

ATTENTION ELASTICITIES AND INVARIANT INFORMATION COSTS*

Dániel Csaba[†]

April, 2021

Abstract

We consider a generalization of rational inattention problems by measuring costs of information through the information radius (Sibson, 1969; Verdú, 2015) of statistical experiments. We introduce a notion of attention elasticity measuring the sensitivity of attention strategies with respect to changes in incentives. We show how the introduced class of cost functions controls attention elasticities while the Shannon model restricts attention elasticity to be unity. We explore further differences and similarities relative to the Shannon model in relation to invariance, posterior separability, consideration sets, and the ability to learn events with certainty. Lastly, we provide an efficient alternating minimization method—analogue to the Blahut-Arimoto algorithm—to obtain optimal attention strategies.

JEL CODES: D83, D90

KEYWORDS: information costs, rational inattention, α -mutual information

*I thank Andrew Caplin for helpful comments and discussions.

[†]QuantCo. Email: csaba.daniel@gmail.com

1 Introduction

The pivotal role of information has long been recognized in economic decisions. Understanding the constraints that shape the beliefs decision makers hold helps us understand and predict their behavior in different environments. The model of rational inattention (Sims, 1998, 2003) has become a widely adopted framework for the analysis of beliefs and behavior when constraints on attention and information acquisition are of primary importance.¹

A central question of the rational inattention (RI) literature has been identifying and understanding the behavioral implications of different classes of information cost functions. Our assumptions on the shape of information costs fully determine patterns in behavior under the assumption of Bayesian rationality and allowing unrestricted preferences.

In the present paper we introduce a notion of *attention elasticity*, a simple behavioral implication of information costs. Attention elasticity captures how sensitively attention strategies change with changes in incentives. One expects that increasing the payoff-difference between two actions in a given state increases the conditional odds of the better action being chosen as the decision maker pays more attention. Similarly, we expect that such a local change in payoff-differences is also reflected on the unconditional odds. The posterior odds links these quantities and is equal to the ratio of the conditional odds over the unconditional odds. Our notion of attention elasticity captures the relative rate of change in posterior odds as we change incentives.

We show that mutual information costs imply an attention elasticity of unity irrespective of the payoff structure and prior beliefs. This puts significant restrictions on the types of beliefs and choice probabilities that one can attain within the Shannon model as introduced in the seminal works of Sims (1998, 2003) and widely used in applications. We show that such patterns in attention elasticity can be understood through a perspective of the mutual information as the average divergence of signal probabilities relative to the barycenter of the Blackwell experiment. This motivates our generalization of the Shannon model.

In particular, we analyze a novel class of information cost functions that relaxes the restrictive elasticity patterns of the Shannon model while maintaining many of its desirable features. The introduced class generalizes mutual information by characterizing costs of statistical experiments based on the average Rényi divergence of conditional signal probabilities relative to the barycenter of the experiment (Sibson, 1969; Verdú, 2015), a class named α -mutual information.² We show that such a view gives rise to information costs that are analogous to constant elasticity of substitution production functions widely used in other domains of economics. This gives clear behavioral meaning to information cost parameters affecting the changes in beliefs and behavior due to changes in incentives. We refer to the corresponding class as α -rational inattention (α -RI). In general symmetric decision problems we provide conditions under which attention strategies exhibit elastic or inelastic patterns in α -RI problems.

We show that α -mutual information satisfies invariance but is not posterior separable providing a novel direction for generalizing the Shannon model. Invariance can be a desirable property in several economic contexts. For instance, in large games with incomplete information invariance of information costs is a prerequisite for the efficiency of equilibria and it prevents non-fundamental volatility as shown by the recent work of Hébert and La'O (2020) and Angeletos and Sastry (2021).

¹For an extensive overview of the rational inattention literature see Maćkowiak, Matějka, and Wiederholt (2020).

²The recent work of Hoiles, Krishnamurthy, and Pattanayak (2020) in the computer science literature has used the same class of information costs in an inverse reinforcement learning task without exploring behavioral implications.

Our approach thus complements several recent papers which have provided generalizations of the Shannon model that result in various posterior separable classes of information cost functions. [Caplin, Dean, and Leahy \(2019b\)](#) show that mutual information is uniquely characterized within the class of uniformly posterior separable cost functions by the invariance under compression property, which is a general implication of invariant cost functions. On the other hand, invariance implies a strong form of symmetry across states and its implications can be restrictive in other contexts. Neighborhood-based cost functions, such as the ones proposed by [Hébert and Woodford \(2021\)](#), and [Pomatto, Strack, and Tamuz \(2019\)](#), aim to accommodate scenarios in which the costs of learning are asymmetric across states. Posterior separable cost functions are natural in sequential learning problems as shown by [Morris and Strack \(2017\)](#), [Hébert and Woodford \(2019\)](#), and [Bloedel and Zhong \(2021\)](#) for instance.

We explore further behavioral implications. We show that under α -mutual information costs the rationally inattentive decision maker might learn payoff-relevant events with certainty, a behavior that is impossible in the Shannon model. We also characterize consideration sets and provide a generalization of the conditions corresponding to the Shannon case ([Caplin, Dean, and Leahy, 2019a](#)).

Lastly, the introduced class of rational inattention problems are tractable and can be efficiently solved through a generalized Blahut-Arimoto-type algorithm, hence have great appeal for empirical applications. We provide guarantees for convergence and discuss details of implementation.

The rest of the paper is organized as follows. [Section 2](#) introduces the notion of attention elasticity and discusses elasticity patterns under the Shannon model. [Section 3](#) introduces α -mutual information costs and presents its main properties. [Section 4](#) derives elasticity patterns in α -RI problems. [Section 5](#) discusses the invariance properties of α -mutual information and its corresponding behavioral implications. [Section 6](#) shows that under the introduced class, learning payoff-relevant events with certainty is possible. [Section 7](#) characterizes consideration sets. [Section 8](#) presents a modification of the Blahut-Arimoto algorithm providing efficient solutions for optimal attention strategies in α -RI problems. [Section 9](#) concludes.

2 Attention Elasticities and the Shannon Model

We consider general rational inattention problems with finite domains. Specifically, we consider decision problems described by a finite set of actions, $a \in \mathcal{A}$, and a finite set of states, $\omega \in \Omega$. The prior belief over states is denoted $\mu \in \Delta(\Omega)$, and the utility function is denoted by $u: \mathcal{A} \times \Omega \rightarrow \mathbb{R}$. The decision maker (DM) can acquire information and base her action on the realization of a signal from a statistical experiment. A statistical experiment is a Markov kernel, $P: \Omega \rightarrow \Delta(\mathcal{S})$, defining a collection of probability distributions over signals $s \in \mathcal{S}$ for each realization of the state. The conditional signal distributions are denoted as $P(\cdot | \omega)$ or as P_ω . A simple experiment has a finite signal support. Keeping the state space fixed, we denote the set of simple experiments as \mathcal{P} . For a given experiment, P , and prior distribution over states, μ , denote the joint distribution as $\mu P \in \Delta(\Omega \times \mathcal{S})$; the marginal distribution over signals—derived from the joint distribution μP —as $\mu P_{\mathcal{S}}$; and the joint distribution obtained from the product of two marginals, $\mu \in \Delta(\Omega)$ and $q \in \Delta(\mathcal{S})$, as $\mu \otimes q$.

We allow information costs to depend on the statistical experiment and the prior, and denote $K: \mathcal{P} \times \Delta(\Omega) \rightarrow \overline{\mathbb{R}}$. We only consider information cost functions that satisfy the convexity property such that, without loss of generality, we can identify the set of signals with the set of actions, $\mathcal{S} \equiv \mathcal{A}$, and use the two sets interchangeably.

The problem of the rationally inattentive DM is then to choose a statistical experiment that maximizes

expected utility net of information costs,³

$$\max_{P \in \mathcal{P}} \sum_{\omega, s} \mu(\omega) P(s | \omega) u(s, \omega) - K(P, \mu). \quad (2.1)$$

Following the seminal works of Sims (1998, 2003) mutual information has been the most widely used information cost in the RI literature—for a large part due to its tractability. The mutual information is defined as the Kullback-Leibler (KL) divergence between the joint distribution and the product distribution derived from the marginals,

$$I(P, \mu) := D(\mu P || \mu \otimes \mu P_S) := \sum_{\omega, s} P(s | \omega) \mu(\omega) \log \frac{P(s | \omega) \mu(\omega)}{\mu P_S(s) \mu(\omega)}. \quad (2.2)$$

We convert units of information costs (nats) to units of expected utility through a generic scalar, κ , to obtain

$$K(P, \mu) = \kappa I(P, \mu). \quad (2.3)$$

With these preliminaries in hand we introduce our notion of attention elasticity.

2.1 Attention Elasticity

The attention strategy of the DM is described by the statistical experiment, P^* , the solution to the RI problem (2.1). This specifies the conditional probability of choosing a certain action for each possible realization of the state. Loosely speaking, we expect that if stakes are higher the inattentive DM pays more attention and chooses the action that maximizes her utility with higher probability for each realization of the state.

Following this reasoning, a fundamental characteristic of attention strategies is how the conditional odds of choosing action a over b , $P^*(a | \omega) / P^*(b | \omega)$, change as we change corresponding payoff differences, $u(a, \omega) - u(b, \omega)$. Such local changes in payoff differences are also reflected in changes on the unconditional odds, $\mu P_S^*(a) / \mu P_S^*(b)$. The posterior odds induced by the optimal experiment links these quantities and is given by the ratio of conditional and unconditional odds,

$$\frac{\gamma^a(\omega)}{\gamma^b(\omega)} = \frac{P^*(a | \omega) \mu P_S^*(b)}{P^*(b | \omega) \mu P_S^*(a)}. \quad (2.4)$$

This connects to an equivalent way of modeling attention strategies through distributions over posteriors such that the posteriors satisfy the martingale property of averaging back to the prior. We capture the local sensitivity of attention strategies by analyzing the sensitivity of posterior odds with changes in incentives. We introduce our notion of attention elasticity under arbitrary changes to the payoff structure.

Definition 1 (Attention elasticity). *The attention elasticity between actions a and b in state ω is defined as the relative rate of change in posterior odds as we change incentives. Letting the posterior odds be denoted by*

$$\Gamma_\omega^{a,b} := \frac{\gamma^a(\omega)}{\gamma^b(\omega)} = \frac{P^*(a | \omega) \mu P_S^*(b)}{\mu P_S^*(a) P^*(b | \omega)}, \quad (2.5)$$

³We assume K is proper and lower-semicontinuous so the optimum is attained.

and letting θ be a parameter of the incentive structure the attention elasticity is given by,

$$\epsilon_{\omega,\theta}^{a,b} := \frac{\partial \Gamma_{\omega}^{a,b}}{\partial \theta} \frac{1}{\Gamma_{\omega}^{a,b}}. \quad (2.6)$$

We consider only a, b and ω such that conditional choice probabilities are interior, $P^*(a | \omega), P^*(b | \omega) > 0$.

The relative rate of change in the posterior odds is equal to the relative rate of change in conditional odds minus that in unconditional odds. Under certain parametrization of the incentive structure, changes in the incentives affecting all actions leave unconditional choice probabilities intact, only affecting conditional choice probabilities. In such cases our notion of attention elasticity is equivalent to the relative rate of change in conditional choice probability odds.

Given a fixed information cost function and a prior the attention strategy is dependent only on the differences in payoffs per units of information costs of the form, $(u(a, \omega) - u(b, \omega))/\kappa$. Any additive constant applied uniformly for each state-dependent payoff has no impact on the attention strategy. Correspondingly, we will consider changes in the incentive structure by changing absolute payoff differences per units of information costs across some actions. A scenario of key interest is in which we change the payoff per unit of information cost of a single action, $\theta = u(a, \omega)/\kappa$, and analyze the changes in the corresponding posterior ratio, $\Gamma_{\omega}^{a,b}$.

It is worth noting the similarities in our definition of attention elasticity and that of elasticity of substitution applied in the context of a profit-maximizing firm. The optimal production plan of the profit-maximizing firm is only dependent on relative price differences across available inputs and the price of the output. We correspondingly analyze the relative rate of change in input ratios with relative changes in price ratios. On the other hand, the attention strategy of a rationally inattentive DM is only dependent on absolute payoff differences per units of information costs. Correspondingly we analyze the relative rate of change in posterior odds with changes in absolute payoff differences.

The Shannon model puts stringent behavioral restrictions on attention elasticities. While immediate we state it as a proposition to emphasize its importance.

Proposition 1. *Under Shannon mutual information costs, for any two actions, a, b , that are chosen with strictly positive probability under the optimal attention strategy of the rationally inattentive DM, the attention elasticity in state ω with $\theta = u(a, \omega)/\kappa$ is identically equal to one.*

$$\epsilon_{\omega,\theta}^{a,b} = 1 \quad \forall \omega \in \Omega. \quad (2.7)$$

For any incentive change under the parametrization $\theta = u(c, \tilde{\omega})/\kappa$ with $(c, \tilde{\omega}) \notin \{(a, \omega), (b, \omega)\}$ the attention elasticity is zero.

Proof. This is an immediate consequence of the invariant likelihood ratio property (Caplin and Dean, 2013) of the Shannon model,

$$\frac{\gamma^a(\omega)}{\gamma^b(\omega)} = \exp\left(\frac{u(a, \omega) - u(b, \omega)}{\kappa}\right). \quad (2.8)$$

□

The restrictions on attention elasticities under the Shannon model is quite strong—the absolute value of elasticity is either one or zero. For instance, in the case where the Shannon model is used in a discrete

choice context proposition 1 has significant implications for elasticity patterns in the corresponding demand system.

We note however, that typically the scalar κ plays a dual role. First, it converts units of information costs (nats) to units of expected utility and allows us to measure all quantities in the same metric. Second, it affects the relative difficulty of acquiring information and hence should affect attention elasticities. With a redefinition of parameters we can consider costs of information where these two roles are treated separately,

$$K(P, \mu) = \kappa \sum_{\omega, s} P(s | \omega) \mu(\omega) \log \left(\frac{P(s | \omega) \mu(\omega)}{\mu P_S(s) \mu(\omega)} \right)^{\frac{1}{\sigma}} = \frac{\kappa}{\sigma} I(P, \mu). \quad (2.9)$$

It is immediate to see that under the above scaled mutual information costs the attention elasticity is constant, $\epsilon_{\omega, \theta}^{a, b} = \sigma$ under the parametrization $\theta = u(a, \omega) / \kappa$. Hence, mutual information implies constant attention elasticity and is the analogue of constant elasticity of substitution (CES) cost functions widely used in production theory. The Shannon model has the distinctive feature that the relative rate of change in posterior odds with changes in incentives is the same irrespective of the initial payoff difference between actions being \$1 or \$1M.

We can gain further insight to this remarkable elasticity property of mutual information costs through the modified logit rule (Matějka and McKay, 2015) that the interior conditional choice probabilities satisfy:

$$P^*(a | \omega) \propto \mu P_S^*(a) \exp \left(\frac{u(a, \omega)}{\kappa} \right). \quad (2.10)$$

Conditional choice probabilities are tilted relative to unconditional choice probabilities through an exponential tilting function dependent on the payoff parameters. This link between the barycenter of the optimal experiment, μP_S^* , and conditional choice probabilities is not unique to the Shannon model. In the following section we introduce a novel class of information cost functions that maintains this relationship while relaxing functional forms.

3 Information Radius and α -Mutual Information

Intuitively, without reference to a decision problem, a signal s corresponding to a statistical experiment is informative in as much as the posterior after observing s deviates from the prior. Since the posterior is proportional to the prior times the evidence, $\gamma^s(\omega) \propto \mu(\omega) P(s | \omega)$, if the likelihood is constant at s such that $P(s | \cdot) = c$ for some c , the corresponding posterior coincides with the prior and we obtain no information at s .

In fact, most characterizations of information costs can be understood as generalizing how much the likelihoods deviate from a reference constant vector. This requires defining the reference point and a measure of dispersion. In the case of mutual information the reference point is defined by the vector of unconditional signal probabilities, $\mu P_S(s) = \sum_{\omega} \mu(\omega) P(s | \omega)$, while the dispersion is captured by the relative entropy, or KL divergence:

$$I(P, \mu) = \sum_{\omega, s} P(s | \omega) \mu(\omega) \log \frac{P(s | \omega) \mu(\omega)}{\mu P_S(s) \mu(\omega)} \quad (3.1)$$

$$= \sum_{\omega} \mu(\omega) D(P_{\omega} \| \mu P_S). \quad (3.2)$$

As seen by equation 3.2 the mutual information is equal to the average KL divergence between the conditional distributions, P_ω , and the vector of unconditional signal probabilities, μP_S .

A useful way to think about the reference point is as the *barycenter* of the experiment under the KL divergence and the prior. That is, the reference point is the center of the experiment under the prior in the sense that it is the distribution that minimizes the average KL divergence:

$$\begin{aligned}\mu P_S &= \arg \min_{q \in \Delta(\mathcal{S})} \sum_{\omega} \mu(\omega) D(P_\omega \| q) \\ &= \arg \min_{q \in \Delta(\mathcal{S})} D(\mu P \| \mu \otimes q).\end{aligned}\tag{3.3}$$

Correspondingly, the mutual information is given by,

$$I(P, \mu) = \min_{q \in \Delta(\mathcal{S})} D(\mu P \| \mu \otimes q).\tag{3.4}$$

We can obtain generalizations of the mutual information by considering different divergence measures and calculating the corresponding deviation from the barycenter. The information radius or α -mutual information is a special case of the analogous definition using the Rényi divergence (Rényi, 1961) instead of the KL divergence.

Definition 2 (Rényi divergence). *For two probability distributions, P, Q , defined over the same finite sample space \mathcal{X} , the Rényi divergence of order $\alpha \in (0, 1) \cup (1, \infty)$ is defined as,*

$$D_\alpha(P \| Q) := \frac{1}{\alpha - 1} \log \left(\sum_{x \in \mathcal{X}} P(x)^\alpha Q(x)^{(1-\alpha)} \right).\tag{3.5}$$

For $\alpha > 1$, we interpret $P(x)^\alpha Q(x)^{(1-\alpha)}$ as $P(x)^\alpha / Q(x)^{(\alpha-1)}$ and adopt the conventions that $0/0 = 0$ and $y/0 = \infty$ for $y > 0$.

Definition 3 (α -mutual information). *For $\alpha \in (0, 1) \cup (1, \infty)$ the α -mutual information is defined as,*

$$\begin{aligned}I_\alpha(P, \mu) &:= \min_{q \in \Delta(\mathcal{S})} D_\alpha(\mu P \| \mu \otimes q) \\ &= \min_{q \in \Delta(\mathcal{S})} \frac{1}{\alpha - 1} \log \left(\sum_{\omega, s} \mu(\omega) P(s | \omega)^\alpha q(s)^{1-\alpha} \right).\end{aligned}\tag{3.6}$$

The solution for the minimization problem in (3.6), the barycenter of the experiment, can be expressed in closed form. Suppressing the dependence on P, μ , and α , the barycenter is given by,

$$q^*(s) = \frac{(\sum_{\omega} \mu(\omega) P(s | \omega)^\alpha)^{\frac{1}{\alpha}}}{\sum_{s'} (\sum_{\omega} \mu(\omega) P(s' | \omega)^\alpha)^{\frac{1}{\alpha}}} \quad \forall s \in \mathcal{S}.\tag{3.7}$$

Substituting the above expression for the barycenter the α -mutual information is given by,

$$I_\alpha(P, \mu) = \frac{\alpha}{\alpha - 1} \log \sum_s \left(\sum_{\omega} \mu(\omega) P(s | \omega)^\alpha \right)^{\frac{1}{\alpha}}.\tag{3.8}$$

The information radius was first introduced by Sibson (1969) as a means to generalize measures of dissimilarity among a finite set of weighted probability measures. Verdú (2015) pointed out the connection between the information radius and the generalization of mutual information through Rényi divergences coining the class α -mutual information.

3.1 Properties of α -Mutual Information

It is immediate that for any α , $I_\alpha(P, \mu) \geq 0$ for all P, μ , and is equal to zero for only non-informative experiments—for a non-informative experiment all likelihoods $P(s | \cdot)$ are constant and resulting posteriors coincide with the prior. As $\alpha \rightarrow 1$ the α -mutual information converges to the Shannon mutual information, just as the Rényi-divergence converges to the KL divergence. For a fixed prior and experiment the α -mutual information is increasing in α .⁴

Looking at the statistical experiment as a distribution over posteriors provides another potentially illuminating perspective. How much the posterior at a signal s deviates from the prior only depends on the likelihood ratios over states—how different the vector $P(s | \cdot)$ is from a constant vector. The cost of an information structure can thus be decomposed to two parts: (i) how much posteriors deviate from the prior, and (ii) what are the probabilities of ending up with different posteriors. With the unconditional signal probabilities being denoted by μP_S we can rewrite the cost function as,

$$\begin{aligned} I_\alpha(P, \mu) &= \frac{\alpha}{\alpha - 1} \log \sum_s \left(\sum_\omega \mu(\omega) P(s | \omega)^\alpha \right)^{\frac{1}{\alpha}} \\ &= \frac{\alpha}{\alpha - 1} \log \sum_s \mu P_S(s) \left(\sum_\omega \mu(\omega) \left(\frac{P(s | \omega)}{\mu P_S(s)} \right)^\alpha \right)^{\frac{1}{\alpha}}. \end{aligned} \quad (3.9)$$

Since the posterior is as $\gamma^s(\omega) \propto P(s | \omega)\mu(\omega)$, constant scaling of the likelihood $P(s | \cdot)$ for a given signal s has no impact on the posterior. Hence, from equation (3.9) we see that we price the information structure based on the posterior deviating from the prior and the probability with which we obtain the posterior under consideration. The parameter α determines how costly it is to twist the likelihoods away from the constant vector.

Applying the monotone increasing transformation $t \mapsto \frac{1}{\alpha-1} \exp(t \frac{\alpha-1}{\alpha})$ on the α -mutual information we obtain a cost function that is convex in P and that is posterior separable.

$$\frac{1}{\alpha - 1} \sum_s \left(\sum_\omega \mu(\omega) P(s | \omega)^\alpha \right)^{\frac{1}{\alpha}} = \frac{1}{\alpha - 1} \sum_s \mu P_S(s) \left(\sum_\omega \gamma^s(\omega)^\alpha \mu(\omega)^{1-\alpha} \right)^{\frac{1}{\alpha}} \quad (3.10)$$

For any interior μ the convexity of the map $\phi_\mu: \Delta(\Omega) \rightarrow \mathbb{R}$ defined by,

$$\phi_\mu(\gamma) := \frac{1}{\alpha - 1} \left(\sum_\omega \gamma(\omega)^\alpha \mu(\omega)^{1-\alpha} \right)^{\frac{1}{\alpha}}, \quad (3.11)$$

implies that the α -mutual information is monotone in the Blackwell order.⁵ On the other hand, by the results of [Denti, Marinacci, and Rustichini \(2020\)](#) we have that the α -mutual information is not an experimental cost of information as it approaches zero as the prior becomes more dogmatic. Following [Denti et al. \(2020\)](#) it is immediate to form the experimental version of the α -mutual information by considering the corresponding information cost under the uniform prior.

⁴For further details on α -mutual information see the comprehensive reviews of [Verdú \(2015\)](#) and [Ho and Verdú \(2015\)](#).

⁵For a detailed argument showing Blackwell-monotonicity see appendix [A](#).

As seen from equation (3.8) the α -mutual information function has the CES property (Arrow, Chenery, Minhas, and Solow, 1961) across states for a given signal—this, however, does not translate to an analogous constant attention elasticity property as remarked in the previous section.

Next, we analyze the class of rational inattention problems under α -mutual information costs and refer to the class as α -RI.

4 α -RI and Attention Elasticities

We take the information costs to be the scaled α -mutual information, $K(P, \mu) = \kappa I_\alpha(P, \mu)$, and pose the rational inattention problem.

$$\min_{P \in \mathcal{P}} \kappa I_\alpha(P, \mu) - \sum_{\omega, s} \mu(\omega) P(s | \omega) u(s, \omega). \quad (4.1)$$

Since the α -mutual information is quasi-convex in P (Ho and Verdú, 2015, Theorem 10) and the constraints are convex, the Karush-Kuhn-Tucker (KKT) conditions are necessary and sufficient for characterizing the solution. Denoting the dual variables corresponding to the primal feasibility conditions, $\sum_s P(s | \omega) = 1, \forall \omega \in \Omega$, by λ_ω , and the ones corresponding to the positivity constraints, $P(s | \omega) \geq 0, \forall \omega, s \in \Omega \times \mathcal{S}$, by $\delta_{\omega, s}$ the KKT conditions are given by,

$$\begin{aligned} \frac{\alpha}{\alpha - 1} \left(\underbrace{\sum_{\tilde{s}} \left(\sum_{\tilde{\omega}} P^*(\tilde{s} | \tilde{\omega})^\alpha \mu(\tilde{\omega}) \right)^{\frac{1}{\alpha}}}_{:= C^*} \right)^{-1} \left(\underbrace{\sum_{\tilde{\omega}} P^*(s | \tilde{\omega})^\alpha \mu(\tilde{\omega})}_{:= \tilde{q}^*(s)^{(1-\alpha)}} \right)^{\frac{1-\alpha}{\alpha}} \mu(\omega) P^*(s | \omega)^{(\alpha-1)} \\ - \frac{\mu(\omega) u(s, \omega)}{\kappa} + \lambda_\omega^* - \delta_{\omega, s}^* = 0, \end{aligned} \quad (4.2)$$

$$\sum_s P^*(s | \omega) = 1 \quad \forall \omega, \quad (4.3)$$

$$P^*(s | \omega) \geq 0, \quad \delta_{\omega, s}^* \geq 0, \quad P^*(s | \omega) \delta_{\omega, s}^* = 0 \quad \forall \omega, s. \quad (4.4)$$

For interior points, $P^*(s | \omega) > 0$,

$$P^*(s | \omega) = \left(C^{*\alpha} \frac{\alpha - 1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_\omega^*}{\mu(\omega)} \right) \right)^{\frac{1}{\alpha-1}} \frac{\tilde{q}^*(s)}{C^*}. \quad (4.5)$$

Note that $q^*(s) := \tilde{q}^*(s)/C^*$ is the barycenter of the optimal experiment as given by equation (3.7) and is analogous to the unconditional choice probabilities $\mu P_{\mathcal{S}}^*(s)$ in the Shannon case. Unlike in the Shannon case, we can not express the optimal dual variables λ_ω^* in closed form, however, we can efficiently obtain them numerically as we discuss in section 8 describing the modified Blahut-Arimoto algorithm.

We state the analogy with the Shannon model explicitly. At interior points, in the Shannon model,

$$P^*(a | \omega) \propto \mu P_{\mathcal{S}}^*(a) \exp \left(\frac{u(a, \omega)}{\kappa} - \frac{\lambda_\omega^*}{\mu(\omega)} \right); \quad (4.6)$$

while in the α -RI model,

$$P^*(a | \omega) \propto q^*(a) \left(\frac{\alpha - 1}{\alpha} \left(\frac{u(a, \omega)}{\kappa} - \frac{\lambda_\omega^*}{\mu(\omega)} \right) \right)^{\frac{1}{\alpha-1}}. \quad (4.7)$$

We see that in both cases the optimal conditional choice probabilities are tilted relative to the barycenter of the optimal experiment. In the Shannon model we can express the dual variables in closed form using the primal feasibility conditions,

$$\lambda_\omega^* = \mu(\omega) \log \left(\sum_{s \in \mathcal{S}} \mu P_S^*(a) \exp \left(\frac{u(s, \omega)}{\kappa} \right) \right). \quad (4.8)$$

For ease of notation we denote,

$$x^*(s, \omega) := \frac{u(s, \omega)}{\kappa} - \frac{\lambda_\omega^*}{\mu(\omega)}, \quad (4.9)$$

which determines the amount by which conditional choice probabilities are tilted relative to the barycenter given the optimal dual variables and payoff parameters.

The restrictive elasticity patterns of the Shannon model are relaxed in α -RI problems. The attention elasticities depend on the entire payoff structure and in particular on the baseline payoff difference between the considered actions. To highlight the mechanisms underlying attention elasticities in the α -RI model we analyze the class of symmetric tracking problems.

4.1 Symmetric Problem with Global Incentive Changes

We consider *general symmetric tracking problems* defined as follows. Given a finite set of states $n := |\Omega|$ consider the same number of available actions, $\mathcal{A} := \{\dots, a_\omega, \dots\}$ with $|\mathcal{A}| = n$, such that the payoff structure is given by

$$u(a, \omega) = \begin{cases} h & \text{if } a = a_\omega, \\ l & \text{otherwise,} \end{cases} \quad (4.10)$$

with $h > l$. It is without loss of generality to normalize the low payoff to zero, $l = 0$. Let the prior be uniform over the state space.

Denote $h^* := x^*(a_\omega, \omega)$ corresponding to action-state pairs that yield high payoff, and $l^* := x^*(a_\omega, \tilde{\omega})$ corresponding to pairs that yield low payoff—note that the dual variables λ_ω^* are constant across ω due to symmetry.

Arguably, one of the simplest possible change to the incentive scheme is increasing the payoff per unit of information cost for the high-payoff actions in all states simultaneously, $\theta = h/\kappa$. Since the payoff structure is symmetric and the prior is uniform, from the symmetry of the cost function we know that the optimal barycenter, q^* , and the unconditional choice probabilities over actions, μP_S^* , are both uniform over the action/signal space both before and after the change to the incentive scheme. This implies that changes in the posterior ratios are due only to changes in the conditional choice probabilities while unconditional choice probabilities are left intact. That is, in this case the attention elasticity captures the relative rate of change in conditional choice probabilities.

This global incentive change exhibits a “two-regime” elasticity pattern as stated by the next proposition.

Proposition 2. *In the general symmetric tracking problem consider any state ω and two distinct actions, a_ω, b , that are chosen with strictly positive probability conditional on ω under the optimal strategy of the α -RI problem. The attention elasticity under the incentive change $\theta = h/\kappa$ satisfies the following.*

For any $\alpha > 0$,

$$\epsilon_{\omega, \theta}^{a_\omega, b} \begin{cases} > 1 & \text{if } (\alpha - 1)l^* < 1, \\ < 1 & \text{if } (\alpha - 1)l^* > 1. \end{cases} \quad (4.11)$$

Furthermore, for any $\alpha < 1$,

$$0 < (\alpha - 1)h^* < \alpha < (\alpha - 1)l^*, \quad (4.12)$$

while for $\alpha > 1$,

$$0 < (\alpha - 1)l^* < \alpha < (\alpha - 1)h^*. \quad (4.13)$$

Proof. Appendix B.

Note that the magnitude of $(\alpha - 1)l^* = (1 - \alpha)\lambda_\omega^*/\mu(\omega)$ depends on the payoff difference between the high- and low-payoff state-action pairs, h . While we don't have a closed-form solution for h^* or l^* , by definition we know that their difference is $h^* - l^* = h$. This payoff difference and the curvature of the information cost function determined by α define whether or not the response in attention strategies is elastic or inelastic. In contrast, under the Shannon model the attention elasticity is constant irrespective of the payoff difference.

For cost parameters $\alpha < 1$ the cost of acquiring information is relatively low. When the initial payoff difference is relatively low, the attention elasticity is elastic, that is the relative rate of change in the posterior odds as we increase payoff differences is higher than that under the Shannon case. However, if the initial payoff difference is above a certain threshold the attention elasticity drops below one. We can express the threshold directly in terms of the dual variable, λ_ω^* , which is increasing in the payoff difference h .

For cost parameters $\alpha > 1$ the cost of acquiring information is relatively high and the elasticity patterns are flipped in an analogous manner. When the initial payoff difference is low, the attention elasticity is inelastic, that is the relative rate of change in the posterior odds is less than that under the Shannon model. For initial payoff difference above a certain threshold the attention elasticity is greater than 1.

Figure 1 illustrates the attention elasticity for various levels of initial payoff differences, h/κ , for different values of the cost parameter. The left-hand-side panel shows attention elasticities when acquiring information is “easy”—low values of α . The right-hand-side panel shows attention elasticities when acquiring information is “hard”.

We require the two actions to be chosen with strictly positive conditional probability to rule out the scenario under $\alpha > 1$ in which the payoff difference is large enough that the optimal attention strategy is the fully informative statistical experiment, under which posterior ratios are not well-defined. Note that this can never happen for $\alpha < 1$ as we discuss in detail in section 6.

4.2 Invariant Likelihood Ratio Property

As implied by the first-order conditions for interior points (4.5) the *invariant likelihood ratio* property (Caplin and Dean, 2013) does not hold under α -mutual information costs—except for $\alpha \rightarrow 1$ as we

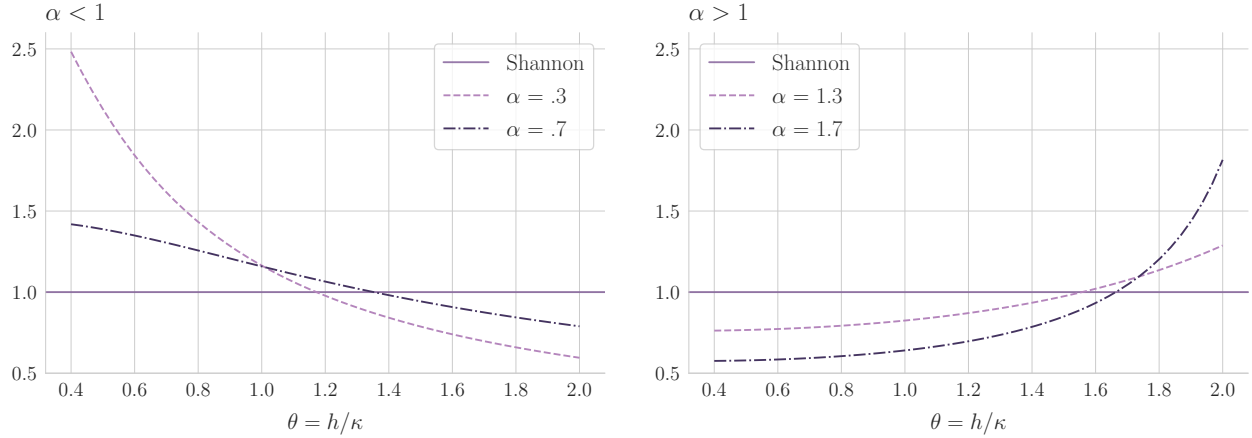


Figure 1: Attentional elasticity, $\epsilon_{\omega, \theta}^{a, b}$

fall back to the Shannon case.

For signals associated with actions a and b that are chosen with positive probability conditional on a given state under the optimal attention strategy we have,

$$\frac{P^*(a | \omega)}{P^*(b | \omega)} = \left(\frac{u(a, \omega)/\kappa - \lambda_{\omega}^*/\mu(\omega)}{u(b, \omega)/\kappa - \lambda_{\omega}^*/\mu(\omega)} \right)^{\frac{1}{\alpha-1}} \frac{q^*(a)}{q^*(b)}. \quad (4.14)$$

We see that the likelihood ratios are twisted relative to the barycenter ratios by a quantity dependent on the payoff parameters (and the optimal dual variables which also depend on the payoff parameters).

4.3 Locally Invariant Posteriors

The *locally invariant posteriors* condition of [Caplin et al. \(2019b\)](#) states that local changes in the prior—conditional on the set of supported actions not changing—do not lead to changes in the optimal posteriors. [Caplin et al. \(2019b\)](#) show that locally invariant posteriors is a characteristic feature of uniformly posterior separable cost functions.

The α -mutual information is not posterior separable nor uniformly posterior separable and hence the locally invariant posteriors property does not hold. However, as noted in equation (3.10) with a monotone increasing transformation we can transform the α -mutual information to a posterior separable cost. If we consider the problem of minimizing information costs under the constraint of obtaining a given level of expected utility this transformation should be innocuous (analogous to the expenditure minimization problem).

5 Invariant Cost Functions

Invariant information costs have important behavioral implications in various economic contexts. In recent work [Caplin et al. \(2019b\)](#) show that invariance is the defining feature of the Shannon model within the class of uniformly posterior separable cost functions. [Hébert and La'O \(2020\)](#) and [Angeletos](#)

and Sastry (2021) show that invariant cost functions play a crucial role in the efficiency of equilibria in large games with incomplete information.

While much of the literature has so far focused on generalizations of the Shannon model within posterior separable classes, the above examples demonstrate that invariance can be a desirable feature in many economic contexts. In this section we show that the α -mutual information is an invariant class, which is not posterior separable. This provides a novel direction for potential further generalization of the Shannon model.

Invariance imposes two conditions on information cost functions. The first condition—often referred to as information monotonicity—requires that coarsening the information acquired within an event should always weakly reduce information costs. The second condition requires a form of independence from the prior. The least costly way to learn about an event should be invariant to priors that only differ in their relative likelihoods over the event but not in the overall probability mass they place on the event. We follow Caplin et al. (2019b) and restate the conditions in terms of statistical experiments.

Take a subset of the state space $\bar{\Omega} \subset \Omega$ and a partition of $\bar{\Omega}$, $\{\bar{\Omega}_z\}_{z \in Z}$. Take any prior $\mu \in \text{int}\Delta(\bar{\Omega})$. Next, given experiment P we construct another experiment P' that carries the same information about the events in the partition as P but no information within any event of the partition.

Specifically, given P and μ define P' as follows:

- (i) The marginal signal distributions coincide:

$$\text{supp}\mu P'_S = \text{supp}\mu P_S \quad \text{and} \quad \mu P'_S(s) = \mu P_S(s) \quad \forall s \in \text{supp}\mu P_S. \quad (5.1)$$

- (ii) For each $z \in Z$,

- The corresponding posteriors induced by P and P' put the same mass on events in the partition:

$$\sum_{\omega \in \bar{\Omega}_z} \frac{P(s | \omega)\mu(\omega)}{\mu P_S(s)} = \sum_{\omega \in \bar{\Omega}_z} \frac{P'(s | \omega)\mu(\omega)}{\mu P'_S(s)}. \quad (5.2)$$

- Conditional on a given event in the partition the posterior induced by P' coincides with the prior. This is equivalent to the likelihood being constant over the event:

$$P'(s | \omega) = c_z \quad \forall \omega \in \bar{\Omega}_z. \quad (5.3)$$

Note that taken together equations (5.1) and (5.3), the condition of (5.2) implies that the constant likelihood has to satisfy:

$$c_z = \frac{\sum_{\omega \in \bar{\Omega}_z} P(s | \omega)\mu(\omega)}{\sum_{\omega \in \bar{\Omega}_z} \mu(\omega)}.$$

It is immediate to check that $\sum_s P'(s | \omega) = 1$ for all $\omega \in \bar{\Omega}$ and all entries are weakly positive hence P' is a valid statistical experiment.

Next take any prior $\bar{\mu}$ such that the priors put equal weight on events in the partition, i.e. $\bar{\mu}(\bar{\Omega}_z) = \mu(\bar{\Omega}_z)$ for all $z \in Z$. Define a statistical experiment \bar{P} in an analogous way as we defined P' but under $\bar{\mu}$. Specifically, given P , μ , and $\bar{\mu}$ define \bar{P} as

- (i) The marginal signal distributions coincide:

$$\text{supp}\bar{\mu}\bar{P}_S = \text{supp}\mu P_S \quad \text{and} \quad \bar{\mu}\bar{P}_S(s) = \mu P_S(s) \quad \forall s \in \text{supp}\mu P_S. \quad (5.4)$$

(ii) For each $z \in Z$,

- The corresponding posteriors induced by (P, μ) and $(\bar{P}, \bar{\mu})$ put the same mass on events in the partition:

$$\sum_{\omega \in \bar{\Omega}_z} \frac{P(s | \omega) \mu(\omega)}{\mu P_S(s)} = \sum_{\omega \in \bar{\Omega}_z} \frac{\bar{P}(s | \omega) \bar{\mu}(\omega)}{\bar{\mu} \bar{P}_S(s)} \quad (5.5)$$

- Conditional on a given event in the partition the posterior induced by $(\bar{P}, \bar{\mu})$ coincides with the prior. This is equivalent to the likelihood being constant over the event:

$$\bar{P}(s | \omega) = \bar{c}_z \quad \forall \omega \in \bar{\Omega}_z \quad (5.6)$$

Similarly, taken together equations (5.4) and (5.6), the condition of (5.5) implies that the constant likelihood has to satisfy:

$$\bar{c}_z = \frac{\sum_{\omega \in \bar{\Omega}_z} P(s | \omega) \mu(\omega)}{\sum_{\omega \in \bar{\Omega}_z} \bar{\mu}(\omega)}.$$

With the above transformations in hand we can define invariant cost functions as in [Caplin et al. \(2019b\)](#).

Definition 4 (Invariance). *A cost function K is invariant if for all finite subsets $\bar{\Omega} \subset \Omega$, all partitions of $\bar{\Omega}$, and all pairs μ and $\bar{\mu}$ that put equal mass on each event of the partition we have,*

$$K(P, \mu) \geq K(P', \mu)$$

and

$$K(P', \mu) = K(\bar{P}, \bar{\mu}).$$

Our next proposition states that the α -mutual information is an invariant cost function.

Proposition 3. *With the convention of $\alpha = 1$ denoting the Shannon mutual information, the α -mutual information cost function is invariant for any $\alpha > 0$.*

Proof. Appendix [C](#).

5.1 Invariance Under Compression

There are important behavioral implications of the cost function being invariant. If there are events for which all actions that have strictly positive probability of being chosen under the optimal policy also have identical payoffs for each state within the events, then conditional action probabilities are constant within such events under the optimal attention strategy. Formally, suppose there is an event E such that,

$$u(a, \omega_i) = u(a, \omega_j) \quad \text{for all } \omega_i, \omega_j \in E, a \in \text{supp} \mu P_S^*. \quad (5.7)$$

Clearly, any two attention strategies, P and P' , that satisfy

$$\sum_{\omega \in E} \mu(\omega) P(a | \omega) = \sum_{\omega \in E} \mu(\omega) P'(a | \omega) \quad \forall a \in \text{supp} \mu P_S^*, \quad (5.8)$$

$$P(a | \omega) = P'(a | \omega) \quad \forall \omega \notin E, \quad (5.9)$$

yield the same expected utility. By invariance of the α -mutual information, among all such attention strategies the one that has minimal cost is such that conditional signal/action probabilities are constant over E ,

$$P(a | \omega) = c_a \quad \forall \omega \in E, a \in \text{supp} \mu P_{\mathcal{S}}^*. \quad (5.10)$$

This in turn implies that posteriors over E are proportional to the prior,

$$\frac{\gamma^a(\omega_i)}{\mu(\omega_i)} = \frac{\gamma^a(\omega_j)}{\mu(\omega_j)} \quad \text{for all } \omega_i, \omega_j \in E, a \in \text{supp} \mu P_{\mathcal{S}}^*. \quad (5.11)$$

Invariance under compression states that if two decision problems have the same basic form (Caplin et al., 2019b), in which we collapse all events like E above, then corresponding attention strategies must coincide. Invariance under compression is the defining feature of the Shannon model (Caplin et al., 2019b) within the class of uniformly posterior separable cost functions. The α -mutual information is not posterior separable yet it satisfies invariance under compression.

Proposition 4. *The solution to the α -RI problem satisfies invariance under compression for any cost parameter $\alpha > 0$.*

Proof. Immediate from the fact that α -mutual information class is invariant for all $\alpha > 0$.

6 The Utility Cost Curve

In Caplin, Csaba, Leahy, and Nov (2020) we introduce the *utility cost curve* (UCC) to capture the attention costs of producing different levels of expected utility in a given decision problem. These curves give a succinct visual demonstration of the difficulty of learning in α -RI models in relation to the standard Shannon model.

Fixing the payoff structure and the prior, denote the expected utility of decisions made based on experiment P as,

$$U(P) := \sum_{\omega, a} P(a | \omega) \mu(\omega) u(a, \omega). \quad (6.1)$$

The minimum and maximum attainable expected utility levels correspond to the non-informative and perfectly informative experiments, respectively. They are given by,

$$\bar{U}^{\min} := \max_{a \in \mathcal{A}} \sum_{\omega} \mu(\omega) u(a, \omega); \quad (6.2)$$

$$\bar{U}^{\max} := \sum_{\omega} \mu(\omega) \max_{a \in \mathcal{A}} u(a, \omega). \quad (6.3)$$

For any potentially attainable expected utility level $v \in [\bar{U}^{\min}, \bar{U}^{\max}]$ the UCC in α -RI problems is defined as,

$$\bar{K}_{\alpha}(v) := \min_{P \in \mathcal{P}: U(P) \geq v} I_{\alpha}(P, \mu). \quad (6.4)$$

Consider the binary symmetric decision problem with uniform prior and payoffs as

	a	b	
ω_1	1	0	,
ω_2	0	1	

Table 1: Binary symmetric problem

which implies that the ex ante expected utility must be between $1/2$ (no information) and 1 (perfectly informative experiment).

The utility cost curve of producing expected utility $v \in [1/2, 1]$ under α -mutual information costs is

$$\bar{K}_\alpha(v) = \frac{\alpha}{\alpha - 1} \log \left(2 \left(\frac{1}{2} (v^\alpha + (1 - v)^\alpha) \right)^{\frac{1}{\alpha}} \right), \quad (6.5)$$

which we can derive from the symmetry of the optimal attention strategy.




Figure 2: The utility cost curve under different values of α

Figure 2 depicts the utility cost curve for different values of α . For each level of attainable expected utility $v \in [1/2, 1]$ the utility cost is increasing in α . In general, the cost of the perfectly informative experiment is equal to the Rényi entropy of order $1/\alpha$ of the prior, which for a uniform prior of n dimensions is equal to $\log n$ independent of α . In the binary symmetric example this implies that the perfectly informative experiments uniformly cost $\log 2$ which we see on figure 2 at $v = 1$.

Figure 2 shows that as we approach the maximum attentional output, the utility cost curves get steeper for the Shannon case and for $\alpha < 1$. In fact, it is well known that in the Shannon model the DM never chooses an attention strategy that resolves all uncertainty.

Formally, the derivative of the utility cost curve at $v \in [1/2, 1)$ is,

$$\bar{K}'_\alpha(v) = \frac{\alpha}{\alpha - 1} \frac{v^{(\alpha-1)} - (1 - v)^{(\alpha-1)}}{v^\alpha + (1 - v)^\alpha}. \quad (6.6)$$

For any $\alpha < 1$ we have

$$\lim_{v \rightarrow 1} \bar{K}'_{\alpha}(v) = \infty.$$

Furthermore, noting that as $\alpha \rightarrow 1$ we obtain the Shannon case for which,

$$\bar{K}_1(v) = v \log v + (1 - v) \log(1 - v) + \log 2, \quad \text{and} \quad \bar{K}'_1(v) = \log \frac{v}{1 - v},$$

we also have $\lim_{v \rightarrow 1} \bar{K}'_1(v) = \infty$. On the other hand, for $\alpha > 1$ we have,

$$\lim_{v \rightarrow 1} \bar{K}'_{\alpha}(v) = \frac{\alpha}{\alpha - 1}.$$

The above features of the cost function imply that in the binary example under α -mutual information costs with parameter values $\alpha \leq 1$ a rationally inattentive DM never learns the realization or non-realization of a state with certainty under the optimal attention strategy. This does not necessarily hold for cost parameters $\alpha > 1$. The implications for learning payoff-relevant events with certainty hold in general.

Proposition 5. *For any α -RI problem with bounded payoffs we have that:*

- for $\alpha \leq 1$ there exists no non-trivial event $E \subsetneq \text{supp}(\mu)$ that the DM optimally learns with certainty;
- for $\alpha > 1$, and if all payoffs are distinct, there exists a positive scalar $\pi > 0$ such that multiplying all payoffs by π the DM optimally learns all events of the form $E = \{\omega : \exists a \text{ s.t. } u(a, \omega) > u(b, \omega) \forall b \neq a\}$ with certainty.

Proof. Appendix D.

Proposition 5 states that under Shannon- or α -mutual information costs with parameter $\alpha < 1$, a rationally inattentive agent never learns to the point where some of her optimal posteriors put zero weight on states that are in the support of the prior. On the other hand, for $\alpha > 1$ we can always increase the stakes by multiplying the payoffs by a positive scalar such that in the resulting decision problem the DM's optimal posteriors only put weight on events under which a single action is optimal, i.e. for all $a \in \text{supp}(\mu P_{\mathcal{S}}^*)$ we have $\gamma^a(E_a) = 1$ where $E_a := \{\omega : u(a, \omega) > u(b, \omega) \forall b \neq a\}$.

6.1 Incentive-based Psychometric Curve

From the utility cost curve we can directly obtain the *incentive-based psychometric curve* (IPC) introduced in Caplin et al. (2020). This captures the attentional output as we scale the incentives linearly. Formally, fixing the prior and the payoff structure, denote the optimal information structure in the modified α -RI problem where all the payoffs are scaled by a positive constant π as P_{π}^* . We capture the changes in the attentional output as we vary the scaling factor π by obtaining the expected utility resulting from P_{π}^* in the original problem. The incentive-based psychometric curve captures the increase in expected utility only due to the changes in the attention strategy but not the increase in payoffs.

$$\bar{U}(\pi) := U(P_{\pi}^*) = \sum_{a, \omega} P_{\pi}^*(a | \omega) \mu(\omega) u(a, \omega). \quad (6.7)$$

As seen in figure 3 scaling the payoffs the attentional output of the DM is weakly increasing. For $\alpha \leq 1$ under no circumstances does the DM produce the maximum attentional output, $\bar{U}^{\max} = 1$. For




Figure 3: The IPC for different values of α

$\alpha > 1$, there is always a threshold such that scaling the payoffs above the threshold, the DM optimally chooses the perfectly informative experiment and the attentional output attains its maximum in line with the general claims of proposition 5.

7 Optimal Consideration Sets

Proposition 5 characterizes optimal posterior beliefs that put zero probability mass on certain events even though the events are ex ante deemed possible under the prior. The present section characterizes the conditions under which certain actions are never chosen under the optimal attention strategy even though they could be the optimal action under the realization of some states. In this sense, RI models provide a framework to study the formation of optimal consideration sets which arise endogenously dependent on the choice problem. In the Shannon model the formation of optimal consideration sets is analyzed by Caplin et al. (2019a).

Note that in the α -RI model the unconditional choice probability is strictly positive if and only if the corresponding barycenter is strictly positive,

$$\mu P_{\mathcal{S}}^*(a) > 0 \iff q^*(a) > 0. \quad (7.1)$$

Correspondingly, whether or not an action is in the optimal consideration set can be expressed in terms of the barycenter of the experiment. Denote the function $x \mapsto \max\{x, 0\} =: [x]^+$ for any $x \in \mathbb{R}$.

Proposition 6. *Under the optimal attention strategy,*

$$\frac{\sum_{\omega} \mu(\omega) \left(\left[\frac{\alpha-1}{\alpha} \left(\frac{u(a, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} \right) \right]^+ \right)^{\frac{\alpha}{\alpha-1}}}{\sum_{\tilde{s}} q^*(\tilde{s}) \sum_{\tilde{\omega}} \mu(\tilde{\omega}) \left(\left[\frac{\alpha-1}{\alpha} \left(\frac{u(\tilde{s}, \tilde{\omega})}{\kappa} - \frac{\lambda_{\tilde{\omega}}^*}{\mu(\tilde{\omega})} \right) \right]^+ \right)^{\frac{\alpha}{\alpha-1}}} \leq 1 \quad \text{for all } a \in \mathcal{A}, \quad (7.2)$$

with equality for actions that are chosen with strictly positive probability, that is for a such that $\mu P_{\mathcal{S}}^*(a) > 0$. Note that q^* is the barycenter of the optimal experiment P^* , while the dual variables λ^* are obtained from the primal feasibility constraints on P^* .

Proof. Appendix E.

As we discuss in appendix E in detail the conditions characterizing consideration sets can be obtained from the KKT conditions. The function thresholding at zero, $[\cdot]^+$, appears in the expression to handle corner-solutions under $\alpha > 1$ discussed in the previous section. For cost parameters $\alpha \leq 1$ strictly positive unconditional choice probabilities imply strictly positive conditional choice probabilities by the first claim of proposition 5.

$$\mu P_{\mathcal{S}}^*(a) > 0 \implies P^*(a | \omega) > 0 \quad \forall \omega \in \text{supp} \mu \quad (\text{for } \alpha \leq 1). \quad (7.3)$$

However, this is not necessarily the case for cost parameters $\alpha > 1$ as stated by the second claim of proposition 5, which states that the DM might learn some payoff-relevant events with certainty, necessarily contradicting the above condition (7.3).

As we show in section 8 the conditions characterizing the consideration sets are also tightly related to the updating rules in the alternating minimization algorithm that we use to obtain the solution to the α -RI problem.

7.1 Example

To illustrate the formation of consideration sets consider again the binary decision problem with symmetric payoffs as in section 6 under varying the prior probability of state ω_1 . As we increase $\mu(\omega_1)$ starting from the uniform prior, action a becomes ex ante optimal with expected utility equal to $\mu(\omega_1)$. The rationally inattentive DM chooses a with higher unconditional probability, while for some threshold value of $\mu(\omega_1)$ the DM exclusively chooses action a , that is action b is not in the consideration set.

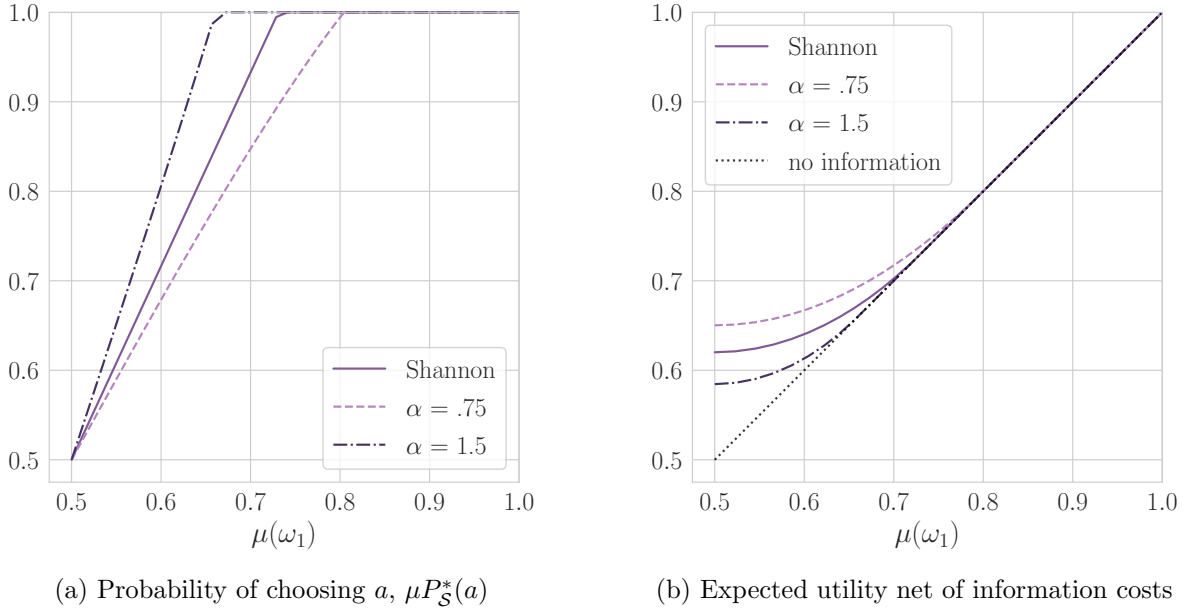


Figure 4: Optimal formation of consideration sets

Panel 4a illustrates the unconditional probability of choosing action a as the prior puts more weight on the state in which it is the better action. We see that the higher the value of α the lower the threshold value of the prior $\mu(\omega_1)$ above which the DM exclusively chooses a . That is, the costlier the attention the earlier the DM drops actions from the consideration set.

On panel 4b we see the evolution of the net expected utility the DM achieves relative to the no information benchmark. With only two actions, once the only action that is considered is a there is no information acquisition as the corresponding lines join with the dotted line on the diagonal.

The formation of optimal consideration sets in α -RI problems is a natural generalization of the Shannon case where besides beliefs and payoffs the curvature of the information cost function controlled by α determines which actions are chosen under the optimal attention strategy.

8 The Modified Blahut-Arimoto Algorithm

In practice we obtain the solution to the α -RI problem through an alternating minimization procedure analogous to the Blahut-Arimoto algorithm (Blahut, 1972; Arimoto, 1972) which is used to obtain solutions for the rational inattention problem under mutual information costs. In the present section we discuss details of implementation and provide proof of convergence.

Anticipating the alternating minimization approach we write the problem using the formulation for the cost function as in equation (3.6).

$$\min_{P \in \mathcal{P}} \min_{q \in \Delta(\mathcal{S})} \kappa D_\alpha(\mu P \| \mu \otimes q) - \sum_{\omega, s} \mu(\omega) P(s | \omega) u(s, w) \quad (8.1)$$

The Rényi-divergence is jointly quasi-convex in its two arguments⁶ and the constraint sets are both convex so the KKT conditions are necessary and sufficient for obtaining the solution.

Letting q^0 be the uniform distribution over all potential signals identified with the action set $\mathcal{S} = \mathcal{A}$ we obtain the sequence $\{(P^t, q^t)\}_{t \geq 1}$ through alternating minimization,

$$P^t \in \arg \min_{P \in \mathcal{P}} \kappa D_\alpha(\mu P \| \mu \otimes q^{t-1}) - \sum_{\omega, s} \mu(\omega) P(s | \omega) u(s, w); \quad (8.2)$$

$$q^t \in \arg \min_{q \in \Delta(\mathcal{S})} \kappa D_\alpha(\mu P^t \| \mu \otimes q) - \sum_{\omega, s} \mu(\omega) P^t(s | \omega) u(s, w). \quad (8.3)$$

The above alternating-minimization method is known as the block coordinate descent method or Gauss-Seidel method. Its convergence properties are well studied and we can verify the conditions of the general results in Tseng (2001) to prove convergence.

Proposition 7. *The sequence $\{(P^t, q^t)\}_{t \geq 1}$ generated by the alternating minimization algorithm is converging to (P^*, q^*) , the critical point of the α -RI problem.*

Proof. Appendix F.

8.1 Practical Implementation

Next, we consider details of the modified Blahut-Arimoto algorithm. The algorithm provides an efficient way for obtaining solutions to α -RI problems, making their use appealing in applications.⁷

⁶For reference see for instance (Van Erven and Harremoës, 2014, Theorem 13).

⁷An implementation of the algorithm in Python is available at https://github.com/danielcsaba/alpha_RI.

8.1.1 Minimization with respect to q

For a given information structure P^t the barycenter of the information structure minimizing the Rényi divergence of order α is given by,

$$q^t(s) = \frac{(\sum_{\omega} \mu(\omega) P^t(s | \omega)^{\alpha})^{\frac{1}{\alpha}}}{\sum_{\tilde{s}} (\sum_{\omega} \mu(\omega) P^t(\tilde{s} | \omega)^{\alpha})^{\frac{1}{\alpha}}}. \quad (8.4)$$

For $q^t(s) > 0$ the KKT conditions for the q -minimization step imply the above expression. $q^t(s) = 0$ only if $P^t(s | \omega) = 0, \forall \omega \in \Omega$, hence the expression is valid for cornersolutions as well.

Obtaining the barycenter under the Rényi divergence is computationally as efficient as under the KL divergence.

8.1.2 Minimization with respect to P

Fixing q^{t-1} the KKT conditions characterizing the minimizer P^t are as follows— λ^t and δ^t are the dual variables corresponding to the linear and positivity constraints on the constraint set.

$$\frac{\alpha}{\alpha - 1} \underbrace{\left(\sum_{\omega, s} \mu(\omega) P^t(s | \omega)^{\alpha} q^{t-1}(s)^{(1-\alpha)} \right)^{-1}}_{:= C^t} \mu(\omega) \left(\frac{P^t(s | \omega)}{q^{t-1}(s)} \right)^{(\alpha-1)} - \frac{\mu(\omega) u(s, \omega)}{\kappa} + \lambda_{\omega}^t - \delta_{\omega, s}^t = 0 \quad \forall \omega, s, \quad (8.5)$$

$$\sum_s P^t(s | \omega) = 1 \quad \forall \omega, \quad (8.6)$$

$$P^t(s | \omega) \geq 0, \quad \delta_{\omega, s}^t \geq 0, \quad P^t(s | \omega) \delta_{\omega, s}^t = 0 \quad \forall \omega, s \quad (8.7)$$

Note that if $q^{t-1}(s) = 0$ then we necessarily have $P^t(s | \omega) = 0$ for all $\omega \in \Omega$.

Suppose that $q^{t-1}(s) > 0$.

- If $\alpha < 1$ then necessarily $P^t(s | \omega) > 0$ for all $\omega \in \Omega$ in line with the first claim of proposition 5. By the complementary slackness condition $\delta_{\omega, s}^t = 0$.
- If $\alpha > 1$ and $P^t(s | \omega) = 0$ then $\lambda_{\omega}^t - \frac{\mu(\omega) u(s, \omega)}{\kappa} = \delta_{\omega, s}^t > 0$. Otherwise $\delta_{\omega, s}^t = 0$ by complementary slackness.

Taking the above arguments together we have,

$$P^t(s | \omega) = \left(C^t \left[\left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^t}{\mu(\omega)} \right) \frac{\alpha - 1}{\alpha} \right]^+ \right)^{\frac{1}{\alpha-1}} q^{t-1}(s). \quad (8.8)$$

Thresholding the inner term at zero is necessary to take care of the case under $\alpha > 1$ in which $q^{t-1}(s) > 0$ and $P^t(s | \omega) = 0$. This can only happen if the term $\left(\frac{u(s,\omega)}{\kappa} - \frac{\lambda_\omega^t}{\mu(\omega)}\right) \frac{\alpha-1}{\alpha} < 0$, so in this case the thresholding function $[\cdot]^+$ returns the correct solution for $P^t(s | \omega)$.

While the constant term C^t is a function of P^t , it is homogeneous of degree α in P^t if we define \bar{C}^t by the mapping $P \mapsto \sum_{\omega,s} \mu(\omega) P(s | \omega)^\alpha q^{t-1}(s)^{(1-\alpha)}$. We can use this fact to factor it out. Define $P_{-C^t}^t$ which is proportional to P^t as,

$$P_{-C^t}^t(s | \omega) := \left(\left[\left(\frac{u(s,\omega)}{\kappa} - \frac{\lambda_\omega^t}{\mu(\omega)} \right) \frac{\alpha-1}{\alpha} \right]^+ \right)^{\frac{1}{\alpha-1}} q^{t-1}(s). \quad (8.9)$$

Then by homogeneity of degree α of \bar{C}^t we have

$$\begin{aligned} C^t &= \bar{C}^t(P^t) = \bar{C}^t(C^t P_{-C^t}^t) = C^{t\alpha} \bar{C}^t(P_{-C^t}^t), \\ C^t &= (\bar{C}^t(P_{-C^t}^t))^{\frac{1}{1-\alpha}}. \end{aligned} \quad (8.10)$$

Finally,

$$P^t(s | \omega) = \frac{\left(\left[\left(\frac{u(s,\omega)}{\kappa} - \frac{\lambda_\omega^t}{\mu(\omega)} \right) \frac{\alpha-1}{\alpha} \right]^+ \right)^{\frac{1}{\alpha-1}}}{\sum_{\tilde{s}, \tilde{\omega}} \mu(\tilde{\omega}) \left(\left[\left(\frac{u(\tilde{s}, \tilde{\omega})}{\kappa} - \frac{\lambda_{\tilde{\omega}}^t}{\mu(\tilde{\omega})} \right) \frac{\alpha-1}{\alpha} \right]^+ \right)^{\frac{\alpha}{\alpha-1}} q^{t-1}(\tilde{s})} q^{t-1}(s). \quad (8.11)$$

It is instructive to compare this expression with proposition 6 characterizing consideration sets. Applying the formula for q^t as in (8.4),

$$q^t(s) = \frac{\left(\sum_{\omega} \mu(\omega) \left(\left[\left(\frac{u(s,\omega)}{\kappa} - \frac{\lambda_\omega^t}{\mu(\omega)} \right) \frac{\alpha-1}{\alpha} \right]^+ \right)^{\frac{\alpha}{\alpha-1}} \right)^{\frac{1}{\alpha}}}{\sum_{\tilde{s}} \left(\sum_{\tilde{\omega}} \mu(\tilde{\omega}) \left(\left[\left(\frac{u(\tilde{s}, \tilde{\omega})}{\kappa} - \frac{\lambda_{\tilde{\omega}}^t}{\mu(\tilde{\omega})} \right) \frac{\alpha-1}{\alpha} \right]^+ \right)^{\frac{\alpha}{\alpha-1}} \right)^{\frac{1}{\alpha}} q^{t-1}(\tilde{s})} q^{t-1}(s). \quad (8.12)$$

We see that either the ratio is one in which case $q^t(s) = q^{t-1}(s)$ or it is less than one in which case $q^t(s)$ converges to zero and the action corresponding to s is not in the consideration set.

8.1.3 Obtaining Dual Variables

Equation (8.11) describes the update step for the information structure given q^{t-1} . The expression depends on the dual variables λ_ω^t corresponding to the primal feasibility constraints, $\sum_s P^t(s | \omega) = 1, \forall \omega$. Using these conditions we can obtain dual variables λ^t through root-finding. To make the dependence on the dual variables explicit write P_λ^t and let the primal feasibility conditions be denoted $R_\omega^t(\lambda) := \sum_s P_\lambda^t(s | \omega) - 1$ where the information structure is defined as in (8.11) for a candidate λ .

Since we have a closed form expression for the Jacobian we can efficiently apply Newton's method.

$$\frac{\partial R_{\omega_i}^t}{\partial \lambda_{\omega_k}} = \begin{cases} (\sum_s P_{\lambda}^t(s | \omega_i)) (\sum_s P_{\lambda}^t(s | \omega_k)) & \text{for } i \neq k, \\ (\sum_s P_{\lambda}^t(s | \omega_k))^2 - \frac{\sum_s \left(\left[\left(\frac{u(s, \omega_k) - \lambda_{\omega_k}}{\mu(\omega_k)} \right)^{\alpha-1} \right]^+ \right)^{\frac{2-\alpha}{\alpha-1}} q^{t-1}(s)}{\sum_{\omega, s} \mu(\omega) \left(\left[\left(\frac{u(s, \omega) - \lambda_{\omega}}{\mu(\omega)} \right)^{\alpha-1} \right]^+ \right)^{\frac{\alpha}{\alpha-1}} q^{t-1}(s)} & \text{for } i = k. \end{cases} \quad (8.13)$$

Furthermore, we have natural bounds on the optimal dual variables depending on the value of α that we enforce in our custom root-finding algorithm. Without enforcing these bounds out-of-the-box root-finding algorithms will typically overshoot and lead to invalid or suboptimal attention strategies. Specifically, in order to satisfy the KKT conditions we need,

- for $\alpha < 1$

$$\lambda_{\omega}^t > \max_{a \in \mathcal{A}} \frac{\mu(\omega) u(a, \omega)}{\kappa} \quad \forall \omega \in \Omega; \quad (8.14)$$

- for $\alpha > 1$

$$\lambda_{\omega}^t < \max_{a \in \mathcal{A}} \frac{\mu(\omega) u(a, \omega)}{\kappa} \quad \forall \omega \in \Omega. \quad (8.15)$$

Taken together the modified Blahut-Arimoto algorithm takes the following form:

Algorithm 1: Modified Blahut-Arimoto algorithm

Input: $(\Omega, \mathcal{A}, u, \mu, \kappa, \alpha)$
Result: P^*, q^*, λ^*

- 1 $q^0 \leftarrow \text{ones}(|\mathcal{A}|/|\mathcal{A}|)$;
- 2 **while** $\text{distance} > \text{tol}$ **do**
- 3 $\lambda^t \leftarrow$ via root-finding on primal feasibility given q^{t-1} ,
- 4 obeying bounds (8.14) and (8.15);
- 5 $P^t \leftarrow$ via equation (8.11) given q^{t-1} and λ^t ;
- 6 $q^t \leftarrow$ via equation (8.4) given P^t ;
- 7 $\text{distance} \leftarrow \|q^t - q^{t-1}\|$
- 8 **end**

9 Conclusion

We consider a generalization of the rational inattention model of Sims (1998, 2003) by measuring costs of information via the information radius of statistical experiments. We show that such costs control elasticities of attention, measuring the sensitivity of attention strategies with changes in incentives. In contrast, the Shannon model restricts attention elasticities to be constant irrespective of the incentives. The resulting class of information costs is invariant but not posterior separable. This provides a novel direction for generalizations of the Shannon model. We explore further behavioral implications relative to the Shannon model and provide an efficient alternating minimization method for obtaining optimal attention strategies.

References

- George-Marios Angeletos and Karthik Sastry. Inattentive economies. Working paper, 2021.
- Suguru Arimoto. An algorithm for computing the capacity of arbitrary discrete memoryless channels. *IEEE Transactions on Information Theory*, 18(1):14–20, 1972.
- Kenneth J. Arrow, Hollis B. Chenery, Bagicha S. Minhas, and Robert M. Solow. Capital-Labor Substitution and Economic Efficiency. *The Review of Economics and Statistics*, pages 225–250, 1961.
- David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, pages 265–272, 1953.
- Richard Blahut. Computation of channel capacity and rate-distortion functions. *IEEE transactions on Information Theory*, 18(4):460–473, 1972.
- Alexander W. Bloedel and Weijie Zhong. The cost of optimally-acquired information. Working paper, 2021.
- Andrew Caplin and Mark Dean. Behavioral Implications of Rational Inattention with Shannon Entropy. Technical report, National Bureau of Economic Research, 2013.
- Andrew Caplin, Mark Dean, and John Leahy. Rational Inattention, Optimal Consideration Sets, and Stochastic Choice. *The Review of Economic Studies*, 86(3):1061–1094, 2019a.
- Andrew Caplin, Mark Dean, and John Leahy. Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy. Technical report, National Bureau of Economic Research, 2019b.
- Andrew Caplin, Dániel Csaba, John Leahy, and Oded Nov. Rational Inattention, Competitive Supply, and Psychometrics. *The Quarterly Journal of Economics*, 135:1681–1724, August 2020.
- Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. 2020.
- Benjamin M. Hébert and Jennifer La’O. Information acquisition, efficiency, and non-fundamental volatility. Technical report, National Bureau of Economic Research, 2020.
- Benjamin M. Hébert and Michael Woodford. Rational inattention when decisions take time. Technical report, National Bureau of Economic Research, 2019.
- Benjamin M. Hébert and Michael Woodford. Neighborhood-based information costs. Technical report, National Bureau of Economic Research, 2021.
- Siu-Wai Ho and Sergio Verdú. Convexity/Concavity of Rényi Entropy and α -mutual Information. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 745–749, 2015.
- William Hoiles, Vikram Krishnamurthy, and Kunal Pattanayak. Rationally inattentive inverse reinforcement learning explains youtube commenting behavior. *Journal of Machine Learning Research*, 21(170):1–39, 2020.
- Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational inattention: A review. *CEPR Discussion Papers*, (15408), 2020.
- Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.

- Stephen Morris and Philipp Strack. The Wald problem and the equivalence of sequential sampling and static information costs. *Unpublished manuscript, June, 2017*.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. The Cost of Information. *arXiv preprint arXiv:1812.04211*, 2019.
- Alfréd Rényi. On Measures of Entropy and Information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, 1961.
- Robin Sibson. Information Radius. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 14(2):149–160, 1969.
- Christopher A. Sims. Stickiness. *Carnegie-Rochester Conference Series on Public Policy*, 49:317 – 356, 1998.
- Christopher A. Sims. Implications of Rational Inattention. *Journal of Monetary Economics*, 50(3): 665–690, 2003.
- Erik Torgersen. *Comparison of Statistical Experiments*. Cambridge University Press, 1991.
- Paul Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494, 2001.
- Tim Van Erven and Peter Harremoës. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- Sergio Verdú. α -Mutual Information. In *2015 Information Theory and Applications Workshop (ITA)*, pages 1–6. IEEE, 2015.

A Monotonicity in the Blackwell order

The following Bayesian criterion is a standard characterization of the Blackwell order—for further reference see [Blackwell \(1953\)](#) and [Torgersen \(1991\)](#).

Definition 5 (Blackwell order). *For two simple statistical experiments, P, Q , defined over the same finite state space Ω , experiment P is more informative than experiment Q if for all $\mu \in \text{int}\Delta(\Omega)$ we have,*

$$\sum_s \phi(\gamma^s) \mu P_S(s) \geq \sum_s \phi(\gamma^s) \mu Q_S(s), \quad (\text{A.1})$$

for all convex functions $\phi: \Delta(\Omega) \rightarrow \mathbb{R}$.

Proposition 8. *The α -mutual information is monotone in the Blackwell order.*

Proof. As introduced in equation (3.10), applying the monotone increasing transformation $t \mapsto \frac{1}{\alpha-1} \exp\left(t \frac{\alpha-1}{\alpha}\right)$ on the α -mutual information we get,

$$\frac{1}{\alpha-1} \sum_s \mu P_S(s) \left(\sum_{\omega} \gamma^s(\omega)^\alpha \mu(\omega)^{1-\alpha} \right)^{\frac{1}{\alpha}}. \quad (\text{A.2})$$

For monotonicity in the Blackwell order we need to show that the map $\phi_\mu: \Delta(\Omega) \rightarrow \mathbb{R}$ defined by,

$$\phi_\mu(\gamma) := \frac{1}{\alpha-1} \left(\sum_{\omega} \gamma(\omega)^\alpha \mu(\omega)^{1-\alpha} \right)^{\frac{1}{\alpha}}, \quad (\text{A.3})$$

for any $\mu \in \text{int}\Delta(\Omega)$ is convex.

We show that the Hessian of ϕ_μ is positive semi-definite. Given an enumeration of the finite states denote $\gamma_i = \gamma(\omega_i)$. Given α, μ and γ , denote $c := \sum_k \gamma_k^\alpha \mu_k^{1-\alpha}$. The Hessian of ϕ_μ is,

$$H_{ij} := \frac{\partial^2 \phi_\mu(\gamma)}{\partial \gamma_i \partial \gamma_j} = \begin{cases} -c^{\frac{1}{\alpha}-2} \left(\frac{\gamma_i}{\mu_i} \right)^{\alpha-1} \left(\frac{\gamma_j}{\mu_j} \right)^{\alpha-1} & \text{if } i \neq j; \\ -c^{\frac{1}{\alpha}-2} \left(\frac{\gamma_i}{\mu_i} \right)^{2(\alpha-1)} + c^{\frac{1}{\alpha}-1} \left(\frac{\gamma_i}{\mu_i} \right)^{\alpha-1} \frac{1}{\gamma_i} & \text{if } i = j. \end{cases} \quad (\text{A.4})$$

For any $x \in \mathbb{R}^{|\Omega|} \setminus \{\mathbf{0}\}$ we have,

$$x' H x = c^{\frac{1}{\alpha}-1} \left(\sum_j \left(\frac{\gamma_j}{\mu_j} \right)^{\alpha-1} \frac{1}{\gamma_j} x_j^2 - \frac{1}{c} \left(\sum_j \left(\frac{\gamma_j}{\mu_j} \right)^{\alpha-1} x_j \right)^2 \right). \quad (\text{A.5})$$

As c is strictly positive the sign of $x' H x$ only depends on the term in the parentheses. Multiplying by c we get,

$$\left(\sum_j \mu_j^{1-\alpha} \gamma_j^\alpha \left(\frac{x_j}{\gamma_j} \right)^2 \right) \left(\sum_j \mu_j^{1-\alpha} \gamma_j^\alpha \right) - \left(\sum_j \mu_j^{1-\alpha} \gamma_j^\alpha \left(\frac{x_j}{\gamma_j} \right) \right)^2. \quad (\text{A.6})$$

For fixed μ, γ , defining the inner product on $\mathbb{R}^{|\Omega|}$ by,

$$\langle x, y \rangle := \sum_j \mu_j^{1-\alpha} \gamma_j^\alpha x_j y_j, \quad (\text{A.7})$$

we have,

$$\text{sign}(x' H x) = \text{sign} \left(\left\langle \frac{x}{\gamma}, \frac{x}{\gamma} \right\rangle \langle \mathbf{1}_{|\Omega|}, \mathbf{1}_{|\Omega|} \rangle - \left\langle \frac{x}{\gamma}, \mathbf{1}_{|\Omega|} \right\rangle^2 \right), \quad (\text{A.8})$$

with $\mathbf{1}_{|\Omega|}$ denoting the vector of ones of length $|\Omega|$. By the Cauchy-Schwarz inequality the right hand side is weakly positive, ϕ_μ is positive semi-definite. □

B Proof of Proposition 2

We first derive the formula for the attention elasticity in terms of changes in the dual variables, λ^* , then analyze the change in the dual via an application of the implicit function theorem. Taking these together we obtain the proof of proposition 2.

To ease notation—keeping the underlying α parameter fixed—we denote the tilting function,

$$\begin{aligned} f: \mathbb{R} &\rightarrow \mathbb{R} \\ x &\mapsto \left(\frac{\alpha - 1}{\alpha} x \right)^{\frac{1}{\alpha-1}}. \end{aligned} \quad (\text{B.1})$$

$f(x^*(s, \omega))$ determines the optimal tilting of the conditional signal probabilities relative to the barycenter,

$$P^*(s \mid \omega) \propto f(x^*(s, \omega)) q^*(s). \quad (\text{B.2})$$

At interior points the quantity $(\alpha - 1)x^*(s, \omega) > 0$ so the function f is well-defined. We deal with boundary points in later sections.

The posterior ratio for a fixed state under the α -mutual information is given by,

$$\Gamma_\omega^{a,b} = \frac{\gamma^a(\omega)}{\gamma^b(\omega)} = \frac{f(x^*(a, \omega)) \sum_{\tilde{\omega}} \mu(\tilde{\omega}) f(x^*(b, \tilde{\omega}))}{f(x^*(b, \omega)) \sum_{\tilde{\omega}} \mu(\tilde{\omega}) f(x^*(a, \tilde{\omega}))}. \quad (\text{B.3})$$

With these notations in hand the general formula for the attention elasticity is given by,

$$\begin{aligned} \epsilon_{\omega, \theta}^{a_\omega, b} &= \frac{f'(x^*(a_\omega, \omega)) x^{*'}(a_\omega, \omega)}{f(x^*(a_\omega, \omega))} - \frac{f'(x^*(b, \omega)) x^{*'}(b, \omega)}{f(x^*(b, \omega))} + \\ &\quad \frac{\sum_{\tilde{\omega}} \mu(\tilde{\omega}) f'(x^*(b, \tilde{\omega})) x^{*'}(b, \tilde{\omega})}{\sum_{\tilde{\omega}} \mu(\tilde{\omega}) f(x^*(b, \tilde{\omega}))} - \frac{\sum_{\tilde{\omega}} \mu(\tilde{\omega}) f'(x^*(a_\omega, \tilde{\omega})) x^{*'}(a_\omega, \tilde{\omega})}{\sum_{\tilde{\omega}} \mu(\tilde{\omega}) f(x^*(a_\omega, \tilde{\omega}))}, \end{aligned} \quad (\text{B.4})$$

with,

$$f'(x) := \frac{df}{dz} \Big|_{z=x} = \frac{1}{\alpha} \left(\frac{\alpha - 1}{\alpha} x \right)^{\frac{2-\alpha}{\alpha-1}}, \quad (\text{B.5})$$

$$x^{*'}(s, \omega) := \frac{\partial x^*}{\partial \theta} = \begin{cases} 1 - \frac{\lambda^{*'}_\omega}{\mu(\omega)} & \text{if } s = s_\omega, \\ -\frac{\lambda^{*'}_\omega}{\mu(\omega)} & \text{otherwise.} \end{cases} \quad (\text{B.6})$$

The derivative of the dual is with respect to the incentive change θ ,

$$\lambda^*{}'_\omega := \frac{\partial \lambda^*_\omega}{\partial \theta},$$

and we write $D_\theta \boldsymbol{\lambda}^*$ to denote the vector of derivatives for all states.

Note that in the general symmetric tracking problem with the global incentive change the last two term of equation B.4 representing changes in the unconditional choice probabilities cancel out—as the unconditional choice probabilities stay uniform.

Note further that,

$$\frac{f'(x)}{f(x)} = \frac{1}{(\alpha - 1)x}. \quad (\text{B.7})$$

Correspondingly we have,

$$\epsilon_{\omega, \theta}^{a_\omega, b} = \frac{\lambda^*{}'_\omega}{\mu(\omega)} \left(\frac{1}{(\alpha - 1)x^*(b, \omega)} - \frac{1}{(\alpha - 1)x^*(a_\omega, \omega)} \right) + \frac{1}{(\alpha - 1)x^*(a_\omega, \omega)}. \quad (\text{B.8})$$

We derive the change in the dual variable via an application of the implicit function theorem on the primal feasibility constraints. We analyze this under a general problem with potentially asymmetric incentive structure.

B.1 Expressing Changes in the Dual Variables

The primal feasibility constraints are given by⁸,

$$R^*_\omega := \frac{\sum_s f(x^*(s, \omega)) q^*(s)}{\left(\sum_s (\sum_{\tilde{\omega}} f(x^*(s, \tilde{\omega}))^\alpha \mu(\tilde{\omega}))^{\frac{1}{\alpha}} q^*(s) \right)^\alpha} = 1 \quad \forall \omega \in \Omega. \quad (\text{B.9})$$

To ease notation we denote the constant in the denominator,

$$\xi := \left(\sum_s \left(\sum_{\tilde{\omega}} f(x^*(s, \tilde{\omega}))^\alpha \mu(\tilde{\omega}) \right)^{\frac{1}{\alpha}} q^*(s) \right)^\alpha,$$

and note that for any $\omega \in \Omega$ the primal feasibility implies,

$$\sum_s f(x^*(s, \omega)) q^*(s) = \xi.$$

To make the dependence on the parameter of interest, θ , explicit we rewrite the vector of primal feasibility constraints as,

$$\mathbf{R}^*(\theta, \boldsymbol{\lambda}^*(\theta), \mathbf{q}^*(\theta)) = \mathbf{1}.$$

By the implicit function theorem, we have that,

$$D_\theta \boldsymbol{\lambda}^* = [D_{\boldsymbol{\lambda}^*} \mathbf{R}^*]_{|\Omega| \times |\Omega|}^{-1} \left(- [D_{\mathbf{q}^*} \mathbf{R}^*]_{|\Omega| \times |S|} D_\theta \mathbf{q}^* - D_\theta \mathbf{R}^* \right). \quad (\text{B.10})$$

⁸We can use the first-order conditions (4.5) and the fact that C^* is homogeneous of degree 1 in the information structure.

We can exploit the structure of the Jacobian with respect to λ^* to obtain the inverse. Note that,

$$D\lambda^* \mathbf{R}^*_{\omega, \tilde{\omega}} = \begin{cases} 1 & \text{if } \omega \neq \tilde{\omega}, \\ 1 - \frac{1}{\mu(\omega)} \frac{\sum_s f'(x^*(s, \omega)) q^*(s)}{\xi} & \text{if } \omega = \tilde{\omega}. \end{cases} \quad (\text{B.11})$$

Since all off-diagonal elements are constant we can apply the Sherman-Morrison formula to obtain the inverse of the Jacobian,

$$[D\lambda^* \mathbf{R}^*]^{-1} = I_{|\Omega|} \mathbf{d} - \frac{1}{1 + \langle \mathbf{d}, \mathbf{1}_{|\Omega|} \rangle} \mathbf{d} \mathbf{d}', \quad (\text{B.12})$$

where $I_{|\Omega|}$ is the identity matrix of size $|\Omega| \times |\Omega|$, $\mathbf{1}_{|\Omega|}$ is the vector of ones of length $|\Omega|$, and \mathbf{d} is the vector with entries,

$$\mathbf{d}_\omega = -\frac{\mu(\omega) \xi}{\sum_s f'(x^*(s, \omega)) q^*(s)}. \quad (\text{B.13})$$

Further,

$$D_{\mathbf{q}^*} \mathbf{R}^*_{\omega, s} = \frac{f(x^*(s, \omega))}{\xi} - A, \quad (\text{B.14})$$

for some constant A . Since $D_{\theta} \mathbf{q}^*$ sums to zero we can ignore A .

Analogous derivations pertaining to the implicit function theorem hold for arbitrary incentive structures.

Continuing with the proof of the current section for the case of the general symmetric tracking problem with $\theta = h/\kappa$, denote $n := |\Omega|$ and

$$x^*(s, \omega) = \begin{cases} h^* & \text{if } s = a_\omega, \\ l^* & \text{otherwise.} \end{cases} \quad (\text{B.15})$$

We have,

$$D_{\theta} \mathbf{R}^* = \mathbf{1}_{|\Omega|} \frac{(f'(h^*) - f'(l^*)) q^*(a_\omega)}{\xi}. \quad (\text{B.16})$$

Note that in the current case $D_{\theta} \mathbf{q}^* = \mathbf{0}$ and \mathbf{d}_ω is constant across ω so we have,

$$[D\lambda^* \mathbf{R}^*]^{-1} \mathbf{1}_{|\Omega|} = \mathbf{1}_{|\Omega|} \frac{\xi}{(f(h^*) - f'(h^*)) + (n-1)(f(l^*) - f'(l^*))}. \quad (\text{B.17})$$

Combining we have,

$$\lambda^*_{\omega} = \frac{(f(h^*) - f'(h^*)) q^*(a_\omega)}{(f(h^*) - f'(h^*)) + (n-1)(f(l^*) - f'(l^*))}. \quad (\text{B.18})$$

We can express the difference,

$$f(x) - f'(x) = f(x) \left(1 - \frac{1}{(\alpha - 1)x} \right) \quad (\text{B.19})$$

Noting that both q^* and μ are uniform we arrive at,

$$\begin{aligned} \epsilon_{\omega, \theta}^{a_{\omega, b}} &= \frac{f(h^*) \left(1 - \frac{1}{(\alpha-1)h^*}\right) \left(\frac{1}{(\alpha-1)l^*} - \frac{1}{(\alpha-1)h^*}\right)}{f(h^*) \left(1 - \frac{1}{(\alpha-1)h^*}\right) + (n-1)f(l^*) \left(1 - \frac{1}{(\alpha-1)l^*}\right)} + \frac{1}{(\alpha-1)h^*} \\ &= \frac{f(h^*) \left(1 - \frac{1}{(\alpha-1)h^*}\right) \frac{1}{(\alpha-1)l^*} + (n-1)f(l^*) \left(1 - \frac{1}{(\alpha-1)l^*}\right) \frac{1}{(\alpha-1)h^*}}{f(h^*) \left(1 - \frac{1}{(\alpha-1)h^*}\right) + (n-1)f(l^*) \left(1 - \frac{1}{(\alpha-1)l^*}\right)} \end{aligned} \quad (\text{B.20})$$

We next prove a lemma on the magnitudes of $(\alpha-1)h^*$ and $(\alpha-1)l^*$ and then finish the proof.

Lemma 1. *In the general symmetric tracking problem for $\alpha < 1$ the optimal solution satisfies,*

$$(\alpha-1)h^* < \alpha < (\alpha-1)l^*; \quad (\text{B.21})$$

for $\alpha > 1$ the optimal interior solution satisfies,

$$(\alpha-1)h^* > \alpha > (\alpha-1)l^*. \quad (\text{B.22})$$

Proof. First, take the stationarity first-order condition (4.2) and rearrange,

$$\frac{(\sum_{\tilde{\omega}} P^*(s | \tilde{\omega})^\alpha \mu(\tilde{\omega}))^{-1}}{\left(\sum_{\tilde{s}} (\sum_{\tilde{\omega}} P^*(\tilde{s} | \tilde{\omega})^\alpha \mu(\tilde{\omega}))^{\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}}} P^*(\omega)^\alpha = \left(\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_\omega^*}{\mu(\omega)} + \frac{\delta_{\omega, s}^*}{\mu(\omega)}\right)\right)^{\frac{\alpha}{\alpha-1}}. \quad (\text{B.23})$$

Multiplying by $\mu(\omega)$ and summing across all states yields,

$$\begin{aligned} \Psi &:= \left(\sum_{\tilde{s}} \left(\sum_{\tilde{\omega}} P^*(\tilde{s} | \tilde{\omega})^\alpha \mu(\tilde{\omega})\right)^{\frac{1}{\alpha}}\right)^{\frac{\alpha}{1-\alpha}} = \\ &\sum_{\omega} \mu(\omega) \left(\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_\omega^*}{\mu(\omega)} + \frac{\delta_{\omega, s}^*}{\mu(\omega)}\right)\right)^{\frac{\alpha}{\alpha-1}} \quad \forall s \in \mathcal{S}. \end{aligned} \quad (\text{B.24})$$

Note that the above condition holds for all actions s irrespective of whether or not they are chosen with strictly positive probability under the optimal strategy. The constant Ψ is strictly positive and real—it is equal to $\Psi = \exp(-I_\alpha(P^*, \mu))$.

Using the primal feasibility constraints and noting that as q^* is uniform ξ is equal to the right-hand side of equation (B.24) we have that,

$$\sum_s f(x^*(s, \omega)) q^*(s) = \xi = \exp(-I_\alpha(P^*, \mu)) \leq 1 \quad \forall \omega \in \Omega; \quad (\text{B.25})$$

$$\sum_{\omega} f(x^*(s, \omega))^\alpha \mu(\omega) = \xi = \exp(-I_\alpha(P^*, \mu)) \leq 1 \quad \forall s \text{ with } q^*(s) > 0. \quad (\text{B.26})$$

We obtain the inequalities by the fact that the α -mutual information is always weakly positive, $I_\alpha(P^*, \mu) \geq 0$.

In the symmetric problem the uniform barycenter q^* in (B.25) implies,

$$\frac{1}{n}f(h^*) + \frac{n-1}{n}f(l^*) \leq 1$$

$$f(l^*) \left(\underbrace{\frac{f(h^*)}{f(l^*)}}_{>1} + (n-1) \right) \leq n \quad (\text{B.27})$$

Next, use the equality,

$$\sum_{\tilde{s}} f(x^*(\tilde{s}, \omega)) q^*(\tilde{s}) = \sum_{\tilde{\omega}} f(x^*(s, \tilde{\omega}))^\alpha \mu(\tilde{\omega}) \quad \text{for any } \omega \text{ and } s \text{ with } q^*(s) > 0, \quad (\text{B.28})$$

which in the symmetric tracking problem implies,

$$\underbrace{\left(\frac{\alpha-1}{\alpha} h^* \right)^{\frac{\alpha}{\alpha-1}}}_{>0} \left(1 - \left(\frac{\alpha-1}{\alpha} h^* \right)^{-1} \right) = \underbrace{(n-1) \left(\frac{\alpha-1}{\alpha} l^* \right)^{\frac{\alpha}{\alpha-1}}}_{>0} \left(\left(\frac{\alpha-1}{\alpha} l^* \right)^{-1} - 1 \right). \quad (\text{B.29})$$

If $\alpha < 1$ then necessarily $(\alpha-1)l^* > (\alpha-1)h^* > 0$ since $l < h$ and conditional probabilities are interior. This implies $f(h^*) > f(l^*) > 0$ and the term in the bracket of equation (B.27) is greater than n . To satisfy the equation we need

$$f(l^*) < 1 \quad \implies \quad \frac{\alpha-1}{\alpha} l^* > 1. \quad (\text{B.30})$$

And to match the signs of the two sides of equation (B.29) we need,

$$\frac{\alpha-1}{\alpha} h^* < 1. \quad (\text{B.31})$$

If $\alpha > 1$ then necessarily $(\alpha-1)h^* > (\alpha-1)l^* > 0$ since $h < l$ and conditional probabilities are interior. This implies $f(h^*) > f(l^*) > 0$ and the term in the bracket of equation (B.27) is greater than n . To satisfy the equation we need

$$f(l^*) < 1 \quad \implies \quad \frac{\alpha-1}{\alpha} l^* \leq 1. \quad (\text{B.32})$$

And to match the signs of the two sides of equation (B.29) we need,

$$\frac{\alpha-1}{\alpha} h^* > 1. \quad (\text{B.33})$$

□

Lastly, note that equation (B.29) implies that under $\alpha < 1$ the denominator of equation (B.20) is always negative, while under $\alpha > 1$ it is always positive.

To complete the proof of proposition 2 we use lemma 1 and equation (B.20).

Case I. $\alpha < 1$.

By lemma 1 we have $(\alpha - 1)l^* > \alpha > (\alpha - 1)h^* > 0$ and $f(h^*) > f(l^*) > 0$. Given the denominator of (B.20) being negative,

$$\begin{aligned} \epsilon_{\omega, \theta}^{a_{\omega, b}} &> 1 && \text{if } (\alpha - 1)l^* < 1, \\ \epsilon_{\omega, \theta}^{a_{\omega, b}} &< 1 && \text{if } (\alpha - 1)l^* > 1. \end{aligned} \quad (\text{B.34})$$

Case II. $\alpha > 1$.

By lemma 1 we have $(\alpha - 1)h^* > \alpha > (\alpha - 1)l^* > 0$ and $f(h^*) > f(l^*) > 0$. Given the denominator of (B.20) being positive,

$$\begin{aligned} \epsilon_{\omega, \theta}^{a_{\omega, b}} &> 1 && \text{if } (\alpha - 1)l^* < 1, \\ \epsilon_{\omega, \theta}^{a_{\omega, b}} &< 1 && \text{if } (\alpha - 1)l^* > 1. \end{aligned} \quad (\text{B.35})$$

□

C Proof of Proposition 3

Note that the α -mutual information for $\alpha \in (0, 1) \cup (1, \infty)$ is given by,

$$I_{\alpha}(P, \mu) = \frac{\alpha}{\alpha - 1} \log \sum_s \left(\sum_{\omega} \mu(\omega) P(s | \omega)^{\alpha} \right)^{\frac{1}{\alpha}}. \quad (\text{C.1})$$

First, we prove information monotonicity. With the transformations introduced in section 5 defining P' , \bar{P} , and $\bar{\mu}$, we need to show,

$$I_{\alpha}(P, \mu) \geq I_{\alpha}(P', \mu).$$

Take the monotone increasing transformation $\phi: t \mapsto \frac{1}{\alpha - 1} \exp(t \frac{\alpha - 1}{\alpha})$, information monotonicity is equivalent to,

$$\phi(I_{\alpha}(P, \mu)) \geq \phi(I_{\alpha}(P', \mu))$$

$$\frac{1}{\alpha - 1} \sum_s \left(\sum_{\omega} \mu(\omega) P(s | \omega)^{\alpha} \right)^{\frac{1}{\alpha}} \geq \frac{1}{\alpha - 1} \sum_s \left(\sum_{\omega} \mu(\omega) P'(s | \omega)^{\alpha} \right)^{\frac{1}{\alpha}}.$$

Note that the transformed α -mutual information is posterior separable as previously remarked in (3.10).

Given that likelihoods of P' are constant over elements of the partition we can write,

$$\phi(I_{\alpha}(P', \mu)) = \frac{1}{\alpha - 1} \sum_s \left(\sum_{z \in Z} \left(c_z^{\alpha} \sum_{\omega \in \bar{\Omega}_z} \mu(\omega) \right) \right)^{\frac{1}{\alpha}} \quad (\text{C.2})$$

$$= \frac{1}{\alpha - 1} \sum_s \left(\sum_{z \in Z} \left(\sum_{\omega \in \bar{\Omega}_z} P(s | \omega) \mu(\omega) \right)^{\alpha} \left(\sum_{\omega \in \bar{\Omega}_z} \mu(\omega) \right)^{(1 - \alpha)} \right)^{\frac{1}{\alpha}}. \quad (\text{C.3})$$

On the other hand,

$$\phi(I_\alpha(P, \mu)) = \frac{1}{\alpha - 1} \sum_s \left(\sum_{z \in Z} \sum_{\omega \in \bar{\Omega}_z} (P(s | \omega) \mu(\omega))^\alpha \mu(\omega)^{(1-\alpha)} \right)^{\frac{1}{\alpha}}. \quad (\text{C.4})$$

Denote the inner function in the parentheses as $f_\alpha(x, y) := x^\alpha y^{(1-\alpha)}$. Note that $f_\alpha(x, y)$ is the ‘‘conic transform’’ of $\tilde{f}_\alpha(x) := x^\alpha$, that is $f_\alpha(x, y) = \tilde{f}_\alpha(x/y)y$ for $y > 0$. Since the prior is $\mu \in \text{int}\Delta(\bar{\Omega})$ we have no issues applying the transformation. Next, we will use the fact that the conic transform preserves convexity/concavity.

We have two cases to consider. If $\alpha < 1$,

- \tilde{f}_α is concave, so the conic transform f_α is also concave. By Jensen’s inequality we have,

$$\frac{1}{|\bar{\Omega}_z|} \sum_{\omega \in \bar{\Omega}_z} (P(s | \omega) \mu(\omega))^\alpha \mu(\omega)^{(1-\alpha)} \leq \frac{1}{|\bar{\Omega}_z|} \left(\sum_{\omega \in \bar{\Omega}_z} P(s | \omega) \mu(\omega) \right)^\alpha \left(\sum_{\omega \in \bar{\Omega}_z} \mu(\omega) \right)^{(1-\alpha)}.$$

- The inequality holds for all $z \in Z$ so it is preserved for the sum over the partitions.
- Since the mapping $x \mapsto x^{\frac{1}{\alpha}}$ is monotone increasing for $x > 0$ and the scalar $\frac{1}{\alpha-1} < 0$, we have that information monotonicity holds,

$$I_\alpha(P, \mu) \geq I_\alpha(P', \mu).$$

For $\alpha > 1$ we can apply an analogous argument with \tilde{f}_α being convex, and so f_α also being convex. Applying Jensen’s inequality and noting that $\frac{1}{\alpha-1} > 0$ we have,

$$I_\alpha(P, \mu) \geq I_\alpha(P', \mu).$$

Next we show that

$$I_\alpha(P', \mu) = I_\alpha(\bar{P}, \bar{\mu}).$$

Note, that since by construction $\mu(\bar{\Omega}_z) = \bar{\mu}(\bar{\Omega}_z)$ for all $z \in Z$ we have that,

$$c_z = \frac{\sum_{\omega \in \bar{\Omega}_z} P(s | \omega) \mu(\omega)}{\sum_{\omega \in \bar{\Omega}_z} \mu(\omega)} = \frac{\sum_{\omega \in \bar{\Omega}_z} P(s | \omega) \mu(\omega)}{\sum_{\omega \in \bar{\Omega}_z} \bar{\mu}(\omega)} = \bar{c}_z \quad \forall z \in Z.$$

Now use the decomposition as in equation (C.2),

$$\begin{aligned} \phi(I_\alpha(P', \mu)) &= \frac{1}{\alpha - 1} \sum_s \left(\sum_{z \in Z} \left(c_z^\alpha \sum_{\omega \in \bar{\Omega}_z} \mu(\omega) \right) \right)^{\frac{1}{\alpha}} \\ &= \frac{1}{\alpha - 1} \sum_s \left(\sum_{z \in Z} \left(\bar{c}_z^\alpha \sum_{\omega \in \bar{\Omega}_z} \bar{\mu}(\omega) \right) \right)^{\frac{1}{\alpha}} = \phi(I_\alpha(\bar{P}, \bar{\mu})). \end{aligned}$$

This completes our proof. □

D Proof of Proposition 5

The proof for the claims of proposition 5 hinges on the directional derivatives of the α -mutual information evaluated at points on the boundary of the constraint set.

Note that the partial derivative of I_α with respect to $P(a | \omega)$ is

$$\frac{\partial I_\alpha(P, \mu)}{\partial P(a | \omega)} = \frac{\alpha}{\alpha - 1} \frac{(\sum_{\tilde{\omega}} \mu(\tilde{\omega}) P(a | \tilde{\omega})^\alpha)^{\frac{1-\alpha}{\alpha}}}{\sum_{\tilde{s}} (\sum_{\tilde{\omega}} \mu(\tilde{\omega}) P(\tilde{s} | \tilde{\omega})^\alpha)^{\frac{1}{\alpha}}} \mu(\omega) P(a | \omega)^{\alpha-1}. \quad (\text{D.1})$$

We treat the two claims separately.

Claim 1. For $\alpha \leq 1$ there exists no non-trivial event $E \subsetneq \text{supp}(\mu)$ that the DM optimally learns with certainty.

Suppose there exists such a set. Then necessarily there exists $a \in \mathcal{A}$ such that $P^*(a | \omega_i) = 0$ for some i while $P^*(a | \omega_j) > 0$ for some j . Furthermore, since signal probabilities have to add up to one, there exists some b such that $P^*(b | \omega_i) > 0$. Moving in the feasible direction that results in a new experiment $(P^*(a | \omega_i) + \epsilon, P^*(b | \omega_i) - \epsilon)$ with all other conditional signal probabilities being the same as in P^* reduces information costs by infinity at the margin as shown by equation (D.1) while only affecting payoffs by a finite amount. This contradicts the optimality of P^* .

Claim 2. For $\alpha > 1$ take an event of the form $E_a := \{\omega : u(a, \omega) > u(b, \omega) \forall b \neq a\}$. Suppose that the optimal attention strategy under the scaled payoff, $\pi u(\cdot, \cdot)$, is such that $\gamma^a(E_a) < 1$. This implies that $P^*(a | \omega_k) > 0$ for some $\omega_k \notin E_a$. Since by assumption all payoffs are distinct, $\omega_k \in E_b$ for some other action $b \neq a$ which implies $u(b, \omega_k) > u(a, \omega_k)$. Moving in the feasible direction that results in a new experiment $(P^*(b | \omega_k) + \epsilon, P^*(a | \omega_k) - \epsilon)$ with all other conditional signal probabilities being the same as in P^* changes information costs by a finite amount on the margin as shown by equation (D.1), so there must be π such that the considered move has greater utility increase on the margin $\pi(u(b, \omega_k) - u(a, \omega_k))$. This would contradict the optimality of the information structure inducing $\gamma^a(E_a) < 1$. □

E Proof of Proposition 6

First, take the stationarity first-order condition (4.2) and rearrange,

$$\frac{(\sum_{\tilde{\omega}} P^*(s | \tilde{\omega})^\alpha \mu(\tilde{\omega}))^{-1}}{\left(\sum_{\tilde{s}} (\sum_{\tilde{\omega}} P^*(\tilde{s} | \tilde{\omega})^\alpha \mu(\tilde{\omega})^{\frac{1}{\alpha}})\right)^{\frac{\alpha}{\alpha-1}}} P^*(\omega)^\alpha = \left(\frac{\alpha - 1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_\omega^*}{\mu(\omega)} + \frac{\delta_{\omega, s}^*}{\mu(\omega)}\right)\right)^{\frac{\alpha}{\alpha-1}}. \quad (\text{E.1})$$

Multiplying by $\mu(\omega)$ and summing across all states yields,

$$\begin{aligned} \Psi &:= \left(\sum_{\tilde{s}} \left(\sum_{\tilde{\omega}} P^*(\tilde{s} | \tilde{\omega})^\alpha \mu(\tilde{\omega}) \right)^{\frac{1}{\alpha}} \right)^{\frac{\alpha}{1-\alpha}} = \\ &\sum_{\omega} \mu(\omega) \left(\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} + \frac{\delta_{\omega, s}^*}{\mu(\omega)} \right) \right)^{\frac{\alpha}{\alpha-1}} \quad \forall s \in \mathcal{S}. \end{aligned} \quad (\text{E.2})$$

Note that the above condition holds for all actions s irrespective of whether or not they are chosen with strictly positive probability under the optimal strategy. The constant Ψ is strictly positive and real.

We consider two cases depending on the value of α .

Case 1. $\alpha < 1$.

If $\mu P_{\mathcal{S}}^*(a) > 0$ (or equivalently $q^*(a) > 0$) then $P^*(a | \omega) > 0$ for all ω by proposition 5. By the complementary slackness conditions we have $\delta_{\omega, a}^* = 0$ for all ω . Hence we have,

$$\sum_{\omega} \mu(\omega) \left(\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} \right) \right)^{\frac{\alpha}{\alpha-1}} = \Psi \quad \forall s \text{ such that } q^*(s) > 0. \quad (\text{E.3})$$

If $\mu P_{\mathcal{S}}^*(a) = 0$ (or equivalently $q^*(a) = 0$) then $P^*(a | \omega) = 0$ for all ω . By the complementary slackness and dual feasibility conditions $\delta_{\omega, a}^* > 0$. Since $\alpha - 1 < 0$ we have that,

$$\sum_{\omega} \mu(\omega) \left(\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} \right) \right)^{\frac{\alpha}{\alpha-1}} < \Psi \quad \forall s \text{ such that } q^*(s) = 0. \quad (\text{E.4})$$

Note that irrespective of $P^*(s | \omega)$ being strictly positive $\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} \right)$ is positive otherwise the stationarity condition would be violated.

Case 2. $\alpha > 1$.

Under $\alpha > 1$ we might have $P^*(a | \omega) = 0$ while $\mu P^*(a) > 0$. Consider such a case. By the stationarity first-order condition we have,

$$\frac{u(a, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} = -\frac{\delta_{\omega, a}^*}{\mu(\omega)}, \quad (\text{E.5})$$

and since $\alpha - 1 > 0$,

$$\left[\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} \right) \right]^+ = \frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} + \frac{\delta_{\omega, s}^*}{\mu(\omega)} \right). \quad (\text{E.6})$$

Correspondingly we arrive at,

$$\sum_{\omega} \mu(\omega) \left(\left[\frac{\alpha-1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} \right) \right]^+ \right)^{\frac{\alpha}{\alpha-1}} = \Psi \quad \forall s \text{ such that } q^*(s) > 0. \quad (\text{E.7})$$

Note that without thresholding the expression at zero for the terms where $P^*(s | \omega) = 0$ would render the exponentiation ill-defined as we would be taking the fractional power of a negative number.

If $\mu P_{\mathcal{S}}^*(a) = 0$ (or equivalently $q^*(a) = 0$) then $P^*(a | \omega) = 0$ for all ω . By the complementary slackness and dual feasibility conditions $\delta_{\omega,a}^* > 0$. Since $\alpha - 1 > 0$ we have that,

$$\sum_{\omega} \mu(\omega) \left(\left[\frac{\alpha - 1}{\alpha} \left(\frac{u(s, \omega)}{\kappa} - \frac{\lambda_{\omega}^*}{\mu(\omega)} \right) \right]^+ \right)^{\frac{\alpha}{\alpha-1}} < \Psi \quad \forall s \text{ such that } q^*(s) = 0. \quad (\text{E.8})$$

Now taking the convex combination of these expressions with the convex weights given by the barycenter of the optimal experiment,

$$\sum_{\tilde{s}} q^*(\tilde{s}) \sum_{\tilde{\omega}} \mu(\tilde{\omega}) \left(\left[\frac{\alpha - 1}{\alpha} \left(\frac{u(\tilde{s}, \tilde{\omega})}{\kappa} - \frac{\lambda_{\tilde{\omega}}^*}{\mu(\tilde{\omega})} \right) \right]^+ \right)^{\frac{\alpha}{\alpha-1}} = \Psi. \quad (\text{E.9})$$

By the fact that $\mu P_{\mathcal{S}}^*(s) > 0 \iff q^*(s) > 0$ proposition 6 holds if we substitute the unconditional signal probabilities for the barycenter. \square

F Proof of Proposition 7

Since the objective is bounded from below and the objective function is decreasing at each step the algorithm is converging. That does not necessarily imply that the limit point is the global optimum of the problem. We can use the fact that algorithm 1 is a special case of the block non-linear Gauss-Seidel method—also called block coordinate descent method—with quasi-convex objective function, two blocks, and convex constraints. We can directly apply theorem 5.1 of Tseng (2001).

The conditions that we need to verify in order to apply the results of Tseng (2001) are:

- The Rényi divergence, D_{α} is continuous on its effective domain, which holds for finite distributions.
- D_{α} is quasi-convex in both arguments—in fact it is jointly quasi-convex (Van Erven and Harremoës, 2014, Theorem 13)—and hemivariate in both arguments.
- D_{α} is lower semi-continuous (Van Erven and Harremoës, 2014, Theorem. 15).
- The effective domain of D_{α} is a product of subsets of Euclidean spaces.

The only property that we need to show is that D_{α} is hemivariate in both arguments separately on its effective domain. A function is hemivariate in an argument if it is not constant on any line segment belonging to the effective domain. That is for a given Q take distinct P_0 and P_1 such that the line segment $\lambda(P_0, Q) + (1 - \lambda)(P_1, Q)$ for any $\lambda \in [0, 1]$ is in the effective domain of D_{α} . Suppose that D_{α} is not hemivariate. Then with $P_{\lambda} := \lambda P_0 + (1 - \lambda)P_1$,

$$D_{\alpha}(P_{\lambda} \| Q) := \frac{1}{\alpha - 1} \log \left(\sum_{x \in \mathcal{X}} P_{\lambda}(x)^{\alpha} Q(x)^{(1-\alpha)} \right) = C \quad \forall \lambda \in [0, 1]. \quad (\text{F.1})$$

We can derive a contradiction by considering

$$\sum_{x \in \mathcal{X}} P_{\lambda}(x)^{\alpha} Q(x)^{(1-\alpha)} = \exp((\alpha - 1)C) \quad \forall \lambda \in [0, 1], \quad (\text{F.2})$$

and using the strict convexity or strict concavity of the function $z \mapsto z^\alpha$ depending on the value of α . The case for Q is analogous.

As a result theorem 5.1 of [Tseng \(2001\)](#) applies and the modified Blahut-Arimoto algorithm converges to a stationary point of the objective function at which all KKT conditions are satisfied.

□