (pooling case (iii.c)) depending on the prior.
2. If $\frac{s_{13}}{s_{23}} > 0$ or $s_{13} = 0 \wedge s_{23} < 0$ or $s_{23} = 0 \wedge s_{13} < 0$, then $x < 0$. In this case, we need to distinguish further:
 - (a) If the optimal split goes along the direction $(-1, 1)$, it is of the form $(q_1, 0, 1 - q_1), (0, q_2, 1 - q_2)$ (pooling case (iii.b)).
 - (b) If the optimal split goes along direction $(x, 1)$ with $x < -1$, it is either of the form $(q_1, 0, 1 - q_1), (0, q_2, 1 - q_2)$ (pooling case (iii.b)) or of the form $(q_1, 1 - q_1, 0), (0, q_2, 1 - q_2)$ (pooling case (iii.c)) depending on the prior.
 - (c) If the optimal split goes along direction $(x, 1)$ with $x > -1$, it is either of the form $(q_1, 0, 1 - q_1), (0, q_2, 1 - q_2)$ (pooling case (iii.b)) or of the form $(q_1, 0, 1 - q_1), (q_2, 1 - q_2, 0)$ (pooling case (iii.a)) depending on the prior.

Case (v). Proposition 3.4.1 applies and under Assumption 3.1 yields non-disclosure. \square

3.B Comment on Assumption 3.1

The structure of function g (see (3.12)) uncovered in Section 3.A.1 implies that for $n \geq 4$, there always exists a direction along which g is linear. Therefore, even when g is concave and non-disclosure is optimal, it is never uniquely optimal for $n \geq 4$. In particular, the sender is indifferent between sticking to the prior and splitting it to some posteriors from the space determined by the linear directions of g (and the prior), possibly all the way to the boundaries of the original simplex. Moreover, if g is concave, it is also concave on the boundary simplexes and we can repeat the same argument, proceeding downward in dimensions. For $n = 3$, by Proposition 3.4.1, g is concave only if it is linear in one direction. Hence, even for $n = 3$, non-disclosure is not uniquely optimal and the sender is indifferent between choosing a non-informative signal (keeping the belief at the prior) and splitting the prior into posteriors along the linear direction, all the way to the edges of the simplex. Therefore, pairwise signals (i.e., signals leading to posteriors supported on at most two states) are also always optimal.²⁴

²⁴This result is reminiscent of the result of Kolotilin and Wolitzky [2020] that there is no loss of generality from focusing on pairwise signals in their setup.

In the main text, we impose Assumption 3.1, which resolves indifference in favor of non-disclosure of states. It is a natural assumption that can be justified by the sender not wasting resources (time and energy) on communication when it is not needed (although the cost of communication is not featured explicitly in our model). This selection criterion simplifies the analysis. First, it enables us to avoid imposing some ad hoc assumptions about the selection of specific partial disclosure patterns from the indifference set. Second, a different natural assumption might be that the sender resolves her indifference in favor of splitting. However, this assumption would require us to impose some additional ad hoc assumptions about the selection of specific directions along which to split (for higher n) in order to deliver concrete predictions. Moreover, such a resolution of indifference would be very sensitive to the prior (even in terms of the predicted pooling structure), so we would need to keep track of the specific directions of indifference, which would render the analysis much more cumbersome.²⁵

3.C Demonstration of the procedure for discovery of the state-pooling structure of the optimal signal

We demonstrate the application of the procedure for discovery of the state-pooling structure of the optimal signal (presented in Section 3.5) to the example introduced in Figure 3.3 (for convenience, we reproduce it in Figure 3.4 in this section). This demonstration is accompanied by Figure 3.5. *Red color* in Figure 3.5 represents highlighting as defined in Section 3.5 – final pools in full and candidate pools in dashed. *Green color* denotes cliques chosen for application of the non-disclosure test (Step 5 of the procedure).




Figure 3.4: Preference misalignment function ρ considered for the demonstration of the graph procedure

²⁵To illustrate the dependence on the prior, for $n = 3$ under linear ρ (which is sufficient for global concavity or convexity), the direction of linearity is $(-\frac{\omega_3 - \omega_2}{\omega_3 - \omega_1}, 1)'$. Since the first component is strictly between 0 and -1, we can see that, while the non-disclosure is also optimal, the state-pooling structure (defined in Section 3.5) of the optimal informative signal can be either $\{\{\omega_1, \omega_3\}, \{\omega_2, \omega_3\}\}$ or $\{\{\omega_1, \omega_3\}, \{\omega_1, \omega_2\}\}$, depending on the prior.

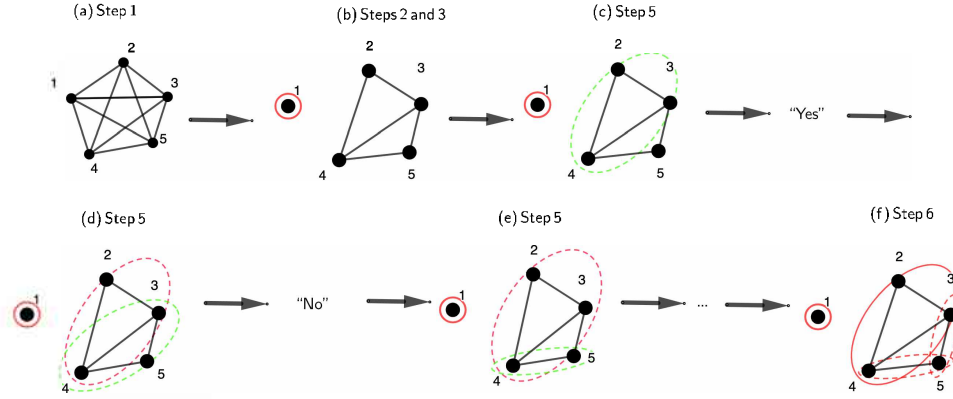


Figure 3.5: Illustration of the execution of the procedure, applied to the input from Figure 3.4; the output is in (f); red color represents highlighting as defined in Section 3.5 – final pools in full and candidate pools in dashed; green color denotes cliques chosen for application of the non-disclosure test

The inputs to the procedure are the values of ω and $\rho(\omega)$ from Figure 3.4. From formula (3.7), we obtain the values of all s_{ij} : $s_{12} = 0.5$, $s_{13} = 1.5$, $s_{14} = \frac{5.5}{3}$, $s_{15} = \frac{2.5}{4}$, $s_{23} = 2.5$, $s_{24} = 2.5$, $s_{25} = \frac{2}{3}$, $s_{34} = 2.5$, $s_{35} = -\frac{1}{4}$, $s_{45} = -3$.

In (a) in Figure 3.5, we start with a fully connected graph on five nodes ($n = 5$) corresponding to states 1, 2, ..., 5.

In (b) in Figure 3.5, we observe the same graph after the application of Steps 2 and 3 of the procedure. We removed all edges ij such that $s_{ij} \in (0, 2)$. As a result, node 1 became isolated, so we highlighted it in full. Hence, we can leave out node 1 from further analysis and focus on nodes 2, 3, 4, and 5.

In (c) in Figure 3.5, we proceed to Steps 4 and 5 of the procedure. It is easily seen that there are two maximal cliques: one formed by nodes 2, 3, and 4 and one formed by nodes 3, 4, and 5. First, we inspect the maximal clique formed by 2, 3, and 4 (highlighted in green) and we apply the non-disclosure test. The non-disclosure condition holds, so we highlight the maximal clique $\{2, 3, 4\}$ in dashed (as illustrated in (d)). Hence, we do not need to consider any more of its subsets in Step 5 and we can move our focus to the other maximal clique.

In (d) in Figure 3.5, we inspect the maximal clique formed by nodes 3, 4, and 5 (highlighted in green). The non-disclosure condition does not hold, so we denote the maximal clique $\{3, 4, 5\}$ as inspected and proceed to consider its subsets of cardinality 2.

In (e) in Figure 3.5, we first consider the clique formed by nodes 4 and 5. As the non-disclosure condition is satisfied, we highlight this clique in dashed. Proceeding with the iteration, we test clique $\{3, 5\}$. Again, the non-disclosure condition is satisfied, so we highlight it in dashed. Finally, clique $\{3, 4\}$ is a subset of the highlighted set $\{2, 3, 4\}$, so we do not test it.

In (f) in Figure 3.5, we proceed to Step 6 of the procedure: as node 2 belongs to only one highlighted clique, $\{2, 3, 4\}$, we highlight that clique in full. The output of the procedure is depicted in (f) in Figure 3.5: the singleton pool $\{1\}$ and pool

$\{2, 3, 4\}$ highlighted in full and pools $\{3, 5\}$ and $\{4, 5\}$ highlighted in dashed. Hence, the posteriors induced by the optimal signal certainly include a posterior supported on states $\omega_2 = 2, \omega_3 = 3, \omega_4 = 4$ and the posterior δ_{ω_1} . Moreover, the optimal signal will induce at least one posterior supported on $\omega_3 = 3, \omega_5 = 5$ or $\omega_4 = 4, \omega_5 = 5$.

Bibliography

- Anat R Admati and Paul Pfleiderer. Robust financial contracting and the role of venture capitalists. *The Journal of Finance*, 49(2):371–402, 1994.
- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, third edition, 2006.
- Ricardo Alonso and Odilon Câmara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.
- Ricardo Alonso and Niko Matouschek. Optimal delegation. *The Review of Economic Studies*, 75(1):259–293, 2008. ISSN 00346527, 1467937X. URL <http://www.jstor.org/stable/4626195>.
- Ricardo Alonso, Wouter Dessein, and Niko Matouschek. Centralization Versus Decentralization: An Application to Price Setting by a Multi-Market Firm. *Journal of the European Economic Association*, 6(2-3):457–467, 2008. ISSN 1542-4774. doi: 10.1162/JEEA.2008.6.2-3.457.
- Rossella Argenziano, Sergei Severinov, and Francesco Squintani. Strategic Information Acquisition and Transmission. *American Economic Journal: Microeconomics*, 8(3):119–155, August 2016. ISSN 1945-7669, 1945-7685. doi: 10.1257/mic.20140135. URL <http://pubs.aeaweb.org/doi/10.1257/mic.20140135>.
- Itai Arieli, Yakov Babichenko, Rann Smorodinsky, and Takuro Yamashita. Optimal persuasion via bi-pooling. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 641, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379755. doi: 10.1145/3391403.3399468. URL <https://doi.org/10.1145/3391403.3399468>.
- Mark Armstrong and David E.M. Sappington. Recent Developments in the Theory of Regulation. In *Handbook of Industrial Organization*, volume 3, pages 1557–1700. Elsevier, 2007. ISBN 978-0-444-82435-6. doi: 10.1016/S1573-448X(06)03027-5.
- Robert J. Aumann. Agreeing to Disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
- Martha J Bailey and Nicolas J Duquette. 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, 2014.
- Ian Ball. Dynamic information provision: Rewarding the past and guiding the future. *Available at SSRN 3103127*, 2019.
- Ian Ball and Xin Gao. Benefiting from bias. Working paper, 2021.
- Abhijit Banerjee and Rohini Somanathan. A simple model of voice. *The Quarterly Journal of Economics*, 116(1):189–227, 2001.

- Nicholas C Barberis. Thirty Years of Prospect Theory in Economics: A Review and Assessment. *Journal of Economic Perspectives*, 27(1):173–196, February 2013. ISSN 0895-3309. doi: 10.1257/jep.27.1.173. URL <https://pubs.aeaweb.org/doi/10.1257/jep.27.1.173>.
- Deepal Basak and Zhen Zhou. Panics and early warnings. *PBCSF-NIFR Research Paper*, 2020.
- Dirk Bergemann and Ulrich Hege. Venture capital financing, moral hazard, and learning. *Journal of Banking & Finance*, 22(6-8):703–735, 1998.
- Dirk Bergemann and Ulrich Hege. The financing of innovation: Learning and stopping. *RAND Journal of Economics*, pages 719–752, 2005.
- Giacomo Bonanno and Klaus Nehring. 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), 1997.
- Andrew Caplin, Mark Dean, and John Leahy. Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, 86(3):1061–1094, 2019.
- Yeon-Koo Che and Navin Kartik. Opinions as incentives. *Journal of Political Economy*, 117(5):815–860, 2009.
- Yeon-Koo Che, Wouter Dessein, and Navin Kartik. Pandering to Persuade. *The American Economic Review*, 103(1):47–79, 2013. ISSN 0002-8282. URL <https://www.jstor.org/stable/23469636>.
- Francesca Cornelli and Oved Yosha. Stage financing and the role of convertible securities. *The Review of Economic Studies*, 70(1):1–32, 2003.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982a.
- Vincent P Crawford and Joel Sobel. Strategic Information Transmission. *Econometrica*, 50(6):1431–1451, 1982b.
- Gregorio Curello and Ludvig Sinander. Screening for breakthroughs. *arXiv preprint arXiv:2011.10090*, 2020.
- Joel S. Demski and David EM Sappington. Delegated expertise. *Journal of Accounting Research*, pages 68–89, 1987.
- Wouter Dessein. Authority and communication in organizations. *The Review of Economic Studies*, 69(4):811–838, 2002.
- Piotr Dworzak and Giorgio Martini. The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019.

- Ronald A Dye. Disclosure of nonproprietary information. *Journal of Accounting Research*, pages 123–145, 1985.
- Georgy Egorov and Konstantin Sonin. Dictators and their viziers: Endogenizing the loyalty–competence trade-off. *Journal of the European Economic Association*, 9(5):903–930, October 2011. ISSN 15424766. doi: 10.1111/j.1542-4774.2011.01033.x. URL <https://academic.oup.com/jeea/article-lookup/doi/10.1111/j.1542-4774.2011.01033.x>.
- Emre Ekinçi and Nikolaos Theodoropoulos. Determinants of delegation: Evidence from british establishment data. *Bogazici Journal: Review of Social, Economic & Administrative Studies*, 35(1), 2021.
- Mehmet Ekmekci, Leandro Gorno, Lucas Maestri, Jian Sun, and Dong Wei. Learning from manipulable signals. *arXiv preprint arXiv:2007.08762*, 2020.
- Jeffrey C Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.
- Jeffrey C Ely and Martin Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128(2):468–506, 2020.
- Andrés Espitia. Confidence and organizations. Working paper, 2023.
- Nicolai J Foss and Keld Laursen. Performance pay, delegation and multitasking under uncertainty and innovativeness: An empirical investigation. *Journal of Economic Behavior & Organization*, 58(2):246–276, 2005.
- Simone Galperti. Persuasion: The art of changing worldviews. *American Economic Review*, 109(3):996–1031, 2019.
- Matthew Gentzkow and Emir Kamenica. A Rothschild-Stiglitz approach to Bayesian persuasion. *American Economic Review*, 106(5):597–601, 2016.
- Paul A Gompers. Optimal investment, monitoring, and the staging of venture capital. *The Journal of Finance*, 50(5):1461–1489, 1995.
- John R Graham, Campbell R Harvey, and Manju Puri. Capital allocation and delegation of decision-making authority within firms. *Journal of financial economics*, 115(3):449–470, 2015.
- Brett Green and Curtis R Taylor. Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects. *American Economic Review*, 106(12):3660–99, 2016.
- Katherine Gundolf, Johanna Gast, and Mickael Géraudel. Startups’ innovation behaviour: An investigation into the role of entrepreneurial motivations. *International Journal of Innovation Management*, 21(07):1750054, 2017.
- Jeanne Hagenbach, Frédéric Koessler, and Eduardo Perez-Richet. Certifiable pre-play communication: Full disclosure. *Econometrica*, 82(3):1093–1131, 2014.
- Benjamin Hébert and Michael Woodford. Rational Inattention when Decisions Take Time. Technical Report w26415, National Bureau of Economic Research, Cambridge, MA, 2019. URL <http://www.nber.org/papers/w26415.pdf>.

- Emeric Henry and Marco Ottaviani. Research and the approval process: The organization of persuasion. *American Economic Review*, 109(3):911–55, 2019.
- Bengt Holmström. On the theory of delegation. Discussion Papers 438, Northwestern University, Center for Mathematical Studies in Economics and Management Science, 1980. URL <https://EconPapers.repec.org/RePEc:nwu:cmsems:438>.
- Patrick Hummel, John Morgan, and Phillip C Stocken. A model of voluntary managerial disclosure. Unpublished manuscript, 2018.
- Alessandro Ispano and Péter Vida. Good cop-bad cop: delegating interrogations. Technical report, THEMA (Théorie Economique, Modélisation et Applications), Université de Cergy-Pontoise, 2022.
- Kinshuk Jerath and Qitian Ren. Consumer rational (in) attention to favorable and unfavorable product information, and firm information design. *Journal of Marketing Research*, 58(2):343–362, 2021.
- Daniel Kahneman, Olivier Sibony, and Cass R Sunstein. *Noise: a flaw in human judgment*. Hachette UK, 2021.
- Emir Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11:249–272, 2019.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Ayca Kaya. Paying with information. Available at SSRN 3661779, 2020.
- Levent Kockesen and Efe A. Ok. Strategic delegation by unobservable incentive contracts. *The Review of Economic Studies*, 71(2):397–424, 2004.
- Anton Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018.
- Anton Kolotilin and Alexander Wolitzky. Assortative information disclosure. UNSW Economics Working Paper 2020-08, 2020.
- Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, December 2009. ISBN 978-1-4008-2945-3. doi: 10.1515/9781400829453. URL <https://www.degruyter.com/document/doi/10.1515/9781400829453/html>.
- Jean-Jacques Laffont and Jean Tirole. The Politics of Government Decision Making: Regulatory Institutions. *Journal of Law, Economics, and Organization*, 6(1):1–32, 1990. URL <https://heinonline.org/HOL/P?h=hein.journals/jleo6&i=11>.
- David E. Lewis. Presidential Appointments and Personnel. *Annual Review of Political Science*, 14(1):47–66, 2011. doi: 10.1146/annurev-polisci-042009-121225. URL <https://doi.org/10.1146/annurev-polisci-042009-121225>.

- Hao Li and Wing Suen. Delegating decisions to experts. *Journal of Political Economy*, 112(S1):S311–S335, 2004.
- Assar Lindbeck and Jörgen Weibull. Delegation of investment decisions, and optimal remuneration of agents. *European Economic Review*, 129:103559, 2020.
- Elliot Lipnowski, Laurent Mathevet, and Dong Wei. Attention management. *American Economic Review: Insights*, 2(1):17–32, 2020.
- Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273, 2023.
- Erik Madsen. Designing deadlines. *American Economic Review*, 112(3):963–97, 2022.
- Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298, 2015.
- Andrei Matveenko and Sergei Mikhailishchev. Attentional role of quota implementation. *Journal of Economic Theory*, 198:105356, 2021. ISSN 0022-0531. doi: <https://doi.org/10.1016/j.jet.2021.105356>. URL <https://www.sciencedirect.com/science/article/pii/S0022053121001733>.
- Jeffrey Mensch. Cardinal representations of information. *Available at SSRN 3148954*, 2018.
- Paul Milgrom and John Roberts. Relying on the information of interested parties. *The RAND Journal of Economics*, pages 18–32, 1986.
- Paul R Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.
- Shintaro Miura. Prudence in persuasion. Unpublished manuscript, 2018.
- Stephen Morris. The common prior assumption in economic theory. *Economics & Philosophy*, 11(2):227–253, 1995.
- Kristoffer P Nimark and Savitar Sundaresan. Inattention and belief polarization. *Journal of Economic Theory*, 180:203–228, 2019.
- Dmitry Orlov, Andrzej Skrzypacz, and Pavel Zryumov. Persuading the principal to wait. *Journal of Political Economy*, 128(7):2542–2578, 2020.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. The cost of information: The case of constant marginal costs. *American Economic Review*, 2023.
- Canice Prendergast. A theory of "Yes Men". *The American Economic Review*, 83(4):757–770, 1993. ISSN 0002-8282.
- Heikki Rantakari. Governing Adaptation. *The Review of Economic Studies*, 75(4):1257–1285, 2008. ISSN 0034-6527.

- Jérôme Renault, Eilon Solan, and Nicolas Vieille. Optimal dynamic information provision. *Games and Economic Behavior*, 104:329–349, 2017.
- Kenneth Rogoff. The optimal degree of commitment to an intermediate monetary target. *The quarterly journal of economics*, 100(4):1169–1189, 1985.
- Björn Segendorff. Delegation and Threat in Bargaining. *Games and Economic Behavior*, 23(2):266–283, May 1998. ISSN 0899-8256. doi: 10.1006/game.1997.0611. URL <https://www.sciencedirect.com/science/article/pii/S0899825697906114>.
- Daniel J Seidmann and Eyal Winter. Strategic information transmission with verifiable messages. *Econometrica: Journal of the Econometric Society*, pages 163–169, 1997.
- Hyun Song Shin. The burden of proof in a game of persuasion. *Journal of Economic Theory*, 64(1):253–264, 1994.
- Lynn M Shore, Beth G Chung-Herrera, Michelle A Dean, Karen Holcombe Ehrhart, Don I Jung, Amy E Randel, and Gangaram Singh. Diversity in organizations: Where are we now and where are we going? *Human resource management review*, 19(2):117–133, 2009.
- Jakub Steiner and Colin Stewart. Perceiving Prospects Properly. *American Economic Review*, 106(7):1601–1631, July 2016. ISSN 0002-8282. doi: 10.1257/aer.20141141. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20141141>.
- Sergey Stepanov. Biased performance evaluation in a model of career concerns: incentives versus ex-post optimality. *Journal of Economic Behavior & Organization*, 179:589–607, November 2020. ISSN 0167-2681. doi: 10.1016/j.jebo.2020.09.024. URL <https://www.sciencedirect.com/science/article/pii/S0167268120303589>.
- Dezső Szalay. The economics of clear advice and extreme options. *The Review of Economic Studies*, 72(4):1173–1198, 2005.
- Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, October 1992. ISSN 1573-0476. doi: 10.1007/BF00122574. URL <https://doi.org/10.1007/BF00122574>.
- Christoph Wolf. Informative milestones in experimentation. *University of Mannheim, Working Paper*, 2017.
- Michael Woodford. Inattentive valuation and reference-dependent choice. Working paper, 2012.