Endogenous Information Acquisition in Cheap-Talk Games*

Sophie Kreutzkamp†

November 3, 2022

Job Market Paper

Link to latest version

Abstract

This paper studies costly information acquisition and transmission. An expert communicates with a decision-maker about a state of nature by sending a cheap-talk message. In efficient equilibria, the expert generally reveals all acquired information to the decision-maker. I show the existence of efficient equilibria under general conditions. For the class of posterior separable cost structures, I derive properties of efficient experiments. Under posterior-mean preferences, any cheap-talk problem is solved by a convex combination of two *bi-pooling policies*. The best bi-pooling policies are characterized for the uniform-quadratic case. Contrary to existing cheap-talk models, monotone partitions are not always optimal.

Keywords: Cheap-talk, Information acquisition, Communication, Bi-pooling

JEL Codes: D82, D83

*I am grateful to Dirk Bergemann, Nina Bobkova, Francesc Dilmé, Patrick Lahr, Stephan Lauermann, Benny Moldovanu, Philipp Strack, Roland Strausz, and Dezsö Szalay as well as seminar participants at Bonn, Yale and various conferences for valuable comments and suggestions. Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1 – 390838866 and through CRC TR 224 (Project B04) is gratefully acknowledged.

†Department of Economics, University of Bonn. email: sophie.kreutzkamp@uni-bonn.de

1 Introduction

Information plays a central role in many decision processes. Nevertheless, decision-makers often face a lack of information, preventing them from making good choices. Therefore, effective information transmission is crucial. A simple, but functional form of information sharing is costless, non-binding and unverifiable communication: cheap talk. Communication is key in our daily life: For instance, the success of big organizations depends on the interaction among its divisions. Another example is the communication strategy of central banks such as the Fed or ECB. Since it affects people's behavior and has thus an impact on the economy, it is of high interest and a common topic in the news media.

This paper investigates a communication game between a decision-maker and a biased expert with learning. In many situations, decision-makers consult experts not because they are omniscient, but because they can gather the desired information more easily by means of their prior knowledge or skills. For example, a firm who contemplates launching a new product might contact a market researcher, who then conducts a survey among potential customers in order to evaluate the product's expected profitability.

Even though the decision-maker delegates the information acquisition to the expert, she typically has an incentive to monitor the learning process in order to increase her benefits of the advice. In the above firm-researcher application, for example, the firm might want to have a look at the survey's questionnaire in advance to ensure that its outcome provides valuable information for the firm.

In this paper, I consider a sender-receiver game à la Crawford and Sobel (1982)[1] with endogenous learning: A sender gathers private, costly information about an unknown state by publicly[2] choosing a statistical experiment à la Blackwell (1953). Then, he communicates via a cheap-talk message with a receiver who takes an action affecting both agents' payoff.

There are two reasons why the receiver might remain partially uninformed: First, there might be a conflict of interest between the two agents, preventing the sender from sharing all his information with the receiver. Second, the sender may not acquire all information if this is too costly for him. To better understand the overall effect and interplay of these two forces, the paper examines experiments that are optimal, meaning implementable in equilibrium and Pareto efficient among all implementable experiments. There typically exist multiple optimal experiments because the two agents asymmetrically benefit from information: While the receiver is best off from the most informative experiments, the sender also has to bear the cost of learning. If more informative experiments incur higher costs, the sender may be

[1] They study a cheap-talk game with a perfectly informed sender and an uninformed receiver.

[2] Hence, the receiver observes which experiment the sender selects, but not its outcome.

better off from experiments providing limited information content.

My model is a general version of Crawford and Sobel (1982) going beyond single-peaked and single-crossing preferences. Apart from usual continuity assumptions on utility and cost functions, and compactness assumptions on state and action spaces, I impose a monotonicity condition on the cost function, ensuring that the costs of an experiment are proportional to its Blackwell informativeness[3]. Cost functions commonly used in the literature on rational inattention, such as the entropy-based cost function, satisfy this requirement. Also, free learning, i.e., zero cost, is a special case of the model.

I begin with a result that simplifies the identification of the set of optimal experiments: It is without loss of generality to restrict attention to equilibria in which the sender tells the truth, meaning that he fully reveals all information he learns to the receiver. The intuition for this is straightforward: Since the receiver has the authority over decision making, the sender does not directly benefit from any information which he does not transmit to the receiver. Therefore, it is more efficient if he only gathers information that he actually wants to forward to the receiver due to the cost caused by information gathering. In addition to that, the paper provides an existence theorem for optimal equilibria in the general framework.

Having too much information can be harmful. To see this most clearly, consider the zero-cost case. Even if the sender has the option to choose a most informative[4] experiment, he typically does not do so in an optimal equilibrium. Since the sender's message is not verifiable to the receiver, and the sender cannot commit to a communication strategy ex ante, the sender can transmit information to the receiver more credibly if he himself knows less: With less information, the sender has less incentives to misreport.

I estimate the efficiency of the communication game's decision-making process by comparing the optimal equilibrium outcomes to best feasible outcomes, i.e., Pareto efficient outcomes among all feasible ones. What are feasible outcomes? Suppose a social planner chooses an experiment in lieu of the sender and takes an action instead of the receiver based on the acquired information. An outcome that can be implemented this way is feasible.

A common concern about cheap talk is that the lack of verifiability and commitment implies efficiency losses. It seems natural that frictions arise in the communication game if the two agents have unequal interests. But where exactly do these inefficiencies come from? Under certain regularity conditions, I show that a best feasible outcome is implementable if and only if there is a unique best feasible outcome. So if there is no disagreement about what is the best outcome, the two agents manage to coordinate in equilibrium to achieve

[3]An experiment is more informative in the sense of Blackwell (1953) than another experiment if the latter is a Blackwell garbling of the former.

[4]An experiment is most informative if no other available experiment is Blackwell more informative.

this outcome. On the other hand, if the agents prefer different feasible outcomes, cheap talk involves inefficiencies. Consequently, different preferences per se do not lead to inefficiencies. Only if there is disagreement about best possible outcomes, frictions are inevitable under cheap talk. This finding provides an explanation why cheap talk is prevalent in real life.

Under posterior separable cost[5], the cheap-talk problem is solved by an experiment that is a convex combination of two experiments whose outcomes satisfy an analogous condition as the convex independence condition of Crémer and McLean (1988, p.1251). Such experiments are optimal because they produce a certain expedient amount of information using the least possible number of different outcomes. If the state space is finite, the number of outcomes of an optimal experiment is bounded by twice the number of states. In standard cheap-talk settings with a fully informed sender, optimal outcomes are typically not implementable if the number of messages is bounded by (twice) the number of states.

With posterior-mean preferences[6], optimal experiments are so-called bi-pooling policies[7]. These experiments induce a monotone partition[8] of the state space such that each outcome is associated with exactly one element of the partition, and at most two different outcomes belong to the same partitioning element. Similar results have been derived in the literature on Bayesian persuasion. This is an interesting observation on the relation between cheap talk and persuasion: While the two problems do not necessarily admit the same solution[9], they are equivalent in the sense that they admit a solution within the same class of experiments.

Finally, I apply the bi-pooling result to the classical setting with quadratic preferences, a uniformly distributed state, an additive sender-bias and zero cost in order to solve for the optimal experiment. I find that monotone partitions are not always optimal. If the bias is sufficiently small, it is either a uniform partition[10], a non-uniform monotone partition or a bi-pooling policy with exactly two bi-pooling elements. This is an interesting finding because optimality of monotone partitions is prevalent in the uniform-quadratic case of related cheap-talk games (see Crawford and Sobel (1982), Pei (2015) or Ivanov (2010)). Comparative statics show that the outcome of my model is closest to the Pareto frontier of the set of feasible

[5]This is a standard assumption in the literature on rational inattention (cf. Maćkowiak, Matějka, and Wiederholt (2018) or Matějka and McKay (2015)).

[6]Such preferences do not depend on the whole conditional distribution of the state, but only on its expected value at any stage of the game.

[7]This term has been introduced by Arieli, Babichenko and Smorodinsky (2020).

[8]A monotone partition is a partition into convex subsets.

[9]In general, optimal cheap talk and optimal Bayesian persuasion are not outcome-equivalent due to the commitment constraint that is imposed under cheap talk, but absent in the persuasion problem. Lipnowski (2020) derives outcome-equivalence of the two problems under certain conditions (namely finiteness of the action space and continuity of the sender's value function), which render the commitment constraint non-binding.

[10]A uniform partition is a monotone partition into equally sized subsets.

outcomes as compared to the related cheap-talk models (i.e., perfect information, covert learning, overt learning restricted to experiments being monotone partitions, mediation, etc.). This result suggests that overt and flexible[11] information acquisition is valuable. The main takeaway of this is that agents should use these factors if they are available: In terms of the firm-researcher application, this means, for instance, that the firm should insist on getting access to the survey that the researcher plans to conduct.

The cheap-talk model with overt information acquisition can also be interpreted as a persuasion model with partial commitment: The sender can commit to the experiment he chooses, but not to the message he sends. So this case is in between full commitment, i.e, Bayesian persuasion (cf. Kamenica and Gentzkow (2011)), and no commitment. A natural question to ask in this context is whether the outcome of partial commitment is closer to no or full commitment. That is, is commitment with respect to the information choice valuable on its own, or only in conjunction with commitment with respect to the communication strategy? I provide comparative statics for the uniform quadratic model suggesting that commitment with respect to the information choice has indeed an intrinsic value.

The paper is structured as follows: Section 2 introduces the model and equilibrium concept. Section 3 contains the recommendation principle, the existence proof, and the comparison of best implementable versus best feasible outcomes. In Section 4, I derive the optimality of independent outcomes under posterior separable cost. Section 5 covers the bi-pooling result for posterior-mean preferences. Section 6 contains the analysis of the uniform-quadratic model, and Section 7 discusses model variants and provides comparative statics. Finally, Section 8 concludes. Proofs are deferred to the appendix.

Related Literature. This paper contributes to the literature on strategic information transmission. Since the seminal work by Crawford and Sobel (1982), the embedment of information acquisition into communication games has been an active subfield of this literature.

Pei (2015) discusses a cheap-talk game with costly, covert information acquisition. For the uniform-quadratic case, he shows that monotone partitions are chosen in equilibrium if the set of available experiments is the set of all finite partitions. In the limit case when costs converge to zero, the equilibrium outcomes of this model correspond to the equilibrium outcomes of the model without information acquisition by Crawford and Sobel (1982). This highlights the difference between overt and covert learning: The option for endogenous learning does not necessarily improve information transmission. It may be crucial that the

[11] Flexible information acquisition means that the sender can choose arbitrary experiments. It is a standard assumption in the literature on information design (cf. Kamenica and Gentzkow (2011)) and coordination games (cf. Yang (2015)).

receiver can observe which information the sender acquires.

Ivanov (2010) examines a model with costless information procurement in which the receiver selects the experiment instead of the sender. He derives optimal monotone partitions for the uniform quadratic model. The equilibrium outcomes of this model are equal to those of my model. Hence, it does not matter whether the sender or the receiver chooses the experiment, as long as this choice is publicly known.

Deimen and Szalay (2019) investigate overt information acquisition in a communication game in which the conflict of interest between the sender and the receiver evolves endogenously through the learning process. To make information transmission the most effective, the sender acquires information that is equally beneficial for both agents because this reduces the endogenous bias and increases cooperation between both parties. This relates to the intuition behind the implementability of best feasible outcomes in my model: Information transmission is best if the two agents are not biased towards different best feasible outcomes because this allows them to cooperate.

This paper adds to the literature on strategic information transmission by investigating a general setting beyond single-peaked and single-crossing preferences or the uniform-quadratic case.

There are several devices other than endogenous learning that potentially improve the outcome of cheap-talk games: multiple rounds of communication, mediation or delegation (see Aumann and Hart (2003), Krishna and Morgan (2004), Goltsman, Hörner, Pavlov and Squintani (2009), and Dessein (2002)).

My work is also related to the literature on Bayesian persuasion as it is a model with partial commitment. Following the pioneering work by Kamenica and Gentzkow (2011)[12], the main tool of this literature is the geometric characterization of solutions via concave closures of the agent's valuation function.[13] The concavification approach is also a useful technique to study cheap-talk games in which the sender has state-independent preferences (see Lipnowski and Ravid (2020)). With state-dependent sender preferences, however, the characterization of concave closures turns out to become an intractable task in general because of additional sender incentive constraints to be incorporated.

An exception is Lyu and Suen (2022). They study a cheap-talk game with overt information acquisition using an extended version of the concavification approach respecting incentive compatibility of the sender. They focus on a binary state space and the zero cost case. Feasibility of their concavification approach crucially hinges on the binary-state

[12]They study a communication game with endogenous information acquisition in which the sender can commit to both the information and the message choice.

[13]See Gentzkow and Kamenica (2014), Gentzkow and Kamenica (2016), Ely (2017) for further applications of the concavification approach.

assumption.[14] Allowing for richer state spaces requires a different approach, namely the characterization of extreme points: Optimal experiments are extreme points of the set of available experiments.

Last but not least, this paper relates to the literature on extreme points and majorization. Since the fundamental work by Kleiner, Moldovanu and Strack (2021) on extreme points of monotonic functions under majorization constraints, this technique has become increasingly popular in the persuasion literature. Kleiner et al. (2021) and Arieli et al. (2020) show that optimal experiments in the standard persuasion problem are bi-pooling policies. For further applications of the extreme-point approach in the persuasion literature, see Matysková and Montes (2021) or Candogan and Strack (2021), for instance.

As for the concavification approach, the incorporation of sender incentive compatibility constraints is a crucial step when applying the extreme-point method to cheap-talk settings. This is considerably easier in the extreme-point approach. Hence, the extreme-point method is a useful tool to determine optimal experiments in the cheap-talk model beyond the binary state case.

2 Model

2.1 Setting

There is a sender S (he), a receiver R (she), and a state of the world ω. The state space Ω is a compact and convex subset of $\mathbb{R}$[15], and the state follows a distribution $\mu_0 \in \triangle\Omega$,[16] which is common knowledge. The state realization ω is initially unknown to both players.

The game proceeds as follows: After the realization of the state, the sender publicly chooses an experiment π, i.e., a distribution over posterior beliefs of the state $\mu \in \triangle\Omega$, at cost $c(\pi)$. The set of available experiments Π is a compact subset of $\triangle(\triangle\Omega)$. The cost function c is continuous on the convex hull of Π. Next, the sender transmits a message $m \in \triangle\Omega$ to the receiver. Then, the receiver takes an action a from a compact and convex action space $A \subset \mathbb{R}$. Payoffs are determined by the agents' von Neumann-Morgenstern utility functions $u_S : A \times \Omega \to \mathbb{R}$ and $u_R : A \times \Omega \to \mathbb{R}$, which are continuous on $A \times \Omega$.

Figure 1 summarizes the course of the game: I impose additional mild assumptions on the cost of information acquisition and the set of available experiments:

[14] With a binary state space, incentive compatibility reduces to a single inequality constraint.

[15] I assume that the state is real-valued for tractability. All results in Section 2.2 and 4 also apply when $\Omega \subset \mathbb{R}^n$.

[16] For a compact metrizable space X, let $\triangle X$ denote the set of distributions over X, endowed with the topology of weak convergence. For each $\chi \in \triangle X$, let $\mathrm{supp}(\chi)$ denote the support of χ.

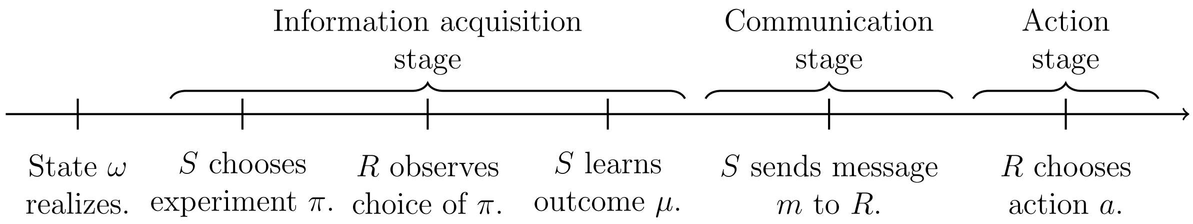

Figure 1: Timeline of the game

Assumption 1. The cost function $c : \mathcal{I} \rightarrow \mathbb{R}$ has the following properties:

1. (*Zero-normalization*): $c(\pi) = 0$ if $\text{supp}(\pi) = \{\mu_0\}$

2. (*Monotonicity*): If π' is a Blackwell garbling[17] of π, then $c(\pi) \geq c(\pi')$.

Zero-normalization requires that learning nothing, i.e., choosing the uninformative experiment π^0 whose sole outcome is the prior distribution, is costless. Monotonicity captures the idea that information which is more precise in the sense of Blackwell should be costlier: Since the information provided by a Blackwell garbling π' can also be generated by the original experiment π through appropriate mixing over its outcome, the former experiment should be less costly than the latter one.

Assumption 2. If $\pi \in \Pi$ and π' is a Blackwell garbling of π, then $\pi' \in \Pi$.

This means that the set of available experiments is sufficiently rich in the sense that for any available experiment, the sender could also choose an arbitrary Blackwell garbling of it. Intuitively, the sender has enough flexibility in acquiring information about the state: He can take any desired experiment — up to a certain level of precision.

2.2 Equilibrium Characterization

A sender strategy $(\sigma_{\mathcal{I}}, \sigma_{\mathcal{M}})$ consists of an *information rule* $\sigma_{\mathcal{I}} \in \triangle\Pi$ and a *communication rule* $\sigma_{\mathcal{M}} : \Pi \times \triangle\Omega \rightarrow \triangle(\triangle\Omega)$. A receiver strategy is an *action rule* $\sigma_{\mathcal{A}} : \Pi \times \triangle\Omega \rightarrow \triangle A$. The sender's belief after choosing the experiment π and observing its outcome μ is μ. The receiver has a belief function $\mu_R : \Pi \times m \rightarrow \triangle\Omega$ with $\mu_R(\cdot|\pi, m)$ indicating her belief after observing the sender's choice of the experiment π and receiving his message m.

As is standard in the cheap-talk literature, I consider perfect Bayesian equilibria.[18]

[17] Formally, π' is a Blackwell garbling of π if for all $\mu' \in \text{supp}(\pi')$, there exists some $\lambda_{\mu'} \in \triangle(\triangle\Omega)$ such that $\mu' = \int_{\text{supp}(\pi)} \mu \, d\lambda_{\mu'}(\mu)$.

[18] Restricting attention to perfect Bayesian equilibria is just for tractability. In fact, for any perfect Bayesian equilibrium, there exists an outcome-equivalent Bayesian equilibrium. Hence, all results in this paper also apply to the latter equilibrium concept.

Definition 1. A *perfect Bayesian equilibrium* (PBE) $E = \{((\sigma_{\mathcal{I}}, \sigma_{\mathcal{M}}), \sigma_{\mathcal{A}}), \mu_R)\}$ consists of a strategy profile $((\sigma_{\mathcal{I}}, \sigma_{\mathcal{M}}), \sigma_{\mathcal{A}})$ with beliefs μ_R such that

1. the information rule $\sigma_{\mathcal{I}}$ is optimal given $(\sigma_{\mathcal{M}}, \sigma_{\mathcal{A}})$, that is,

$$\text{supp}(\sigma_{\mathcal{I}}) \subseteq \arg\max_{\pi \in \Pi} \int\limits_{\triangle\Omega} \int\limits_{\triangle\Omega} \int\limits_{\Omega} \int\limits_{A} u_S(a, \omega) \; d\sigma_{\mathcal{A}}(a|I, m) \; d\mu(\omega) \; d\sigma_{\mathcal{M}}(m|I, \mu) \; d\pi(\mu)$$

2. the communication rule $\sigma_{\mathcal{M}}$ is optimal given $\sigma_{\mathcal{A}}$, i.e., for all $(\pi, \mu) \in \Pi \times \triangle\Omega$,

$$\text{supp}(\sigma_{\mathcal{M}}(\pi, \mu)) \subseteq \arg\max_{m \in \triangle\Omega} \int\limits_{\Omega} \int\limits_{A} u_S(a, \omega) \; d\sigma_{\mathcal{A}}(a|\pi, m) \; d\mu(\omega)$$

3. the action rule $\sigma_{\mathcal{A}}$ is optimal given μ_R, that is, for all $(\pi, m) \in \Pi \times \triangle\Omega$,

$$\text{supp}(\sigma_{\mathcal{A}}(\pi, m)) \subseteq \arg\max_{a \in A} \int\limits_{\Omega} u_R(a, \omega) \; d\mu_R(\omega|\pi, m)$$

4. the receiver's beliefs are derived from Bayes' rule, whenever possible.

The ex-ante expected payoffs of an equilibrium E are

$$\mathcal{U}_S(E) \equiv \int\limits_{\Pi} \int\limits_{\triangle\Omega} \int\limits_{\triangle\Omega} \int\limits_{\Omega} \int\limits_{A} u_S(a, \omega) \; d\sigma_{\mathcal{A}}(a|\pi, m) \; d\mu(\omega) \; d\sigma_{\mathcal{M}}(m|\pi, \mu) \; d\pi(\mu) - c(\pi) \; d\sigma_{\mathcal{I}}(\pi)$$

for the sender and

$$\mathcal{U}_R(E) \equiv \int\limits_{\Pi} \int\limits_{\triangle\Omega} \int\limits_{\triangle\Omega} \int\limits_{\Omega} \int\limits_{A} u_R(a, \omega) \; d\sigma_{\mathcal{A}}(a|\pi, m) \; d\mu_R(\omega|\pi, m) \; d\sigma_{\mathcal{M}}(m|\pi, \mu) \; d\pi(\mu) \; d\sigma_{\mathcal{I}}(\pi)$$

for the receiver.

2.3 Babbling Equilibria

Existence of a PBE is guaranteed because it is always possible to construct an equilibrium in which the sender babbles, i.e, transmits no information, and the receiver chooses an ex-ante optimal action.

Definition 2. A PBE is a *babbling* equilibrium if $\mu_R(\cdot|\pi, m) = \mu_0$ for all $\pi, m \in \triangle\Omega$. The set of babbling equilibria is denoted by $\mathcal{E}_0$.

When the sender babbles, there is no belief-updating from the receiver's perspective because her on-path beliefs coincide with the prior distribution. Hence, she does not learn anything.

Proposition 1. *There exists an equilibrium. Moreover, $\mathcal{E}_0$ is nonempty.*

While equilibrium existence holds true, uniqueness generically fails. Indeed, there can even be multiple babbling equilibria unless the receiver's best response to the prior belief is unique: One can sustain her choosing any priorly optimal action in equilibrium.

Nevertheless, the payoffs of babbling equilibria serve as a lower bound for the agents' ex-ante expected payoffs of any PBE.

Proposition 2. *For any PBE E, it holds that*

$$\mathcal{U}_S(E) \geq \inf_{E' \in \mathcal{E}_0} \mathcal{U}_S(E') \quad \textit{and} \quad \mathcal{U}_R(E) \geq \sup_{E' \in \mathcal{E}_0} \mathcal{U}_R(E').$$

Both agents' equilibrium payoffs are bounded below by some babbling outcome: The receiver can always choose an action that is optimal for her given the prior distribution, thus securing the payoff of any arbitrary babbling equilibrium. On the contrary, the sender's payoff is not necessarily the same across all babbling equilibria.[19]. Nevertheless, he can secure the lowest babbling payoff by choosing the uninformative experiment π^0.

3 Best Implementable Outcomes

This section studies optimal equilibria. I will use the following terminology:

Definition 3. A payoff profile $(\mathcal{U}_R, \mathcal{U}_S)$ is *implementable* if there is some PBE E such that $\mathcal{U}_i = \mathcal{U}_i(E)$ for all $i \in \{S, R\}$. Any such PBE E is said to *generate* $(\mathcal{U}_R, \mathcal{U}_S)$.

A payoff profile $(\mathcal{U}_R, \mathcal{U}_S)$ *dominates* another one $(\mathcal{U}'_R, \mathcal{U}'_S)$ if $\mathcal{U}_i \geq \mathcal{U}'_i$ for all i with at least one inequality holding strictly. A PBE *dominates* another PBE if the payoff profile generated by the former PBE dominates the payoff profile generated by the latter one.

A payoff profile is *best implementable* if it is implementable and not dominated by any other implementable payoff profile. Any PBE generating a best implementable payoff profile is called *optimal*. An experiment π is called *optimal* if $\pi \in \Pi$ of some optimal PBE.

Best implementability means a payoff profile can be attained in a PBE, and is undominated by the ex-ante expected payoffs of all other PBE in the sense of the Pareto criterion.[20]

[19] This can be case if $\arg\max_{a \in A} \int_\Omega u_R(a, \omega)\, d\mu_0(\omega)$ is not a singleton.

[20] The literature on Bayesian persuasion and cheap-talk games usually focuses on sender- or receiver-

3.1 Fully Revealing Equilibria

This section provides a useful tool towards determining the set of best implementable payoff profiles: It is without loss of generality to focus on equilibria in which the sender fully reveals his acquired information to the receiver.

Definition 4. A PBE is a *fully revealing* equilibrium if $\mu_R(\cdot|\pi, m) = \mu$ for every $m \in \text{supp}(\sigma_{\mathcal{M}}(\pi, \mu))$, all $\mu \in \text{supp}(\pi)$ and all $\pi \in \text{supp}(\sigma_{\mathcal{I}})$. The communication rule in a fully revealing PBE is denoted by $\sigma_{\mathcal{M}}^{FR}$.

In a fully revealing equilibrium, the sender transmits all his information to the receiver such that their equilibrium beliefs coincide.

The argument of the recommendation principle involves two steps: First, it suffices to focus on equilibria where the sender uses a pure strategy in the information acquisition stage for the purpose of identifying optimal experiments.

Definition 5. An information rule $\sigma_{\mathcal{I}}$ is called *pure-strategy* if $|\text{supp}(\sigma_{\mathcal{I}})| = 1$.

Under a pure-strategy information rule, the sender chooses one experiment with probability 1 in equilibrium.

Lemma 1. *Any best implementable payoff profile is generated by a PBE with a pure-strategy information rule. Moreover, any optimal experiment is chosen by the sender in some PBE with a pure-strategy information rule.*

The proof of this lemma is constructive: For any optimal experiment π that the sender chooses in some PBE E with a mixed-strategy information rule, one can find a payoff-equivalent PBE E_π in which the sender takes π only. The sender gets the same ex-ante expected payoff in the two PBE because he is indifferent between all experiments he chooses in the PBE E. The receiver's payoff in the PBE E is the average over her payoffs in the PBE

optimal equilibria (cf. Kamenica and Gentzkow (2011), Gentzkow and Kamenica (2014) or Ambrus, Azevedo and Kamada (2013)). So why is it worth studying all best implementable payoff profiles? If the receiver's best response in the action stage is guaranteed to be unique, the sender-optimal equilibrium is most reasonable: Supposing that the two players try to coordinate on mutually most beneficial outcomes even off the equilibrium path in case the sender mistakenly chooses an experiment not prescribed by the information rule, the model prediction is exactly the sender-optimal outcome. This argumentation breaks down if the receiver is indifferent between different actions: She could announce to choose an action being least favorable to the sender among those between which she is indifferent in case the sender deviates to an experiment off path. This is a credible threat. Hence, the model prediction is no longer the sender-optimum necessarily under the assumption of off-path coordination. It is not clear what the actual model prediction is, but it is a best implementable outcome, most likely in between the sender- and receiver-optimum. Analyzing all best implementable outcomes ensures to find the characteristics of the optimal experiment under this equilibrium selection criterion.

E_π in which the sender chooses one of the experiments in the support of the mixed strategy with probability 1. If the receiver does not obtain the same payoff in all latter PBE, there is one in which she gets a higher payoff than in the PBE E. So the latter one cannot generate a best implementable payoff profile.

Second, fully revealing PBE dominate other equilibria.

Lemma 2. *Any PBE with a pure-strategy information rule is dominated by a fully revealing PBE with a pure-strategy information rule.*

Since information is costly, it is inefficient if the sender acquires more information than he is actually willing to transmit to the receiver. Any additional private information is of no use to him as he can influence her action choice only via the amount of information that he reveals to her.

Endowed with these two lemmata, the recommendation principle can be established:

Theorem 1. *Any best implementable payoff profile is generated by some fully revealing PBE with a pure-strategy information rule.*

Next, I prove the existence of best implementable payoff profiles.

Theorem 2. *There exists a best implementable payoff profile. Moreover, the set of best implementable payoff profiles is compact.*

Existence follows from an application of Berge's maximum theorem.[21] Assumption 1.2 and compactness of Π ensure that the corresponding optimization problem has a continuous objective function over a compact set. Let

$$\tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_{\mathcal{A}}) \equiv \int_{\triangle\Omega} \int_{\Omega} \int_{A} u_S(a, \omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\pi, \mu)\ d\mu(\omega)\ d\pi(\mu) - c(\pi)$$

$$\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) \equiv \int_{\triangle\Omega} \int_{\Omega} \int_{A} u_R(a, \omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu)\ d\mu(\omega)\ d\pi(\mu)$$

be the agents' ex-ante expected payoff if the sender chooses experiment π, communicates its outcome fully to the receiver, and the receiver chooses her action according to the strategy $\tilde{\sigma}_{\mathcal{A}} : \triangle\Omega \to \triangle A$. A payoff profile $(\mathcal{U}_S^*, \mathcal{U}_R^*)$ is best implementable if and only if there exists some π^* and some $\tilde{\sigma}_{\mathcal{A}}^*$ with $\mathcal{U}_i^* = \tilde{\mathcal{U}}_i(\pi^*, \tilde{\sigma}_{\mathcal{A}}^*)$ for each $i \in \{S, R\}$ solving

$$\max_{(\pi, \tilde{\sigma}_{\mathcal{A}})} \quad \tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_{\mathcal{A}})$$

[21] See Aliprantis and Border (2006), p.570.

$$\text{s.t.} \quad \operatorname{supp}\left(\tilde{\sigma}_{\mathcal{A}}\left(\mu\right)\right) \subseteq \arg\max_{a \in A} \int_{\Omega} u_R\left(a, \omega\right) \, d\mu\left(\omega\right) \quad \text{for all } \mu \in \triangle\Omega \tag{1}$$

$$\begin{aligned} &\int_{\Omega} \int_A u_S(a, \omega) \, d\tilde{\sigma}_{\mathcal{A}}\left(a|\mu\right) \, d\mu(\omega) \\ &\qquad \geq \int_{\Omega} \int_A u_S(a, \omega) \, d\tilde{\sigma}_{\mathcal{A}}\left(a|\mu'\right) \, d\mu(\omega) \end{aligned} \quad \text{for all } \mu, \mu' \in \operatorname{supp}\left(\pi\right) \tag{2}$$

$$\tilde{\mathcal{U}}_S\left(\pi, \tilde{\sigma}_{\mathcal{A}}\right) \geq \mathcal{U}_S^0 \tag{3}$$

$$\tilde{\mathcal{U}}_R\left(\pi, \tilde{\sigma}_{\mathcal{A}}\right) \geq \mathcal{U}_R \tag{4}$$

for some $\mathcal{U}_R \geq \mathcal{U}_R^0$, where $\mathcal{U}_i^0$ denotes agent i's minimal ex-ante expected payoff in a PBE.[22] This maximization problem determines the highest possible ex-ante expected payoff the sender can obtain in a fully revealing PBE with a pure-strategy information rule provided that the receiver's ex-ante expected payoff does not fall below a threshold value $\mathcal{U}_R$: The constraints (1) guarantee the optimality of the receiver's action rule. The inequality constraints (2) ensure that the sender has an incentive to report the experiment's observed outcome always truthfully. The participation constraint (3) makes sure that the sender is willing to acquire costly information, that is, his ex-ante expected payoff must not be lower than the minimal ex-ante expected payoff he would obtain if he acquired no information. Condition (4) makes sure that the receiver's ex-ante expected payoff is at least $\mathcal{U}_R$.

The equilibrium that implements the best implementable payoff profile $(\mathcal{U}_S^*, \mathcal{U}_R^*)$ looks as follows: It is a fully revealing PBE with a pure-strategy information rule in which the sender chooses the experiment π^*. For any other experiment off the equilibrium path, the agents agree on a babbling outcome which yields the minimal ex-ante expected payoff to the sender that he can obtain in a babbling PBE.

3.2 Best Implementable versus Best Feasible Outcomes

The set of best implementable payoff profiles depends on both the model's parameters (utility functions, set of available actions/experiments, etc.) and the game's structure (information acquisition and transmission by the sender, decision authority of the receiver): To better understand the interplay of these two forces, I compare the best possible outcomes that do not depend on the game's structure to the best implementable payoff profiles.

Suppose there is a social planner, i.e., a third neutral party, who chooses an experiment in lieu of the sender according to a mixed strategy $\bar{\sigma}_{\mathcal{I}} \in \triangle\Pi$ and, based on the experiment's outcome, an action in place of the receiver according to the strategy $\bar{\sigma}_{\mathcal{A}} : \Pi \times \triangle\Omega \to \triangle A$.

[22] According to Proposition 2, $\mathcal{U}_S^0 = \inf_{E' \in \mathcal{E}_0} \mathcal{U}_S\left(E'\right)$ and $\mathcal{U}_R^0 = \sup_{E' \in \mathcal{E}_0} \mathcal{U}_S\left(E'\right)$.

Payoffs that can be generated via such an intervention by a social planner are of the form

$$\bar{\mathcal{U}}_S\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right) \equiv \int_{\Pi}\int_{\triangle\Omega}\int_{\Omega}\int_{A} u_S(a,\omega)\ d\bar{\sigma}_{\mathcal{A}}(a|\pi,\mu)\ d\mu(\omega)\ d\pi\left(\mu\right) - c\left(\pi\right)\ d\bar{\sigma}_{\mathcal{I}}(\pi)$$

$$\bar{\mathcal{U}}_R\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right) \equiv \int_{\Pi}\int_{\triangle\Omega}\int_{\Omega}\int_{A} u_R(a,\omega)\ d\bar{\sigma}_{\mathcal{A}}(a|\pi,\mu)\ d\mu(\omega)\ d\pi\left(\mu\right)\ d\bar{\sigma}_{\mathcal{I}}(\pi).$$

Definition 6. A payoff profile $(\mathcal{U}_S,\mathcal{U}_R)$ is *feasible* if there exist some $\bar{\sigma}_{\mathcal{I}}$ and some $\bar{\sigma}_{\mathcal{A}}$ such that $\mathcal{U}_i = \bar{\mathcal{U}}_i\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right)$ for all i. It is *best feasible* if it is feasible and dominates all other feasible payoff profiles.

Existence of best feasible payoff profiles is guaranteed:

Theorem 3. *There exists a best feasible payoff profile. Moreover, the set of best feasible payoff profiles is compact.*

The proof proceeds analogously to the proof of Theorem 2. Formally, a payoff profile $(\mathcal{U}^*_S,\mathcal{U}^*_R)$ is best feasible if and only if there exists some $\bar{\sigma}^*_{\mathcal{I}}$ and some $\bar{\sigma}^*_{\mathcal{A}}$ with $\mathcal{U}^*_i = \bar{\mathcal{U}}_i\left(\bar{\sigma}^*_{\mathcal{I}},\bar{\sigma}^*_{\mathcal{A}}\right)$ fr all i solving

$$\max_{(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}})} \quad \bar{\mathcal{U}}_S\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right)$$

$$\text{s.t.} \quad \bar{\mathcal{U}}_R\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right) \geq \mathcal{U}_R. \tag{5}$$

for some $\mathcal{U}_R \in \mathbb{R}$. While every implementable payoff profile is feasible, the converse does not hold true. So best implementable payoff profiles can, but need not, be best feasible. The next theorem states a sufficient condition for the equivalence of best implementable and best feasible outcomes:

Theorem 4. *Suppose there exists a unique best feasible payoff profile. Then, it is the unique best implementable one.*

Uniqueness of the best feasible payoff profile implies that the sender- and receiver-optimal feasible payoff profile coincide. So both players agree on what is the best attainable outcome. This closes the gap between best feasible and best implementable payoff profiles because the sender and the receiver have an incentive to coordinate in equilibrium to achieve the best feasible outcome.

Allowing for multiplicity annihilates the equivalence result:

Theorem 5. *Suppose there exist at least two best feasible payoff profiles. Then, a best feasible payoff profile* $(\mathcal{U}_S,\mathcal{U}_R)$ *is not best implementable* $\text{supp}(\sigma_{\mathcal{I}}) = \{\pi^{full}\}$ *for any* $(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}})$ *with* $\mathcal{U}_i = \bar{\mathcal{U}}_i\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right)$ *for each* i.

Consider first the receiver-optimal feasible payoff profile. Since it is feasible only with the fully informative experiment π^{full}, the sender would have an incentive to misreport in the communication stage to achieve the receiver-optimal feasible payoff profile, which is distinct by multiplicity. For all other best feasible payoff profiles, the receiver would have an incentive to choose another action rule after the sender fully revealed the fully informative outcome in order to generate the receiver-optimal feasible payoff profile instead.

To conclude, the frictions caused by information transmission from the acquirer to the decision-maker affect the best implementable outcomes beyond the model parameters whenever the sender and the receiver do not concur about the best achievable outcome, that is, whenever there are multiple best feasible payoff profiles.

The subsequent finding reinforces the idea that disagreement over optimal feasible outcomes leads to inefficiencies: It provides conditions under which uniqueness of the best feasible payoff profile is a necessary and sufficient condition for the equivalence of best implementability and best feasibility. Let $A^*(\omega, u) = \arg\max_{a \in A} u\,(a, \omega)$ be the set of optimal actions if the state is ω given the utility function is u.

Corollary 1. *If $A^*(\omega, \alpha u_S + (1-\alpha)u_R) \cap A^*(\omega', \alpha u_S + (1-\alpha)u_R) = \emptyset$ for all $\alpha \in [0,1]$ and all $\omega \neq \omega'$, and $c = 0$, then a best feasible payoff profile is best implementable if and only if it is unique.*

Remarkably, the set of best feasible and best implementable payoff profiles either coincide or are disjoint — depending on whether the former one is a singleton or not. In that sense, the potential frictions introduced by the cheap-talk game influence all best implementable payoff profiles equally. The section concludes with an illustrating example:

Example 1. Consider the uniform-quadratic case à la Crawford and Sobel (1982) with an additive bias $b = 0.126$ and zero cost. Note that $A^*(\omega, \alpha u_S + (1-\alpha)u_R) = \omega + (1-\alpha)b$ for all $\alpha, \omega \in [0,1]$. There is a continuum of best feasible payoff profiles and a unique best implementable payoff profile. Showing the receiver's ex-ante expected payoff on the x-axis and the sender's ex-ante expected payoff on the y-axis, Figure 2 compares the best feasible payoff profiles (the points on the blue line) to the best implementable payoff profile $(\mathcal{U}_S^*, \mathcal{U}_R^*)$, the red point. Since there are multiple best feasible payoff profiles, the red point cannot lie on the blue curve by Corollary 1.

4 Posterior Separable Cost

This section analyses optimal experiments for posterior separable cost functions.

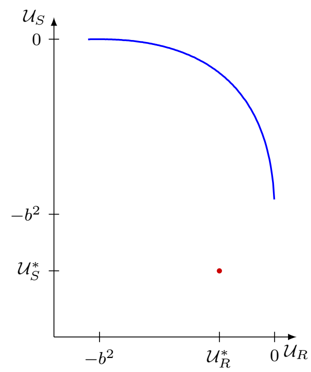

Figure 2: Payoff comparison in the uniform-quadratic setting with bias $b = 0.126$

Definition 7. A cost function is *posterior separable* if there exists a convex function $k : \triangle\Omega \to \mathbb{R}$ so that $c(\pi) = -k(\mu_0) + \int_{\triangle\Omega} k(\mu)\ d\pi(\mu)$.

Any posterior separable cost function satisfies Assumption 1. Posterior separability is the standard assumption on the cost structure in the literature on rational inattention (cf. Caplin and Dean (2013) or Caplin, Dean and Leahy (2019)). Examples include the entropy-based cost as well as zero cost. For the sake of tractability, I suppose that the sender can choose any arbitrary experiment — a condition being in line with Assumption 2. Hence, all results derived in the previous sections remain applicable under the following assumption:

Assumption 3. The cost function c is posterior separable and $\Pi = \triangle(\triangle\Omega)$.

Under posterior separable cost and flexible learning, optimal experiments can be derived from the following class of experiments:

Definition 8. An experiment π exhibits *weakly independent* outcomes if

$$\mu = \int_{\triangle\Omega} \nu\ d\lambda(\nu) \quad \Rightarrow \quad \lambda = \delta_\mu \text{ a.e.,}$$

where δ_μ is the Dirac measure at μ, for all $\mu \in \Pi_I$, except a set of measure zero.

An experiment π exhibits *strongly independent* outcomes if for all $\lambda_1, \lambda_2 \in \triangle(\triangle\Omega)$ with $p\lambda_i + (1-p)\lambda_i' = \pi$ for some $p \in (0,1)$ and some $\lambda_i' \in \triangle(\triangle\Omega)$ for $i \in \{1,2\}$, it holds that

$$\int_{\triangle\Omega} \nu\ d\lambda_1(\nu) = \int_{\triangle\Omega} \nu\ d\lambda_2(\nu) \quad \Rightarrow \quad \lambda_1 = \lambda_2 \text{ a.e..}$$

Outcomes are weakly independent if no outcome can be replicated as a convex combination of other outcomes. Strong independence means that no convex combination of outcomes can be represented by another convex combination of outcomes.[23] Strong independence implies weak independence, but not vice versa.

Every optimal experiment is outcome-equivalent to an experiment being a convex combination of two experiments with strongly independent outcomes.

Theorem 6. *Take any best implementable payoff profile* $(\mathcal{U}_S, \mathcal{U}_R)$.

(i) There exist some $\pi, \pi \in \Pi'$ *with strongly independent outcomes and some* $\alpha \in [0, 1]$ *such that* $(\mathcal{U}_S, \mathcal{U}_R)$ *is generated by a fully revealing PBE with a pure-strategy information rule in which the sender chooses the experiment* $\alpha\pi + (1 - \alpha)\pi'$.

(ii) If $c = 0$, $(\mathcal{U}_S, \mathcal{U}_R)$ *is generated by a fully revealing PBE with a pure-strategy information rule where the sender chooses an experiment with weakly independent outcome.*

Any experiment π without strongly independent outcomes corresponds to a convex combination of two experiments π_1 and π_2. If either one, say π_2, does not generate strongly independent outcomes, it is again a convex combination of some experiments π_3 and π_4. But then, there is some convex combination either of π_1 and π_3, π_1 and π_4, or π_3 and π_4 that is implementable and dominates π. The last fact follows from linearity of payoffs, and is illustrated in Figure 3: The 2-simplex shows all convex combinations of the experiments

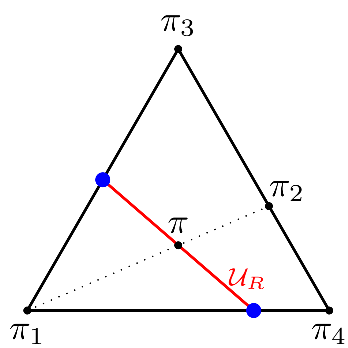

Figure 3: Payoff comparison in a simplex of three experiments

π_1, π_3 and π_4. The original experiment π lies in the interior of that simplex, on the dotted line between π_1 and π_2 as it is a convex combination of those two points. Similarly, π_2 lies on the line between π_3 and π_4. The red line represents all experiments that yield the same ex-ante expected payoff $\mathcal{U}_R$ to the receiver as the original experiment π. By linearity, the

[23] Geometrically, strong independence means that the set of outcomes $\mathrm{supp}(\pi)$ corresponds to the set of extreme points of a simplex within the space $\triangle(\triangle\Omega)$. That is, the experiment's outcomes span a simplex in the space $\triangle(\triangle\Omega)$.

sender's ex-ante expected payoff is monotone along this line so that the optimum is attained at one of its blue intersections with the edges of the simplex, which corresponds to a convex combination of only two of the three experiments π_1, π_3 and π_4.

Applying Theorem 6 to settings with a finite state space provides a bound on the number of different outcomes of optimal experiments: It suffices to consider experiments whose number of outcomes do not exceed twice the number of states. Especially for environments with a small state space, this result can be quite helpful towards finding an optimal experiment as the result rules out a huge number of potential candidates for such an experiment.

Corollary 2. *Suppose* $|\Omega| < \infty$. *For any best implementable payoff profile* $(\mathcal{U}_S, \mathcal{U}_R)$,

(i) there exists some π *with* $\text{supp}(\pi) \leq 2|\Omega|$ *such that* $(\mathcal{U}_S, \mathcal{U}_R)$ *is generated by a fully revealing PBE with a pure-strategy information rule in which the sender chooses* I.

(ii) If $c = 0$ *and* $|\Omega| = 2$, *there is a* π *with* $\text{supp}(\pi) = 2$ *so that* $(\mathcal{U}_S, \mathcal{U}_R)$ *is generated by a fully revealing PBE with a pure-strategy information rule where the sender takes* π.

Under the finiteness assumption, any set of $|\Omega| + 1$ outcomes is linearly dependent.

Remark 1. The finding of Corollary 2 is a distinguishing feature of cheap-talk models with endogenous learning: In cheap-talk settings, where the sender is fully informed about the state of nature, optimal outcomes cannot be implemented in general if the number of messages may not exceed the number of states, as long as the latter one is finite.

5 Posterior-mean preferences

This section studies optimal equilibria for posterior-mean preferences. All assumptions made in previous sections remain valid.

Definition 9. A von Neumann-Morgenstern utility function u is *partially separable* if there are continuous functions $u_1 : A \to \mathbb{R}$, $u_2 : A \to \mathbb{R}$ and $u_3 : \Omega \to \mathbb{R}$ such that

$$u(a, \omega) = u_1(a) + u_2(a) \cdot \omega + u_3(\omega).$$

Partially separable utility functions are additively separable in a and ω, except for a component that is quasi-linear in ω. This is a broad class of preferences containing those frequently studied in the literature on cheap-talk and Bayesian persuasion[24], namely quadratic

[24]For the use of quadratic preferences, see Crawford and Sobel (1982), Krishna and Morgan (2004), Ivanov (2010), or Pei (2015), for instance. For quasi-linear utility, see Gentzkow and Kamenica (2016), Kolotilin, Mylovanov, Zapechelnyuk and Li (2017), or Candogan and Strack (2021), for example.

preferences $(u(a,\omega) = -(a-(\omega+b))^2$ where $b \in \mathbb{R})$, and quasi-linear utility functions $(u(a,\omega) = u_1(a) + u_2(a)\cdot\omega)$. From now on, I consider posterior-mean preferences:

Assumption 4. The agents' von Neumann-Morgenstern utility functions are partially separable, and the cost function is of the form

$$c(\pi_I) = -k\left(\int_\Omega \omega \, d\mu_0(\omega)\right) + \int_{\triangle\Omega} k\left(\int_\Omega \omega \, d\mu(\omega)\right) d\pi(\mu).$$

Under these preferences, the sender's best response to his belief in the communication stage as well as the receiver's best response to her belief in the action stage depend on the expected state conditional on the respective belief only. Moreover, both agents' ex-ante expected payoffs are additively separable with respect to the prior μ_0 and the distribution over posteriors π.

Lemma 3. *Take any $\mu, \mu' \in \triangle\Omega$ with $\int_\Omega \omega \, d\mu(\omega) = \int_\Omega \omega \, d\mu'(\omega)$. Then, it holds that*

$$\arg\max_{a\in A} \int_\Omega u_R(a,\omega)\, d\mu(\omega) = \arg\max_{a\in A} \int_\Omega u_R(a,\omega)\, d\mu'(\omega),$$

and for any $a, a' \in A$,

$$\int_\Omega u_S(a,\omega)\, d\mu(\omega) \geq \int_\Omega u_S(a,\omega)\, d\mu(\omega) \;\Leftrightarrow\; \int_\Omega u_S(a,\omega)\, d\mu'(\omega) \geq \int_\Omega u_S(a,\omega)\, d\mu'(\omega).$$

For any $\sigma_{\mathcal{A}} : \triangle(\triangle\Omega) \times \triangle\Omega \to \triangle A$ and $\pi \in \triangle(\triangle\Omega)$, there exist functions $v_{i,1} : \triangle\Omega \to \mathbb{R}$ and $v_{i,2} : \Omega \to \mathbb{R}$ such that

$$\int_{\triangle\Omega}\int_\Omega\int_A u_i(a,\omega)\, d\sigma_{\mathcal{A}}(a|\pi,\mu)\, d\mu(\omega)\, d\pi(\mu) = \int_{\triangle\Omega} v_{i,1}(\mu)\ d\pi(\mu) + \int_\Omega v_{i,2}(\omega)\ d\mu_0(\omega). \quad (6)$$

In any fully revealing PBE, both agents' decisions during the game only depend on the conditional mean of the state given the experiment's outcome — not on the outcome itself. Since learning is costly, it is without loss of generality to restrict attention to experiments whose outcomes imply different conditional means. Let $\bar{\Pi}$ be the set of such experiments.

Next, I introduce bi-pooling policies:[25] These experiments divide the state space into subintervals and assign at most two outcome realizations to each of those subintervals — therefore the name. The following definition formalizes the notion of a state space division:

[25] The term "bi-pooling policy" is borrowed from Arieli et al. (2020). My definition extends their definition to distributions of the state with atoms.

Definition 10. A *partitioning* of an experiment $\pi \in \bar{\Pi}$ is a collection of closed intervals in Ω, denoted by $\{[\underline{\omega}_j, \bar{\omega}_j] | j \in J\}$, endowed with the corresponding collection of open intervals $\{(\underline{\omega}_j, \bar{\omega}_j) | j \in J\}$ such that

1. $\bigcup_{j \in J} [\underline{\omega}_j, \bar{\omega}_j] = \Omega$,
2. $(\underline{\omega}_j, \bar{\omega}_j) \cap (\underline{\omega}_{j'}, \bar{\omega}_{j'}) = \emptyset$ for all $j, j' \in J : j \neq j'$, and
3. for all $\mu \in \text{supp}(\pi)$, there is exactly one $j \in J$ so that $\Pr\left(\tilde{\omega} \in [\underline{\omega}_j, \bar{\omega}_j] | \mu, \pi\right) = 1$.

A partitioning $\{[\underline{\omega}_j, \bar{\omega}_j] | j \in J\}$ of I is called *finest* if there is no $j^* \in J$ and $\omega^* \in [\underline{\omega}_{j^*}, \bar{\omega}_{j^*}]$ so that the collection $\{[\underline{\omega}_j, \bar{\omega}_j] | j \in J \backslash \{j^*\}\} \cup \{[\underline{\omega}_j, \omega^*], [\omega^*, \bar{\omega}_j]\}$ is a partitioning of π, too.

Verbally, a partitioning is a collection of closed and open intervals so that the union of closed intervals covers the state space, the open intervals are pairwise disjoint, and each outcome μ is associated with one closed interval of the collection.[26] Endowed with the definition of a partitioning, I can now specify bi-pooling policies:

Definition 11. Take an experiment $\pi \in \bar{\Pi}$ with finest partitioning $\{[\underline{\omega}_j, \bar{\omega}_j] | j \in J\}$. Then, π is a *bi-pooling policy* if for any $j \in J$ with $(\underline{\omega}_j, \bar{\omega}_j) \cap \text{supp}(\omega) \neq \emptyset$, there are at most two outcomes $\mu \in \text{supp}(\pi)$ that realize if $\omega \in \left(\underline{\omega}_j, \bar{\omega}_j\right)$. Two outcomes form a *2-partition* if they realize with positive probability on the same open interval $(\underline{\omega}_j, \bar{\omega}_j)$. All other outcomes form a *1-partition*.

Monotone partitions form a subset of the set of bi-pooling policies:

Definition 12. A bi-pooling policy is a *monotone partition* if all its outcomes form a 1-partition.

Under posterior-mean preferences, optimal experiments lie within the set of bi-pooling policies, which are a subset of the set of experiments with strongly independent outcomes.

Corollary 3. *Any best implementable payoff profile is generated by a fully revealing PBE with a pure-strategy information rule in which the sender chooses a convex combination of two bi-pooling policies. If $c = 0$, it is generated by a fully revealing PBE with a pure-strategy information rule in which the sender chooses a bi-pooling policy.*

Bi-pooling policies are optimal as they are the most informative experiments (they cannot be replicated by a mixture of other experiments) among all experiments for which both agents have a common interest in the revelation of their information content, that is, among

[26] Each experiment has a partitioning: the singleton collection $\{\Omega\}$. Furthermore, the finest partitioning always exists and is unique.

all experiments the sender is willing to choose in a fully revealing PBE. Reducing the class of experiments to bi-pooling policies is a useful step towards finding the optimal experiments as demonstrated in the next section where I compute optimal bi-pooling policies.

6 The Uniform-Quadratic Case

This section characterizes the optimal experiments in the uniform-quadratic model with an additive bias $b > 0$ and zero cost: The state is uniformly distributed on the interval $[0, 1]$, the agents' von Neumann-Morgenstern utility functions are given by $u_R(a, \omega) = -(a - \omega)^2$ and $u_S(a, \omega) = -(a - (\omega + b))^2$, and the action space is $A = [0, 1]$.

By Corollary 3, some optimal experiment must be a bi-pooling policy, which can be determined using the general maximization problem on page 12.

Lemma 4. *A tuple* $(I, \tilde{\sigma}_{\mathcal{A}})$ *fulfills* (1) *and* (2) *if and only if* $\operatorname{supp}(\sigma_{\mathcal{A}}(\mu)) = \{\int_\Omega \omega \, d\mu(\omega)\}$ *for all* $\mu \in \operatorname{supp}(\pi)$ *and* $\left|\int_\Omega \omega \, d\mu(\omega) - \int_\Omega \omega \, d\mu'(\omega)\right| \geq 2b$ *for all* $\mu, \mu' \in \operatorname{supp}(\pi)$*.*

Due to the quadratic preferences, the receiver's unique best response to outcome μ is the conditional mean $\int_\Omega \omega \, d\mu(\omega)$ in a fully revealing PBE. The sender's incentive compatibility constraints reduce to the distance between any two induced posterior means of the state exceeding some constant, namely twice the bias.

Remark 2. With a perfectly informed sender, incentive compatibility requires that the distance between any two messages is not constant, but increasing because equilibria exhibit intervals of increasing length (cf. Crawford and Sobel (1982)). Where does the difference come from? The recommendation principle is responsible for the constant distance between induced posterior means: This paper's model can be interpreted as one where the sender is perfectly informed, but the states of interest are the conditional means, which are in the middle of the different intervals. In Crawford and Sobel (1982), intervals are of increasing length because the sender is inclined towards exaggerating on the right end of an interval, which are the states where he has the highest incentive to misreport. In my model variant, the sender's incentives towards exaggerating/undermining are equally distributed because the relevant state is not on the right end of an interval, but in the middle.

Ex-ante expected payoffs in a fully revealing PBE differ by a constant between agents: $\tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_{\mathcal{A}}) = \tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) - b^2$ for all $(\pi, \tilde{\sigma}_{\mathcal{A}})$ satisfying the conditions stated in Lemma 4. So there is a unique best implementable payoff profile, making (3) and (4) redundant.

The sender's incentive compatibility constraints together with the fact that Ω is bounded imply that the optimal experiment has a finite support $\{\mu_1, \ldots, \mu_n\}$, where $n \in \mathbb{N}$. Define $\bar{\omega}_i \equiv \int_\Omega \omega \, d\mu_i(\omega)$, and let p_i be the probability that the experiment's outcome is μ_i.

The optimal experiment depends on the value of the bias: If $b > \frac{1}{4}$, no information is revealed.[27] Before deriving the exact structure of an optimal bi-pooling policy for $b \leq \frac{1}{4}$, it is useful to think about its size, i.e., the optimal number of outcome realizations n. The underlying trade-off is that experiments with more outcome realizations are more informative, thus payoff superior, but also less likely implementable in equilibrium. The best experiment is thus of a sufficiently large size that is still consistent with incentive compatibility. The size of an optimal experiment can only take two different values:

Lemma 5. *If $b \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right]$ for some $n \geq 3$, there is an optimal experiment being*

- *the uniform partition*[28] *of size $n-1$, or*
- *a bi-pooling policy of size n so that $\bar{\omega}_{i+1} - \bar{\omega}_i = 2b$ for all $i \in \{1, \ldots, n-1\}$, and both μ_1 and μ_n are 1-partitions.*

Neglecting all incentive compatibility constraints, the best experiment of a specific size is the uniform partition of that size. Moreover, uniform partitions of larger sizes dominate uniform partitions of smaller sizes. Consequently, the uniform partition of size $n-1$ dominates all experiments of size $n-1$ or smaller, and it satisfies incentive compatibility. Since experiments of size $n+1$ or larger are not incentive compatible, the optimal policy is thus either the uniform partition of size $n-1$, or a bi-pooling policy of size n.

6.1 Small Bias: $b \leq \frac{1}{12}$

For a bias $b \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right] \leq \frac{1}{12}$, the optimal bi-pooling policy of size n is either a monotone, non-uniform partition whose 1-partitions are of alternating size, or a bi-pooling policy with exactly two 2-partitions — one between the second and third outcome realization, and another one between the second-last and third-last realization — and equally-sized 1-partitions in-between.

Lemma 6. *Let $n \geq 7$ and $b \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right]$. The optimal bi-pooling policy of size n satisfies $\bar{\omega}_i = \frac{1}{2} + b(2i - n - 1)$ for all i, and there exists some $\hat{b}_n \in \left[\frac{1}{2n}, \frac{1}{2(n-1)}\right)$ so that*

[27] In this case, the distance between any two induced posterior means must be strictly larger than $\frac{1}{2}$, implying that at most two different actions on $[0, 1]$ can be implemented. However, the distance between these two means may be no larger than $\frac{1}{2}$. Due to the uniform distribution of the state, the maximal distance $\frac{1}{2}$ can essentially be achieved if and only if the lower action is induced whenever $\tilde{\omega} \in (0, x)$ and the upper action is induced whenever $\tilde{\omega} \in (x, 1)$ for some $x \in (0, 1)$, that is, whenever the experiment is essentially a partition.

[28] A uniform partition of size $n' \in \mathbb{N}$ is a monotone partition with $p_1 = \ldots = p_{n'}$.

- *if* $b \leq \hat{b}_n$*, it is a monotone partition with*

$$p_i = \begin{cases} 1 - 2b(n-1) & \text{, if } i \text{ is odd} \\ 2b(n+1) - 1 & \text{, if } i \text{ is even} \end{cases} \tag{7}$$

- *if* $b \geq \hat{b}_n$*,* μ_2 *and* μ_3 *as well as* μ_{n-2} *and* μ_{n-1} *form a 2-partition, respectively, and all other* μ_i *are 1-partitions such that*

$$p_i = \begin{cases} 1 - 2b(n-1) & \text{, if } i \in \{1, n\} \\ \frac{(n+4)b}{2} - \frac{1}{4} & \text{, if } i \in \{2, 3, n-2, n-1\} \\ 2b & \text{, else} \end{cases} \tag{8}$$

Furthermore, it holds that $\hat{b}_n \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right)$ *if* n *is odd, and* $\hat{b}_n = \frac{1}{2n}$ *if* n *is even.*

To give an intuition for this result, notice that uniform, monotone partitions would be best, but are infeasible. Therefore, one either has to give up on uniformity by choosing a non-uniform, monotone partition, or one can restore parity for most outcome realizations by choosing two 2-partitions close to the boundary of the state space $[0, 1]$.

If n is even, the optimal experiment of size n cannot be a monotone partition: Any monotone partition for which all incentive constraints of adjacent posterior means are binding has alternatingly sized 1-partitions: $p_1 = p_3 = \ldots = p_{n-1}$ and $p_2 = p_4 = \ldots = p_n$. Moreover, the posterior mean $\bar{\omega}_i$ is always in the middle of the interval of states its outcome is associated with. This implies that $p_1 + p_2 = p_3 + p_4 = \ldots = p_{n-1} + p_n = 4b$ and thus $\sum_{i=1}^{n} p_i = \frac{n}{2} \cdot 4b = 2bn > 1$ — a contradiction because $\sum_{i=1}^{n} p_i = 1$.

Figure 4 depicts an optimal bi-pooling policy of size 12 when the bias is $b = 0.0435$: With values of the state realization on the x-axis and values of the state's density function on the y-axis, the outer rectangle represents the area under the state's density function. It is divided into several sub-rectangles, each of which belongs to one outcome of the bi-pooling policy, which is specified by the circled values. As stated in Lemma 5 and Lemma 6, the first and last outcome form a 1-partition, respectively, the second and third as well as the tenth and the eleventh outcome form a 2-partition, respectively, and the third to ninth outcome form 1-partitions of equal size. Why is this bi-pooling policy optimal? First, note that 2-partitions are involve a loss of information. By separating two outcomes forming one 2-partition into two 2-partitions, keeping their probabilities equal, one could construct a better experiment. However, this is not possible because no monotone partition of size 12 is feasible, as argued above. Hence, 2-partitions are necessary, but due to their inefficiency, it is best to hold

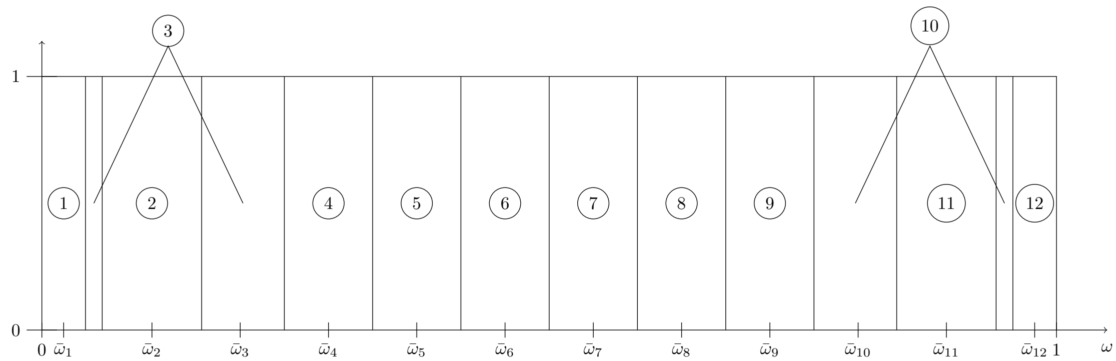

Figure 4: Optimal bi-pooling policy of size $n = 12$ with underlying bias $b = 0.0435$

their number small. Due to the symmetric distribution of the state, symmetric experiments are optimal, therefore, there are two 2-partitions. Since the first and last outcome form 1-partitions, the 2-partitions cannot be directly at the margins of the state space. But why are they not further to the center? If so, there would be 1-partitions of alternating sizes at the margins, and less 1-partitions of equal size in the center. This is worse because uniform partitions dominate monotone partitions of alternating sizes.

6.2 Large Bias: $\frac{1}{12} < b \leq \frac{1}{4}$

If the bias exceeds $\frac{1}{12}$, the optimal bi-pooling policy of maximal size n can be constructed straightforwardly: If $n = 3$, Lemma 5 implies that the optimal experiment must be a monotone partition because both the first and last outcome form a 1-partition, so does the intermediate one. For $n = 4$, there are two options: Again, the first and last outcome form a 1-partition so that the two intermediate outcomes either form a 2-partition or two separate 1-partitions. Recall that monotone partitions are not feasible if the number of outcomes is even. Consequently, the two intermediate outcomes must form a 2-partition. If $n = 5$, three different scenarios are possible: Either the second and third or the third and fourth outcome form a 2-partition, or all outcomes form a 1-partition. Due to the symmetry of the uniform distribution, it turns out that the optimal partition is also symmetric, hence it is a monotone partition. For $n = 6$, one obtains by a similar symmetry-argument and by infeasibility of monotone partitions that both the second and the third as well as the fourth and the fifth outcome form a 2-partition, respectively.

Figure 5 illustrates an optimal policy of size 4 on the left side, which is a bi-pooling policy with a proper 2-partition, and an optimal experiment of size 5 on the right side, which is a monotone partition. The following lemma formalizes the above observations:

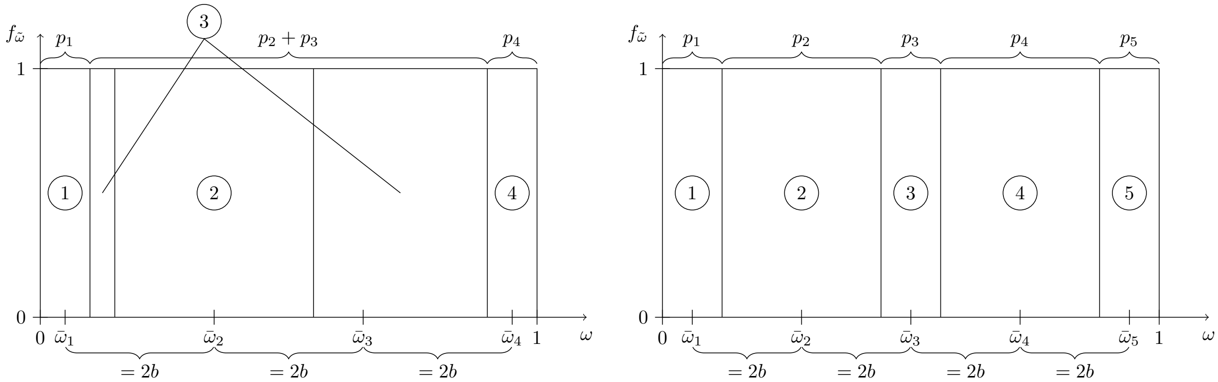

Figure 5: Optimal bi-pooling policy of size $n = 4$ with underlying bias $b = 0.15$ on the left side and of size $n = 5$ with bias $b = 0.11$ on the right side

Lemma 7. *Let $n \in \{3, 4, 5, 6\}$ and $b \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right]$. The optimal bi-pooling policy of size n satisfies $\bar\omega_i = \frac{1}{2} + b(2i - n - 1)$ for all i, and*

- *if n is odd, it is a monotone partition with (7).*
- *if n is even, μ_2 and μ_3 as well as μ_{n-2} and μ_{n-1} form a 2-partition with (8).*

6.3 Globally optimal experiments

The underlying trade-off when comparing the optimal bi-pooling policy of size n with the optimal experiment of size $n - 1$, i.e., the uniform partition of that size, is the following: The optimal bi-pooling policy of size n has an additional outcome realization coming at the expense of non-uniform 1-partitions or 2-partitions to satisfy incentive compatibility. The latter effect dominates as the bias increases implying that the best bi-pooling policy of size n is globally optimal for small bias values in the interval $\left(\frac{1}{2n}, \frac{1}{2(n-1)}\right]$, while the uniform partition of size $n - 1$ is globally optimal for larger biases in that interval.

Proposition 3. *If $b \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right]$ for some $n \geq 3$, there exists some $\bar{b}_n \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right)$ such that the optimal experiment is the best bi-pooling policy of size n if $b \leq \bar{b}_n$ and it is the uniform partition of size $n - 1$ if $b \geq \bar{b}_n$. Furthermore, if $n \geq 7$ and n is odd, then $\bar{b}_n \leq \hat{b}_n$.*

7 Model Variants

7.1 Information Acquisition by the Receiver

Consider a variation of the cheap-talk model in which the receiver chooses the experiment instead of the sender: First, the receiver publicly chooses an experiment (whose costs are borne by the sender). Then, the sender privately observes its outcome and sends a cheap-talk message to the receiver. Third, the receiver takes an action.

The recommendation principle from Section 3.1 remains valid: To see this, recall that the best implementable payoff profiles of the baseline model can be generated by a fully revealing PBE in which the sender chooses a pure-strategy information rule, and a babbling outcome is implemented everywhere off-path. Since both agents' expected payoffs on-path exceed their expected payoffs under the babbling outcome, it does not matter whether the sender or the receiver chooses the experiment. Consequently, the model predictions do not change:

Proposition 4. *The set of best implementable payoff profiles in the model where the receiver publicly chooses the experiment coincides with the one of the original model.*

7.2 Covert Information Acquisition

Suppose now the sender chooses the experiment covertly instead of overtly, that is, he cannot commit to the choice of the experiment.

A version of recommendation principle can be established for this model variant, too. Any PBE of the model variant is a PBE of the baseline model, but the converse does not hold true: If the sender chooses a certain experiment when he cannot commit to it, he would also choose it if he had commitment power. Hence, model predictions may change:

Proposition 5. *Any best implementable payoff profile in the model where the sender learns covertly is dominated by an implementable payoff profile of the original one.*

7.3 Bayesian Persuasion

Under Bayesian persuasion, the sender can commit both to the experiment he chooses and full revelation of its outcome. Therefore, higher payoff can be achieved in equilibrium:

Proposition 6. *Any best implementable payoff profile of the original model is dominated by an implementable payoff profile of the model where the sender can commit to full revelation.*

7.4 Mediation

Suppose the sender is perfectly informed about the state, but does not directly communicate with the receiver. Instead, he sends a cheap-talk message to a neutral mediator who then transmits a message to the receiver. Can better equilibrium outcomes be attained under endogenous learning or mediation?

In general, endogenous, overt learning and mediation cannot be ranked: This is because the sender's incentives to report his information are affected at different stages: Under endogenous learning, the sender has less incentives to misreport as he can decide to acquire less than perfect information. On the other hand, the mediator can enforce outcomes that are beneficial to both the sender and the receiver: If the sender communicates with the receiver directly, he would always send a message that implements the best possible action among all that the receiver chooses in equilibrium given his information. A mediator is not restricted to that (see the mediation solution by Krishna and Morgan (2004), for instance).

In the uniform-quadratic case, however, endogenous learning always outperforms mediation:

Proposition 7. *Consider the uniform-quadratic setting with zero cost. For any* $b \leq \frac{1}{2}$*, the unique best implementable payoff profile under endogenous, overt learning dominates any payoff profile that can be implemented by a mediator.*

As a consequence, endogenous, overt learning also outperforms long cheap talk (see Aumann and Hart (2003)) in this setting because mediation dominates long cheap talk (see Krishna and Morgan (2004)).

7.5 Comparative statics

This section illustrates the effects of the discussed model variants using the example of the uniform-quadratic model.

Ivanov (2010) characterizes the optimal monotone partition of the model variant from Section 7.1. From Proposition 3, one can therefore infer the following:

1. If n is odd, then the optimal policy of size n is always a monotone partition and corresponds to the solution of Ivanov (2010).

2. If n is even, the optimal policy of size n is a bi-pooling policy. Since no monotone partition of size n is feasible, the best experiment in the analysis of Ivanov (2010) is the uniform partition of size $n-1$.

Finally, let's compare the best implementable payoff profile in my model to the optimal payoffs in related models. How does my model (overt and flexible information acquisition by sender or receiver) compare to Crawford and Sobel (1982) (full information acquisition by default), Pei (2015) (covert information acquisition) and Ivanov (2010) (overt information acquisition by sender or receiver among the set of monotone partitions). Figure 6 illustrates the different payoff profiles of these models if $b = 0.126$: The blue line is the Pareto frontier

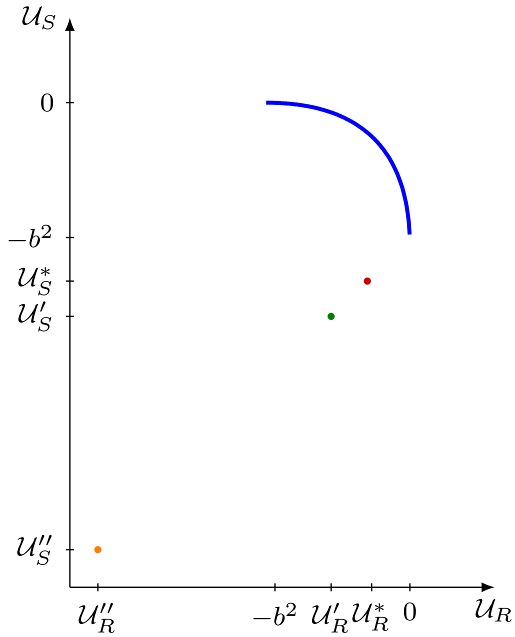

Figure 6: Payoff comparison in the uniform-quadratic setting with bias $b = 0.126$

of feasible payoff profile, the red point ($\mathcal{U}_S^*$,$\mathcal{U}_R^*$) is the solution of my model, the green point ($\mathcal{U}_S'$,$\mathcal{U}_R'$) is the solution of Ivanov (2010) and the yellow point ($\mathcal{U}_S''$,$\mathcal{U}_R''$) is the solution to Crawford and Sobel (1982) and Pei (2015). Comparing the red and the yellow point, note that the red one is considerably closer to the blue Pareto frontier. This suggests that with zero cost, covert information acquisition has no effect, while overt information has a significant effect. By comparing the red point to the green point, one can infer that flexible information acquisition (i.e., the option to choose non-monotone partitions as well) is valuable, too.

On the other hand, how does my model compare to Kamenica and Gentzkow (2011) (i.e., Bayesian persuasion)? Bayesian persuasion means that the sender can commit to both the information as well as the communication rule. In the uniform-quadratic setting, this means that the sender can commit to the fully informative experiment, and the receiver chooses her optimal action. So the solution of Kamenica and Gentzkow is the right endpoint of the blue curve. My setting is a model of partial commitment (only commitment with respect to information, not with respect to communication), and Crawford and Sobel (1982) can be interpreted as an environment with no commitment at all. The partial commitment solution is closer to the full commitment solution than to the no commitment solution. This suggests that the value of commitment with respect to information is higher than the value of commitment with respect to messages. Or in other words, the efficiency loss of giving up on the latter commitment device while keeping the former one is not too large. This is good news to the related literature: A common critique about Bayesian persuasion is that commitment with respect to the communication strategy is hard to apply to real-world settings. On the other hand, there are applications for which commitment with respect to information makes sense (see the example of a survey from the introduction). Since the latter commitment device dominates the former (in terms of value) my model of partial commitment could be an adequate alternative to Bayesian persuasion.

8 Conclusion

Summing up, this paper explores the effects of costly information procurement in a model of strategic information transmission from an expert to a decision-maker. The first main finding is that in an optimal equilibrium, the sender generically only acquires information that he is also willing to forward to the receiver.

Under posterior separable cost, optimal information structures are determined by the class of experiments with independent outcomes. In settings with a finite state space, applying this fact considerably reduces the set of experiments among which one can find an optimal information structure. Besides, I have derived an equivalence of the characteristics of optimal experiments in cheap-talk and Bayesian persuasion problems, and I have solved the uniform-quadratic model.

References

Aliprantis, C. and K. Border (2006). *Infinite Dimensional Analysis*. Springer, Berlin.

Ambrus, A., E. M. Azevedo, and Y. Kamada (2013). "Hierarchical cheap talk". In: *Theoretical Economics* 8.1, pp. 233–261.

Arieli, I., Y. Babichenko, and R. Smorodinsky (2020). "Optimal Persuasion via Bi-Pooling". Working paper.

Aumann, R. J. and S. Hart (2003). "Long cheap talk". In: *Econometrica* 71.6, pp. 1619–1660.

Billingsley, P. (1999). *Convergence of probability measures*. John Wiley & Sons.

Blackwell, D. (1953). "Equivalent comparisons of experiments". In: *The Annals of Mathematical Statistics* 24.2, pp. 265–272.

Candogan, O. and P. Strack (2021). "Optimal disclosure of information to a privately informed receiver". Working paper.

Caplin, A. and M. Dean (2013). *Behavioral implications of rational inattention with shannon entropy*. Tech. rep. National Bureau of Economic Research.

Caplin, A., M. Dean, and J. Leahy (2019). "Rational inattention, optimal consideration sets, and stochastic choice". In: *The Review of Economic Studies* 86.3, pp. 1061–1094.

Crawford, V. P. and J. Sobel (1982). "Strategic information transmission". In: *Econometrica* 50.6, pp. 1431–1451.

Crémer, J. and R. P. McLean (1988). "Full extraction of the surplus in Bayesian and dominant strategy auctions". In: *Econometrica: Journal of the Econometric Society*, pp. 1247–1257.

Deimen, I. and D. Szalay (2019). "Delegated expertise, authority, and communication". In: *American Economic Review* 109.4, pp. 1349–74.

Dessein, W. (2002). "Authority and communication in organizations". In: *The Review of Economic Studies* 69.4, pp. 811–838.

Dudley, R. M. (1989). "Real analysis and probability. Wadsworth & Brooks". In: *Cole, Pacific Groves, California* 8, p. 39.

Ely, J. C. (2017). "Beeps". In: *American Economic Review* 107.1, pp. 31–53.

Gentzkow, M. and E. Kamenica (2014). "Costly persuasion". In: *American Economic Review* 104.5, pp. 457–62.

— (2016). "A Rothschild-Stiglitz approach to Bayesian persuasion". In: *American Economic Review* 106.5, pp. 597–601.

Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani (2009). "Mediation, arbitration and negotiation". In: *Journal of Economic Theory* 144.4, pp. 1397–1420.

Ivanov, M. (2010). "Informational control and organizational design". In: *Journal of Economic Theory* 145.2, pp. 721–751.

Kamenica, E. and M. Gentzkow (2011). "Bayesian persuasion". In: *American Economic Review* 101.6, pp. 2590–2615.

Kleiner, A., B. Moldovanu, and P. Strack (2021). "Extreme points and majorization: Economic applications". In: *Econometrica* 89.4, pp. 1557–1593.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017). "Persuasion of a privately informed receiver". In: *Econometrica* 85.6, pp. 1949–1964.

Krishna, V. and J. Morgan (2004). "The art of conversation: eliciting information from experts through multi-stage communication". In: *Journal of Economic theory* 117.2, pp. 147–179.

Lipnowski, E. (2020). *Equivalence of Cheap Talk and Bayesian Persuasion in a Finite Continuous Model*.

Lipnowski, E. and D. Ravid (2020). "Cheap talk with transparent motives". In: *Econometrica* 88.4, pp. 1631–1660.

Lyu, Q. and W. Suen (2022). "Information design in cheap talk". Working paper.

Maćkowiak, B., F. Matějka, and M. Wiederholt (2018). "Survey: Rational inattention, a disciplined behavioral model". Working paper.

Matějka, F. and A. McKay (2015). "Rational inattention to discrete choices: A new foundation for the multinomial logit model". In: *American Economic Review* 105.1, pp. 272–98.

Matysková, L. and A. Montes (2021). "Bayesian Persuasion With Costly Information Acquisition". Working paper.

Pei, H. D. (2015). "Communication with endogenous information acquisition". In: *Journal of Economic Theory* 160, pp. 132–149.

Yang, M. (2015). "Coordination with flexible information acquisition". In: *Journal of Economic Theory* 158, pp. 721–738.

Appendix

Proofs to Section 2: Model

Proof of Proposition 1 The proof proceeds by construction of a babbling equilibrium.

Since Π is nonempty (by implicit assumption), and π^0 is a Blackwell garbling of any experiment, it follows from Assumption 2 that $\pi^0 \in \Pi$. Assumption 1 implies that $c(\pi^0) = 0$.

Fix some $m_0 \in \triangle\Omega$ and some $a_0 \in A_0 \equiv \arg\max_{a \in A} \int_\Omega u_R(a, \omega)\, d\mu_0(\omega)$. Notice that a_0 is well-defined: Since $A \times \Omega$ is compact and u_R is continuous on $A \times \Omega$, u_R is bounded on $A \times \Omega$. By the Dominated Convergence Theorem, continuity of $u_R(\cdot, \omega)$ on A for each $\omega \in \Omega$ implies continuity of $\int_\Omega u_R(\cdot, \omega)\, d\mu_0(\omega)$ on A. Consequently, $\int_\Omega u_R(\cdot, \omega)\, d\mu_0(\omega)$ attains a maximum on the compact set A, i.e., A_0 is nonempty, by the Weierstrass extreme value theorem.

Consider the strategy profile $((\sigma^0_{\mathcal{I}}, \sigma^0_{\mathcal{M}}), \sigma^0_{\mathcal{A}})$ and beliefs μ^0_R with

$$\begin{aligned}
&\operatorname{supp}\left(\sigma^0_{\mathcal{I}}\right) = \{\pi^0\},\\
&\operatorname{supp}\left(\sigma^0_{\mathcal{M}}(\pi, \mu)\right) = \{m_0\} \quad \text{for all } \mu \in \triangle\Omega \text{ and all } \pi \in \Pi,\\
&\operatorname{supp}\left(\sigma^0_{\mathcal{A}}(\pi, m)\right) = \{a_0\} \quad \text{for all } m \in \triangle\Omega \text{ and all } \pi \in \Pi, \text{ and}\\
&\mu^0_R(\cdot|\pi, m) = \mu_0 \quad \text{for all } m \in \triangle\Omega \text{ and all } \pi \in \Pi.
\end{aligned}$$

It can be easily verified that these profiles constitute a PBE E^0: First, the receiver's beliefs are always consistent with Bayes' rule on the equilibrium path, and it is F in any subgame in the action stage. Hence, it is — by definition of a_0 — optimal for her to choose a_0 always. Likewise, it is a best response for the sender to transmit the uninformative message m_0 always because the receiver does not condition her action decision on the sender's message. Third, note that both agents' expected utilities (excluding the sender's cost) are constant across all subgames after the information acquisition stage. Therefore, the sender optimally chooses an experiment which induces the lowest possible cost in the information acquisition stage. Since $c(\pi^0) = 0$, it is optimal to take experiment π^0.

By definition of μ^0_R, E^0 is a babbling equilibrium.

Proof of Proposition 2 Fix some PBE $E' = \{((\sigma'_{\mathcal{I}}, \sigma'_{\mathcal{M}}), \sigma'_{\mathcal{A}}), \mu'_R\}$.

(a) First, note that $\mu_R(\cdot|\pi, m) = \mu_0$ for all $m \in \operatorname{supp}(\sigma'_{\mathcal{M}}(\pi, \mu))$, $\mu \in \operatorname{supp}(\pi)$ and $\pi \in \operatorname{supp}(\sigma_{\mathcal{I}})$ for any $E \in \mathcal{E}_0$. By optimality of the receiver's action rule in all such PBE, her expected payoff is the same in all on-path subgames after the information acquisition stage of any such PBE. So her ex-ante expected payoff is constant across all babbling PBE:

$$\inf_{E \in \mathcal{E}_0} \mathcal{U}_R(E') = \mathcal{U}_R\left(E^0\right) = \sup_{E' \in \mathcal{E}_0} \mathcal{U}_R(E')$$

where E^0 is defined as in the proof of Proposition 1.

Next, I demonstrate that $\mathcal{U}_R(E') \geq \mathcal{U}_R(E^0)$: If the receiver deviated in the PBE E' to $\sigma^0_{\mathcal{A}}$, she would attain the ex-ante expected payoff $\mathcal{U}_R(E^0)$. But as $\sigma'_{\mathcal{A}}$ is a best response for the receiver in E', it follows that $\mathcal{U}_R(E') \geq \mathcal{U}_R(E^0)$.

(b) Consider the strategy profile $((\sigma_{\mathcal{I}}^0, \sigma_{\mathcal{M}}^0), \sigma_{\mathcal{A}}^{00})$ and beliefs μ_R^0 where

$$\text{supp}\left(\sigma_{\mathcal{A}}^{00}\left(\pi^0, m\right)\right) \in A_{00} \equiv \arg\min_{a \in A_0} \int\limits_{\Omega} u_S(a,\omega)\, d\mu_0(\omega) \tag{9}$$

for all $m \in \triangle\Omega$. Recall from the proof of Proposition 1 that $\int_\Omega u_S(a,\omega)\, d\mu_0(\omega)$ is continuous on A. This together with the compactness of A implies that A_0 is compact by Berge's maximum theorem. Consequently, A_{00} is nonempty by the Weierstrass extreme value theorem.

Using analogous arguments as in the proof of Proposition 1, one can check that the above profiles constitute a PBE $E^{00} \in \mathcal{E}^0$.

Next, I show that the sender's ex-ante expected payoff in the PBE E^{00} serves as a lower bound for his ex-ante expected payoff in any PBE: For any $\pi \in \text{supp}(\sigma_{\mathcal{I}}')$, it must hold that

$$\begin{aligned}
&\int\limits_{\triangle\Omega}\int\limits_{\triangle\Omega}\int\limits_{\Omega}\int\limits_{A} u_S(a,\omega)\ d\sigma_{\mathcal{A}}'(a|\pi,m)\ d\mu(\omega)\ d\sigma_{\mathcal{M}}'(m|\pi,\mu)\ d\pi(\mu) - c(\pi)\\
&\geq \int\limits_{\triangle\Omega}\int\limits_{\triangle\Omega}\int\limits_{\Omega}\int\limits_{A} u_S(a,\omega)\ d\sigma_{\mathcal{A}}'\left(a|\pi^0,m\right)\ d\mu(\omega)\ d\sigma_{\mathcal{M}}'\left(m|\pi^0,\mu\right)\ d\pi^0(\mu) - c\left(\pi^0\right)\\
&\geq \mathcal{U}_S\left(E^{00}\right).
\end{aligned}$$

The first inequality follows from the fact that $\sigma_{\mathcal{I}}'$ is a best response for the sender given $(\sigma_{\mathcal{M}}', \sigma_{\mathcal{A}}')$, and the second inequality holds true by construction of E^{00} due to (9).

This yields

$$\mathcal{U}_S(E) \geq \mathcal{U}_S\left(E^{00}\right), \tag{10}$$

completing the proof.

Proofs to Section 3: Best Implementable Outcomes

Proof of Lemma 1 Fix some best implementable payoff profile generated by the PBE $E = \{((\sigma_{\mathcal{I}}, \sigma_{\mathcal{M}}), \sigma_{\mathcal{A}}), \mu_R\}$. For any $\pi \in \text{supp}(\sigma_{\mathcal{I}})$, let $E^\pi = \{((\sigma_{\mathcal{I}}^\pi, \sigma_{\mathcal{M}}), \sigma_{\mathcal{A}}), \mu_R\}$ where $\text{supp}(\sigma_{\mathcal{I}}^\pi) = \{\pi\}$, and note that E^π is a PBE, too. By optimality of $\sigma_{\mathcal{I}}$ for the sender in the PBE E, it follows for all $\pi, \pi' \in \text{supp}(\sigma_{\mathcal{I}})$ that $\mathcal{U}_S(E^\pi) = \mathcal{U}_S\left(E^{\pi'}\right)$ because

$$\begin{aligned}
&\int\limits_{\triangle\Omega}\int\limits_{\triangle\Omega}\int\limits_{\Omega}\int\limits_{A} u_S(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi,m)\ d\mu(\omega)\ d\sigma_{\mathcal{M}}(m|\pi,\mu)\ d\pi(\mu) - c(\pi)\\
&\quad = \int\limits_{\triangle\Omega}\int\limits_{\triangle\Omega}\int\limits_{\Omega}\int\limits_{A} u_S(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi',m)\ d\mu(\omega)\ d\sigma_{\mathcal{M}}(m|\pi',\mu)\ d\pi'(\mu) - c(\pi').
\end{aligned}$$

Best implementability of $(\mathcal{U}_S(E), \mathcal{U}_R(E))$ yields $\mathcal{U}_R(E) \geq \mathcal{U}_R(E^\pi)$ for all $\pi \in \text{supp}(\sigma_{\mathcal{I}})$. By construction, the former one is the mean of the latter ones, $\mathcal{U}_R(E) = \int_\Pi \mathcal{U}_R(E^\pi)\ d\sigma_{\mathcal{I}}(\pi)$, so that $\mathcal{U}_R(E) = \mathcal{U}_R(E^\pi)$ for all π. Hence, $(\mathcal{U}_S(E), \mathcal{U}_R(E))$ is generated by any PBE E^π.

Proof of Lemma 2 Fix some PBE E with supp$(\sigma_{\mathcal{I}}) = \{\pi^*\}$ for some $\pi^* \in \Pi$. Now consider the experiment π^{**} with outcomes $\{\mu_R(\cdot|\pi^*, m)\}_{m \in \bigcup_{\mu \in \text{supp}(\pi^*)} \text{supp}(\sigma_{\mathcal{M}}(\pi^*,\mu))}$ distributed according to $\int_{\mu \in \text{supp}(\pi^*)} \sigma_{\mathcal{M}}(\cdot|\pi^*, \mu) \; d\pi^*(\mu)$. In words, the distribution over outcomes under the experiment π^{**} corresponds to the distribution over the receiver's on-path beliefs in the PBE E. By Assumption 2, $\pi^{**} \in \Pi$ as π^{**} is a Blackwell garbling of π^*. By Assumption 1, one gets that $c(\pi^{**}) \leq c(\pi^*)$.

One can find a fully revealing PBE E' in which the sender chooses π^{**}. The agents' expected utilities are the same as in the PBE E, but the sender's cost are smaller in E'. Hence, E' dominates E. In order to construct the fully revealing PBE, consider the strategy profile $((\sigma'_{\mathcal{I}}, \sigma'_{\mathcal{M}}), \sigma'_{\mathcal{A}})$ and beliefs μ'_R with supp$(\sigma'_{\mathcal{I}}) = \{\pi^{**}\}$,

$$\sigma'_{\mathcal{M}}(\cdot|\pi, \mu) = \begin{cases} \sigma^{\text{FR}}_{\mathcal{M}}(\cdot|\pi^{**}, \mu) & , \pi = \pi^{**} \\ \sigma^0_{\mathcal{M}}(\cdot|\pi, \mu) & , \text{else} \end{cases} \quad \text{for all } \mu \in \triangle\Omega$$

$$\sigma'_{\mathcal{A}}(\cdot|\pi, m) = \begin{cases} \sigma_{\mathcal{A}}(\cdot|\pi^*, \tau(m)) & , m \in \text{supp}\left(\sigma'_{\mathcal{M}}(\pi^{**}, \mu_R(\cdot|\pi^*, \tau(m)))\right) \text{ and } \pi = \pi^{**} \\ \sigma^0_{\mathcal{A}}(\cdot|\pi, m) & , m \in \triangle\Omega \text{ and } \pi \neq \pi^{**} \end{cases}$$

$$\mu'_R(\cdot|\pi, m) = \begin{cases} \mu & , m \in \text{supp}\left(\sigma'_{\mathcal{M}}(\pi^{**}, \mu)\right) \text{ and } \pi = \pi^{**} \\ F & , m \in \triangle\Omega \text{ and } \pi \neq \pi^{**} \end{cases},$$

where $((\sigma^0_I, \sigma^0_m), \sigma^{00}_a)$ and (μ^0_s, μ^0_R) are defined as in the proofs of Proposition 1 and 2, and where $\tau : \bigcup_{\mu \in \text{supp}(\pi^{**})} \text{supp}\left(\sigma'_{\mathcal{M}}(\pi^{**}, \mu)\right) \to \bigcup_{\mu \in \text{supp}(\pi^*)} \text{supp}\left(\sigma_{\mathcal{M}}(\pi^*, \mu)\right)$ is a bijective function defined as follows: If the sender sends message m after choosing π^* in the PBE E, he sends message $\tau^{-1}(m)$ after taking π^{**} in the equilibrium candidate E' if the experiment's outcome is $\mu_R(\cdot|\pi^*, m)$.[29] The above profiles constitute a PBE E': First, the agents' beliefs are updated according to Bayes' rule whenever possible. Moreover, since $\int_{\Omega} u_R(a, \omega) \;\; d\mu'_R(\omega|\pi^{**}, m) = \int_{\Omega} u_R(a, \omega) \;\; d\mu_R(\omega|\pi^*, \tau(m))$ for all $a \in A$ and all $m \in \text{supp}\left(\sigma'_{\mathcal{M}}(\pi^{**}, \mu_R(\cdot|\pi^*, \tau(m)))\right)$, and since $\sigma_{\mathcal{A}}$ is a best response given μ_R in any subgame after the sender chooses π^*, one can infer that $\sigma'_{\mathcal{A}}$ is optimal given μ'_R in any subgame after he takes π^{**}. Besides, $\sigma'_{\mathcal{A}}$ is also a best response given μ'_R in any subgame after the sender chooses an experiment $\pi \neq \pi^{**}$ (see the proof of Proposition 2). Hence, $\sigma'_{\mathcal{A}}$ is optimal

[29]Bijectivity of τ (or rather τ^{-1}) ensures that E' is fully revealing: If $\mu_R(\cdot|\pi^*, m) \neq \mu_R(\cdot|\pi^*, m')$ for some $m, m' \in \bigcup_{\mu \in \text{supp}(\pi^*)} \text{supp}\left(\sigma_{\mathcal{M}}(\pi^*, \mu)\right)$, but $\tau^{-1}(m) = \tau^{-1}(m')$, the sender would not fully reveal to the receiver whether he observed outcome $\mu_R(\cdot|\pi^*, m)$ or $\mu_R(\cdot|\pi^*, m')$ after choosing π^{**} in E'. Moreover, $\sigma'_{\mathcal{A}}(\cdot|\pi^{**}, m)$ is well-defined for all $m \in \bigcup_{\mu \in \text{supp}(\pi^{**})} \text{supp}\left(\sigma'_{\mathcal{M}}(\pi^{**}, \mu)\right)$ since the communication rule is fully revealing on the equilibrium path and the receiver updates her beliefs according to Bayes' rule: For any such m, which the sender sends after choosing π^{**} in the candidate for a fully revealing equilibrium E', there exists a message $\tau(m)$ which he sends in the PBE E after choosing π^* so that the receiver's belief after observing π^* and $\tau(m)$ in E equals the outcome of the experiment π^{**} conditional on which the sender transmits m in E'.

given μ'_R. Furthermore, one obtains that

$$\begin{aligned}
&\int_\Omega \int_A u_S(a,\omega)\ d\sigma'_{\mathcal{A}}(a|\pi^{**},m)\ d\mu(\omega) \\
&\qquad = \int_\Omega \int_A u_S(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi^{*},\tau(m))\ d\mu(\omega) \\
&\qquad = \int_{\triangle\Omega} \underbrace{\int_\Omega \int_A u_S(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi^{*},\tau(m))\ d\mu(\omega)}_{\geq \int_\Omega \int_A u_S(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi^{*},\tau(m'))\ d\mu(\omega) \text{ for all } \mu' \in \operatorname{supp}(\pi^*)}\ d\Pr(\mu'|\pi^*,\tau(m),\sigma_{\mathcal{M}}) \\
&\qquad \geq \int_\Omega \int_A u_S(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi^{*},\tau(m'))\ d\mu(\omega) \\
&\qquad = \int_\Omega \int_A u_S(a,\omega)\ d\sigma'_{\mathcal{A}}(a|\pi^{**},m')\ d\mu(\omega)
\end{aligned}$$

for all $m \in \operatorname{supp}(\sigma'_{\mathcal{M}}(\pi^{**},\mu))$, $m' \in \triangle\Omega$ and $\mu \in \operatorname{supp}(\pi^{**})$. The first equality holds true since $\sigma'_{\mathcal{A}}(\cdot|\pi^{**},m) = \sigma_{\mathcal{A}}(\cdot|\pi^{*},\tau(m))$ for all $m \in \bigcup_{\mu\in\operatorname{supp}(\pi^{**})} \operatorname{supp}(\sigma'_{\mathcal{M}}(\pi^{**},\mu))$. The second equality is obtained by the law of iterated expectations, with $\Pr(\cdot|\pi^*,\tau(m),\sigma_{\mathcal{M}})$ being the distribution over outcomes after the sender chooses π^* and given he takes message $\tau(m)$ according to $\sigma_{\mathcal{M}}$. The weak inequality is due to the fact that $\tau(m)$ is optimal for the sender in the PBE E after taking π^* and learning its outcome μ so that $\tau(m) \in \operatorname{supp}(\sigma_{\mathcal{M}}(\pi^*,\mu))$. The last two equalities follow from the same arguments as for the first two equalities applied in reverse order. So $\sigma'_{\mathcal{M}}$ is a best response given $\sigma'_{\mathcal{A}}$ in any subgame after the sender chooses π^{**}. Also, $\sigma'_{\mathcal{M}}$ is a best response given $\sigma'_{\mathcal{A}}$ in any subgame after the sender chooses an $\pi \neq \pi^{**}$ (see the proof of Proposition 2). Hence, $\sigma'_{\mathcal{M}}$ is optimal given $\sigma'_{\mathcal{A}}$. By the law of iterated expectations, it holds for all $i \in \{S,R\}$ that

$$\begin{aligned}
&\int_{\triangle\Omega}\int_{\triangle\Omega}\int_\Omega \int_A u_i(a,\omega)\ d\sigma'_{\mathcal{A}}(a|\pi^{**},m)\ d\mu(\omega)\ d\sigma'_{\mathcal{M}}(m|\pi^{**},\mu)\ d\pi^{**}(\mu) \qquad (11)\\
&= \int_{\triangle\Omega}\int_{\triangle\Omega}\int_{\triangle\Omega}\int_\Omega \int_A u_i(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi^{*},\tau(m))\ d\mu(\omega)\ d\Pr(\mu'|\pi^*,\tau(m),\sigma_{\mathcal{M}})\ d\Pr(\tau(m)|\pi^*) \\
&= \int_{\triangle\Omega}\int_{\triangle\Omega}\int_\Omega \int_A u_i(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi^{*},\tau(m))\ d\mu(\omega)\ d\sigma_{\mathcal{M}}(\tau(m)|\pi^*,\mu)\ d\pi^*(\mu),
\end{aligned}$$

so each agent's expected utilities after π^{**} in E' and after π^* in E coincide. This yields

$$\begin{aligned}
&\int_{\triangle\Omega}\int_{\triangle\Omega}\int_{\Omega}\int_{A} u_S(a,\omega)\ d\sigma'_{\mathcal{A}}(a|\pi^{**},m)\ d\mu(\omega)\ d\sigma'_{\mathcal{M}}(m|\pi^{**},\mu)\ d\pi^{**}(\mu) - c(\pi^{**}) \\
&\geq \int_{\triangle\Omega}\int_{\triangle\Omega}\int_{\Omega}\int_{A} u_S(a,\omega)\ d\sigma_{\mathcal{A}}(a|\pi^{*},\tau(m))\ d\mu(\omega)\ d\sigma_{\mathcal{M}}(\tau(m)|\pi^{*},\mu)\ d\pi^{*}(\mu) - c(\pi^{*}) \\
&\geq \mathcal{U}_S\left(E^{00}\right) \\
&= \int_{\triangle\Omega}\int_{\triangle\Omega}\int_{\Omega}\int_{A} u_S(a,\omega)\ d\sigma'_{\mathcal{A}}\left(a|\pi^{0},m\right)\ d\mu(\omega)\ d\sigma'_{\mathcal{M}}\left(m|\pi^{0},\mu\right)\ d\pi^{0}(\mu) - c\left(\pi^{0}\right) \\
&\geq \int_{\triangle\Omega}\int_{\triangle\Omega}\int_{\Omega}\int_{A} u_S(a,\omega)\ d\sigma'_{\mathcal{A}}(a|\pi,m)\ d\mu(\omega)\ d\sigma'_{\mathcal{M}}(m|\pi,\mu)\ d\pi(\mu) - c(\pi)
\end{aligned}$$

for all $\pi \in \Pi$. The first inequality follows from $c(\pi^{**}) \leq c(\pi^*)$. Since E is a PBE where the sender chooses π^*, the term in the second line equals $\mathcal{U}_S(E)$. The second inequality immediately follows from the definition of E^{00} (cf. the proof of Proposition 2). Since the sender's expected utility is constant across all subgames after the sender chooses some $\pi \neq \pi^{**}$ (as the babbling outcome is implemented in any such subgame), the last inequality results from the fact that $c(\pi) \geq c(\pi^0)$ by Assumption 1 as π^0 is a Blackwell garbling of π. As a consequence, the $\sigma'_{\mathcal{I}}$ is optimal given $(\sigma'_{\mathcal{M}}, \sigma'_{\mathcal{A}})$. This concludes the proof that $\{((\sigma'_{\mathcal{I}}, \sigma'_{\mathcal{M}}), \sigma'_{\mathcal{A}}), \mu'_R\}$ forms a PBE E', which is fully revealing and has a pure-strategy information rule by construction. (11) implies that $\mathcal{U}_R(E') = \mathcal{U}_R(E)$, and that $\mathcal{U}_S(E') \geq \mathcal{U}_S(E)$ since $c(\pi_{**}) \leq c(\pi^*)$. The PBE E is thus dominated by the fully revealing PBE E', which has a pure-strategy information rule.

Proof of Theorem 1 Fix a best implementable payoff profile. By Lemma 1, there is a PBE with a pure-strategy information rule generating this payoff profile. Suppose none of these PBE is fully revealing. Fix such a PBE E, and note that by Lemma 2, the PBE E is dominated by some fully revealing PBE E' with a pure-strategy information rule. Hence, either $(\mathcal{U}_S,\mathcal{U}_R)$ is generated by E' or $(\mathcal{U}_S,\mathcal{U}_R)$ is not best implementable — a contradiction.

Proof of Theorem 2 The proof proceeds by using Berge's maximum theorem for the optimization problem stated on page 12. Let $\mathcal{S}$ be the set of continuous mappings $\tilde{\sigma}_{\mathcal{A}} : (\triangle\Omega, d_{\mathrm{LP}}) \to (\triangle A, d_{\mathrm{LP}})$, where d_{LP} is the Lévy–Prokhorov metric[30], endowed with the sup metric, and let $C : \mathbb{R} \rightrightarrows \Pi \times \mathcal{S}$ be the correspondence defined by the constraint set of the optimization problem: $C(\mathcal{U}_R) \equiv \{(\pi, \tilde{\sigma}_{\mathcal{A}}) \in \Pi \times \mathcal{S} | (\pi, \tilde{\sigma}_{\mathcal{A}}) \text{ satisfies } (1)-(4).\}$ for each $\mathcal{U}_R \in \mathbb{R}$. By Assumption 2.3, Π is compact. Moreover, the set of continuous functions $\mathcal{S}$ endowed with the sup metric constitutes a compact space. So by Tychonoff's theorem, $\Pi \times \mathcal{S}$ is compact. To apply the maximum theorem, I need the following auxiliary lemmata. The first one is on the compactness of the values of the correspondence C:

Lemma A.1. *If $C(\mathcal{U}_R) \neq \emptyset$, then $C(\mathcal{U}_R)$ is compact.*

[30]See Dudley (1989), p.394 for the definition.

Proof. Suppose $C(\mathcal{U}_R) \neq \emptyset$. Let's show that $C(\mathcal{U}_R)$ is closed: Fix some sequence $(\pi_n, \tilde{\sigma}_{\mathcal{A},n})_{n\in\mathbb{N}}$ in $C(\mathcal{U}_R)$ with limit $(\pi, \tilde{\sigma}_{\mathcal{A}})$. Let $\tilde{a}_n(\mu)(\tilde{a}(\mu))$ denote the random variable with distribution $\tilde{\sigma}_{\mathcal{A},n}(\cdot|\mu)(\tilde{\sigma}_{\mathcal{A}}(\cdot|\mu))$. I verify in sequence that $(\pi, \tilde{\sigma}_{\mathcal{A}})$ satisfies (1)–(4):

(1) Suppose there is an $a' \in \text{supp}(\tilde{\sigma}_{\mathcal{A}}(\mu))$ with $a' \notin \arg\max_{a\in A} \int_\Omega u_R(a,\omega)\ d\mu(\omega)$. Note that $\Pr(\tilde{a}(\mu) \in (a'-\epsilon, a'+\epsilon)) = \tilde{\sigma}_{\mathcal{A}}(a'+\epsilon|\mu) - \tilde{\sigma}_a(a'-\epsilon|\mu) > 0$ for any $\epsilon > 0$ as $a' \in \text{supp}(\tilde{\sigma}_{\mathcal{A}}(\mu))$. Since the metric space $(\triangle A, d_{\text{LP}})$ is endowed with the Lévy-Prokhorov metric, $(\tilde{\sigma}_{\mathcal{A},n}(\cdot|\mu))$ converges in distribution to $\tilde{\sigma}_{\mathcal{A}}(\cdot|\mu)$ (see Billingsley (1999), p.72). Take any $\epsilon > 0$ with $\Pr(\tilde{a}(\mu) \in \{a'-\epsilon, a'+\epsilon\}) = 0$, and note that $(a'-\epsilon, a'+\epsilon)$ is a continuity set of $\tilde{a}(\mu)$. Let ϵ^* denote the set of such ϵ. As $(\tilde{\sigma}_{\mathcal{A},n}(\cdot|\mu))$ converges in distribution to $\tilde{\sigma}_{\mathcal{A}}(\cdot|\mu)$, one gets $\Pr(\tilde{\sigma}_{\mathcal{A},n}(\cdot|\mu) \in (a'-\epsilon, a'+\epsilon)) \to \Pr(\tilde{\sigma}_{\mathcal{A}}(\cdot|\mu) \in (a'-\epsilon, a'+\epsilon))$ as $n \to \infty$ for any $\epsilon \in \epsilon^*$. As $\Pr(\tilde{\sigma}_{\mathcal{A}}(\cdot|\mu) \in (a'-\epsilon, a'+\epsilon)) > 0$, for all $\epsilon \in \epsilon^*$, there is an $n(\epsilon) \in \mathbb{N}$ with $\Pr(\tilde{\sigma}_{a,n}(\cdot|\mu) \in (a'-\epsilon, a'+\epsilon)) > 0$. So there is some $a'(\epsilon) \in (a'-\epsilon, a'+\epsilon)$ so that $a'(\epsilon) \in \arg\max_{a\in A} \int_\Omega u_R(a,\omega)\ d\mu(\omega)$ because $\left(\pi_{n(\epsilon)}, \tilde{\sigma}_{\mathcal{A},n(\epsilon)}\right)$ satisfies (1). By construction, $a'(\epsilon) \to a'$ as $\epsilon \to 0$.[31] Since $\arg\max_{a\in A} \int_\Omega u_R(a,\omega)\ d\mu(\omega)$ is closed, this yields $a' \in \arg\max_{a\in A} \int_\Omega u_R(a,\omega)\ d\mu(\omega)$.

(2) Suppose $\int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu)\ d\mu(\omega) < \int_\Omega \int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu')\ d\mu(\omega)$ for some $\mu, \mu' \in \text{supp}(\pi)$. So for all $n \in \mathbb{N}$, $\mu \notin \text{supp}(\pi_n)$ or $\mu' \notin \text{supp}(\pi_n)$. Take a subsequence $(\pi_{n_k}, \tilde{\sigma}_{\mathcal{A},n_k})_{k\in\mathbb{N}}$ of $(\pi_n, \tilde{\sigma}_{\mathcal{A},n})_{n\in\mathbb{N}}$ so that $\mu \notin \text{supp}(\pi_{n_k})$ for any k.[32] By analogous arguments as in the proof of (1), there exists a subsubsequence $\left(\pi_{n_{k_l}}, \tilde{\sigma}_{\mathcal{A},n_{k_l}}\right)_{l\in\mathbb{N}}$ and a sequence $(\mu_l)_l((\mu'_l)_l)$ with $\mu_l \in \text{supp}(\pi_{n_{k_l}})$ $(\mu_l \in \text{supp}(\pi_{n_{k_l}}))$ for all l converging to μ (μ'). Since $\int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu)\ d\mu(\omega) - \int_\Omega \int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu')\ d\mu(\omega)$ is continuous in (μ, μ') and $\int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu)\ d\mu_l(\omega) - \int_\Omega \int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu')\ d\mu_l(\omega) \geq 0$ for each l because $\left(\pi_{n_{k_l}}, \tilde{\sigma}_{\mathcal{A},n_{k_l}}\right)$ satisfies condition (2), one obtains for the limit that $\int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu)\ d\mu(\omega) - \int_\Omega \int_A u_S(a,\omega)\ d\tilde{\sigma}_{\mathcal{A}}(a|\mu')\ d\mu(\omega) \geq 0$.

(3) As $\tilde{\mathcal{U}}_S$ is continuous, $\tilde{\mathcal{U}}_S(\pi_n, \tilde{\sigma}_{\mathcal{A},n}) \geq \mathcal{U}_S^0$ for all n implies $\tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_{\mathcal{A}}) \geq \mathcal{U}_S^0$.

(4) By continuity of $\tilde{\mathcal{U}}_R$, $\tilde{\mathcal{U}}_R(\pi_n, \tilde{\sigma}_{\mathcal{A},n}) \geq \mathcal{U}_R$ for all n yields $\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) \geq \mathcal{U}_R$.
Hence, $C(\mathcal{U}_R)$ is a closed subset of the compact set $\Pi \times \mathcal{S}$, so $C(\mathcal{U}_R)$ is compact. □

The next auxiliary lemma is about non-emptiness of the values of the correspondence C:

Lemma A. 2. *There is some $\mathcal{U}_R^{max} \in [\mathcal{U}_R^0, \infty)$ so that $C\left(\bar{\mathcal{U}}_R\right) \neq \emptyset$ iff $\bar{\mathcal{U}}_R \leq \mathcal{U}_R^{max}$.*

Proof. Note that $C(\mathcal{U}_R(E^0)) \neq \emptyset$ since $(\pi^0, \sigma_{\mathcal{A}}^0(\pi^0, \cdot)) \in C(\mathcal{U}_R(E^0))$, where σ_a^0 is defined as in the proof of Lemma 2. Recall that E^0 is a fully revealing PBE with a pure-strategy information rule, so $(\pi^0, \sigma_{\mathcal{A}}^0(\pi^0, \cdot))$ satisfies (1) and (2). Also, $\tilde{\mathcal{U}}_S(\pi^0, \sigma_{\mathcal{A}}^0(\pi^0, \cdot)) = \mathcal{U}_S(E^0) \geq \mathcal{U}_S^0$

[31] There exists a sequence in ϵ^* converging to zero because $\tilde{a}(\mu)$ has at most countably many realizations that occur with positive probability.

[32] If such a subsequence does not exist, take any infinite subsequence $(\pi_{n_k}, \tilde{\sigma}_{\mathcal{A},n_k})_{k\in\mathbb{N}}$ of $(I_n, \tilde{\sigma}_{a,n})_{n\in\mathbb{N}}$ such that $\mu' \notin \text{supp}(\pi_{n_k})$ for any k, which must then exist, and switch the labels of μ and μ'.

and $\tilde{\mathcal{U}}_R(\pi^0, \sigma^0_{\mathcal{A}}(\pi^0,\cdot)) = \mathcal{U}_R(E^0) = \mathcal{U}^0_R$, that is, $(\pi^0, \sigma_{\mathcal{A}}(\pi^0,\cdot))$ fulfills (3) and (4). Second, $C(\mathcal{U}_R) \neq \emptyset$ implies $C(\mathcal{U}'_R) \neq \emptyset$ for all $\mathcal{U}'_R \leq \mathcal{U}_R$: If $(\pi^0, \sigma_{\mathcal{A}}(\pi^0,\cdot)) \in C(\mathcal{U}_R)$, then $\tilde{\mathcal{U}}_R(\pi^0, \sigma_{\mathcal{A}}(\pi^0,\cdot)) \geq \mathcal{U}_R \geq \mathcal{U}'_R$. Since the constraints (1)–(3) do not depend on $\mathcal{U}_R$, this yields $(\pi^0, \sigma_{\mathcal{A}}(\pi^0,\cdot)) \in C(\mathcal{U}'_R)$. Third, continuity of u_R and compactness of $A \times \Omega$ yield $\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) \leq \max_{(a,\omega)\in A\times\Omega} u_R(a,\omega) < \infty$ for any $(\pi, \tilde{\sigma}_{\mathcal{A}}) \in \Pi \times \mathcal{S}$. Hence, there exists some $\mathcal{U}^{\max}_R \in [\mathcal{U}^0_R, \infty)$ such that $C(\mathcal{U}_R) \neq \emptyset$ if $\mathcal{U}_R < \mathcal{U}^{\max}_R$ and $C(\mathcal{U}_R) = \emptyset$ if $\mathcal{U}_R > \mathcal{U}^{\max}_R$. Finally, let's verify that $C(\mathcal{U}^{\max}_R) \neq \emptyset$: Consider an increasing sequence $(\mathcal{U}_{R,n})_{n\in\mathbb{N}}$ with limit $\mathcal{U}^{\max}_R$, and a sequence $(\pi_n, \tilde{\sigma}_{\mathcal{A},n})_{n\in\mathbb{N}}$ satisfying $\tilde{\mathcal{U}}_R(\pi_n, \tilde{\sigma}_{\mathcal{A},n}) = \mathcal{U}_{R,n}$ and (1)–(3) for all n. Such a sequence exists since $C(\mathcal{U}_R) \neq \emptyset$ for all $\mathcal{U}_R < \mathcal{U}^{\max}_R$. By compactness of $\Pi \times \mathcal{S}$, there exists a convergent subsequence of $(\pi_n, \tilde{\sigma}_{\mathcal{A},n})_{n\in\mathbb{N}}$ with limit $(\pi, \tilde{\sigma}_{\mathcal{A}})$. By continuity of $\tilde{\mathcal{U}}_R$, it follows that $\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) = \mathcal{U}^{\max}_R$. From the proof of Lemma A. 1, $(\pi, \tilde{\sigma}_{\mathcal{A}})$ satisfies (1)–(3) so that $(\pi, \tilde{\sigma}_{\mathcal{A}}) \in C(\mathcal{U}^{\max}_R)$. □

The third auxiliary lemma establishes continuity of C:

Lemma A. 3. *C is continuous in $\mathcal{U}_R$ on $(-\infty, \mathcal{U}^{max}_R]$.*

Proof. I show that C is both upper and lower hemicontinuous at any $\mathcal{U}_R \in (-\infty, \mathcal{U}^{\max}_R]$: Take any sequence $(\mathcal{U}_{R,n})_n$ converging to $\mathcal{U}_R$, any $(\pi, \tilde{\sigma}_{\mathcal{A}}) \in \Pi \times \mathcal{S}$, and any sequence $(\pi_n, \tilde{\sigma}_{\mathcal{A},n})_n$ so that $(\pi_n, \tilde{\sigma}_{\mathcal{A},n}) \in C(\mathcal{U}_{R,n})$ for all n converging to $(\pi, \tilde{\sigma}_{\mathcal{A}})$. Suppose $(\pi, \tilde{\sigma}_{\mathcal{A}}) \notin C(\mathcal{U}_{R,n})$. By Lemma A.1, since $(\pi_n, \tilde{\sigma}_{\mathcal{A},n})_n$ satisfies (1)–(3), so does its limit. Consequently, $(\pi, \tilde{\sigma}_{\mathcal{A}}) \notin C(\mathcal{U}_R)$ means that $\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) < \mathcal{U}_R$. Continuity of $\tilde{\mathcal{U}}_R$ yields $\tilde{\mathcal{U}}_R(\pi_n, \tilde{\sigma}_{\mathcal{A},n}) < \mathcal{U}_{R,n}$ for large n — a contradiction. This proves upper hemicontinuity. Take again a sequence $(\mathcal{U}_{R,n})_n$ converging to $\mathcal{U}_R$ and any $(\pi, \tilde{\sigma}_{\mathcal{A}}) \in C(\mathcal{U}_R)$. If $\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) > \mathcal{U}_R$, then $(\pi, \tilde{\sigma}_{\mathcal{A}}) \in C(\mathcal{U}_{R,n})$ for large n. So there is a subsequence $(\mathcal{U}_{R,n_l})_l$ and a sequence $(\pi_l, \tilde{\sigma}_{\mathcal{A},l})_l$ with $(\pi_l, \tilde{\sigma}_{\mathcal{A},l}) = (\pi, \tilde{\sigma}_{\mathcal{A}})$ and $\tilde{\mathcal{U}}_R(\pi_l, \tilde{\sigma}_{\mathcal{A},l}) > (\mathcal{U}_{R,l})$ for all l, which, by construction, converges to $(\pi, \tilde{\sigma}_{\mathcal{A}})$. If $\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) = \mathcal{U}_R$, there is, by continuity of $\tilde{\mathcal{U}}_R$, some $(\pi_n, \tilde{\sigma}_{\mathcal{A},n})_n$ satisfying (1)–(3) and $\tilde{\mathcal{U}}_R(\pi_n, \tilde{\sigma}_{\mathcal{A},n}) = \mathcal{U}_{R,n}$ for all n with limit $(\pi, \tilde{\sigma}_{\mathcal{A}})$. This proves lower hemicontinuity. □

So the objective function $\tilde{\mathcal{U}}_S$ is continuous in $(\pi, \tilde{\sigma}_{\mathcal{A}})$, and the correspondence C is compact-valued, non-empty and continuous on $(-\infty, \mathcal{U}^{\max}_R]$. By Berge's maximum theorem, $\arg\max_{(\pi,\tilde{\sigma}_{\mathcal{A}})\in C(\mathcal{U}_R)} \tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_{\mathcal{A}})$ is non-empty for all $\mathcal{U}_R \in (-\infty, \mathcal{U}^{\max}_R]$, and $\mathcal{U}^*_S(\bar{\mathcal{U}}_R) \equiv \max_{(\pi,\tilde{\sigma}_{\mathcal{A}})\in C(\mathcal{U}_R)} \tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_{\mathcal{A}})$ is continuous in $\mathcal{U}_R$ on $(-\infty, \mathcal{U}^{\max}_R]$. This proves the existence part of the theorem. To prove the compactness part, notice first that the set of best implementable payoff profiles is $\bigcup_{\mathcal{U}_R\in(-\infty,\mathcal{U}^{\max}_R]} \mathcal{U}^*_S(\mathcal{U}_R)$. Moreover, observe that $\min_{(a,\omega)\in A\times\Omega} u_S(a,\omega) \leq \mathcal{U}^*_S(\mathcal{U}_R) \leq \max_{(a,\omega)\in A\times\Omega} u_S(a,\omega)$ for all $\mathcal{U}_R \in (-\infty, \mathcal{U}^{\max}_R]$, where the min- and max-function are well-defined by continuity of u_S and compactness of $A \times \Omega$. As a result, continuity of $\mathcal{U}^*_S$ implies that $\bigcup_{\mathcal{U}_R\in(-\infty,\mathcal{U}^{\max}_R]} \mathcal{U}^*_S(\mathcal{U}_R)$ is compact.

Proof of Theorem 3 Analogously to Theorem 2, the proof proceeds by applying Berge's maximum theorem to the optimization problem on page 13. Details are therefore omitted.

Proof of Theorem 4 Let $(\mathcal{U}^*_S, \mathcal{U}^*_R)$ be the unique best feasible payoff profile. By feasibility, there is some $(\bar{\sigma}_{\mathcal{I}}, \bar{\sigma}_{\mathcal{A}})$ so that $\mathcal{U}^*_i = \bar{\mathcal{U}}_i(\bar{\sigma}_{\mathcal{I}}, \bar{\sigma}_{\mathcal{A}})$ for all i. First, the strategy profile $\left(\left(\bar{\sigma}_I, \sigma^{\mathrm{FR}}_m\right), \bar{\sigma}_a\right)$ together with some beliefs μ_R form a PBE E: Suppose the receiver has an incentive to deviate. Then, there is some $\sigma_{\mathcal{A}}$ with $\mathcal{U}_R\left(\left\{\left(\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{M}}^{\text{FR}}\right),\sigma_{\mathcal{A}}\right),\mu_R\right\}\right) > \mathcal{U}_R^*$. As $\sigma_{\mathcal{M}}^{\text{FR}}$ is fully revealing, $\mathcal{U}_i\left(\left\{\left(\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{M}}^{\text{FR}}\right),\sigma_{\mathcal{A}}\right),\mu_R\right\}\right) = \bar{\mathcal{U}}_i\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{A}}\right)$ for all i, i.e., the payoff profile induced by the receiver's deviation is feasible. This yields $\bar{\mathcal{U}}_S\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{A}}\right) < \mathcal{U}_S^*$ as $\left(\bar{\mathcal{U}}_S\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{A}}\right),\bar{\mathcal{U}}_R\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{A}}\right)\right)$ is feasible and would otherwise strictly dominate $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ — a contradiction since $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is best feasible. By compactness of the set of best feasible payoff profiles (cf. Theorem 3), there exists thus a best feasible payoff profile $(\mathcal{U}_S',\mathcal{U}_R')$ with $\mathcal{U}_R' \geq \bar{\mathcal{U}}_R\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{A}}\right) > \mathcal{U}_R^*$ and $\mathcal{U}_S' \leq \bar{\mathcal{U}}_S\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{A}}\right) < \mathcal{U}_S^*$, contradicting uniqueness of the best feasible payoff profile. Now, suppose the sender has a profitable deviation $(\sigma_{\mathcal{I}}\sigma_{\mathcal{M}})$, that is, $\mathcal{U}_S\left(\left\{\left(\left(\sigma_{\mathcal{I}},\sigma_{\mathcal{M}}\right),\bar{\sigma}_{\mathcal{A}}\right),\mu_R\right\}\right) > \mathcal{U}_S^*$. Notice that there is some $\hat{\sigma}_{\mathcal{A}}$ so that $\mathcal{U}_i\left(\left\{\left(\left(\sigma_{\mathcal{I}},\sigma_{\mathcal{M}}\right),\bar{\sigma}_{\mathcal{A}}\right),\mu_R\right\}\right) = \bar{\mathcal{U}}_i\left(\sigma_{\mathcal{I}},\hat{\sigma}_{\mathcal{A}}\right)$ for all i, i.e., the payoff profile induced by the sender's deviation is feasible. One can conclude from this that either $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is not best feasible, or there exists another best feasible payoff profile, contradicting uniqueness. By construction, it holds that $\mathcal{U}_i\left(E\right) = \mathcal{U}_i^*$ for all i. Hence, $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is implementable. Since $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is also best feasible, it is best implementable. Finally, suppose $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is not the unique best implementable payoff profile, i.e., there is another best implementable payoff profile $(\mathcal{U}_S'',\mathcal{U}_R'')$. But then, since $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is also feasible, there exists another best feasible payoff profile by Theorem 3 — a contradiction.

Proof of Theorem 5 Fix some best feasible payoff profile $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ and suppose that the conditions stated in Theorem 5 hold true. In particular, $\text{supp}(\bar{\sigma}_{\mathcal{I}}) = \left\{\pi^{\text{full}}\right\}$ implies that the best feasible payoff profile for agent i is generated if the social planner fully learns the state and chooses the optimal action for agent i per state, i.e., this payoff profile is $\left(\int_\Omega u_S\left(a_i^*\left(\omega\right),\omega\right)\, d\mu_0\left(\omega\right) - c\left(\pi^{\text{full}}\right), \int_\Omega u_R\left(a_i^*\left(\omega\right),\omega\right)\, d\mu_0\left(\omega\right)\right)$. Since there are at least two best feasible payoff profiles, the sender-optimal and receiver-optimal ones do not coincide so that the set of states with $a_S^*\left(\omega\right) \neq a_R^*\left(\omega\right)$ is of positive measure. If $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is implementable, there is a PBE $\left\{\left(\left(\bar{\sigma}_{\mathcal{I}},\sigma_{\mathcal{M}}^{\text{FR}}\right),\bar{\sigma}_{\mathcal{A}}\right),\mu_R\right\}$ for some such $(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}})$. However, if $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is the receiver-optimal feasible payoff profile, the sender has an incentive to misreport in the communication stage to generate the sender-optimal outcome. Similarly, if $(\mathcal{U}_S^*,\mathcal{U}_R^*)$ is any other best feasible payoff profile, the receiver has an incentive to deviate in the action stage to generate the receiver-optimal outcome.

Proof of Corollary 1 The "if" direction holds due to Theorem 4. The "only if" direction follows from Theorem 5: Fix a best feasible payoff profile $(\mathcal{U}_S^*,\mathcal{U}_R^*)$. From the optimization problem on page 13, it follows that there is some $(\bar{\sigma}_{\mathcal{I}}^*,\bar{\sigma}_{\mathcal{A}}^*)$ solving $\max_{(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}})} \alpha\bar{\mathcal{U}}_S\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right) + (1-\alpha)\bar{\mathcal{U}}_R\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right)$, with $\alpha \in [0,1]$, so that $\mathcal{U}_i^* = \bar{\mathcal{U}}_i\left(\bar{\sigma}_{\mathcal{I}}^*,\bar{\sigma}_{\mathcal{A}}^*\right)$ for all i. Since $c = 0$, one obtains that

$$
\begin{aligned}
&\alpha\bar{\mathcal{U}}_S\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right) + (1-\alpha)\bar{\mathcal{U}}_R\left(\bar{\sigma}_{\mathcal{I}},\bar{\sigma}_{\mathcal{A}}\right)\\
&\quad = \int_\Pi \int_{\triangle\Omega} \int_\Omega \int_A \alpha u_S(a,\omega) + (1-\alpha)u_R(a,\omega)\; d\bar{\sigma}_{\mathcal{A}}(a|\pi,\mu)\; d\mu(\omega)\; d\pi\left(\mu\right)\; d\bar{\sigma}_{\mathcal{I}}(\pi).
\end{aligned}
$$

Since cost are zero and $A^*(\omega,\alpha u_S + (1-\alpha)u_R) \cap A^*(\omega',\alpha u_S + (1-\alpha)u_R) = \emptyset$ for all $\omega \neq \omega'$, the objective function is maximized if and only if $\text{supp}(\bar{\sigma}_{\mathcal{I}}) = \{\pi^{\text{full}}\}$.

Proofs to Section 4: Posterior Separable Cost

Proof of Theorem 6 Let $(\mathcal{U}_S, \mathcal{U}_R)$ be best implementable.

(i) Suppose in any fully revealing PBE E with a pure-strategy information rule, the sender chooses an experiment π that is not a convex combination of two experiments with strongly independent outcomes. In particular, π does not have strongly independent outcomes as $\pi = \alpha \cdot \pi + (1-\alpha) \cdot \pi$ for any $\alpha \in [0,1]$. Consequently, there is some $p \in (0,1)$ and some $\lambda_1, \lambda_2, \lambda_1', \lambda_2' \in \triangle(\triangle\Omega)$ with $p\lambda_i + (1-p)\lambda_i' = \pi$ for each i so that λ_1 and λ_2 are not equal almost everywhere, but $\int_\Pi \nu \; d\lambda_1(\nu) = \int_\Pi \nu \; d\lambda_2(\nu)$. Any experiment without strongly independent outcomes is a convex combination of to other experiments: Let $\pi_1 \equiv p\lambda_2 + (1-p)\lambda_1'$ be the experiment resulting from π by replacing mass p of the distribution over outcomes λ_1 by mass p of λ_2, and $\pi_2 \equiv p\lambda_1 + (1-p)\lambda_2'$. Note that π can be expressed as a convex combination[33] of π_1 and π_2 as $\frac{\pi_1+\pi_2}{2} = \frac{p\lambda_1 + (1-p)\lambda_1' + \lambda_2 + (1-p)\lambda_2'}{2} = \pi$. By assumption, π_1 or π_2 does not have strongly independent outcomes, so for concreteness, suppose π_2 does so. By the same procedure as above, there are two other experiments $\pi_3 \equiv q\lambda_4 + (1-q)\lambda_3'$ and $\pi_4 \equiv q\lambda_3 + (1-q)\lambda_4'$ for some $q \in (0,1)$ and some $\lambda_3, \lambda_4, \lambda_3', \lambda_4' \in \triangle(\triangle\Omega)$ with λ_3 and λ_4 not being equal a.e. so that $\pi_2 = \frac{\pi_3+\pi_4}{2}$. Thus, $\pi = \frac{2\pi_1+\pi_3+\pi_4}{4}$. Let Π^{conv} be the set of all convex combinations $\sum_{i\in\{1,3,4\}} \alpha_i \pi_i$. Implementability of $(\mathcal{U}_S, \mathcal{U}_R)$ implies existence of some $\tilde{\sigma}_{\mathcal{A}}$ so that $\mathcal{U}_i = \tilde{\mathcal{U}}_i(\pi, \tilde{\sigma}_{\mathcal{A}})$ for all i. Moreover, it implies implementability of $\left(\tilde{\mathcal{U}}_S(\pi', \tilde{\sigma}_{\mathcal{A}}), \tilde{\mathcal{U}}_R(\pi', \tilde{\sigma}_{\mathcal{A}})\right)$ as $\text{supp}(\pi') \subseteq \text{supp}(\pi)$ ensures that (2) holds true for any $\pi' \in \Pi^{\text{conv}}$. Denote by $\Pi(\mathcal{U}_R)$ the set of experiments $\pi' \in \Pi^{\text{conv}}$ with $\tilde{\mathcal{U}}_R(\pi', \tilde{\sigma}_{\mathcal{A}}) = \mathcal{U}_R$. By linearity of $\tilde{\mathcal{U}}_R(\cdot, \tilde{\sigma}_{\mathcal{A}})$ on Π^{conv}, this set forms a straight line in the 2-simplex Π^{conv}. Linearity of $\tilde{\mathcal{U}}_S(\cdot, \tilde{\sigma}_{\mathcal{A}})$ on Π^{conv} implies that one of the endpoints of this line π'' maximizes $\tilde{\mathcal{U}}_S(\cdot, \tilde{\sigma}_{\mathcal{A}})$ on Π^{conv} on that line. Moreover, $\alpha_1\alpha_3\alpha_4 = 0$ at any such endpoint, that is, π'' a convex combination of π_1 and π_3, π_1 and π_4, or π_3 and π_4. By construction, $\tilde{\mathcal{U}}_S(\pi'', \tilde{\sigma}_{\mathcal{A}}) \geq \mathcal{U}_S$ as $\pi'' \in \Pi(\mathcal{U}_R)$. If $\tilde{\mathcal{U}}_S(\pi'', \tilde{\sigma}_{\mathcal{A}}) > \mathcal{U}_S$, $(\mathcal{U}_S, \mathcal{U}_R)$ is dominated by the implementable payoff profile $\left(\tilde{\mathcal{U}}_S(\pi'', \tilde{\sigma}_{\mathcal{A}}), \mathcal{U}_R\right)$ contradicting best implementability of the former one. If $\tilde{\mathcal{U}}_S(\pi'', \tilde{\sigma}_{\mathcal{A}}) = \mathcal{U}_S$, then $(\mathcal{U}_S, \mathcal{U}_R)$ is generated by a fully revealing PBE with a pure-strategy information rule in which the sender chooses the experiment π'' being a convex combination of two experiments. If these two experiments have strongly independent outcomes, the proof is complete. If not, one can repeat the above argumentation taking $\pi^1 = \pi''$ as the new initial experiment instead of π. One can rerun this procedure as many times until either $(\mathcal{U}_S, \mathcal{U}_R)$ is shown to be dominated or until one finds an experiment π''^n satisfying the properties mentioned in (i) of the theorem's statement. In particular, this process ends eventually: If $(\mathcal{U}_S, \mathcal{U}_R)$ is not shown to be dominated in any (finite) round n and the constructed experiment π''^n does not satisfy (i), then $\tilde{\mathcal{U}}_S(\cdot, \tilde{\sigma}_{\mathcal{A}})$ and $\tilde{\mathcal{U}}_R(\cdot, \tilde{\sigma}_{\mathcal{A}})$ are both constant on $\triangle(\triangle\Omega)$. Consequently, $(\mathcal{U}_S, \mathcal{U}_R)$ is generated by a fully revealing PBE with a pure-strategy information rule in which the sender chooses any experiment in $\triangle(\triangle\Omega)$, so particularly any experiment with strongly independent outcomes (such as π^0).

[33] A convex combination of two experiment π and π' is defined as $\alpha\pi + (1-\alpha)\pi'$ for some $\alpha \in [0,1]$.

(ii) Suppose the sender chooses an experiment π not having weakly independent outcomes in any fully revealing PBE E with a pure-strategy information rule. Let μ^* be the set of outcomes $\mu \in \text{supp}(\pi)$ with $\mu = \int_\Pi \nu \, d\lambda_\mu(\nu)$ for some λ_μ not being equal to δ_μ a.e., and let λ_{μ^*} be the conditional distribution over outcomes of π given $\mu \in \mu^*$. Since μ^* is of positive measure $p \in (0,1]$, there exists some $\lambda' \in \triangle(\triangle\Omega)$ with $\pi = p\lambda_{\mu^*} + (1-p)\lambda'$. Let $\pi' \equiv p \int_{\mu^*} \lambda_\mu \, d\lambda_{\mu^*}(\mu) + (1-p)\lambda'$ be the experiment that results from replacing mass p of outcomes in μ^* by the corresponding distribution over outcomes λ_μ. By construction, π' has weakly independent outcomes. Furthermore, π is a Blackwell garbling of π'. Implementability of $(\mathcal{U}_S, \mathcal{U}_R)$ yields existence of some $\tilde{\sigma}_a$ so that $\mathcal{U}_i = \tilde{\mathcal{U}}_i(\pi, \tilde{\sigma}_\mathcal{A})$ for all i and implementability of $\left(\tilde{\mathcal{U}}_S(\pi', \tilde{\sigma}_\mathcal{A}), \tilde{\mathcal{U}}_R(\pi', \tilde{\sigma}_\mathcal{A})\right)$ because $\text{supp}(\pi') \subseteq \text{supp}(\pi)$. Incentive compatibility (2) of $(\pi, \tilde{\sigma}_\mathcal{A})$ implies $\tilde{\mathcal{U}}_S(\pi', \tilde{\sigma}_\mathcal{A}) = \tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_\mathcal{A})$ since cost are zero. Since π' is a Blackwell garbling of π, one gets $\tilde{\mathcal{U}}_R(\pi', \tilde{\sigma}_\mathcal{A}) \geq \tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_\mathcal{A})$. So $(\mathcal{U}_S, \mathcal{U}_R)$ is either dominated by $(\tilde{\mathcal{U}}_S(\pi', \tilde{\sigma}_\mathcal{A}), \tilde{\mathcal{U}}_R(\pi', \tilde{\sigma}_\mathcal{A}))$, or it is generated by a fully revealing PBE with a pure-strategy information rule in which the sender chooses the experiment π'.

Proof of Corollary 2 Suppose $\Omega = \{\omega_1, \ldots, \omega_n\}$ for some $n \in \mathbb{N}$.

(i) Verify that $|\text{supp}(\pi)| \leq |\Omega| = n$ for all π with strongly independent outcomes. If not, there exists some π with strongly independent outcomes and $|\text{supp}(\pi)| \geq n+1$. There are at least $n+1$ different distributions over outcomes $\lambda_1, \ldots, \lambda_{n+1} \in \triangle(\triangle\Omega)$ of positive measure: $\sum_{i=1}^{n+1} \alpha_i \lambda_i + \left(1 - \sum_{i=1}^{n+1} \alpha_i\right) \lambda' = \pi$ for some $\alpha_1, \ldots, \alpha_{n+1} \geq 0$ with $\sum_{i=1}^{n+1} \alpha_i \leq 1$. Define $v_i \equiv \left(\int_{\Pi_I} \mu(\omega_1) \, d\lambda_i(\mu), \ldots, \int_{\Pi_I} \mu(\omega_n) \, d\lambda_i(\mu)\right)'$ for all $i \in \{1, \ldots, n+1\}$. As $\int_\Pi \mu(\omega) \, d\lambda_i(\mu) = 0$ for all $\omega \notin \Omega$, the posterior distribution of the state given the distribution over outcomes λ_i is fully characterized by v_i. Since $n+1$ n-dimensional vectors are linearly dependent, there exist some numbers $\beta_1, \ldots, \beta_{n+1} \in \mathbb{R}$, not all zero, so that $\sum_{i=1}^{n+1} \beta_i \cdot v_i = 0_n$. Note that $\sum_{i=1}^{n+1} \beta_i = 0$ as $\sum_{j=1}^{n} \int_\Pi \mu(\omega_j) \, d\lambda_i(\mu) = 1$ for all i. In particular, there is some i with $\beta_i < 0$. Rearranging $\sum_{i=1}^{n+1} \beta_i \cdot v_i = 0_n$ yields $\sum_{i:\beta_i<0} \frac{\beta_i}{-\sum_{j:\beta_j<0} \beta_j} \cdot v_i = \sum_{i:\beta_i \geq 0} \frac{\beta_i}{\sum_{j:\beta_j<0} \beta_j} \cdot v_i$, where $\frac{\beta_i}{-\sum_{j:\beta_j<0} \beta_j} \in [0,1]$ for all $\beta_i < 0$ and $\frac{\beta_i}{\sum_{j:\beta_j<0} \beta_j} \in [0,1]$ for all $\beta_i \geq 0$. Hence, there are two convex combinations (over convex combinations) over outcomes which can be represented by one another, that is, π does not exhibit strongly independent outcomes. With that, it follows from Theorem 6 (i) that any best implementable payoff profile is generated by a fully revealing PBE in which the sender chooses a experiment with at most $n + n = 2n$ different outcomes.

(ii) If $n = 2$, any experiment π with weakly independent outcomes can have most 2 different outcomes. If not, there are at least 3 different outcomes $\mu_1, \mu_2, \mu_3 \in \text{supp}(\pi)$. Define $w_i \equiv (\mu_i(\omega_1), \mu_i(\omega_2))'$ for each i, and notice that $\mu_i(\omega) = 0$ for all $\omega \notin \Omega$. The 2-dimensional vectors w_1, w_2 and w_3 are linearly independent implying existence of some $\beta_1, \beta_2, \beta_3 \in \mathbb{R}$, not all zero, so that $\sum_{i=1}^{3} \beta_i \cdot w_i = 0_2$. Without loss of generality, suppose $\beta_1 < 0$ and $\beta_2, \beta_3 \geq 0$. Rearranging $\sum_{i=1}^{3} \beta_i \cdot w_i = 0_2$ yields $w_1 = \sum_{i=2}^{3} \frac{\beta_i}{-\beta_1} \cdot w_i$, where $\frac{\beta_i}{-\beta_1} \geq 0$ for $i = 2, 3$ and $\sum_{i=2}^{3} \frac{\beta_i}{-\beta_1} = 1$ as $\sum_{i=1}^{3} \beta_i = 0$. Since this holds true for any three outcomes, there is a positive measure of outcomes in $\text{supp}(\pi)$ that can be represented as a convex combination of other outcomes in supp(π) — a contradiction. If cost are zero, an immediate consequence of Theorem 6 (ii) is that any best implementable payoff profile is generated in a fully revealing PBE in which the sender chooses some π with supp(π) ≤ 2.

Proofs to Section 5: Partially Separable Utility

Proof of Lemma 3 Denote $\bar{\omega}(\mu) = \int_\Omega \omega \, d\mu(\omega)$. By Assumption 4, it holds that

$$\int_\Omega u_i(a,\omega) \; d\mu(\omega) = u_{i,1}(a) + u_{i,2}(a) \cdot \bar{\omega}(\mu) + \int_\Omega u_{i,3}(\omega) \, d\mu(\omega) \tag{12}$$

for all $\mu \in \triangle\Omega$ and each i. Now fix any $\mu, \mu' \in \triangle\Omega$ with $\bar{\omega}(\mu) = \bar{\omega}(\mu')$. Condition (12) implies that $\arg\max_{a\in A} \int_\Omega u_R(a,\omega) \; d\mu(\omega) = \arg\max_{a\in A} (u_{R,1}(a) + u_{R,2}(a) \cdot \bar{\omega}(\mu))$ and thus $\arg\max_{a\in A} \int_\Omega u_R(a,\omega) \; d\mu(\omega) = \arg\max_{a\in A} \int_\Omega u_R(a,\omega) \; d\mu'(\omega)$. Next, take any $a, a' \in A$, and notice that by (12), it follows that $\int_\Omega u_S(a,\omega) \; d\mu(\omega) \geq \int_\Omega u_S(a,\omega) \; d\mu(\omega)$ is equivalent to $u_{S,1}(a) + u_{S,2}(a) \cdot \bar{\omega}(\mu) \geq u_{S,1}(a') + u_{S,2}(a') \cdot \bar{\omega}(\mu)$, and thus also equivalent to $\int_\Omega u_S(a,\omega) \; d\mu'(\omega) \geq \int_\Omega u_S(a,\omega) \; d\mu'(\omega)$. Finally, for all $\mu \in$ supp(π) and $\omega \in \Omega$, let $v_{i,1}(\mu) \equiv \int_A u_{i,1}(a) + u_{i,2}(a) \cdot \bar{\omega}(\mu) \; d\sigma_{\mathcal{A}}(a|\pi,\mu)$ and $v_{i,2}(\omega) \equiv u_{i,3}(\omega)$. With that, (6) can be easily verified.

Proof of Corollary 3 The proof proceeds similarily as the proof of Theorem 6. Let $\left(\tilde{\mathcal{U}}_S(\pi,\tilde{\sigma}_{\mathcal{A}}), \tilde{\mathcal{U}}_R(\pi,\tilde{\sigma}_{\mathcal{A}})\right)$ be best implementable, where $\pi \in \bar{\Pi}$.

(i) Suppose that π is not a convex combination of two bi-pooling policies, so in particular, π is not a bi-pooling policy. That is, for its finest partitioning $\{[\underline{\omega}_j, \bar{\omega}_j] | j \in J\}$, there is some $j \in J$ with $(\underline{\omega}_j, \bar{\omega}_j) \cap \text{supp}(\omega) \neq \emptyset$ such that the set of outcomes μ_j^* that realize if $\tilde{\omega} \in (\underline{\omega}_j, \bar{\omega}_j)$ is of cardinality 3 or larger. Let $\bar{\lambda}$ be the distribution over posterior means of the state under experiment π. As shown by Kleiner et al. (2021) and Arieli et al. (2020), $\bar{\lambda}$ is not an extreme point within the set of possible[34] distributions over posterior means, that is, there exist some $\bar{\lambda}_1$, $\bar{\lambda}_2$ with $\bar{\lambda} = \frac{1}{2} \cdot \bar{\lambda}_1 + \frac{1}{2} \cdot \bar{\lambda}_2$. Denote π_i an experiment generating the distribution over posterior means $\bar{\lambda}_i$ for any i. In addition, if π_2 is no bi-pooling policy either, there exist experiments π_3 and π_4 such that $\bar{\lambda}_2 = \frac{1}{2} \cdot \bar{\lambda}_3 + \frac{1}{2} \cdot \bar{\lambda}_4$. This yields $\bar{\lambda} = \frac{1}{2} \cdot \bar{\lambda}_1 + \frac{1}{4} \cdot \bar{\lambda}_3 + \frac{1}{4} \cdot \bar{\lambda}_4$. Since both agent's incentives and payoffs depend on the distribution over posterior means only by Lemma 3, there is some $\tilde{\sigma}'_{\mathcal{A}}$ (which is constructed by $\tilde{\sigma}'_{\mathcal{A}}(\mu) = \tilde{\sigma}_{\mathcal{A}}(\mu')$ for some $\mu' \in$ supp(π) with $\int_\Omega \omega \; d\mu(\omega) = \int_\Omega \omega \; d\mu'(\omega)$ for all $\mu \in \bigcup_{i\in\{1,3,4\}} \text{supp}(\pi_i)$) such that $\left(\tilde{\mathcal{U}}_S(\pi',\tilde{\sigma}'_{\mathcal{A}}), \tilde{\mathcal{U}}_R(\pi',\tilde{\sigma}'_{\mathcal{A}})\right)$ is implementable for all $\pi' \in \Pi^{\text{conv}}$ and $\tilde{\mathcal{U}}_i\left(\frac{1}{2} \cdot \pi_1 + \frac{1}{4} \cdot \pi_3 + \frac{1}{4} \cdot \pi_4, \tilde{\sigma}'_{\mathcal{A}}\right) = \tilde{\mathcal{U}}_i(\pi,\tilde{\sigma}_{\mathcal{A}})$ for all $i \in \{S, R\}$. The rest of the proof follows exactly the same steps as the proof of Theorem 6(i).

(ii) As argued in part (i), if π is no bi-pooling policy, there exist some π_1, π_2 such that $\bar{\lambda} = \frac{1}{2} \cdot \bar{\lambda}_1 + \frac{1}{2} \cdot \bar{\lambda}_2$. Moreover, $\bar{\lambda}$ is a Blackwell garbling of one of these two distributions, say of $\bar{\lambda}_1$ (see Arieli et al. (2020)). Using the arguments from the proof of part (i) and following

[34] given the prior distribution of the state

the same steps as in the proof of Theorem 6(ii), one obtains that $\left(\tilde{\mathcal{U}}_S(\pi_1, \tilde{\sigma}'_{\mathcal{A}}), \tilde{\mathcal{U}}_R(\pi_1, \tilde{\sigma}'_{\mathcal{A}})\right)$ is implementable and either dominates the payoff profile $\left(\tilde{\mathcal{U}}_S(\pi, \tilde{\sigma}_{\mathcal{A}}), \tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}})\right)$ or is equal to the latter one.

Proofs to Section 6: Uniform Quadratic Case

Proof of Lemma 4 Solving $\max_{a\in A} \int_\Omega -(a-\omega)^2 \; d\mu(\omega)$ yields the first-order condition $a = \bar{\omega}(\mu)$, which is the global maximizer by strict concavity of the objective function in a.

Take any $\mu, \mu' \in \Pi_I$, and notice that the incentive compatibility constraints (2) of the sender to fully reveal μ instead of misreporting μ' reduce to

$$\begin{aligned}
0 &\le \int\limits_\Omega -(\bar{\omega}(\mu) - \omega - b)^2 + (\bar{\omega}(\mu') - \omega - b)^2 \; d\mu(\omega) \\
&= \int\limits_\Omega -(\bar{\omega}(\mu) - \omega - b)^2 + (\bar{\omega}(\mu') - \bar{\omega}(\mu) + \bar{\omega}(\mu) - \omega - b)^2 \, d\mu(\omega) \\
&= \int\limits_\Omega 2\,(\bar{\omega}(\mu') - \bar{\omega}(\mu))(\bar{\omega}(\mu) - \omega - b) + (\bar{\omega}(\mu') - \bar{\omega}(\mu))^2 \, d\mu(\omega) \\
&= -2b\,(\bar{\omega}(\mu') - \bar{\omega}(\mu)) + (\bar{\omega}(\mu') - \bar{\omega}(\mu))^2 \\
&= (\bar{\omega}(\mu') - \bar{\omega}(\mu) - 2b) \cdot (\bar{\omega}(\mu') - \bar{\omega}(\mu)) \,.
\end{aligned}$$

Hence, either $\int_\Omega \omega \; d\mu'(\omega) - \int_\Omega \omega \; d\mu(\omega) \ge 2b$ or $\int_\Omega \omega \; d\mu'(\omega) - \int_\Omega \omega \; d\mu(\omega) \le 0$ must hold. The sender has an incentive to truthfully reveal μ' instead of deviating to μ if $0 \le \int_\Omega -(\bar{\omega}(\mu') - \omega - b)^2 + (\bar{\omega}(\mu) - \omega - b)^2 \; d\mu'(\omega) = (\bar{\omega}(\mu) - \bar{\omega}(\mu') - 2b)\,(\bar{\omega}(\mu) - \bar{\omega}(\mu'))$, implying that either $\int_\Omega \omega \; d\mu(\omega) - \int_\Omega \omega \; d\mu'(\omega) \ge 2b$ or $\int_\Omega \omega \; d\mu(\omega) - \int_\Omega \omega \; d\mu'(\omega) \le 0$ must hold. These two incentive compatibility constraints yield $|\bar{\omega}(\mu) - \bar{\omega}(\mu')| \ge 2b$.

Ex-ante expected payoffs Notice that

$$\tilde{\mathcal{U}}_R(\pi, \tilde{\sigma}_{\mathcal{A}}) = \int_{\triangle\Omega} \int_\Omega -(\bar{\omega}(\mu) - \omega)^2 \; d\mu(\omega) \; d\pi(\mu) = -\sum_{i=1}^{n} p_i \mathrm{Var}(\omega|\mu_i)\,,$$

where $\mathrm{Var}(\omega|\mu_i)$ is the conditional distribution of the state given μ_i. If the experiment is a bi-pooling policy, these conditional variances take simple analytic forms: If μ_i forms a 1-partition the state is uniformly distributed on the interval $\left[\bar{\omega}_i - \frac{p_i}{2}, \bar{\omega}_i + \frac{p_i}{2}\right]$, so one obtains $\mathrm{Var}(\omega|\mu_i) = \frac{p_i^2}{12}$. Let $\Omega_i \subseteq \Omega$ be set of states associated with the 1-partition μ_i. To determine the conditional variance if μ_i and μ_{i+1} form a 2-partition, let $\Omega_{i,i+1} \subseteq \Omega$ be set of states associated with outcomes μ_i and μ_{i+1} (i.e., if the state lies in $\Omega_{i,i+1}$, the outcome is μ_i or μ_{i+1}) and let F_i (F_{i+1}) be the conditional distribution of the state given outcome μ_i (μ_{i+1}). The weighted sum $p_i F_i + p_{i+1} F_{i+1}$ is the conditional distribution of the state given $\omega \in \Omega_{i,i+1}$, which is the uniform distribution on $\Omega_{i,i+1}$. Consequently, its mean is $\inf\{\Omega_{i,i+1}\} + \frac{p_i + p_{i+1}}{2}$. Let $\underline{d}_i = \bar{\omega}_i - \inf\{\Omega_{i,i+1}\}$ and $\bar{d}_{i+1} = \sup\{\Omega_{i,i+1}\} - \bar{\omega}_{i+1}$. Simple algebraic transformations yield for the conditional variance of the state given the experiment's outcome is μ_i or μ_{i+1}

$$\begin{aligned}\text{Var}\left(\omega|\mu_i \text{ or } \mu_{i+1}\right) &= \frac{p_i \int (\omega - \bar{\omega}_i)^2 \, dF_i(\omega) + p_{i+1} \int (\omega - \bar{\omega}_{i+1})^2 \, dF_{i+1}(\omega)}{p_i + p_{i+1}} \\ &= \frac{(p_i + p_{i+1})^2}{12} - (\frac{p_i + p_{i+1}}{2} - \underline{d}_i)(\frac{p_i + p_{i+1}}{2} - \bar{d}_{i+1}).\end{aligned}$$

The proofs of Lemma 5, Lemma 6, Lemma 7 and Proposition 3 are deferred to an Online Appendix.

Proofs to Section 7: Model Variants

The proofs of Proposition 4-6 directly follow from the arguments in the main text.

Proof of Proposition 7 For any $b \leq \frac{1}{2}$, the receiver's ex-ante expected payoff under any mediation rule is bounded above by $-\frac{1}{3}b(1-b)$ (see Goltsman et al. (2009)). By Lemma 5, her ex-ante expected payoff in the unique best equilibrium under endogenous learning is bounded below by $-\frac{1}{12(n-1)^2}$, where n is the unique integer such that $b \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right]$. It can be easily verified that this lower bound exceeds the above upper bound.

Online Appendix

Omitted Proofs to Section 6: Uniform-Quadratic Case

Proof of Lemma 6 Fix $n \geq 3$ and $b \in \left(\frac{1}{2n}, \frac{1}{2(n-1)}\right]$.

First, I show that ignoring all incentive compatibility constraints, uniform partition are the best bi-pooling policies:

Fact 1. *The payoff-maximizing bi-pooling policy of size $k \in \mathbb{N}$ is the uniform partition of size k.*

Proof. Consider some non-monotone bi-pooling policy with outcomes $\mu_1, \ldots, \mu_k$ and probabilities $p_1, \ldots, p_k$. By non-monotonicity, there is some i such that μ_i and μ_{i+1} form a 2-partition. Now consider a new experiment with outcomes $\mu'_1, \ldots, \mu'_k$ and probabilities $p'_1, \ldots, p'_k$ that is constructed from the original experiment by splitting the 2-partition into two separate 1-partitions while keeping their weights equal: $p'_i = p_i$ and $p'_{i+1} = p_{i+1}$. Note that

$$
\begin{aligned}
&\sum_{i=1}^{k} p'_i \text{Var}\left(\omega|\mu'_i\right) - \sum_{i=1}^{k} p_i \text{Var}\left(\omega|\mu_i\right) \\
&= \frac{(p_i + p_{i+1})^3}{12} - (p_i + p_{i+1})(\frac{p_i + p_{i+1}}{2} - \frac{p_i}{2})(\frac{p_i + p_{i+1}}{2} - \frac{p_{i+1}}{2}) \\
&\quad - \left(\frac{(p_i + p_{i+1})^3}{12} - (p_i + p_{i+1})(\frac{p_i + p_{i+1}}{2} - \underline{d}_i)(\frac{p_i + p_{i+1}}{2} - \bar{d}_{i+1})\right) \\
&= -(p_i + p_{i+1})(\frac{p_i + p_{i+1}}{2} - \frac{p_i}{2})(\frac{p_i + p_{i+1}}{2} - \frac{p_{i+1}}{2}) + (p_i + p_{i+1})(\frac{p_i + p_{i+1}}{2} - \underline{d}_i)(\frac{p_i + p_{i+1}}{2} - \bar{d}_{i+1}) \\
&< 0
\end{aligned}
$$

since $\underline{d}_i > \frac{p_i}{2}$ and $\bar{d}_{i+1} > \frac{p_{i+1}}{2}$ as μ_i and μ_{i+1} form a 2-partitions. Hence, the payoff-maximizing bi-pooling policy must be monotone. Now suppose it was not uniform. But then, it holds that

$$\sum_{i=1}^{k} p_i \text{Var}\left(\omega|\mu_i\right) = \sum_{i=1}^{k} \frac{p_i^3}{12} > \sum_{i=1}^{k} \frac{(\frac{1}{k})^3}{12},$$

where the last term is the weighted conditional variance of the uniform partition of size k. The inequality follows from convexity of x^3 and the fact that $\sum_{i=1}^{k} p_i = \sum_{i=1}^{k} \frac{1}{k} = 1$. □

Next, notice that $\sum_{i=1}^{k-1} \frac{(\frac{1}{k})^3}{12} > \sum_{i=1}^{n-1} \frac{(\frac{1}{n-1})^3}{12}$ for all $k < n-1$, that is, the uniform partition of size $n-1$ yields strictly higher payoff than the uniform partition of size $k < n-1$. Moreover, the uniform partition of size $n-1$ is implementable because the incentive compatibility constraints are satisfied: The distance between any two adjacent posterior means of the state is $\frac{1}{n-1} \geq 2b$ as $b \leq \frac{1}{2(n-1)}$. Consequently, the optimal experiment is either of size n or the uniform partition of size $n-1$.

It remains to show that under an optimal bi-pooling policy of size n, the outcomes μ_1 and μ_n form 1-partitions and the incentive compatibility constraints of all adjacent posterior means are binding. The proof is completed by the following two facts:

Fact 2. *For any optimal bi-pooling policy, both* μ_1 *and* μ_n *are* 1*-partitions.*

Proof. Suppose first that μ_1 is no 1-partition, that is, μ_1 and μ_2 form a 2-partition. Hence, it holds that $p_1 + p_2 > 2 \cdot (\bar{\omega}_2 - \bar{\omega}_1)$ and $\max\left\{\underline{d}_1, \bar{d}_2\right\} < \frac{p_1+p_2}{2}$. The conditional variance of the state given μ_1 or μ_2 is

$$\text{Var}\left(\omega|\mu_1 \text{ or } \mu_2\right) = \frac{\left(p_1+p_2\right)^2}{12} - \left(\frac{p_1+p_2}{2} - \underline{d}_1\right) \cdot \left(\frac{p_1+p_2}{2} - \bar{d}_2\right).$$

Now construct a new information structure by decreasing $\bar{\omega}_1$ by some $\epsilon > 0$, decreasing p_1 by $\delta = \frac{p_1 \epsilon}{\bar{\omega}_2 - \bar{\omega}_1 + \epsilon} \in (0, p_1)$ and increasing p_2 by δ , keeping everything else unmodified.[35] Note that all incentive compatibility constraints remain satisfied because all posterior means except $\bar{\omega}_1$ are the same as under the original experiment, and $\bar{\omega}_1$ decreases, that is, the first incentive compatibility constraint becomes even less binding. As long as $\epsilon < \frac{p_1+p_2}{2} - (\bar{\omega}_2 - \bar{\omega}_1)$, μ_1 and μ_2 continue forming a 2-partition. Hence, the conditional variance of the state given μ_1 or μ_2 becomes

$$\widehat{\text{Var}}\left(\omega|\mu_1 \text{ or } \mu_2\right) = \frac{\left(p_1+p_2\right)^2}{12} - \left(\frac{p_1+p_2}{2} - (\underline{d}_1 - \epsilon)\right) \cdot \left(\frac{p_1+p_2}{2} - \bar{d}_2\right).$$

For all $\epsilon \in \left(0, \frac{p_1+p_2}{2} - (\bar{\omega}_2 - \bar{\omega}_1)\right)$, this variance is strictly smaller than under the original information structure because

$$\begin{aligned}
&\widehat{\text{Var}}\left(\omega|\mu_1 \text{ or } \mu_2\right) - \text{Var}\left(\omega|\mu_1 \text{ or } \mu_2\right) \\
&\quad = -\left(\frac{p_1+p_2}{2} - (\underline{d}_1 - \epsilon)\right) \cdot \left(\frac{p_1+p_2}{2} - \bar{d}_2\right) + \left(\frac{p_1+p_2}{2} - \underline{d}_1\right) \cdot \left(\frac{p_1+p_2}{2} - \bar{d}_2\right) \\
&\quad = -\epsilon \cdot \left(\frac{p_1+p_2}{2} - \bar{d}_2\right) < 0.
\end{aligned}$$

Since the conditional variance of the state given $\omega \notin \Omega_{1,2}$ and the set $\Omega_{1,2}$ (and thus the probability that $\tilde{\omega} \in \Omega_{1,2}$)) are the same under both information structures, the original bi-pooling policy is not optimal.

The proof for the fact that μ_n must be a 1-partition, too, works analogously. □

Fact 3. *Under any optimal bi-pooling policy of size* $n \in \left(\frac{1}{2b}, \frac{1}{2b} + 1\right] \cap \mathbb{N}$*, all incentive compatibility constraints are binding.*

Proof. Suppose not, i.e., $\bar{\omega} - \bar{\omega}_{i-1} > 2b$ for some $i \in \{2, \ldots, n\}$. There are four different cases to be considered:

- Case 1: (a) μ_{i-1} belongs to a 1-partition, and μ_i belongs to a 2-partition.
  (b) μ_{i-1} belongs to a 2-partition, and μ_i belongs to a 1-partition.
- Case 2: μ_{i-1} and μ_i belong to different 2-partitions.
- Case 3: μ_{i-1} and μ_i belong to the same 2-partition.
- Case 4: μ_{i-1} and μ_i both belong to a 1-partition.

[35] δ is chosen such that the conditional mean of the state given μ_1 or μ_2 remains unchanged, that is, it satisfies $p_1\bar{\omega}_1 + p_2\bar{\omega}_2 = (p_1 - \delta)(\bar{\omega}_1 - \epsilon) + (p_2 + \delta)\bar{\omega}_2$.

Proof of Case 1*(a)*: Since μ_i and μ_{i+1} form a 2-partition, it holds that $p_i+p_{i+1} > 2\cdot(\bar{\omega}_{i+1} - \omega_i)$ and $\max\left\{\underline{d}_i, \bar{d}_{i+1}\right\} < \frac{p_i+p_{i+1}}{2}$.

Now construct a new experiment by decreasing $\bar{\omega}_i$ by some $\epsilon > 0$, decreasing p_i by $\delta = \frac{p_i \epsilon}{\bar{\omega}_{i+1}-\bar{\omega}_i+\epsilon} \in (0, p_i)$ and increasing p_{i+1} by δ such that the incentive compatibility constraints remain satisfied, that is,

$$(\bar{\omega}_i - \epsilon) - \bar{\omega}_{i-1} \geq 2b \quad \Leftrightarrow \quad \epsilon \leq (\bar{\omega}_i - \bar{\omega}_{i-1}) - 2b,$$

and such that μ_i and μ_{i+1} continue forming a 2-partition, i.e.,

$$p_i + p_{i+1} > 2 \cdot (\bar{\omega}_{i+1} - (\bar{\omega}_i - \epsilon)) \quad \Leftrightarrow \quad \epsilon < \frac{p_i + p_{i+1}}{2} - (\bar{\omega}_{i+1} - \bar{\omega}_i).$$

The total decrease of the conditional variance of the state given μ_i or μ_{i+1} is

$$\epsilon \cdot \left(\frac{p_i + p_{i+1}}{2} - \bar{d}_{i+1}\right) > 0.$$

Proof of Case 1*(b)*: This works analogously the the proof of Case 1(a): One can find a better experiment by increasing $\bar{\omega}_{i-1}$ by some $\epsilon \in (0, \min\{(\bar{\omega}_{i-1} - \bar{\omega}_{i-2}) - 2b, \frac{p_{i-2}i+p_{i-1}}{2} - (\bar{\omega}_{i-1} - \bar{\omega}_{i-2})\})$, decreasing p_{i-1} by $\delta = \frac{p_{i-1}\epsilon}{\bar{\omega}_{i-1}-\bar{\omega}_{i-2}+\epsilon}$ and increasing p_{i-2} by δ.

Proof of Case 2: The proof of this case proceeds as in Case 1(a):

Decreasing $\bar{\omega}_i$ by some $\epsilon \in \left(0, \min\left\{(\bar{\omega}_i - \bar{\omega}_{i-1}) - 2b, \frac{p_i+p_{i+1}}{2} - (\bar{\omega}_{i+1} - \bar{\omega}_i)\right\}\right)$, decreasing p_i by $\delta = \frac{p_i\epsilon}{\bar{\omega}_{i+1}-\bar{\omega}_i+\epsilon}$ and increasing p_{i+1} by δ reduces the conditional variance of the state given μ_i or μ_{i+1} by

$$\epsilon \cdot \left(\frac{p_i + p_{i+1}}{2} - \bar{d}_{i+1})\right) > 0.$$

Proof of Case 3: First, note that $\bar{\omega}_{i-1} - \bar{\omega}_{i-2} = \bar{\omega}_{i+1} - \bar{\omega}_i = 2b$ by Case 1 and 2. From the fact that μ_{i-1} and μ_i form a 2-partition, one can conclude that $\underline{d}_{i-1} + \bar{d}_i > \bar{\omega}_i - \bar{\omega}_{i-1} > 2b$, implying that $\underline{d}_{i-1} > b$ or $\bar{d}_i > b$. Furthermore, since $p_{i-1} + p_i = \underline{d}_{i-1} + (\bar{\omega}_i - \bar{\omega}_{i-1}) + \bar{d}_i$, it follows that $p_{i-1} + p_i > 4b$ and $\underline{d}_{i-1} + \bar{d}_i > \frac{p_{i-1}+p_i}{2}$.

Let's now consider the case where $\underline{d}_{i-1} \leq \bar{d}_i$, so in particular it holds that $\bar{d}_i > b$:

If μ_{i+1} is a 1-partition, one can construct a new information structure by shifting the interval of states $\left(\bar{\omega}_i + \bar{d}_i - \epsilon, \bar{\omega}_i + \bar{d}_i\right)$ from outcomes μ_{i-1} or μ_i to outcome μ_{i+1} for some $\epsilon > 0$. Since μ_{i+1} is a 1-partition, $\bar{\omega}_{i+1}$ decreases by exactly $\frac{\epsilon}{2}$. Hence, $\bar{\omega}_i$ must decrease by at least $\frac{\epsilon}{2}$ in order to fulfill the ith incentive compatibility constraint. For concreteness, suppose that $\bar{\omega}_i$ decreases by exactly $\frac{\epsilon}{2}$ as well. To satisfy the $(i-1)$th incentive compatibility constraint, it is necessary that

$$\left(\bar{\omega}_i - \frac{\epsilon}{2}\right) - \bar{\omega}_{i-1} \geq 2b \quad \Leftrightarrow \quad \epsilon \leq \frac{\bar{\omega}_i - \bar{\omega}_{i-1}}{2} - b.$$

The conditional variance of the state given μ_{i-1} or μ_i or μ_{i+1} is given by

$$\begin{aligned}
&\frac{1}{\sum_{j=i-1}^{i+1} p_j} \cdot \left[\frac{(p_{i-1}+p_i-\epsilon)^3}{12} + \frac{(p_{i+1}+\epsilon)^3}{12} \right. \\
&\qquad \left. - (p_{i-1}+p_i-\epsilon) \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \underline{d}_{i-1} \right) \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \left(\bar{d}_i - \frac{\epsilon}{2} \right) \right) \right] \\
= &\frac{1}{\sum_{j=i-1}^{i+1} p_j} \cdot \left[\frac{(p_{i-1}+p_i-\epsilon)^3}{12} + \frac{(p_{i+1}+\epsilon)^3}{12} \right. \\
&\qquad \left. - (p_{i-1}+p_i-\epsilon) \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \underline{d}_{i-1} \right) \cdot \left(\frac{p_{i-1}+p_i}{2} - \bar{d}_i \right) \right]
\end{aligned}$$

Differentiating this term with respect to ϵ yields

$$\frac{1}{p_{i-1}+p_i+p_{i+1}} \cdot \left[-\frac{(p_{i-1}+p_i-\epsilon)^2}{4} + \frac{(p_{i+1}+\epsilon)^2}{4} + \left(\frac{p_{i-1}+p_i}{2} - \bar{d}i \right) \cdot (p_{i-1}+p_i-\epsilon-\underline{d}_{i-1}) \right].$$

At $\epsilon = 0$, this derivative reduces to

$$\frac{1}{p_{i-1}+p_i+p_{i+1}} \cdot \left[-\frac{(p_{i-1}+p_i)^2}{4} + \frac{p_{i+1}^2}{4} + \left(\frac{p_{i-1}+p_i}{2} - \bar{d}_i \right) \cdot (p_{i-1}+p_i-\underline{d}_{i-1}) \right].$$

Notice that $\underline{d}_{i-1} > \frac{p_{i-1}+p_i}{2} - \bar{d}_i$ holds due to the fact that μ_{i-1} and μ_i are a 2-partition, and that $p_{i+1} = 2 \cdot (2b - \bar{d}_i)$ because μ_{i+1} is a 1-partition, which implies that $p_{i+1} < 2b$ as $\bar{d}_i > b$. As a consequence, one obtains that the derivative is strictly negative at $\epsilon = 0$ because

$$\begin{aligned}
&-\frac{(p_{i-1}+p_i)^2}{4} + \frac{p_{i+1}^2}{4} + \left(\frac{p_{i-1}+p_i}{2} - \bar{d}_i \right) \cdot (p_{i-1}+p_i-\underline{d}_{i-1}) \\
&\quad < -\frac{(p_{i-1}+p_i)^2}{4} + \frac{p_{i+1}^2}{4} + \left(\frac{p_{i-1}+p_i}{2} - \bar{d}_i \right) \cdot \left(\frac{p_{i-1}+p_i}{2} + \bar{d}_i \right) = \frac{p_{i+1}^2}{4} - \bar{d}_i^2 < \frac{(2b)^2}{4} - b^2 = 0.
\end{aligned}$$

Since the derivative is continuous in ϵ, it is thus strictly negative on some non-empty open set around zero. So for any $\epsilon > 0$ sufficiently close to zero, the conditional variance of the state given μ_{i-1} or μ_i or μ_{i+1} is strictly smaller under the new information structure compared to the original bi-pooling policy (which corresponds to the case when $\epsilon = 0$), which therefore cannot be optimal.

If μ_{i+1} and μ_{i+2} form a 2-partition and the (i+1)th incentive compatibility constraint is binding, i.e., $\bar{\omega}_{i+2} - \bar{\omega}_{i+1} = 2b$, construct a new information structure by shifting the interval $(\bar{\omega}_i + \bar{d}_i - \epsilon, \bar{\omega}_i + \bar{d}_i)$ from outcomes μ_{i-1} or μ_i to outcomes μ_{i+1} or μ_{i+2} for some $\epsilon > 0$ leaving $\bar{\omega}_i$ and $\bar{\omega}_{i+1}$ unchanged. As long as

$$\underline{d}_{i-1} + \bar{d}_i - \epsilon \geq \frac{p_{i-1}+p_i-\epsilon}{2} \quad \Leftrightarrow \quad \epsilon \leq 2 \cdot \left(\underline{d}_{i-1} + \bar{d}_i - \frac{p_{i-1}+p_i}{2} \right) = 2 \cdot (\underline{d}_{i-1} + \bar{d}_i) - (p_{i-1}+p_i),$$

μ_{i+1} and μ_{i+2} continue being a 2-partition. The conditional variance of the state given μ_{i-1},

μ_i, μ_{i+1} or μ_{i+2} becomes

$$\frac{1}{\sum_{j=i-1}^{i+2} p_j} \cdot \Bigg[\frac{(p_{i-1}+p_i-\epsilon)^3}{12} + \frac{(p_{i+1}+p_{i+2}+\epsilon)^3}{12} - (p_{i-1}+p_i-\epsilon) \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \underline{d}_{i-1}\right) \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \left(\bar{d}_i - \epsilon\right)\right) - (p_{i+1}+p_{i+2}+\epsilon) \cdot \left(\frac{p_{i+1}+p_{i+2}+\epsilon}{2} - \left(\underline{d}_{i+1}+\epsilon\right)\right) \cdot \left(\frac{p_{i+1}+p_{i+2}+\epsilon}{2} - \bar{d}_{i+2}\right) \Bigg].$$

Differentiating this term with respect to ϵ yields

$$\frac{1}{\sum_{j=i-1}^{i+2} p_j} \cdot \Bigg[\frac{(p_{i-1}+p_i-\epsilon)^2}{4} - \frac{(p_{i+1}+p_{i+2}+\epsilon)^2}{4} + \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \underline{d}_{i-1}\right) \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \left(\bar{d}_i-\epsilon\right)\right) + \frac{p_{i-1}+p_i-\epsilon}{2} \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \left(\bar{d}_i-\epsilon\right)\right) - \frac{p_{i-1}+p_i-\epsilon}{2} \cdot \left(\frac{p_{i-1}+p_i-\epsilon}{2} - \underline{d}_{i-1}\right) - \left(\frac{p_{i+1}+p_{i+2}+\epsilon}{2} - \left(\underline{d}_{i+1}+\epsilon\right)\right) \cdot \left(\frac{p_{i+1}+p_{i+2}+\epsilon}{2} - \bar{d}_{i+2}\right) + \frac{p_{i+1}+p_{i+2}+\epsilon}{2} \cdot \left(\frac{p_{i+1}+p_{i+2}+\epsilon}{2} - \bar{d}_{i+2}\right) - \frac{p_{i+1}+p_{i+2}+\epsilon}{2} \cdot \left(\frac{p_{i+1}+p_{i+2}+\epsilon}{2} - \left(\underline{d}_{i+1}+\epsilon\right)\right) \Bigg].$$

At $\epsilon = 0$, this term becomes

$$\begin{aligned}
&\frac{1}{\sum_{j=i-1}^{i+2} p_j} \cdot \Bigg[\frac{(p_{i-1}+p_i)^2}{4} - \frac{(p_{i+1}+p_{i+2})^2}{4} + \left(\frac{p_{i-1}+p_i}{2} - \underline{d}_{i-1}\right) \cdot \left(\frac{p_{i-1}+p_i}{2} - \bar{d}_i\right) \\
&\quad + \frac{p_{i-1}+p_i}{2} \cdot \left(\frac{p_{i-1}+p_i}{2} - \bar{d}_i\right) - \frac{p_{i-1}+p_i}{2} \cdot \left(\frac{p_{i-1}+p_i}{2} - \underline{d}_{i-1}\right) \\
&\quad - \left(\frac{p_{i+1}+p_{i+2}}{2} - \left(\underline{d}_{i+1}\right)\right) \cdot \left(\frac{p_{i+1}+p_{i+2}}{2} - \bar{d}_{i+2}\right) \\
&\quad + \frac{p_{i+1}+p_{i+2}}{2} \cdot \left(\frac{p_{i+1}+p_{i+2}}{2} - \bar{d}_{i+2}\right) - \frac{p_{i+1}+p_{i+2}}{2} \cdot \left(\frac{p_{i+1}+p_{i+2}}{2} - \underline{d}_{i+1}\right) \Bigg] \\
&= \frac{1}{\sum_{j=i-1}^{i+2} p_j} \cdot \Bigg[-\bar{d}_i \cdot \left(\frac{p_{i-1}+p_i}{2} - \underline{d}_{i-1}\right) - \frac{p_{i-1}+p_i}{2} \cdot \bar{d}_i \\
&\quad + \underline{d}_{i+1} \cdot \left(\frac{p_{i+1}+p_{i+2}}{2} - \bar{d}_{i+2}\right) + \frac{p_{i+1}+p_{i+2}}{2} \cdot \underline{d}_{i+1} \Bigg] \\
&= \frac{1}{\sum_{j=i-1}^{i+2} p_j} \cdot \left[-\bar{d}_i \cdot \left((p_{i-1}+p_i) - \underline{d}_{i-1}\right) + \underline{d}_{i+1} \cdot \left((p_{i+1}+p_{i+2}) - \bar{d}_{i+2}\right) \right] \\
&= \frac{1}{\sum_{j=i-1}^{i+2} p_j} \cdot \left[-\bar{d}_i \cdot \left((a_i - a_{i-1}) + \bar{d}_i\right) + \underline{d}_{i+1} \cdot \left(\underline{d}_{i+1} + (a_{i+2} - a_{i+1})\right) \right] \\
&< \frac{1}{\sum_{j=i-1}^{i+2} p_j} \cdot \left[-b \cdot (2b+b) + b \cdot (b+2b) \right] = 0.
\end{aligned}$$

By continuity of this term in ϵ, there exists some non-empty open interval around zero such that the derivative is strictly negative on this interval. So for any $\epsilon > 0$ sufficiently close to zero, the conditional variance of the state given μ_{i-1}, μ_i, μ_{i+1} or μ_{i+2} is strictly smaller than under the original experiment.

As a consequence, the original information structure can only be optimal if μ_{i+1} and μ_{i+2} form a 2-partition and $\bar{\omega}_{i+2} - \bar{\omega}_{i+1} > 2b$. Besides, since $\underline{d}_{i+1} = (\bar{\omega}_{i+1} - \bar{\omega}_i) - \bar{d}_i < 2b - b = b$, this implies that $\bar{d}_{i+2} > b$. So by applying the same reasoning as above, this can only be optimal if μ_{i+3} and μ_{i+4} form a 2-partition and $\bar{\omega}_{i+4} - \bar{\omega}_{i+3} > 2b$. Iterating forward, it is thus necessary that for all $j \in \{i, i+2, \ldots, n-2, n\}$, μ_{j-1} and μ_j form a 2-partition and $\bar{\omega}_j - \bar{\omega}_{j-1} > 2b$. So in particular, $n - i$ must be even. However, since μ_n forms a 1-partition by Fact 2, this is not possible under an optimal bi-pooling policy.

The proof for the case where $\underline{d}_{i-1}+ > \bar{d}_i$ works analogously: One can show that μ_{i-2} cannot form a 1-partition, but for all $j \in \{2, 4, \ldots, i-2, i\}$, μ_j must belong to a 2-partition together with μ_{j-1} such that the $(j-1)$th incentive compatibility constraint is binding – implying that i must be even. However, this contradicts the fact that μ_1 forms a 1-partition by Fact 2.

Proof of Case 4: If both μ_i and μ_{i-1} are a 1-partition, it follows from $\bar{\omega}_i - \bar{\omega}_{i-1} > 2b$ that $p_{i-1} + p_i = 2 \cdot (\bar{\omega}_i - \bar{\omega}_{i-1}) > 4b$.

Let's now focus on the case $p_i \geq p_{i-1}$. Consequently, it holds that $p_i > 2b$.

Suppose first that $i = n$, i.e., $p_n > 2b$. Since μ_1 and μ_n form a 1-partition, respectively, it holds that $\underline{d}_1 = \frac{p_1}{2}$ and $\bar{d}_n = \frac{p_n}{2}$. Hence, one obtains that

$$\frac{p_1 + p_n}{2} = \underline{d}_1 + \bar{d}_n = 1 - \sum_{j=2}^{n} \bar{\omega}_j - \bar{\omega}_{j-1} \leq 1 - (n-1) \cdot 2b < 2b,$$

where the first, weak inequality holds true due to the fact that all incentive compatibility constraints are satisfied, and the second, strong inequality follows from the fact that $b > \frac{1}{2n}$. One thus gets that $p_1 < 2b < p_n$.

Consider now the following alternative information structure: Increase p_1 by some $\epsilon > 0$ and decrease p_n by the same ϵ such that

$$\frac{p_{n-1} + p_n - \epsilon}{2} \geq 2b \quad \Leftrightarrow \quad \epsilon \leq 4b - (p_{n-1} + p_n)\,.$$

The conditional variance of the state given μ_1 or μ_n is

$$\frac{1}{p_1 + p_n} \cdot \left[\frac{(p_1 + \epsilon)^3}{12} + \frac{(p_n - \epsilon)^3}{12} \right].$$

Its derivative with respect to ϵ is

$$\frac{(p_1 + \epsilon)^2}{4} - \frac{(p_n - \epsilon)^2}{4},$$

which is strictly negative at $\epsilon = 0$, and thus by continuity strictly negative on some non-empty, open interval around $\epsilon = 0$. Hence, the conditional variance of the state given μ_1

or μ_n is strictly smaller under the alternative experiment for any $\epsilon > 0$ sufficiently close to zero, implying that the original bi-pooling policy cannot be optimal.

Consequently, it must be that $i < n$.

Suppose now that μ_{i+1} is a 1-partition. If $p_{i+1} \geq p_i$, one can conclude that $\bar{\omega}_{i+1} - \bar{\omega}_i = \frac{p_{i+1}+p_i}{2} > \frac{2b+2b}{2} = 2b$. But then, one can find a better information structure by decreasing p_i by some $\epsilon > 0$ and increasing $\min\{p_1, p_n\} < 2b$ by the same ϵ such that all incentive compatibility constraints remain satisfied, that is, $p_{i+1} + p_i - \epsilon \geq 4b$ and $p_i + p_{i-1} - \epsilon \geq 4b$, or, equivalently,

$$\epsilon \leq \min\{4b - p_i - p_{i+1}, 4b - p_{i-1} - p_i\} = 4b - p_i - \max\{p_{i-1}, p_{i+1}\} = 4b - p_i - p_{i+1}.$$

Compared to the original information structure, the conditional variance of the state given μ_1 or μ_i or μ_n reduces by

$$\frac{1}{p_1 + p_i + p_n}\left[\frac{(\min\{p_1, p_n\})^3}{12} + \frac{p_i^3}{12} - \left(\frac{(\min\{p_1, p_n\} + \epsilon)^3}{12} + \frac{(p_i - \epsilon)^3}{12}\right)\right].$$

This term is strictly positive for any $\epsilon \in (0, \frac{p_i - \min\{p_1, p_n\}}{2})$, implying that the original experiment is not optimal.

Hence, it must be that $p_{i+1} < p_i$. But then, there is a better information structure under which p_{i+1} is increased by some $\epsilon > 0$ while p_i is reduced by the same ϵ such that all incentive compatibility constraints remain valid, i.e.,

$$p_{i-1} + p_i - \epsilon \geq 4b \quad \Leftrightarrow \quad \epsilon \leq 4b - p_{i-1} - p_i;$$

and such that the conditional variance of the state given μ_i or μ_{i+1} shrinks, that is,

$$\frac{1}{p_i + p_{i+1}}\left[\frac{p_i^3}{12} + \frac{p_{i+1}^3}{12} - \left(\frac{(p_i - \epsilon)^3}{12} + \frac{(p_{i+1} + \epsilon)^3}{12}\right)\right] > 0,$$

which is the case if $\epsilon \in (0, \frac{p_i - p_{i+1}}{2})$.

Hence, μ_{i+1} must belong to a 2-partition. Let j be the smallest integer larger than $i+1$ such that μ_j is a 1-partition. This is well-defined as μ_n is a 1-partition by Fact 2. Then, $j - (i+1)$ is even, and for all $k \in \{1, \ldots, \frac{j-(i+1)}{2}\}$, μ_{2k+i-1} and μ_{2k+i} form a 2-partition. Moreover, one obtains that $\bar{\omega}_l - \bar{\omega}_{l-1} = 2b$ for all $l \in \{i+1, \ldots, j\}$ by Cases 1-3. Consider the following alternative information structure: For some $\epsilon \in (0, b)$, shift the interval $(\sup\{\Omega_i\} - \epsilon, \sup\{\Omega_i\})$ from outcome μ_i to outcomes μ_{i+1} or μ_{i+2}, for each $k \in \{1, \ldots, \frac{j-(i+1)}{2} - 1\}$, shift $(\sup\{\Omega_{2k+i-1,2k+i}\} - \epsilon, \sup\{\Omega_{2k+i-1,2k+i}\})$ from $\Omega_{2k+i-1,2k+i}$ to $\Omega_{2k+i+1,2k+i+2}$, and shift $(\sup\{\Omega_{j-2,j-1}\} - \epsilon, \sup\{\Omega_{j-2,j-1}\})$ from $\Omega_{j-2,j-1}$ to Ω_j such that the posterior means $\bar{\omega}_i, \bar{\omega}_{i+1}, \ldots, \bar{\omega}_{j-1}, \bar{\omega}_j$ decrease by $\frac{\epsilon}{2}$, respectively.[36] The fact that both $\bar{\omega}_i$ and $\bar{\omega}_j$ decrease by $\frac{\epsilon}{2}$ yields that μ_i and μ_j remain 1-partitions. Furthermore, the construction

[36] Choosing $\epsilon < b$ ensures that intervals are shifted from one 1- or 2-partition to the next one only, that is, $(\sup\{\Omega_i\} - \epsilon, \sup\{\Omega_i\}) \subseteq \Omega_i$ and $(\sup\{\Omega_{2k+i-1,2k+i}\} - \epsilon, \sup\{\Omega_{2k+i-1,2k+i}\}) \subseteq \Omega_{2k+i-1,2k+i}$ for every $k \in \{1, \ldots, \frac{j-(i+1)}{2}\}$.

ensures that μ_{2k+i-1} and μ_{2k+i} continue forming a 2-partition because

$$\widehat{p_{2k+i-1}} + \widehat{p_{2k+i}} = p_{2k+i-1} + p_{2k+i} > 2\left(a_{2k+i} - a_{2k+i-1}\right) = 2\left(a_{2k+i} - \frac{\epsilon}{2} - \left(a_{2k+i-1} - \frac{\epsilon}{2}\right)\right)$$

for all $k \in \{1, \ldots, \frac{j-(i+1)}{2}\}$. Additionally, as long as

$$\bar{\omega}_i - \bar{\omega}_{i-1} - \frac{\epsilon}{2} \geq 2b \quad \Leftrightarrow \quad \epsilon \leq 2\left(\bar{\omega}_i - \bar{\omega}_{i-1} - 2b\right),$$

all incentive compatibility constraints remain valid. Besides, note that $\underline{d}_{2k+i-1} < b$ and $\bar{d}_{2k+i} > b$ holds for all $k \in \{1, \ldots, \frac{j-(i+1)}{2}\}$, which can be shown by induction on k: For $k = 1$, this is true because

$$\underline{d}_{i+1} = \left(\bar{\omega}_{i+1} - \bar{\omega}_i\right) - \bar{d}_i = 2b - \frac{p_i}{2} < 2b - \frac{2b}{2} = b,$$

and thus

$$\bar{d}_{i+2} = p_{i+1} + p_{i+2} - \left(\bar{\omega}_{i+2} - \bar{\omega}_{i+1}\right) - \underline{d}_{i+1} = p_{i+1} + p_{i+2} - 2b - \underline{d}_{i+1} > 4b - 2b - b = b.$$

So for $k > 1$, one obtains that

$$\underline{d}_{2k+i-1} = \left(\bar{\omega}_{2(k-1)+i+1} - \bar{\omega}_{2(k-1)+i}\right) - \bar{d}_{2(k-1)+i} = 2b - \bar{d}_{2(k-1)+i} < b,$$

as $\bar{d}_{2(k-1)+i} > b$ by the induction hypothesis $(k-1)$, and hence

$$\bar{d}_{2k+i} = p_{2k+i-1} + p_{2k+i} - \left(\bar{\omega}_{2k+i} - \bar{\omega}_{2k+i-1}\right) - \underline{d}_{2k+i-1} = p_{2k+i-1} + p_{2k+i} - 2b - \underline{d}_{2k+i-1} > 4b - 2b - b = b.$$

For any $k \in \{1, \ldots, \frac{j-(i+1)}{2}\}$, the conditional variance of the state given $\tilde{\omega} \in \Omega_{2k+i-1} \cup \Omega_{2k+i}$ becomes

$$\frac{\left(p_{2k+i-1} + p_{2k+i}\right)^2}{12} - \left(\frac{p_{2k+i-1} + p_{2k+i}}{2} - \left(\underline{d}_{2k+i-1} + \frac{\epsilon}{2}\right)\right) \cdot \left(\frac{p_{2k+i-1} + p_{2k+i}}{2} - \left(\bar{d}_{2k+i} - \frac{\epsilon}{2}\right)\right).$$

At $\epsilon = 0$, its derivative with respect to ϵ is

$$\frac{1}{2} \cdot \left(\left(\frac{p_{2k+i-1} + p_{2k+i}}{2} - \left(\bar{d}_{2k+i}\right)\right) - \left(\frac{p_{2k+i-1} + p_{2k+i}}{2} - \left(\underline{d}_{2k+i-1}\right)\right)\right) = \frac{\underline{d}_{2k+i-1} - \bar{d}_{2k+i}}{2} < \frac{b - b}{2} = 0.$$

The conditional variance of the state given μ_i or μ_j is

$$\frac{1}{p_i + p_j}\left[\frac{\left(p_i - \epsilon\right)^3}{12} + \frac{\left(p_j + \epsilon\right)^3}{12}\right],$$

and its derivative with respect to ϵ is

$$\frac{1}{p_i + p_j}\left[-\frac{p_i^2}{4} + \frac{p_j^2}{4}\right] < 0$$

at $\epsilon = 0$, as $p_i > 2b$ and $p_j = 2\underline{d}_j = 2\left(\bar{\omega}_j - \bar{\omega}_{j-1} - j \bar{-} 1\right) = 2\left(2b - j \bar{-} 1\right) < 2(2b - b) = 2b$.

Therefore, one can conclude by the same reasoning as in the above three cases that there is a better experiment for some $\epsilon > 0$ sufficiently close to zero. $\square$

Proof of Lemma 7, Lemma 8 and Proposition 3

Optimal Partial Bi-pooling Policies

Consider the following auxiliary problem: Fix some $b > 0$, some $j \in \mathbb{N}$ and some interval $(\underline{\omega}, \bar{\omega}) \subset [0, 1]$. The aim is to find an optimal *partial bi-pooling policy* of size j on the interval $(\underline{\omega}, \bar{\omega})$, that is, a policy which minimizes the conditional variance of the state given $\tilde{\omega} \in (\underline{\omega}, \bar{\omega})$. A partial bi-pooling policy is defined by a partition $\{\Omega_1, \ldots, \Omega_j\}$ of $(\underline{\omega}, \bar{\omega})$ together with a sequence of values $(\bar{\omega}_i)_{i=1}^j$ such that $\bar{\omega}_i$ is the conditional mean of the state given $\tilde{\omega} \in \Omega_i$ and Ω_i either forms a 1-partition, meaning that the outcome of a corresponding experiment on the restricted interval $(\underline{\omega}, \bar{\omega})$ forms a 1-partition, or belongs to a 2-partition for all $i \in \{1, \ldots, j\}$.

Next, I characterize the optimal partial bi-pooling policy among all those with the following properties: $\bar{\omega}_i - \bar{\omega}_{i-1} = 2b$ for all $i \in \{2, \ldots, j\}$, $\underline{d}_1 \equiv \bar{\omega}_1 - \underline{\omega} \in (0, 2b)$ and $\bar{d}_j \equiv \bar{\omega} - \bar{\omega}_j \in (0, 2b)$.

Optimal Partial Bi-pooling Policy Given $\underline{d}_1$ and $\bar{d}_j$

Toward this goal, I first characterize the optimal policy for any fixed $\underline{d}_1 \in (0, 2b)$ and $\bar{d}_j \in (0, 2b)$ and for any $j \in \mathbb{N}$. More precisely, I determine under which conditions Ω_1 forms a 1-partition or a 2-partition with Ω_2 under the optimal experiment. Then, I compute the optimal conditional variance, which turns out to depend on $\underline{d}_1$, $\bar{d}_j$ and j only, that is, it can be expressed as

$$\mathrm{Var}(\omega | \omega \in (\underline{\omega}, \bar{\omega})) = v^*(\underline{d}_1, \bar{d}_j, j)$$

for some function $v^* : (0, 2b) \times (0, 2b) \times \mathbb{N} \to \mathbb{R}$.

$j = 1$**:**

Proof. In this case, Ω_1 forms a 1-partition, and it must be that $\underline{d}_1 = \bar{d}_1$. The conditional variance of the state given $\tilde{\omega} \in (\underline{\omega}, \bar{\omega})$ is

$$\mathrm{Var}(\omega | \omega \in (\underline{\omega}, \bar{\omega})) = \frac{\left(\underline{d}_1 + \bar{d}_1\right)^2}{12} = v^*(\underline{d}_1, \bar{d}_1, 1).$$

$\square$

$j = 2$**:**

- Ω_1 forms a 1-partition iff $\underline{d}_1 = 2b - \bar{d}_j$.
- Ω_1 and Ω_2 form a 2-partition iff $\underline{d}_1 > 2b - \bar{d}_j$.

Proof. If Ω_1 forms a 1-partition, then Ω_2 is a 1-partition as well. But this is possible if and only if $p_1 + p_2 = 2(a_2 - a_1)$. Since $p_1 + p_2 = \underline{d}_1 + (a_2 - a_1) + \bar{d}_2$ and $\bar{\omega}_2 - \bar{\omega}_1 = 2b$, this is equivalent to $\underline{d}_1 + \bar{d}_2 = 2b$. On the other hand, Ω_1 and Ω_2 form a 2-partition if and only if $p_1 + p_2 > 2(\bar{\omega}_2 - \bar{\omega}_1)$, or, equivalently, $\underline{d}_1 + \bar{d}_2 > 2b$. The conditional variance of the state is

$$\begin{aligned}
\mathrm{Var}(\omega|\omega \in (\underline{\omega}, \bar{\omega})) &= \begin{cases} \frac{1}{\underline{d}_1+2b+\bar{d}_2} \cdot \left(2\underline{d}_1 \cdot v^*(\underline{d}_1, \underline{d}_1, 1) + 2\bar{d}_2 \cdot v^*(\bar{d}_2, \bar{d}_2, 1)\right) & \text{, if } \underline{d}_1 + \bar{d}_2 = 2b \\ \frac{(\underline{d}_1+2b+\bar{d}_2)^2}{12} - \left(\frac{\underline{d}_1+2b+\bar{d}_2}{2} - \underline{d}_1\right) \cdot \left(\frac{\underline{d}_1+2b+\bar{d}_2}{2} - \bar{d}_2\right) & \text{, if } \underline{d}_1 + \bar{d}_2 > 2b \end{cases} \\
&= \frac{1}{\underline{d}_1 + 2b + \bar{d}_2} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_2^3}{3} + b \cdot (\underline{d}_1^2 + \bar{d}_2^2) - \frac{4b^3}{3} \right) \\
&= v^*(\underline{d}_1, \bar{d}_2, 2)
\end{aligned}$$

if $\underline{d}_1 + \bar{d}_2 \geq 2b$. □

$j = 3$**:**

- Ω_1 forms a 1-partition iff $\underline{d}_1 \leq \bar{d}_j$.
- Ω_1 and Ω_2 form a 2-partition iff $\underline{d}_1 > \bar{d}_j$.

Proof. Suppose that Ω_1 forms a 1-partition. Then, $p_1 = 2\underline{d}_1$. Moreover, Ω_2 and Ω_3 must form either a 2-partition or two separate 1-partitions. Using the result of the case $j = 2$, this is possible if and if $p_2 + p_3 \geq 2\,(\bar{\omega}_3 - \bar{\omega}_2)$. Since $p_1 + p_2 + p_3 = \underline{d}_1 + (\bar{\omega}_3 - \bar{\omega}_1) + \bar{d}_3$ and $\bar{\omega}_2 - \bar{\omega}_1 = \bar{\omega}_3 - \bar{\omega}_2 = 2b$, this inequality holds true if and only if $\underline{d}_1 \leq \bar{d}_3$. If Ω_1 forms a 2-partition together with Ω_2, then Ω_3 is a 1-partition and $p_3 = 2\bar{d}_3$. Furthermore, this is possible if and only if $p_1 + p_2 > 2(\bar{\omega}_2 - \bar{\omega}_1)$, or, equivalently, $\underline{d}_1 > \bar{d}_3$. Then, the conditional variance of the state becomes

$$\begin{aligned}
&\mathrm{Var}(\omega|\omega \in (\underline{\omega}, \bar{\omega})) \\
&= \begin{cases} \frac{1}{\underline{d}_1+4b+\bar{d}_3} \cdot \left(2\underline{d}_1 \cdot v^*(\underline{d}_1, \underline{d}_1, 1) + \left((2b - \underline{d}_1) + 2b + \bar{d}_3\right) \cdot v^*(2b - \underline{d}_1, \bar{d}_3, 2)\right) & \text{, if } \underline{d}_1 < \bar{d}_3 \\ \frac{1}{\underline{d}_1+4b+\bar{d}_3} \cdot \left(2\underline{d}_1 \cdot v^*(\underline{d}_1, \underline{d}_1, 1) + \left((2b - \underline{d}_1) + (2b - \bar{d}_3)\right) \cdot v^*(2b - \underline{d}_1, 2b - \bar{d}_3, 1)\right. & \text{, if } \underline{d}_1 = \bar{d}_3 \\ \qquad \left. +2\bar{d}_3 \cdot v^*(\bar{d}_3, \bar{d}_3, 1)\right) & \\ \frac{1}{\underline{d}_1+4b+\bar{d}_3} \cdot \left(\left(\underline{d}_1 + 2b + (2b - \bar{d}_3)\right) \cdot v^*(\underline{d}_1, 2b - \bar{d}_3, 2) + 2\bar{d}_3 \cdot v^*(\bar{d}_3, \bar{d}_3, 1)\right) & \text{, if } \underline{d}_1 > \bar{d}_3 \end{cases} \\
&= \begin{cases} \frac{1}{\underline{d}_1+4b+\bar{d}_3} \cdot \left(\frac{\underline{d}_1^3+\bar{d}_3^3}{3} + 3b\underline{d}_1^2 + b\bar{d}_3^2 - 8b^2\underline{d}_1 + \frac{16b^3}{3}\right) & \text{, if } \underline{d}_1 \leq \bar{d}_3 \\ \frac{1}{\underline{d}_1+4b+\bar{d}_3} \cdot \left(\frac{\underline{d}_1^3+\bar{d}_3^3}{3} + b\underline{d}_1^2 + 3b\bar{d}_3^2 - 8b^2\bar{d}_3 + \frac{16b^3}{3}\right) & \text{, if } \underline{d}_1 \geq \bar{d}_3 \end{cases} \\
&= v^*(\underline{d}_1, \bar{d}_3, 3).
\end{aligned}$$

□

$j = 4$**:**

- Ω_1 forms a 1-partition iff $\underline{d}_1 \leq \max\{b, 2b - \bar{d}_j\}$.
- Ω_1 and Ω_2 form a 2-partition iff $\underline{d}_1 > \max\{b, 2b - \bar{d}_j\}$.

Proof. Suppose that Ω_1 forms a 1-partition under an optimal partial bi-pooling policy. Using the result of the case $j = 3$, this policy is uniquely determined and has the following properties: Ω_2 forms a 1-partition if and only if $2b - \underline{d}_1 \leq \bar{d}_4$, i.e., $d_1 \geq 2b - \bar{d}_4$, and Ω_2 and Ω_3 form a 2-partition if and only if $d_1 < 2b - \bar{d}_4$.

On the other hand, if Ω_1 and Ω_2 form a 2-partition under an optimal partial bi-pooling policy, then the conditional variance of the state given that $\omega \in (\underline{\omega}, \bar{\omega})$ equals the optimal value of the minimization problem

$$\begin{aligned}\min_{d} \quad & \frac{1}{\underline{d}_1 + 6b + \bar{d}_4} \cdot \left[(\underline{d}_1 + 2b + d) \cdot v^*(\underline{d}_1, d, 2) + (2b - d + 2b + \bar{d}_4) \cdot v^*(2b - d, \bar{d}_4, 2)\right] \\ \text{s.t.} \quad & \underline{d}_1 + d \geq 2b \text{ and } (2b - d) + \bar{d}_4 \geq 2b.\end{aligned}$$

In particular, the inequality constraints ensure that Ω_1 and Ω_2 as well as Ω_3 and Ω_4 either form a 2-partition or two separate 1-partitions, respectively (cf. the case $j = 2$).[37] Besides, they are equivalent to $\bar{d}_4 \geq d \geq 2b - \underline{d}_1$ so that the constraint set is non-empty if and only if $\bar{d}_4 \geq 2b - \underline{d}_1$. Differentiating the objective function with respect to d yields

$$\frac{1}{\underline{d}_1 + 6b + \bar{d}_4} \cdot \left[d^2 + 2bd - (2b - d)^2 - 2b(2b - d)\right] = \frac{8b(d - b)}{\underline{d}_1 + 6b + \bar{d}_4}.$$

Hence, the solution of the optimization problem is

$$d^* = \begin{cases} b & \text{, if } \bar{d}_4 \geq b \geq 2b - \underline{d}_1 \\ \bar{d}_4 & \text{, if } b \geq \bar{d}_4 \geq 2b - \underline{d}_1 \\ 2b - \underline{d}_1 & \text{, if } \bar{d}_4 \geq 2b - \underline{d}_1 \geq b \end{cases} = \begin{cases} b & \text{, if } \underline{d}_1 \geq b \text{ and } \bar{d}_4 \geq b \\ \bar{d}_4 & \text{, if } \underline{d}_1 \geq 2b - \bar{d}_4 \text{ and } \bar{d}_4 \leq b \\ 2b - \underline{d}_1 & \text{, if } \underline{d}_1 \leq b \text{ and } \underline{d}_1 \geq 2b - \bar{d}_4 \end{cases}$$

If $\underline{d}_1 < 2b - \bar{d}_4$, no policy under which Ω_1 and Ω_2 form a 2-partition is feasible such that Ω_1 must be a 1-partition under the optimal policy.

If $\underline{d}_1 \geq 2b - \bar{d}_4$, the only possible 1-partition also belongs to the constraint set of the minimization problem: It corresponds to $d^* = 2b - \underline{d}_1$. Hence, Ω_1 and Ω_2 form a 2-partition in the optimum if and only if $d^* \neq 2b - \underline{d}_1$, that is, if and only if

$$\begin{aligned}\underline{d}_1 > b \text{ and } \bar{d}_4 \geq b \quad \text{or} \quad \underline{d}_1 > 2b - \bar{d}_4 \text{ and } \bar{d}_4 \leq b \quad & \Leftrightarrow \quad \underline{d}_1 > b \text{ and } \underline{d}_1 > 2b - \bar{d}_4 \\ & \Leftrightarrow \quad \underline{d}_1 > \max\{b, 2b - \bar{d}_4\}.\end{aligned}$$

[37]That is, the constraint set corresponds to the set of policies with these properties.

One obtains that

$$\begin{aligned}
&\operatorname{Var}(\omega|\omega\in(\underline{\omega},\bar{\omega}))\\
&=\begin{cases}
\frac{1}{\underline{d}_1+6b+\bar{d}_4}\cdot\left(2\underline{d}_1\cdot v^*(\underline{d}_1,\underline{d}_1,1)+(2b-\underline{d}_1+4b+\bar{d}_4)\cdot v^*(2b-\underline{d}_1,\bar{d}_4,3)\right) & \text{, if } \underline{d}_1\le\max\{b,2b-\bar{d}_4\}\\
\frac{1}{\underline{d}_1+6b+\bar{d}_4}\cdot\big((\underline{d}_1+2b+\bar{d}_4)\cdot v^*(\underline{d}_1,\bar{d}_4,2) & \text{, if } \underline{d}_1\ge 2b-\bar{d}_4 \text{ and } \bar{d}_4\le b\\
\qquad +(2b-\bar{d}_4+2b+\bar{d}_4)\cdot v^*(2b-\bar{d}_4,\bar{d}_4,2)\big) & \\
\frac{1}{\underline{d}_1+6b+\bar{d}_4}\cdot\left((\underline{d}_1+2b+b)\cdot v^*(\underline{d}_1,b,2)+(b+2b+\bar{d}_4)\cdot v^*(b,\bar{d}_4,2)\right) & \text{, if } \underline{d}_1\ge b \text{ and } \bar{d}_4\ge b
\end{cases}\\
&=\begin{cases}
\frac{1}{\underline{d}_1+6b+\bar{d}_4}\cdot\left(\frac{\underline{d}_1^3+\bar{d}_4^3}{3}+3b\left(\underline{d}_1^2+\bar{d}_4^2\right)-8b^2\left(\underline{d}_1+\bar{d}_4\right)+12b^3\right) & \text{, if } \underline{d}_1\le 2b-\bar{d}_4\\
\frac{1}{\underline{d}_1+6b+\bar{d}_4}\cdot\left(\frac{\underline{d}_1^3+\bar{d}_4^3}{3}+5b\underline{d}_1^2+b\bar{d}_4^2-8b^2\underline{d}_1+4b^3\right) & \text{, if } b\ge\underline{d}_1\ge 2b-\bar{d}_4\\
\frac{1}{\underline{d}_1+6b+\bar{d}_4}\cdot\left(\frac{\underline{d}_1^3+\bar{d}_4^3}{3}+b\underline{d}_1^2+5b\bar{d}_4^2-8b^2\bar{d}_4+4b^3\right) & \text{, if } \underline{d}_1\ge 2b-\bar{d}_4 \text{ and } \bar{d}_4\le b\\
\frac{1}{\underline{d}_1+6b+\bar{d}_4}\cdot\left(\frac{\underline{d}_1^3+\bar{d}_4^3}{3}+b\left(\underline{d}_1^2+\bar{d}_4^2\right)\right) & \text{, if } \underline{d}_1\ge b \text{ and } \bar{d}_4\ge b
\end{cases}\\
&=v^*(\underline{d}_1,\bar{d}_4,4).
\end{aligned}$$

□

$j=5$:

- Ω_1 forms a 1-partition iff $\underline{d}_1\le\min\{\frac{4b}{3},\max\{b,\bar{d}_j\}\}$.
- Ω_1 and Ω_2 form a 2-partition iff $\underline{d}_1>\min\{\frac{4b}{3},\max\{b,\bar{d}_j\}\}$.

Proof. Suppose that Ω_1 forms a 1-partition under an optimal partial bi-pooling policy. Using the result of the case $j=4$, the best policy among those ones is uniquely determined and has the following properties: Ω_2 forms a 1-partition if and only if $2b-\underline{d}_1\le\max\{b,2b-\bar{d}_j\}$, or, equivalently, $\underline{d}_1\ge\min\{b,\bar{d}_5\}$. Similarly, Ω_2 and Ω_3 form a 2-partition if and only if $\underline{d}_1<\min\{b,\bar{d}_5\}$. On the other hand, if Ω_1 and Ω_2 form a 2-partition under an optimal partial bi-pooling policy, then the conditional variance of the state given that $\omega\in(\underline{\omega},\bar{\omega})$ equals the optimal value of the minimization problem

$$\begin{aligned}
\min_{d}\quad & \frac{1}{\underline{d}_1+8b+\bar{d}_5}\cdot\left((\underline{d}_1+2b+d)\cdot v^*(\underline{d}_1,d,2)+(2b-d+4b+\bar{d}_5)\cdot v^*(2b-d,\bar{d}_5,3)\right)\\
\text{s.t.}\quad & \underline{d}_1+d\ge 2b
\end{aligned}$$

As before, the inequality constraint ensures that Ω_1 and Ω_2 either form a 2-partition or two separate 1-partitions (cf. the case $j=2$). Since $v^*(2b-d,\bar{d}_5,3)$ is differentiable everywhere except if $2b-d=\bar{d}_5$, the objective function is differentiable with respect to d everywhere except at $d=2b-\bar{d}_5$, and the derivative is given by

$$\frac{1}{\underline{d}_1+8b+\bar{d}_5}\cdot\begin{cases}8b(d-b) & \text{, if } 2b-d>\bar{d}_5\\ 4b(3d-2b) & \text{, if } 2b-d<\bar{d}_5\end{cases}.$$

Hence, the solution of the optimization problem is

$$d^* = \begin{cases} b & \text{, if } 2b - b \geq \bar{d}_5 \text{ and } \underline{d}_1 + b \geq 2b \\ \frac{2b}{3} & \text{, if } 2b - \frac{2b}{3} \leq \bar{d}_5 \text{ and } \underline{d}_1 + \frac{2b}{3} \geq 2b \\ 2b - \bar{d}_5 & \text{, if } 2b - b \leq \bar{d}_5 \text{ and } 2b - \frac{2b}{3} \geq \bar{d}_5 \text{ and } \underline{d}_1 + 2b - \bar{d}_5 \geq 2b \\ 2b - \underline{d}_1 & \text{, else} \end{cases}$$

$$= \begin{cases} b & \text{, if } \underline{d}_1 \geq b \geq \bar{d}_5 \\ \frac{2b}{3} & \text{, if } \underline{d}_1 \geq \frac{4b}{3} \text{ and } \bar{d}_5 \geq \frac{4b}{3} \\ 2b - \bar{d}_5 & \text{, if } \underline{d}_1 \geq \bar{d}_5 \text{ and } \frac{4b}{3} \geq \bar{d}_5 \geq b \\ 2b - \underline{d}_1 & \text{, } \underline{d}_1 \leq \min\{\frac{4b}{3}, \max\{b, \bar{d}_5\}\} \end{cases}$$

So if $\underline{d}_1 \leq \min\{\frac{4b}{3}, \max\{b, \bar{d}_5\}\}$, any policy under which Ω_1 and Ω_2 form a 2-partition yields a higher conditional of the state than some policy under which Ω_1 and Ω_2 form two separate 1-partitions as the solution to the optimization problem is $d^* = 2b - \underline{d}_1$. Consequently, the optimal policy must be the best one among those under which Ω_1 is a 1-partition, as specified above.

Notice that $\min\{b, \bar{d}_5\} \leq \min\{\frac{4b}{3}, \max\{b, \bar{d}_5\}\}$. So on the other hand, if $\underline{d}_1 \geq \min\{\frac{4b}{3}, \max\{b, \bar{d}_5\}\}$, implying that $\underline{d}_1 \geq \min\{b, \bar{d}_5\}$, the best policy under which Ω_1 is a 1-partition also belongs to the constraint set of the minimization problem. But this policy is not optimal as $d^* \neq 2b - \underline{d}_1$. It results that the optimal policy is thus one under which Ω_1 belongs to a 2-partition. It follows that the conditional variance of the state given $\tilde{\omega} \in (\underline{\omega}, \bar{\omega})$ is

$$
\begin{aligned}
&\mathrm{Var}\left(\omega | \omega \in (\underline{\omega}, \bar{\omega})\right) \\
&= \begin{cases}
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot (2\underline{d}_1 \cdot v^*(\underline{d}_1, \underline{d}_1, 1) & , \text{if } \underline{d}_1 \leq \min\{\frac{4b}{3}, \max\{b, \bar{d}_j\}\} \\
\qquad + (2b - \underline{d}_1 + 6b + \bar{d}_5) \cdot v^*(2b - \underline{d}_1, \bar{d}_5, 4)\big) & \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \big((\underline{d}_1 + 2b + \frac{2}{3} \cdot b) \cdot v^*(\underline{d}_1, \frac{2}{3} \cdot b, 2) & , \text{if } \underline{d}_1 \geq \frac{4b}{3} \text{ and } \bar{d}_5 \geq \frac{4b}{3} \\
\qquad + (\frac{4}{3} \cdot b + 4b + \bar{d}_5) \cdot v^*(\frac{4}{3} \cdot b, \bar{d}_5, 3)\big) & \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \big((\underline{d}_1 + 2b + 2b - \bar{d}_5) \cdot v^*(\underline{d}_1, 2b - \bar{d}_5, 2) & , \text{if } \underline{d}_1 \geq \bar{d}_5 \text{ and } \frac{4b}{3} \geq \bar{d}_5 \geq b \\
\qquad + (\bar{d}_5 + 4b + \bar{d}_5) \cdot v^*(\bar{d}_5, \bar{d}_5, 3)\big) & \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot ((\underline{d}_1 + 2b + b) \cdot v^*(\underline{d}_1, b, 2) & , \text{if } \underline{d}_1 \geq b \geq \bar{d}_5 \\
\qquad + (b + 4b + \bar{d}_5) \cdot v^*(b, \bar{d}_5, 3)\big) &
\end{cases} \\
&= \begin{cases}
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_5^3}{3} + 5b\underline{d}_1^2 + 3b\bar{d}_5^2 - 8b^2(\underline{d}_1 + \bar{d}_5) + \frac{32b^3}{3}\right) & , \text{if } b \geq \underline{d}_1 \geq \bar{d}_5 \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_5^3}{3} + 3b\underline{d}_1^2 + 5b\bar{d}_5^2 - 8b^2(\underline{d}_1 + \bar{d}_5) + \frac{32b^3}{3}\right) & , \text{if } b \geq \bar{d}_5 \geq \underline{d}_1 \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_5^3}{3} + 3b\underline{d}_1^2 + b\bar{d}_5^2 - 8b^2\underline{d}_1 + \frac{20b^3}{3}\right) & , \text{if } \bar{d}_5 \geq b \geq \underline{d}_1 \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_5^3}{3} + 7b\underline{d}_1^2 + b\bar{d}_5^2 - 16b^2\underline{d}_1 + \frac{32b^3}{3}\right) & , \text{if } \bar{d}_5 > \underline{d}_1 \text{ and } \frac{4b}{3} \geq \underline{d}_1 \geq b \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_5^3}{3} + b\left(\underline{d}_1^2 + \bar{d}_5^2\right)\right) & , \text{if } \underline{d}_1 \geq \frac{4b}{3} \text{ and } \bar{d}_5 \geq \frac{4b}{3} \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_5^3}{3} + b\underline{d}_1^2 + 7b\bar{d}_5^2 - 16b^2\bar{d}_5 + \frac{32b^3}{3}\right) & , \text{if } \underline{d}_1 \geq \bar{d}_5 \text{ and } \frac{4b}{3} \geq \bar{d}_5 \geq b \\
\frac{1}{\underline{d}_1 + 8b + \bar{d}_5} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_5^3}{3} + b\underline{d}_1^2 + 3b\bar{d}_5^2 - 8b^2\bar{d}_5 + \frac{20b^3}{3}\right) & , \text{if } \underline{d}_1 \geq b \geq \bar{d}_5
\end{cases} \\
&= v^*(\underline{d}_1, \bar{d}_5, 5).
\end{aligned}
$$

under any optimal policy. □

Lemma A. 1. *If $j \geq 6$, any optimal partial bi-pooling policy on $(\underline{\omega}, \bar{\omega})$ has the following properties:*
(i) If j is even, Ω_1 forms a 1-partition if and only if $\underline{d}_1 \leq \min\left\{\frac{j-2}{j-3} \cdot b, \max\{b, 2b - \bar{d}_j\}\right\}$, while Ω_1 and Ω_2 are a 2-partition if and only if $\underline{d}_1 > \min\left\{\frac{j-2}{j-3} \cdot b, \max\{b, 2b - \bar{d}_j\}\right\}$. The

optimal conditional variance of the state given $\tilde{\omega} \in (\underline{\omega}, \bar{\omega})$ satisfies

$$
\begin{aligned}
&\left(\underline{d}_1 + (j-1)\cdot 2b + \bar{d}_j\right) \cdot Var(\omega | \omega \in (\underline{\omega}, \bar{\omega})) \\
&= \begin{cases}
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + 3b\left(\underline{d}_1^2 + \bar{d}_j^2\right) - 8b^2\left(\underline{d}_1 + \bar{d}_j\right) + \frac{2(j+14)b^3}{3} & , \textit{if } \underline{d}_1 \leq b \textit{ and } \bar{d}_j \leq b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + 3b\underline{d}_1^2 + (2j-5)b\bar{d}_j^2 - 8b^2\underline{d}_1 - 4(j-2)b^2\bar{d}_j + \frac{4(2j+1)b^3}{3} & , \textit{if } \underline{d}_1 + \bar{d}_j \leq 2b \textit{ and } \frac{j-2}{j-3}\cdot b \geq \bar{d}_j \geq b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + 3b\underline{d}_1^2 + b\bar{d}_j^2 - 8b^2\underline{d}_1 + \left(\frac{2(j+5)}{3} - \frac{2}{j-3}\right)b^3 & , \textit{if } \underline{d}_1 \leq \frac{j-4}{j-3}\cdot b \textit{ and } \bar{d}_j \geq \frac{j-2}{j-3}\cdot b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + (2j-3)b\underline{d}_1^2 + b\bar{d}_j^2 - 4(j-2)b^2\underline{d}_1 + \frac{4(2j-5)b^3}{3} & , \textit{if } \underline{d}_1 + \bar{d}_j \geq 2b \textit{ and } b \geq \underline{d}_1 \geq \frac{j-4}{j-3}\cdot b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + (2j-5)b\underline{d}_1^2 + 3b\bar{d}_j^2 - 4(j-2)b^2\underline{d}_1 - 8b^2\bar{d}_j + \frac{4(2j+1)b^3}{3} & , \textit{if } \underline{d}_1 + \bar{d}_j \leq 2b \textit{ and } \frac{j-2}{j-3}\cdot b \geq \underline{d}_1 \geq b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + b\underline{d}_1^2 + 3b\bar{d}_j^2 - 8b^2\bar{d}_j + \left(\frac{2(j+5)}{3} - \frac{2}{j-3}\right)\cdot b^3 & , \textit{if } \underline{d}_1 \geq \frac{j-2}{j-3}\cdot b \textit{ and } \bar{d}_j \leq \frac{j-4}{j-3}\cdot b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + b\underline{d}_1^2 + (2j-3)b\bar{d}_j^2 - 4(j-2)b^2\bar{d}_j + \frac{4(2j-5)b^3}{3} & , \textit{if } \underline{d}_1 + \bar{d}_j \geq 2b \textit{ and } b \geq \bar{d}_j \geq \frac{j-4}{j-3}\cdot b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + b\cdot\left(\underline{d}_1^2 + \bar{d}_j^2\right) + \frac{2(j-4)b^3}{3} & , \textit{if } \underline{d}_1 \geq b \textit{ and } \bar{d}_j \geq b
\end{cases} \\
&= \left(\underline{d}_1 + (j-1)\cdot 2b + \bar{d}_j\right) \cdot v^*(\underline{d}_1, \bar{d}_j, j).
\end{aligned}
$$

(ii) If j is odd, Ω_1 forms a 1-partition if and only if $\underline{d}_1 \leq \min\left\{\frac{j-1}{j-2}\cdot b, \max\{b, \bar{d}_j\}\right\}$, while Ω_1 and Ω_2 are a 2-partition if and only if $\underline{d}_1 > \min\left\{\frac{j-1}{j-2}\cdot b, \max\{b, \bar{d}_j\}\right\}$. The optimal conditional variance of the state given $\omega \in (\underline{\omega}, \bar{\omega})$ fulfills

$$
\begin{aligned}
&\left(\underline{d}_1 + (j-1)\cdot 2b + \bar{d}_j\right) \cdot Var(\omega | \omega \in (\underline{\omega}, \bar{\omega})) \\
&= \begin{cases}
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + 3b\left(\underline{d}_1^2 + \bar{d}_j^2\right) - 8b^2\left(\underline{d}_1 + \bar{d}_j\right) + \left(\frac{2(j+14)}{3} - \frac{2}{j-4}\right)b^3 & , \textit{if } \underline{d}_1 \leq \frac{j-5}{j-4}\cdot b \textit{ and } \bar{d}_j \leq \frac{j-5}{j-4}\cdot b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + 3b\underline{d}_1^2 + (2j-5)b\bar{d}_j^2 - 8b^2\underline{d}_1 - 4(j-3)b^2\bar{d}_j + \frac{8(j-1)b^3}{3} & , \textit{if } \underline{d}_1 \leq \bar{d}_j \textit{ and } b \geq \bar{d}_j \geq \frac{j-5}{j-4}\cdot b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + 3b\underline{d}_1^2 + b\bar{d}_j^2 - 8b^2\underline{d}_1 + \frac{2(j+5)b^3}{3} & , \textit{if } \underline{d}_1 \leq b \leq \bar{d}_j \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + (2j-3)b\underline{d}_1^2 + b\bar{d}_j^2 - 4(j-1)b^2\underline{d}_1 + \frac{8(j-1)b^3}{3} & , \textit{if } \underline{d}_1 \geq \bar{d}_j \textit{ and } b \geq \underline{d}_1 \geq \frac{j-5}{j-4}\cdot b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + (2j-3)b\underline{d}_1^2 + 5b\bar{d}_j^2 - 4(j-1)b^2\underline{d}_1 + \frac{8(j-1)b^3}{3} & , \textit{if } \underline{d}_1 \leq \bar{d}_j \textit{ and } \frac{j-1}{j-2}\cdot b \geq \underline{d}_1 \geq b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + b\underline{d}_1^2 + 3b\bar{d}_j^2 - 8b^2\bar{d}_j + \frac{2(j+5)b^3}{3} & , \textit{if } \underline{d}_1 \leq b \leq \bar{d}_j \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + b\underline{d}_1^2 + (2j-3)b\bar{d}_j^2 - 4(j-1)b^2\bar{d}_j + \frac{8(j-1)b^3}{3} & , \textit{if } \underline{d}_1 \geq \bar{d}_j \textit{ and } \frac{j-1}{j-2}\cdot b \geq \bar{d}_j \geq b \\
\frac{\underline{d}_1^3+\bar{d}_j^3}{3} + b\cdot\left(\underline{d}_1^2 + \bar{d}_j^2\right) + \frac{2(j-1)(j-5)b^3}{3(j-2)} & , \textit{if } \underline{d}_1 \geq \frac{j-3}{j-2}\cdot b \textit{ and } \bar{d}_j \geq \frac{j-3}{j-2}\cdot b
\end{cases} \\
&= \left(\underline{d}_1 + (j-1)\cdot 2b + \bar{d}_j\right) \cdot v^*(\underline{d}_1, \bar{d}_j, j).
\end{aligned}
$$

Proof. This result is proven by induction on j with two base cases ($j = 6, 7$) and two induction steps, one for j odd and the other for j even:

Base case 1: $j = 6$:

First, notice that among all experiments where Ω_1 forms a 1-partition, the best one has the following characteristics (see the results of the case $j = 5$): Ω_2 forms a 1-partition if and only if it holds that $2b - \underline{d}_1 \leq \min\{\frac{4b}{3}, \max\{b, \bar{d}_j\}\}$, i.e.,

$$\underline{d}_1 \geq \max\left\{\frac{2}{3} \cdot b, \min\left\{b, 2b - \bar{d}_j\right\}\right\} = \min\left\{b, \max\left\{\frac{2}{3} \cdot b, 2b - \bar{d}_j\right\}\right\},$$

while Ω_2 and Ω_3 are a 2-partition if and only if

$$\underline{d}_1 < \min\left\{b, \max\left\{\frac{2}{3} \cdot b, 2b - \bar{d}_j\right\}\right\}.$$

If Ω_1 belongs to a 2-partition in the optimum, the conditional variance of the state given $\omega \in (\underline{\omega}, \bar{\omega})$ equals the optimal value of

$$\begin{aligned} \min_{d} \quad & \frac{1}{\underline{d}_1 + (j-1) \cdot 2b + \bar{d}_j} \cdot \big((\underline{d}_1 + 2b + d) \cdot v^*(\underline{d}_1, d, 2) + (2b - d + (j-3) \cdot 2b + \bar{d}_j) \cdot v^*(2b - d, \bar{d}_j, j-2)\big) \\ \text{s.t.} \quad & \underline{d}_1 + d \geq 2b. \end{aligned}$$

Differentiation of the objective function – whenever this is possible – yields

$$\frac{1}{\underline{d}_1 + (j-1) \cdot 2b + \bar{d}_j} \cdot \begin{cases} 4b(3d - 2b) & \text{, if } 2b - d + \bar{d}_j < 2b \\ 16b(d - b) & \text{, if } 2b - d < b \text{ and } 2b - d + \bar{d}_j > 2b \\ 8b(d - b) & \text{, if } 2b - d > \min\{b, 2b - \bar{d}_j\} \end{cases}.$$

The solution of the optimization problem is thus

$$\begin{aligned} d^* &= \begin{cases} \frac{2b}{3} & \text{, if } 2b - \frac{2b}{3} + \bar{d}_j \leq 2b \text{ and } \underline{d}_1 + \frac{2b}{3} \geq 2b \\ b & \text{, if } 2b - b + \bar{d}_j \geq 2b \text{ and } \underline{d}_1 + b \geq 2b \\ \bar{d}_j & \text{, if } 2b - \frac{2b}{3} + \bar{d}_j \geq 2b \text{ and } 2b - b + \bar{d}_j \leq 2b \text{ and } \underline{d}_1 + \bar{d}_j \geq 2b \\ 2b - \underline{d}_1 & \text{, else} \end{cases} \\ &= \begin{cases} \frac{2b}{3} & \text{, if } \underline{d}_1 \geq \frac{4b}{3} \text{ and } \bar{d}_j \leq \frac{2b}{3} \\ b & \text{, if } \underline{d}_1 \geq b \text{ and } \bar{d}_j \geq b \\ \bar{d}_j & \text{, if } \underline{d}_1 \geq 2b - \bar{d}_j \text{ and } b \geq \bar{d}_j \geq \frac{2b}{3} \\ 2b - \underline{d}_1 & \text{, if } \underline{d}_1 \leq \min\{\frac{4b}{3}, \max\{b, 2b - \bar{d}_j\}\} \end{cases}. \end{aligned}$$

If $\underline{d}_1 \leq \min\{\frac{4b}{3}, \max\{b, 2b - \bar{d}_j\}\}$, any policy under which Ω_1 and Ω_2 form a 2-partition yields a higher conditional variance of the state than some policy under which Ω_1 and Ω_2 form two separate 1-partitions as the solution to the minimization problem is $d^* = 2b - \underline{d}_1$. Consequently, the best policy must be the best one among those where Ω_1 is a 1-partition, as specified above. Note that $\min\{\frac{4b}{3}, \max\{b, 2b - \bar{d}_j\}\} \geq \min\{b, \max\{\frac{2b}{3}, 2b - \bar{d}_j\}\}$. So if $\underline{d}_1 > \min\{\frac{4b}{3}, \max\{b, 2b - \bar{d}_j\}\}$, it follows that $\underline{d}_1 > \min\{b, \max\{\frac{2b}{3}, 2b - \bar{d}_j\}\}$. Therefore, the best policy under which Ω_1 forms a 1-partition also belongs to the constraint set of the minimization problem. However, this policy is not optimal as $d^* \neq 2b - \underline{d}_1$. Therefore, the optimal policy is the one under which Ω_1 and Ω_2 form a 2-partition, as specified by the minimization problem. The optimal conditional variance of the state given $\tilde{\omega} \in (\underline{\omega}, \bar{\omega})$ is thus given by

$$\begin{aligned}
&\left(\underline{d}_1 + (j-1)\cdot 2b + \bar{d}_j\right)\cdot \mathrm{Var}\left(\tilde{\omega}|\tilde{\omega}\in(\underline{\omega},\bar{\omega})\right)\\
&= \begin{cases}
2\underline{d}_1\cdot v^*(\underline{d}_1,\underline{d}_1,1) & ,\text{if } \underline{d}_1 \le \min\{\frac{4b}{3},\max\{b,\bar{d}_j\}\}\\
\quad +(2b-\underline{d}_1+(j-2)2b+\bar{d}_j)\cdot v^*(2b-\underline{d}_1,\bar{d}_j,j-1) & \\
(\underline{d}_1+2b+\frac{2}{3}\cdot b)\cdot v^*(\underline{d}_1,\frac{2}{3}\cdot b,2) & ,\text{if } \underline{d}_1 > \frac{4}{3}\cdot b \text{ and } \bar{d}_j > \frac{4}{3}\cdot b\\
\quad +(\frac{4}{3}\cdot b+(j-3)2b+\bar{d}_j)\cdot v^*(\frac{4}{3}\cdot b,\bar{d}_j,j-2) & \\
(\underline{d}_1+2b+\bar{d}_j)\cdot v^*(\underline{d}_1,\bar{d}_j,2) & ,\text{if } \underline{d}_1 \ge 2b-\bar{d}_j \text{ and } b\ge \bar{d}_j \ge \frac{2b}{3}\\
\quad +(2b-\bar{d}_j+(j-3)2b+\bar{d}_j)\cdot v^*(2b-\bar{d}_j,\bar{d}_j,j-2) & \\
(\underline{d}_1+2b+b)\cdot v^*(\underline{d}_1,b,2) & ,\text{if } \underline{d}_1 \ge b \text{ and } \bar{d}_j \ge b\\
(b+(j-3)2b+\bar{d}_j)\cdot v^*(b,\bar{d}_j,j-2) &
\end{cases}
\end{aligned}$$

By plugging in the terms of the functions $v^*(\cdot,\cdot,j-1)$ and $v^*(\cdot,\cdot,j-2)$, one obtains the expressions listed in the lemma.

Base case 2: $j = 7$:

First, note that among all experiments where Ω_1 forms a 1-partition, the best one has the following characteristics (cf. the results of case $j = 6$): Ω_2 is a 1-partition if and only if it holds that $2b - \underline{d}_1 \le \min\{\frac{4b}{3}, \max\{b, 2b - \bar{d}_j\}\}$, i.e.,

$$\underline{d}_1 \ge \max\left\{\frac{2}{3}\cdot b, \min\left\{b, \bar{d}_j\right\}\right\} = \min\left\{b, \max\left\{\frac{2}{3}\cdot b, \bar{d}_j\right\}\right\},$$

while Ω_2 and Ω_3 are a 2-partition if and only if

$$\underline{d}_1 < \min\left\{b, \max\{\frac{2}{3}\cdot b, \bar{d}_j\}\right\}.$$

If Ω_1 belongs to a 2-partition in the optimum, the conditional variance of the state given $\omega \in (\underline{\omega}, \bar{\omega})$ equals the optimal value of

$$\begin{aligned}
\min_{d}\quad & \frac{1}{\underline{d}_1+(j-1)\cdot 2b+\bar{d}_j}\cdot\left((\underline{d}_1+2b+d)\cdot v^*(\underline{d}_1,d,2)+(2b-d+(j-3)\cdot 2b+\bar{d}_j)\cdot v^*(2b-d,\bar{d}_j,j-2)\right)\\
\text{s.t.}\quad & \underline{d}_1+d\ge 2b.
\end{aligned}$$

Differentiation of the objective function – whenever this is possible – yields

$$\frac{1}{\underline{d}_1+(j-1)\cdot 2b+\bar{d}_j}\cdot\begin{cases}
16b(d-b) & ,\text{if } b > 2b-d > \bar{d}_j\\
4b(3d-2b) & ,\text{if } 2b-d < \min\{b,\bar{d}_j\}\\
4b(5d-4b) & ,\text{if } 2b-d < \bar{d}_j \text{ and } \frac{4b}{3} > 2b-d > b\\
8b(d-b) & ,\text{if } 2b-d > \min\{\frac{3b}{3},\max\{b,\bar{d}_j\}\}
\end{cases}.$$

The solution of the optimization problem is thus

$$d^* = \begin{cases} \frac{4b}{5} & \text{, if } 2b - \frac{4b}{5} \le \bar{d}_j \text{ and } \frac{4b}{3} \ge 2b - \frac{4b}{5} \ge b \text{ and } \underline{d}_1 + \frac{4b}{5} \ge 2b \\ b & \text{, if } \left(2b - b \ge \bar{d}_j \text{ or } \frac{4b}{3} \le 2b - b \text{ or } 2b - b \le b\right) \\ & \quad \text{ and } 2b - b \ge \min\{b, \bar{d}_j\} \text{ and } \underline{d}_1 + b \ge 2b \\ 2b - \bar{d}_j & \text{, if } \left(2b - \frac{4b}{5} \ge \bar{d}_j \text{ or } \frac{4b}{3} \le 2b - \frac{4b}{5} \text{ or } 2b - \frac{4b}{5} \le b\right) \\ & \quad \text{ and } 2b - b \le \bar{d}_j \text{ and } \frac{4b}{3} \ge 2b - b \ge b \text{ and } \underline{d}_1 + 2b - \bar{d}_j \ge 2b \\ 2b - \underline{d}_1 & \text{, else} \end{cases}$$

$$= \begin{cases} \frac{4b}{5} & \text{, if } \underline{d}_1 \ge \frac{6b}{5} \text{ and } \bar{d}_j \ge \frac{6b}{5} \\ b & \text{, if } \underline{d}_1 \ge b \text{ and } \bar{d}_j \le b \\ 2b - \bar{d}_j & \text{, if } \underline{d}_1 \ge \bar{d}_j \text{ and } \frac{6b}{5} \ge \bar{d}_j \ge b \\ 2b - \underline{d}_1 & \text{, if } \underline{d}_1 \le \min\{\frac{6b}{5}, \max\{b, \bar{d}_j\}\} \end{cases}.$$

Notice that $\frac{2b}{3}$ never solves the optimization problem because the derivative of the objective function is $4b(3d-2b)$ only if $d > b$, making it strictly positive whenever $2b-d < \min\{b, \bar{d}_j\}$.

If $\underline{d}_1 \le \min\{\frac{6b}{5}, \max\{b, \bar{d}_j\}\}$, any policy under which Ω_1 and Ω_2 form a 2-partition yields a higher conditional variance of the state than some policy under which Ω_1 and Ω_2 form two separate 1-partitions as the solution to the minimization problem is $d^* = 2b - \underline{d}_1$. Consequently, the best policy must be the best one among those where Ω_1 is a 1-partition, as specified above.

Note that $\min\{\frac{6b}{5}, \max\{b, \bar{d}_j\}\} \ge \min\{b, \max\{\frac{2b}{3}, \bar{d}_j\}\}$. So if $\underline{d}_1 > \min\{\frac{6b}{5}, \max\{b, \bar{d}_j\}\}$, it follows that $\underline{d}_1 > \min\{b, \max\{\frac{2b}{3}, \bar{d}_j\}\}$. Therefore, the best policy under which Ω_1 forms a 1-partition also belongs to the constraint set of the minimization problem. However, this policy is not optimal as $d^* \neq 2b - \underline{d}_1$. Therefore, the optimal policy is the one under which Ω_1 and Ω_2 form a 2-partition, as specified by the minimization problem.

The optimal conditional variance of the state given $\omega \in (\underline{\omega}, \bar{\omega})$ is thus given by

$$\left(\underline{d}_1 + (j-1)\cdot 2b + \bar{d}_j\right) \cdot \text{Var}\left(\tilde{\omega} | \tilde{\omega} \in (\underline{\omega}, \bar{\omega})\right)$$

$$= \begin{cases} 2\underline{d}_1 \cdot v^*(\underline{d}_1, \underline{d}_1, 1) + (2b - \underline{d}_1 & \text{, if } \underline{d}_1 \le \min\{\frac{6b}{5}, \max\{b, \bar{d}_j\}\} \\ \quad +(j-2)2b + \bar{d}_j) \cdot v^*(2b - \underline{d}_1, \bar{d}_j, j-1) & \\ (\underline{d}_1 + 2b + \frac{4}{5}\cdot b) \cdot v^*(\underline{d}_1, \frac{4}{5}\cdot b, 2) & \text{, if } \underline{d}_1 \ge \frac{6b}{5} \text{ and } \bar{d}_j \ge \frac{6b}{5} \\ \quad +(\frac{6b}{5} + (j-3)2b + \bar{d}_j) \cdot v^*(\frac{6b}{5}, \bar{d}_j, j-2) & \\ (\underline{d}_1 + 2b + 2b - \bar{d}_j) \cdot v^*(\underline{d}_1, 2b - \bar{d}_j, 2) & \text{, if } \underline{d}_1 \ge \bar{d}_j \text{ and } \frac{6b}{5} \ge \bar{d}_j \ge b \\ \quad +(\bar{d}_j + (j-3)2b + \bar{d}_j) \cdot v^*(\bar{d}_j, \bar{d}_j, j-2) & \\ (\underline{d}_1 + 2b + b) \cdot v^*(\underline{d}_1, b, 2) + (b + (j-3)2b + \bar{d}_j) \cdot v^*(b, \bar{d}_j, j-2) & \text{, if } \underline{d}_1 \ge b \text{ and } \bar{d}_j \le b \end{cases}$$

By plugging in the terms of the functions $v^*(\cdot,\cdot,j-1)$ and $v^*(\cdot,\cdot,j-2)$, one obtains the expressions listed in the lemma.

Induction step 1: $j-1 \to j$, j even

Among all experiments where Ω_1 forms a 1-partition, the best one has the following characteristics (cf. the results of case j): Ω_2 is a 1-partition if and only if $2b - \underline{d}_1 \leq \min\{\frac{(j-2)b}{j-3}, \max\{b, \bar{d}_j\}\}$, i.e.,

$$\underline{d}_1 \geq \min\left\{b, \max\left\{\frac{j-4}{j-3} \cdot b, 2b - \bar{d}_j\right\}\right\},$$

while Ω_2 and Ω_3 are a 2-partition if and only of

$$\underline{d}_1 < \min\left\{b, \max\left\{\frac{j-4}{j-3} \cdot b, 2b - \bar{d}_j\right\}\right\}.$$

If Ω_1 belongs to a 2-partition in the optimum, the conditional variance of the state given $\omega \in (\underline{\omega}, \bar{\omega})$ equals the optimal value of

$$\begin{aligned} \min_d \quad & \frac{1}{\underline{d}_1 + (j-1) \cdot 2b + \bar{d}_j} \cdot \big((\underline{d}_1 + 2b + d) \cdot v^*(\underline{d}_1, d, 2) + (2b - d + (j-3) \cdot 2b + \bar{d}_j) \cdot v^*(2b - d, \bar{d}_j, j-2)\big) \\ \text{s.t.} \quad & \underline{d}_1 + d \geq 2b. \end{aligned}$$

Differentiation of the objective function – whenever this is possible – yields

$$\frac{1}{\underline{d}_1 + (j-1) \cdot 2b + \bar{d}_j} \cdot \begin{cases} 4b(3d - 2b) & \text{, if } 2b - d < \min\{b, \max\{\frac{j-6}{j-5} \cdot b, 2b - \bar{d}_j\}\} \\ 4b(j-2)(d-b) & \text{, if } 2b - d + \bar{d}_j > 2b \text{ and } b > 2b - d > \frac{j-4}{j-3} \cdot b \\ 4b((j-3)d(j-4)b) & \text{, if } 2b - d + \bar{d}_j < 2b \text{ and } \frac{j-4}{j-5} \cdot b > 2b - d > b \\ 8b(d-b) & \text{, if } 2b - d > \min\{\frac{j-4}{j-5} \cdot b, \max\{b, 2b - \bar{d}_j\}\} \end{cases}.$$

The solution of the optimization problem is thus

$$\begin{aligned} d^* &= \begin{cases} \frac{(j-4)b}{j-3} & \text{, if } 2b - \frac{(j-4)b}{j-3} + \bar{d}_j \leq 2b \text{ and } \frac{j-4}{j-5} \cdot b \geq 2b - \frac{(j-4)b}{j-3} \geq b \text{ and } \underline{d}_1 + \frac{(j-4)b}{j-3} \geq 2b \\ b & \text{, if } \left(2b - b + \bar{d}_j \geq 2b \text{ or } \frac{j-4}{j-5} \cdot b \leq 2b - b \text{ or } 2b - b \leq b\right) \\ & \quad \text{and } 2b - b \geq \min\{b, \max\{\frac{j-6}{j-5} \cdot b, 2b - \bar{d}_j\}\} \text{ and } \underline{d}_1 + b \geq 2b \\ \bar{d}_j & \text{, if } \left(2b - \frac{(j-4)b}{j-3} + \bar{d}_j \geq 2b \text{ or } \frac{j-4}{j-5} \cdot b \leq 2b - \frac{(j-4)b}{j-3} \text{ or } 2b - \frac{(j-4)b}{j-3} \leq b\right) \\ & \quad \text{and } 2b - b + \bar{d}_j \leq 2b \text{ and } \frac{j-4}{j-5} \cdot b \geq 2b - b \geq b \text{ and } \underline{d}_1 + \bar{d}_j \geq 2b \\ 2b - \underline{d}_1 & \text{, else} \end{cases} \\ &= \begin{cases} \frac{(j-4)b}{j-3} & \text{, if } \underline{d}_1 \geq \frac{(j-2)b}{j-3} \text{ and } \bar{d}_j \leq \frac{(j-4)b}{j-3} \\ b & \text{, if } \underline{d}_1 \geq b \text{ and } \bar{d}_j \geq b \\ \bar{d}_j & \text{, if } \underline{d}_1 \geq 2b - \bar{d}_j \text{ and } b \geq \bar{d}_j \geq \frac{(j-4)b}{j-3} \\ 2b - \underline{d}_1 & \text{, if } \underline{d}_1 \leq \min\{\frac{(j-2)b}{j-3}, \max\{b, 2b - \bar{d}_j\}\} \end{cases}. \end{aligned}$$

In particular, note that $\frac{2b}{3}$ is never the solution because the derivative of the objective function is $4b(3d - 2b)$ only if $d > b$, making it strictly positive whenever $2b - d < \min\{b, \max\{\frac{j-6}{j-5} \cdot b, 2b - \bar{d}_j\}\}$.

If $\underline{d}_1 \leq \min\{\frac{(j-2)b}{j-3}, \max\{b, 2b - \bar{d}_j\}\}$, any policy under which Ω_1 and Ω_2 form a 2-partition yields a higher conditional variance of the state than some policy under which Ω_1 and Ω_2

form two separate 1-partitions as the solution to the minimization problem is $d^* = 2b - \underline{d}_1$. Consequently, the best policy must be the best one among those where Ω_1 is a 1-partition, as specified above.

Notice that $\min\{\frac{(j-2)b}{j-3}, \max\{b, 2b-\bar{d}_j\}\} \geq \min\{b, \max\{\frac{(j-4)b}{j-3}, 2b-\bar{d}_j\}\}$ holds true. Hence, if $\underline{d}_1 > \min\{\frac{(j-2)b}{j-3}, \max\{b, 2b-\bar{d}_j\}\}$, it follows that $\underline{d}_1 > \min\{b, \max\{\frac{(j-4)b}{j-3}, 2b-\bar{d}_j\}\}$. Therefore, the best policy under which Ω_1 forms a 1-partition also belongs to the constraint set of the minimization problem. However, this policy is not optimal as $d^* \neq 2b - \underline{d}_1$. Therefore, the optimal policy is the one under which Ω_1 and Ω_2 form a 2-partition, as specified by the minimization problem.

The optimal conditional variance of the state given $\omega \in (\underline{\omega}, \bar{\omega})$ is thus given by

$$\begin{aligned}
&\left(\underline{d}_1 + (j-1)\cdot 2b + \bar{d}_j\right) \cdot \mathrm{Var}\left(\omega | \omega \in (\underline{\omega}, \bar{\omega})\right) \\
&= \begin{cases}
2\underline{d}_1 \cdot v^*(\underline{d}_1, \underline{d}_1, 1) + & \text{, if } \underline{d}_1 \leq \min\{\frac{(j-2)b}{j-3}, \max\{b, \bar{d}_j\}\} \\
\quad +(2b - \underline{d}_1 + (j-2)2b + \bar{d}_j) \cdot v^*(2b - \underline{d}_1, \bar{d}_j, j-1) & \\
(\underline{d}_1 + 2b + \frac{(j-4)b}{j-3}) \cdot v^*(\underline{d}_1, \frac{(j-4)b}{j-3}, 2) & \text{, if } \underline{d}_1 > \frac{(j-2)b}{j-3} \text{ and } \bar{d}_j > \frac{(j-2)b}{j-3} \\
\quad +(\frac{(j-2)b}{j-3} + (j-3)2b + \bar{d}_j) \cdot v^*(\frac{(j-2)b}{j-3}, \bar{d}_j, j-2) & \\
(\underline{d}_1 + 2b + \bar{d}_j) \cdot v^*(\underline{d}_1, \bar{d}_j, 2) & \text{, if } \underline{d}_1 \geq 2b - \bar{d}_j \text{ and } b \geq \bar{d}_j \geq \frac{(j-4)b}{j-3} \\
\quad +(2b - \bar{d}_j + (j-3)2b + \bar{d}_j) \cdot v^*(2b - \bar{d}_j, \bar{d}_j, j-2) & \\
(\underline{d}_1 + 2b + b) \cdot v^*(\underline{d}_1, b, 2) & \text{, if } \underline{d}_1 \geq b \text{ and } \bar{d}_j \geq b \\
\quad +(b + (j-3)2b + \bar{d}_j) \cdot v^*(b, \bar{d}_j, j-2) &
\end{cases}
\end{aligned}$$

By plugging in the terms of the functions $v^*(\cdot, \cdot, j-1)$ and $v^*(\cdot, \cdot, j-2)$, one obtains the expressions listed in the lemma. □

Induction step 2: $j - 1 \to j$, j odd

The proof of this step works analogously to the proof of the second base case and is therefore omitted.

Optimal Partial Bi-pooling Policy Given $\underline{d}_1 + \bar{d}_j$

As a next step, I determine the optimal policy for any fixed $\underline{d}_1 + \bar{d}_j \in (0, 2b]$ and any j by comparing the optimal policies for fixed $\underline{d}_1$ and fixed $\bar{d}_j$ from the previous section. It turns out that the optimal policy is symmetric in the sense that $\underline{d}_1 = \bar{d}_j$:

Lemma A.2. *For any $j \in \mathbb{N}$, given $\underline{d}_1 + \bar{d}_j \in (0, 2b]$, the optimal bi-pooling policy is the one with $\underline{d}_1 = \bar{d}_j$, that is,*

$$v^*(\underline{d}_1, \bar{d}_j, j) \geq v^*\left(\frac{\underline{d}_1 + \bar{d}_j}{2}, \frac{\underline{d}_1 + \bar{d}_j}{2}, j\right) \tag{13}$$

for any $\underline{d}_1, \bar{d}_j \in (0, 2b)$ such that $\underline{d}_1 + \bar{d}_j \in (0, 2b]$.

Proof.

$j=1$**:** The only feasible policy is the one with $\underline{d}_1 = \bar{d}_1$, which is thus trivially optimal.

$j=2$**:** Since both the quadratic function x^2 and the cubic function x^3 are convex on $\mathbb{R}_+$, it follows by Jensen's inequality that

$$\begin{aligned} v^*(\underline{d}_1, \bar{d}_2, 2) &= \frac{1}{\underline{d}_1 + 2b + \bar{d}_2} \cdot \left(\frac{\underline{d}_1^3 + \bar{d}_2^3}{3} + b \cdot (\underline{d}_1^2 + \bar{d}_2^2) - \frac{4b^3}{3} \right) \\ &\geq \frac{1}{\frac{\underline{d}_1 + \bar{d}_2}{2} + 2b + \frac{\underline{d}_1 + \bar{d}_2}{2}} \cdot \left(\frac{\left(\frac{\underline{d}_1+\bar{d}_2}{2}\right)^3 + \left(\frac{\underline{d}_1+\bar{d}_2}{2}\right)^3}{3} + b \cdot (\left(\frac{\underline{d}_1 + \bar{d}_2}{2}\right)^2 + \left(\frac{\underline{d}_1 + \bar{d}_2}{2}\right)^2) - \frac{4b^3}{3} \right) \\ &= v^* \left(\frac{\underline{d}_1 + \bar{d}_2}{2}, \frac{\underline{d}_1 + \bar{d}_2}{2}, 2 \right) \end{aligned}$$

$j=3$**:** First, I show that the optimal policy with $\underline{d}_1 \leq \bar{d}_3$ is the one with $\underline{d}_1 = \bar{d}_3$: For this, notice that the conditional variance can be written as

$$v^*(\underline{d}_1, \bar{d}_3, 3) = v^* \left(\frac{\underline{d}_1 + \bar{d}_3}{2} - \epsilon, \frac{\underline{d}_1 + \bar{d}_3}{2} + \epsilon, 3 \right)$$

for some $\epsilon \in [0, b]$ for any such $\underline{d}_1, \bar{d}_3$. Its derivative with respect to ϵ is

$$\begin{aligned} &\frac{1}{\underline{d}_1 + 4b + \bar{d}_3} \left(-\left(\frac{\underline{d}_1 + \bar{d}_3}{2} - \epsilon \right)^2 + \left(\frac{\underline{d}_1 + \bar{d}_3}{2} + \epsilon \right)^2 - 6b \left(\frac{\underline{d}_1 + \bar{d}_3}{2} - \epsilon \right) + 2b \left(\frac{\underline{d}_1 + \bar{d}_3}{2} + \epsilon \right) + 8b^2 \right) \\ &= \frac{1}{\underline{d}_1 + 4b + \bar{d}_3} \left(2\epsilon \left(\underline{d}_1 + \bar{d}_3 \right) - 3b \left(\underline{d}_1 + \bar{d}_3 \right) + 6b\epsilon + b \left(\underline{d}_1 + \bar{d}_3 \right) + 2b\epsilon + 8b^2 \right) \\ &= \frac{1}{\underline{d}_1 + 4b + \bar{d}_3} \left(2(\epsilon - b) \left(\underline{d}_1 + \bar{d}_3 \right) + 8b(b + \epsilon) \right) \\ &\geq \frac{1}{\underline{d}_1 + 4b + \bar{d}_3} \left(2(\epsilon - b) \cdot 2b + 8b(b + \epsilon) \right) \\ &= \frac{1}{\underline{d}_1 + 4b + \bar{d}_3} \left(12b\epsilon + 4b^2 \right) > 0 \end{aligned}$$

for all $\epsilon \in (0, b)$, where the weak inequality follows from the fact that $\underline{d}_1 + \bar{d}_3 \leq 2b$. This yields

$$v^* \left(\frac{\underline{d}_1 + \bar{d}_3}{2} - \epsilon, \frac{\underline{d}_1 + \bar{d}_3}{2} + \epsilon, 3 \right) > v^* \left(\frac{\underline{d}_1 + \bar{d}_3}{2}, \frac{\underline{d}_1 + \bar{d}_3}{2}, 3 \right)$$

for all $\epsilon > 0$.

Due to the symmetry of $v^*(\cdot, \cdot, 3)$ for $\underline{d}_1 \leq \bar{d}_3$ and $\underline{d}_1 \geq \bar{d}_3$, the proof for the fact that the optimal policy with $\underline{d}_1 \geq \bar{d}_3$ is the one with $\underline{d}_1 = \bar{d}_3$ works analogously and is therefore omitted.

$j=4$**:** Since $\underline{d}_1 + \bar{d}_4 \leq 2b$, that is, $\underline{d}_1 \leq 2b - \bar{d}_4$, it results from the convexity of the functions

x^2 and x^3 that
$$v^*(\underline{d}_1, \bar{d}_4, 4) \geq v^*\left(\frac{\underline{d}_1 + \bar{d}_4}{2}, \frac{\underline{d}_1 + \bar{d}_4}{2}, 4\right)$$
by Jensen's inequality.

$j = 5$**:** Since $\underline{d}_1 + \bar{d}_5 \leq 2b$, only the following four cases need to be analyzed: $\underline{d}_1 \geq b \geq \bar{d}_5$, $b \geq \underline{d}_1 \geq \bar{d}_5$, $b \geq \bar{d}_5 \geq \underline{d}_1$ and $\bar{d}_5 \geq b \geq \underline{d}_1$.

First, let's verify that the optimal policy among all with $\underline{d}_1 \geq b \geq \bar{d}_5$ is the one with $\underline{d}_1 = b$: Note that
$$v^*(\underline{d}_1, \bar{d}_5, 5) = v^*(b + \epsilon, \underline{d}_1 + \bar{d}_5 - b - \epsilon, 3)$$
for some $\epsilon \in [0, b]$ for any such $\underline{d}_1, \bar{d}_5$. Its derivative with respect to ϵ is

$$\begin{aligned}
&\frac{1}{\underline{d}_1 + 8b + \bar{d}_5}\left((b+\epsilon)^2 - \left(\underline{d}_1 + \bar{d}_5 - b - \epsilon\right)^2 + 2b(b+\epsilon) - 6b\left(\underline{d}_1 + \bar{d}_5 - b - \epsilon\right) + 8b^2\right)\\
&\geq \frac{1}{\underline{d}_1 + 8b + \bar{d}_5}\left((b+\epsilon)^2 - (2b - b - \epsilon)^2 + 2b(b+\epsilon) - 6b\,(2b - b - \epsilon) + 8b^2\right)\\
&= \frac{1}{\underline{d}_1 + 8b + \bar{d}_5}\left(12b\epsilon + 4b^2\right) > 0
\end{aligned}$$

for all $\epsilon \in (0, b)$, where the weak inequality follows from the fact that $\underline{d}_1 + \bar{d}_5 - b - \epsilon \geq 0$ and $\underline{d}_1 + \bar{d}_5 \leq 2b$. This yields

$$v^*b + \epsilon, \underline{d}_1 + \bar{d}_5 - b - \epsilon, 3) > v^*(b, \underline{d}_1 + \bar{d}_5 - b, 3)$$

for all $\epsilon > 0$.

Next, I show that the optimal policy given $b \geq \underline{d}_1 \geq \bar{d}_5$ is the one satisfying $\underline{d}_1 = \bar{d}_5$. The proof works similarly as the proof for case $j = 3$: The conditional variance can be expressed as
$$v^*(\underline{d}_1, \bar{d}_5, 5) = v^*\left(\frac{\underline{d}_1 + \bar{d}_5}{2} + \epsilon, \frac{\underline{d}_1 + \bar{d}_5}{2} - \epsilon, 5\right)$$
for some $\epsilon \in [0, b]$ for any such $\underline{d}_1, \bar{d}_5$, whose derivative with respect to ϵ is

$$\begin{aligned}
&\frac{1}{\underline{d}_1 + 8b + \bar{d}_5}\left(\left(\frac{\underline{d}_1 + \bar{d}_5}{2} + \epsilon\right)^2 - \left(\frac{\underline{d}_1 + \bar{d}_5}{2} - \epsilon\right)^2 + 2b\left(\frac{\underline{d}_1 + \bar{d}_5}{2} + \epsilon\right) - 6b\left(\frac{\underline{d}_1 + \bar{d}_5}{2} - \epsilon\right) + 8b^2\right)\\
&= \frac{1}{\underline{d}_1 + 8b + \bar{d}_5}\left(2\epsilon\left(\underline{d}_1 + \bar{d}_5\right) + b\left(\underline{d}_1 + \bar{d}_5\right) + 2b\epsilon - 3b\left(\underline{d}_1 + \bar{d}_5\right) + 6b\epsilon + 8b^2\right)\\
&\geq \frac{1}{\underline{d}_1 + 8b + \bar{d}_5}\left(12b\epsilon + 4b^2\right) > 0
\end{aligned}$$

for all $\epsilon[0, b]$, which yields

$$v^*\left(\frac{\underline{d}_1 + \bar{d}_5}{2} - \epsilon, \frac{\underline{d}_1 + \bar{d}_5}{2} + \epsilon, 5\right) > v^*\left(\frac{\underline{d}_1 + \bar{d}_5}{2}, \frac{\underline{d}_1 + \bar{d}_5}{2}, 5\right)$$

for all $\epsilon > 0$.

Since v^* is continuous in both $\underline{d}_1$ and $\bar{d}_5$ everywhere, the optimal policy with $b \geq \underline{d}_1 \geq \bar{d}_5$ (which always exists for any $\underline{d}_1 + \bar{d}_5 \leq 2b$) is better than the optimal policy with $\underline{d}_1 \geq b \geq \bar{d}_5$ (which exists only if $\underline{d}_1 + \bar{d}_5$ is sufficiently large) because the latter one also satisfies $b \geq \underline{d}_1 \geq \bar{d}_5$.

As a result, the policy with $\underline{d}_1 = \bar{d}_5$ is the best one among all with $\underline{d}_1 \geq \bar{d}_5$.

By symmetry, one can deduce that this policy is also optimal among all policies with $\underline{d}_1 \leq \bar{d}_5$ (by analyzing the last two cases $b \geq \bar{d}_5 \geq \underline{d}_1$ and $\bar{d}_5 \geq b \geq \underline{d}_1$) so that it is overall optimal.

$j \geq 6$**:** The proof works analogously and is therefore omitted.

□

Optimal Bi-Pooling Policies

The aim of this section is to determine the optimal bi-pooling policy for a given bias $b \in \left[\frac{1}{2n}, \frac{1}{2(n-1)}\right)$, where $n \geq 2$.

By Lemma 6, all incentive compatibility constraints are binding in the optimal bi-pooling policy of size n. Notice that $\underline{d}_1 + (n-1)2b + \bar{d}_n = 1$, which yields that

$$\underline{d}_1 + \bar{d}_n = 1 - (n-1)2b \in \left(0, \frac{1}{n}\right] \subset (0, 2b).$$

Therefore, the results on the optimal structure of partial bi-pooling policies also apply to the complete bi-pooling policies on the whole state space (simply let $\underline{\omega} = 0$ and $\bar{\omega} = 1$).

So the optimal policy of size n is the one characterized in Lemma 7 and Lemma 8.

With that, one can show that the payoff of the optimal bi-pooling policy of size n is weakly decreasing in b on the interval $\left[\frac{1}{2n}, \frac{1}{2(n-1)}\right)$. Finally, it turns out that for given $b \in \left[\frac{1}{2n}, \frac{1}{2(n-1)}\right)$, the overall optimal policy is the best one of size n if b is sufficiently large, while it is the best one of size $n-1$ if b is sufficiently small.