Working Paper Series 614
(ISSN 1211-3298)

Bayesian Persuasion with Costly Information Acquisition

Ludmila Matysková

CERGE-EI
Prague, March 2018

ISBN 978-80-7343-421-2 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-457-0 (Národohospodářský ústav AV ČR, v. v. i.)

Bayesian Persuasion with Costly Information Acquisition*

Ludmila Matysková†

March 13, 2018

Abstract

A sender who chooses a signal to reveal to a receiver can often influence the receiver's subsequent actions. Is persuasion more difficult when the receiver has her own sources of information? Does the receiver benefit from having additional information sources? We consider a Bayesian persuasion model extended to a receiver's endogenous acquisition of information under an entropy-based cost commonly used in rational inattention. A sender's optimal signal can be computed from standard Bayesian persuasion subject to an additional constraint: the receiver never gathers her own costly information. We further determine a finite set of the sender's signals satisfying the additional constraint in which some optimal signal must be contained. The set is characterized by linear conditions using the receiver's utility and information cost parameters. The new method is also applicable to a standard Bayesian persuasion model and can simplify, sometimes dramatically, the search for a sender's optimal signal (as opposed to a standard concavification technique used to solve these models). We show that the 'threat' of additional learning weakly decreases the sender's expected equilibrium payoff. However, the outcome can be worse not only for the sender, but also for the receiver.

Keywords: Bayesian persuasion, Rational inattention, Costly information acquisition, Information design
JEL classification: D72, D81, D82, D83

*I thank the following individuals for their valuable comments: Andrew Caplin, Yeon-Koo Che, Mark Dean, Navin Kartik, Andrei Maatvenko, Filip Matějka, Laurent Mathevet, Konrad Mierendorff, Vladimír Novák, Pietro Ortoleva, Avner Shaked, Jakub Steiner, Michael Woodford, Jan Zápal, participants in the Cognition and Decision Lab Meeting and the Microtheory Colloquium at Columbia University, Brown Bag seminar participants at CERGE-EI, and participants in the 28th Jerusalem Summer School in Economic Theory. The author was supported by the H2020-MSCA-RISE project GEMCLIME-2020 GA No. 681228, by Charles University project GA UK no. 178715, and by the ERC-2015-STG (H2020) project no. 678081.

†CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences, Politických vězňů 7, 111 21 Prague, Czech Republic;
email: lmatysko@cerge-ei.cz

1 Introduction

Does a buyer/a politician benefit from the ability to acquire her own information in addition to a seller's/a lobbyist's information? A decision maker often relies on free information provided by an interested party, but she may also be able to obtain her own information at a costly effort. Will she choose to acquire any? Will she benefit from the *threat* of acquiring additional information?

We consider a Bayesian persuasion model (Kamenica and Gentzkow, 2011; henceforth KG) extended to an *endogenous* acquisition of costly information. As in KG, a sender chooses a signal conveying information about an unknown state of the world to disclose to a receiver, a decision-maker. However, unlike in KG, before taking action, the receiver further chooses her own signal under an entropy-based cost, as in rational inattention (Sims, 2003). We show that the possibility of additional learning reduces the sender's persuasive power (lowering his expected equilibrium utility). However, the outcome can also be worse for the receiver, as the sender can strategically prefer to disclose significantly less information when the receiver has her own learning option. For instance, the sender's strategic manipulation of a receiver's consideration set or his dislike for particular actions can lead to such a scenario, see Section 5.

We exploit a similarity in Bayesian persuasion and rational inattention allowing for tractability: any signal is feasible as long as it is consistent with prior beliefs. Signals can thus be explicitly modeled by posterior distributions over unknown states under a martingale property (KG; Caplin and Dean, 2013; henceforth CD). An optimal signal is then found by a *concavification* of an underlying value function of posterior beliefs related to the expected utility of the information designer. The concavification method remains valid in our model, but it requires solving the receiver's maximization problems for an entire space of beliefs first, which quickly becomes intractable.

We propose a new method not relying on concavification; it is sufficient to search through a relatively small finite set of the sender's signals, characterized by a series of linear conditions. The method is also applicable to a KG model, which is a limiting case of our model, and can simplify, sometimes dramatically, the search for a sender's optimal signal. First, our model can be solved as a standard Bayesian persuasion under an additional constraint: the receiver never costly learns, captured in a Never-Learning Lemma. This results from both the sender's and the receiver's information technology being unconstrained (apart from the martingale property) and certain properties of the receiver's cost function. Second, we construct a finite set of the sender's signals satisfying the additional constraint and in which some optimal strategy must be contained. The new method complements the result of Lipnowski and Mathevet (2017) who show, in the standard Bayesian persuasion model, sufficiency to consider a properly chosen subset of the sender's signals. While they provide general abstract conditions on the subset, we provide exact linear conditions which follow from the entropy-based cost, but are also valid for the standard Bayesian persuasion at the limit.

We conclude the paper by examining the robustness to variations of the receiver's cost. The key simplification step (the Never-Learning Lemma) holds for a whole class of *posterior-separable* cost functions, for which the entropy-based cost is a prime example. Once the information technology is less flexible, this simplification need not hold; the sender can take advantage of the restriction on the set of feasible receiver's signals. However, the possibility that the receiver can be hurt by having the option to learn is not unique to posterior-separable cost functions.

The paper is organized as follows. Section 2 sets up a motivating example. Section 3 provides a general model. Section 4 states the main simplification result and describes the new solution method. Section 5 describes comparative statics, giving examples in which the receiver does not benefit from having the learning option. Section 6 discusses the assumption of the receiver's information technology. Section 7 gives an an overview of the relevant literature and Section 8 concludes.

2 Motivation: Simple model

A seller (he) is persuading a buyer (she) to purchase his product (e.g., a music CD), which can be either a good match ($\omega = 1$) or a bad match ($\omega = 0$). The buyer can either buy ($a = 1$) or not buy ($a = 0$). $u(a, \omega)$ and $v(a)$ are the buyer's and seller's utilities,

$$u(a,\omega) = \begin{cases} 1 & a = 1 \ \wedge \ \omega = 1 \\ -1 & a = 1 \ \wedge \ \omega = 0 \\ 0 & \text{otherwise} \end{cases}, \qquad v(a) = \begin{cases} 1 & a = 1 \\ 0 & a = 0 \end{cases}.$$

We identify the beliefs with probability that $\omega = 1$. A common prior belief is $\mu_0 := \Pr[\omega = 1] < 0.5$ (under which the buyer's optimal action is not to buy).

The seller may persuade the buyer to take his preferred action (to buy) by providing further information (e.g., let her listen to a song). The buyer then updates her priors to an interim belief $\mu := \Pr[\omega = 1|\text{seller's information}]$. The seller's information strategy is a choice of a lottery $\tau \in \Delta([0, 1])$ over interim beliefs with mean μ_0.

After the buyer updates to a particular interim belief μ, she can gather additional information at a costly effort (e.g. search on the Internet), further updating her beliefs to a posterior belief $\gamma = \Pr[\omega = 1|\text{seller's and buyer's information}]$. Her information strategy is a choice of a lottery $\phi \in \Delta([0, 1])$ over posterior beliefs with mean μ. If the optimal lottery satisfies $\text{supp}(\phi) = \mu$, we say she *does not learn* at μ. Otherwise, we say she *learns* at μ.

While the seller's information is costless, the buyer bears a cost for her information (e.g. opportunity cost of time). Given μ, the cost of a lottery ϕ is $\lambda \left(H(\mu) - \mathbb{E}_\phi H(\gamma)\right)$, where $\lambda \geq 0$ is an information cost parameter and $H(\mu) - \mathbb{E}_\phi H(\gamma) \geq 0$ states how much uncertainty about the match, as measured by Shannon entropy[1] $H(\cdot)$, is expected to be reduced by ϕ.

[1] The Shannon entropy at belief $p \in [0, 1]$ is $H(p) = -(p \ln p + (1-p) \ln(1-p))$ where $0 \ln 0 = 0$.

To solve the game, we exploit one feature of the buyer's optimal behavior: once she obtains her chosen information (updates to a particular posterior γ), she never wishes to engage in another round of learning even if given a chance. This stems from a set of the buyer's information strategies being unconstrained (apart from the consistency requirement that a mean is preserved) and from her cost function being *posterior-separable* (see Section 6). The latter guarantees that the cost is increasing in Blackwell informativeness and that it is invariant to intermediate stages[2]. As the seller's set of information strategies is also unconstrained, he can always skip the buyer's potential learning and directly 'send' her to the corresponding posteriors where she would have ended up by herself, without changing the outcome of the game. Since the seller's information is costless, it is thus sufficient to focus on a specific class of the seller's strategies under which the buyer never decides to further costly learn.

Buyer's optimal behavior

We follow an approach of CD to solve for the buyer's optimal behavior. Given μ, the buyer maximizes

$$\begin{aligned} \max_{\phi\in\Delta([0,1])} \quad & \mathbb{E}_\phi[B(\gamma)] - \lambda\left(H(\mu) - \mathbb{E}_\phi[H(\gamma)]\right) \\ s.t. \quad & \mathbb{E}_\phi[\gamma] = \mu, \end{aligned} \tag{1}$$

where the expectation is taken over posterior beliefs induced by ϕ and $B(\gamma)$ is the buyer's gross expected utility at posterior γ given that her subsequent action is optimal. Hence, $B(\gamma) = 0$ for $\gamma < 1/2$ (not buying) and $B(\gamma) = 2\gamma - 1$ otherwise (buying).

[2]The cost of achieving a particular distribution of posterior beliefs would be the same regardless of whether the learning occurs in one or more stages.

Figure 1: Buyer's value function $\hat{u}(\gamma)$ and its concavification $U(\gamma)$

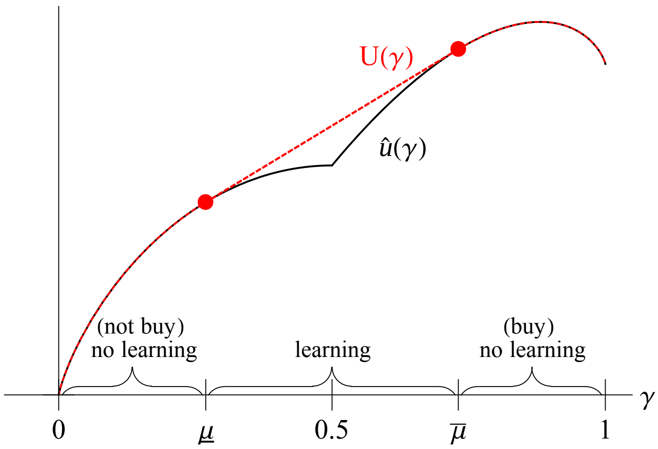

Note that the problem (1) can be rewritten as

$$\begin{aligned} \max_{\phi \in \Delta([0,1])} \quad & \mathbb{E}_\phi[\hat{u}(\gamma)] - \underbrace{\lambda H(\mu)}_{=const.} \\ s.t. \quad & \mathbb{E}_\phi[\gamma] = \mu, \end{aligned} \tag{2}$$

where $\hat{u}(\gamma) = B(\gamma) + \lambda H(\gamma)$ is the buyer's value function at posterior γ. The problem (2) has a geometric interpretation. Let $U(\gamma)$ be a *concavification* of $\hat{u}(\gamma)$ defined as the smallest concave function that is everywhere weakly greater than $\hat{u}(\gamma)$. CD showed that the support of an optimal lottery ϕ^* are those posterior beliefs that support the tangent hyperplane to the lower epigraph of the concavification above the interim belief μ, $U(\mu)$. Hence, whenever $\hat{u}(\mu) = U(\mu)$, the receiver does not learn at μ and whenever $\hat{u}(\mu) < U(\mu)$, she learns at μ, where the support of the optimal lottery is always the same: $\text{supp}(\phi^*) = \{\underline{\mu}, \overline{\mu}\}$, see Fig. 1.

Hence, there are two threshold interim beliefs[3] $0 \leq \underline{\mu} \leq \overline{\mu} \leq 1$ that divide the space of interim beliefs into two *non-learning* and one *learning regions*. A non-learning region of a particular action are all interim beliefs at which the buyer does not learn and optimally takes that action. A non-learning region of not buying is $[0, \underline{\mu}]$

[3]Solving the buyer's maximization problem, we obtain $\underline{\mu} = \frac{1}{1+e^{\frac{1}{\lambda}}}$ and $\overline{\mu} = \frac{e^{\frac{1}{\lambda}}}{1+e^{\frac{1}{\lambda}}}$ (see Appendix A).

Figure 2: Function $\hat{v}(\mu)$ and sender's optimal information strategy

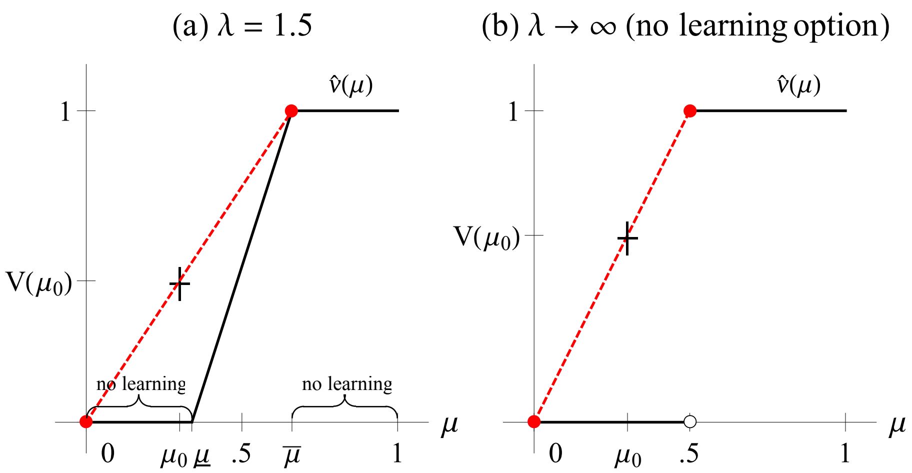

and that of buying is $[\overline{\mu}, 1]$. For intermediate values of μ, when the buyer is very uncertain about what the right thing to do is, she learns and only sometimes buys (in the case of favorable information). Note that once she obtains her information (updates her beliefs either to a posterior $\underline{\mu}$ or $\overline{\mu}$), she does not wish to engage in another round of learning even if given the chance[4].

Bayesian persuasion s.t. never-learning constraint

For each μ, let $\hat{v}(\mu)$ be the seller's expected utility which already accounts for the optimal buyer's behavior at μ.[5] Let $V(\mu)$ be a *concavification* of $\hat{v}(\mu)$ defined as the smallest concave function that is everywhere weakly greater than $\hat{v}$. Then the seller's expected equilibrium utility is the concavification evaluated at the prior, $V(\mu_0)$, and the support of the optimal sender's lottery can be found from the graph in the same fashion as in the buyer's problem. See Fig. 2 for an example of $\hat{v}(\mu)$ and the resulting optimal sender's strategy with $\lambda \to \infty$ (equivalent to KG's setting with a buyer who cannot gather her own information) and with $\lambda = 1.5$.

[4] As the buyer's value function $\hat{u}(\gamma)$ and its concavification $U(\gamma)$ coincide at values $\underline{\mu}$ and $\overline{\mu}$, no learning is optimal either at $\underline{\mu}$ or at $\overline{\mu}$.

[5] Hence, $\hat{v}(\mu) = 0$ if $\mu \leq \underline{\mu}$ (the buyer does not learn and does not buy), $\hat{v}(\mu) = 1$ if $\mu \geq \overline{\mu}$ (the buyer does not learn and buys), and $\hat{v}(\mu) = \frac{1}{\overline{\mu}-\underline{\mu}}\mu - \frac{\underline{\mu}}{\overline{\mu}-\underline{\mu}}$ for $\mu \in (\underline{\mu}, \overline{\mu})$.

The support of the optimal lottery is $\text{supp}(\tau^*) = \{0, \overline{\mu}\}$ (when $\lambda \to \infty$, $\underline{\mu} = \overline{\mu} = 1/2$). These are beliefs that belong to non-learning regions. We show that it is generally sufficient to only consider the seller's information strategies under which the buyer never costly learns, i.e., the lotteries with the support over interim beliefs from non-learning regions only. This is because the buyer never wishes to have more than one round of costly learning, the set of the seller's information strategies is unconstrained, and his information is costless. The seller can then skip the receiver's learning part with his information without changing the outcome of the game.

Further, note that the support of the optimal lottery, $\text{supp}(\tau^*) = \{0, \overline{\mu}\}$, are *extreme points*[6] of the non-learning regions. We show that it is sufficient to consider the sender's strategies under which only extreme points of non-learning regions are chosen. Hence, in this example, one only needs to consider lotteries with support over interim beliefs from the set $\{0, \underline{\mu}, \overline{\mu}, 1\}$, where the thresholds $\underline{\mu}$, $\overline{\mu}$ are specified by particular linear equations resulting from the characterization of the optimal receiver's strategy.

3 General model

A receiver (she) chooses an action a from a finite set A. A payoff-relevant state ω is drawn from a finite set Ω according to an interior prior distribution $\mu_0 \in \Delta(\Omega)$. Before choosing her action, the receiver obtains free information about ω provided by a sender (he) and rationally updates her beliefs from the prior μ_0 to interim belief $\mu \in \Delta(\Omega)$. A sender's (information) strategy is a choice of a distribution $\tau \in \Delta(\Delta(\Omega))$ over the (updated) interim beliefs s.t. $\mathbb{E}_\tau[\mu] = \mu_0$ (martingale property).

After updating to a particular μ and before her move, the receiver can acquire additional costly information about ω, further rationally updating her beliefs from μ to a posterior belief $\gamma \in \Delta(\Omega)$. A receiver's strategy consists of an information strategy,

[6] An extreme point of a convex set S is a point in S which does not lie in any open line segment joining two points of S. Hence, extreme points of non-learning region $[0, \underline{\mu}]$ are $\{0, \underline{\mu}\}$ and extreme points of non-learning region $[\overline{\mu}, 1]$ are $\{\overline{\mu}, 1\}$.

which is a choice of a distribution $\phi \in \Delta(\Delta(\Omega))$ over the (further updated) posterior beliefs s.t. $\mathbb{E}_\phi[\gamma] = \mu$ (martingale property), and an action strategy $\sigma : \Delta(\Omega) \to A$, where $\sigma(\gamma)$ indicates the choice of action at a posterior belief γ. Let S be the set of all action strategies. We focus on sender-preferred subgame perfect equilibria: in case of indifference, the receiver uses a strategy that is (weakly) preferred by the sender[7].

The sender bears no information costs and derives utility $v(a, \omega)$. The *value* of the sender's strategy is the equilibrium expectation of $v(a, \omega)$ under that strategy profile. The *sender benefits from persuasion* if there exists τ whose value is strictly larger than the equilibrium expectation of $v(a, \omega)$ under no sender information, defined as τ_0 with $\text{supp}(\tau_0) = \mu_0$.

The receiver derives gross utility $u(a, \omega)$, where the term 'gross' indicates that information costs are not included. As is standard in RI literature, we assume Shannon-entropy based cost. For a random variable T with finite support distributed according to $\mu \in \Delta(\text{supp}(T))$, the Shannon entropy is given by

$$H(T|\mu) = -\sum_{\theta \in \text{supp}(T)} \mu(\theta) \ln \mu(\theta), \tag{3}$$

which is a measure of uncertainty about T (where $0 \log 0 = 0$ by convention). We assume the cost is proportional to the conditional mutual information[8]

$$I(\phi, \omega|\mu) = H(\omega|\mu) - \mathbb{E}_\phi[H(\omega|\gamma)] \tag{4}$$

between a receiver's information strategy ϕ and the state ω. Given μ, it captures how much uncertainty about ω is expected to be reduced by ϕ. The receiver solves the following problem.

Definition 1. Given interim belief μ, the *receiver's rational inattention problem*

[7]See Appendix A for a sufficient assumption for a unique optimal receiver's strategy.

[8]For more on the Shannon entropy and mutual information, see Cover and Thomas (2006).

(henceforth the receiver's RI problem) is

$$\begin{aligned} \max_{\phi\in\Delta(\Delta(\Omega)),\sigma\in S} \quad & \mathbb{E}_\phi\left[\sum_{\omega\in\Omega}\gamma(\omega)u\left(\sigma(\gamma),\omega\right)\right] - \lambda I(\phi,\omega|\mu) \\ s.t. \quad & \mathbb{E}_\phi[\gamma]=\mu, \end{aligned} \tag{5}$$

where $\lambda \geq 0$ is an information marginal cost parameter, and the expectation is over posterior beliefs γ distributed according to ϕ; $\gamma(\omega)$ denotes the probability of state ω at belief γ.

Matějka and McKay (2015) prove the existence of a solution to (5), which we denote by (ϕ^*_μ, σ^*)[9]. For a characterization of the receiver's optimal behavior, see Appendix A. Say *the receiver does not learn* at μ if $\text{supp}(\phi^*_\mu) = \mu$. Otherwise, say *the receiver learns* at μ.

Applying backward induction, we can express a sender's conditional expected utility for each μ, denoted by $\hat{v}(\mu)$, where

$$\hat{v}(\mu) := \mathbb{E}_{\phi^*_\mu}\left[\sum_{\omega\in\Omega}\gamma(\omega)v(\sigma^*(\gamma),\omega)\right], \tag{6}$$

where the expectation is over posterior beliefs γ distributed according to ϕ^*_μ. $\hat{v}(\mu)$ is the sender's expected utility at an interim belief μ already accounting for the subsequent optimal receiver's behavior at μ. The sender solves the following problem.

Definition 2. Given prior μ_0, the *sender's maximization problem* is

$$\begin{aligned} \max_{\tau\in\Delta(\Delta(\Omega))} \quad & \mathbb{E}_\tau[\hat{v}(\mu)] \\ s.t. \quad & \mathbb{E}_\tau[\mu]=\mu_0, \end{aligned} \tag{7}$$

where the expectation is over interim beliefs μ distributed according to τ.

[9]The optimal action strategy is independent of the interim belief as $\arg\max_a \mathbb{E}_\gamma u(a,\omega)$ is independent of the intermediate steps as to how one arrives at having posterior belief γ.

Say *the receiver never learns* if the optimal sender's strategy τ^* satisfies: $\forall \mu \in \text{supp}(\tau^*)$, the receiver does not learn at μ.

4 Bayesian persuasion s.t. never-learning

While $\hat{v}(\mu)$ is straightforward when the receiver has no additional learning option[10], it becomes complicated once she can learn. It requires finding the receiver's optimal behavior for an entire space of interim beliefs, $\{(\phi^*_\mu, \sigma^*)\}_{\mu \in \Delta(\Omega)}$, which is intractable already for a small state space. We provide an approach that avoids such calculations.

First, we show that the game can be solved as a standard Bayesian persuasion model subject to an additional constraint: the receiver never costly learns. Determining the interim beliefs at which the receiver does not learn is sufficient. At such beliefs, the receiver' behavior is deterministic and $\hat{v}(\mu) = \sum_{\omega \in \Omega} \mu(\omega) v(\sigma^*(\mu), \omega)$. We then further specify a finite set of these beliefs on which some optimal sender's strategy must be supported.

4.1 Never-learning constraint

When the receiver is fairly uncertain about what the right thing to do is, she first learns to refine her beliefs before acting. However, when her interim belief is precise enough, she does not learn. Let us formalize the subsets of such interim beliefs.

Definition 3. A *non-learning region of action $a \in A$* is

$$NL^a \;\; := \;\; \{\mu \in \Delta(\Omega) : \text{supp}(\phi^*_\mu) = \mu \;\wedge\; a \in \arg\max_{a' \in A} \sum_{\omega \in \Omega} \mu(\omega)[u(a', \omega)]\}. \quad (8)$$

[10] With no additional learning option, the receiver's optimal action is always deterministic at μ. In that case, $\hat{v}(\mu)$ is a piecewise-linear (upper semi-continuous) function: $\forall \mu \in \Delta(\Omega)$ we have $\hat{v}(\mu) = \sum_{\omega \in \Omega} \mu(\omega) v(\sigma^*(\mu), \omega)$.

The non-learning region of some action are all interim beliefs μ at which no learning and taking that action are optimal[11]. In the introductory example, a non-learning region of *not buy* is $\mu \in [0, \underline{\mu}]$, and that of *buy* is $\mu \in [\overline{\mu}, 1]$.

The following Never-Learning Lemma states that it is sufficient to focus on a subset of the sender's strategies, under which the receiver never learns. The game can be solved as a standard Bayesian persuasion problem subject to a *never-learning* constraint[12].

Lemma 1 (Never-Learning)**.** Let τ be a sender's information strategy of value v. Then there exists a sender's strategy τ' of value v where $\forall \mu \in \text{supp}(\tau')$: $\mu \in \cup_a NL^a$.

In a proof of Never-Learning Lemma, we use a specific feature of the receiver's optimal behavior captured in Lemma 2.

Lemma 2. The receiver wants to costly learn at most once, even if more rounds of costly learning were possible: $\forall \mu \in \Delta(\Omega)$, a receiver's optimal information strategy ϕ^*_μ satisfies: if $\gamma \in \text{supp}(\phi^*_\mu)$ then $\gamma \in \cup_{a \in A} NL^a$.

Lemma 2 follows from the set of the receiver's information strategies being unconstrained, the cost being increasing in Blackwell informativeness and invariant to intermediate stages[13]. Then, as the set of the sender's information strategies is also unconstrained, he can incorporate, at no cost, any receiver information strategy. Lemma 2 implies that doing so does not change the particular outcome of the game (a distribution of the receiver's actions conditional on the state) with respect to what outcome the original sender's information strategy induced. Thus, there is no need to solve the receiver's RI problems for an entire space of interim beliefs; finding non-learning regions is sufficient. The Never-Learning Lemma is an analogy to a revelation principle in mechanism design problems.

[11] When non-learning regions overlap, i.e. $\exists \mu \in NL^a$: $|\arg\max_{a' \in A} \sum_{\omega \in \Omega} \mu(\omega)[u(a', \omega)]| > 1$, the optimal action strategy, $\sigma^*(\mu)$, follows a sender-preferred equilibrium assumption.

[12] For any $\lambda \in \mathbb{R}$, $\hat{v}(\mu)$ differs from $\hat{v}(\mu)$ when the receiver has no additional learning option ($\lambda \to \infty$) only at μ not belonging to some non-learning region.

[13] The cost of achieving a particular distribution of posterior beliefs would be the same regardless of whether the learning occurs in one or more stages.

Figure 3: Extreme points of non-learning regions and a candidate sender's optimal strategy ($a, \omega \in \{1, 2, 3\}$, $u(a, \omega) = a$ if $a = \omega$ and 0 otherwise, $\lambda = 2$).

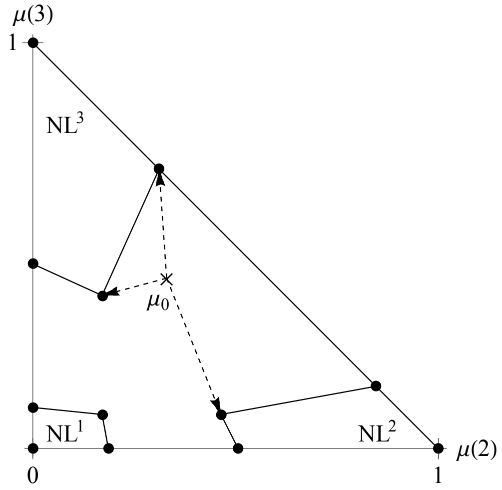

4.2 Extreme-points solution method

Recall that in the introductory example, there is a (unique) seller's strategy, under which only the extreme points of non-learning regions are induced. While the uniqueness property is not general (see Section 4.3), an optimality of some such strategy is generally guaranteed, which is captured in Prop. 1.

Let EP^a denote the *set of extreme points*[14] *of a non-learning region* NL^a.

Proposition 1. The set $\cup_{a \in A} EP^a$ is non-empty and finite. Furthermore, whenever a sender's problem (7) has a solution, there exists an optimal sender's strategy τ^*, for which $|\text{supp}(\tau^*)| \leq |\Omega|$ and $\forall \mu \in \text{supp}(\tau^*)$: $\mu \in \cup_{a \in A} EP^a$.

Note that a sender-preferred equilibrium assumption implies an upper semi-continuity of $\hat{v}$, which guarantees the existence of the equilibrium. Prop. 1 then says that we can solve the game by comparing values of a finite number of the sender's strategies. These strategies induce (at most $|\Omega|$ of) extreme points of non-learning regions. See Fig. 3 for an illustration of one such candidate sender's strategy. Once $v(a, \omega)$ is

[14] An extreme point of a convex set B is a point in B which does not lie in any open line segment joining two points of B.

specified, we can determine the optimal one. The proof of Prop. 1 shows that for any sender strategy τ under which the receiver never learns, there exists a sender's strategy τ', inducing only the extreme points, that has weakly higher value (where the sender-preferred assumption is used for a limit case of $\lambda \to \infty$). Carathéodory Theorem is then used in restricting the size of the support of an optimal sender's strategy.

The following lemmas characterize the set $\cup_a EP^a$ and the values of the candidate sender's strategies from Prop. 1. First, Lemma 3 characterizes the non-learning regions as being either a closed convex set determined by a finite series of linear inequalities (where the receiver's primitives—$u(a,\omega)$, λ—are the parameters) or an empty set. The linear conditions result from taking Shannon entropy as a measure of uncertainty in a *posterior-separable* cost function, see Section 6. The conditions follow from eq. (14) in Appendix A, where the receiver's optimal behavior is characterized.

Lemma 3. For any $\lambda \geq 0$ we have $\cup_{a\in A} NL^a \neq \emptyset$. Furthermore,

$$NL^a = \left\{ \mu \in \Delta(\Omega) : \sum_{\omega\in\Omega} \mu(\omega) \left(\frac{e^{\frac{u(a',\omega)}{\lambda}}}{e^{\frac{u(a,\omega)}{\lambda}}} \right) \leq 1 \quad \forall a' \neq a \right\} \quad \forall a \in A. \tag{9}$$

Whenever $NL^a \neq \emptyset$, Lemma 3 implies that NL^a has (finitely many) extreme points (Krein-Milman Theorem). Lemma 4 states that an extreme point of a non-learning region is a belief in NL^a for which $|\Omega|$ of constraints from Lemma 3 are binding.

Lemma 4. An extreme point of NL^a is $\mu \in \mathbb{R}^{|\Omega|}$ where $\sum_\omega \mu(\omega) = 1$ that satisfies (i) $\forall \omega \in \Omega$: $\mu(\omega) \geq 0$ and $\mu(\omega) \leq 1$; and (ii) $\sum_{\omega\in\Omega} \mu(\omega) \left(\frac{e^{\frac{u(a',\omega)}{\lambda}}}{e^{\frac{u(a,\omega)}{\lambda}}} \right) \leq 1 \quad \forall a' \neq a$, of which $|\Omega| - 1$ affine independent constraints are binding.

Lemma 5 determines the value $\hat{v}(\mu)$ when μ is an extreme point of some non-learning region. When an extreme point belongs to more non-learning regions, the sender-preferred equilibrium assumption applies.

Lemma 5. Let $\mu \in EP^a$ and $a = \sigma^*(\mu)$. Then $\hat{v}(\mu) = \sum_{\omega\in\Omega} \mu(\omega) v(a,\omega)$.

Hence, to find an optimal sender's strategy, it suffices to:
(i) determine $\cup_{a\in A} EP^a$ (using Lemma 4);
(ii) evaluate $\hat{v}(\mu)$ at those beliefs (using Lemma 5);
(iii) compare the values of the sender's strategies that satisfy Bayes' law and induce at most $|\Omega|$ beliefs from the set $\cup_{a\in A} EP^a$.

Note that Lemmas 3, 4, 5 and Prop. 1 are applicable to KG, the case of a receiver with no additional learning option, by taking $\lambda \to \infty$.

4.3 Equilibrium with learning?

In a setting with binary action and state spaces, a setting used in a number of recent papers[15], the receiver never costly learns in an equilibrium as long as the sender benefits from persuasion, see Prop. 2. In a more general setting, however, this does not necessarily hold in the case of multiple equilibria (even when the sender benefits from persuasion), see Example 1. However, if we further assumed that the sender incurs a strictly positive cost whenever the receiver learns (e.g., waiting cost), equilibria with additional learning would disappear.

Let us first slightly restrict the sender's preferences, a necessary and sufficient condition for a unique equilibrium in a binary action and state spaces case. We rule out pathological cases that can lead to situations in which two actions a and a' are both induced (under no learning) by a sender's optimal strategy, but at the belief at which the receiver takes a, the sender is exactly indifferent between a and a'.

Assumption 1. There exists no action $a \in A$ s.t. (i) $\forall \mu \in \Delta(\Omega) : \hat{v}(\mu) \leq \sum_{\omega\in\Omega} \mu(\omega) v(a,\omega)$, and (ii) $\exists \mu \in NL^{a'}$ where $a' \neq a$ and $\hat{v}(\mu) = \sum_{\omega\in\Omega} \mu(\omega) v(a,\omega)$.

Proposition 2. Suppose $|A| = |\Omega| = 2$ and the sender benefits from persuasion. Then, (i) the receiver never learns in any equilibrium; and (ii) A1 holds if and only if there exists unique equilibrium.

[15] Standard Bayesian persuasion was applied to bank regulation (Gick and Pausch, 2012), electoral manipulation (Gehlbach and Simpser, 2015), investment decision (Bizzotto, Rüdiger and Vigier, 2015), and forecasting of disasters (Aoyagi, 2014).

Figure 4: (a) Equilibrium with no learning; (b) Equilibrium with learning: the receiver learns at μ'^2 optimally inducing posteriors $\{\gamma^1, \gamma^2\}$.

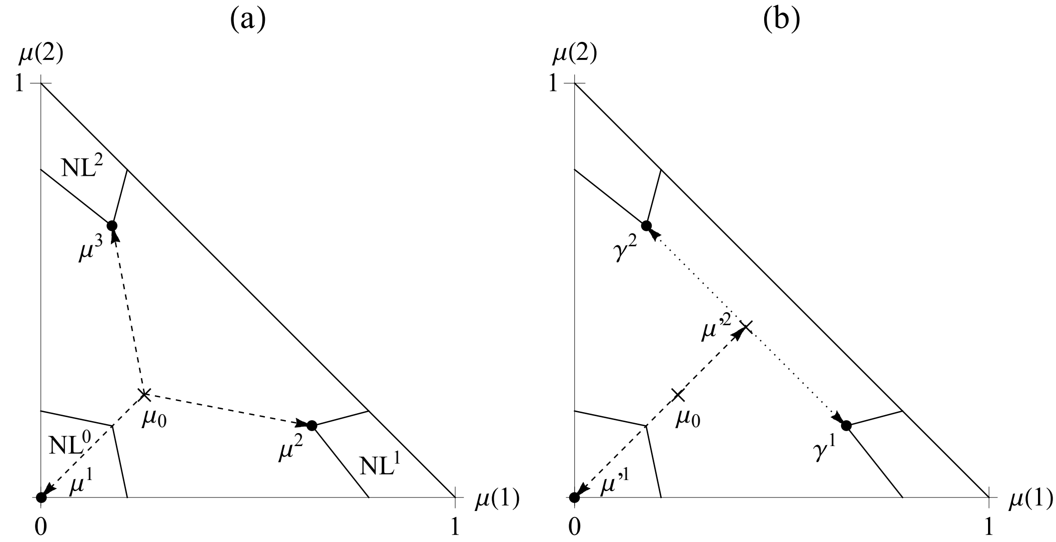

Example 1. *Let* $A = \Omega = \{0, 1, 2\}$*;* $u(a, \omega) = 1$ *if* $a = \omega$ *and* $u(a, \omega) = 0$ *otherwise;* $v(a, \omega) = 1$ *if* $a \neq 0, a \neq \omega$*, and* $v(a, \omega) = 0$ *otherwise;* $\lambda = 0.75$*; prior belief:* $\{\mu_0(0) = 0.5, \mu_0(1) = \mu_0(2) = 0.25\}$.

Fig. 4 depicts two optimal sender strategies for Example 1. In part a), an optimal sender's strategy τ^*: $\text{supp}(\tau^*) = \{\mu^1, \mu^2, \mu^3\}$ from Prop. 1 is shown. Under this strategy, the receiver never learns. In part b), a different sender's strategy τ'^*: $\text{supp}(\tau'^*) = \{\mu'^1, \mu'^2\}$ is considered. At μ'^2, the receiver learns and optimally induces posteriors $\{\gamma^1, \gamma^2\} = \{\mu^2, \mu^3\}$ with appropriate probabilities. Here, it is no longer true that the receiver never learns. As the outcome of the game (distribution of the receiver's strategies conditional on the state) is the same under both τ^* and τ'^*, τ'^* is also optimal. Note that the assumption A1 is satisfied in this example and hence A1 generally is not a sufficient assumption for uniqueness of equilibria.

4.4 Characterization of the sender's optimal strategies

Let us provide further characterization of candidate strategies from Prop. 1. Recall that in the introductory example, the optimal sender's strategy has $\text{supp}(\tau^*) = \{0, \overline{\mu}\}$; the buyer does not buy when $\mu = 0$ and buys when $\mu = \overline{\mu}$. Note two properties. First, whenever the buyer chooses the *least-preferred action* (*not to buy*), *she is certain of the state*, $\mu = 0$; she never rejects a good match. Second, whenever the buyer chooses an *action that is not the seller's least-preferred* (*to buy*), her beliefs are at a *border of a non-learning region*[16]. If her belief was inside a non-learning region when she buys ($\mu > \overline{\mu}$), the seller can increase the probability of buying by slightly decreasing μ.

The first of these properties holds in general. Say an action $\underline{a}$ is a *worst action* if $v(\underline{a}, \omega) < v(a, \omega)$ for all $a \neq \underline{a}$ and ω. Let $EP^{\Delta(\Omega)}$ denote the set of extreme points of the probability simplex $\Delta(\Omega)$.

Proposition 3. If an optimal sender's strategy from Prop. 1 induces a belief $\mu \in NL^{\underline{a}}$, where $\underline{a}$ is a worst action, and $\underline{a} = \sigma^*(\mu)$, then $\mu \in EP^{\Delta(\Omega)}$.

Prop. 3 states that whenever the receiver takes a worst action in an equilibrium (under the sender's strategy from Prop. 1), the state is fully revealed[17]. In the introductory example, the action *not to buy* is a worst action. When the receiver takes it, she is certain that the state is *bad*.

The second property holds under restriction on sender's preference. When A1 is not satisfied, the sender can be exactly indifferent to a change in the probability mass between two different actions that are both induced (under no learning) by an optimal sender's strategy. This may break the second property.

Let $\text{bd}(B)$ denote a boundary of the set B.

[16]With binary state, a border coincides with an extreme point

[17]If the receiver takes a worst action in an equilibrium under a different strategy than that of Prop. 1, there is no uncertainty left in the sense that the probability of states for which $\underline{a}$ is not optimal to choose is zero. Prop. 3 is an analogy of Prop. 4 in KG.

Proposition 4. Let A1 hold and suppose the sender benefits from persuasion. If an optimal sender's strategy induces $\mu \in NL^a$, then $\mu \in \text{bd}(NL^a)$.

Note that NL^a has a piecewise-linear boundary. Consider two different extreme points of NL^a from different 'linear segments' of its boundary. Prop. 4 implies that, under A1, both beliefs cannot be optimally induced (otherwise an optimal sender's strategy inducing an interim belief from the interior of NL^a also exists)[18].

5 Comparative statics

In this section, we examine the relationship between agents' expected equilibrium utilities and the receiver's information cost parameter λ. We show that the access to information has a disciplinary effect on the sender (decreases his expected equilibrium utility), but it is not necessarily beneficial for the receiver either.

5.1 Sender

Proposition 5. The sender's expected equilibrium utility (weakly) *increases* in λ.

As the receiver's access to her own information represents an additional *never-learning* constraint for the sender, it can only hurt him. The non-learning regions, and hence the set of sender's strategies under which the receiver never learns, do not shrink as λ increases. The receiver's potential learning is less threatening as her information becomes more expensive.

[18] Prop. 4 is a modification of Prop. 5 in KG which states that any optimally induced interior belief leads the receiver to being indifferent between at least two different actions (given the analogy of A1 is satisfied).

5.2 Receiver

Proposition 6. Assume A1, $|A| = |\Omega| = 2$, and the sender benefits from persuasion $\forall\lambda > 0$. Then the receiver's expected equilibrium utility (weakly) *decreases* in λ.

In a binary setting, the receiver (weakly) benefits from cheaper information, see Prop. 6. In general, however, the receiver does not necessarily gain from having the threat of learning; for intermediate cost, she can prefer commitment to not having the option to learn at all. Ceteris paribus, an agent would benefit from information being cheaper. However, in a strategic setting, the opponent responds to how expensive the information of the other agent is. With conflict of interest, the sender's choice under intermediate λ can be less informative (Blackwell sense[19]) than his choice under higher λ, making the receiver strictly prefer high to intermediate cost. We illustrate this in two examples. In Example 2, the sender targets a specific consideration set of the receiver (the set of actions chosen with strictly positive probability), and in Example 3, the sender dislikes a particular set of actions.

Example 2. $\Omega = \{0, 1\}$, $A = \{l, r, s\}$, *the prior belief* $\mu_0 := \Pr[\omega = 1] = 0.1$, *and*

$$u(a,\omega) = \begin{cases} 0.9 & a = l \wedge \omega = 0 \\ 1.5 & a = r \wedge \omega = 1 \\ 0.7 & a = s \\ 0 & otherwise \end{cases}, \quad v(a,\omega) = \begin{cases} 4 & a = r \wedge \omega = 0 \\ 0.9 & a = s \\ 0 & otherwise \end{cases}.$$

Consider a receiver with three actions, two risky (l, r) and one safe (s). At her prior, she would choose l, the sender's least preferred action. The sender can design an informative experiment inducing one other action to be chosen (upon favorable realization), where the 'amount' determines which one. With enough information, a sender's most preferred action, r, is chosen upon favorable realization. With less

[19] An information strategy τ is more Blackwell-informative than τ' if and only if obtaining information via τ is preferred to information via τ' by all expected utility maximizers. Equivalently, τ is more Blackwell-informative than τ' if and only if $\text{supp}(\tau')$ lie inside the convex hull of $\text{supp}(\tau)$ (Thm 12.2.2. in Blackwell and Girshick, 1954).

Figure 5: Example 2—Manipulation of the receiver's consideration set: $\hat{v}$ over non-learning regions and an optimal sender's strategy. The sender targets: (a) actions $\{l, s\}$ (less information); (b) actions $\{l, r\}$ (more information).

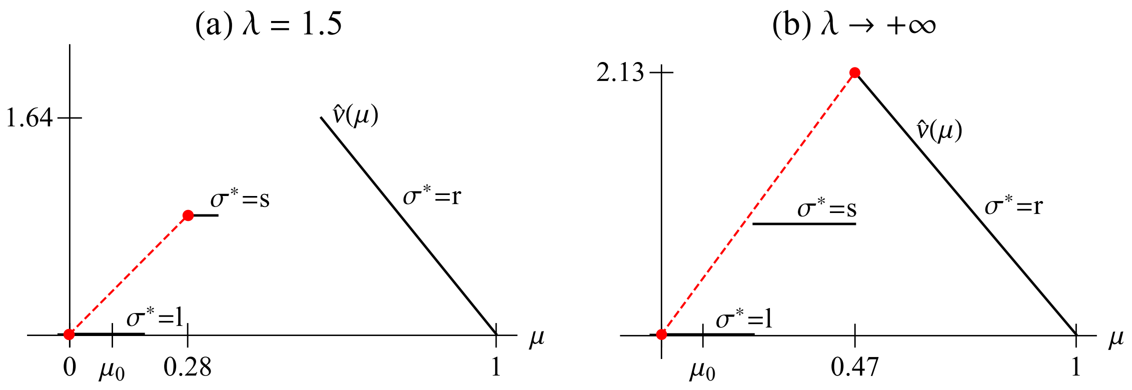

Figure 6: Example 2—Equilibrium expected utilities as a function of λ.

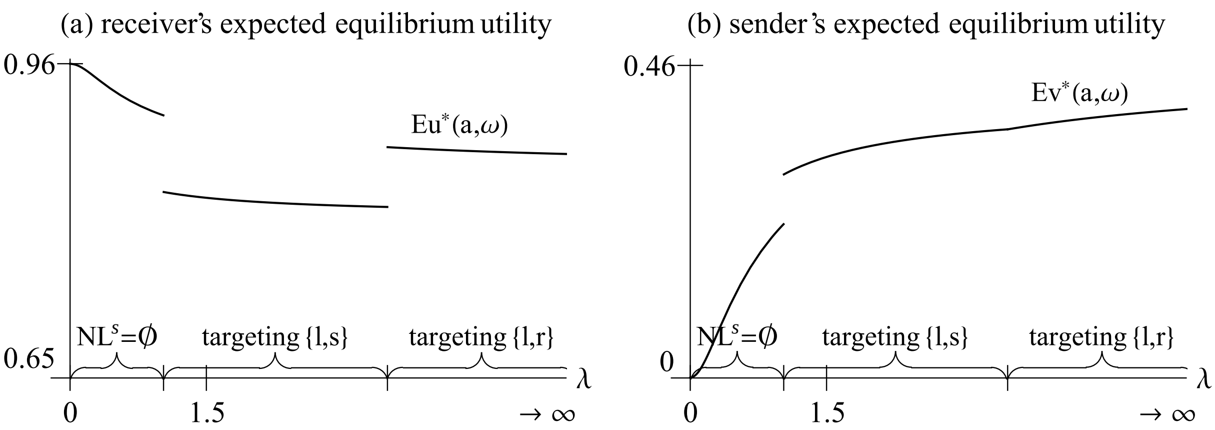

information, s is chosen upon favorable realization, but that happens with higher probability. Fig. 5 depicts the sender's optimal choice; μ_0 and μ denote the prior and interim probabilities of $\omega = 1$, respectively. When $\lambda \to \infty$, the sender targets r. However, when $\lambda = 1.5$, too much information is now needed to target r and the sender finds it optimal to give less information and to be satisfied with targeting s, but with higher probability. Fig. 6 depicts the agents' equilibrium expected utilities. For intermediate values of λ, the sender finds it optimal to target action s.

Figure 7: Example 3—$\hat{v}(\mu)$ over non-learning regions and an optimal sender's strategy. The sender provides more information when $\lambda = 1250$ than when $\lambda = 1.25$.

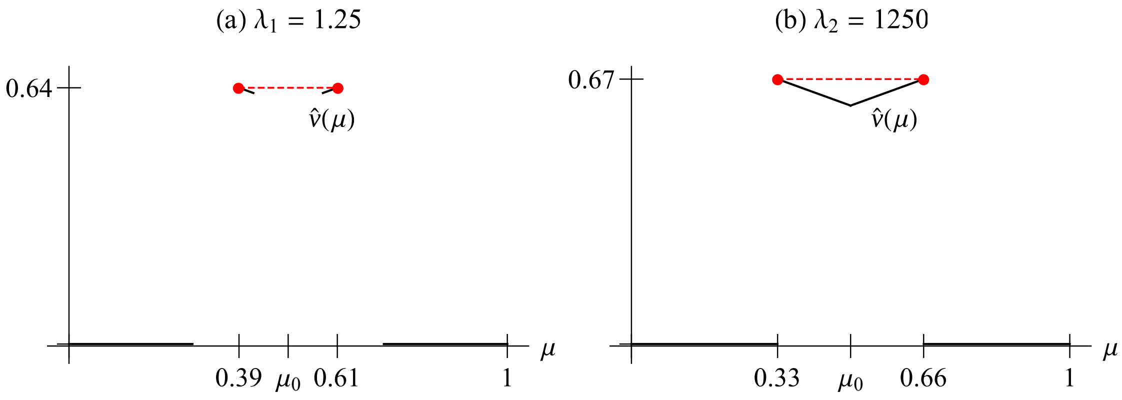

Example 3. $\Omega = \{0, 1\}$, $A = \{L, R, l, r\}$. *Let* $u(L, 0) = u(R, 1) = 1$, $u(l, 0) = u(r, 1) = 0.8$, $u(l, 1) = u(r, 0) = 0.2$, *and* $u(a, \omega) = 0$ *otherwise. Let* $v(a, \omega) = u(a, \omega)$ *if* $a \in \{l, r\}$ *and* $v(a, \omega) = 0$ *otherwise. Let the prior* $\mu_0 := \Pr[\omega = 1] = 1/2$. *Consider* $\lambda_1 = 1.25$ *and* $\lambda_2 = 1250$.

A receiver has four risky actions (L, R, l, r) and a prior belief at which she learns and possibly takes either l or r. A sender dislikes actions L, R, but cares about determining which of the two actions l, r is optimal for the receiver. He wants to give as much information as possible to distinguish which of the actions l, r is better, under the constraint that neither L nor R is chosen, happening if too much information is given. As the receiver's information becomes more expensive, this constraint is less restrictive and the sender is able to give 'more' information (Blackwell sense) than before, see Fig. 7; μ_0 and μ denote the prior and interim probabilities of $\omega = 1$, respectively.

6 Discussion of the cost function

The assumed cost is a *posterior-separable* cost (Caplin, Dean, and Leahy, 2017):

Definition 4. Given μ, a *posterior-separable* cost function is

$$c(\phi; \omega|\mu) = F(\omega|\mu) - \mathbb{E}_\phi[F(\omega|\gamma)] \tag{10}$$

where $F : \Delta(\Omega) \longrightarrow \mathbb{R}_+$ is a concave function and the expectation is over posteriors γ induced by ϕ.

Lemma 2 and the Never-Learning Lemma hold under any posterior-separable cost because of the following properties. The cost is invariant to intermediate stages: the cost of achieving a particular distribution of posterior beliefs is the same if the learning occurs in one or more stages. The cost also does not impose any restriction on the set of feasible information strategies. The marginal cost of any receiver's information strategy is independent of the interim belief, i.e., of the starting point of the receiver's RI problem, and it is increasing in Blackwell informativeness. These properties guarantee Lemma 2, which is the core of the Never-Learning Lemma. Prop. 1 applies as well when the finiteness of $\cup_a EP^a$ is omitted from the statement. Lemmas 3 and 4 are specific for Shannon entropy, $F(\cdot) = \lambda H(\cdot)$.

The Never-Learning Lemma does not hold for all possible cost functions. For instance, a receiver choosing a precision of a normally distributed signal at some cost can wish to costly learn once more upon some signal realizations. However, in such a setting, if the receiver were allowed to engage in as many learning rounds as she wanted, an analogy of the Never-Learning Lemma would be obtained.

The possibility that the receiver can be hurt by having access to better information technology is not unique to posterior-separable costs. In Appendix B, we solve the introductory example under a different cost function: by paying $c \geq 0$, the buyer obtains a binary signal $s \in \{good, bad\}$ of fixed precision $p := \Pr[s = \omega|\omega] > 0.5$. The seller exploits the restrictiveness of the buyer's set of signals, adding an additional effect in play. *First*, *the key simplification step*–the Never-Learning Lemma–*fails*. As the buyer cannot vary the precision of her signal with interim beliefs, she may wish to engage in more than one round of learning. The *seller* can take advantage of this invariance and *can strictly prefer to target buying through learning. Second, the seller's expected equilibrium utility is non-monotone in c.* Whenever he induces learning, he prefers a receiver with lower cost, because then inducing learning requires less provided information. *Third, the buyer's expected equilibrium utility is non-monotone in c.* The buyer prefers *intermediate* to low or high cost. Whenever the seller induces learning, he gives just enough information so that the buyer is indifferent between learning and not. Hence, any benefit from costly learning is exactly offset by paying cost c. With sufficiently high cost, however, the seller targets *buying* directly without the buyer's learning, who thus obtains valuable information without paying c.

7 Related literature

Our model extends KG by enabling a receiver to *endogenously* acquire her own costly information. Extensions with an exogenously privately informed receiver have already been examined in Kolotilin (2015, 2017) and Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017), and summarized in Bergemann and Morris (2016). Our result showing that the receiver's expected equilibrium utility can be non-monotone in her information cost parameter has a similar intuition as Kolotilin (2017), showing that the receiver's expected equilibrium utility can be non-monotone in the precision of her (costless) exogenous private information. Our model is more general than the above papers because we neither restrict the set of actions to be binary, as they all do, nor do we assume linear environments as Kolotilin et al. (2017) do. The notion that an agent in a strategic setting can be hurt by having access to better information technology is not unique for Bayesian persuasion, e.g., see Roesler and Szentes (2017) and Kessler (1998) for such a case in a contracting environment.

There is a rapidly growing Bayesian persuasion literature. Rayo and Segal (2010), Perez-Richet and Prady (2011), Alonso and Câmara (2014), and Hedlund (2017) explore Bayesian persuasion with a privately informed sender. Gentzkow and Kamenica (2014) model situations where the sender bears a cost associated with his information. They provide a class of cost functions (including entropy-based cost) that are compatible with the concavification approach. Similarly, we work with an entropy-based cost, but the receiver is the bearer of the cost. Other extensions include competition (Gentzkow and Kamenica, 2017a,b; Li and Norman, 2017), heterogeneous priors (Alonso and Câmara, 2016), or dynamic framework (Au, 2015; Ely, Frankel, and Kamenica, 2015; Ely, 2017).

The assumptions about the receiver's cost function fall under rational inattention (Sims, 2003). Single-agent rational inattention decision problems have been studied for investment decisions (van Nieuwerburgh and Veldkamp, 2009), rare events (Maćkowiak and Wiederholt, 2011), static stochastic choice (Caplin and Dean, 2013, 2015; Caplin, Dean, and Leahy, 2016; Denti, Mihm, de Oliveira, and Ozbek, (2016); Matějka and McKay, 2015), and dynamic stochastic choice (Steiner, Steward, and Matějka, 2017). Yang (2011), Martin (2017), and Ravid (2014) examine rational inattention in strategic situations. While the overall framework of our model is strategic, the particular rational inattention problem is essentially a single-agent decision problem since the costly information acquisition occurs at the last stage of the game. We thus follow the papers on static stochastic choice when solving a receiver's rational inattention problem.

8 Conclusion

We extend a model of Bayesian persuasion to a possibility of additional costly information acquisition by the receiver, modeled as in rational inattention. We exploit common features of Bayesian persuasion and rational inattention, resulting in a tractable model which can be used as a building block for applied problems. Based on the characterization of an optimal receiver's strategy, we offer an alternative solving algorithm characterized by a series of linear conditions. The new algorithm is also applicable to a standard Bayesian persuasion model and can simplify, sometimes dramatically, the search for an optimal sender's strategy. We further show that the receiver does not necessarily benefit from having additional sources of information and can prefer commitment to not having any.

References

Alonso, R., and Câmara, O. (2014), "On the value of persuasion by experts", mimeo.

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

Aoyagi, M. (2014), "Strategic obscurity in the forecasting of disasters", *Games and Economic Behavior* 87, 485-496.

Au, P.H. (2015), "Dynamic information disclosure", *RAND Journal of Economics* 46, 791-823.

Bergemann, D., and Morris, S. (2016), "Information design, Bayesian persuasion, and Bayes correlated equilibrium", *The American Economic Review* 106(5), 586-591.

Bizzotto, J., Rüdiger, J., and Vigier, A. (2015), "The optimal timing of persuasion", mimeo.

Blackwell, D. and Girshick, M.A. (1954), *Theory of Games And Statistical Decisions*, John Wiley and Sons, NY.

Caplin, A., and Dean, M. (2013), " Behavioral implications of rational inattention with shannon entropy", National Bureau of Economic Research Working Paper no.19318.

Caplin, A., and Dean, M. (2015), "Revealed preference, rational inattention, and costly information acquisition", *The American Economic Review* 105(7), 2183-2203.

Caplin, A., Dean, M., and Leahy, J. (2016), "Rational inattention, optimal consideration sets, and stochastic choice", NYU working paper.

Caplin, A., Dean, M., and Leahy, J. (2017), "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy", working paper.

Cover, T. M., and Thomas, J.A (2006), *Elements of information theory 2nd edition*, Wiley, Hoboken, NJ.

Denti, T., Mihm, M., de Oliveira, H., and Ozbek, K. (2016), "Rationally inattentive preferences and hidden information costs", *Theoretical Economics*.

Ely, J. C., Frankel, A., and Kamenica, E. (2015), "Suspense and surprise", *Journal of Political Economy* 123(1), 215-260.

Ely, J. C. (2017), "Beeps", *The American Economic Review* 107(1), 31-53.

Gehlbach, S., and Simpser, A. (2015), "Electoral manipulation as bureaucratic control", *American Journal of Political Science* 59(1), 212-224.

Gentzkow, M., and Kamenica, E. (2014), "Costly persuasion", *The American Economic Review* 104(5), 457-462.

Gentzkow, M., and Kamenica, E. (2017a), "Bayesian persuasion with multiple senders and rich signal spaces", *Games and Economic Behavior* 104, 411-429.

Gentzkow, M., and Kamenica, E. (2017b), "Competition in persuasion", *Review of Economic Studies* 84, 300-322.

Gick, W., and Pausch, T. (2012), "Persuasion by Stress Testing: Optimal Disclosure of Supervisory Information in the Banking Sector", mimeo.

Hedlund, J. (2017), "Bayesian persuasion by a privately informed sender", *Journal of Economic Theory* 167, 229-268.

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

Kessler, A. (1998), "The value of ignorance", *RAND Journal of Economics* 29, 339-354.

Kolotilin, A. (2015), "Experimental design to persuade", *Games and Economic Behavior* 90, 215-226.

Kolotilin, A. (2017), "Optimal information disclosure: A linear programming approach", *Theoretical Economics*, forthcoming.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., and Li, M. (2017), "Persuasion of a privately informed receiver", *Econometrica*, forthcoming.

Li, F., and Norman, P. (2017), "On Bayesian persuasion with multiple senders", mimeo.

Lipnowski, E., and Mathevet, L. (2017), "Simplifying Bayesian persuasion", mimeo.

Maćkowiak and Wiederholt (2011), "Inattention to Rare Events", CEPR Discussion Papers 8626.

Martin, D. (2017), "Strategic pricing with rational inattention to quality", *Games and Economic Behavior* 104, 131-145.

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

Perez-Richet, E., and Prady, D. (2011), "Complicating to persuade", mimeo.

Ravid, D. (2014), "Bargaining with rational inattention", Working paper.

Rayo, L., and Segal, I. (2010), "Optimal information disclosure", *Journal of Political Economy* 118(5), 949-987.

Roesler, A.-K., and Szentes, B. (2017), "Buyer-optimal learning and monopoly pricing", *The American Economic Review* 107(7), 1072-2080.

Sims, C.A. (2003), "Implications of rational inattention", *Journal of Monetary Economics* 50(3), 665-690.

Steiner, J., Stewart, C., and Matějka, F. (2017), "Rational inattention dynamics: inertia and delay in decision-making", *Econometrica* 85(2), 521-553.

van Nieuwerburgh, S., and Veldkamp, L. (2009), "Information Immobility and the Home Bias Puzzle", *The Journal of Finance* 64(3), 1187-1215.

Yang (2011), "Coordination with rational inattention", mimeo.

Appendix A: Receiver's RI problem

A1: Geometric interpretation—Concavification

The receiver's RI problem (5) can be rewritten as

$$\begin{array}{cl} \max_{\phi\in\Delta(\Delta(\Omega))} & \mathbb{E}_\phi[\hat{u}(\gamma)] - \underbrace{\lambda H(\omega|\mu)}_{=const.} \\ s.t. & \mathbb{E}_\phi[\gamma] = \mu \end{array} \tag{11}$$

where $\hat{u}(\gamma) := B(\gamma) + \lambda H(\omega|\gamma)$ is a receiver's value function at posterior γ and $B(\gamma) := \mathbb{E}_\gamma[u(\sigma^*(\gamma), \omega)]$ is a receiver's expected utility at posterior γ under her optimal action strategy $\sigma^*(\gamma) \in \arg\max_{a\in A} \mathbb{E}_\gamma[u(a,\omega)]$. The problem (11) has a geometric interpretation. Let

$$U(\gamma) := \sup\{z \in \mathbb{R} | (\gamma, z) \in co(\hat{u})\} \tag{12}$$

where sup denotes supremum and $co(\hat{u})$ denotes the convex hull[20] of the graph $\hat{u}$, be the *concavification* of $\hat{u}$. U is the smallest concave function that is everywhere weakly greater than $\hat{u}$. CD showed that $U(\mu) - \lambda H(\omega|\mu)$ is the receiver's expected utility under her optimal behavior, the receiver learns at μ if and only if $U(\mu) > \hat{u}(\mu)$, and the support of the optimal information strategy, $\text{supp}(\phi^*_\mu)$, are the posterior beliefs that support the tangent hyperplane to the lower epigraph of the concavification U above μ. See Fig. 1 in Sec. 2 for this interpretation in the introductory example.

[20]A convex hull of a set X is the smallest convex set that contains X.

A2: Characterization of the solution

Convexity of the entropy-based cost function implies that strictly more informative strategies are strictly more costly than less informative such strategies. Hence optimization is inconsistent with the choice of the same action in two distinct posteriors (receiving distinct signals that lead to the same action is inefficient as information is acquired but not acted upon). This implies that given μ, for purposes of optimization, an optimal receiver's strategy (ϕ^*_μ, σ^*) can be specified as a subset of available actions $C_\mu \subseteq A$ (a consideration set) chosen with strictly positive unconditional probabilities $\mathcal{P}^a_\mu := \mathbb{E}_{\phi^*_\mu}(\Pr[\sigma^*(\gamma) = a]) > 0$ and corresponding act-specific posteriors $\gamma^a_\mu := \{\gamma \in \Delta(\Omega) : \gamma \in \text{supp}(\phi^*_\mu) \wedge \sigma^*(\gamma) = a\}$, see Matějka and McKay (2015). Caplin et al. (2016) provide characterization of the receiver's optimal strategy, which is captured in Definition 5.

Definition 5. Given interim belief μ, a *rational inattentive strategy* at μ (henceforth, RI strategy)—a solution to problem (5)—consists of tuples $\{\mathcal{P}^a_\mu\}_{a\in A}$ and $\{\gamma^a_\mu\}_{a\in C_\mu}$, where each action is chosen in at most one posterior, such that $\forall \omega \in \Omega : \mu(\omega) = \sum_{a\in A} \mathcal{P}^a_\mu \gamma^a_\mu(\omega)$, and:

1. *Invariant Likelihood Ratio Equations for Chosen Actions:* given $a, a' \in C_\mu$, and $\omega \in \Omega$,

$$\frac{\gamma^a_\mu(\omega)}{e^{\frac{u(a,\omega)}{\lambda}}} = \frac{\gamma^{a'}_\mu(\omega)}{e^{\frac{u(a',\omega)}{\lambda}}} \tag{13}$$

2. *Likelihood Ratio Inequalities for Unchosen Actions:* given $a \in C_\mu$ and $a'' \in A \setminus C_\mu$,

$$\sum_{\omega\in\Omega} \gamma^a_\mu(\omega) \left(\frac{e^{\frac{u(a'',\omega)}{\lambda}}}{e^{\frac{u(a,\omega)}{\lambda}}} \right) \leq 1. \tag{14}$$

Applying Lemma 4 to the motivating example in Section 2, we can find the (unknown) extreme points of non-learning regions $\underline{\mu}, \overline{\mu}$ by solving $(1 - \underline{\mu})e^{\frac{-1}{\lambda}} + \underline{\mu}e^{\frac{1}{\lambda}} = 1$ and $(1 - \overline{\mu})\frac{1}{e^{\frac{-1}{\lambda}}} + \overline{\mu}\frac{1}{e^{\frac{1}{\lambda}}} = 1$, respectively. Hence, $\underline{\mu} = \frac{1}{1+e^{\frac{1}{\lambda}}}$ and $\overline{\mu} = \frac{e^{\frac{1}{\lambda}}}{1+e^{\frac{1}{\lambda}}}$.

Uniqueness

Generally, the receiver's RI strategy may not be unique at all $\mu \in \Delta(\Omega)$.

Assumption 2. $\{e^{\frac{u(a,\omega)}{\lambda}}, a \in A\}$ are affine independent. That is, one cannot find scalars α_a, not all zero, such that $\sum_{a\in A} \alpha_a = 0$ and $\sum_{a\in A} \alpha_a e^{\frac{u(a,\omega)}{\lambda}} = 0$.

Matějka and McKay (2015) and CD show that A2 is a sufficient condition for uniqueness.

Lemma 6. If A2 holds, then the receiver's optimal strategy is always unique.

A2 rules out cases such as a receiver with two duplicate actions giving her the same state-dependent payoffs. It is not very restrictive: when A2 fails, it holds under a slight perturbation of $u(a, \omega)$. When A2 holds, all equilibria are sender-preferred since the receiver is never indifferent between two strategies in an equilibrium.

Appendix B: Different cost function assumption

Here, we solve the introductory example under a different cost function. We assume the receiver can obtain a partially revealing binary signal at a fixed cost $c \geq 0$.

Let $\Omega = \{0, 1\}$, $A = \{0, 1\}$, $v(a, \omega) = a$, $u(a, \omega) = 1$ if $a = \omega = 1$, $u(a, \omega) = -1$ if $a = 1$ and $\omega = 0$ and 0 otherwise. For the purpose of this part of the appendix, we identify all the beliefs with the probability of state $\omega = 1$. Let $\mu_0, \mu, \gamma \in [0, 1]$ be the probability of $\omega = 1$ at a prior, interim, and posterior belief, respectively.

Given μ, the receiver can obtain a binary signal $s \in \{0, 1\}$ of precision $p := \Pr[s = \omega|\omega] > 0.5$ by paying $c \geq 0$. Say *the receiver learns* if she pays c and gets the signal.

Receiver's maximization problem

Given μ, if the receiver learns, she updates her beliefs to a posterior $\gamma_s(\mu) := \Pr[\omega = 1|s, \mu]$ with probability $\phi_s(\mu) := \Pr[s|\mu]$, where

$$\gamma_1(\mu) = \frac{p\mu}{\phi_1}, \qquad \phi_1(\mu) = p\mu + (1-p)(1-\mu),$$
$$\gamma_0(\mu) = \frac{(1-p)\mu}{\phi_0}, \qquad \phi_0(\mu) = (1-p)\mu + p(1-\mu).$$

The receiver takes action $a = 1$ if and only if $\gamma_s \geq 1/2$. Her expected utility from learning is

$$U^L(\mu) = \max\{0, 2\gamma_1(\mu) - 1\}\phi_1(\mu) + \max\{0, 2\gamma_0(\mu) - 1\}\phi_0(\mu) - c.$$

If she does not learn, she takes action $a = 1$ if and only if $\mu \geq 1/2$, obtaining expected utility

$$U^{NL}(\mu) = \max\{0, 2\mu - 1\}.$$

For sufficiently low cost c, there are two interim beliefs at which the receiver is indifferent between learning and not: $\underline{\mu} < 1/2$ such that upon $s = 1$, the receiver switches to action $a = 1$, but the expected marginal benefit is exactly c: $U^{NL}(\underline{\mu}|\underline{\mu} < 1/2) = U^L(\underline{\mu}|\underline{\mu} < 1/2, \gamma_1(\underline{\mu}) \geq 1/2)$; and $\overline{\mu} \geq 1/2$ such that upon $s = 0$, the receiver switches to action $a = 0$, but the expected marginal benefit is exactly c: $U^{NL}(\overline{\mu}|\overline{\mu} \geq 1/2) = U^L(\overline{\mu}|\overline{\mu} > 1/2, \gamma_0(\overline{\mu}) < 1/2)$. The first equation is $0 = \left(2\gamma_1(\underline{\mu}) - 1\right)\phi_1(\underline{\mu}) - c$ and the second equation is $2\overline{\mu} - 1 = (2\gamma_1(\overline{\mu}) - 1)\,\phi_1(\overline{\mu}) - c$, yielding

$$\underline{\mu} = 1 - p + c, \quad \overline{\mu} = p - c,$$

where $p - c \geq 1/2$ must hold.

Hence, if $c \leq p - 1/2$, there are two non-learning, $[0, \underline{\mu})$, $[\overline{\mu}, 1]$, and one learning, $[\underline{\mu}, \overline{\mu})$, regions. In contrast to the original model, a non-learning region need not

Figure 8: $\hat{v}(\mu)$ and the sender's optimal strategy with $p = 0.8$. The sender targets (a) learning (less information) and (b) no learning (more information).

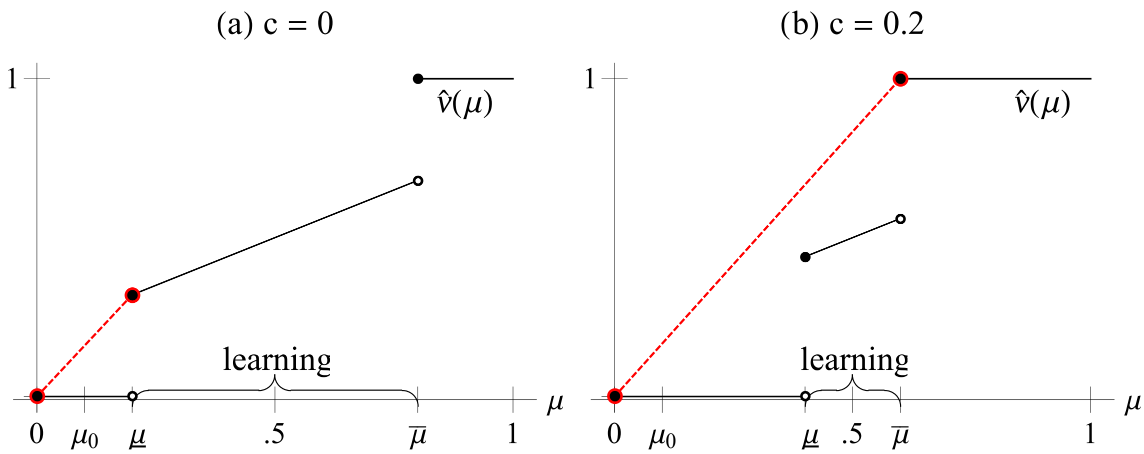

be closed, as the sender-preferred equilibrium assumption puts the belief $\underline{\mu}$ to the learning region. If $c > p - 1/2$, the receiver never learns for any μ.

Sender's maximization problem

Suppose $c \leq p - 1/2$. A seller's conditional expected utility $\hat{v}(\mu)$ is

$$\hat{v}(\mu) = \begin{cases} 0 & 0 \leq \mu < \underline{\mu} \\ p\mu + (1-p)(1-\mu) & \underline{\mu} \leq \mu < \overline{\mu} \\ 1 & \overline{\mu} \leq \mu \leq 1 \end{cases}.$$

A sender's optimal strategy can be found by concavification V of $\hat{v}$. Fig. 8 depicts $\hat{v}(\mu)$ and an optimal sender's strategy with $p = 0.8$ when $c \to 0$ and $c = 0.2$.

Comparison with the original model

First, the key simplification step–the Never-Learning Lemma–*does not hold.* The posterior γ_s can fall into a learning region. If the sender then 'sent' the receiver to γ_s directly, the receiver would learn instead of acting right away, thus changing the outcome of the game. *Second, $\hat{v}(\mu)$ is discontinuous* at the indifference points

Figure 9: Equilibrium expected utilities as a function of c with $p = 0.8$.

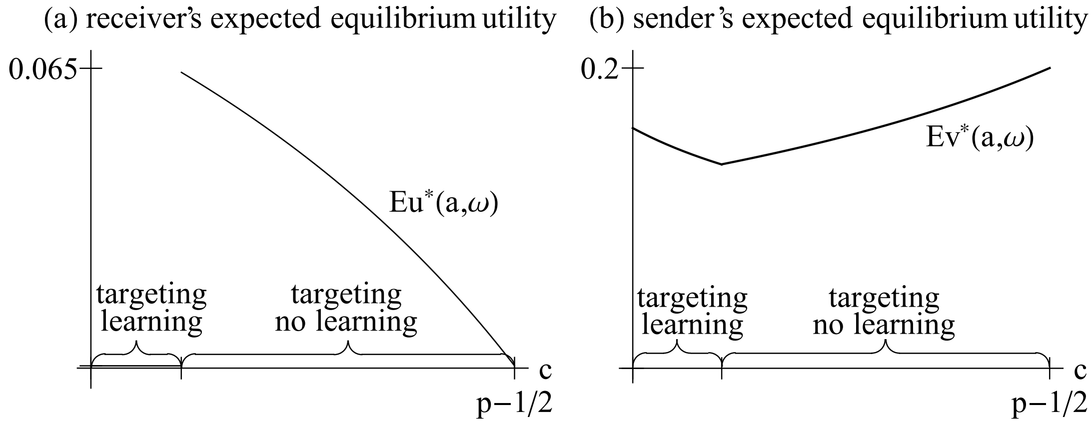

of learning/not, i.e., at $\{\underline{\mu}, \overline{\mu}\}$. In the original model, such discontinuities do not exist in the introductory example as the threshold beliefs $\{\underline{\mu}, \overline{\mu}\}$ have a different interpretation. They are points toward which the optimal amount of learning gradually shrinks, but at which it is *strictly* optimal not to learn. The sender-preferred assumption is not needed there, because the RI strategy is always unique. *Third, the sender can strictly prefer to induce* $a = 1$ *indirectly through the receiver's learning* and provide her with just enough information so that she learns ($\text{supp}(\tau^*) = \{0, \underline{\mu}\}$ in Fig. 8 (a)). Under entropy-based cost, the receiver can vary the precision of her information, but here it is fixed, which is taken advantage of. *Fourth, the sender's expected equilibrium utility* $\mathbb{E}v^*(a, \omega)$ *is non-monotone in* c: he strictly prefers low to intermediate cost, see Fig. 9 (b). When the sender targets $a = 1$ indirectly through the receiver's learning, the required amount of information to induce learning increases with c. *Fifth, the receiver strictly prefers intermediate to high and low cost.* Under low cost, the sender targets the receiver's learning. He provides just enough information so that the receiver is indifferent between learning and not; any benefit is exactly offset by paying c. With sufficiently high cost, however, the sender induces $a = 1$ directly without the receiver's learning; the receiver obtains valuable information without paying any cost. In contrast, under entropy-based cost, the receiver's expected equilibrium utility is always the highest when her information cost parameter $\lambda = 0$. As she can decide on the amount of her information, she always becomes fully informed when $\lambda = 0$, not leaving any room for sender's manipulation.

Appendix C: Proofs

Proof of Lemma 1

Proof. Let τ be the sender's strategy of value v. Suppose $\exists \mu' \in \text{supp}(\tau)$ and $\mu' \neq \cup_a NL^a$. Then the receiver's optimal information strategy at μ', $\phi^*_{\mu'}$, satisfies: $\mu' \notin \text{supp}(\phi^*_{\mu'})$, μ' lies in the convex hull of $\text{supp}(\phi^*_{\mu'})$ and $\text{supp}(\phi^*_{\mu'}) \subseteq \cup_a NL^a$ (Lemma 2). Then there exists the sender's strategy τ' where $\text{supp}(\tau') = (\text{supp}(\tau)) \setminus \mu') \cup \text{supp}(\phi^*_{\mu'})$, which does not change the distribution of the receiver's actions conditional on the state (since $\text{supp}(\phi^*_{\mu'}) \subseteq \cup_a NL^a$). Hence, the value of τ' is also v. Formally,

$$\tau'(\mu) = \begin{cases} \tau(\mu) & \mu \notin \{\mu'\} \cup \text{supp}(\phi^*_{\mu'}) \\ 0 & \mu = \mu' \\ \tau(\mu) + \tau(\mu')\phi^*_{\mu'}(\mu) & \mu \in \text{supp}(\phi^*_{\mu'}) \end{cases} .$$

□

Proof of Lemma 2

Proof. Note that the functions used here are defined in Appendix A. Let ϕ^*_μ be the receiver's optimal information strategy at μ and suppose that $\exists \gamma' \in \text{supp}(\phi^*_\mu)$ for which $\gamma' \notin \cup_a NL^a$. That is, there exists distribution $\phi_{\gamma'}$ of posterior beliefs γ,

$\mathbb{E}_{\phi_{\gamma'}}[\gamma] = \gamma'$, such that

$$\begin{aligned} \mathbb{E}_{\phi_{\gamma'}}[B(\gamma) + \lambda H(\gamma)] - \lambda H(\gamma') &> B(\gamma'), \\ \mathbb{E}_{\phi_{\gamma'}}[B(\gamma) + \lambda H(\gamma)] &> B(\gamma') + \lambda H(\gamma'), \\ \mathbb{E}_{\phi_{\gamma'}}[\hat{u}(\gamma)] &> \hat{u}(\gamma'). \end{aligned}$$

Consider a different receiver's information strategy ϕ', where

$$\phi'(\gamma) = \begin{cases} \phi^*_\mu(\gamma) & \gamma \notin \{\gamma'\} \cup \text{supp}(\phi_{\gamma'}) \\ 0 & \gamma = \gamma' \\ \phi^*_\mu(\gamma) + \phi^*_\mu(\gamma')\phi_{\gamma'}(\gamma) & \gamma \in \text{supp}(\phi_{\gamma'}) \end{cases}.$$

But then ϕ' gives the receiver strictly higher expected utility than ϕ^*_μ, contradicting the optimality of ϕ^*_μ. □

Proof of Lemma 3

Proof. For $\lambda > 0$, the receiver's marginal cost of becoming fully informed is infinity (property of Shannon entropy). Hence, for any $\lambda > 0$, the receiver never decides to become fully informed, implying $\cup_a NL^a \neq \emptyset$ (if $\lambda \to 0$, the non-learning regions are the generic interim beliefs at which the state is fully revealed). The rest of the statement follows from the equation (14) of the solution to the receiver's RI problem. □

Let us further state another Lemma that will be used throughout the following proofs.

Lemma 7. Let $V(\mu)$ be *concavification* of $\hat{v}(\mu)$ defined as the smallest concave function that is everywhere weakly greater than $\hat{v}$.

i) If $\mu \in \text{supp}(\tau^*)$, then $V(\mu) = \hat{v}(\mu)$.

ii) The sender benefits from persuasion if and only if $\hat{v}(\mu_0) < V(\mu_0)$.

Lemma 7 is an analogy of Lemma 2 and Corollary 2 from KG, which is applicable to our setting when $\hat{v}(\mu)$ modified to our setting is considered.

Proof of Proposition 1

Proof. First, let us show that the set $\cup_a EP^a$ is non-empty and finite. Lemma 3 shows $\cup_a NL^a \neq \emptyset$. Let a be an action for which $NL^a \neq \emptyset$. By the Krein-Milman Theorem and (9)—showing convexity of NL^a—NL^a is a closed convex hull of its extreme points; $EP^a \neq \emptyset$. As NL^a is an intersection of the simplex $\Delta(\Omega)$ and a collection of half-spaces $\left\{\mu \in \mathbb{R}^\Omega : \sum_{\omega\in\Omega} \mu(\omega) \left(\frac{e^{\frac{u(a',\omega)}{\lambda}}}{e^{\frac{u(a,\omega)}{\lambda}}}\right) \leq 1 \;\; \forall a' \neq a\right\}$, the set EP^a is finite. As the action space A is finite, then $\cup_a EP^a \neq \emptyset$ and is also finite.

Second, let τ^* be the sender's optimal strategy under which the receiver never learns. Suppose $\exists \mu' \in \text{supp}(\tau^*)$: $\mu' \in NL^a \setminus EP^a$. Then there exists a subset $X \subseteq EP^a$ where μ' lies in the convex hull of X. Hence, there exists another sender's strategy τ' where $\text{supp}(\tau') = (\text{supp}(\tau') \setminus \mu') \cup X$. If the chosen action at some belief of X is different from $\sigma^*(\mu')$, the action chosen at μ', it can only lead to the sender's higher expected utility by the sender-preferred assumption.

Third, let τ^* be the sender's optimal strategy and suppose $|\text{supp}(\tau^*)| > |\Omega|$. Then $\text{supp}(\tau^*)$ supports the tangent hyperplane to the lower epigraph of the concavification above prior. Such hyperplane is defined by any $|\Omega|$ different points it contains. By the Carathéodory Theorem, there exists a subset $C \subset |\text{supp}(\tau^*)|$ with $|C| \leq |\Omega|$ such that the prior belief μ_0 lies in the convex hull of C. Hence, there exists a sender's strategy τ' with $\text{supp}(\tau') = C$. As $\text{supp}(\tau')$ supports the tangent hyperplane to the lower epigraph of the concavification above prior, τ' is thus also optimal. □

Proof of Proposition 2

Proof. Let A1 hold, $A = \Omega = \{0, 1\}$, and suppose the sender benefits from persuasion. Let $\mu_0, \mu \in [0, 1]$ be the probability of $\omega = 1$ at the prior and interim belief, respectively. For any $\lambda > 0$, there are two non-learning regions with $EP^0 = \{0, \underline{\mu}\}$, $EP^1 = \{\overline{\mu}, 1\}$ where $0 < \underline{\mu} \leq \overline{\mu} < 1$. Note that $\hat{v}(\mu)$ is a piecewise-linear function, with linear segments over $[0, \underline{\mu}]$, $[\underline{\mu}, \overline{\mu}]$, and $[\overline{\mu}, 1]$. Without the loss of generality, we can consider the sender's strategies that induce at most 2 different interim beliefs.

Part i)

The concavification $V(\mu)$ of $\hat{v}(\mu)$ can have four forms:

i) $v(\mu) = V(\mu)$ if $\mu \in [\overline{\mu}, 1] \cup \{0\}$ and $v(\mu) < V(\mu)$ otherwise,

ii) $v(\mu) = V(\mu)$ if $\mu \in [0, \underline{\mu}] \cup \{1\}$ and $v(\mu) < V(\mu)$ otherwise,

iii) $v(\mu) = V(\mu)$ if $\mu \in \{0\} \cup \{1\}$ and $v(\mu) < V(\mu)$ otherwise,

iv) $v(\mu) = V(\mu)$ if $\mu \in [0, 1]$.

Note that the in iv), the sender does not benefit from persuasion for any prior $\mu_0 \in [0, 1]$. Suppose, contrary to the proposition, that there exists a sender's optimal strategy τ^* with $\tilde{\mu} \in \text{supp}(\tau^*)$ and $\tilde{\mu} \in (\underline{\mu}, \overline{\mu})$—the receiver learns at $\tilde{\mu}$. Then $\hat{v}(\tilde{\mu}) = V(\tilde{\mu})$ (Lemma 7). Since $\hat{v}(\mu)$ is linear over $[\underline{\mu}, \overline{\mu}]$, this implies that $\hat{v}(\mu) = V(\mu)$ for all $\mu \in [\underline{\mu}, \overline{\mu}]$. But then, only case iv) can happen. In particular, $\hat{v}(\mu_0) = V(\mu_0)$, which contradicts with the sender benefitting from persuasion.

Part ii)

Let us prove the *only if* part. From part i) the receiver never learns in an equilibrium. Let A1 hold and let τ^* be an optimal sender's strategy. From Prop. 4, the shape of $\hat{v}(\mu)$ and the fact that the sender benefits from persuasion, we have $\text{supp}(\tau^*) = \{\mu_l, \mu_r\} \in \{\{0, \overline{\mu}\}, \{\underline{\mu}, 1\}, \{0, 1\}\}$. Note that each pair is a point on the frontier of NL^0 and NL^1. Suppose, contrary to the proposition, there are two different optimal sender's strategies. Then there is a non-learning region of one of the actions such that one strategy induces belief on one frontier and the other strategy induces belief on the other frontier of that non-learning region. But then a new strategy that would instead, ceteris paribus, induce a convex combination of these two beliefs would also be optimal (since the convex combination still leads to the same action). However, such a new belief lies inside the non-learning region, which contradicts Prop. 4.

Let us prove the *if* part. Suppose there is a unique equilibrium, let τ^* be the optimal sender's strategy, but, contrary to the proposition, A1 does not hold. Without the loss of generality, let $a = 0$ be the action that does not satisfy A1. That is, $\forall \mu \in [0, 1]$: $\hat{v}(\mu) \leq (1 - \mu)v(0, 0) + \mu v(0, 1)$ and $\exists \mu \in [\overline{\mu}, 1]$: $\hat{v}(\mu) = (1 - \mu)v(0, 0) + \mu v(0, 1)$. Then it is either $\hat{v}(\overline{\mu}) = (1 - \overline{\mu})v(0, 0) + \overline{\mu}v(0, 1)$, $\hat{v}(1) = v(0, 1)$, or both. Since the equilibrium is unique and the sender benefits from persuasion, it must be $\text{supp}(\tau^*) = \{\mu_l, \mu_r\} \in \{\{0, \overline{\mu}\}, \{\underline{\mu}, 1\}, \{0, 1\}\}$ (from Prop. 1).

Suppose $\hat{v}(\overline{\mu}) = (1 - \overline{\mu})v(0, 0) + \overline{\mu}v(0, 1)$. Then $\hat{v}(\mu)$ is linear over the whole $[0, \overline{\mu}]$ and $\text{supp}(\tau^*) \neq \{0, \overline{\mu}\}$ (since then the sender would not benefit from persuasion from Lemma 7). Furthermore, as $\hat{v}(1) \leq v(0, 1)$, then $\text{supp}(\tau^*) \notin \{\{\underline{\mu}, 1\}, \{0, 1\}\}$, because under such strategies he cannot benefit from persuasion.

Suppose $\hat{v}(\overline{\mu}) < (1 - \overline{\mu})v(0, 0) + \overline{\mu}v(0, 1)$ and $\hat{v}(1) = v(0, 1)$. Then the values $\hat{v}(0)$, $\hat{v}(\underline{\mu})$, $\hat{v}(1)$ lie on the same line, which is thus the concavification V. Therefore, $\hat{v}(0) = V(0)$, $\hat{v}(\underline{\mu}) = V(\underline{\mu})$, $\hat{v}(1) = V(1)$ and $\hat{v}(\overline{\mu}) < V(\overline{\mu})$. Hence, $\text{supp}(\tau^*) \neq \{0, \overline{\mu}\}$ from Lemma 7. But then $\text{supp}(\tau^*) = \{\underline{\mu}, 1\}$ if and only if $\text{supp}(\tau^*) = \{0, 1\}$, contradicting the uniqueness of an equilibrium. □

Proof of Proposition 3

Proof. We can follow proof of Proposition 4 in KG applied to our setting. □

Proof of Proposition 4

Proof. Let A1 hold and suppose the sender benefits from persuasion. Proposition 5 of KG, applied to our setting, implies that for any optimal sender's strategy τ^* we have $\text{supp}(\tau^*) \cap \cup_{a \in A} \text{int}(NL^a) = \emptyset$, where $\text{int}(NL^a)$ denotes the interior of NL^a. $\square$

Proof of Proposition 5

Proof. Focusing on the sender's strategies under which the receiver never learns is sufficient (Lemma 1). We show that the non-learning regions do not shrink as λ increases. As the sender can choose from the same (or possibly even bigger) set of strategies, he never becomes strictly worse off as λ increases.

Given μ, let s be a particular receiver's strategy at μ, $EV(s)$ be a gross expected receiver's utility under s, and $I(s;\omega|\mu)$ be mutual information based on Shannon entropy associated with s. Let $\lambda \geq 0$. Suppose the receiver does not learn at μ. Let s_{NL} be the receiver's optimal non-learning strategy at μ. Let s_L be an arbitrary strategy with strictly positive learning at μ. We have

$$EV(s_{NL}) - \lambda I(s_{NL};\omega|\mu) \geq EV(s_L) - \lambda I(s_L;\omega|\mu).$$

Let $\lambda' > \lambda$. Then $\lambda I(s_{NL};\omega|\mu) = \lambda' I(s_{NL};\omega|\mu) = 0$ (no-learning costs zero) and $\lambda I(s_L;\omega|\mu) < \lambda' I(s_L;\omega|\mu)$. Hence, $EV(s_{NL}) - \lambda' I(s_{NL};\omega|\mu) > EV(s_L) - \lambda' I(s_L;\omega|\mu)$, showing that no-learning strategy remains optimal at μ. $\square$

Proof of Proposition 6

Proof. Let A1 hold, $A = \Omega = \{0,1\}$, and suppose the sender benefits from persuasion $\forall\lambda > 0$. Let $\mu_0, \mu \in [0,1]$ be the probability of $\omega = 1$ at the prior and interim belief, respectively. For $\lambda > 0$, there are two non-learning regions with $EP^0 = \{0, \underline{\mu}(\lambda)\}$, $EP^1 = \{\overline{\mu}(\lambda), 1\}$ where $0 < \underline{\mu}(\lambda) \leq \overline{\mu}(\lambda) < 1$. For each

$\lambda > 0$ there is a unique sender's optimal strategy τ^*_λ, where $\text{supp}(\tau^*_\lambda) = \{\mu_l, \mu_r\} \in \{\{0, \overline{\mu}(\lambda)\}, \{\underline{\mu}(\lambda), 1\}, \{0, 1\}\}$ (Prop.'s 1 and 2). The receiver never learns under τ^*_λ.

1. If $\text{supp}(\tau^*_\lambda) = \{0, 1\}$, the receiver obtains complete information. If there is any change in the sender's strategy as a result of an increase in λ, the receiver can only be worse off.

2. Since the non-learning regions do not shrink as λ increases (proof of Prop. 5), we have $\frac{\partial \underline{\mu}(\lambda)}{\partial \lambda} \geq 0$ and $\frac{\partial \overline{\mu}(\lambda)}{\partial \lambda} \leq 0$. Therefore, the sender's strategies inducing either always $\{0, \overline{\mu}(\lambda)\}$ or always $\{\underline{\mu}(\lambda), 1\}$ are (weakly) Blackwell less informative as λ increases. Without the loss of generality, assume $\text{supp}(\tau^*_{\lambda_1}) = \{0, \overline{\mu}(\lambda_1)\}$. Consider any $\lambda_2 > \lambda_1$. We show that inducing $\{0, \overline{\mu}(\lambda_2)\}$ remains optimal when the receiver's marginal cost of information is λ_2. Let $\hat{v}_{\lambda_i}, V_{\lambda_i}$ denote $\hat{v}$ and its concavification V, respectively, when the marginal cost of information is λ_i. We have $0 < \underline{\mu}(\lambda_1) \leq \underline{\mu}(\lambda_2) \leq \overline{\mu}(\lambda_2) \leq \overline{\mu}(\lambda_1) < 1$. From Lemma 7 and the uniqueness of the sender's strategy, it must be the case that the line connecting $\hat{v}_{\lambda_1}(0)$ and $\hat{v}_{\lambda_1}(\overline{\mu}(\lambda_1))$ is strictly above the line connecting $\hat{v}_{\lambda_1}(\underline{\mu}(\lambda_1))$ and $\hat{v}_{\lambda_1}(\overline{\mu}(\lambda_1))$. Based on the shape of $\hat{v}$ in this setting and the fact that $0 < \underline{\mu}(\lambda_2) \leq \overline{\mu}(\lambda_2)$, this then implies that the line connecting $\hat{v}_{\lambda_1}(0)$ and $\hat{v}_{\lambda_2}(\overline{\mu}(\lambda_2))$ will also be strictly above the line connecting $\hat{v}_{\lambda_2}(\underline{\mu}(\lambda_2))$ and $\hat{v}_{\lambda_2}(\overline{\mu}(\lambda_2))$. Hence $\hat{v}_{\lambda_2}(\underline{\mu}(\lambda_2)) \neq V_{\lambda_2}(\underline{\mu}(\lambda_2))$. A similar logic applies to showing that inducing $\{0, 1\}$ is not optimal either.

$\square$

Abstrakt

V této práci rozšiřujeme model Bayesovského přesvědčování o možnost, že příjemce si může za určitou cenu navíc vyhledat dodatečné informace z vlastních zdrojů, což je modelováno jako v racionální nepozornosti. Díky tomu, že využíváme společných rysů, které mezi sebou mají modely Bayesovkého přesvědčování a racionálni nepozornosti, jsme schopni poskytnout řešitelný model, jež lze použít jako základní stavební kámen pro aplikované problémy. Na základě charakterizace optimální strategie příjemce nabízíme alternativní algoritmus pro řešení daného modelu, který je charakterizován sérií lineárních podmínek. Tento nový algoritmus je aplikovatelný na standardní model Bayesovského přesvědčování a v řadě případů zjednodušuje, někdy až dramaticky, hledání jeho řešení. Dále ukazujeme, že příjemce nutně nemusí mít prospěch z možnosti vyhledat své vlastní nákladné informace, a může tak upřednostňovat závazek k nečinnosti, co se sbírání vlastních informací týče.

Working Paper Series
ISSN 1211-3298
Registration No. (Ministry of Culture): E 19443

Individual researchers, as well as the on-line and printed versions of the CERGE-EI Working Papers (including their dissemination) were supported from institutional support RVO 67985998 from Economics Institute of the CAS, v. v. i.

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

Published by
Charles University, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute of the CAS, v. v. i. (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, tel.: +420 224 005 153, Czech Republic.
Printed by CERGE-EI, Prague
Subscription: CERGE-EI homepage: http://www.cerge-ei.cz

Phone: + 420 224 005 153
Email: office@cerge-ei.cz
Web: http://www.cerge-ei.cz

Editor: Byeongju Jeong

The paper is available online at http://www.cerge-ei.cz/publications/working_papers/.

ISBN 978-80-7343-421-2 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-457-0 (Národohospodářský ústav AV ČR, v. v. i.)