



# Robust inattentive discrete choice

Lars Peter Hansen<sup>a,b,1,2</sup>, Jianjun Miao<sup>c,d,e,1</sup>, and Hao Xing<sup>f,1</sup>

Contributed by Lars Peter Hansen; received August 27, 2024; accepted December 20, 2024; reviewed by Tommaso Denti and Benjamin Hebert

Rational inattention models characterize optimal decision-making in data-rich environments. In such environments, it can be costly to look carefully at all of the information. Some information is much more salient for the decision at hand and merits closer scrutiny. The inattention decision model formalizes this choice and deduces how best to navigate through the potentially vast array of data when making decisions. In the rational formulation, the decision-maker commits fully to a subjective prior distribution over the possible states of the world that could be realized. We relax this assumption and look for a robustly optimal solution to the inattention problem by allowing the decision-maker to be ambiguity averse with respect to this prior. We feature a setup that is deliberately simple by a) assuming a discrete set of choices, b) using Shannon's mutual information to quantify attention costs, and c) imposing relative entropy with respect to a baseline probability distribution to quantify prior divergence. We provide necessary and sufficient conditions for the robust solution and develop numerical methods to solve it. In comparison to the rational solution with no prior uncertainty, our decision-maker slants priors in more cautious or pessimistic directions when deducing how to allocate attention over the range of available information. This approach implements a form of robustness to prior misspecification, or equivalently, a form of ambiguity aversion. We explore some examples that show how the robust solution differs from the rational solution with a commitment to a subjective prior distribution and how it differs from imposing risk aversion.

rational inattention | robustness | information acquisition

People often make choices with limited information. In so doing, they determine where to focus their attention when the available information is vast. To study such decision problems, Sims (1, 2) introduces to economics a rational inattention (RI) framework with close ties to information theory and, in particular, to the rate distortion problem. The information costs therein are constructed using the Shannon's (3) entropy-based mutual information (henceforth the Shannon model). While Sims (1) focuses on continuous choices in dynamic settings, Caplin and Dean (4), Matějka and McKay (5), and Caplin et al. (6, 7) adopt the RI framework to study discrete choice problems in static settings. We follow these latter authors by studying discrete choice problems.\* Such problems have applications to such diverse fields as labor economics, industrial organization, macroeconomics, and political economy. This framework also has links to model selection problems in statistics and control theory. Those literatures study the ability to statistically discriminate among models using existing datasets. This becomes an attention allocation problem when the decision maker is unsure which statistical model is correct and explores some of a vast amount information to help with this determination.

While the Shannon model is a standard approach in the literature, it may suffer from a problem of prior misspecification. Specifically, the Shannon mutual information is defined as the difference between the entropies of the prior and the posterior. The prior is assumed to be exogenously given to the decision maker (DM), and the DM has full confidence in this prior. The perspective adopted in this paper is that of a DM with prior ambiguity who is looking to make a robustness adjustment.

Prior uncertainty can have important ramifications for the Shannon model because, as Caplin et al. (6) show, the different specifications of the prior can generate different so called "consideration sets." Consideration sets are collections of alternatives chosen with positive probabilities. The implications of prior ambiguity carry over to these sets and change how the DM views alternatives within these sets.

Operationally, we incorporate prior robustness into a static, discrete-choice setting by integrating a version of refs. 8–10. As Hansen and Sargent (8) show, there is a robust prior interpretation of the smooth ambiguity model of Klibanoff et al. (11). While the smooth ambiguity model has multiple axiomatic defenses, Hansen and Sargent (9) argue

\*An earlier version of this research was presented at the American Economic Association meetings in January of 2022 in a session on "Expectations and Macro-Finance."

## Significance

When people make decisions they often have a large volume of information at their disposal. Faced with this challenge, in practice they focus their attention on a more limited amount of salient information. To assess what information should be exploited in decision-making, we extend a modeling approach that is called "rational inattention." This approach assumes that the decision maker has a subjective prior that is used in determining how to direct attention across the vast set of information at the decision-maker's disposal. The resulting solution can be highly sensitive, however, to the choice of the prior. We relax this assumption and allow the decision-maker to confront subjective uncertainty when taking discrete actions, thereby making the decisions more robust.

Author affiliations: <sup>a</sup>Economics Department, University of Chicago, Chicago, IL 60637; <sup>b</sup>Booth School of Business, University of Chicago, Chicago, IL 60637; <sup>c</sup>School of Economics, Zhejiang University, Hangzhou 310058, China; <sup>d</sup>Tsingshan Institute for Advanced Business Studies, Zhejiang University, Hangzhou 310058, China; <sup>e</sup>Department of Economics, Boston University, Boston, MA 02215; and <sup>f</sup>Finance Department, Questrom School of Business, Boston University, Boston, MA 02215

Author contributions: L.P.H., J.M., and H.X. designed research; performed research; and wrote the paper.

Reviewers: T.D., New York University; and B.H., Stanford University.

The authors declare no competing interest.

Copyright © 2025 the Author(s). Published by PNAS. This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

<sup>1</sup>L.P.H., J.M., and H.X. contributed equally to this work.

<sup>2</sup>To whom correspondence may be addressed. Email: lhansen@uchicago.edu.

Published February 7, 2025.

that the robust prior interpretation opens the door to the axiomatic formulations of refs. 12 and 13.<sup>†</sup> Intuitively, the DM does not have a single prior about the state of the world but instead has a baseline or reference prior. The DM explores the consequences of alternative priors that may be different from a baseline prior. The deviation of any prior from the baseline prior is penalized by a cost modeled by the relative entropy. The cost is scaled in utility units by a robustness parameter that dictates the degree of ambiguity aversion. Alternatively, the DM may impose a relative entropy constraint on the family of priors, giving rise to a version of Gilboa and Schmeidler's (16) max-min utility preferences.

## 1. Model

We first introduce the model setup and then provide an equivalent formulation to simplify the decision problem.

**1.1. Setup.** In what follows, we use a bold lowercase letter  $\mathbf{x}$  to denote a random variable and a letter  $x$  to denote its realization. The set  $\Delta(Y)$  denotes the collection of probability distributions on any finite set  $Y$ . One such set is the finite state space  $X = \{x_1, x_2, \dots, x_M\}$  and a prior  $\mu \in \Delta(X)$ . Another is the set  $S$ , which denotes a finite set of signal realizations.

The DM does not observe the state but can acquire a signal about the state. Based on the signal, the DM selects an element from a finite action set  $A$  to maximize expected utility  $u: X \times A \rightarrow \mathbb{R}$ . The DM also chooses an optimal information structure by paying a cost modeled by the Shannon mutual information. A strategy is a pair  $(d, \sigma)$  composed of

1. an information strategy,  $d$ , consisting of a system of signal distributions  $d(s|x)$ , for all  $s \in S, x \in X$ ;
2. an action strategy,  $\sigma: S \rightarrow A$ , specifying an action  $a = \sigma(s)$  when observing a signal  $s$ .

Let  $\Sigma$  denote the set of all strategies. While the information strategy,  $d$ , is an object of choice, the prior over states,  $\mu$ , is given with full commitment in the standard Shannon RI specification.

An information strategy and prior over states induce a joint distribution  $d \otimes \mu$  over  $X \times S$ :

$$(d \otimes \mu)(x, s) = d(s|x) \mu(x).$$

In what follows, we let  $\mathbb{E}_{d \otimes \mu}$  denote the expectation computed using  $d(s|x)\mu(x)$ . The implied marginal distribution over signals is

$$\tau(s) = \sum_x \mu(x) d(s|x),$$

and the conditional distribution over states is

$$\mu_s(x) \stackrel{\text{def}}{=} \frac{d(s|x) \mu(x)}{\tau(s)}$$

for signals  $s$  for which  $\tau(s) > 0$ , and  $\mu_s(x) = \mu(x)$  for signals  $s$  with  $\tau(s) = 0$ .

We consider Shannon entropy for both the marginal and the conditional distribution of states:

$$\mathbb{H}(\mu) = - \sum_x \mu(x) \log \mu(x)$$

<sup>†</sup>For other statistical motivations for smooth ambiguity, see refs. 14 and 15.

$$\mathbb{H}(\mu_s) = - \sum_x \mu_s(x) \log \mu_s(x).$$

Entropy  $\mathbb{H}(\mu)$  measures the amount of uncertainty embedded in the prior about the underlying state. Mutual information measures the reduction of uncertainty after observing signals:

$$\mathbb{I}(d \otimes \mu) \stackrel{\text{def}}{=} \sum_s [\mathbb{H}(\mu) - \mathbb{H}(\mu_s)] \tau(s).$$

Equivalently,

$$\begin{aligned} \mathbb{I}(d \otimes \mu) &= \sum_{x,s} d(s|x) \mu(x) (\log [d(s|x) \mu(x)] - \log [\mu(x) \tau(s)]), \end{aligned}$$

which is the Kullback–Leibler (KL) divergence between the joint distribution of states and signals relative to the product of the marginals.<sup>‡</sup>

As a precursor to our analysis, we first state the standard RI problem as a benchmark for comparison:

**Problem 1.** (*Signal RI problem*)

$$V(\mu) = \max_{(d,\sigma) \in \Sigma} \mathbb{E}_{d \otimes \mu}(u[\mathbf{x}, \sigma(\mathbf{s})]) - \lambda \mathbb{I}(d \otimes \mu).$$

In this problem, we interpret  $\lambda$  as the shadow cost of information expressed in utility units. It scales mutual information and hence acts as a per unit measure of cost. We refer to it as a shadow cost because we are not claiming that it has an actual market counterpart.

As is well known in the literature, the solution to Problem 1 is sensitive to the specification of the prior  $\mu$ . Since prior misspecification is a concern, we consider a robust alternative in which the single prior,  $\mu$ , is replaced by baseline prior,  $\hat{\mu}$ , and the DM explores the adverse consequences of prior misspecification subject to a scaled version of a relative entropy or KL cost:

$$\mathbb{R}(\mu || \hat{\mu}) := \sum_x \mu(x) \log \frac{\mu(x)}{\hat{\mu}(x)}.$$

Our robust alternative to Problem 1 is

**Problem 2.** (*Robust signal RI problem*)

$$\begin{aligned} W(\hat{\mu}) &= \max_{(d,\sigma) \in \Sigma} \min_{\mu \in \Delta(X)} \mathbb{E}_{d \otimes \mu}(u[\mathbf{x}, \sigma(\mathbf{s})]) - \lambda \mathbb{I}(d \otimes \mu) \\ &\quad + \xi \mathbb{R}(\mu || \hat{\mu}), \end{aligned} \tag{1}$$

where  $\xi > 0$  denotes a robustness parameter.

Since the constraint sets are compact and the objective function is continuous, the robust RI problem has a solution. The minimization in [1] reflects the DM's aversion to prior ambiguity as parameterized by  $\xi$ .<sup>§</sup>

As specified,  $\xi$  governs a smooth tradeoff between the utility maximization and the cost of prior ambiguity as in multiplier

<sup>‡</sup>See ref. 17 for a textbook treatment of information theory.

<sup>§</sup>Mutual information depends on the prior. An alternative formulation of information cost was proposed in ref. 18 using channel capacity, which is the largest mutual information among all possible priors. In our Problem 2, the prior contributes to the expected utility. Notice that if we set  $\xi = 0$ , the minimization over priors would compute the channel capacity as the implied measure of information were we to ignore how the prior influences the expected utility.

preferences of Hansen and Sargent (19) and the variational preferences of Maccheroni et al. (12).<sup>‡</sup> Large values of  $\xi$  induce large costs in deviating from the baseline prior and hence a small amount of aversion to ambiguity. Alternatively, we may suppose that there is a constraint on relative entropy:

$$\mathbb{R}(\mu || \hat{\mu}) \leq \eta$$

for some  $\eta > 0$ . In this case  $\eta$  governs the ambiguity aversion as in ref. 16. Consider a Lagrangian specification of the problem with  $\xi$  as a multiplier as in refs. 19 and 20. Using a standard duality argument:

$$\begin{aligned} W(\hat{\mu}) &= \max_{(d, \sigma) \in \Sigma} \max_{\xi \geq 0} \min_{\mu \in \Delta(X)} \mathbb{E}_{d \otimes \mu} (u[\mathbf{x}, \sigma(\mathbf{s})]) - \lambda \mathbb{I}(d \otimes \mu) \\ &\quad + \xi \mathbb{R}(\mu || \hat{\mu}) - \xi \eta \\ &= \max_{\xi \geq 0} \max_{(d, \sigma) \in \Sigma} \min_{\mu \in \Delta(X)} \mathbb{E}_{d \otimes \mu} (u[\mathbf{x}, \sigma(\mathbf{s})]) - \lambda \mathbb{I}(d \otimes \mu) \\ &\quad + \xi \mathbb{R}(\mu || \hat{\mu}) - \xi \eta. \end{aligned}$$

The inner max-min problem of the second representation is of the same form as the smooth version posed in Problem 2.

An entirely analogous argument applies if we impose an information constraint that

$$\mathbb{I}(d \otimes \mu) \leq \kappa$$

we could add the term  $\lambda \kappa$  to the objective and treat  $\lambda$  as a Lagrange multiplier. More generally, for each choice of  $(\lambda, \xi)$ , we could deduce an implied information constraint and relative entropy constraint. When imposing numerical parameter values in applications, the magnitudes of the implied constraints can be informative ways to assess parameter plausibility.

**Example 1:** We now consider a model selection problem familiar from statistics. Let states be models and an action be a guess of a model with utility function:

$$u(x, a) = \begin{cases} 1 & a = x \\ 0 & a \neq x. \end{cases}$$

While we posited a symmetric utility function, the rewards for the correct identification of a model and penalties for mistakes could be distinct across partitions. This would be the case, for instance, when the alternative statistical models might have different implications for future courses of action. Robustness considerations would come into play in the choice of prior over states.

Now suppose there is a hypothetical dataset that defines the upper bound of information, captured by a distribution  $\vec{d}(y | x)$  (or a data-rich likelihood when viewed as a function of  $\mathbf{x}$ ), where  $y$  is a realization of the potentially available data. This upper bound could be formalized using Blackwell's (21) ordering and information quantified using the same Shannon approach applied here by comparing posteriors to the prior. In other words, a potentially interesting extension of our formulation could impose this upper bound on what can be revealed by the signals in conjunction with an information cost in a way that preserves convexity as in Blackwell's formulation. Robustness considerations could come into play, not only in the choice of prior over states but also in the choice of the probabilistic upper bound on the available information about the models.<sup>#</sup>

<sup>‡</sup> Refs. 12 and 19 feature model misspecification, but subsequent literature has extended their approach to the case of prior misspecification. See ref. 9.

<sup>#</sup> The paper (22) uses a related information restriction, but applied coherently across multiple decision-makers and combined with an  $f$ -divergence bound. They deduce implications of games with imperfect information from the perspective of an outside observer.

**1.2. Simplifying the Problem Solution.** When solving Problem 2., there is no reason to have signals provide more information than what is necessary to incorporate into decision making since signals are costly to obtain. This suggests an alternative formulation with a probabilistic decision,  $p(a | x)$ , and an information cost imposed directly on  $p$ . In effect, the choice  $p$  now serves as both an action and a signal distribution conditioned on the state. In this subsection, we pose and study this alternative formulation.

Observe that a strategy  $(d, \sigma) \in \Sigma$  generates a choice rule  $p \in \Delta(A|X)$ , where  $\Delta(A|X)$  denotes the set of conditional distributions  $p(\cdot | x)$  on  $A$  given by

$$p(a|x) = \Pr(\sigma(s) = a|x) = \sum_{\{s: \sigma(s)=a\}} d(s|x). \quad [2]$$

Instead of the strategy  $(d, \sigma)$ , we now let actions play the role of signals in terms of the attention costs. We suppose that the DM specifies a choice rule that is a conditional probability of actions given the states. This conditional distribution along with the prior,  $\mu$ , over states imply a joint distribution  $p \otimes \mu$  over  $X \times A$ :

$$(p \otimes \mu)(x, a) \equiv p(a|x) \mu(x).$$

Given this joint distribution, we define the mutual information as

$$\mathbb{I}(p \otimes \mu) \stackrel{\text{def}}{=} \sum_a [\mathbb{H}(\mu) - \mathbb{H}(\mu_a)] q(a),$$

where  $\mu_a \in \Delta(X)$  and  $q \in \Delta(A)$  denote the posterior and marginal distributions implied by the joint distribution  $p \otimes \mu$ :

$$q(a) = \sum_x p(a|x) \mu(x), \quad [3]$$

$$\mu_a(x) = \frac{p(a|x) \mu(x)}{\sum_x p(a|x) \mu(x)}, \text{ if } q(a) > 0. \quad [4]$$

**Problem 3.** (Choice-based robust RI problem)

$$\begin{aligned} J(\hat{\mu}) &\stackrel{\text{def}}{=} \max_{p \in \Delta(A|X)} \min_{\mu \in \Delta(X)} \mathbb{E}_{p \otimes \mu} [u(\mathbf{x}, \mathbf{a})] - \lambda \mathbb{I}(p \otimes \mu) \\ &\quad + \xi \mathbb{R}(\mu || \hat{\mu}), \end{aligned} \quad [5]$$

where  $\mathbb{E}_{p \otimes \mu}$  is an expectation operator given distribution  $p \otimes \mu$ .

In the Online Appendix, we provide a Recommendation Lemma to establish the equivalence of Problems 2 and 3. Before providing solutions to Problem 3 in the next section, we first present some special limiting cases. First, when  $\xi = \infty$ , Problem 3 is reduced to the standard RI problem. In this case the worst-case prior is the baseline prior. Second, when  $\xi = 0$ , Problem 3 is reduced to the following one:

$$\max_{p \in \Delta(A|X)} \min_{\mu \in \Delta(X)} \mathbb{E}_{p \otimes \mu} [u(\mathbf{x}, \mathbf{a})] - \lambda \mathbb{I}(p \otimes \mu).$$

This is the extreme case in which the DM thinks any prior in the feasible set  $\Delta(X)$  is possible and the penalty cost is zero. Third, when  $\lambda = 0$ , the DM can acquire signals to fully observe the state. The DM will select the highest-payoff action with probability one conditional on each state. Finally, when  $\lambda = \infty$ , the DM does not acquire any information about the state and selects an action to maximize expected utility given the worst-case prior.

**Remark 1:** For the rational intention specification in which there is a commitment to the baseline prior, the distribution  $q$  has meaning as a statement of the ex ante probability of the different actions. Under robustness there is ambiguity about the prior. The choice of a worst-case prior is not the subjective belief of the DM, but rather it is a device to obtain a robustly optimal choice of  $p^*(a | x)$ . Since the distribution,  $q^*(a)$ , inherits the worst-case prior, it is not interpretable as the ex ante probability of the alternative actions. Nevertheless, it may still provide a useful summary of the robust choices  $p^*(a | x)$  in comparison to an average using the baseline prior distribution.

Notice that randomization among the discrete options remains a possibility, as a choice of  $p$  that is independent of the state is feasible for the DM.

## 2. Model Solution

We include a seemingly superfluous contribution to the maximization problem by including the marginal  $q$  over actions. At the same time, we ignore the constraint linking this marginal to the conditional  $p(a | x)$  and the prior  $\mu(x)$ . That link will follow directly from the optimization that we investigate.<sup>11</sup> Along with the inclusion of  $q$  in the optimization, we introduce the corresponding consideration sets defined as the set of actions chosen with positive probabilities. Formally, for any  $q \in \Delta(A)$ , the associated consideration set is defined as  $B(q) = \{a \in A : q(a) > 0\}$ .

To help understand why we may include  $q$  in the optimization, a straightforward derivation shows that

$$\min_{q \in \Delta(A)} \sum_x \mu(x) \sum_a [\log p(a|x) - \log q(a)] p(a|x)$$

has

$$q^*(x) = \sum_x \mu(x) p(a|x)$$

as its solution, which is simply the marginal distribution over  $A$  implied by  $p$ . By adding and subtracting  $\log \mu(x)$  inside the summations, observe that the minimizing objective is the information measure  $\mathbb{I}(p \otimes \mu)$ .

We solve Problem 3 by computing:

$$J(\hat{\mu}) \stackrel{\text{def}}{=} \max_{q \in \Delta(A), p \in \Delta(A|X)} \min_{\mu \in \Delta(X)} F(p, q, \mu), \quad [6]$$

where

$$F(p, q, \mu) \stackrel{\text{def}}{=} \sum_x \mu(x) \left[ G(p, q)(x) + \xi \log \frac{\mu(x)}{\hat{\mu}(x)} \right], \quad [7]$$

and

$$G(p, q)(x) \stackrel{\text{def}}{=} \sum_a p(a|x) \left[ u(x, a) - \lambda \log \frac{p(a|x)}{q(a)} \right]$$

The function  $F$  is concave in  $(p, q)$  and convex in  $\mu$ . By the Minimax Theorem, we can exchange the extremization without affecting the optimized value:

$$J(\hat{\mu}) = \min_{\mu \in \Delta(X)} \max_{q \in \Delta(A), p \in \Delta(A|X)} F(p, q, \mu). \quad [8]$$

<sup>11</sup>This equivalent formulation utilizes special features of the Shannon entropy cost function. For general uniformly posterior-separable (UPS) cost functions, Caplin et al. (6) and Caplin and Dean (23) introduce a posterior-based approach. Dynamic problems with UPS cost functions are analyzed by Miao and Xing (24) and a numerical algorithm is also proposed therein.

This also gives an alternative way to compute the robust solution to the RI problem. In what follows, we implement a hybrid approach.

Consider the inner minimization in the problem Eq. 6, taking  $p = p^*$  and  $q = q^*$  as given. This has a well-known solution from robust control theory, large deviation theory, and other applications of relative entropy (for example, see refs. 20, 25, and 26):

$$\mu^*(x) = \frac{\exp \left[ - \left( \frac{1}{\xi} \right) v(x) \right] \hat{\mu}(x)}{\sum_y \exp \left[ - \left( \frac{1}{\xi} \right) v(y) \right] \hat{\mu}(y)} \quad [9]$$

for

$$v(x) \stackrel{\text{def}}{=} G(p^*, q^*)(x). \quad [10]$$

This displays exponential tilting of  $\mu^*$  relative to  $\hat{\mu}$  toward values of  $x$  for which the values of  $v$  are relatively low. The minimized objective is known to be

$$J(\hat{\mu}) = -\xi \log \sum_x \hat{\mu}(x) \exp \left[ - \left( \frac{1}{\xi} \right) v(x) \right]. \quad [11]$$

Consider the inner maximization in the problem Eq. 8. This takes  $\mu = \mu^*$  as given, and as consequence, the last term in the objective  $F$  can be ignored. It suffices to study:

$$\max_{q \in \Delta(A), p \in \Delta(A|X)} \sum_x \mu^*(x) G(p, q)(x) \quad [12]$$

The maximizing solution for  $(p, q)$  coincide with those of a standard RI problem. Following ref. 5, the first-order conditions for  $p$  imply that

$$p^*(a | x) = \frac{q(a) \exp [u(x, a)/\lambda]}{\sum_b q(b) \exp [u(x, b)/\lambda]} \quad \mu^*(x) > 0. \quad [13]$$

While this formula can be deduced by ignoring a nonnegativity constraint on  $p$ , the resulting  $p^*$  will be nonnegative provided that  $q$  is.

Next, we follow ref. 6 by substituting formula Eq. 13 into the objective  $F$  for Problem 8 and maximizing the resulting objective as a function of  $q$ . Thus, we form:

$$\widehat{G}(q)(x) \stackrel{\text{def}}{=} \lambda \log \sum_a q(a) \exp [u(x, a)/\lambda], \quad [14]$$

and  $q^*$  solves:

$$\max_{q \in \Delta(A)} \lambda \sum_x \mu^*(x) \log \sum_a q(a) \exp [u(x, a)/\lambda].$$

In particular,  $v = \widehat{G}(q^*)$ , where  $v$  is defined in Eq. 10.

We now impose that  $q$  be nonnegative, as this constraint could well bind. Following ref. 6, we obtain the relation:

$$\sum_x \mu^*(x) \left( \frac{\exp [u(x, a)/\lambda]}{\sum_b q^*(b) \exp [u(x, b)/\lambda]} \right) \begin{cases} \leq 1 & \forall a, \\ = 1 & a \in B(q^*) \end{cases} \quad [15]$$

as a set of first-order conditions. Since  $F$  is concave in  $(p, q)$ , conditions Eqs. 13 and 15 are both necessary and sufficient for the maximization given  $\mu = \mu^*$ .

We combine these results in the following proposition.

**Proposition 1.** The triple  $(p^*, q^*, \mu^*)$  is the solution to the robust RI problem if, and only if, Eqs. 9, 13, and 15 are satisfied. The resulting value function  $J$  is given by

$$J(\hat{\mu}) = -\xi \log \sum_x \hat{\mu}(x) \exp \left[ -\left(\frac{1}{\xi}\right) v(x) \right].$$

for  $v$  given by Eq. 10.

The robust decision is the conditional distribution  $p^*(a | x)$ , capturing the signal distribution. Recall that we “normalized” the problem so that the action coincides with the signal. While our decision maker chooses  $\mu^*$  along with  $p^*$ , the worst-case prior,  $\mu^*$ , is a vehicle for constructing  $p^*$  and together these determine  $q^*$ , and the worst-case posterior distribution is given by

$$\mu_a^*(x) = \frac{p^*(a | x) \mu^*(x)}{q^*(a)},$$

provided that  $q^*(a) > 0$ .

We interpret  $v(x)$  as the ex post payoff (utility) derived from a standard RI problem for a given prior  $\mu^*$ . Moreover, imposing  $\mu^*$  on the maximization problem over  $(p, q)$  gives the standard RI solution. Thus one interpretation is that the robust solution replaces  $\hat{\mu}$  with  $\mu^*$  in an otherwise standard RI problem where  $\mu^*$  tilts the prior distribution by putting more weight on states for which ex post payoff is relatively lower. An alternative robust preference interpretation has the analogous slanting for possible pair  $(p, q)$  in forming a ranking over such admissible pairs. Under either interpretation, the resulting value function is smaller as it reflects the cost of prior misspecification or ambiguity aversion.

**Remark 2:** As Caplin et al. (6) note, in the absence of robustness considerations, we could rephrase the problem to be a choice of  $(\mu, q, \mu_a(\cdot))$  subject to

$$\sum_a \mu_a(x) q(a) = \mu(x)$$

for all states  $x$ . The necessary and sufficient conditions are again given by Eqs. 9, 13, and 15 substituting

$$\frac{\mu_a^*(x) q^*(a)}{\mu^*(x)}$$

for  $p^*(a | x)$ . This same insight extends to our analysis with robustness concerns with the following caveat. In the problem analyzed here, it is the twisted or worst-case probability that adjusts for prior ambiguity and limited attention that would be chosen under this strategy. In contrast, the chosen signal distributions that we feature do not require this qualification.

### 3. Numerical Method

An algorithm of refs. 27 and 28 finds numerical solutions to the standard RI problem.

This algorithm is an application of the general block coordinate descent algorithm, which applies to multivariate optimization problems in which the constraint set has a Cartesian product property. The key idea is that at each iteration one solves the optimization problem with respect to each of the block coordinates taken in cyclic order. Based on this idea, we propose the following generalized Arimoto–Blahut algorithm to solve the robust RI Problem 3:

1. Start with a guess  $\mu^{(0)} \in \Delta(X)$  with  $\mu^{(0)}(x) > 0$  for all  $x$  and a guess  $p^{(0)} \in \Delta(A|X)$  with  $p^{(0)}(a|x) > 0$  for all  $(x, a)$ .
2. Given  $(p^{(k)}, q^{(k)}, \mu^{(k)})$  for step size  $0 < s \leq 1$  compute:

$$\bar{q}^{(k+1)}(a) = \sum_x \mu^{(k)}(x) p^{(k)}(a|x),$$

$$q^{(k+1)}(a) = q^{(k)}(a) + s [\bar{q}^{(k+1)}(a) - q^{(k)}(a)]$$

$$v^{(k+1)}(x) = \lambda \log \sum_a q^{(k+1)}(a) \exp(u(x, a) / \lambda)$$

3. Given  $(q^{(k+1)}, v^{(k+1)}, p^{(k)}, \mu^{(k)})$ , for step size  $0 < s \leq 1$  construct:

$$\bar{p}^{(k+1)}(a|x) = \frac{q^{(k+1)}(a) \exp(u(x, a) / \lambda)}{\sum_b q^{(k+1)}(b) \exp(u(x, b) / \lambda)},$$

$$p^{(k+1)}(a | x) = p^{(k)}(a | x)$$

$$+ s [\bar{p}^{(k+1)}(a | x) - p^{(k)}(a | x)]$$

$$\bar{\mu}^{(k+1)}(x) = \frac{\exp(-v^{(k+1)}(x) / \xi) \hat{\mu}(x)}{\sum_y \exp(-v^{(k+1)}(y) / \xi)},$$

$$\mu^{(k+1)}(x) = \mu^{(k)}(x) + s [\bar{\mu}^{(k+1)}(x) - \mu^{(k)}(x)].$$

4. Iterate over integer  $k \geq 0$  until convergence.

We applied this algorithm to compute the numerical solutions for some illustrations reported in the next two sections.\*\*

### 4. Finding the Good Alternative

We begin with a consumer choice problem analyzed previously by Caplin et al. (6). There is a discrete set of  $M$  possible consumption goods for the consumer to select among. Alternatively, the choices might be over different investment opportunities. An action is a choice of a good, and the value of the good depends on the underlying state. For simplicity, we make the action space and the state space identical in this example. Consumer preferences are captured by the utility function:

$$u(x, a) = \begin{cases} u_g & \text{if } x = a \\ u_b & \text{if } x \neq a \end{cases} \quad [16]$$

where  $u_g > u_b$ . For convenience, we parameterize  $u_b = \lambda \log \bar{u}$  and  $u_g = \lambda(\log \bar{u} + \log(1 + \delta))$ . The numerical magnitude of  $\bar{u}$  turns out to be inconsequential to the analysis.††

Let  $\hat{\mu}(x)$  be the baseline prior probability that option  $x$  yields the good prize. Without loss of generality, we order states according to  $\hat{\mu}(x_i) \geq \hat{\mu}(x_{i+1}) \geq \hat{\mu}(x_M) > 0$ , for  $i \in \{1, \dots, M - 1\}$ .

The DM can learn about the state by paying a mutual information cost. The DM also has concerns about prior misspecification and seeks robust decision making that performs well when their prior  $\mu$  may deviate from  $\hat{\mu}$ . The decision problem can be formalized as Problem 3.

\*\*The Github repository for the computational code used in this paper can be found at [https://github.com/lphansen/rational\\_inattention](https://github.com/lphansen/rational_inattention). In particular, a user-friendly notebook can be accessed at [https://lphansen.github.io/rational\\_inattention/intro.html](https://lphansen.github.io/rational_inattention/intro.html).

††It is straightforward to reinterpret this decision problem as one of model selection.

The solution to this decision problem will have a threshold whereby some actions may lie outside the consideration set. A central part of the solution is the characterization of this threshold. As we will show, introducing robustness into the analysis can expand the consideration set.

To construct a threshold, introduce:

$$\rho_k \stackrel{\text{def}}{=} \left[ \sum_{i=1}^k [\widehat{\mu}(x_i)]^\psi \left( \frac{1}{\delta + k} \right) \right]^{\frac{1}{\psi}}$$

where

$$\psi \stackrel{\text{def}}{=} \frac{\xi}{\lambda + \xi}. \tag{17}$$

Notice that

$$\rho_1 < \widehat{\mu}(x_1).$$

Find the largest  $k^* \leq M$  such that

$$\rho_k < \widehat{\mu}(x_k) \quad \forall 1 \leq k \leq k^*.$$

Only the  $k \leq k^*$  are in the consideration set. Define:

$$\rho^* \stackrel{\text{def}}{=} \rho_{k^*}.$$

**Proposition 2.** *For the robust solution for the consumer choice problem, the choice rule is given by*

$$p^*(a|x_k) = \frac{1 + \delta}{\delta} \left[ \left( \frac{\widehat{\mu}(x_k)}{\rho^*} \right)^\psi - 1 \right] \left( \frac{\rho^*}{\widehat{\mu}(x_k)} \right)^\psi, \quad a = x_k,$$

$$p^*(a|x_k) = \frac{1}{\delta} \left[ \left( \frac{\widehat{\mu}(x_\ell)}{\rho^*} \right)^\psi - 1 \right] \times \left( \frac{\rho^*}{\widehat{\mu}(x_k)} \right)^\psi, \quad \forall a = x_\ell, k \neq \ell \leq k^*,$$

$$p^*(a|x_k) = 0, \quad \forall a = x_\ell, \ell > k^*$$

for  $1 \leq k \leq k^*$ , and

$$p^*(a|x_k) = \frac{1}{\delta} \left[ \left( \frac{\widehat{\mu}(x_\ell)}{\rho^*} \right)^\psi - 1 \right], \quad \forall a = x_\ell, \ell \leq k^*,$$

$$p^*(a|x_k) = 0, \quad \forall a = x_\ell, \ell > k^*$$

for  $M \geq k > k^*$ .

The worst-case prior is given by

$$\mu^*(x_k) = \frac{(\rho^*)^{1-\psi} \widehat{\mu}(x_k)^\psi}{(\rho^*)^{1-\psi} \sum_{i=1}^{k^*} \widehat{\mu}(x_i)^\psi + \sum_{i=k^*+1}^M \widehat{\mu}(x_i)}$$

for  $1 \leq k \leq k^*$ , and

$$\mu^*(x_k) = \frac{\widehat{\mu}(x_k)}{(\rho^*)^{1-\psi} \sum_{i=1}^{k^*} \widehat{\mu}(x_i)^\psi + \sum_{i=k^*+1}^M \widehat{\mu}(x_i)}$$

for  $M \geq k > k^*$ .

Proposition 2 also implies a worst-case posterior distribution for the underlying states given the signal (action), and a worst-case marginal distribution for the actions. As is the case for the worst-case prior, these are not intended to depict the DM's actual beliefs.

As  $\xi \rightarrow \infty$ , a specification studied by Caplin et al. (6), deviating from the baseline prior is increasingly costly and thus

the worst-case prior approximates the baseline prior itself. In this case, Proposition 2 converges to the corresponding findings in Theorem 1 of ref. 6.

The robust solution is qualitatively similar to that when  $\xi = \infty$ . Specifically, the consideration set is determined by a threshold strategy: the decision maker will consider only alternatives with a prior probability that exceeds an endogenously determined threshold  $\rho^*$ .

Quantitatively, increasing prior robustness concerns may enlarge the consideration set, leading the decision maker to consider more options. Formally, decreasing the penalty  $\xi$  increases the prior ambiguity concerns. We verify in the Online Appendix that it decreases  $\rho_k$  for each  $k$ , and thus  $k^*$  is larger. In the extreme, as  $\xi \downarrow 0$ ,  $\rho_k$  diminishes to zero,  $k^* = M$ ,  $\mu^*(x) = \frac{1}{M}$ ,

$$\mu_a^*(x) \rightarrow \frac{1 + \delta}{\delta + M} \text{ if } x = a,$$

$$\mu_a^*(x) \rightarrow \frac{1}{\delta + M} \text{ if } x \neq a,$$

for any  $x \in X$ .

More generally, notice that the solution depends on  $\xi$  and  $\lambda$  through the construction of  $\psi$  given in Eq. 17. Thus, there is a one-dimensional curve in the  $(\xi, \lambda)$  space for which the solution remains the same. For instance, reduction in  $\xi$  can be offset by reductions in  $\lambda$ , the attention cost, without altering the solutions for  $p^*(a|x)$  and  $\mu^*(x)$ .<sup>‡‡</sup>

## 5. Correlated Options

We now consider a static investment problem with three options and two states. Options one and two are negatively correlated while option three is constant across states and hence risk-free. The two states are equally likely under the baseline prior  $\widehat{\mu}(x_1) = \widehat{\mu}(x_2) = 1/2$ . Let

$$u(x_i, j) = \frac{[c(x_i, j)]^{1-\alpha} - 1}{1 - \alpha}$$

for  $x_i \in X, j \in A$ , and  $0 \leq \alpha < 1$ . Table 1 presents the payoffs  $c(x_i, j)$  for each of three options as a function of the two different states.

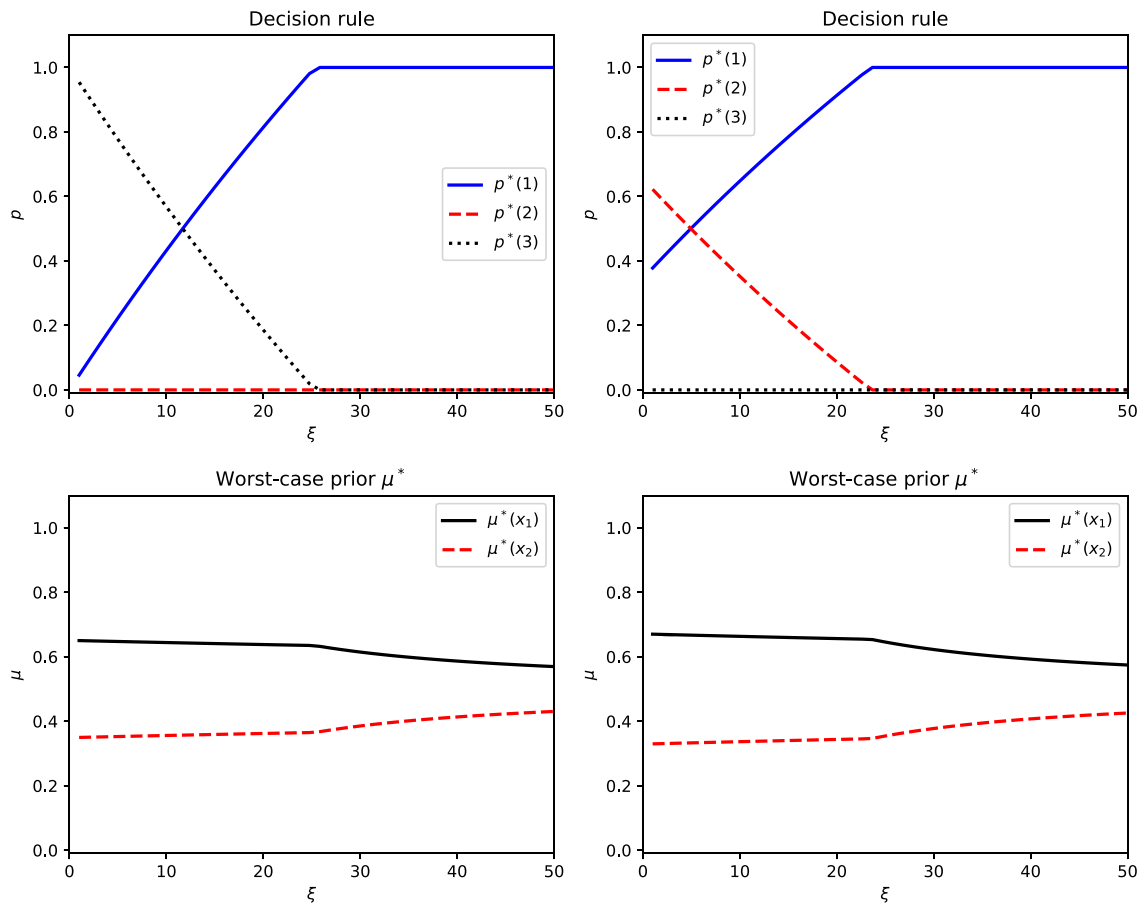
Notice that option one has the highest expected payoff and option two has the lowest under the baseline prior when  $5 < r < 10$ . We initially set  $\alpha = 0$ , implying risk neutrality. Consider first the special case in which  $\lambda = \infty$ , so that there is only prior ambiguity but no possibility for the DM to look for information about the underlying states. Randomization across actions independent of states is still allowed as the implied

**Table 1. This table displays payoffs  $c(x_i, j)$  for each state  $x_i \in X$  and each option  $j \in A$**

| Option \ State | $x_1$ | $x_2$        |
|----------------|-------|--------------|
| 1              | 0     | $2 \times r$ |
| 2              | $r$   | 0            |
| 3              | 5     | 5            |

We will consider different values of  $r > 0$  in the computations that follow.

<sup>‡‡</sup>This explicit link between  $\xi$  and  $\lambda$  is reminiscent of an observationally equivalent discussion in ref. 29 for a class of linear, Gaussian models.



**Fig. 1.** This figure explores the sensitivity to changes in  $\xi$ . The attention cost parameter,  $\lambda = \infty$ ; utility curvature parameter,  $\alpha = 0$ ; and the payoff on option two in state one is  $r = 7$  for the *Left* figure and  $r = 7.5$  for the *Right*.

mutual information will be zero. In this case, we simply write  $p(j|x_i) = p(j)$  for any  $j \in A$  and  $x_i \in X$ .

Consider first the case in which  $r = 7$  as depicted in the left side of Fig. 1, and  $r = 7.5$  on the right side of this same figure. In the absence of prior ambiguity ( $\xi = \infty$ ), action one maximizes expected utility in both cases. For  $r = 7$ , when  $\xi$  is 25 or less, the DM randomizes across actions one and three. Option two is never included as part of the solution. For  $r = 7.5$ , randomization only occurs at higher levels of ambiguity aversion (smaller values of  $\xi$ ).

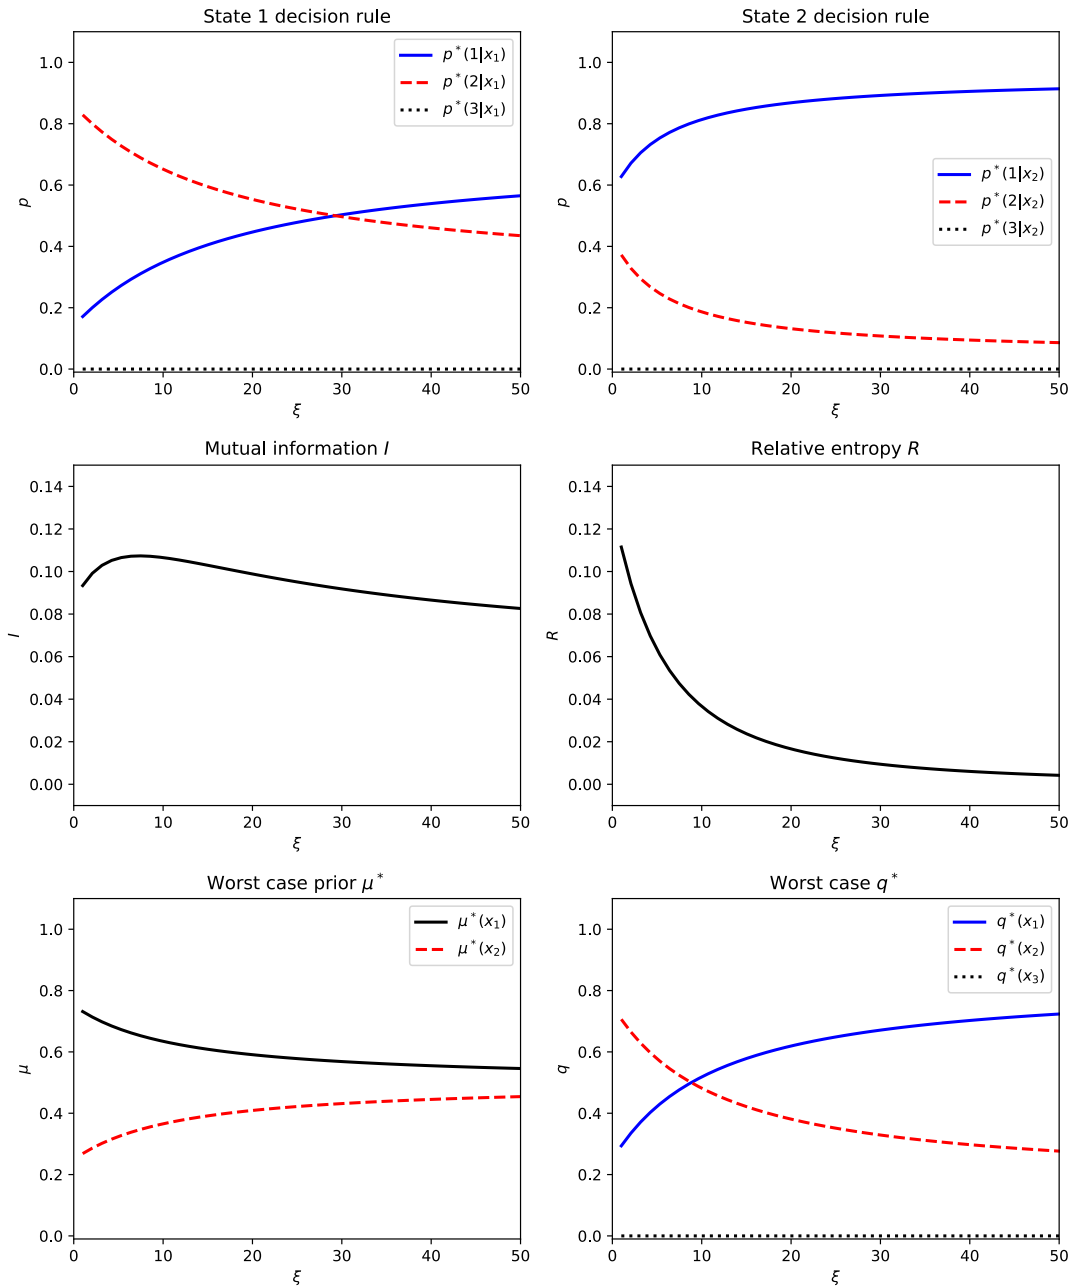
As  $\xi$  declines to zero, the randomized choices converge to a solution to the limiting  $\xi = 0$  case. In the  $\xi = 0$  limit, all priors over states are entertained. In this case, we recover a solution that is familiar from robust statistical decision theory. We equalize the expected payoffs conditioned on the states, which restricts  $p(2) = 2p(1)$ . We then maximize subject to this restriction, by choosing  $p(1) \in [0, 1/3]$  and setting  $p(2)$  accordingly. When  $r = 7.5$  all such probability choices agree. Thus there are multiple choices including the one depicted in Fig. 1. When  $r > 7.5$ , the optimized solution is  $p^*(1) = 1/3$ ; and when  $r < 7.5$ ,  $p^*(1) = p^*(2) = 0$ .

We now explore attention allocation. We again report results for  $r = 7$  and set  $\lambda = 10$  for illustration purposes. We report the findings in Fig. 2. For lower values of  $\xi$ , option two becomes attractive. In contrast to the  $\lambda = \infty$  case, this is even true for  $r < 7.5$ . In the large  $\xi$  limiting case, the solution coincides with the RI solution. In this limiting case, Caplin et al. (6) note that option two becomes attractive because it helps the decision-maker learn about the true state of nature in contrast to option three.

Robustness considerations about the prior further enhance the attractiveness of option two. This may be expected given what happens in the absence of attention considerations for  $r \geq 7.5$ . In the presence of attention costs, increasing  $r$ , to say  $r = 7.5$ , makes option two all the more attractive for  $\lambda = 10$ , as reported in the Online Appendix. Option three is never chosen in this illustration.

Notice in Fig. 2 that the worst-case prior is tilted toward state 1 and more so for small values of  $\xi$  (more prior ambiguity). As expected, the resulting relative entropy diminishes with  $\xi$ . Also observe that the implied mutual information is relatively stable across different values of  $\xi$ .

**Remark 3:** Proposition 3 of ref. 5 shows that if an action becomes more attractive according to prior beliefs, then the rationally inattentive DM will select that action with a higher unconditional probability. Prior robustness concerns provide a different perspective on this result. Without attention considerations, we see from Fig. 1 that even if option one is preferred under the benchmark probabilities, this can be reversed by entertaining robustness concerns. This carries over to Fig. 2 and to the  $r = 7.5$  counterpart as option one is clearly preferred under the baseline probabilities when  $\xi = \infty$ , but this gets reversed for low values of  $\xi$ . Matějka and McKay's (5) monotonicity will hold for the endogenously determined worst-case probabilities. But our DM has prior ambiguity and the worst-case probabilities are just a device to impute robust attention choices and are not intended as the actual beliefs.

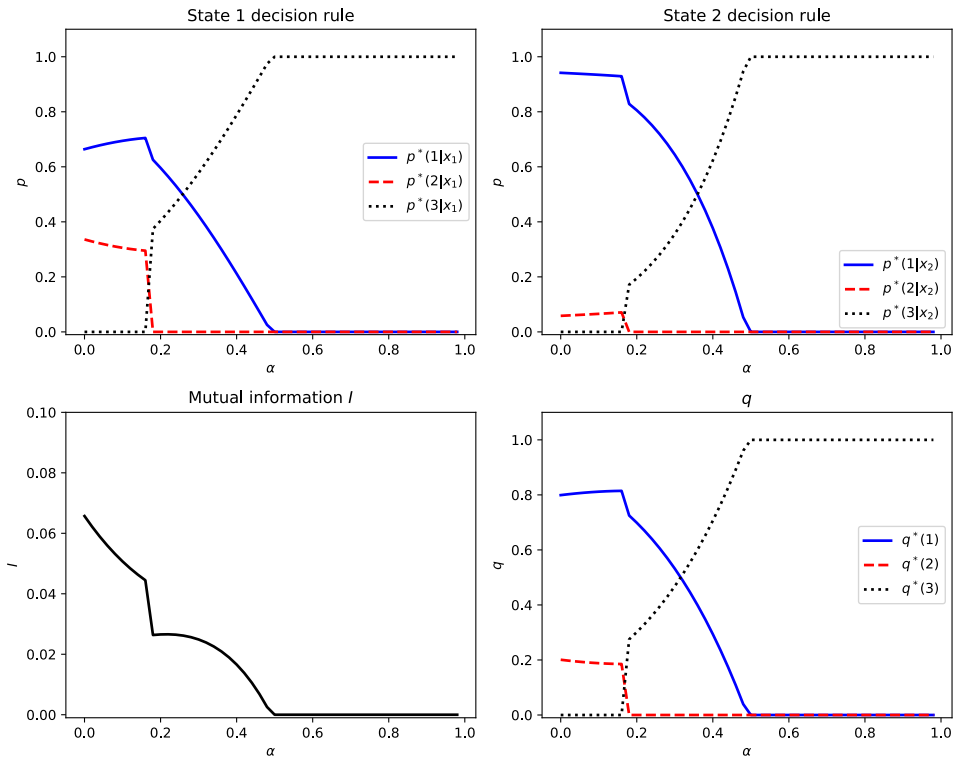


**Fig. 2.** This figure explores the sensitivity to changes in  $\xi$ . The attention cost parameter,  $\lambda = 10$ ; utility curvature parameter,  $\alpha = 0$ , and  $r = 7$ .

We next explore whether the impact of ambiguity aversion is similar to that of risk aversion. Recall that increases in  $\xi$  are associated with less aversion to ambiguity over priors. Perhaps the most interesting comparison is between relatively large  $\alpha$  and small  $\xi$ .

Consider the case in which  $\alpha = 1$ . This limit results in  $u(x_i, j) = \log c(x_i, j)$ . When  $r > 5$ , we see immediately that option one is chosen only if it is known that the realized state is  $x_1$ , and option two is chosen only if it is known that the realized state is  $x_1$ . These options only happen if the attention allocation is sufficient to reveal the actual state. Otherwise, option 3 is chosen with probability one. Thus attention choice will be one of two extremes, pay no attention or pay enough attention to reveal underlying state. In what follows, we explore more modest utility curvature.

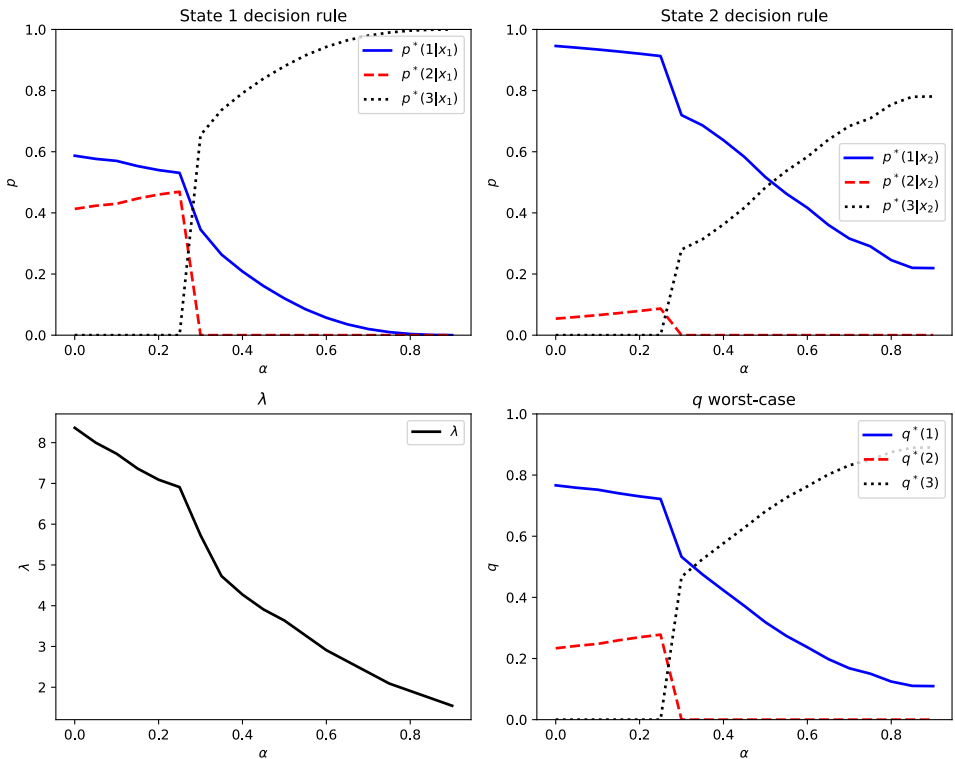
Fig. 3 investigates the implications for more modest values of  $\alpha$ , while omitting prior robustness concerns. The attention cost remains at  $\lambda = 10$ . Option two is considerably less attractive under risk rather than aversion to prior ambiguity. Indeed, under risk aversion, it is set to zero for values of  $\alpha$  that exceed 0.2. At this same threshold, option three begins to be considered and it becomes more prominent for larger  $\alpha$ . For  $\alpha$  in excess of 0.5, only the risk-less option (option three) is considered. Given this, the mutual information drops to zero. This is in contrast to the risk-neutral case in Fig. 2 where option three is not chosen for any of the values  $\xi$ . Overall, increasing  $\alpha$  makes option two less attractive. Recall from Fig. 2 that increases in  $\xi$  also make option two less attractive, albeit in a smoother way; but these increases imply reductions, not increases, in ambiguity aversion.



**Fig. 3.** This figure explores the sensitivity to changes in  $\alpha$ . The attention cost parameter,  $\lambda = 10$ ; robustness parameter,  $\xi = \infty$ , and  $r = 7.0$ .

Since attention costs are depicted in utility units, changing  $\alpha$  impacts how we should view the magnitude of the attention costs. For a different perspective on the impact of risk aversion, we also report results where attention is imposed as a constraint,

$\mathbb{I} \leq 0.1$ , in Fig. 4. The quantitative magnitudes are quite different, as should be expected since the attention allocation in Fig. 3 dropped to zero for the larger choices of  $\alpha$ . Now option one remains part of the optimized solution even for larger values



**Fig. 4.** This figure explores the sensitivity to changes in  $\alpha$ . The mutual information is constrained by  $\mathbb{I} \leq \kappa = 0.1$ ;  $\xi = \infty$  and  $r = 7.0$ .

of  $\alpha$ . Option two is a little bit more prominent than in the fixed  $\lambda$  case. Recall that probabilistic choice of signals and actions are explicitly linked in our computation of a solution as describe in Section 1.2.

To summarize our findings, we find important differences between the impact of risk and that of ambiguity aversion (preference for prior robustness). When the decision-maker becomes more risk averse, the riskless option three comes to dominate, as we might expect. With a fixed attention constraint, option one remains as part of the solution when we impose an information constraint in place of a cost. In contrast, when the decision-maker becomes more averse to ambiguity, depending on the payoff disparity, option two can remain attractive in contrast to options three. This outcome depends on the magnitude of the payoff disparity. By including attention choice subject to cost, option two remains even more attractive due to its role in conveying information.

We have two broad lessons to take away from this example.

- i) While in the previous RI literature, information revelation impacts choices, this example illustrates how robustness considerations amplify that role.
- ii) This example illustrates an aspect of Knight's (30) assertion that

Uncertainty must be taken in a sense radically distinct from the familiar notion of risk, ... and there are far-reaching and crucial differences in the bearings of the phenomena depending on which of the two is really present and operating.

## 6. Conclusion

Our examples are only meant to be illustrative. Many applications would seem to call for a large number of potential states. Our proposed computational algorithm could still be analyzed even with many more states. Recall that the robustly, optimal RI solution will have the same number of signals as decision possibilities, which can be exploited in computation. Prior ambiguity may well be even more prominent for larger state spaces as it may be even more difficult to specify the priors over such spaces with full confidence.

While many contributions in the existing literature embrace mutual information to model attention cost, some of its behavioral implications are inconsistent with experimental evidence (see, e.g., refs. 18, 23, 31, and 32). Motivated by this evidence, more flexible cost functions for information acquisition are proposed by Caplin et al. (7), Caplin and Dean (23), Hébert and Woodford (33), Pomatto et al. (34), and Bloedel and Zhong (35). These cost functions continue to depend on the DM's prior beliefs and so a robustness analysis over the priors, as we have illustrated here, continues to be a potentially important ingredient to incorporate.

In empirical or experimental studies, the observable data are state-dependent choice probabilities. To test theories of discrete choices, one has to make some additional assumptions. Caplin and Dean (4), Caplin et al. (7), and Caplin and Martin (36) make an important assumption that the DM and the econometrician share the same exogenously specified subjective

prior over states of the world. Under these assumptions, there are unique posteriors over states which can be used when testing implications. Specifically, they construct tests based on the ratio of posteriors conditioned on different actions. In addition, Caplin and Dean (4) point to two axioms, no improving attention cycles (NIAC) and no improving action switches (NIAS), as necessary and sufficient conditions for a costly information acquisition representation of state-dependent probabilities. In our setting, prior ambiguity implies posterior ambiguity, which brings into question many approaches to test the decision-making model. Section 2 of the Online Appendix shows that NIAC and NIAS can either hold or be violated in the prior ambiguity case. It remains an interesting challenge as to how best to test robust versions of the inattention model in the presence of ambiguity.

Here, we consider only the potential misspecification of priors over states, but not of the chosen signal distributions given the state. (For instance, see ref. 9 for an elaboration of potentially distinct forms of prior robustness and likelihood robustness with connections to decision theory, control theory, and statistics.) When a decision maker chooses to allocate attention, there may well be uncertainty about the resulting  $p(a | x)$ . This, too, would be a valuable extension of RI decision theory. While we adopt a static perspective in this paper, we suspect that many applications would be better suited for a dynamic or recursive extension to the formulation we investigate in this paper. This could open an additional motivation for robust information acquisition that would operate through continuation values since "tomorrow's prior" is "today's posterior" in a dynamic setting.

In other settings, the analog to our robustness perspective results in a representation of preferences that is a special case of smooth ambiguity aversion. Within this context, we find this "special case" to be particularly interesting because of it by providing a formalization of prior robustness pertinent to robust Bayesian analysis. In this paper, our treatment of prior ambiguity extends to construction of the information cost, which would not be the case had we followed ref. 37. Denti et al. (37) replace the attention cost used by Matějka and McKay (5) and others with a measure of the cost of statistical experimentation, arguing that the latter should not depend on the decision-maker prior. For instance, in the special case of information based on relative entropy, Denti et al. (37) consider a reference probability, say uniform, that is distinct from the decision-maker prior. Both probabilities show up as part of the solution to the decision-makers optimization problem, and there remains a potentially interesting scope for assessing prior robustness. For an entirely different formulation than ours of a model with smooth ambiguity and limited attention, see ref. 38. This paper adopts a different starting point and presumes that the ambiguity adjustment in ref. 38 applies to the utility contribution only.

**Data, Materials, and Software Availability.** There are no data underlying this work.

**ACKNOWLEDGMENTS.** We gratefully acknowledge comments by Maryam Farboodi, Diana Petrova, Monika Piazzesi, Doron Ravid, and Judy Yue, and research assistance by Chun Hei Hung. We are very appreciative of feedback provided on an earlier draft by Tommaso Denti and Benjamin Hebert.

1. C. A. Sims, "Stickiness" in *Carnegie-Rochester Conference Series on Public Policy* (1998), vol. 49, pp. 317-356.
2. C. A. Sims, Implications of rational inattention. *J. Mon. Econ.* **50**, 665-690 (2003).
3. C. E. Shannon, A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379-423 (1948).
4. A. Caplin, M. Dean, Revealed preference, rational inattention, and costly information acquisition. *Am. Econ. Rev.* **105**, 2183-2203 (2015).

5. F. Matějka, A. McKay, Rational inattention to discrete choices: A new foundation for the multinomial logit model. *Am. Econ. Rev.* **105**, 272-298 (2015).
6. A. Caplin, M. Dean, J. Leahy, Rational inattention, optimal consideration sets, and stochastic choice. *Rev. Econ. Stud.* **86**, 1061-1094 (2019).
7. A. Caplin, M. Dean, J. Leahy, Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Eur. J. Polit. Econ.* **130**, 1676-1715 (2022).

8. L. P. Hansen, T. J. Sargent, Recursive robust estimation and control without commitment. *J. Econ. Theor.* **136**, 1–27 (2007).
9. L. P. Hansen, T. J. Sargent, Risk, ambiguity, and misspecification: Decision theory, robust control, and statistics. *J. Appl. Econ.* **39**, 969–999 (2023).
10. L. P. Hansen, J. Miao, Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9163–9168 (2018).
11. P. Klibanoff, M. Marinacci, S. Mukerji, A smooth model of decision making under uncertainty. *Econometrica* **73**, 1849–1892 (2005).
12. F. Maccheroni, M. Marinacci, A. Rustichini, Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* **74**, 1147–1498 (2006).
13. T. Strzalecki, Axiomatic foundations of multiplier preferences. *Econometrica* **79**, 47–73 (2011).
14. S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, L. Montrucchio, Ambiguity and robust statistics. *J. Econ. Theor.* **148**, 974–1049 (2013).
15. T. Denti, L. Pomatto, Model and predictive uncertainty: A foundation for smooth ambiguity preferences. *Econometrica* **90**, 551–584 (2022).
16. I. Gilboa, D. Schmeidler, Maxmin expected utility with non-unique prior. *J. Math. Econ.* **18**, 141–153 (1989).
17. T. M. Cover, J. A. Thomas, *Elements of Information Theory* (Wiley, 2012).
18. M. Woodford, "Inattentive valuation and reference-dependent choice" (Working Paper, 2012).
19. L. P. Hansen, T. J. Sargent, Robust control and model uncertainty. *Am. Econ. Rev.* **91**, 60–66 (2001).
20. I. R. Petersen, M. R. James, P. Dupuis, Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Trans. Autom. Control* **45**, 398–412 (2000).
21. D. Blackwell, "Comparison of experiments" in *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability* (1951), vol. 2, pp. 93–101.
22. B. Brooks, S. Du, A. Haberman, "Robust predictions with bounded information" (Working Paper, 2004).
23. A. Caplin, M. Dean, "Behavioral implications of rational inattention with Shannon entropy" (NBER Working Paper No. 19318, 2013).
24. J. Miao, H. Xing, Dynamic discrete choice under rational inattention. *Econ. Theor.* **77**, 597–652 (2024).
25. M. D. Donsker, S. S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time, I–IV. *Commun. Pure Appl. Math.* **28**, 1–47 (1975).
26. P. Dupuis, R. Ellis, A., *Weak Convergence Approach to the Theory of Large Deviations* (John Wiley and Sons, Inc., 1997).
27. S. Arimoto, An algorithm for computing the capacity of arbitrary discrete memoryless channels. *IEEE Trans. Inf. Theory* **18**, 14–20 (1972).
28. R. Blahut, Computation of channel capacity and rate-distortion functions. *IEEE Trans. Inf. Theory* **18**, 460–473 (1972).
29. K. Kasa, Robustness and information processing. *Rev. Econ. Dyn.* **9**, 1–33 (2006).
30. F. H. Knight, *Risk, Uncertainty, and Profit* (Houghton Mifflin, 1921).
31. M. Dean, N. Neligh, Estimating information cost functions in models of rational inattention. *J. Econ. Theor.* **187**, 1–32 (2020).
32. M. Dean, N. Neligh, Experimental tests of rational inattention. *Eur. J. Polit. Econ.* **131**, 3415–3461 (2023).
33. B. Hébert, M. Woodford, Neighborhood-based information costs. *Am. Econ. Rev.* **111**, 3225–3255 (2021).
34. L. Pomatto, P. Strack, O. Tamuz, The cost of information: The case of constant marginal costs. *Am. Econ. Rev.* **113**, 1360–1393 (2023).
35. A. Bloedel, W. Zhong, "The cost of optimally-acquired information" (Working Paper, 2021).
36. A. Caplin, D. Martin, A testable theory of imperfect perception. *Econ. J.* **125**, 184–202 (2015).
37. T. Denti, M. Marinacci, A. Rustichini, Experimental cost of information. *Am. Econ. Rev.* **112**, 3106–3123 (2022).
38. F. Fabbri, "Rational inattention with ambiguity aversion" (Working Paper, 2024).