Robust Rationally Inattentive Discrete Choice

(Preliminary and not ready for circulation)

Lars Peter Hansen* Jianjun Miao† Hao Xing‡

January 7, 2022

Abstract

We introduce robustness to the rational inattention model with Shannon mutual information costs in a discrete choice setting when the decision maker is concerned about model misspecification/ambiguity. We provide necessary and sufficient conditions for the robust solution and develop numerical methods to solve it. We show that the decision maker slants their beliefs pessimistically toward worse outcomes. As a result, their choice behavior can be qualitatively different from that in the standard rational inattention model with risk aversion. We apply our model to three consumer problems and show that tests based on misspecified models can lead to type I errors.

Keywords: Rational Inattention, Robustness, Information acquisition

JEL Classifications: C6, D8, E2.

*Departments of Economics and Statistics and the Booth School, University of Chicago, lhansen@chicago.edu

†Department of Economics, Boston University, miaoj@bu.edu

‡Department of Finance, Boston University, haoxing@bu.edu

1 Introduction

People often make choices with limited information. People also pay limited attention to process different pieces of information. To study such decision problems, Sims (1998, 2003) introduces a rational inattention (RI) framework by modeling information costs using the Shannon (1948) entropy-based mutual information (henceforth the Shannon model). While Sims focuses on continuous choices in dynamic settings, Caplin and Dean (henceforth CD) (2013, 2015), Matějka and McKay (henceforth MM) (2015), and Caplin, Dean, and Leahy (henceforth CDL) (2019, 2020) adopt the RI framework to study discrete choice problems in static settings.

Discrete choice problems have wide applications in labor economics, industrial organization, macroeconomics, and political economy. Understanding how RI affects discrete choices is important in both theory and application. While the Shannon model is a standard approach in the literature, it may suffer from a problem of model 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). This assumption is similar to the rational expectations hypothesis that imposes a communism of models: the people being modeled, the econometrician, and nature share the same model, i.e., the same probability distribution over the state of the world (Hansen and Sargent (2010b)).

Model misspecification is not innocuous in the standard Shannon model because CDL (2019) show that different specifications of the prior can generate very different consideration sets (sets of alternatives chosen with positive probabilities). Once acknowledging that all models are approximations to the real world, the DM may be concerned about model misspecification. They then want to seek robust decision making under model uncertainty. The goal of our paper is to robustifying the Shannon model by integrating the robust control approach of Hansen and Sargent (2001, 2008) and Hansen and Miao (2018) into the RI framework in a static discrete choice setting.

The robust control approach has a decision theoretic foundation axiomatized by Klibanoff, Marinacci, and Mukerji (2005), Maccheroni, Marinacci, and Rustichini (2006), and Strzalecki (2011). Intuitively, the DM does not have a single prior about the state of the world. They have other priors in mind that may be different from a baseline or reference 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. This parameter can also be interpreted as an ambiguity aversion parameter related to the literature on ambiguity motivated by the Ellsberg paradox (Ellsberg (1961)). The DM adopts the worst-case prior that minimizes the sum of expected utility evaluated at that prior and the associated penalty cost. In this paper we adopt both robust control and ambiguity aversion interpretations interchangeably.

We make three contributions to the literature. First, we provide necessary and sufficient conditions to characterize the solution to the robust RI problem. Mathematically, this is a minimax problem that admits a zero-sum game interpretation. We show that the worst-case prior distorts the baseline prior by putting more weight on the state with lower payoffs. These payoffs are determined by a standard RI problem as in MM (2015), CD (2013, 2015), and CDL (2019) when the worst-case prior is taken as given.

Applying the posterior-based approach of CD (2013) and CDL (2019, 2020), we provide a geometric interpretation of the robust RI solution. We find that the invariant likelihood ratio (ILR) property discussed in those papers also holds for the robust Shannon model. We argue that testing the ILR property using misspecified priors can generate Type I errors. CD (2013) also show that posterior beliefs are invariant to local changes in prior beliefs – the locally invariant posteriors (LIP) property. We find a similar invariance property but in a different sense: as long as local changes of the baseline prior still leave the worst-case prior in the convex hull of the tangency points of the supporting hyperplane and net utilities, the optimal posteriors do not change.

Our second contribution is to propose an efficient algorithm to solve the robust RI problem numerically. This algorithm generalizes the Arimoto-Blahut algorithm in the engineering literature (Arimoto (1972) and Blahut (1972)) and is an application of the general block coordinate descent (BCD) algorithm (Bertsekas (2016)). The BCD algorithm solves general multivariate optimization problems in which the constraint sets have a Cartesian product property. Our robust RI problems can be transformed into this form.

Our final contribution is to apply our model to the three consumer problems studied by CDL (2019). We show that introducing robustness to these problems generates some interesting and surprising results not available in the literature.

In consumer problem 1, the consumer chooses between a number of alternatives, one of which is of high quality and the remainder are of low quality. The identity of the high quality alternative is unknown ex ante, but can be learned by the consumer. We show that, when the consumer is concerned about model misspecification, consideration sets are determined by a threshold strategy as in CDL (2019): the consumer will consider only alternatives with a baseline prior probability of being high quality which is above an endogenously determined threshold. Alternatives below this threshold will never be chosen even though there is a chance that this set includes the high-quality good. Unlike in CDL (2019), the threshold in our model depends on the ambiguity aversion parameter and ambiguity aversion lowers this threshold. The size of the consideration set increases with the degree of ambiguity aversion. When the consumer becomes infinitely ambiguity averse as in the maxmin expected utility model of Gilboa and Schmeidler (1989), all alternatives are in the consideration set and the worst-case prior becomes a uniform distribution. Moreover, each alternative is chosen with an equal probability for any positive information cost.

In consumer problem 2 we assume that the payoffs to different alternatives are independent under the baseline prior. We show that they are not independent under the worst-case prior. Nevertheless, the consideration set is still determined by a threshold strategy. The ranking of alternatives is determined by the expectation of a convex transformation of the payoffs, taken with respect to the worst-case prior, but not the baseline prior as in CDL (2019). Ambiguity aversion lowers this expectation, causing more alternatives to be included in the consideration set. Using numerical examples we show that ambiguity aversion acts like risk aversion in determining the consideration set for any fixed information cost parameter: an ambiguity averse DM wants to consider a safe alternative.

In consumer problem 3 we study an example with arbitrarily correlated payoffs under the baseline prior. There are three options: a safe option with constant payoffs, two risky options with negatively correlated payoffs. In this case no simple threshold strategy determines the consideration set. There is a hedging motive even under risk or ambiguity neutrality: the rationally inattentive consumer chooses the two risky options only. Introducing risk aversion or ambiguity aversion changes the result dramatically. With risk aversion only, the consumer will choose the safe option and one of the risky options only.

By contrast, with ambiguity aversion only, the consumer will always choose the two risky options only. As the degree of ambiguity aversion gradually rises to a sufficiently high level, the consumer reverses their unconditional choice probabilities of the two risky options: the consumer is more likely to choose the risky option with lower (higher) payoffs, when their degree of ambiguity aversion is sufficiently high (low). Intuitively, the consumer's specification doubts induce them to slant their beliefs pessimistically toward worse outcomes. They are willing to choose the worse risky option with a higher probability when they believe that this option is much more likely to pay off.

Based on consumer problem 3, we study an example to illustrate that an ambiguity averse DM may not select an action with a higher unconditional probability if that action is improved in the sense of first-order-stochastic dominance under the baseline prior. This result is in contrast to Proposition 3 of MM (2015). The intuition is that the ambiguity averse DM slants their beliefs toward to the state with worse outcomes and hence may select another action that gives a higher payoff in that state with a higher unconditional probability.

Our paper is related to two strands of literature. First, in addition to the aforementioned studies, our paper is related to a growing literature on RI. Most studies on discrete choices under RI are set up in static models. Steiner, Stewart, and Matějka (2017) extend the static model of MM (2015) to a dynamic setting. Woodford (2009) studies a dynamic binary choice problem under RI (the problem of a firm that decides each period whether to reconsider its price). CD (2013) and CDL (2020) extend the static Shannon model to general information cost functions. Miao and Xing (2020) extend their models to a dynamic setting using the posterior-based approach of CD (2013).

Most existing work on RI has focused on models with continuous choices, which are typically set up in the linear-quadratic-Gaussian framework (e.g., Peng and Xiong (2006), Luo (2008), Maćkowiak and Wiederholt (2009), Mondria (2010), Van Nieuwerburgh and Veldkamp (2010), Kőszegi and Matějka (2019), Miao (2019), and Miao, Wu, and Young (2019)). Jung et al (2018) show that rationally inattentive agents can constrain themselves voluntarily to a discrete choice set even when the initial choice set is continuous. See Sims (2010) and Maćkowiak, Matějka and Wiederholt (2018) for surveys and references cited therein.

Our paper is also related to the large literature on robustness and ambiguity. Most studies in this literature focus on applications in finance and macroeconomics with continuous choices. Instead of citing all relevant papers, we refer the reader to Epstein and Schneider (2010) and Hansen and Sargent (2008, 2010b) for surveys. Here we mention most closely related papers by Hansen (2007), Hansen and Sargent (2007, 2010a), and Hansen and Miao (2018) who combine robustness and learning. These papers introduce robustness to both models given hidden states and priors about these states. In this case there are two pessimistic belief adjustments. In both robust control and RI models, Shannon entropy is an important ingredient.

To the best of our knowledge, our paper is the first that introduces robustness to the Shannon model in a discrete choice setting. Previous applications of robust control models often suffer from the critique that ambiguity aversion and risk aversion have qualitatively similar effects on behavior. By contrast, we show that ambiguity aversion and risk aversion can generate qualitatively different choice behaviors under RI.

2 Model

We first introduce the model setup and then provide a recommendation lemma to simplify the decision problem.

2.1 Setup

We start by presenting the standard Shannon model in a discrete choice setting studied by CD (2013), MM (2015), and CDL (2019). Uncertainty is represented by a finite state space $X = \{x_1, x_2, ..., x_M\}$ and a prior $\mu \in \Delta(X)$, where $\Delta(Y)$ denotes the set of probability distributions on any finite set Y. We use a bold letter $\mathbf{x}$ to denote a random variable and a normal letter x to denote its realization.

The DM does not observe the state, but can acquire a signal about the state. Based on the signal they select an element from a finite action set A to maximize expected utility of $u(x, a)$, where $u : X \times A \to \mathbb{R}$ is a bounded function. They can also choose an optimal information structure by paying a cost modeled by the Shannon mutual information.

Formally, let S denote a finite set of signal realizations. For simplicity suppose that $|A| = |S|$. A strategy is a pair (d, σ) 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 $\sigma : S \to A$, which gives an action $a = \sigma(s)$ when observing a signal s.

Let Σ denote the set of all strategies. An information strategy d and the prior μ induce a joint distribution $d \otimes \mu$ over $X \times S : (d \otimes \mu)(x, s) \equiv d(s|x)\,\mu(x)$. Given this joint distribution, we can define the mutual information as

$$I_{d\otimes\mu}(\mathbf{x}; \mathbf{s}) = \sum_{x,s} d(s|x)\,\mu(x) \ln \frac{d(s|x)}{\sum_x d(s|x)\,\mu(x)} = H(\mu) - H(\mu(\cdot|\mathbf{s})), \tag{1}$$

where $(\mathbf{x}, \mathbf{s})$ is a random vector on $X \times S$ with distribution $d \otimes \mu$, $\mu(x|s)$ denotes the posterior given by

$$\mu(x|s) = \frac{\mu(x)\,d(s|x)}{\sum_x \mu(x)\,d(s|x)} \quad \text{if } \nu(s) \equiv \sum_x \mu(x)\,d(s|x) > 0,$$

and $H(\mu)$ $(H(\mu(\cdot|\mathbf{s})))$ denotes the Shannon (conditional) entropy defined as

$$H(\mu) \equiv -\sum_x \mu(x) \ln \mu(x)\,, \;\; H(\mu(\cdot|\mathbf{s})) = -\sum_s \nu(s)\,H(\mu(\cdot|s))\,.$$

Equation (1) gives two equivalent definitions of the mutual information. See Cover and Thomas (2006) for a textbook introduction.

Entropy $H(\mu)$ measures the amount of prior uncertainty. After observing signals $\mathbf{s}$, the amount of uncertainty becomes $H(\mu(\cdot|\mathbf{s}))$. The mutual information measures the reduction of uncertainty after observing signals. Since entropy is a concave function, the mutual information is always nonnegative.

We are ready to state the standard RI problem in the signal-based form:

Problem 1 *(Signal-based standard RI problem)*

$$V(\mu) \equiv \max_{(d,\sigma)\in\Sigma} \mathbb{E}_{d\otimes\mu}\left[u(\mathbf{x}, \sigma(\mathbf{s}))\right] - \lambda I_{d\otimes\mu}(\mathbf{x}; \mathbf{s})\,,$$

where $\lambda > 0$ denotes the shadow cost of information and $\mathbb{E}_{d\otimes\mu}$ is an expectation operator given distribution $d \otimes \mu$.

We use $V(\mu)$ to denote the value function for the standard RI problem. As is well known in the literature, the solution to this problem is sensitive to the specification of the prior μ. If the decision maker is concerned about model misspecification, they want to seek robust decision making.

We follow Hansen and Sargent (2008) and Hansen and Miao (2018) to introduce robustness. Let $\widehat{\mu} \in \Delta(X)$ denote a baseline prior. For simplicity, suppose that $\widehat{\mu}(x) > 0$ for any $x \in X$. This prior may not be the 'true' distribution of the state of nature. The DM thinks another prior $\mu \in \Delta(X)$ as possible. Whenever μ and $\widehat{\mu}$ are different, the DM is subject to a penalty measured by the relative entropy

$$R(\mu||\widehat{\mu}) \equiv \sum_{x} \mu(x) \ln \frac{\mu(x)}{\widehat{\mu}(x)}.$$

We can then formulate the robust RI problem in the signal-based form:

Problem 2 *(Signal-based robust RI problem)*

$$W(\widehat{\mu}) \equiv \min_{\mu \in \Delta(X)} \max_{(d,\sigma) \in \Sigma} \mathbb{E}_{d \otimes \mu}[u(\mathbf{x}, \sigma(\mathbf{s}))] - \lambda I_{d \otimes \mu}(\mathbf{x}; \mathbf{s}) + \theta R(\mu||\widehat{\mu}), \tag{2}$$

where $\theta > 0$ 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 (2) reflects the DM's pessimistic behavior. They want to consider the worst-case scenario. The parameter θ plays a critical role. There are several interpretations. First, we may assume that the deviation of μ from $\widehat{\mu}$ is restricted by the constraint $R(\mu||\widehat{\mu}) \leq \eta$ for some $\eta > 0$. The DM solves the following so called constraint problem (Hansen and Sargent (2008)):

$$\min_{\mu \in \Delta(X)} \max_{(d,\sigma) \in \Sigma} \mathbb{E}_{d \otimes \mu}[u(\mathbf{x}, \sigma(\mathbf{s}))]$$

subject to $R(\mu||\widehat{\mu}) \leq \eta$. Then θ is related to the Lagrange multiplier of this constraint. Treating θ as a fixed penalty parameter, the minimization problem in (2) is a multiplier problem (Hansen and Sargent (2008)). We can interpret θ as the shadow cost of belief distortions. A larger value of θ means that the DM is penalized more when their beliefs deviate from $\widehat{\mu}$ and thus the DM will trust the baseline prior $\widehat{\mu}$ more.

Second, we follow Maccheroni, Marinacci, and Rustichini (2006) and interpret $1/\theta$ as an ambiguity aversion parameter. The idea is that the DM has ambiguous beliefs about the state of nature and has a set of priors in mind with $\widehat{\mu}$ being one of them. The DM with a larger value of $1/\theta$ is more ambiguity averse. Maccheroni, Marinacci, and Rustichini (2006) provide a decision theoretic foundation for this interpretation. In this paper we use both interpretations interchangeably.

2.2 Revelation Principle

Problem 2 is difficult to analyze. To simplify the analysis, we transform it into an equivalent full information problem in which the choice is random and represented by a (state-dependent) choice rule.

Formally, we say 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 $x \in X$, if

$$p(a|x) = \Pr(\sigma(s) = a|x) = \sum_{\{s \in S: \sigma(s)=a\}} d(s|x). \tag{3}$$

Then p and μ induce 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 can define the mutual information as

$$I_{p\otimes\mu}(\mathbf{x}; \mathbf{a}) = \sum_{x,a} p(a|x)\mu(x) \ln \frac{p(a|x)}{\sum_x p(a|x)\mu(x)} = H(\mu) - H(\mu(\cdot|\mathbf{a})), \tag{4}$$

where $H(\mu(\cdot|\mathbf{a}))$ is the conditional entropy defined as

$$H(\mu(\cdot|\mathbf{a})) = \sum_a q(a) H(\mu(\cdot|a)).$$

Here $\mu(\cdot|a) \in \Delta(X)$ and $q \in \Delta(A)$ denote the posterior and marginal distributions induced by the joint distribution $p \otimes \mu$:

$$q(a) = \sum_x p(a|x)\mu(x), \tag{5}$$

$$\mu(x|a) = \frac{p(a|x)\mu(x)}{\sum_x p(a|x)\mu(x)} \text{ if } q(a) > 0. \tag{6}$$

Conversely, a choice rule $p \in \Delta(A|X)$ can induce a strategy (d, σ). Specifically, let $|S| = |A|$ and fix any bijection $\phi : A \to S$. Define

$$d(s|x) = p(a|x), \ \sigma(s) = a, \ \text{for } s = \phi(a). \tag{7}$$

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

$$J(\widehat{\mu}) \equiv \min_{\mu \in \Delta(X)} \max_{p \in \Delta(A|X)} \mathbb{E}_{p\otimes\mu}[u(\mathbf{x}, \mathbf{a})] - \lambda I_{p\otimes\mu}(\mathbf{x}; \mathbf{a}) + \theta R(\mu || \widehat{\mu}), \tag{8}$$

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

We have the following recommendation lemma or the revelation principle similar to MM (2015).

Lemma 1 *Let (μ^*, d^*, σ^*) solve Problem 2. Then (d^*, σ^*) generates a choice rule p such that (μ^*, p) solves Problem 3. Conversely, let (μ^*, p^*) solve Problem 3. Then p^* induces a strategy (d, σ) such that (μ^*, d, σ) solves Problem 2. Moreover, $W(\widehat{\mu}) = J(\widehat{\mu})$.*

By this lemma, we will focus our analysis on Problem 3. Before providing solutions to this problem in the next section, we first present some special limiting cases. First, when $\theta = \infty$, Problem 3 is reduced to the standard RI problem. In this case the worst-case prior is the baseline prior. Second, when $\theta = 0$, Problem 3 is reduced to the following one:

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

This is the extreme case in which the DM thinks any prior in the feasible set $\Delta\left(X\right)$ 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.

3 Model Solution

We first present the solution to the standard RI problem with $\theta = \infty$ and then analyze the robust RI problem with $0 < \theta < \infty$. We also discuss algorithms to solve these problems.

3.1 A Benchmark

MM (2015) present necessary conditions for the standard RI problem. CD (2013) and CDL (2019) characterize the complete necessary and sufficient conditions and argue that these conditions are important for identifying the consideration set.

The consideration set is defined as the set of actions chosen with positive probabilities. Formally, for any $q \in \Delta\left(A\right)$, the associated consideration set is defined as $B\left(q\right) = \left\{a \in A : q\left(a\right) > 0\right\}$.

Proposition 1 *Fix a prior $\mu \in \Delta\left(X\right)$. The choice rule $p^* \in \Delta\left(A|X\right)$ is an optimal solution to the standard RI problem if and only if it satisfies the following conditions*

$$p^*\left(a|x\right) = \frac{q^*\left(a\right) \exp\left(u\left(x,a\right)/\lambda\right)}{\sum_b q^*\left(b\right) \exp\left(u\left(x,b\right)/\lambda\right)} \text{ for } \mu\left(x\right) > 0, \tag{9}$$

and

$$\sum_x \frac{\mu\left(x\right) \exp\left(u\left(x,a\right)/\lambda\right)}{\sum_b q^*\left(b\right) \exp\left(u\left(x,b\right)/\lambda\right)} \leq 1 \text{ for all } a \in A, \tag{10}$$

with equality if $a \in B\left(q^*\right)$, where

$$q^*\left(a\right) = \sum_x \mu\left(x\right) p^*\left(a|x\right) \text{ for all } a \in A. \tag{11}$$

The value function is given by

$$V\left(\mu\right) = \sum_x \mu\left(x\right) v\left(x\right),\ v\left(x\right) \equiv \lambda \ln \sum_a q^*\left(a\right) \exp\left(u\left(x,a\right)/\lambda\right), \tag{12}$$

and $V\left(\mu\right)$ is convex and satisfies

$$\frac{\partial V(\mu)}{\partial \mu(x)} = v(x) - v(x_M),\ x = x_1, x_2, ..., x_{M-1}. \tag{13}$$

We may interpret the above solution in two ways. First, equation (9) shows that the optimal choice probabilities resemble the logit rule (MM (2015)). The action weights q^* shift the choice probabilities toward those actions that appeared to be good candidates a priori. These weights depend on the prior μ and the cost of information λ. Second, following Steiner, Stewart, and Matějka (2017), the weights q^* can be interpreted as a default rule. The DM's actual choice probabilities $p^*(a|x)$ differ from the default rule and tilts toward the action with a higher payoff. The default rule itself is endogenous.

By (9), we can compute the posterior given any chosen action $a \in B(q^*)$:

$$\mu(x|a) = \frac{\exp(u(x,a)/\lambda)\,\mu(x)}{\sum_b q^*(b)\exp(u(x,b)/\lambda)}. \tag{14}$$

In this case (10) holds with equality, meaning that the expression for $\mu(\cdot|a)$ given above is a valid probability distribution. This expression for the posterior implies the invariant likelihood ratio (ILR) property discussed in CD (2013):

$$\frac{\mu(x|a)}{\mu(x|b)} = \frac{\exp(u(x,a)/\lambda)}{\exp(u(x,b)/\lambda)}. \tag{15}$$

for any two distinct chosen actions $a, b \in B(q^*)$ and for any $x \in X$. By this property, we can compute the cost of information:

$$\lambda = \frac{u(x,a) - u(x,b)}{\ln \mu(x|a) - \ln \mu(x|b)}. \tag{16}$$

CD (2013) design experiments to test the ILR property using (16).

3.2 Robustness

We now turn to the analysis of the robust RI Problem 3 in which the prior μ may be misspecified. We need to solve the worst-case prior and the associated choice rule.

The critical step to solve Problem 3 is to transform it into the following simpler problem:

$$J(\widehat{\mu}) \equiv \min_{\mu \in \Delta(X)} \max_{p \in \Delta(A|X), q \in \Delta(A)} F(p,q,\mu), \tag{17}$$

where

$$F(p,q,\mu) = \sum_{x,a} \mu(x)\, p(a|x) \left(u(x,a) - \lambda \ln \frac{p(a|x)}{q(a)} \right) + \theta \sum_x \mu(x) \ln \frac{\mu(x)}{\widehat{\mu}(x)}. \tag{18}$$

Here q is a default rule and must satisfy (5). A somewhat surprising result is that the first-order condition with respect to q for (17) gives (5) (Blahut (1972, Theorem 4)). Thus we do not need to impose constraint (5) when solving problem (17). Then the choice set in (17) becomes a Cartesian product of compact convex sets. We can check that the objective function F is convex in μ and jointly concave in p and q. By the minimax theorem, we can exchange the extremization without affecting the optimized value:

$$J(\widehat{\mu}) = \max_{p\in\Delta(A|X), q\in\Delta(A)} \min_{\mu\in\Delta(X)} F(p,q,\mu). \tag{19}$$

The inside minimization problem gives the multiplier utility model of Hansen and Sargent (2008). The maxmin expected utility model of Gilboa and Schmeidler (1989) is the limit as $\theta \to 0$.

Now we can state the necessary and sufficient conditions for optimality as follows:

Proposition 2 *The pair $(p^*, \mu^*) \in \Delta(A|X) \times \Delta(X)$ is an optimal solution to the robust RI problem if and only if the following conditions are satisfied:*

$$p^*(a|x) = \frac{q^*(a)\exp(u(x,a)/\lambda)}{\sum_b q^*(b)\exp(u(x,b)/\lambda)} \quad \text{for } \mu^*(x) > 0, \tag{20}$$

$$\sum_x \frac{\mu^*(x)\exp(u(x,a)/\lambda)}{\sum_b q^*(b)\exp(u(x,b)/\lambda)} \leq 1 \quad \text{for all } a \in A, \tag{21}$$

with equality if $a \in B(q^*)$, where

$$q^*(a) = \sum_x \mu^*(x) p^*(a|x), \quad a \in A, \tag{22}$$

$$\mu^*(x) = \frac{\exp(-v(x)/\theta)\widehat{\mu}(x)}{\sum_x \exp(-v(x)/\theta)\widehat{\mu}(x)}, \quad x \in X. \tag{23}$$

The value function is given by

$$J(\widehat{\mu}) = -\theta \ln \sum_x \widehat{\mu}(x) \exp\left(-\frac{v(x)}{\theta}\right), \tag{24}$$

where $v(x)$ is given in (12).

By (12), we can interpret $v(x)$ as the ex post payoff (utility) derived from a standard RI problem for a given prior μ^*. Comparing Propositions 1 and 2 reveals that the robust choice and default rules differ from those in the standard model in that the prespecified prior μ is replaced by the worst-case prior μ^*. The worst-case prior tilts the baseline prior $\widehat{\mu}$ by putting more weights on the state with a lower payoff $v(x)$. At the same time, $v(x)$ is generated from a standard RI problem the given prior μ^*. The resulting value function is smaller than that in the standard model $J(\widehat{\mu}) < V(\widehat{\mu})$. The loss in utility reflects the cost of model misspecification or ambiguity aversion.

We can interpret the solution to the robust RI Problem 3 as an ex post Bayesian solution. Once the worst-case prior μ^* is determined, the DM solves a standard RI problem taking the prior μ^* as given.

Importantly, the ex post payoff $v(x)$ in state x depends on the default rule q^* by (12), which in turn is endogenously chosen from a standard RI problem given the worst-case prior μ^*. In this sense, both μ^* and (p^*, q^*) are jointly determined in the robust RI problem.

By (20), the robust posterior $\mu^*(x|a)$ also satisfies (14). The difference is that the exogenous prior μ is replaced by the worst-case prior μ^* and the default rule q^* is derived from the robust RI solution. The ILR property (15) also holds for the robust posterior $\mu^*(x|a)$ and $\mu^*(x|b)$ for any $x \in X$ and $a, b \in B(q^*)$. Moreover, equation (16) also holds.

There is an alternative interpretation of the robust RI Problem 3 and its solution. We may view this problem as a zero-sum game (Hansen and Sargent (2008)). The minimizing player chooses a prior μ first and then the maximizing player chooses a choice rule p next. By the minimax theorem as in (19), we can change the timing protocol without affecting the solution by assuming that the maximizing player chooses first and then the minimizing player chooses next. The robust RI solution can then be interpreted as a Nash equilibrium. This equilibrium is the same as the Stacklberg equilibrium for the zero-sum game.

3.3 Posterior-Based Approach

We can rewrite the robust RI problem as

$$J(\widehat{\mu}) = \max_{\mu \in \Delta(X)} V(\mu) + \theta R(\mu || \widehat{\mu}),$$

where $V(\mu)$ is the value function for the standard RI problem with the prior μ. Given $V(\mu)$, this is a robust control problem studied by Hansen and Sargent (2008) and Hansen and Miao (2019).

CD (2013) and CDL (2019) propose a posterior-based approach to solve the standard Shannon model for any given prior μ and determine $V(\mu)$:

$$V(\mu) = \max_{q \in \Delta(A), \mu(\cdot|\cdot) \in \Delta(X|A)} \sum_a q(a) N^a(\mu(\cdot|a)) - \lambda H(\mu) \tag{25}$$

subject to

$$\mu(x) = \sum_a \mu(x|a) q(a), \ x \in X, \tag{26}$$

where N^a denotes the net utility function associated with action a defined as

$$N^a(\mu(\cdot|a)) \equiv \sum_x \mu(x|a) u(x,a) + \lambda H(\mu(\cdot|a)).$$

We can then reformulate the choice-based robust RI problem as an equivalent posterior-based problem of choosing $\mu \in \Delta(X), q \in \Delta(A)$, and $\mu(\cdot|\cdot) \in \Delta(X|A)$. Once q and $\mu(\cdot|\cdot)$ are determined, we can recover the choice rule

$$p(a|x) = \frac{q(a)\mu(x|a)}{\mu(x)} \text{ if } \mu(x) > 0.$$

Notice that the net utility function $N^a(\mu(\cdot|a))$ is concave in $\mu(\cdot|a)$ due to the concavity of entropy $H(\mu(\cdot|a))$. However, the objective function in (25) is not jointly concave in q and $\mu(\cdot|\cdot)$.

Thus one has to solve a concavification problem. This posterior-based approach has a nice geometric interpretation of the solution, which helps understand economic intuition. Specifically, for any given prior μ, one graphs the net utilities associated with all actions and finds the point on the convex hull directly above the prior μ; the optimal posteriors are the points of tangency of the supporting hyperplane at this point and the net utility functions (see Figure 1). This generates the value function $V(\mu)$ for the standard RI problem. On top of this problem, the robust strategy determines the worst-case prior μ^* that minimizes $V(\mu)$ plus a penalty cost measured by the entropy relative to the baseline prior $\widehat{\mu}$. The robust posteriors and default rule are optimal relative to μ^*.

[Insert Figure 1 Here]

By Propositions 1 and 2, we have

$$V(\mu^*) = \sum_x \mu^*(x)\, v(x) = \sum_x \mu^*(x)\, \widehat{V}(x) + \lambda \sum_x \mu^*(x) \ln \mu^*(x),$$

where $\widehat{V}(x) \equiv v(x) - \lambda \ln \mu^*(x)$. Miao and Xing (2019) show that $\widehat{V}(x)$ is the height of the supporting hyperplane at the point with $\mu^*(x) = 1$ and $\mu^*(x') = 0$ for all $x' \neq x$. Replacing $v(x)$ in (23), we can show that

$$\mu^*(x) = \frac{\widehat{\mu}(x)^{\frac{\theta}{\lambda+\theta}} \exp\left(-\frac{1}{\lambda+\theta}\widehat{V}(x)\right)}{\sum_{x'} \widehat{\mu}(x')^{\frac{\theta}{\lambda+\theta}} \exp\left(-\frac{1}{\lambda+\theta}\widehat{V}(x')\right)}.$$

Thus the worst-case prior μ^* puts more weight on a lower payoff $\widehat{V}(x)$. The payoff $\widehat{V}(x)$ is related to $v(x)$ and has a better geometric interpretation. Notice that $\widehat{V}(x)$ or $v(x)$ tends to be low when $u(x,a)$ is low for a chosen action a.

The left panel of Figure 1 illustrates the case with two states x_1 and x_2 and one action a. For the standard Shannon model, the tangency point of the supporting plane lies directly above the prior $\widehat{\mu}$ so that the prior is the optimal posterior. For the robust Shannon model, the optimal posterior is the same as the worst-case prior μ^*. Since $q^*(a) = 1$ with only one action a, we have $v(x) = u(x,a)$ and

$$\mu^*(x) = \frac{\exp(-u(x,a)/\theta)\, \widehat{\mu}(x)}{\sum_x \exp(-u(x,a)/\theta)\, \widehat{\mu}(x)}, \quad x = x_1, x_2.$$

In the figure, since $u(x_1,a) < u(x_2,a)$, we have $\mu^*(x_1) > \widehat{\mu}(x_1)$ and the tangent hyperplane tilts downward toward the state x_1. The right panel of Figure 1 illustrates the case with two states and two actions. The worst-case prior $\mu^*(x)$ lies in the convex hull of the tangency points, which give the optimal posteriors $\mu^*(x_1|a)$ and $\mu^*(x_1|b)$.

As in CDL (2019), for an action to be in the consideration set, its net utility must touch the supporting hyperplane. Except in cases of indifference, this means that the net utility function associated with this action would pierce the hyperplane associated with a problem that did not include this action. This is more likely if the net utility associated with this action is higher (i.e. the payoffs are higher) or the hyperplane is lower (i.e. the payoffs to the other actions are lower). The latter case reflects a hedging motive: an action is more likely to be considered when it pays off more in states in which other actions pay off less.

For the standard Shannon model, CD (2013) establish a locally invariant posteriors (LIP) property; that is, optimal posterior distributions are locally invariant to changes in priors in the convex hull of the optimal posteriors. This result is intuitive using the geometric interpretation discussed earlier: local changes of priors in the convex hull of the tangency points of the hyperplane and net utility functions do not change these tangency points. In our robust RI model, these priors are not the primitives of the model. They are the worst-case priors endogenously derived from a robust control problem given a baseline prior $\widehat{\mu}$ and a penalty parameter θ. Similar to CD (2013), we have the following invariance result. The proof is similar to that of Corollary 1 in CD (2013) and hence is omitted here.

Corollary 1 *Let B and $\mu(\cdot|\cdot)$ be the optimal consideration set and posterior distribution for the robust RI problem with the baseline prior $\widehat{\mu}$ and the penalty parameter θ. If $B' \subset B$ and $\mu(\cdot|a) = \mu'(\cdot|a)$ for any $a \in B'$, then B' and $\mu'(\cdot|a)$ are optimal for a robust RI problem with some baseline prior $\widehat{\mu}'$ and penalty parameter θ'.*

It merits emphasis that the LIP of CD (2013) does not hold in our robust RI model. More specifically, our invariance result only holds for a combination of $\widehat{\mu}'$ and θ'. Local changes of baseline priors $\widehat{\mu}$ in the convex hull of the tangency points may affect the optimal posteriors. But a local change of $\widehat{\mu}$ such that the associated worst-case prior μ^* remains in that convex hull will not change the optimal posteriors. See the right panel of Figure 1 for the intuition and Section 4.1 for a specific example.

CDL (2019) establish a converse result that finds exogenous priors associated with any given consideration set. We are unable to provide a similar result. In our model, a baseline prior and a robustness parameter are primitives. Any given consideration set must be generated by both some baseline prior and some robustness parameter. In Section 4 we will study some examples that permit a clear comparative statics analysis.

3.4 Numerical Methods

There is no closed-form solution in general even for the standard RI problem. Arimoto (1972) and Blahut (1972) develop the following efficient algorithm to find numerical solutions:

1. Start with a guess $p^{(0)} \in \Delta(A|X)$ with $p^{(0)}(a|x) > 0$ for all $(x, a) \in X \times A$.

2. Given $p^{(k)}$, compute

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

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

3. Iterate over $k \geq 0$ until convergence.

To understand the intuition for this algorithm, we combine the above two equations to yield

$$q^{(k+1)}(a) = \left[\sum_x \frac{\mu(x) \exp(u(x,a)/\lambda)}{\sum_a q^{(k)}(a) \exp(u(x,a)/\lambda)}\right] q^{(k)}(a).$$

The term in brackets is the left side of (10) and determines if $q^{(k+1)}(a)$ rises or falls. Because of condition (10), the algorithm can converge to two possible limit points. One is a positive limit with $q(a) > 0$ and the term in brackets is equal to one, a case that includes $q(a) = 1$. The other limit is such that $q(a) = 0$, which happens when (10) holds with strict inequality.

The Arimoto-Blahut 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 coordinate taken in cyclic order. Under some convexity assumptions, the algorithm converges (Bertsekas (2016)). 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 $\left(q^{(k)}, p^{(k)}, \mu^{(k)}\right)$, compute

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

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

$$\mu^{(k+1)}(x) = \frac{\exp\left(-v^{(k+1)}(x)/\theta\right) \widehat{\mu}(x)}{\sum_x \exp\left(-v^{(k+1)}(x)/\theta\right) \widehat{\mu}(x)},$$

where

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

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

We will apply this algorithm to derive numerical solutions in the next section.

4 Applications

In this section we revisit the three consumer choice problems studied by CDL (2019). We introduce robustness to these problems and study the implications of model uncertainty for the consideration set and choice behavior.

4.1 Consumer Problem 1: Finding the Good Alternative

We begin by a canonical consumer problem that permits a closed-form solution. The consumer is faced with a range of possible goods identified as a set $A = \{1, 2, ..., M\}$. One of these options is good. The others are bad. The utilities of the good and bad options are u_G and u_B respectively, with $u_G > u_B$. We define the state space to be the same as the action space $X = A$, with the interpretation that state $i \in X$ is the state in which option i is of high quality and all others are of low quality. Thus

$$u(x, a) = \begin{cases} u_G & \text{if } x = a \\ u_B & \text{otherwise} \end{cases} . \tag{27}$$

Let $\widehat{\mu}(i)$ be the baseline prior probability that option i yields the good prize. Without loss of generality, we order states according to $\widehat{\mu}(i) \geq \widehat{\mu}(i+1) \geq \widehat{\mu}(M) > 0$. The DM can acquire information about the state by paying entropy-based information cost. The DM also has concerns about model misspecification and seeks robust decision making that performs well when their prior μ may deviate from $\widehat{\mu}$. The decision problem can be formalized as Problem 3.

Let

$$\exp(u_B/\lambda) \equiv \overline{u} \text{ and } \exp(u_G/\lambda) \equiv \overline{u}(1+\delta).$$

We will show that the solution does not depend on $\overline{u}$.

Proposition 3 *For robust consumer problem 1, define $K = M$ if*

$$\widehat{\mu}(M) > \left[\frac{\sum_{i=1}^{M} \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta + M}\right]^{\frac{\lambda+\theta}{\theta}} . \tag{28}$$

Otherwise, define $K < M$ as the unique integer such that

$$\widehat{\mu}(K) > \rho \geq \widehat{\mu}(K+1), \tag{29}$$

where

$$\rho \equiv \left[\frac{\sum_{i=1}^{K} \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta + K}\right]^{\frac{\lambda+\theta}{\theta}} . \tag{30}$$

Then the worst-case prior is

$$\mu^*(x) = \frac{\rho^{\frac{\lambda}{\lambda+\theta}}\widehat{\mu}(x)^{\frac{\theta}{\lambda+\theta}}}{\rho^{\frac{\lambda}{\lambda+\theta}}\sum_{i=1}^K \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} + \sum_{i=K+1}^M \widehat{\mu}(i)}, \quad \textit{for } x = 1, \ldots, K, \tag{31}$$

$$\mu^*(x) = \frac{\widehat{\mu}(x)}{\rho^{\frac{\lambda}{\lambda+\theta}}\sum_{i=1}^K \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} + \sum_{i=K+1}^M \widehat{\mu}(i)}, \quad \textit{for } x = K+1, \ldots, M. \tag{32}$$

The optimal default rule is

$$q^*(a) = \frac{1}{\delta}\left[\left(\frac{\widehat{\mu}(a)}{\rho}\right)^{\frac{\theta}{\lambda+\theta}} - 1\right], \quad \textit{for } a = 1, \ldots, K, \tag{33}$$

$$q^*(a) = 0, \quad \textit{for } a = K+1, \ldots, M.$$

The consideration set is given by $B = \{1, 2, ..., K\}$. *The posterior distribution for any* $a \leq K$ *is*

$$\mu^*(x|a) = \begin{cases} \frac{(1+\delta)\rho}{\rho^{\frac{\lambda}{\lambda+\theta}}\sum_{i=1}^K \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} + \sum_{i=K+1}^M \widehat{\mu}(i)}, & x = a \leq K, \\ \frac{\rho}{\rho^{\frac{\lambda}{\lambda+\theta}}\sum_{i=1}^K \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} + \sum_{i=K+1}^M \widehat{\mu}(i)}, & x \neq a, x \leq K, \\ \mu^*(x), & x > K. \end{cases}$$

By this proposition, we can derive the following robust choice rule for $x = 1, ..., K$:

$$p^*(a|x) = \frac{1+\delta}{\delta}\left[\left(\frac{\widehat{\mu}(a)}{\rho}\right)^{\frac{\theta}{\lambda+\theta}} - 1\right]\left(\frac{\rho}{\widehat{\mu}(x)}\right)^{\frac{\theta}{\lambda+\theta}}, \text{ if } x = a$$

$$p^*(a|x) = \frac{1}{\delta}\left[\left(\frac{\widehat{\mu}(a)}{\rho}\right)^{\frac{\theta}{\lambda+\theta}} - 1\right]\left(\frac{\rho}{\widehat{\mu}(x)}\right)^{\frac{\theta}{\lambda+\theta}}, \text{ if } x \neq a \leq K,$$

$$p^*(a|x) = 0, \text{ if } a > K.$$

For $x > K$ and for all $a \in A$, $p^*(a|x) = q^*(a)$.

In the case of the uniform baseline prior $\widehat{\mu}(x) = 1/M$ for any $x \in X$, condition (28) is satisfied and hence $B = A$. Then we have $q^*(a) = 1/M$ and $\mu^*(x) = \widehat{\mu}(x) = 1/M$ for any $a \in A$ and $x \in X$. The decision problem is symmetric. There is no belief distortion because the ex post payoff $v(x)$ in each state is the same.

In general, as $\theta \to \infty$, deviating from the baseline model is increasingly costly and thus the worst-case prior is the baseline prior itself. In this case Proposition 3 converges to Theorem 1 of CDL (2019).

The robust solution is similar to that of CDL (2019). Specifically, the consideration set is determined by a threshold strategy: the consumer will consider only alternatives with a prior probability of being high quality that is above an endogenously determined threshold ρ. Alternatives below this threshold will never be chosen even though there is a chance that this set includes the high-quality good. Among considered alternatives, attention is allocated in such a way that the posterior probability of any alternative being of high-quality conditional on being chosen is the same, regardless of prior belief. Unlike the solution of CDL (2019), the ambiguity aversion parameter θ affects the endogenous threshold ρ and hence the consideration set. It also affects the default and choice rules.

To see the impact of this parameter, we establish the following comparative statics and limiting results.

Proposition 4 *For robust consumer problem 1, the number of chosen actions $|B|$ increases with $1/\theta$. Let $\mu(M) > 0$. Then, as $\theta \to 0$, $K \to M$, $\mu^*(x) \to 1/M$, $q^*(a) \to 1/M$, and*

$$\begin{aligned} p(a|x) &\to \frac{1+\delta}{\delta+M} \text{ if } x = a, \\ p(a|x) &\to \frac{1}{\delta+M} \text{ if } x \neq a, \end{aligned}$$

for any $x, a = 1, \ldots, M$.

This proposition shows that the consumer considers more options when they are more averse to model uncertainty or misspecification, i.e., when $1/\theta$ is larger. Intuitively, a more ambiguity averse consumer puts more weight to worse outcomes. As a result they are willing to consider more alternatives even though these alternatives are of low quality; that is, the threshold becomes lower. By contrast, CDL (2019) use numerical examples to show that the size of the consideration set decreases with the information cost parameter λ for the standard Shannon model.[1]

As $\theta \to 0$, the consumer becomes infinitely ambiguity averse as in the maxmin model of Gilboa and Schmeidler (1989). Then the consumer considers all alternatives as possible and puts equal weight on each alternative. As a result, the unconditional probability of choosing any alternative is equal. The conditional probability of making the right choice given any state is also the same.

We use Proposition 3 to illustrate the LIP property. Consider another baseline prior $\widetilde{\mu}$ that satisfies the assumptions of the proposition. Moreover, let $\widetilde{\mu}$ satisfy

$$\sum_{i=K+1}^{M} \widetilde{\mu}(i) = \sum_{i=K+1}^{M} \widehat{\mu}(i), \ \sum_{i=1}^{K} \widetilde{\mu}(i)^{\frac{\theta}{\lambda+\theta}} = \sum_{i=1}^{K} \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}.$$

Then the optimal posteriors $\mu^*(\cdot|a)$, for $a = 1, 2, ..., K$, remain the same under $\widetilde{\mu}$. However, this prior $\widetilde{\mu}$ may not lie in the convex hull spanned by $\mu^*(\cdot|a)$ for $a = 1, 2, ..., K$ in the sense that $\widetilde{\mu}(x) = \sum_a \mu^*(x|a) q(a)$ for some $q(a) \in \Delta(A)$. Instead this equation can hold when $\widetilde{\mu}(x)$ is replaced by the worst-case prior μ^*.

We use a numerical example to illustrate Propositions 3 and 4. Let $M = 3$, $\lambda = 1$, $\widehat{\mu}(1) = 0.5$, $\widehat{\mu}(2) = 0.35$, $\widehat{\mu}(3) = 0.15$, $u_G = 1$, and $u_B = 0$. Figure 2 plots the solutions for different values of θ. When θ is sufficiently large, the worst-case prior approaches the baseline prior and the robust

[1] We can prove this result using a similar method for Proposition 4 in the appendix.

solution approaches the standard solution for which only alternatives 1 and 2 are considered. When θ is sufficiently small, alternative 3 also enters the consideration set. The worst-case prior puts more weight to alternative 3. As θ approaches zero, the worst-case prior approaches a uniform distribution over all alternatives.

[Insert Figure 2 Here]

4.2 Consumer Problem 2: Independent Valuation

We next consider a consumer choice problem with independent valuation. The consumer is faced with the choice between a number of different alternatives, the values of which are uncertain, but independently distributed. For example, the consumer decides which of several different cars to buy. Each car has a distribution of possible utilities it can deliver, depending on its price, fuel efficiency, reliability, and so on. We make the assumption that the utility associated with one car is independent of that of any other under some baseline belief. The consumer faces model uncertainty and does not trust their baseline belief. They may consider other priors different from the baseline as possible. The consumer must decide which cars to consider, and what to learn about each considered car prior to purchase.

Formally, the action set is $A = \{1, 2, ..., M\}$. Let $Z \subset \mathbb{R}$ be a finite set of possible utility levels for all actions. We define the state space as $X = Z^M$. A typical state is therefore a vector of realized utilities for each possible action: $x = (x_1, ..., x_M)'$, where $x_i \in Z$ for all $i \in A$. The utility of a state/action pair is then given by $u(x, i) = x_i$. The assumption of independence under the baseline prior implies that there exist baseline probability distributions $\widehat{\mu}_1, ..., \widehat{\mu}_M \in \Delta(Z)$ such that, for every $x \in X$,

$$\widehat{\mu}(x) = \prod_{i=1}^{M} \widehat{\mu}_i(x_i). \tag{34}$$

Proposition 5 *For robust consumer problem 2 with independent valuation, there exists some constant c such that action i is chosen with $q^*(i) > 0$ if*

$$\mathbb{E}^*[\exp(u(\mathbf{x}, i)/\lambda)] \equiv \sum_{x \in X} \mu^*(x) \exp(x_i/\lambda) > c,$$

and $q^*(i) = 0$ otherwise, where $\mathbb{E}^*$ denotes the expectation operator with respect to the worst-case prior μ^*.

This proposition shows that the robust solution also involves a threshold strategy determining a consideration set of alternatives about which the consumer will learn and from which they will make their eventual choice. However, in this case, the threshold is in terms of the expectation of the normalized utilities $\exp(u(x, i)/\lambda)$, which are a convex transformation of the utility payoff

$u(x,i)$. Unlike the solution of CDL (2019), the expectation is taken with respect to the worst-case prior instead of the exogenously given baseline prior. Since an ambiguity averse consumer tilts their belief toward worse outcomes, this will lower their expected normalized utilities. Moreover, ambiguity aversion also affects the endogenous threshold c.

Unlike the baseline prior, the worst-case prior violates the independence property (34). Nevertheless, in the appendix we show that the worst-case prior satisfies a special independence property in the sense that the component x_j of the state associated with any unchosen action j is independent of the remaining components of the state. The worst-case belief only distorts the probability of the components associated with the chosen actions. This property is critical for the proof of Proposition 5.

As CDL (2019) argue, the identification of the consideration set is not trivial. We thus take a numerical example in their paper to illustrate the robust solution. Let $M=6$ and $Z=\{0,\ 5.5,\ 10\}$. There are six actions. Under the baseline belief, the first action has a value of 5.5 for sure and so is the "safe" option. The other five actions have an ex-ante 50% chance of having value 10 and a 50% chance of having value zero, and so can be seen as "risky" options. Formally, we have $\widehat{\mu}_1(x_1=5.5)=1$, $\widehat{\mu}_1(x_1=0)=\widehat{\mu}_1(x_1=10)=0$, $\widehat{\mu}_i(x_i=10)=0.5$, $\widehat{\mu}_i(x_i=0)=0.5$, and $\widehat{\mu}_i(x_i=5.5)=0$, for $i=2,3,...,6$. Let $\widehat{\mu}(x)=\Pi_{i=1}^{6}\widehat{\mu}_i(x_i)$ for $x=(x_1,...,x_6)\in Z^6$.

There could be three possible consideration sets: only the safe action, only the five risky actions, or all options. Which consideration set is optimal depends on the model parameter values. We will focus on the information cost parameter λ and the ambiguity aversion parameter θ. Before presenting numerical solutions, we first establish conditions under which the "risky only" consideration set is optimal.

Lemma 2 *For robust consumer problem 2 with independent valuation, let $\{1,2,..,N\}=B\subset A$ be the set of ex ante identical actions, i.e., $\widehat{\mu}_i(x_0)=\widehat{\mu}_B(x_0)\in(0,1)$, for all $x_0\in Z$ and $i\in B$. Then the default rule $q^*(i)=1/N$ for $i\in B$ is optimal if for any $j\notin B$,*

$$\sum_{x_j\in Z}\exp\left(x_j/\lambda\right)\widehat{\mu}_j\left(x_j\right)\leq\frac{1}{N}\left\{\mathbb{E}^*\left[\frac{1}{\sum_{n=1}^{N}\exp\left(\mathbf{x}_n/\lambda\right)}\right]\right\}^{-1},\tag{35}$$

where $\mathbb{E}^*$ is the expectation operator with respect to the worst-case prior μ^*.

This result is similar to Lemma 1 of CDL (2019). There are two differences: First, the expectation on the left-hand side of the inequality is taken with respect to the baseline prior $\widehat{\mu}_j$. Second, the expectation on the right-hand side is taken with respect to the worst-case prior μ^*. The proof relies on the special independence property of μ^* discussed earlier. By this property the marginal distribution for x_j is actually the baseline prior $\widehat{\mu}_j$ for any unchosen action j.

Under model uncertainty, condition (35) is more likely to be violated, compared to that in the standard RI model. Because the worst-case prior puts more weight on worse outcomes, ambiguity aversion lowers the expression on the right-hand side of the inequality in (35).

[Insert Figure 3 Here]

We now discuss numerical solutions for the example presented earlier. Figure 3 plots the default rules $q^*(a)$ against the information cost parameter λ for the standard ($\theta = \infty$) and robust ($\theta = 1$) Shannon models. This figure replicates Figure 2 of CDL (2019) when $\theta = \infty$ and shows that the default rule is not monotonic with λ. There are four regions of values for λ : (i) When λ is very small, all options are chosen with positive probabilities. (ii) As λ increases to a higher value, only risky actions are chosen as shown in Lemma 2. (iii) When λ is even higher, the safe action enters the consideration set. (iv) When λ is sufficiently high, only the safe action is chosen. This non-monotonic relation comes from the result that the ranking of alternatives is determined by the expectation of a convex transformation of the payoffs shown in Proposition 5.

By contrast, for the robust Shannon model with $\theta = 1$, Figure 3 shows that all actions are chosen when λ is small and only the safe action is chosen when λ is sufficiently large. In this case condition (35) in Lemma 2 is violated. The intuition is that an ambiguity averse DM wants to consider the low-payoff safe action to avoid model uncertainty even if the information cost λ is small.

In this sense ambiguity aversion is similar to risk aversion. To further explore the implications of these two attitudes toward uncertainty, we introduce risk aversion and set

$$u(x, i) = \frac{x_i^{1-\alpha}}{1-\alpha},\ i \in A,$$

where $\alpha > 0$ is the relative risk aversion parameter. Following CDL (2019), we set the payoff to the safe option to 5. The left panel of Figure 4 shows the consideration sets for many different values of the robustness parameter θ and the information cost parameter λ with risk neutral utility ($\alpha = 0$). The right panel of Figure 4 shows the consideration sets for many different values of the robustness parameter θ and risk aversion parameter α with fixed information cost $\lambda = 2$. These two panels show that ambiguity aversion is similar to risk aversion in the sense that they generate a qualitatively identical comparative statics result for the consideration sets. The difference is the quantitative impact. In the next subsection we will show that this result does not hold generally.

[Insert Figure 4 Here]

4.3 Consumer Problem 3: Correlated Valuation

We now study the general case in which the value of the different alternatives may be arbitrarily correlated. There will be no analytical results for this case. Thus we use numerical examples taken

Action \ State	x_1	x_2	$\mathbb{E}_{\widehat{\mu}}\left[e^{u(x,j)/\lambda}\right]$
a	0	15	2.74
b	6	0	1.41
c	5	5	1.65

Table 1: This table displays payoffs c_{ij} for each state $i \in X$ and each action $j \in A$. The baseline prior $\widehat{\mu}$ assigns equal probabilities to the two states. Set $\alpha = 0$ and $\lambda = 10$.

from CDL (2019) to illustrate the robust solution.

Suppose that there are two states $X = \{x_1, x_2\}$ and three actions $A = \{a, b, c\}$. The two states are equally likely under the baseline prior $\widehat{\mu}(x_1) = \widehat{\mu}(x_2) = 1/2$. Let

$$u(i,j) = \frac{c_{ij}^{1-\alpha}}{1-\alpha}, \quad i \in X, j \in A, \ \alpha > 0.$$

Table 1 presents the payoffs c_{ij} for $\lambda = 10$ and $\alpha = 0$.

Figure 5 presents the robust solutions for different values of θ. For all values of θ, the safe option c is never chosen, even though it dominates the risky option b ex ante. As discussed in CDL (2019), the DM considers the risky option b instead of the safe option c because option b helps the DM better learn about the true state of nature. Option b pays off in state 1 when option a does not pay off and provides a better hedge than the safe option. This means that the net utility associated with option b is more likely to pierce the hyperplane when only action a is available, as illustrated in Figure 1. In other words, the net utility associated with action c (not shown in Figure 1) is below the tangent hyperplane and will never be chosen.

When θ is sufficiently large, the robust solution approaches the standard solution in CDL (2019). In particular, option a is chosen with a higher unconditional probability than option b. Interestingly, when θ is sufficiently small, we get the reverse result even though option b appears to be worse than option a in terms of payoffs. The intuition is that the ambiguity averse DM shifts their belief toward the state with lower payoffs. As shown on the right panel of Figure 1, $\mu^*(x_1)$ rises when θ decreases because $\widehat{V}(x_1) < \widehat{V}(x_2)$. As long as $\mu^*(x_1)$ is in the convex hull of the optimal posteriors $\mu^*(x_1|a)$ and $\mu^*(x_1|b)$, these posteriors do not vary with θ. However, as $\mu^*(x_1)$ moves toward $\mu^*(x_1|b)$ when θ decreases, the weight $q^*(b)$ rises and eventually crosses 0.5 and becomes higher than $q^*(a)$.

[Insert Figure 5 Here]

Is the impact of ambiguity aversion similar to that of risk aversion? The answer is yes for the example with independent valuation in the previous subsection. Here we paint a totally different picture with correlated valuation. Figure 6 presents the solutions for different values of the risk aversion parameter α without robustness $\theta = \infty$.

[Insert Figure 6 Here]

In the risk-neutral case of $\alpha = 0$, the safe action c is never chosen as we have discussed above. As long as $\alpha > 0$, the risk averse DM will never choose option b. Instead, the DM will choose both a and c with positive probabilities when α is not too large. In this case the safe option c is a better hedge against risk. When α is sufficiently large, the DM will choose the safe option c only.

Comparing Figures 5 and 6, we find that the impact of ambiguity aversion is qualitatively different from that of risk aversion for the example studied here. The key difference is that ambiguity aversion manifest itself through a pessimistic shift of belief toward the state with low payoffs. Rationally inattentive choice given this pessimistic prior belief may be very different from that given risk aversion. Risk aversion reflects the aversion of a mean-preserving spread of payoffs from choices.

Notice that this result does not hold in the full information case ($\lambda = 0$). In this case the DM will choose the best option conditional on the realized state and belief distortions do not affect choice behavior. By contrast, with limited information ($\lambda > 0$), the rationally inattentive DM always makes 'mistakes' by choosing the worse risky option b with some positive probability. Beliefs about this uncertainty plays an important role. When the ambiguity averse DM believes the worse option b is more likely to pay off, they will choose option b with a higher probability than option a.

4.4 Non-monotonicity

Proposition 3 of MM (2015) shows that if an action becomes more attractive then the rationally inattentive DM will select that action with a higher unconditional probability. In this subsection we give an example based on consumer problem 3 to show that this result may not hold true when the DM is averse to model uncertainty.

Suppose that there are two states $X = \{x_1, x_2\}$ and two actions $A = \{a, b\}$ as illustrated on the right panel of Figure 1. Let $u(x_1, b) = 1$, $u(x_2, b) = 0$, $u(x_2, a) = 1 + \epsilon$, $u(x_1, a) = \epsilon$, for $\epsilon \in (0, 1)$. Let the baseline prior be $\widehat{\mu}(x_1) = \widehat{\mu}(x_2) = 0.5$. We are interested in the comparative statics result for a change of $\epsilon \in (0, 1)$. The key feature of this example is that action b gives a higher payoff in state x_1 than action a, but a lower payoff in state x_2. Moreover, the payoff in state x_2 for action a is higher than that in state x_1 for action b.

An increase in ϵ shifts up the net utilities associated with action a. Then the tangent line tilts downward in state x_1, resulting in a higher ratio $\widehat{V}(x_2)/\widehat{V}(x_1)$, and hence state x_1 becomes relatively worse. The ambiguity averse DM will put more weight on state x_1 and less weight on state x_2 in the worst-case prior. Thus the DM is more likely to choose action b that gives a higher payoff in state x_1. On the other hand, when taken the worst-case prior as given, the rationally inattentive DM is more likely to choose action a with improved payoffs as ϵ is increased, a result implied by Proposition 3 of MM (2015). The net effect depends on which force dominates.

Figure 7 presents the solutions for the robust RI problem with $\theta = 1$ and the standard RI problem with $\theta = \infty$. The top (bottom) panels present the case with $\lambda = 0.1$ ($\lambda = 10$). To gain intuition, consider the limiting case as $\lambda \to 0$. Then information becomes free and we have $p(b|x_1) = 1$ and $p(a|x_2) = 1$ so that $q^*(a) = \mu^*(x_2)$. Since $u(x_2, a) = 1 + \epsilon > u(x_1, b) = 1$, the ambiguity averse DM attaches a lower weight to state x_2 in the worst-case prior μ^*. Moreover, the weight $\mu^*(x_2) = q^*(a)$ decreases as ϵ increases. This result also holds true when the information cost λ is sufficiently small because the impact of ambiguity aversion dominates. By contrast, when λ is sufficiently large, information is so costly that the DM cannot perfectly distinguish between states x_1 and x_2. In this case, the DM will 'mistakenly' choose action a in state x_1 with a positive probability. The impact of rational inattention dominates so that $q^*(a)$ increases with ϵ as in MM (2015).

4.5 Implications for Experimental Studies

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 Martin (2015), CD (2015), and CDL (2020) make an important assumption that the DM, the econometrician, and nature share the same probability model as in the rational expectations hypothesis. Then the prior about the state of the world is taken as observable and exogenously given.

In this subsection we discuss the limits of this assumption and the implications if it is violated once model misspecification/ambiguity is taken into account. We use some numerical examples to illustrate that tests based on the rational expectations hypothesis can lead to type I error.

We start with the example based on consumer problem 1, which is also related to the experiments in CD (2013). Suppose that there are two states and two actions. The baseline prior $\widehat{\mu}$ is uniform. The DM is ambiguity averse and risk neutral. For this symmetric case, Proposition 3 implies that the worst-case prior μ^* is the same as the uniform baseline prior. Suppose that the state-dependent choice probabilities p from the experimental studies are generated by our robust Shannon model.

To test the Shannon model, CD (2013, 2015) suggest to use the choice probabilities p and the baseline prior $\widehat{\mu}$ to compute the so called revealed posteriors

$$\mu(x|a) = \frac{p(a|x)\,\widehat{\mu}(x)}{\sum_x p(a|x)\,\widehat{\mu}(x)} \text{ for } x \in X \text{ and } a \in A. \tag{36}$$

After specifying payoffs $u(x, a)$ as in (27), one has all the data to test the ILR property using (16). In particular, one can test whether the right-hand side of (16) is independent of the state x. In this symmetric case, this procedure is robust to model misspecification because the worst-case prior μ^*

is the same as the baseline prior $\widehat{\mu}$.

Next we modify the above example by assuming that the baseline prior is not uniform. Let $\widehat{\mu}(1) = 0.6$ and $\widehat{\mu}(2) = 0.4$. Let $\lambda = 5$, $\theta = 5$, $u_B = 0$, and $u_G = 10$. Then our robust Shannon model generated choice probabilities are

$$p(a = 1|x = 1) = 0.9061, \quad p(a = 1|x = 2) = 0.1502.$$

If we use the baseline prior $\widehat{\mu}$ to compute the revealed posteriors, we would have

$$\mu(x = 1|a = 1) = 0.9005, \quad \mu(x = 1|a = 2) = 0.1422$$

Now we can see that

$$\frac{u(1,1) - u(1,2)}{\ln \mu(1|1) - \ln \mu(1|2)} = 5.4180 \neq 4.6421 = \frac{u(2,1) - u(2,2)}{\ln \mu(2|1) - \ln \mu(2|2)}.$$

Thus one might conclude that the ILR property fails and hence reject the Shannon model. The problem of this procedure is that the prior is misspecified. One has to use the worst-case prior μ^* instead of $\widehat{\mu}$ to compute the revealed posteriors in (36). Then ILR still holds.

Finally, we suppose that the payoffs are asymmetric, but the baseline prior is still uniform. Let $u(1,1) = 10$, $u(2,2) = 5$, and $u(2,1) = u(1,2) = 0$. Let $\lambda = 5$ and $\theta = 5$. Then the model generated choice probabilities are

$$p(a = 1|x = 1) = 0.8516, \quad p(a = 1|x = 2) = 0.2222.$$

If we use the baseline prior to compute revealed posteriors, we would have

$$\mu(x = 1|a = 1) = 0.7930, \quad \mu(x = 1|a = 2) = 0.1602.$$

Again the ILR property would be falsely rejected in this case. If we use the worst-case prior to compute revealed posteriors, the ILR property will still hold.

5 Conclusion

The standard RI model with Shannon information costs assumes that the prior is exogenously given. We introduce robustness to such a model in which the DM is concerned about misspecification of priors. We characterize the robust solution and develop efficient numerical methods. We apply our model to three consumer problems and find that under RI risk aversion and ambiguity aversion can generate qualitatively different choice behaviors.

Appendix

A Proofs

Proof of Lemma 1: Because the entropy penalty term is the same in (2) and (8), it suffices to prove

$$\max_{p\in\Delta(A|X)} \mathbb{E}_{p\otimes\mu}\left[u\left(\mathbf{x},\mathbf{a}\right)\right]-\lambda I_{p\otimes\mu}\left(\mathbf{x};\mathbf{a}\right)=\max_{(d,\sigma)\in\Sigma} \mathbb{E}_{d\otimes\mu}\left[u\left(\mathbf{x},\sigma\left(\mathbf{s}\right)\right)\right]-\lambda I_{d\otimes\mu}\left(\mathbf{x};\mathbf{s}\right), \tag{A.1}$$

for any fixed prior $\mu\in\Delta\left(X\right).$

Given any strategy $(d,\sigma)\in\Sigma$, we can construct a choice rule p as in (3) and hence define $I\left(\mathbf{x};\mathbf{a}\right)$ as in (4). We will prove that $I_{p\otimes\mu}\left(\mathbf{x};\mathbf{a}\right)\leq I_{d\otimes\mu}\left(\mathbf{x};\mathbf{s}\right).$ By (1) and (4), we only need to show that $H(\mu(\cdot|\mathbf{a}))\geq H(\mu(\cdot|\mathbf{s}))$; that is,

$$\sum_a q(a)\,H(\mu(\cdot|a))\geq\sum_s\nu(s)\,H(\mu(\cdot|s)).$$

Since $a=\sigma\left(s\right)$, we have

$$\mu(x|a)=\sum_s\mu(x|s)\mathrm{Pr}(s|a),\ \ x\in X.$$

Since Shannon entropy $H:\Delta\left(X\right)\rightarrow\mathbb{R}$ is a concave function, it follows from Jensen's inequality that

$$H(\mu(\cdot|a))\geq\sum_s\mathrm{Pr}(s|a)H(\mu(\cdot|s)).$$

Multiplying both sides by $q(a)$ and summing over a, we obtain

$$\sum_a q(a)\,H(\mu(\cdot|a))\geq\sum_s\sum_a\mathrm{Pr}(s|a)q(a)\,H(\mu(\cdot|s))=\sum_s\nu(s)\,H(\mu(\cdot|s)),$$

as desired. Since $\mathbb{E}_{d\otimes\mu}\left[u\left(\mathbf{x},\sigma\left(\mathbf{s}\right)\right)\right]=\mathbb{E}_{p\otimes\mu}\left[u\left(\mathbf{x},\mathbf{a}\right)\right]$ by construction, we have

$$\mathbb{E}_{d\otimes\mu}\left[u\left(\mathbf{x},\sigma\left(\mathbf{s}\right)\right)\right]-\lambda I_{d\otimes\mu}\left(\mathbf{x};\mathbf{s}\right)\leq\ \mathbb{E}_{p\otimes\mu}\left[u\left(\mathbf{x},\mathbf{a}\right)\right]-\lambda I_{p\otimes\mu}\left(\mathbf{x};\mathbf{a}\right).$$

Let the strategy (d^*,σ^*) achieve the maximum of the problem on the right-hand side of (A.1). Let p be the induced choice rule. We then have

$$\max_{(d,\sigma)\in\Sigma}\ \mathbb{E}_{d\otimes\mu}\left[u\left(\mathbf{x},\sigma\left(\mathbf{s}\right)\right)\right]-\lambda I_{d\otimes\mu}\left(\mathbf{x};\mathbf{s}\right)\leq\ \mathbb{E}_{p\otimes\mu}\left[u\left(\mathbf{x},\mathbf{a}\right)\right]-\lambda I_{p\otimes\mu}\left(\mathbf{x};\mathbf{a}\right). \tag{A.2}$$

Conversely, given any choice rule p, we can construct a strategy (d,σ) as in (7). This construction implies

$$\mathbb{E}_{d\otimes\mu}\left[u\left(\mathbf{x},\sigma\left(\mathbf{s}\right)\right)\right]-\lambda I_{d\otimes\mu}\left(\mathbf{x};\mathbf{s}\right)=\mathbb{E}_{p\otimes\mu}\left[u\left(\mathbf{x},\mathbf{a}\right)\right]-\lambda I_{p\otimes\mu}\left(\mathbf{x};\mathbf{a}\right).$$

Let p^* achieve the maximum of the problem on the left-hand side of (A.1) and (d,σ) be the induced strategy. Then we have

$$\max_{p\in\Delta(A|X)}\ \mathbb{E}_{p\otimes\mu}\left[u\left(\mathbf{x},\mathbf{a}\right)\right]-\lambda I_{p\otimes\mu}\left(\mathbf{x};\mathbf{a}\right)=\mathbb{E}_{d\otimes\mu}\left[u\left(\mathbf{x},\sigma\left(\mathbf{s}\right)\right)\right]-\lambda I_{d\otimes\mu}\left(\mathbf{x};\mathbf{s}\right) \tag{A.3}$$

Combining (A.2) and (A.3), we obtain the desired result. Q.E.D.

Proof of Proposition 2: It is straightforward to check that the objective function F is convex in μ and jointly concave in p and q. We first solve the inside maximization problem for a fixed μ. By Proposition 1, we have the necessary and sufficient conditions for p and q, (9), (10), and (11). The value function for the inside problem $V(\mu)$ satisfies (12).

Now we consider the outside minimization problem, which can be written as

$$\min_{\mu\in\Delta(X)} V(\mu) + \theta \sum_x \mu(x) \ln \frac{\mu(x)}{\widehat{\mu}(x)}. \tag{A.4}$$

Since V and the relative entropy are convex in μ, this is a convex optimization problem. Replace $\mu(x_M)$ by

$$\mu(x_M) = 1 - \sum_{i=1}^{M-1} \mu(x_i). \tag{A.5}$$

Using (13) and taking first-order conditions with respect to $\mu(x)$ yields

$$v(x) - v(x_M) + \theta \left[\ln \frac{\mu(x)}{\widehat{\mu}(x)} - \ln \frac{\mu(x_M)}{\widehat{\mu}(x_M)} \right] = 0, \tag{A.6}$$

for $x = x_1, ..., x_{M-1}$. Solving for $\mu(x)$, summing over $x = x_1, x_2, ..., x_{M-1}$, and using (A.5), we can derive

$$\mu(x_M) = \frac{\exp(-v(x_M)/\theta)\, \widehat{\mu}(x_M)}{\sum_x \exp(-v(x)/\theta)\, \widehat{\mu}(x)}.$$

Plugging this expression into (A.6), we can derive

$$\mu(x) = \frac{\exp(-v(x)/\theta)\, \widehat{\mu}(x)}{\sum_{x'} \exp(-v(x')/\theta)\, \widehat{\mu}(x')}, x = x_1, ..., x_{M-1}.$$

We then obtain (23). Replacing μ with μ^* in (9), (10), (11) yields (20), (21), and (22). Plugging μ^* into (A.4) and using (12), we can derive the value function $J(\widehat{\mu})$ in the proposition. Q.E.D.

Proof of Proposition 3: It follows from (12) that

$$v(x) = \lambda \ln \sum_a q^*(a) \exp(u(x,a)/\lambda) = \lambda \ln (\overline{u}(1 + \delta q^*(x))).$$

Plugging the above expression into (23), we obtain

$$\mu^*(x) = \frac{(1+\delta q^*(x))^{-\frac{\lambda}{\theta}} \widehat{\mu}(x)}{\sum_{x'} (1+\delta q^*(x'))^{-\frac{\lambda}{\theta}} \widehat{\mu}(x')}. \tag{A.7}$$

Let B be the consideration set. For any chosen action $a \in B$, (21) is an equality. Plugging (A.7) into this equality yields

$$\begin{aligned}1 =& \sum_x \frac{\mu^*(x)\exp(u(x,a)/\lambda)}{\sum_b q^*(b)\exp(u(x,b)/\lambda)} \\ =& \frac{(1+\delta)\mu^*(a)}{1+\delta q^*(a)} + \sum_{a'\in B\setminus\{a\}} \frac{\mu^*(a')}{1+\delta q^*(a')} + \sum_{b\in A\setminus B} \mu^*(b) \\ =& \frac{\delta\mu^*(a)}{1+\delta q^*(a)} + \sum_{a'\in B} \frac{\mu^*(a')}{1+\delta q^*(a')} + \sum_{b\in A\setminus B} \mu^*(b). \end{aligned} \tag{A.8}$$

Observe that the last two terms on the right-hand side are independent of a. Therefore the first term $\frac{\delta\mu^*(a)}{1+\delta q^*(a)}$ is identical for any chosen action a. Using (A.7) to replace $\mu^*(a)$, we deduce that $(1+\delta q^*(a))^{-\frac{\lambda}{\theta}-1}\widehat{\mu}(a)$ is identical for any $a \in B$ as the denominator of the term on the right-hand side of (A.7) is independent of a. Therefore, we denote

$$\rho \equiv [1+\delta q^*(a)]^{-\frac{\lambda}{\theta}-1}\,\widehat{\mu}(a), \qquad \text{for any } a \in B. \tag{A.9}$$

Combining (A.8) and (A.9), we obtain

$$\frac{\rho(\delta+|B|)}{\sum_x (1+\delta q^*(x))^{-\frac{\lambda}{\theta}}\widehat{\mu}(x)} = 1 - \sum_{b\in A\setminus B}\mu^*(b) = \sum_{a\in B}\mu^*(a). \tag{A.10}$$

It then follows from (A.9) and (A.7) that

$$1+\delta q^*(a) = \left[\frac{\rho}{\widehat{\mu}(a)}\right]^{-\frac{\theta}{\lambda+\theta}} \quad \text{and} \quad \mu^*(a) = \frac{\rho^{\frac{\lambda}{\lambda+\theta}}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}}{\sum_x (1+\delta q^*(x))^{-\frac{\lambda}{\theta}}\widehat{\mu}(x)}. \tag{A.11}$$

Combining (A.10) and (A.11), we obtain

$$(\delta+|B|)\rho^{\frac{\theta}{\lambda+\theta}} = \sum_{a\in B}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}. \tag{A.12}$$

Equation (A.9) implies $\rho < \widehat{\mu}(a)$ for any chosen action $a \in B$ with $q^*(a) > 0$. It follows from (A.12) that

$$\widehat{\mu}(a) > \rho = \left[\frac{\sum_{a\in B}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}}{\delta+|B|}\right]^{\frac{\lambda+\theta}{\theta}}, \qquad \text{for any } a \in B. \tag{A.13}$$

For $a \notin B$, (21) is an inequality

$$\begin{aligned}1 \geq& \sum_x \frac{\mu^*(x)\exp(u(x,a)/\lambda)}{\sum_b q^*(b)\exp(u(x,b)/\lambda)} \\ =& \sum_{a'\in B}\frac{\mu^*(a')}{1+\delta q^*(a')} + \sum_{b\in A\setminus B}\mu^*(b) + \delta\mu^*(a),\end{aligned}$$

which is equivalent to

$$\delta\mu^*(a) \leq \sum_{a' \in B}\left[\mu^*(a') - \frac{\mu^*(a')}{1+\delta q^*(a')}\right] = \sum_{a' \in B}\frac{\delta q^*(a')}{1+\delta q^*(a')}\mu^*(a'). \tag{A.14}$$

Because a is not chosen, $q^*(a) = 0$. Then (A.7) implies that

$$\mu^*(a) = \frac{\widehat{\mu}(a)}{\sum_x (1+\delta q^*(x))^{-\frac{\lambda}{\theta}}\widehat{\mu}(x)}, \quad \text{for any } a \in A \setminus B. \tag{A.15}$$

Plugging the previous equation and (A.7) into (A.14), we obtain

$$\widehat{\mu}(a) \leq \sum_{a' \in B} q^*(a')\left[1+\delta q^*(a')\right]^{-\frac{\lambda}{\theta}-1}\widehat{\mu}(a') = \sum_{a' \in B} q^*(a')\,\rho = \rho, \quad \text{for any } a \in A \setminus B, \tag{A.16}$$

where the equality follows from $\sum_{a' \in B} q^*(a') = 1$.

By (A.13) and (A.16), we have

$$\widehat{\mu}(a) > \rho, \text{ for any } a \in B; \quad \widehat{\mu}(a) \leq \rho, \text{ for any } a \in A \setminus B,$$

where

$$\rho = \left[\frac{\sum_{a \in B}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}}{\delta + |B|}\right]^{\frac{\lambda+\theta}{\theta}}.$$

If

$$\widehat{\mu}(M) > \left[\frac{\sum_{a=1}^{M}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}}{\delta + M}\right]^{\frac{\lambda+\theta}{\theta}}, \tag{A.17}$$

then

$$\widehat{\mu}(x) > \left[\frac{\sum_{a=1}^{M}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}}{\delta + M}\right]^{\frac{\lambda+\theta}{\theta}}, \quad \text{for any } x = 1, \ldots, M,$$

which implies $B = A$, $K = M$, and all actions are chosen.

If (A.17) is violated, we claim that there must exist a unique integer K such that

$$\widehat{\mu}(K) > \left[\frac{\sum_{a=1}^{K}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}}{\delta + K}\right]^{\frac{\lambda+\theta}{\theta}} \geq \widehat{\mu}(K+1). \tag{A.18}$$

To see this, note that

$$\hat{\mu}(1) > \left[\frac{\widehat{\mu}(1)^{\frac{\theta}{\lambda+\theta}}}{\delta + 1}\right]^{\frac{\lambda+\theta}{\theta}} = \frac{\widehat{\mu}(1)}{(\delta+1)^{\frac{\lambda+\theta}{\theta}}}.$$

On the other hand,

$$\widehat{\mu}(M) \leq \left[\frac{\sum_{a=1}^{M}\widehat{\mu}(a)^{\frac{\theta}{\lambda+\theta}}}{\delta + M}\right]^{\frac{\lambda+\theta}{\theta}},$$

because (A.17) is violated.

If $\left[\widehat{\mu}(1)^{\frac{\theta}{\lambda+\theta}}/(\delta+1)\right]^{\frac{\lambda+\theta}{\theta}} \geq \widehat{\mu}(2)$, then we set $K=1$. Otherwise, $\left[\widehat{\mu}(1)^{\frac{\theta}{\lambda+\theta}}/(\delta+1)\right]^{\frac{\lambda+\theta}{\theta}} < \widehat{\mu}(2)$, which implies $\widehat{\mu}(1)^{\frac{\theta}{\lambda+\theta}} < (\delta+1)\widehat{\mu}(2)^{\frac{\theta}{\lambda+\theta}}$. Adding $\widehat{\mu}(2)^{\frac{\theta}{\lambda+\theta}}$ on both sides of the preceding inequality, we obtain $\sum_{i=1}^{2}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} < (\delta+2)\widehat{\mu}(2)^{\frac{\theta}{\lambda+\theta}}$, which is equivalent to

$$\left[\frac{\sum_{i=1}^{2}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+2}\right]^{\frac{\lambda+\theta}{\theta}} < \widehat{\mu}(2).$$

If

$$\left[\frac{\sum_{i=1}^{2}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+2}\right]^{\frac{\lambda+\theta}{\theta}} \geq \widehat{\mu}(3),$$

we set $K=2$. Otherwise, we continue this process by counting up. In general, if

$$\widehat{\mu}(k) > \left[\frac{\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+k}\right]^{\frac{\lambda+\theta}{\theta}}, \text{ but } \left[\frac{\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+k}\right]^{\frac{\lambda+\theta}{\theta}} < \widehat{\mu}(k+1),$$

we must have

$$\left[\frac{\sum_{i=1}^{k+1}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+k+1}\right]^{\frac{\lambda+\theta}{\theta}} < \mu(k+1)$$

for any integer $k \leq M-1$. In the end, because of assumption (28), the process must stop to find the smallest $K \leq M-1$ such that (A.18) holds.

To show that there is a unique K satisfying (A.18), we write the second inequality as $\sum_{i=1}^{K}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} \geq (\delta+K)\,\widehat{\mu}(K+1)^{\frac{\theta}{\lambda+\theta}}$. Suppose that there is another integer $K' \geq K+1$ also satisfying (A.18). Adding $\sum_{i=K+1}^{K'}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}$ to both sides of the preceding inequality, we obtain

$$\sum_{i=1}^{K'}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} \geq (\delta+K)\widehat{\mu}(K+1)^{\frac{\theta}{\lambda+\theta}} + \sum_{i=K+1}^{K'}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} \geq (\delta+K')\widehat{\mu}(K')^{\frac{\theta}{\lambda+\theta}},$$

where the second inequality follows from the decreasing property of $\widehat{\mu}(x)$ in x. We then obtain a contradiction.

Then we deduce that the consideration set is $B=\{1,2,...,K\}$ and

$$\rho = \left[\frac{\sum_{i=1}^{K}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+K}\right]^{\frac{\lambda+\theta}{\theta}}.$$

We can now use (A.11) and (A.15) to identify μ^* and use (A.11) to identify q^* as in the proposition. The optimal posterior can also be determined. When $x=a$ and $a \leq K$,

$$\mu^*(x|a) = \frac{\exp(u(x,a)/\lambda)\mu^*(x)}{\sum_b q^*(b)\exp(u(x,b)/\lambda)} = \frac{(1+\delta)\mu^*(x)}{1+\delta q^*(a)} = \frac{(1+\delta)\rho}{\sum_{i=1}^{K}\rho^{\frac{\lambda}{\lambda+\theta}}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} + \sum_{i=K+1}^{M}\widehat{\mu}(i)}.$$

When $x \neq a$ and $x \leq K$,

$$\mu^*(x|a) = \frac{\mu^*(x)}{1+\delta q^*(a)} = \frac{\rho}{\sum_{i=1}^{K}\rho^{\frac{\lambda}{\lambda+\theta}}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} + \sum_{i=K+1}^{M}\widehat{\mu}(i)}.$$

When $x > K$, $\mu^*(x|a) = \mu^*(x)$. The proof is completed. Q.E.D.

Proof of Proposition 4: (i) Define

$$\rho(k,\theta) \equiv \left[\frac{\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+k}\right]^{\frac{\lambda+\theta}{\theta}}$$

for any $k>0$. We claim that $\rho(k,\theta)$ is increasing in θ. To see this, consider $\theta<\theta'$. Then

$$\left(\frac{\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta+k}\right)^{\frac{\lambda+\theta}{\theta}} = \left[\left(\frac{\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{1}{\lambda/\theta+1}}}{\delta+k}\right)^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}}\right]^{\lambda/\theta'+1}$$

$$= \left[\left(\frac{k}{\delta+k}\right)^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}}\left(\frac{1}{k}\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{1}{\lambda/\theta+1}}\right)^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}}\right]^{\lambda/\theta'+1}$$

$$< \left[\frac{k}{\delta+k}\left(\frac{1}{k}\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{1}{\lambda/\theta+1}}\right)^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}}\right]^{\lambda/\theta'+1}$$

$$\leq \left[\frac{k}{\delta+k}\left(\frac{1}{k}\sum_{i=1}^{k}\widehat{\mu}(i)^{\lambda/\theta'+1}\right)\right]^{\lambda/\theta'+1} = \left(\frac{\sum_{i=1}^{k}\widehat{\mu}(i)^{\frac{\theta'}{\lambda+\theta'}}}{\delta+k}\right)^{\frac{\lambda+\theta'}{\theta'}},$$

where the first inequality follows from $\left(\frac{k}{\delta+k}\right)^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}} < \frac{k}{\delta+k}$, because $\frac{\lambda/\theta+1}{\lambda/\theta'+1}>1$ and $\frac{k}{\delta+k}<1$, and the second inequality holds follows from Jensen's inequality $\left[\frac{1}{k}\sum_{i=1}^{k}\ell(i)\right]^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}} \leq \frac{1}{k}\sum_{i=1}^{k}\ell(i)^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}}$ with $\ell(i)=\widehat{\mu}(i)^{\frac{1}{\lambda/\theta+1}}$ as $x^{\frac{\lambda/\theta+1}{\lambda/\theta'+1}}$ is convex in x.

Let K_θ be the unique integer satisfying

$$\widehat{\mu}(K_\theta) > \rho(K_\theta,\theta) \geq \widehat{\mu}(K_\theta+1).$$

For $\theta'>\theta$, because $\rho(K_\theta,\theta')>\rho(K_\theta,\theta)$, we have $\rho(K_\theta,\theta')>\widehat{\mu}(K_\theta+1)$ as well. The previous inequality and the definition of $\rho(K_\theta,\theta')$ imply that

$$\sum_{i=1}^{K_\theta}\widehat{\mu}(i)^{\frac{\theta'}{\lambda+\theta'}} \geq (\delta+K_\theta)\,\widehat{\mu}(K_\theta+1)^{\frac{\theta'}{\lambda+\theta'}}.$$

For any $K'>K_\theta$, adding $\sum_{i=K_\theta+1}^{K'}\widehat{\mu}(i)^{\frac{\theta'}{\lambda+\theta'}}$ on both sides and using the property that $\widehat{\mu}(x)$ is decreasing in x, we obtain

$$\sum_{i=1}^{K'}\widehat{\mu}(i)^{\frac{\theta'}{\lambda+\theta'}} \geq (\delta+K_\theta)\,\widehat{\mu}(K_\theta+1)^{\frac{\theta'}{\lambda+\theta'}} + \sum_{i=K_\theta+1}^{K'}\widehat{\mu}(i)^{\frac{\theta'}{\lambda+\theta'}} \geq (\delta+K')\widehat{\mu}(K')^{\frac{\theta'}{\lambda+\theta'}},$$

which is equivalent to $\rho(K',\theta')\geq\widehat{\mu}(K')$. Therefore any integer $K'>K_\theta$ cannot satisfy $\widehat{\mu}(K')>\rho(K',\theta')$. However, we have shown in Proposition 3 that there exists a unique $K_{\theta'}$ such that

$$\widehat{\mu}(K_{\theta'}) > \rho\left(K_{\theta'},\theta'\right) \geq \widehat{\mu}(K_{\theta'}+1).$$

Therefore we must have $K_{\theta'} \leq K_{\theta}$.

(ii) Observe that $\widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}} \leq 1$ because $\widehat{\mu}(i) \leq 1$. Therefore

$$\rho(M, \theta) = \left[\frac{\sum_{i=1}^{M} \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}{\delta + M}\right]^{\frac{\lambda+\theta}{\theta}} \leq \left(\frac{M}{\delta + M}\right)^{\frac{\lambda+\theta}{\theta}}.$$

As $\theta \to 0$, the right-hand side of the inequality above decreases to zero, and hence $\rho(M, \theta)$ converges to zero as well. Since $\mu(M) > 0$, inequality (28) is satisfied and hence $K = M$, when $\theta > 0$ is small enough. It then follows from (31) that

$$\mu^*(x) = \frac{\rho^{\frac{\lambda}{\lambda+\theta}} \widehat{\mu}(x)^{\frac{\theta}{\lambda+\theta}}}{\sum_{i=1}^{M} \rho^{\frac{\lambda}{\lambda+\theta}} \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}} = \frac{\widehat{\mu}(x)^{\frac{\theta}{\lambda+\theta}}}{\sum_{i=1}^{M} \widehat{\mu}(i)^{\frac{\theta}{\lambda+\theta}}}, \quad x = 1, \ldots, M,$$

when $\theta > 0$ is small enough. Sending $\theta \to 0$, the numerator converges to 1 and the denominator converges to M, and hence $\mu^*(x)$ converges to $1/M$.

For the statement on q^*, observe from (33) that

$$\frac{1 + \delta q^*(a)}{1 + \delta q^*(a')} = \left[\frac{\widehat{\mu}(a')}{\widehat{\mu}(a)}\right]^{-\frac{\theta}{\lambda+\theta}}, \quad \text{for any } a \neq a'.$$

As $\theta \to 0$, the right-hand side converges to 1, hence the difference of $q^*(a)$ and $q^*(a')$ converges to zero for any pair of a and a'. Since $\sum_{a=1}^{M} q^*(a) = 1$, it follows that each $q^*(a)$ converges to $1/M$. Q.E.D.

Proof of Proposition 5: We prove by contradiction. Suppose that there exist two actions k and i such that $\mathbb{E}^*[\exp(u(\mathbf{x}, k)/\lambda)] < \mathbb{E}^*[\exp(u(\mathbf{x}, i)/\lambda)]$, but k is chosen with positive probability and i is not. By the necessary and sufficient conditions in Proposition 2, we have

$$\mathbb{E}^*\left[\frac{\exp(\mathbf{x}_k/\lambda)}{\sum_a q^*(a)\exp(u(\mathbf{x}, a)/\lambda)}\right] = 1 \geq \mathbb{E}^*\left[\frac{\exp(\mathbf{x}_i/\lambda)}{\sum_a q^*(a)\exp(u(\mathbf{x}, a)/\lambda)}\right]. \tag{A.19}$$

Let B denote the consideration set.

We now show that the worst-case prior μ^* satisfies a certain independence property. By the independence assumption of $\widehat{\mu}$, we have

$$\mu^*(x) = \frac{\exp(-v(x)/\theta)\widehat{\mu}(x)}{\sum_x \exp(-v(x)/\theta)\widehat{\mu}(x)} = \frac{\exp(-v(x)/\theta)\prod_{j=1}^{M}\widehat{\mu}_j(x_j)}{\sum_x \exp(-v(x)/\theta)\widehat{\mu}(x)}. \tag{A.20}$$

It follows from (12) that

$$v(x) = \lambda \ln \sum_a q^*(a)\exp(u(x, a)/\lambda) = \lambda \ln \sum_{j \in B} q^*(j)\exp(x_j/\lambda) \tag{A.21}$$

is independent of $\exp(x_i/\lambda)$ as $i \notin B$. Let x_{-i} denote the vector x without the component x_i. Then we rewrite (A.20) as

$$\mu^*(x) = \widehat{\mu}_i(x_i)\,\nu(x_{-i}), \text{ if } i \notin B, \tag{A.22}$$

where we can check that

$$\nu(x_{-i}) \equiv \frac{\exp(-v(x)/\theta)\prod_{j=1,j\neq i}^{M}\widehat{\mu}_j(x_j)}{\sum_x \exp(-v(x)/\theta)\,\widehat{\mu}(x)}$$

is a probability distribution on Z^{M-1}.

Since action i is not chosen,

$$\sum_a q^*(a)\exp(u(x,a)/\lambda) = \sum_{j\in B} q^*(j)\exp(x_j/\lambda)$$

is independent of $\exp(u(x,i)/\lambda) = \exp(x_i/\lambda)$. Thus, by Proposition 2,

$$\begin{aligned}
\mathbb{E}^*\left[\frac{\exp(\mathbf{x}_i/\lambda)}{\sum_a q^*(a)\exp(u(\mathbf{x},a)/\lambda)}\right] &= \sum_{(x_1,\dots,x_M)\in Z^M}\frac{\exp(x_i/\lambda)\,\mu^*(x_1,\dots,x_M)}{\sum_{j\in B}q^*(j)\exp(x_j/\lambda)} \\
&= \mathbb{E}^*[\exp(\mathbf{x}_i/\lambda)]\,\mathbb{E}^*\left[\left(\sum_{j\in B}q^*(j)\exp(\mathbf{x}_j/\lambda)\right)^{-1}\right],
\end{aligned}$$

where the last equality follows from the special independence property of μ^* derived earlier. Thus, by the supposition and (A.19), we have

$$\begin{aligned}
\mathbb{E}^*\left[\frac{\exp(\mathbf{x}_k/\lambda)}{\sum_a q^*(a)\exp(u(\mathbf{x},a)/\lambda)}\right] &= \mathbb{E}^*\left[\frac{\exp(\mathbf{x}_k/\lambda)}{\sum_{j\in B}q^*(j)\exp(\mathbf{x}_j/\lambda)}\right] \\
&> \mathbb{E}^*[\exp(\mathbf{x}_k/\lambda)]\,\mathbb{E}^*\left[\left(\sum_{j\in B}q^*(j)\exp(\mathbf{x}_j/\lambda)\right)^{-1}\right].
\end{aligned}$$

Now we rewrite this inequality as

$$\begin{aligned}
0 &< \mathbb{E}^*\left[\frac{\exp(\mathbf{x}_k/\lambda)-\mathbb{E}^*[\exp(\mathbf{x}_k/\lambda)]}{\sum_{j\in B}q^*(j)\exp(\mathbf{x}_j/\lambda)}\right] \\
&= \mathbb{E}^*\left[\frac{\exp(\mathbf{x}_k/\lambda)-\mathbb{E}^*[\exp(\mathbf{x}_k/\lambda)]}{q^*(k)\exp(\mathbf{x}_k/\lambda)+\sum_{j\in B\backslash\{k\}}q^*(j)\exp(\mathbf{x}_j/\lambda)}\right] \\
&= \mathbb{E}^*\left[\mathbb{E}^*\left[\left.\frac{\exp(\mathbf{x}_k/\lambda)-\mathbb{E}^*[\exp(\mathbf{x}_k/\lambda)]}{q^*(k)\exp(x_k/\lambda)+\sum_{j\in B\backslash\{k\}}q^*(j)\exp(\mathbf{x}_j/\lambda)}\right|\sum_{j\in B\backslash\{k\}}q^*(j)\exp(\mathbf{x}_j/\lambda)\right]\right] \\
&= \mathbb{E}^*\left\{Cov\left(\exp(\mathbf{x}_k/\lambda),\left.\frac{1}{q^*(k)\exp(\mathbf{x}_k/\lambda)+\sum_{j\in B\backslash\{k\}}q^*(j)\exp(\mathbf{x}_j/\lambda)}\right|\sum_{j\in B\backslash\{k\}}q^*(j)\exp(\mathbf{x}_j/\lambda)\right)\right\}.
\end{aligned}$$

The last conditional covariance is negative, which is a contradiction. Q.E.D.

Proof of Lemma 2: First we show that a strategy of choosing $i \in B$ with probability $1/N$ satisfies the necessary and sufficient conditions of Proposition 2. Let $q^*(j) = 1/N$ for $j \in B$. Then

$$\sum_{x \in X} \frac{\exp(u(x,i)/\lambda)\,\mu^*(x)}{\sum_{j \in B} q^*(j)\exp(u(x,j)/\lambda)} = \sum_{x \in X} \frac{\exp(x_i/\lambda)\,\mu^*(x)}{\frac{1}{N}\sum_{j \in B}\exp(x_j/\lambda)}, \tag{A.23}$$

where μ^* is the worst-case prior given in the proposition. By (12), we have

$$v(x) = \lambda \ln \sum_a q^*(a)\exp(u(x,a)/\lambda) = \lambda \ln \sum_{j \in B} \frac{1}{N}\exp(x_j/\lambda). \tag{A.24}$$

Let $\widetilde{v}(\widetilde{x})$ denote the last expression, where $\widetilde{x} \equiv (x_1, x_2, ..., x_N)$. By (23), we have

$$\begin{aligned} \mu^*(x) &= \frac{\exp(-v(x)/\theta)\,\widehat{\mu}(x)}{\sum_x \exp(-v(x)/\theta)\,\widehat{\mu}(x)} = \frac{\exp(-v(x)/\theta)\prod_{i=1}^{M}\widehat{\mu}_i(x_i)}{\sum_x \exp(-v(x)/\theta)\,\widehat{\mu}(x)} \\ &= \frac{\exp(-\widetilde{v}(\widetilde{x})/\theta)\prod_{i=1}^{M}\widehat{\mu}_i(x_i)}{\sum_{\widetilde{x} \in Z^N}\exp(-\widetilde{v}(\widetilde{x})/\theta)\prod_{i=1}^{N}\widehat{\mu}_B(x_i)}. \end{aligned}$$

Substituting this expression in (A.23) yields

$$\sum_{x \in X} \frac{\exp(u(x,i)/\lambda)\,\mu^*(x)}{\sum_{j \in B} q^*(j)\exp(u(x,j)/\lambda)} = \sum_{\widetilde{x} \in Z^N} \frac{\exp(x_i/\lambda)}{\frac{1}{N}\sum_{j \in B}\exp(x_j/\lambda)} \frac{\exp(-\widetilde{v}(\widetilde{x})/\theta)\prod_{i=1}^{N}\widehat{\mu}_B(x_i)}{\sum_{\widetilde{x} \in Z^N}\exp(-\widetilde{v}(\widetilde{x})/\theta)\prod_{i=1}^{N}\widehat{\mu}_B(x_i)}.$$

Except for $\exp(x_i/\lambda)$, other terms are invariant under permutation. By the same permutation argument as in CDL (2019), we can show that

$$\sum_{x \in X} \frac{\exp(u(x,i)/\lambda)\,\mu^*(x)}{\sum_{j \in B} q^*(j)\exp(u(x,j)/\lambda)} = 1 \text{ for } i \in B.$$

Now we show that for all $k \notin B$,

$$\sum_{x \in X} \frac{\exp(u(x,k)/\lambda)\,\mu^*(x)}{\sum_{j \in B} q^*(j)\exp(u(x,j)/\lambda)} \leq 1.$$

By the special independence property in (A.22) and the condition in the lemma, we obtain the desired result. Q.E.D.

Reference

Arimoto, Suguru, 1972, An Algorithm for Computing the Capacity of Arbitrary Discrete Memoryless Channels, *IEEE Transactions on Information Theory* 18, 14-20.

Bertsekas, Dimitri, 2016, *Nonlinear Programming*, Athena Scientific.

Blahut, Richard, E., 1972, Computation of Channel Capacity and Rate-Distortion Functions, *IEEE Transactions on Information Theory* 18, 460-473.

Caplin, Andrew, and Mark Dean, 2013, Behavioral Implications of Rational Inattention with Shannon Entropy, NBER working paper 19318.

Caplin, Andrew, and Mark Dean, 2015, Revealed Preference, Rational Inattention, and Costly Information Acquisition, *American Economic Review* 105, 2183-2203.

Caplin, Andrew, Mark Dean, and John Leahy, 2019, Rational Inattention, Optimal Consideration Sets, and Stochastic Choice, *Review of Economic Studies* 86, 1061-1094.

Caplin, Andrew, Mark Dean, and John Leahy, 2020, Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy, working paper, Columbia University.

Caplin, Andrew, and Daniel Martin, 2015, A Testable Theory of Imperfect Perception, *Economic Journal* 125, 184–202.

Cover, Thomas M., and Joy A. Thomas, 2006, *Elements of Information Theory*, 2nd edition, John Wiley & Sons, Inc.

Ellsberg, Daniel, 1961, Risk, Ambiguity and the Savage Axiom, *Quarterly Journal of Economics* 75, 643–669.

Epstein, Larry G., and Martin Schneider, 2010, Ambiguity and Asset Markets, *Annual Review of Financial Economics*, Vol. 2, 315-346.

Gilboa, Itzhak, and David Schmeidler, 1989, Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics* 18, 141-153.

Hansen, Lars, P., 2007, Beliefs, Doubts and Learning: Valuing Macroeconomic Risk, *American Economic Review* 97, 1-30.

Hansen, Lars, P., and Jianjun Miao, 2018, Aversion to Ambiguity and Model Misspecification in Dynamic Stochastic Environments, *Proceedings of National Academy of Sciences* 115, 9163-9168.

Hansen, Lars P., and Thomas J. Sargent, 2001, Robust Control and Model Uncertainty, *American Economic Review* 91, 60–66.

Hansen, Lars P., and Thomas J. Sargent, 2007, Recursive Robust Estimation and Control without Commitment, *Journal of Economic Theory* 136, 1–27.

Hansen, Lars P., and Thomas J. Sargent, 2008, *Robustness*, Princeton University Press, Princeton, New Jersey.

Hansen, Lars P., and Thomas J. Sargent, 2010a, Fragile Beliefs and Price of Uncertainty, *Quantitative Economics* 1, 129-162.

Hansen, Lars P., and Thomas J. Sargent, 2010b, Wanting Robustness in Macroeconomics, in *Handbook of Monetary Economics*, V. 3B, Benjamin M. Friedman and Michael Woodford (ed.), North-Holland.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005, A Smooth Model of Decision Making under Ambiguity, *Econometrica* 73, 1849–1892.

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini, 2006, Ambiguity Aversion, Robustness, and the Variational Representation of Preferences, *Econometrica* 74, 1447–1498.

Jung, Junehyuk, Jeong-Ho Kim, Filip Matějka, and Christopher A. Sims, 2018, Discrete Actions in Information-constrained Decision Problems, working paper, Princeton University.

Kőszegi, Botand and Filip Matějka, 2019, Choice Simplification: A Theory of Mental Budgeting and Naive Diversification, forthcoming in *Quarterly Journal of Economics*.

Luo, Yulei, 2008, Consumption Dynamics under Information Processing Constraints, *Review of Economic Dynamics* 11, 366-385.

Maćkowiak, Bartosz, and Mirko Wiederholt, 2009, Optimal Sticky Prices under Rational Inattention, *American Economic Review* 99, 769–803.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt, 2018, Survey: Rational Inattention, a Disciplined Behavioral Model, working paper, Sciences Po.

Matějka, Filip, and Alisdair McKay, 2015, Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model, *American Economic Review* 105, 272-298.

Miao, Jianjun, 2019, Multivariate LQG Control under Rational Inattention in Continuous Time, working paper, Boston University.

Miao, Jianjun, Jieran Wu, and Eric Young, 2019, Multivariate Rational Inattention, working paper, Boston University.

Miao, Jianjun, and Hao Xing, 2020, Dynamic Discrete Choice under Rational Inattention, working paper, Boston University.

Mondria, Jordi, 2010, Portfolio Choice, Attention Allocation, and Price Comovement, *Journal of Economic Theory* 145, 1837-1864.

Peng, Lin, and Wei Xiong, 2006, Investor Attention, Overconfidence and Category Learning, *Journal of Financial Economics* 80, 563-602.

Shannon, Claude, E., 1948, A Mathematical Theory of Communication, *Bell System Technical Journal* 27, 379–423.

Sims, Christopher A., 1998, Stickiness, *Carnegie-Rochester Conference Series On Public Policy* 49, 317–356.

Sims, Christopher A., 2003, Implications of Rational Inattention, *Journal of Monetary Economics* 50 (3), 665–690.

Sims, Christopher A., 2010, Rational Inattention and Monetary Economics, in *Handbook of Monetary Economics*, V. 3A, Benjamin M. Friedman and Michael Woodford (ed.), North-Holland.

Steiner, Jakub, Colin Stewart, and Filip Matějka, 2017, Rational Inattention Dynamics: Inertia and Delay in Decision-Making, *Econometrica* 85, 521-553.

Strzalecki, Tomasz, 2011, Axiomatic Foundations of Multiplier Preferences, *Econometrica* 97, 47-73.

Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2010, Information Acquisition and Under-diversification, *Review of Economic Studies* 77, 571-626.

Woodford, Michael, 2009, Information Constrained State Dependent Pricing. *Journal of Monetary Economics* 56(S), S100–S124.

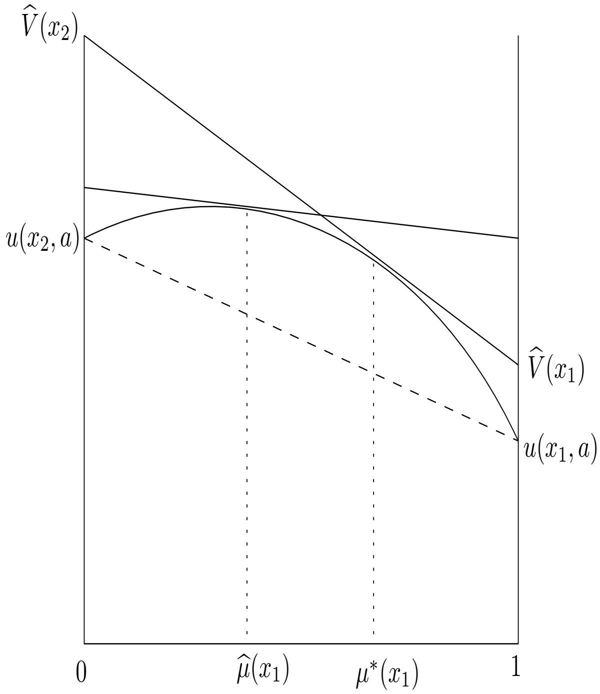

Figure 1: Geometric illustration of robust solutions.

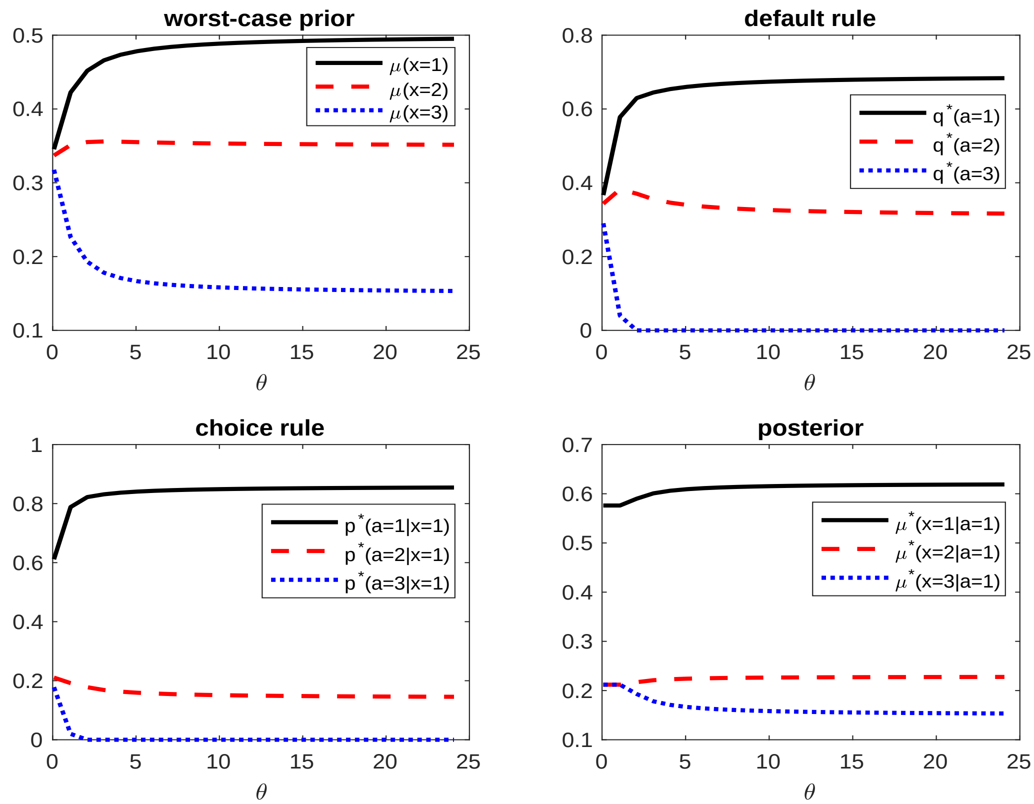

Figure 2: Solutions for consumer problem 1 with different degrees of ambiguity aversion.

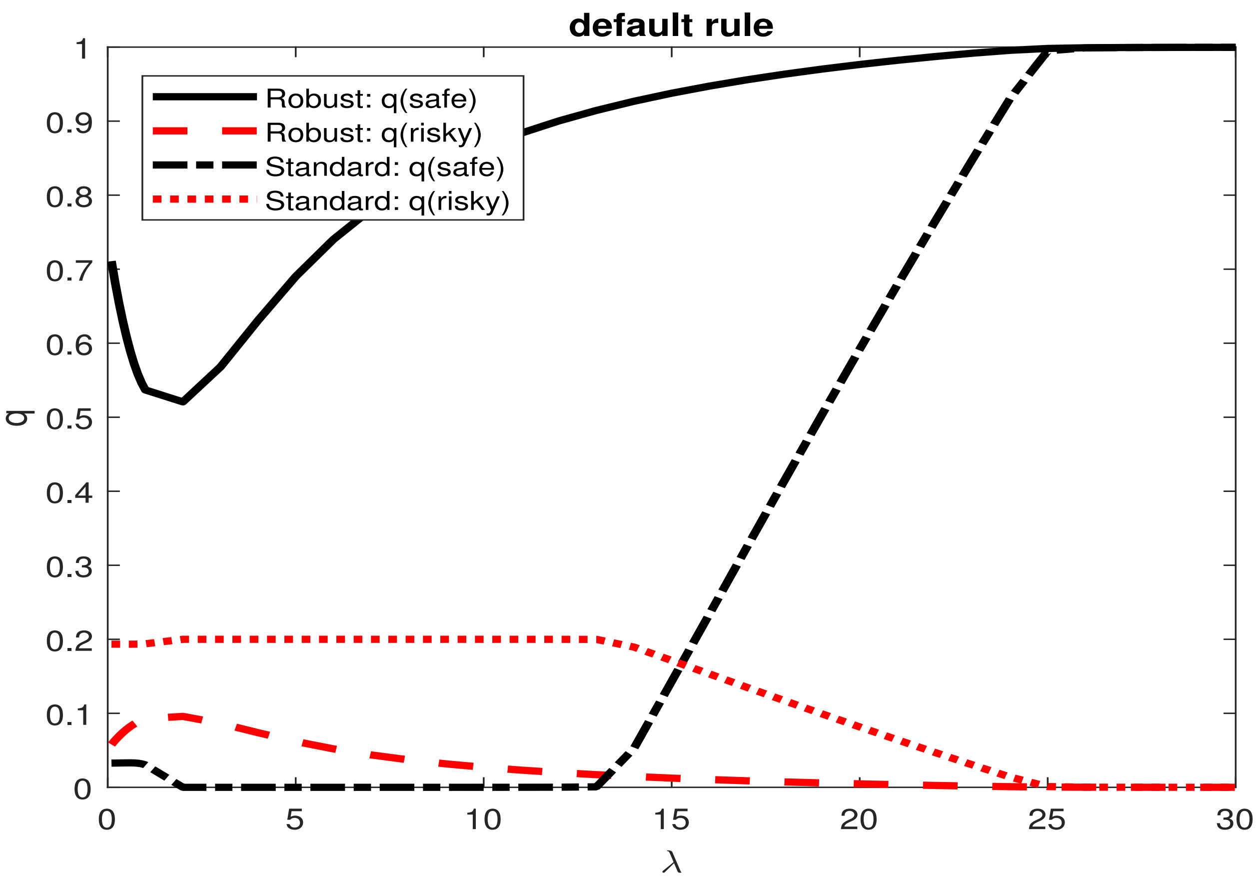

Figure 3: The impact of the information cost parameter λ on the default rule in the standard and robust Shannon models for consumer problem 2.

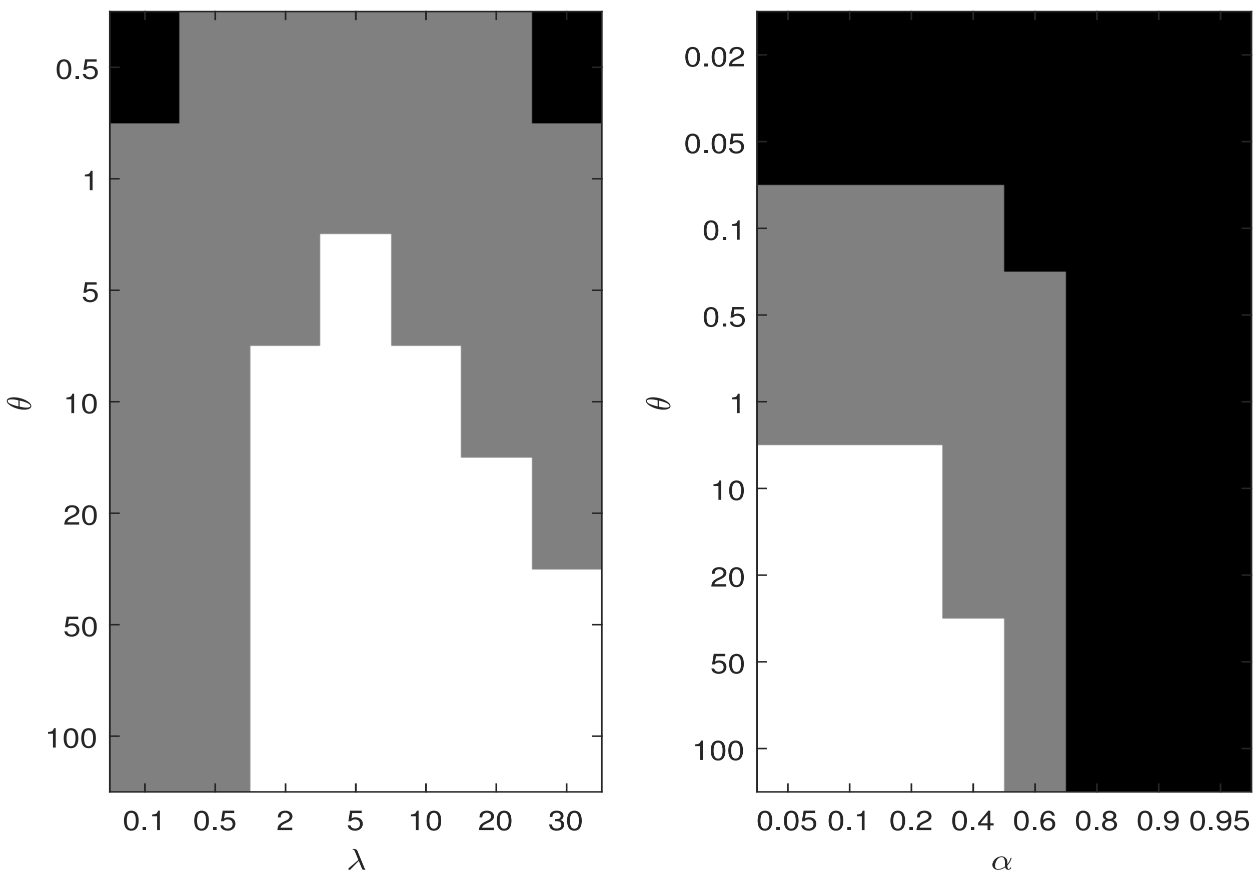

Figure 4: The left (right) panel plots the consideration sets for different values of the robustness parameter θ and the information cost (risk aversion) parameter λ (α) for consumer problem 2 with $\alpha = 0$ ($\lambda = 2$). White area: risky option only. Grey area: risky and safe options. Black area: safe option only.

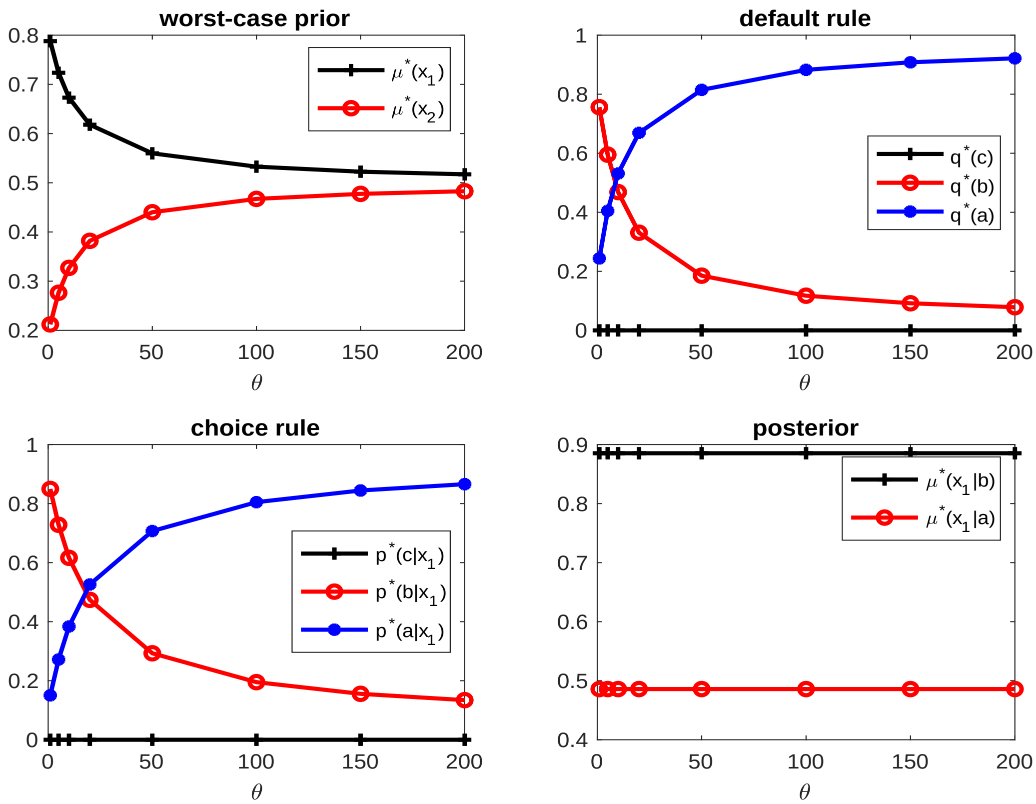

Figure 5: Solutions for consumer problem 3 with different degrees of ambiguity aversion.

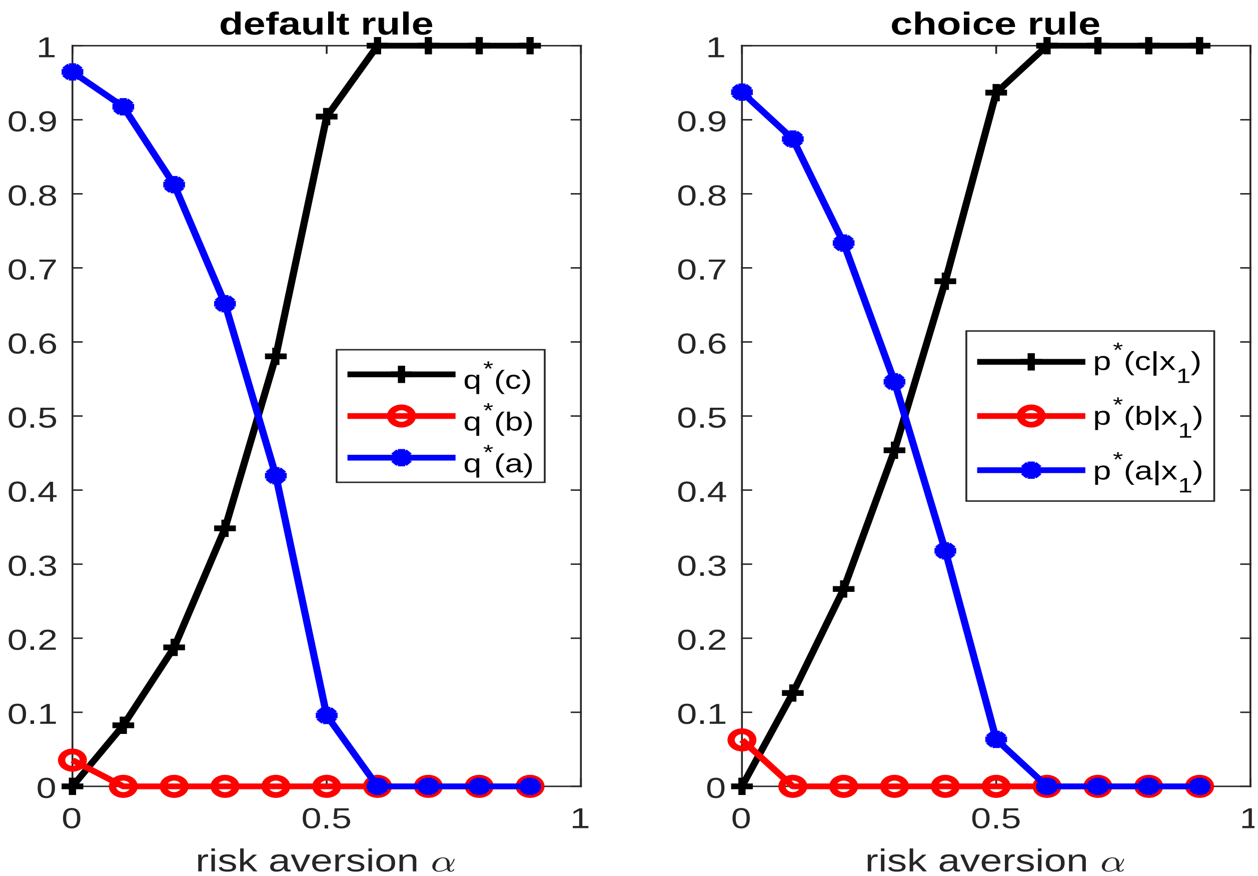

Figure 6: Solutions for consumer problem 3 with different degrees of risk aversion.

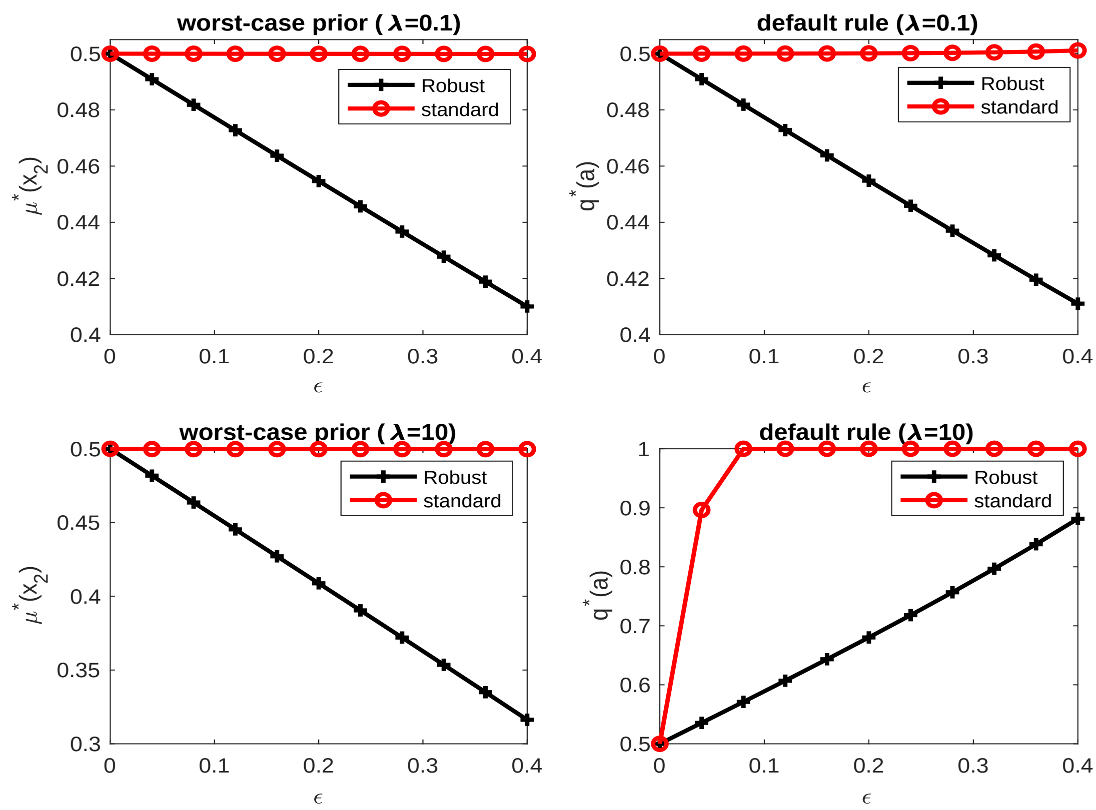

Figure 7: Comparative statics.