2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Almost-Bayesian quadratic persuasion with a scalar prior

Olivier Massicot and Cédric Langbort

Abstract— In this article, we consider a problem of strategic communication between a sender (Alice) and a receiver (Bob) akin to the now-traditional model of Bayesian Persuasion introduced by Kamenica & Gentzkow, with the crucial difference that Bob is not assumed Bayesian. In lieu of the Bayesian assumption, Alice assumes that Bob behaves "almost like" a Bayesian agent, in some sense, without resorting to any specific model.

Under this assumption, we study Alice's strategy when both utilities are quadratic and the prior is scalar. We show that, contrary to the Bayesian case, Alice's optimal response may be more subtle than revealing "all or nothing." More precisely, Alice reveals the state of the world when it lies outside a specific interval, and nothing otherwise. This interval increases (and the amount of information shared decreases) as Bob further departs from Bayesianity, much to his detriment.

I. INTRODUCTION

In recent years, the relevance and interest in problems related to strategic information transmission (SIT) have expanded beyond the field of Information Economics and reached communities in decision & control, information theory, and computer science. These communities have explored new applications of SIT ideas, concepts, and modeling paradigms, including adversarial sensing and estimation [1]–[3], persuasive interactions between humans and autonomous agents/vehicles [4]–[6], and congestion mitigation [7]–[13]. Furthermore, these fields have facilitated the investigation of more complex SIT problem formulations, such as communication over limited channels [14], [15] and algorithmic approaches [16].

The Bayesian Persuasion model introduced by Kamenica & Gentzkow [17] involves two actors: a Sender (henceforth referred to as "Alice") who has knowledge of the state of the world and seeks to convince an uninformed Receiver ("Bob") to take actions that benefit her. The model in [17] has two critical features. First, Alice commits to a signaling strategy, creating a Stackelberg game in which she is the leader. This distinguishes it from the perfect Bayesian equilibria considered in the cheap-talk formulation of [18]. This commitment assumption defines the Bayesian Persuasion framework and is present in all extensions of [17], which includes multiple receivers [19], [20] and/or senders [21], costly messages [22], online settings [20], [23], [24], and the possibility of Bob gaining additional information [25]–[27].

This work was supported by the ARO MURI grant W911NF-20-0252 (76582 NSMUR).

O. Massicot and C. Langbort are with the Coordinated Science Lab, University of Illinois Urbana–Champaign, Urbana, IL 61801, USA, om3@illinois.edu and langbort@illinois.edu

The second crucial component of [17] is the assumption that Bob uses Bayes' rule to update his belief after receiving Alice's message. This not only defines the situations the model captures, but also plays a central role in enabling the computation of Alice's signaling policy. In later scenarios where Bob's response and Alice's reward are determined by the mean of the posterior, such as [28]–[30], Gentzkow & Kamenica utilize a result from Blackwell [31], [32] in [29] to parametrize the set of posterior mean distributions that Bob can have upon receiving Alice's message. This reformulation and, hence, Bob's Bayesianity, makes Alice's program theoretically tractable and has been instrumental in most methods for determining her policy (e.g., [16], [25], [29], [30]).

Given the instrumental importance of the way in which Bob is assumed to update his prior in [17], multiple recent studies (e.g., [26], [27], [33]–[35]) have attempted to reconcile the framework of [17] with the empirical fact (confirmed in many behavioral economics experiments such as [36], [37]) that human decision makers can and often do fail to be perfectly Bayesian, either through lack of access to a correct prior, or by accessing or incorrectly (according to Bayes' rule) processing information.

The present work is a continuation of a study of our own [38] which overcomes its main shortcomings by offering an exact solution to the signaling problem of interest, albeit in a simpler situation. Much like [33], we consider Bob to be non-Bayesian, however we do not make any explicit assumption regarding the process replacing Bayes rule. Instead, we model Bob's possible posteriors via a generic robust hypothesis, in a manner resembling the notion of an almost-maximizing agent [39]. More precisely, we assume that, upon receiving Alice's message, Bob's posterior lies within an adequately defined neighborhood of the correct Bayesian posterior, regardless of how it was computed. This notion of "almost-Bayesianity", hinted at in [33] and formalized in [38], sets us apart from other models which either rely on parametric uncertainty (which assume Bob's thought process is known to Alice, save for few parameters such as unknown mismatched prior [34], [35]) or make Alice account for the fact that Bob may receive private side information (before [27] or after [26] messaging).

While this robust hypothesis approach could model Bob's default to apply Bayes' rule in general persuasion and SIT problems, we focus on the particular linear quadratic setting introduced by [40]. Even this relatively simple case presents interesting non-trivial features: much like the celebrated Witsenhausen's counterexample

979-8-3503-0123-6/23/$31.00 ©2023 IEEE

[41], it presents a "linear-quadratic-Gaussian" situation in which linear policies may not be optimal. In addition, and in contrast with Witsenhausen's counterexample, finding the optimal linear policy is itself challenging. These difficulties motivate the present article, in which the prior is scalar.

We show that it is optimal to disclose $\boldsymbol{x}$ when it lies outside of a specific interval, and nothing otherwise. This interval is the unique interval J of length 2α—a parameter that scales linearly with the size of the neighborhood defining the robust hypothesis and otherwise only depends on Alice's utility—such that

$$\mathbb{E}[\boldsymbol{x} \,|\, \boldsymbol{x} \in J] = t_0.$$

This policy provides a subtle compromise between full-disclosure, corresponding to $J = \varnothing$, and no-disclosure, corresponding to $J = \mathbb{R}$. J grows with α, thus Bob receives less information from Alice as she believes he strays further from Bayes' rule, this can only harm him.

We first present the background necessary to place the problem in context in Section II: Sections II-A and II-B summarize our previous study with its shortcomings, whereas Section II-D is a self-contained tutorial on the notion of information in the sense of Blackwell and its application to strategic information design when Bob's best response is a function of the mean of his posterior. Section III-A is dedicated to solving the problem of interest using the recent developments presented in Section II-D, and deriving structural insights as corollaries. Section IV compares the results obtained with our previous findings and discusses their relative importance.

For the sake of brevity, a few simplifying assumptions on the prior have been made and elementary proofs have been omitted. In this article, bold symbols represent random variables and their regular variant denotes a realization.

II. Background

A. Quadratic persuasion

In this article, we study strategic information transmission games where both the sender (Alice) and the uninformed receiver (Bob) have quadratic utilities (precisely in $(\boldsymbol{x}, \boldsymbol{a})$ with the coming notation). Alice chooses a messaging scheme, communicates this choice to Bob, then observes the state of the world $\boldsymbol{x} \in \mathbb{R}^n$ and messages $\boldsymbol{y} \in \mathbb{R}^p$ to Bob according to the messaging scheme she has committed to earlier. Bob then plays action $\boldsymbol{a} \in \mathbb{R}^m$ according to his posterior belief $\mu_{\boldsymbol{y}}$ over $\boldsymbol{x}$. As both utilities are quadratic, a usual reduction is to assume Bob plays $\boldsymbol{a} = \hat{\boldsymbol{x}} = \mathbb{E}[\boldsymbol{x} \mid \boldsymbol{y}]$, the mean square estimate, modifying Alice's utility as necessary. In dimension $n = 1$, these games resemble the original cheap talk setting [18], save for the commitment element that is peculiar to Bayesian persuasion [17].

To fix the notation, we denote by ν the prior distribution of the random variable $\boldsymbol{x}$, it is a Borel probability measure over $\mathbb{R}^n$ which, with further modifications of Alice's utility, can be centered and of covariance I_n. When ρ is a Borel distribution over $\mathbb{R}^n$, we denote by $\bar{\rho}$ its mean and Σ_ρ its covariance, so that $\bar{\nu} = 0$, $\Sigma_\nu = I_n$ and $\hat{\boldsymbol{x}} = \bar{\mu}_{\boldsymbol{y}}$. Much like when the prior is discrete, it is more fruitful to reason using the induced distribution τ of the Bayesian belief $\mu_{\boldsymbol{y}}$ than the signaling policy itself, [17]. In actuality, only the distribution $\hat{\tau}$ of the estimate $\hat{\boldsymbol{x}}$ is ever relevant to the analysis thanks to the quadratic utilities.

Let us denote Alice's cost under realized state x and action a by

$$v(x, a) = \left(\begin{bmatrix} x \\ a \end{bmatrix} - l \right)^\top Q \left(\begin{bmatrix} x \\ a \end{bmatrix} - l \right) + r,$$

where $r \geq 0$, $Q \succeq 0$. As observed in [40], a simple calculation shows that Alice's expected cost is an affine function of the covariance of the estimate $\hat{\boldsymbol{x}}$, $\Sigma_{\hat{\tau}}$:

$$\mathbb{E}[v(\boldsymbol{x}, \boldsymbol{a})] = \operatorname{Tr}(D\Sigma_{\hat{\tau}}) + c,$$

where

$$\begin{aligned} D &= Q_{12} + Q_{21} + Q_{22} \\ c &= \operatorname{Tr} Q_{11} + l^\top Q l + r. \end{aligned}$$

This, added to the fact that $0 \preceq \Sigma_{\hat{\tau}} \preceq \Sigma_\nu = I_n$ no matter the signaling policy, allows [40] to conclude that Alice's optimal cost is at least

$$\min_{0 \preceq X \preceq I_n} \operatorname{Tr}(DX) + c. \tag{1}$$

Being concave, this program admits an orthogonal projection matrix P^* as a solution. When ν is Gaussian [40] (or even just isotropic [38]), signaling $\boldsymbol{y} = P\boldsymbol{x}$ with P orthogonal projection matrix yields $\hat{\boldsymbol{x}} = P\boldsymbol{x}$ and thus $\Sigma_{\hat{\tau}} = P$. As a result, $\boldsymbol{y} = P^*\boldsymbol{x}$ is an optimal messaging policy. Given its importance, we call the signaling policy $\boldsymbol{y} = P\boldsymbol{x}$ the projective policy of (orthogonal projection) matrix P. It reveals the component of $\boldsymbol{x}$ that belongs to the image of P, and eclipses the one in its kernel.

When $n = 1$ and ν is any distribution, the lower bound (1) can be achieved with a projective policy since the only two projective policies are $\boldsymbol{y} = 0$ and $\boldsymbol{y} = \boldsymbol{x}$, which yield $\hat{\boldsymbol{x}} = 0$ and $\hat{\boldsymbol{x}} = \boldsymbol{x}$ respectively. When $n \geq 2$ and ν is not isotropic, however, there might be a gap between the optimal cost of Alice and the lower bound (1). Even when ν is discrete, an optimal solution remains elusive [42].

B. Almost-Bayesian persuasion

There are many reasons why Bob may not follow Bayes' rule exactly, e.g., if he makes computations errors in doing so, if the computation is costly, or if the representation of the posterior distributions are not accurate in the formula. In [33], different models have been considered to replace Bayes' rule, such as conservative Bayesianism, motivated update or Grether's α–β model. The latter all propose an alternative to Bayes' rule, however they require Alice to know precisely how Bob forms his posterior, which seems excessively demanding.

In an attempt to capture Bob's worst behavior without a specific assumption on how he fails to be Bayesian, it seems natural for Alice to consider that the erroneous posterior of Bob always lies within a safety set around the true posterior. In a previous article [38] where we lay the foundations to the study of quadratic persuasion of almost-Bayesian agents, we show that many generic robust hypotheses (for instance that Bob's posterior $\mu'_{\boldsymbol{y}}$ is close to $\mu_{\boldsymbol{y}}$ in a certain statistical sense, or that he plays almost optimally) amount to assuming the response of Bob lies in an ellipsoid of given shape, centered on the mean square estimate given the message, i.e.,

$$\boldsymbol{a} \in \hat{\boldsymbol{x}} + C\mathcal{B},$$

where $\mathcal{B}$ is the Euclidean unit ball.

This calls for a revisit of Alice's expected cost. It is equal to the Bayesian objective, plus the following penalty term,

$$\mathbb{E}\left[\sup_{\eta\in\mathcal{B}}\ 2((Q_{21}+Q_{22})\hat{\boldsymbol{x}} - (Ql)_2)^\top C\eta + \eta^\top C^\top Q_{22} C\eta\right]. \tag{2}$$

When $n \geq 2$, it is not possible to analytically solve the inner program, let alone take its expectation. In [38], we frame this term between two matching bounds (up to a multiplicative ratio, approximately equal to 0.46 when ν is Gaussian). Using the upper bound, we could derive a pessimistic program in terms of the covariance $\Sigma_{\hat{\tau}}$, much like (1),

$$\min_{0\preceq X\preceq I_n}\ \mathrm{Tr}(DX) + c + \bar{\lambda} + \sqrt{f + \mathrm{Tr}(EX)}, \tag{3}$$

where $\bar{\lambda}, f, E$ are constants.

This program is concave, hence admits a solution Q^* that is an orthogonal projection matrix, which thus corresponds to the projective policy of matrix Q^*. We have developed a numerical method to obtain such a Q^*. Under this pessimistic assumption, we also show that Alice tends to share less information with Bob as ϵ grows when the robust hypothesis is $C = \epsilon C_0$, in the sense that the minimal rank of a solution—that is, the minimal number of active canals required—decreases.

This theorem comes with two caveats. First, it only applies to the solutions where the objective is approximated using the pessimistic bound. Second, the notion of information is only captured by the rank of the solution. To circumvent these two hurdles and to do away with approximations, throughout this article we focus on a special case: the scalar case ($n = 1$).

C. Generic program under real prior

From now on, we assume that $n = 1$. Of course, most matrices and vectors are now scalars in disguise. To make this fact more legible, we denote $\epsilon = C \geq 0$, $q = Q_{22} \geq 0$. Further, we focus on the case $D < 0$[1], that is, scaling all parameters accordingly, $D = -1$. In this special case, the inner program of the penalty term (2) can be exactly evaluated:

$$\begin{aligned}\sup_{\eta\in[-1,1]}\ &\epsilon((q-1)\hat{\boldsymbol{x}} - 2(Ql)_2)\eta + \epsilon^2 q\eta^2 \\ &= \epsilon|(q-1)\hat{\boldsymbol{x}} - 2(Ql)_2| + \epsilon^2 q.\end{aligned}$$

In the end, the program of Alice takes the form

$$c + \epsilon^2 q + \inf_{\hat{\tau}\prec\nu}\ \mathbb{E}_{\hat{\tau}}[\epsilon|(q-1)\hat{\boldsymbol{x}} - 2(Ql)_2| - \hat{\boldsymbol{x}}^2], \tag{4}$$

where $\hat{\tau} \prec \nu$ denotes the fact that $\hat{\tau}$ is a distribution of estimates induced by some signaling policy from the prior ν, and the subscript $\hat{\tau}$ on the expectation emphasizes the dependency on $\hat{\tau}$, now a variable. The case $q = 1$ is uninteresting in that both the pessimistic program and the true program are solved by revealing all the information, we thus assume $q \neq 1$. Having defined

$$\begin{aligned}\alpha &= \epsilon|q-1| \geq 0 \\ t_0 &= \frac{2(Ql)_2}{q-1}\end{aligned}$$

solving (4) amounts to solving the generic program

$$V_\alpha^* = \inf_{\hat{\tau}\prec\nu}\ V_\alpha(\hat{\tau}) \tag{5}$$

where

$$V_\alpha(\hat{\tau}) \triangleq \mathbb{E}_{\hat{\tau}}[\alpha|\hat{\boldsymbol{x}} - t_0| - \hat{\boldsymbol{x}}^2].$$

Most of our analysis is devoted to solving this program and establishing the properties of its solutions. Note that here the entire distribution $\hat{\tau}$ matters, not only its covariance. For this reason, we review next a technique that was first introduced by Blackwell [31] and later systematically employed for Bayesian persuasion when the prior is supported on a compact interval and the objective of Alice is ultimately linear in $\hat{\tau}$, see [29].

D. Information in the sense of Blackwell

In his seminal work [31], [32], Blackwell introduced a partial order among distributions called "sufficiency" which, although our present context is rather different, corresponds to that informally defined earlier. The following definition is a reformulation more suited to our needs.

Definition 1. Given two real Borel probability distributions ρ, υ where υ has finite variance, we say that υ is sufficient for ρ, if there exists a stochastic function $\boldsymbol{f}$ of $\boldsymbol{z} \sim \upsilon$ such that $\mathbb{E}[\boldsymbol{z}\,|\,\boldsymbol{f}(\boldsymbol{z})] \sim \rho$.

This order is interesting on multiple accounts. Firstly, it exactly captures what we set earlier: $\hat{\tau} \prec \nu$ if and only if ν is sufficient for $\hat{\tau}$. Secondly, when $\hat{\tau}_1 \prec \hat{\tau}_2$, any decision maker whose utility u is affine in $\boldsymbol{x}$ will perform worse under $\hat{\tau}_1$ than $\hat{\tau}_2$:

$$\int \max_{a\in\mathcal{A}}\ u(\hat{x}, a)\,\mathrm{d}\hat{\tau}_1(\hat{x}) \leq \int \max_{a\in\mathcal{A}}\ u(\hat{x}, a)\,\mathrm{d}\hat{\tau}_2(\hat{x}),$$

[1] When $D \geq 0$ it is optimal to not reveal any information, and this is also reflected in the pessimistic program (3).

simply as a consequence of Jensen's inequality. This also corresponds to $\hat{\tau}_1$ being "less informative" than $\hat{\tau}_2$ as defined in [31]. Thirdly, and most crucially, the order possesses an operative characterization, at least when υ is compactly supported, which was used in [29].

In our case, however, we cannot readily use this characterization as ν is not compactly supported. Nonetheless, it remains possible to use it as a necessary condition (in the form of Principle 1 to come). To do so, we define—very similarly to [29]—for any ρ real probability measure of finite variance:

$$G_\rho(t) \triangleq \int (t-z)^+ \,\mathrm{d}\rho(z).$$

This function is always convex, and by Jensen's inequality, whenever $\rho \prec \upsilon$ are two distributions,

$$G_{\delta_{\bar{\upsilon}}}(t) \le G_\rho(t) \le G_\upsilon(t), \text{ for all } t \in \mathbb{R}.$$

Because of the importance of these properties, we define

$$\mathcal{G}_\upsilon \triangleq \{G \colon \mathbb{R} \to \mathbb{R}, \text{ convex s.t. } G_{\delta_{\bar{\upsilon}}} \le G \le G_\upsilon\}.$$

Principle 1. Consider the following program

$$\inf_{\hat{\tau} \prec \nu} \ \phi(\hat{\tau}). \tag{6}$$

If there exists a function Φ such that $\phi(\hat{\tau}) = \Phi(G_{\hat{\tau}})$ and if there exists $\hat{\tau}^* \prec \nu$ such that

$$\Phi(G_{\hat{\tau}^*}) = \inf_{G \in \mathcal{G}_\nu} \ \Phi(G),$$

then τ^* is an optimal solution of (6).

Indeed, given the remark we have made earlier,

$$\inf_{\hat{\tau} \prec \nu} \ \phi(\hat{\tau}) \ge \inf_{G \in \mathcal{G}_\nu} \ \Phi(G) = \Phi(G_{\hat{\tau}^*}) = \phi(\hat{\tau}^*).$$

In turn, one could first solve the above program in G, and verify whether a solution corresponds to a distribution for which ν is sufficient.

Now that the importance of G_ρ has been established, we turn to its study. Being convex, we denote by

$$\partial G_\rho(t) = [\partial^- G_\rho(t), \partial^+ G_\rho(t)]$$

its subdifferential at t (see [43] for instance). As we show in Lemma 1, the endpoints of this interval are respectively equal to the cumulative distribution functions excluding and including t:

$$F^-_\rho(t) \triangleq \rho((-\infty, t)), \ F^+_\rho(t) \triangleq \rho((-\infty, t]).$$

In other words, G_ρ is a primitive of the cumulative distribution function (itself a primitive of the probability density function, if it exists). As such, kinks in G_ρ encode atoms of ρ, with mass equal to the change of slope, moreover G_ρ is affine on an interval $[a, b]$ as soon as it has no mass on its interior, i.e. when $\rho((a, b)) = 0$.

Lemma 1. When ρ is a real probability measure of finite variance and $\boldsymbol{z} \sim \rho$, G_ρ is convex, finite on $\mathbb{R}$, such that $G_\rho \ge G_{\delta_{\bar{\rho}}}$ and $F^\pm_\rho = \partial^\pm G_\rho$. For all $t \in \mathbb{R}$,

$$\mathbb{E}[|\boldsymbol{z} - t|] = 2G_\rho(t) + \bar{\rho} - t.$$

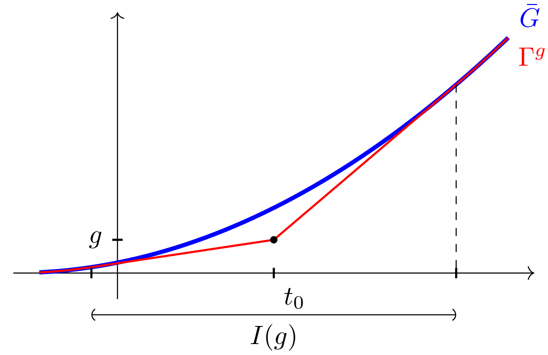

Fig. 1. Graphical representation of Γ^g, solution of (7) under the additional constraint that $G(t_0) = g$.

Noting that $G_{\delta_{\bar{\rho}}}(t) = (t - \bar{\rho})^+$ and denoting the Lebesgue measure by λ, one recovers the variance of ρ from G_ρ:

$$\Sigma_\rho = 2 \int G_\rho - G_{\delta_{\bar{\rho}}} \,\mathrm{d}\lambda.$$

III. Application to Alice's almost-Bayesian problem

In the remainder of this article, we assume ν is continuous with positive density and finite variance, as for instance $\mathcal{N}(0,1)$. Since we will use G_ν and $G_{\delta_{\bar{\nu}}}$ extensively, we define $\bar{G} = G_\nu$ and $\underline{G} = G_{\delta_\nu}$, and since ν is fixed from now on, we drop the dependency in ν and simply denote by $\mathcal{G}$ the set of convex functions framed between $\underline{G}$ and $\bar{G}$. By assumption, $\bar{G}$ is twice continuously differentiable and $\bar{G}'' > 0$.

For the sake of exposition, we devote this section to an informal derivation of our main results, followed by few illustrative examples. We save the rigor of the proofs for the appendix.

A. Main results

Following Principle 1 and relying on Lemma 1, we are brought to define

$$\Phi(G) \triangleq \alpha(2G(t_0) + \bar{\nu} - t_0) - 2 \int G - \underline{G} \,\mathrm{d}\lambda$$

so that

$$\mathbb{E}_{\hat{\tau}}[\alpha|\hat{\boldsymbol{x}} - t_0| - \hat{\boldsymbol{x}}^2] = \Phi(G_{\hat{\tau}}).$$

The following provides a lower bound of program (5)

$$\inf_{G \in \mathcal{G}} \ \Phi(G). \tag{7}$$

The form of Φ is peculiar. Informally, we seek to minimize $G(t_0)$ while maximizing the area between G and $\bar{G}$. We can thus first fix the value $G(t_0)$, find the optimal such convex curve, and later optimize over the value $G(t_0)$. It is quite clear geometrically that the optimal curve G, subject to the additional constraint that $G(t_0) = g$, is none other than the one whose epigraph is the convex hull of the point (t_0, g) and the epigraph of $\bar{G}$. We call Γ^g this curve, see Figure 1.

The function Γ^g is quite remarkable. It is equal to $\bar{G}$ outside a specific interval we denote by $I(g)$, and on $I(g)$

it is equal to the maximum between the tangents to $\bar{G}$ at either end of $I(g)$. Moreover, the tangents meet at the abscissa t_0 with value g, and $I(g_2) \subset I(g_1)$ if $g_1 \leq g_2$.

The distribution of posterior means associated to Γ^g is simple: it is equal to the prior ν outside $I(g)$, and it is a mass at t_0 on $I(g)$. Denoting by ν_A the measure ν restricted to a Borel set A (that is, $\nu_A(B) = \nu(A \cap B)$), and by A^{c} the complement of A, the distribution corresponding to Γ^g is

$$\hat{\tau}^g = \nu_{I(g)^{\mathrm{c}}} + \nu(I(g))\delta_{t_0}.$$

This distribution points to a very simple signaling policy: reveal $\boldsymbol{x}$ when $\boldsymbol{x} \in I(g)$, reveal nothing otherwise; that is, use the signaling policy σ^g defined by

$$\sigma^g(\boldsymbol{x}) \triangleq \begin{cases} \boldsymbol{x} & \text{if } \boldsymbol{x} \notin I(g) \\ t_0 & \text{if } \boldsymbol{x} \in I(g). \end{cases}$$

We verify later that the policy σ^g indeed induces the function Γ^g. All that remains is to optimize $\Phi(\Gamma^g)$ with respect to $g \in [\underline{G}(t_0), \bar{G}(t_0)]$. Using the definition of Γ^g,

$$\frac{\partial}{\partial g}\Phi(\Gamma^g) = 2\alpha - l(g),$$

where $l(g)$ is the diameter of $I(g)$, finite when $g > \underline{G}(t_0)$. The function l is decreasing with $\lim_{\underline{G}(t_0)} l = \infty$ and $l(\bar{G}(t_0)) = 0$. In turn, the unique solution in g is $l^{-1}(2\alpha)$, thus the unique solution of (7) is $\Gamma^{l^{-1}(2\alpha)}$, implementable by $\sigma^{l^{-1}(2\alpha)}$ as stated next.

Theorem 1. An optimal policy to (5) is to send $\boldsymbol{y} = \boldsymbol{x}$ when $\boldsymbol{x} \notin I(l^{-1}(2\alpha))$, nothing otherwise. $I(l^{-1}(2\alpha))$ is the only open interval J of length 2α such that the estimate of $\boldsymbol{x}$ conditioned on the event $\boldsymbol{x} \in J$ is t_0.

This policy is remarkably simple and constitute a subtle compromise between full- and no-information disclosure, corresponding to $\alpha = 0$ and $\alpha = \infty$ respectively. When the prior distribution is symmetric about t_0, the interval J described above in none other that $(t_0 - \alpha, t_0 + \alpha)$, this allows us to state the following.

Corollary 1. When the prior restricted to $(t_0 - \alpha, t_0 + \alpha)$ is symmetric about t_0, an optimal policy to (5) is to send $\boldsymbol{y} = \boldsymbol{x}$ when $|\boldsymbol{x} - t_0| \geq \alpha$, nothing otherwise.

The strongest conclusion we draw from Theorem 1 is that, given the fact that $\Gamma^{l^{-1}(2\alpha_2)} < \Gamma^{l^{-1}(2\alpha_1)}$ when $\alpha_2 > \alpha_1$, Alice shares less information in the sense of Blackwell with Bob as he departs further from Bayesianity, which only harms him. This cements the same conclusion we had previously drawn in [38], albeit in the special case of a scalar prior. Formally put, we have the following.

Corollary 2. If $\hat{\tau}_1$, $\hat{\tau}_2$ are solution of (5) at respectively $\alpha_1 < \alpha_2$, then $\hat{\tau}_2 \prec \hat{\tau}_1$.

B. Illustrative examples

As a first example, let us outline a numerical method to compute an optimal signaling policy. Define

$$T_s(t_0) \triangleq \bar{G}(s) + \bar{G}'(s)(t_0 - s),$$

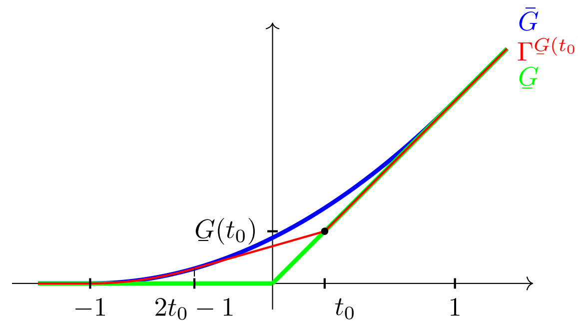

Fig. 2. Solution for $\nu = \mathcal{U}([-1, 1])$ when $0 \leq t_0 < 1 - \alpha$.

so that $I(l^{-1}(2\alpha)) = (s^* - \alpha, s^* + \alpha)$ where s^* is the unique root in s of

$$T_{s+\alpha}(t_0) - T_{s-\alpha}(t_0) = \mathbb{E}[(t_0 - \boldsymbol{x})\mathbb{1}_{(s-\alpha,s+\alpha)}(\boldsymbol{x})].$$

The latter only increases in s and so s^* can be efficiently computed. For instance, considering the standard Gaussian distribution $\nu = \mathcal{N}(0, 1)$ with $\alpha = 1$ and $t_0 = 1$, we find

$$I(l^{-1}(2\alpha)) \approx (0.364, 2.364),$$

whereas for $\alpha = 1.1$, we find

$$I(l^{-1}(2\alpha)) \approx (0.343, 2.543)$$

which is in agreement with Corollary 2.

Although our assumption requires that ν be continuous with positive density over $\mathbb{R}$, our results can be carefully extended to the cases where the density is piecewise continuous and not necessarily positive. More specifically, the unique optimal solution to (7) remains Γ^g for a certain g, which is still induced by σ^g, and $g \mapsto \Phi(\Gamma^g)$ is still convex though not necessarily continuously differentiable or even strictly convex.

As a second example, consider the uniform distribution over $[-1, 1]$. If $t_0 \notin (-1, 1)$, $[\underline{G}(t_0), \bar{G}(t_0)] = \{\bar{G}(t_0)\}$, thus it is optimal for Alice to fully reveal $\boldsymbol{x}$. Assume now that $t_0 \in [0, 1]$ without loss of generality. If $t_0 + \alpha \leq 1$, ν restricted to $(t_0 - \alpha, t_0 + \alpha)$ is symmetric about t_0, hence it is optimal for Alice to reveal $\boldsymbol{x}$ if $|\boldsymbol{x} - t_0| \geq \alpha$, in a manner analogous to Corollary 1. If on the other hand $t_0 + \alpha > 1$, $l(g) < 2\alpha$ for all $g > \underline{G}(t_0)$, and therefore it is optimal to reveal $\boldsymbol{x}$ if $\boldsymbol{x} \leq 2t_0 - 1$, following the policy $\sigma^{\underline{G}(t_0)}$ (see Figure 2).

IV. Conclusion

Focusing on the scalar case (and further assuming the prior is continuous with positive density), we have been able to solve the problem of almost-Bayesian quadratic persuasion. An optimal policy consists in confounding all values $\boldsymbol{x}$ of a given interval and signaling truthfully the others, the interval increasing as Bob departs further from Bayes' rule. This agrees with the intuition as Alice tries to minimize the absolute deviation of $\hat{\boldsymbol{x}}$ from t_0 while maximizing its variance, whereas the weight on the former increases as Bob is less Bayesian.

Additionally, the interval of values Alice does not disclose takes a simple form. When α is the relative coefficient in front of the deviation of $\hat{\boldsymbol{x}}$ from t_0, this specific interval is the unique interval of length 2α for which the conditional expectation yields t_0. Even when no closed-form expression is available, it is possible to compute this interval numerically.

Beyond confirming that linear policies are not optimal, this study shows that under appropriate assumptions the information Alice is willing to share with Bob decreases with his lack of Bayesianity. This reinforces the similar observation made in [38], although this time, thanks to Blackwell's hierarchy of information, this is absolute: Bob is worse off being less Bayesian in the eyes of Alice, no matter his objective.

References

[1] F. Farokhi, A. M. Teixeira, and C. Langbort, "Estimation with strategic sensors," IEEE Transactions on Automatic Control, vol. 62, no. 2, pp. 724–739, 2016.

[2] F. Farokhi, H. Sandberg, I. Shames, and M. Cantoni, "Quadratic gaussian privacy games," in 2015 54th IEEE conference on decision and control (CDC), pp. 4505–4510, IEEE, 2015.

[3] E. Kazıklı, S. Gezici, and S. Yüksel, "Quadratic privacy-signaling games and the mmse information bottleneck problem for gaussian sources," IEEE Transactions on Information Theory, 2022.

[4] V. Hebbar and C. Langbort, "A stackelberg signaling game for human-uav collaboration in a search-and-rescue context," IFAC-PapersOnLine, vol. 53, no. 5, pp. 297–302, 2020.

[5] C. Peng and M. Tomizuka, "Bayesian persuasive driving," in 2019 American Control Conference (ACC), pp. 723–729, IEEE, 2019.

[6] M. Le Treust and T. Tomala, "Information design for strategic coordination of autonomous devices with non-aligned utilities," in 2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 233–242, IEEE, 2016.

[7] S. Das, E. Kamenica, and R. Mirka, "Reducing congestion through information design," in 2017 55th annual allerton conference on communication, control, and computing (allerton), pp. 1279–1284, IEEE, 2017.

[8] Y. Zhu and K. Savla, "On the stability of optimal bayesian persuasion strategy under a mistrust dynamics in routing games," in 2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 92–99, IEEE, 2018.

[9] Y. Zhu and K. Savla, "On routing drivers through persuasion in the long run," in 2019 IEEE 58th Conference on Decision and Control (CDC), pp. 4091–4096, IEEE, 2019.

[10] O. Massicot and C. Langbort, "Public signals and persuasion for road network congestion games under vagaries," IFAC-PapersOnLine, vol. 51, no. 34, pp. 124–130, 2019.

[11] O. Massicot and C. Langbort, "Competitive comparisons of strategic information provision policies in network routing games," IEEE Transactions on Control of Network Systems, 2021.

[12] Y. Zhu and K. Savla, "Information design in non-atomic routing games with partial participation: Computation and properties," IEEE Transactions on Control of Network Systems, 2022.

[13] B. L. Ferguson, P. N. Brown, and J. R. Marden, "Avoiding unintended consequences: How incentives aid information provisioning in bayesian congestion games," arXiv preprint arXiv:2204.06046, 2022.

[14] M. Le Treust and T. Tomala, "Persuasion with limited communication capacity," Journal of Economic Theory, vol. 184, p. 104940, 2019.

[15] A. S. Vora and A. A. Kulkarni, "Information extraction from a strategic sender over a noisy channel," in 2020 59th IEEE Conference on Decision and Control (CDC), pp. 354–359, IEEE, 2020.

[16] S. Dughmi and H. Xu, "Algorithmic bayesian persuasion," in Proceedings of the forty-eighth annual ACM symposium on Theory of Computing, pp. 412–425, 2016.

[17] E. Kamenica and M. Gentzkow, "Bayesian persuasion," American Economic Review, vol. 101, no. 6, pp. 2590–2615, 2011.

[18] V. P. Crawford and J. Sobel, "Strategic information transmission," Econometrica: Journal of the Econometric Society, pp. 1431–1451, 1982.

[19] M. Gentzkow and E. Kamenica, "Bayesian persuasion with multiple senders and rich signal spaces," Games and Economic Behavior, vol. 104, pp. 411–429, 2017.

[20] M. Castiglioni, A. Marchesi, A. Celli, and N. Gatti, "Multi-receiver online bayesian persuasion," in International Conference on Machine Learning, pp. 1314–1323, PMLR, 2021.

[21] Y. Wang, "Bayesian persuasion with multiple receivers," Available at SSRN 2625399, 2013.

[22] A. Nguyen and T. Y. Tan, "Bayesian persuasion with costly messages," Journal of Economic Theory, vol. 193, p. 105212, 2021.

[23] M. Castiglioni, A. Celli, A. Marchesi, and N. Gatti, "Online bayesian persuasion," Advances in Neural Information Processing Systems, vol. 33, pp. 16188–16198, 2020.

[24] Y. Zu, K. Iyer, and H. Xu, "Learning to persuade on the fly: Robustness against ignorance," in Proceedings of the 22nd ACM Conference on Economics and Computation, pp. 927–928, 2021.

[25] L. Matyskova, "Bayesian persuasion with costly information acquisition," cERGE-EI Working Paper Series, no. 614, 2018.

[26] P. Dworczak and A. Pavan, "Preparing for the worst but hoping for the best: Robust (bayesian) persuasion," Econometrica, vol. 90, no. 5, pp. 2017–2051, 2022.

[27] J. Hu and X. Weng, "Robust persuasion of a privately informed receiver," Economic Theory, vol. 72, no. 3, pp. 909–953, 2021.

[28] A. Kolotilin, T. Mylovanov, A. Zapechelnyuk, and M. Li, "Persuasion of a privately informed receiver," Econometrica, vol. 85, no. 6, pp. 1949–1964, 2017.

[29] M. Gentzkow and E. Kamenica, "A Rothschild-Stiglitz approach to Bayesian persuasion," American Economic Review, vol. 106, no. 5, pp. 597–601, 2016.

[30] O. Candogan and P. Strack, "Optimal disclosure of information to privately informed agents," Available at SSRN 3773326, 2021.

[31] D. Blackwell, "Comparison of experiments," in Proceedings of the Second Berkley Symposium on Mathematical Statistics and Probability, pp. 93–102, University of California Press, 1951.

[32] D. Blackwell, "Equivalent comparisons of experiments," The annals of mathematical statistics, pp. 265–272, 1953.

[33] G. de Clippel and X. Zhang, "Non-bayesian persuasion," Journal of Political Economy, vol. 130, no. 10, pp. 2594–2642, 2022.

[34] E. Kazıklı, S. Sarıtaş, S. Gezici, and S. Yüksel, "Optimal signaling with mismatch in priors of an encoder and decoder," arXiv preprint arXiv:2101.00799, 2021.

[35] S. Kosterina, "Persuasion with unknown beliefs," Theoretical Economics, vol. 17, no. 3, pp. 1075–1107, 2022.

[36] C. Camerer, "Bounded rationality in individual decision making," Experimental economics, vol. 1, no. 2, pp. 163–183, 1998.

[37] D. J. Benjamin, "Errors in probabilistic reasoning and judgment biases," Handbook of Behavioral Economics: Applications and Foundations 1, vol. 2, pp. 69–186, 2019.

[38] O. Massicot and C. Langbort, "Almost-bayesian quadratic persuasion (extended version)," 2022. Available at https://arxiv.org/abs/2212.13619.

[39] H. D. Dixon, Surfing Economics. Bloomsbury Publishing, 2017.

[40] W. Tamura, "Bayesian persuasion with quadratic preferences," Available at SSRN 1987877, 2018.

[41] H. S. Witsenhausen, “A counterexample in stochastic optimum control,” SIAM Journal on Control, vol. 6, no. 1, pp. 131–147, 1968.
[42] M. O. Sayin and T. Başar, “Bayesian persuasion with state-dependent quadratic cost measures,” IEEE Transactions on Automatic Control, vol. 67, no. 3, pp. 1241–1252, 2021.
[43] J.-B. Hiriart-Urruty and C. Lemaréchal, Fundamentals of convex analysis. Springer Science & Business Media, 2004.

Appendix

To establish the claims made previously, we first show that Γ^g is the unique solution of the constrained program

$$\begin{aligned} \inf_{G \in \mathcal{G}} \quad & \alpha(2G(t_0) + \bar{\nu} - t_0) - 2\int G - \underline{G}\,\mathrm{d}\lambda \\ \text{s.t.} \quad & G(t_0) = g. \end{aligned} \tag{8}$$

Lemma 2. Γ^g is the unique solution of (8).

Proof. The first term is equal to $\alpha(2g + \bar{\nu} - t_0)$, hence is constant and can leave the minimization. This program then amounts to maximizing the area between G and $\underline{G}$ where $G \in \mathcal{G}$ takes value g at t_0. To fix things, call $\mathcal{E}$ the epigraph of $\bar{G}$ and $\mathcal{E}^g$ the convex hull of $\{(t_0, g)\} \cup \mathcal{E}$. By definition, Γ^g is a valid argument of (8). Moreover, if G is a different valid argument, its epigraph must contain both $\mathcal{E}$ and (t_0, g), so that $G \le \Gamma^g$ and $G \ne \Gamma^g$ (which are both continuous) and in particular $\Phi(\Gamma^g) < \Phi(G)$. ☐

Next, we establish a closed-form expression for Γ^g. To do so, we denote by T_s the tangent to $\bar{G}$ at $s \in [-\infty, \infty]$ (where $T_{-\infty}(t) = 0$ and $T_\infty(t) = t - \bar{\nu}$), and by h the function which to s associates the value $T_s(t_0)$, i.e.,

$$h\colon s \in [-\infty, \infty] \longmapsto T_s(t_0) = \bar{G}(s) + \bar{G}'(s)(t_0 - s) \in [\underline{g}, \bar{g}],$$

where for lighter expressions, we have defined

$$\underline{G}(t_0) = \underline{g},\ \ \bar{G}(t_0) = \bar{g}.$$

The function h is increasing and decreasing on $[-\infty, t_0]$ and $[t_0, \infty]$ respectively, so we may invert it and define

$$\xi_A \triangleq h\big|_{[-\infty, t_0]}^{-1},\ \ \xi_B \triangleq h\big|_{[t_0, \infty]}^{-1}$$

monotonic functions, continuously differentiable on the interior of their domain, $(\underline{g}, \bar{g})$. With this notation,

$$I(g) = (\xi_A(g), \xi_B(g)).$$

Lemma 3. For all $g \in [\underline{g}, \bar{g}]$ and $t \in \mathbb{R}$,

$$\Gamma^g(t) = \begin{cases} \bar{G}(t) & \text{if } t \le \xi_A(g) \\ T_{\xi_A(g)}(t) & \text{if } \xi_A(g) < t \le t_0 \\ T_{\xi_B(g)}(t) & \text{if } t_0 < t < \xi_B(g) \\ \bar{G}(t) & \text{if } \xi_B(g) \le t. \end{cases}$$

Proof. Denote by G the function introduced in the statement. By definition, it belongs to $\mathcal{G}$ and $G(t_0) = g$, thus $G \le \Gamma^g$. Conversely, take a point $D = (t, z)$ in the epigraph of G. If $D \in \mathcal{E}$ or $D = (t_0, g)$, by definition $D \in \mathcal{E}^g$. Assume otherwise, then $t \in I(g)$. If $t \le t_0$, the line that connects D, C^g crosses the graph of $\bar{G}$ at some point E such that $E, D, (t_0, g)$ are aligned in this order.

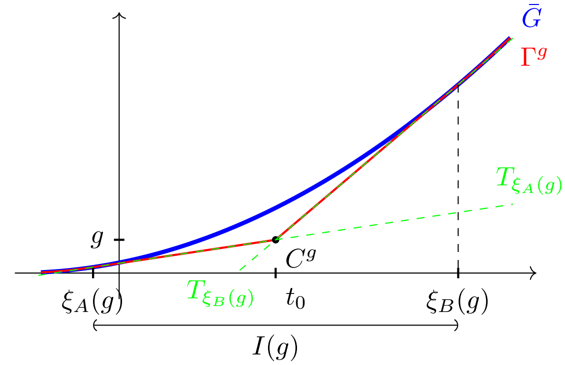

Fig. 3. Graphical representation of the notation introduced.

Then $D \in [(t_0, g), E]$ and similarly if $t \ge t_0$, in all cases D belongs to the adherence of $\mathcal{E}^g$, and thus $G \ge \Gamma^g$. ☐

Using this expression, we can now verify that Γ^g corresponds to the announced distribution of posterior means, τ^g, and that this latter is induced by policy σ^g.

Lemma 4. $\Gamma^g = G_{\tau^g}$ and τ^g is induced by the policy σ^g.

Proof. From the form of σ^g, we can already deduce that the induced distribution of posterior means has the form

$$\hat{\tau} = \nu_{I(g)^{\mathsf{c}}} + \nu(I(g))\delta_{s_0}$$

where $s_0 \in I(g)$ is the conditional expectation of $\boldsymbol{x}$ given that $\boldsymbol{x} \in I(g)$. In turn $G_\tau = \bar{G}$ on $I(g)^{\mathsf{c}}$ and G_τ is tangent to $\bar{G}$ at $\xi_A(g)$ and $\xi_B(g)$. Continuity of G_τ allows us to conclude that $s_0 = t_0$, and so $G_\tau = \Gamma^g$. ☐

The following lemma, whose proof is a simple calculation, omitted for the sake of brevity, is the last piece of the puzzle in establishing Theorem 1.

Lemma 5. The function $g \in [\underline{g}, \bar{g}] \mapsto \Phi(\Gamma^g)$ is continuous, continuously differentiable on $(\underline{g}, \bar{g}]$, strictly concave and

$$\frac{\partial}{\partial g}\Phi(\Gamma^g) = 2\alpha - l(g).$$

Proof of Theorem 1. Combining Lemmas 2 and 5, $\Gamma^{l^{-1}(2\alpha)}$ is the unique solution of (7). Lemma 4 asserts this solution corresponds to the distribution of mean posteriors $\tau^{l^{-1}(2\alpha)}$, product of the signaling policy $\sigma^{l^{-1}(2\alpha)}$. Principle 1 implies that $\sigma^{l^{-1}(2\alpha)}$ is an optimal policy. The interval $I(l^{-1}(2\alpha))$ has length 2α by definition and the associated conditional expectation is t_0. This uniquely characterizes $I(l^{-1}(2\alpha))$ since $s \mapsto \mathbb{E}[\boldsymbol{x} \mid \boldsymbol{x} \in [s, s + 2\alpha]]$ is strictly increasing.

We should emphasize that, although there are multiple optimal signaling policies, the optimal distribution of posteriors means is unique. This is because $\Gamma^{l^{-1}(2\alpha)}$ is the unique solution of (7). ☐

Proof of Corollary 2. Let $\hat{\tau}_1$, $\hat{\tau}_2$ be solution of (5) at respectively $\alpha_1 < \alpha_2$. Considering $\hat{\tau}_1$ as the prior, $\hat{\tau}_2$ is the distribution of posterior means induced by the signaling policy σ^2, therefore $\hat{\tau}_2 \prec \hat{\tau}_1$. ☐