OPTIMAL INQUIRY

Tai-Wei Hu and Andriy Zapechelnyuk

Abstract. We propose a new framework of costly information processing. A decision maker processes information about an uncertain state of the world by a procedure called *inquiry*. An inquiry starts with an initial question about the state, specifies subsequent questions depending on earlier answers, and, eventually, prescribes decisions. The decision maker bears a cost proportional to the expected length of inquiry. Thus, more refined information is costlier. We characterize optimal inquiries and show their dynamic consistency. We also show that optimal inquiry exhibits two behavioural biases: focused consideration (the decision maker restricts attention to a subset of decisions) and confirmation bias (the decision maker seeks evidence through inquiry to confirm her prior guess of which decisions are optimal).

JEL Classification: D81, D83.

Keywords: Rational inattention; bounded rationality; information theory; decision under risk; confirmation bias; consideration set

Preliminary Draft

Date: February 26, 2024

Tai-Wei Hu: School of Economics, University of Bristol, 12 Priory Road, Bristol, BS8 1TU, UK. *E-mail:* taiwei.hu@bristol.ac.uk

Andriy Zapechelnyuk: School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK. *E-mail:* azapech@gmail.com.

The authors are grateful to Tan Gan and Larry Samuelson for helpful comments.

1. Introduction

Inquiry is one of the most frequent and important modes of information processing in our daily life. Examples are abundant. A doctor visit usually consists of a series of questions from reception to actual consultation of the patient's conditions. A crime investigation typically consists of a series of questions and processing their answers. Inquiry about product characteristics and payment schemes is an important aspect of shopping experiences. In all these examples, information to be gathered can be potentially overwhelming, whereas cognitive resources available to process it are limited and precious. In this paper, we propose a theory of optimal inquiry to process information that takes costly cognition seriously, and study its implications to the decision-making processes and outcomes.

We formalize an *inquiry* as the decision maker's strategy of asking questions about the relevant state of the world. It starts with an initial question and a contingent plan that decides which question to ask depending on the answers to the previous ones. As in the standard Bayesian paradigm, the answers to the inquiry hence determine the information set that guides the decision maker's final action. Unlike the standard framework, however, our framework explicitly postulates a cognitive cost associated with the length of the inquiry; more precisely, we assume that the cost is linear in the average number of questions the inquiry entails.

Our framework provides an explicit and intuitive procedure for information processing, and share the same motivation that initiates the rational-inattention literature, as Sims' (2010) example shows, "Finding whether a test well indicates oil is present may cost thousands of dollars, yet provide only the answer to a yes-or-no question [...] Rational inattention theory [...] might explain why an executive in the oil company [...] might after 'looking at' all the reports seem to know the test well report in detail, while having only a vague idea of what was in the other reports." Different from that literature, however, our explicit formulation of inquiry allows us to put the cost directly on the procedure due to the physical and mental acts involved, independent of the content. For example, the cost of performing a blood-sugar test and processing its result (in terms of physical or cognitive resources) would be independent of the doctor's prior knowledge about the patient in our model, whereas the cost would depend on the prior belief of the doctor in the standard rational inattention model.

Moreover, the explicit formulation of the decision-making process allows us to evaluate the dynamic consistency of the process in our framework. While optimality is defined from the ex ante perspective, we show that an optimal inquiry is dynamically consistent in the following sense. Consider a decision maker who processes information following an optimal inquiry, and suppose that she has asked a few questions but not yet ready for a final decision. At this point, she could stop and reconsider her inquiry strategy, taking all the information she already processed so far as given. Dynamic consistency requires it be optimal to stick with the original plan at this interim stage, and we prove this property for any optimal inquiry in our framework.

Optimal inquiry trades off the accuracy of information processed against the cognitive resources needed to achieve it. On the extensive margin, this leads to an endogenous *consideration set*, according to which only a subset of all feasible options is considered, to save cognitive resources. On the intensive margin, this leads an endogenous collection of *categories*. The inquiry determines which category the state belongs to, and then prescribes the decision in the consideration set that is optimal for that category. A novel feature is that the optimal categories are jointly determined with the consideration set and the structure of inquiry.

Our first main result fully characterizes the optimal inquiry. This characterization bridges information theory to economic decision-making, for which we use two well-known results in information theory: the Kraft inequality and the Huffman coding. First, we use the Kraft inequality to fully characterize the set of payoff-relevant outcomes that can be implemented by an inquiry. Such an outcome consists of a consideration set, categories of states induced by the inquiry, and the inquiry length profile corresponding to each category. Kraft inequality is then employed to show that only the size of the consideration set and the length profile matters and they are connected by a simple equation.

This allows us to reduce the overall design problem into two stages: first we choose a consideration set and a length profile, and then we solve for the optimal categories for a given consideration set and length profile. However, the optimal inquiry must also satisfy a fixed-point property: the optimal information partition give rise to a distribution of the corresponding categories, but the length profile must minimize the average length with respect to that distribution. The solution to this minimization problem is the well known Huffman coding scheme for a given distribution, and our framework endogenizes the underlying distribution through optimal categories.

We draw behavioural implications from the optimal trade-off between a more precise posterior for each option considered against the simplicity of the decision process, here measured by the number of questions required to arrive at the posterior. First, there is negative correlation between the inquiry length that leads to a category and the likelihood of the choice of the associated decision. That is, more likely decisions are also easier to arrive at. Second, the consideration set shrinks with the cognitive cost. For a very low cost, all feasible options are considered; for a very high cost, only one of them will be considered and no information is processed. In between, the optimal average length of the inquiry decreases with the cost, and, in environments with i.i.d. valuations according to the uniform distribution, the size of consideration set decreases with the cost.

Our second main result demonstrates that the optimal inquiry exhibits confirmation bias. We show that the decision maker optimally seek information to *confirm* the optimality of the *ex ante* more likely options by enlarging the categories in which such decisions are taken. This is in line with the definition given by Nickerson (1998), "It refers usually to unwitting selectivity in the acquisition and use of evidence." By doing so, the decision maker endogenously increases the probability to end the inquiry with shortest questions and hence decreases the average length of inquiry. We show that such bias always exists weakly, and it holds strictly whenever the optimal inquiry includes an unequal length profile.

Related Literature. This paper provides a rational account of how to use limited cognitive resources, and how this optimal allocation explains two well-known behavioral biases: focused consideration, according to which the decision-maker only focus on a subset of feasible options, and confirmation bias, according to which the decision-maker seeks evidence to confirm her ex ante most likely options. These results contribute to three strands of literature, both conceptually and methodologically.

The first strand includes papers with an explicit formulation of decision procedures connected to the use of cognitive resources. Wilson (2014), following the approach of Cover and Thomas (2006), formulates the decision-making process as a finite automaton. The main result in Wilson (2014) is a dynamic-consistency type of result called multi-self consistency. The cognitive limitation is modelled via an exogenously given number of memory states that capture the decision-maker's memory capacity. In contrast, we prove the dynamic consistency in the conventional sense and endogenize the size of the optimal inquiry via a cognitive cost.

The second strand includes papers that take the rational-inattention approach and model information processing via noisy-signal acquisition initiated by Sims (2003). Methodologically, in this literature the cognitive cost is typically modelled as entropy reduction relative to the prior belief, as in Matějka and McKay (2015) and Jung et al. (2019). In contrast, we have an explicit formulation of the decision process, and our cognitive cost is directly associated with asking questions. The dynamic consistency result allows us to show that our approach is robust to the assumption that the decision-maker has to commit to the information acquisition strategy that is typical in the literature. Moreover, our results provide endogenous categorization directly linked to the consideration sets generated by the optimal inquiry.

The third strand includes papers that rationalize behavioural biases with cognitive frictions. These papers range from axiomatic to constrained optimization approaches. For consideration sets, the former include Masatlioglu et al. (2012) and Manzini and Mariotti (2014) and the latter includes Caplin et al. (2019). While our approach is closer to the latter, we connect the two approaches by showing that our optimal inquiry satisfies certain desirable axioms, such as dynamic consistency and the attention-filter property in Masatlioglu et al. (2012). For confirmation bias, Wilson's (2014) model also generates a form of confirmation bias based on limited memory. However, in her model the decision-maker does not seek evidence but passively processes it. In contrast, our decision-maker actively seeks evidence to confirm her more likely options. Nimark and Sundaresan (2019) also obtain a "confirmation effect" using the rational-inattention approach. Both papers consider binary states and confirmation to the prior. In contrast, we define confirmation bias as the decision-maker seeking evidence to confirm ex ante most likely guesses about which decision is optimal, a definition that is based on the observable choices.

2. The model

2.1. **Primitives.** A decision-maker (DM) needs to process information about a uncertain state of the world before taking an action. The DM's utility $u(a, x)$ depends on her action, $a \in A$, and an uncertain state of the world, $x \in X$. The set of actions A is finite and contains at least two actions. The set of states X is a convex subset of $\mathbb{R}^L$, $L \in \mathbb{N}$. State x is distributed according to a probability distribution G that is absolutely continuous and has full support on X. We will use notation $\mathbb{P}[\cdot]$ and $\mathbb{E}[\cdot]$ to denote the probability and expectation under G, respectively.

We say that action a weakly dominates another action a' if $u(a, x) \geq u(a', x)$ for all $x \in X$. Throughout the paper, we assume:

(A_1) For all $a \in A$, expected utility $\mathbb{E}[u(a, x)]$ is finite.
(A_2) For all $a, a' \in A$, a does not weakly dominate a'.
(A_3) For all $a', a'' \in A$ and any constant $c \in \mathbb{R}$, the set $\{x \in X : u(a', x) - u(a'', x) = c\}$ has empty interior.

Assumption (A_1) is needed for the DM's optimization problem to be well defined. Assumptions (A_2) and A_3) are introduced to simplify exposition. Assumption (A_2) precludes existence of dominated actions. Assumption (A_3) is a generalization of the condition of "thin" indifference curves between each pair of actions. It means that the utility curves of any two actions are almost never parallel to each other. Many usual utility functions satisfy this assumption. For example, (A_3) is satisfied for the following three classes of utility functions.

(U1) Here we have $|A| = L$ and $X \subset \mathbb{R}^L$, and $u(a_l, x) = x_l$, that is, x_l is the value of action a_l.
(U2) Here we have $X \subset \mathbb{R}^L$ and $A \subset \mathbb{R}^L$, and, for each $a \in A$, there is a $(\alpha_a, \beta_a) \in \mathbb{R} \times \mathbb{R}^L$ such that $u(a, x) = \beta_a \cdot x + \alpha_a$. This is essentially the Lancaster model of product characteristics.
(U3) Here $A \subset \mathbb{R}^L$ and $u(a, x)$ is the negative distance between a and x, that is, $u(a, x) = -||a - x||_p$, where $|| \cdot ||$ is the L_p-norm on $\mathbb{R}^L$. This is typically called the tracking problem.

2.2. **Inquiries.** When confronted with a state x, the DM does not observe x directly. Instead, she relies on a series of questions about x. Formally, we consider an *inquiry* as a series of *true/false* questions formulated as propositions about x. A proposition is a statement about certain properties of x, that is, a proposition has the form "$x \in Y$", where $Y \subset X$ is a Borel set. Denote the collections of Borel subsets of X by $\mathcal{B}(X)$.

We identify a proposition with a set $Y \in \mathcal{B}(X)$; we say that the proposition is *true* if $x \in Y$ and it is *false* if $x \notin Y$.

An *inquiry* $Q = \langle N, T, \sigma, \chi, d \rangle$ is a finite binary tree that describes propositions at non-terminal nodes and decisions in terminal nodes of the tree. Specifically:

- a finite set N of nodes contains a root n^o and a nonempty set T of terminal nodes (note that the tree may consist of a single terminal node, i.e., $N = T = \{n^o\}$);
- each non-terminal node $n \in N - T$ is followed by exactly two edges labelled *true* and *false*;
- successor function σ assigns to each non-terminal node $n \in N - T$ and each edge $e = \{true, false\}$ a child $\sigma(n, e) \in N$ of node n following edge e;
- proposition mapping χ assigns to each non-terminal node $n \in N - T$ a proposition $\chi(n) \in \mathcal{B}(X)$;
- decision rule d assigns to each terminal node $t \in T$ an action $d_t \in A$.

We denote the set of all possible inquiries by $\mathcal{Q}_X$.

Given a state of nature $x \in X$, the inquiry Q begins with the proposition $\chi(n^o)$ at the root of the inquiry tree, and it ends whenever a terminal node is reached. It proceeds by following the tree. At a non-terminal node $n \in N - T$, the inquiry asks whether it is true that $x \in \chi(n)$. If true, then the inquiry proceeds to the node $\sigma(n, true)$; otherwise, the inquiry proceeds to the node $\sigma(n, false)$. When a terminal node $t \in T$ is reached, the DM takes action d_t.

2.3. **Information.** The inquiry transforms a quantitative statement, say, "$x \geq r$", into a qualitative one, say, "yes" or "no", eventually leading to a qualitative recommendation of which action to choose. The underlying assumption is that the DM cannot directly digest quantitative information. Knowing that my blood sugar level is 6 *mmol/L* means little to a medical lay person like me, but knowing that it is below the level that would be labelled as "normal" is very useful as it suggests a decision of not going to the GP. Indeed, our theory is aimed at the optimal thresholds for what it means by "normal" (do nothing), "concerning" (see the doctor soon), or "emergency" (call an ambulance). In fact, in many concrete applications we should think of the quantitative state x as a composite score of different dimensions that only experts can read but our DM can and need to digest the qualitative information; our notion of inquiry endogenizes the semantics of the qualitative information.

Formally, the inquiry categorizes states of nature into subsets through a series of questions. When arriving any (terminal or non-terminal) node $n \in N$, the DM's information about the state is summarized by a subset of states, denoted by $I_n(Q)$. That is, given the answers to the questions in previous nodes, the DM can infer that the true state belongs to $I_n(Q)$, recursively defined as follows. Clearly, at the root, all states are possible, and hence $I_{n^o}(Q) = X$. Given a non-terminal node $n \in N - T$, let n^{true} and n^{false} be the successors of n after "true" and "false" answers to the proposition $\chi(n)$, respectively. Then we define

$$I_{n^{true}}(Q) = I_n(Q) \cap \chi(n) \quad \text{and} \quad I_{q^{false}}(Q) = I_n(Q) \cap \big(X - \chi(n)\big). \tag{1}$$

Now, for each $x \in X$, the DM will reach some terminal node t at the end of the inquiry. Thus, the set $I_t(Q)$ consists of all states under which terminal node t is reached. Thus, the collection of sets $\{I_t(Q) : t \in T\}$ forms a partition of X, and it represents the information sets at the end of the inquiry.

As we are not concerned with zero probability events, we adopt and use throughout the paper a measure based notion of partition that disregards sets of measure zero under G. Specifically:

Definition 2.1. A collection of sets $\{X_1, X_2, ..., X_K\}$ is a *partition of* X if $\mathbb{P}(X_k) > 0$ for each k, and $\sum_k \mathbb{P}(X_k) = \mathbb{P}(X) = 1$.

2.4. **Payoffs.** We assume that asking questions is costly. Let the DM's cognitive cost of any single question be $\lambda > 0$. Given an inquiry Q, let $\ell_t(Q)$ be the length of the path from n^o to t in the tree, that is, $\ell_t(Q)$ is the number of questions asked to reach terminal node t. Then, the ex-post cost of inquiry at terminal node t is equal to $\lambda \ell_t(Q)$.

We can now formulate the DM's optimization problem. Given an inquiry Q, if terminal node t is reached, the DM's ex-post payoff net of the cognitive cost at state x is

$$u(d_t, x) - \lambda \ell_t(Q).$$

Because each terminal node $t \in T$ is reached whenever the state x is in $I_t(Q)$, the DM's ex-ante expected utility from inquiry Q is

$$W(Q; \lambda) = \sum_{t \in T} \int_{x \in I_t(Q)} \Big(u(d_t, x) - \lambda \ell_t(Q)\Big) G(\mathrm{d}x). \tag{2}$$

The DM's optimization problem is

$$\max_{Q\in\mathcal{Q}} W(Q;\lambda). \tag{3}$$

The maximization problem (2) resembles the problem studied in the rational inattention literature, e.g. Matějka and McKay (2015), Jung et al. (2019), and Caplin et al. (2019). But this resemblance is more in formality than in substance. Indeed, while from information theory we that the average length of investigation defined here is closely related to entropy and the rational-inattention approach is motivated by measuring cognitive cost as number of questions, the standard approach measures the cost of information in terms of entropy reduction relative to the prior belief. In contrast, here the primitive cost does not depend on the prior—it is simply the asking (and the implied act of processing the answer) itself is costly. Moreover, while in the usual setup the model is silent about the corresponding procedure that the DM uses to arrive her decision, in our model there is an explicit connection between the solution to (3) and the procedure used. In particular, we may say that the realized process is *simpler* for a decision if fewer questions are needed to arrive that decision, that is, ℓ is smaller.

3. Optimal Inquiries

Unlike most of the literature, our model of information processing emphasizes the dynamic nature of the process. In particular, the order of the questions in an inquiry not only affects the final information set for the DM in her decision, but also affects the length of the inquiry, which also matters for the DM's payoff.

This emphasis of process makes the decision problem, (3), a nonstandard one, and it does not allow the use of the standard optimization techniques. In particular, the set of possible inquiries, $\mathcal{Q}_X$, is a infinite set and its elements are discrete in nature, which precludes the standard first-order approach. Below we establish three principles of optimality for the decision-making processes: the first is dynamic consistency, the second is the use of the consideration set, and the third is relationship between the inquiry length and the corresponding category of states for each element in the consideration set.

3.1. **Dynamic Consistency.** In our formulation, we implicitly assume that the DM commits to an inquiry strategy ex ante. Here we show that the DM does not want to change her ex-ante optimal inquiry strategy at any interim stage, so, her choices are not only ex-ante, but also sequentially optimal. This means that the commitment assumption is without loss of generality.

We use the following notion of dynamic consistency. Let $Q = \langle N, T, \sigma, \chi, d\rangle \in \mathcal{Q}_X$ be an inquiry. Consider a node $n \in N$. At that node, the DM infers that the state is in $I_n(Q)$. Observe that every possible play after reaching n is itself an inquiry, whose initial set of states is $I_n(Q)$. Let us refer to it as a *sub-inquiry at node* $n \in N$. The set of all possible sub-inquiries at n given information $I_n(Q)$ is $\mathcal{Q}_{I_n(Q)}$. Denote by Q_n the specific sub-inquiry at n that prescribes to play according to the original inquiry Q.

Suppose that the DM initially follows inquiry Q but, upon reaching node n, she reevaluates her plan: whether to follow the original plan Q_n or to deviate to another sub-inquiry $\hat{Q}$. Let $W_n(\hat{Q}; \lambda)$ be the DM's expected payoff conditional on reaching node n if she chooses sub-inquiry $\hat{Q} \in \mathcal{Q}_{I_n(Q)}$ upon arrival to n. We say that the original inquiry Q is dynamically consistent if no deviation is beneficial at any node.

Definition 3.1. An inquiry $Q = \langle N, T, \sigma, \chi, d\rangle$ is *dynamically consistent* if, for each node $n \in N$,

$$W_n(Q_n; \lambda) = \max_{\hat{Q} \in \mathcal{Q}_{I_n(Q)}} W_n(\hat{Q}; \lambda). \tag{4}$$

Note that dynamic consistency implies that the DM behaves sequentially optimally at each terminal node. Specifically, the DM chooses a decision that maximizes her expected payoff given the information at that node. That is, if Q is dynamically consistent, then, for each terminal node $t \in T$, the action d_t must be a solution of

$$\max_{a \in A} \int_{x \in I_t(Q)} u(a, x) G(\mathrm{d}x | I_t(Q)). \tag{5}$$

We have the following theorem.

Theorem 3.1. *Every optimal inquiry is dynamically consistent.*

3.2. **Outcomes.** We argue that it suffices to describe an optimal inquiry by its payoff relevant outcome. The outcome consist of three components: the set of actions employed at the terminal nodes, the inquiry length profile associated with the terminal nodes, and the information to the DM at the terminal nodes.

First, observe that if an inquiry is optimal, then every node must be reached with positive probability. Indeed, if there was a node n that is only reached with probability zero, then, in some predecessor node n', the proposition $\chi(n')$ or its complement would have had measure zero, so the associated costly question would have been redundant.

Lemma 3.1. *If an inquiry* $Q = \langle N, T, \sigma, \chi, d\rangle$ *is optimal, then every node* $n \in N$ *is reached with positive probability.*

Second, observe that an optimal inquiry cannot induce the same action in two or more terminal nodes. Indeed, if it was the case, it would be no need to distinguish between these terminal nodes, so the number of costly questions in the inquiry could be reduced.

Lemma 3.2. *If an inquiry* $Q = \langle N, T, \sigma, \chi, d\rangle$ *is optimal, then* $d_t \neq d_{t'}$ *for all pairs of distinct terminal nodes* $t, t' \in T$.

An immediate implication of Lemma 3.2 is that each terminal node corresponds to a unique action in A. In what follows, we will identify terminal nodes with actions they induce. Specifically, let $D(Q)$ be the set of actions induced in inquiry Q. We will refer to $D(Q)$ as the *consideration set*. The actions of $D(Q)$ are identified with the terminal nodes of Q. A representative terminal node is denoted by $d \in D(Q)$.

The identification of terminal nodes with their corresponding actions allows us to characterize the payoff-relevant outcomes by the consideration set. To do so, let Q be an inquiry with consideration set $D(Q)$. Observe that each action d in the consideration set $D(Q)$ is associated with the length of inquiry $\ell_d(Q)$ leading to the terminal node where d is chosen, and with the information set $I_d(Q)$ induced by Q in that terminal node. Let $\ell(Q) = (\ell_d(Q))_{d\in D(Q)}$ and $I(Q) = (I_d(Q))_{d\in D(Q)}$. Let us refer to the triple $Z(Q) = (D(Q), \ell(Q), I(Q))$ as the outcome profile induced by Q.

Let $D \subseteq A$, let $(\ell_d)_{d\in D} \in \mathbb{N}^{|D|}$ be a profile of lengths, and let I be a partition of X into $|D|$ elements. Denote by $\mathcal{Z}_X$ the set of triples (D, ℓ, I). We say that an outcome profile $(D, \ell, I) \in \mathcal{Z}_X$ is *implementable* if there exists an inquiry $Q \in \mathcal{Q}_X$ that induces this outcome profile, that is, $(D, \ell, I) = Z(Q)$. The following lemma characterizes implementable outcomes.

Lemma 3.3. *An outcome profile* $(D, \ell, I) \in \mathcal{Z}_X$ *is implementable if and only if*

$$\sum_{d\in D} 2^{-\ell_d} = 1. \tag{6}$$

Observe that outcomes capture all we need to know to evaluate the DM's expected payoff. Indeed, suppose that two different inquiries Q and Q' implement the same outcome (D, ℓ, I). Then, by (2), we have

$$W(Q;\lambda) = W(Q';\lambda) = \sum_{d\in D} \int_{x\in I_d} \Big(u(d,x) - \lambda \ell_d\Big) G(\mathrm{d}x).$$

3.3. **Optimal Inquiries.** We have shown that, without loss of generality, an inquiry can be summarized by an outcome (D, ℓ, I) it induces. Moreover, Lemma 3.3 shows that the partition of information I does not affect whether or not an outcome profile if implementable. This characterization allows us to solve the optimal inquiry in two stages. We first fix an arbitrary (D, ℓ) that satisfies (6), and solve for the optimal I; then, we maximize over all possible (D, ℓ)'s.

In the first stage, taking (D, ℓ) as given, we solve for the optimal action $d \in D$ for each state: under state x, the payoff net of the cognitive cost associated with action d is given by $u(d, x) - \lambda \ell_d$. Accordingly, for a given (D, ℓ), let $I^*_d(D, \ell)$ be the set of states where action d is the unique best-response action among all actions in D when the DM takes into account the cost of inquiry associated with each action:

$$I^*_d(D,\ell) = \big\{x \in X : u(d,x) - \lambda \ell_d > u(a,x) - \lambda \ell_a \text{ for all } a \in D - \{d\}\big\}. \tag{7}$$

Let $I^*(D, \ell) = \{I^*_d(D, \ell)\}_{d\in D}$. Note that $I^*(D, \ell)$ is a partition of X, because, by assumption (A_3), the set $\big(\bigcup_{d\in D} I^*(D,\ell)\big) - X$ has measure zero. The key observation is that $I^*(D, \ell)$ is the optimal information partition given (D, ℓ), as the DM chooses the unique best-response action for each state $x \in X$, except for a measure zero of states.

Lemma 3.4. *If (D, ℓ, I) is the outcome of an optimal inquiry, then I is identical to $I^*(D, \ell)$ up to a measure zero set.*

The second stage is the choice of (D, ℓ). Let $\mathcal{Z}^*$ be the set of all pairs (D, ℓ) with $D \subseteq A$ and ℓ satisfying (6). The DM chooses $(D, \ell) \in \mathcal{Z}^*$, and the outcome is determined by $(D, \ell, I^*(D, \ell))$. By Lemmas 3.2–3.4, we obtain the following characterization of optimal inquiries.

Theorem 3.2. *An inquiry Q is a solution of* (3) *if and only if the pair* $(D(Q), \ell(Q))$ *is a solution of*

$$\max_{(D,\ell)\in\mathcal{Z}^*} \sum_{d\in D} \int_{x\in I_d^*(D,\ell)} \Big(u(d,x) - \lambda \ell_d\Big) G(\mathrm{d}x). \tag{8}$$

Because $\mathcal{Z}^*$ is a finite set, and the expected utility is bounded for each $d \in D$ by assumption (A_1), we establish the existence of optimal inquiry.

Corollary 3.1. *An optimal inquiry exists.*

We conclude this section by pointing out a useful property of optimal inquiry. Let (D, ℓ, I) be the outcome of an optimal inquiry. The optimal length profile ℓ is determined through the Huffman coding applied to the probability distribution $\{\mathbb{P}(I_d)\}_{d\in D}$ over D. The algorithm follows a simple rule to generate a binary tree, using the probabilities $\{\mathbb{P}(I_d)\}_{d\in D}$ as the input. To describe the algorithm, denote $p_d^0 = \mathbb{P}(I_d)$ and order the probabilities according to the their values. In the first stage, take the last two, say, d_1 and d_2, and add up $p^1_{\{d_1,d_2\}} = p^0_{d_1} + p^0_{d_2}$; for other $d \neq d_1, d_2$, take $p_d^1 = p_d^0$. Next, we add up the least two values in p_C^1's to obtain $p^2_{C'}$'s. Then repeat the same exercise for p_C^1's. Note that at each stage, we obtain a distribution, p_C^j with the C's form a partition of D. in the end of the process, we obtain $p_{C_1}^{|D|-2}$ and $p_{C_2}^{|D|-2}$, where $C_1 \cup C_2 = D$. The tree is then generated as follows. From the initial node we branch into C_1 and C_2. Now, $p_{C_1}^{|D|-2}$ is the sum of $p_{C_1'}^{|D|-3} + p_{C_2'}^{|D|-3}$ with $C_1' \cup C_2' = C_1$, and we branch C_1 into C_1' and C_2', and so on. This ends until we reach the singleton d's and those will be the terminal nodes. For each I_d, the corresponding length is the length of the path leading from the initial node to the terminal node corresponding to d. We have the following proposition.

Proposition 3.1. *If (D, ℓ, I) is the outcome of an optimal inquiry, Then:*

(a) ℓ is obtained from the Huffman coding w.r.t. the distribution $\{\mathbb{P}(I_d)\}_{d\in D}$;

(b) for all $d, d' \in D$, if $\ell_d < \ell_{d'}$, then $\mathbb{P}(I_d) \geq \mathbb{P}(I_{d'})$.

According to Proposition 3.1, for any candidate consideration set D, the optimal length profile ℓ is endogenously determined by the partition $I^*(D, \ell)$. Moreover, decisions that take longer to reach are less likely to be chosen. Since the information partitions need to adjust with the length profile ℓ and vice versa, a fixed-point argument is needed to determine the two jointly. The adjustment of the optimal categories reflect a distortion coming from the cognitive cost, as compared to the standard Bayesian analysis where the DM learns the state at zero cost, and then chooses the optimal action in each state.

A positive cognitive cost brings about two effects. First, it may be optimal to have a smaller consideration set than A. Second, the categories that are reached faster under optimal inquiry have greater probabilities. So, the DM is willing to sacrifice the precision in the sense of taking the precise action by enlarging the set of parameters under which an action is taken with a shorter inquiry. This preference generates a "bias" if we compare the categories thus generated to the ones that would be used by a Bayesian DM. We examine these two effects in the following two sections.

4. Optimal Consideration Sets

Now we show that, as the cognitive cost increases, it is optimal for the DM to reduce her consideration set, and the expected length of inquiry decreases.

Given an outcome $Z = (D, \ell, I) \in \mathcal{Z}_X$, let $\bar{\ell}(Z)$ be the *expected length of inquiry*:

$$\bar{\ell}(Z) = \sum_{d \in D} \ell_d \mathbb{P}(I_d).$$

We say that an inquiry with outcome Z is *uniform* if ℓ_d is the same for all $d \in D$. That is the inquiry is uniform if it ends with the same number of questions for all states of nature. Note that this can only happen if $|D| = 2^k$ for some $k \in \{0, 1, ...\}$.

Proposition 4.1. *Given λ, let Z_λ be the outcome of an optimal inquiry. Then:*

(a) The average inquiry length $\bar{\ell}(Z_\lambda)$ is decreasing in λ, and it strictly decreasing in λ whenever Z_λ is not uniform.

(b) There exist two thresholds $\lambda_2 > \lambda_1 > 0$ such that for all $\lambda < \lambda_1$, the optimal consideration set is $D_\lambda = A$; and for all $\lambda > \lambda_2$, the optimal consideration set is a singleton, $|D_\lambda| = 1$, and $\bar{\ell}(Z_\lambda) = 0$.

Proposition 4.1 gives the comparative statics result w.r.t. the cognitive cost. The main trade-off faced by the DM is between more precise information induced by longer inquiries and the cognitive cost. But more precise information is useful only if it leads to differential actions that generate higher differential payoffs under different states of nature. In contrast, if two actions are similar, it will not be worthwhile to differentiate them. Formally, let $\delta(a', a'')$ measure how close actions a' and a'' are in the payoff space:

$$\delta(a', a'') = \sup_{x \in X} |u(a', x) - u(a'', x)|.$$

Proposition 4.2. *When two actions are sufficiently close, then only one of them will be in the optimal consideration set.*

The literature, however, is also interested in the comparative statics w.r.t. the set of actions A. In particular, a popular property is called "attention filter" (Masatlioglu et al., 2012), and is defined as follows. Suppose that DM's consideration set D is a strict subset of A. Then, the attention filter property requires that, under a smaller action set A' that contains D, the optimal consideration set is still D. The following theorem shows that this property holds under optimal inquiry.

Proposition 4.3. *If (D, ℓ, I) is the outcome for an optimal inquiry for A, then it is also the outcome for an optimal inquiry for each A' such that $D \subseteq A' \subsetneq A$.*

5. Endogenous Confirmation Biases

We have seen from the previous section that, under the optimal inquiry, it is optimal for the DM to choose a consideration set and ask questions that would lead to decisions in that set only. Together with the dynamic consistency result (Theorem 3.1), the DM will never consider any action outside that set even if she can re-optimize at some point in her inquiry. This may be interpreted as a form of confirmation bias in the *extensive margin*, as the DM only searches for evidence to support decisions within the chosen consideration set.

Now we turn to confirmation bias in the intensive margin. Namely, under a given optimal consideration set $D \subseteq A$, the DM searches for evidence to confirm the desirability of the actions in D that are most likely to be optimal from ex ante perspective.

Given the consideration set D, consider what would have been the optimal partition of X in case of costless information processing, $\lambda = 0$. This case corresponds to the standard Bayesian analysis where the DM knows the state. Let I^B be the optimal information partition for a Bayesian DM (without the cost) when the set of feasible actions given by D. Formally, for each $d \in D$, $I^B_d = I^*_d(D, \ell)$ as defined by (7) for any ℓ, with $\lambda = 0$. Note that, by (A_2) and (A_3), I^B_d has a nonempty interior for each $d \in D$.

We define confirmation bias against the zero-cost benchmark, given the same consideration set. To do so, let us order the actions in D, so $D = \{d_k\}_{k=1}^{|D|}$, such that

$$\begin{gathered}\mathbb{P}(I_{d_1}) \geq \mathbb{P}(I_{d_2}) \geq ... \geq \mathbb{P}(I_{d_K}) \\ \text{with a tie-breaking rule } \mathbb{P}(I_{d_k}) = \mathbb{P}(I_{d_{k+1}}) \implies \ell_{d_k} \leq \ell_{d_{k+1}}.\end{gathered} \tag{9}$$

We have the following definition.

Definition 5.1. An inquiry Q with outcome (D, ℓ, I) has *confirmation bias* if for every order $\{d_k\}_{k=1}^{|D|}$ that satisfies (9)

$$\bigcup\nolimits_{t=1}^{K} I_{d_k}^B \subseteq \bigcup\nolimits_{k=1}^{K} I_{d_k} \quad \text{all } K = 1, ..., |D|, \tag{10}$$

It has *strict confirmation bias* if (10) holds and there exists $K \in \{1, ..., |D|\}$ such that

$$\left(\bigcup\nolimits_{k=1}^{K} I_{d_k}\right) - \left(\bigcup\nolimits_{t=1}^{K} I_{d_k}^B\right) \quad \text{has nonempty interior.} \tag{11}$$

Note that the inclusion (10) implies probability ranking for each $K \in \{1, ..., |D|\}$:

$$\sum_{k=1}^{K} \mathbb{P}(I_{d_k}) \geq \sum_{k=1}^{K} \mathbb{P}(I_{d_k}^B). \tag{12}$$

Moreover, because the distribution G has full support on X, strict confirmation bias implies that (12) is strict for some K.

According to Definition 5.1, confirmation bias means that an outside observer would conclude that the DM take the accumulative frequencies of the more likely actions would be higher than those of a corresponding Bayesian DM with identical preferences but has not cognitive cost. The reason for that to happen in our model is that the DM optimally adjust her inquiry strategy to seek information that would confirm those more likely actions.

The following theorem shows that all optimal inquiries have confirmation bias. Recall that an inquiry with outcome (D, ℓ, I) is uniform if the inquiry lengths ℓ_d are the same for all $d \in D$.

Theorem 5.1. *Every optimal inquiry has confirmation bias. Moreover, an optimal inquiry has strict confirmation bias if and only if it is not uniform.*

Theorem 5.1 then shows that the optimal inquiry would always prefer to confirm the more likely actions, and strictly so whenever the inquiry is not uniform. This strictness happens whenever the total number of categories is not a power of two, or, even when it is, the probabilities of different categories are sufficiently heterogeneous.

This result has various implications. First, it shows that whenever the DM has limited cognitive resources and prefers to minimize the use of them, she would bias toward more ex ante more likely actions to obtain a simpler cognitive process. This seems to match with the findings in the "tunnel vision" theory of wrongful convictions, according to which the investigators would direct the instigation and process information aiming at confirming ex ante more likely hypotheses (c.f., O'Brien (2009), Findley and Scott (2006)). Similarly, papers that study misdiagnoses find that doctors often have bias toward more likely conditions and inquire along the line to confirm those hypotheses (c.f. Croskerry (2013)). Our theory then provides a potential explanation based on costly cognition for such phenomena.

Appendix A. Proofs

A.1. **Proof of Theorem 3.1.** To prove Theorem 3.1, we use the results presented in Sections 3.2 and 3.3.

Let $Q^* = \langle N, T, \sigma, \chi, d\rangle$ be an optimal inquiry, and let $Z^* = (D^*, \ell^*, I^*)$ be the outcome implemented by Q^*. By Lemma 3.2, $D^* = T$. By Lemma 3.4, for each $d \in D^*$, category I^*_d is given by (7). To simplify notation, let $X_n = I_n(Q^*)$ for each $n \in N$. By Lemma 3.1,

$$\mathbb{P}(X_n) > 0 \ \text{ for all } n \in N. \tag{13}$$

Fix a node $n \in N$. Let $T_n \subset T$ be the set of terminal nodes that can be reached from n under Q^*. Note that if n is terminal (that is, if $n \in T$), then $T_n = \{n\}$. Let $\ell_n(Q^*)$ be the length of the path from n^o to n. Let Q^*_n be the sub-inquiry at n induced by the optimal inquiry Q^*. Conditional on reaching n, the DM's expected payoff from a sub-inquiry $\hat{Q} = \langle \hat{N}, \hat{T}, \hat{\sigma}, \hat{\chi}, \hat{d}\rangle \in \mathcal{Q}_{X_n}$ is given by

$$W_n(\hat{Q};\lambda) = \frac{1}{\mathbb{P}(X_n)}\left(\sum_{t\in\hat{T}}\int_{x\in I_t(\hat{Q})} u(\hat{d}_t, x) - \lambda \ell_t(\hat{Q})\right) G(\mathrm{d}x|X_n), \tag{14}$$

where $\{I_t(\hat{Q})\}_{t\in\hat{T}}$ is a partition of X_n induced by $\hat{Q}$, and $\ell_t(\hat{Q})$ is the length of inquiry beginning from node n and terminating at node $t \in \hat{T}$. Recall that Q^*_n is the sub-inquiry at n that prescribes to follow the optimal inquiry Q^*, so the DM's expected payoff from Q^* conditional on reaching n is given by (14) with $\hat{Q} = Q^*_n$.

Let us prove (4). Clearly, $W_n(Q_n^*;\lambda) \leq \max_{\hat{Q}\in\mathcal{Q}(X_n)} W_n(\hat{Q};\lambda)$. Suppose by contradiction that this inequality is strict. That is, there is a deviation $\hat{Q} \in \mathcal{Q}(X_n)$ in node n such that $W_n(Q_n^*;\lambda) < W_n(\hat{Q};\lambda)$, or equivalently, by (14),

$$\sum_{t\in T_n}\left(\int_{x\in I_t(Q_n^*)} u(d_t,x) - \lambda\ell_t(Q_n^*)\right) G(\mathrm{d}x|X_n) < \sum_{t\in\hat{T}}\left(\int_{x\in I_t(\hat{Q})} u(\hat{d}_t,x) - \lambda\ell_t(\hat{Q})\right) G(\mathrm{d}x|X_n). \quad (15)$$

Let $\tilde{T} = (T - T_n)\cup\hat{T}$, and construct an outcome $\tilde{Z} = (\tilde{I}_t,\tilde{\ell}_t,\tilde{d}_t)_{t\in\tilde{T}}$ as follows:

$$(\tilde{I}_t,\tilde{\ell}_t,\tilde{d}_t) = \begin{cases}(I_t(Q^*),\ell_t(Q^*),d_t), & \text{for each } t\in T-T_n,\\ (I_t(\hat{Q}),\ell_n(Q^*)+\ell_t(\hat{Q}),\hat{d}_t), & \text{for each } t\in\hat{T}.\end{cases}$$

By construction, $\tilde{Z}$ is an implementable outcome by an inquiry in $\mathcal{Q}_X$. Namely, inquiry $\tilde{Q}$ that implements $\tilde{Z}$ is obtained from Q^* by replacing the branch that follows node n with $\hat{Q}$. Then, we have

$$\begin{aligned}W(\tilde{Q};\lambda)-W(Q^*;\lambda) &= \mathbb{P}(X_n)\big(W_n(\hat{Q};\lambda) - W_n(Q_n^*;\lambda)\big)\\ &= \mathbb{P}(X_n)\left[\sum_{t\in\hat{T}}\left(\int_{x\in I_t(\hat{Q})} u(d_t,x) - \lambda(\ell_n(Q^*)+\ell_t(\hat{Q}))\right) G(\mathrm{d}x|X_n)\right.\\ &\quad\left.-\sum_{t\in T_n}\left(\int_{x\in I_t(Q_n^*)} u(d_t^*,x) - \lambda\ell_t(Q^*)\right) G(\mathrm{d}x|X_n)\right] > 0.\end{aligned}$$

The first equality is by definition of W and that $\tilde{Q}$ and Q^* differ only in the branch at node n. The second equality is by definition of W_n and the fact that the total length of path from n^o to t under $\tilde{Q}$ is the sum of the length from n^o to n under Q^* and the length from n to t under $\hat{Q}$. The inequality is by (13) and (14). Thus, we reached a contradiction to the assumption that Q^* is optimal. □

A.2. **Proof of Lemma 3.1.** Let Q be an optimal inquiry. By contradiction, let $n' \in N$ be a node that is reached with probability zero, but all the predecessors are reached with positive probability. Let n be the immediate predecessor of n', and let n'' be the second successor of n. Consider now a new inquiry $\hat{Q}$ obtained by modifying Q as follows. The question at node n and the entire branch following n' are removed. Instead, upon reaching node n, the inquiry $\hat{Q}$ will follow the branch of Q starting from the node n''. Clearly, every terminal node $t \in T$ that is reached with positive probability under Q is reached with the same probability under $\hat{Q}$, and the DM's expected payoff conditional on reaching any such node is unchanged. But the length of inquiry for the terminal nodes in the branch that starts from n'' is shorter under $\hat{Q}$. This contradicts the optimality of Q. □

A.3. **Proof of Lemma 3.2.** To prove Lemma 3.2, we use the following three claims.

Claim A.1. *Let (N, T, σ) be a binary tree with a set of nodes N, a set of terminal nodes $T \subset N$, and a successor function σ. For each $t \in T$, let ℓ_t be the length of the path from the root to t. Then $\sum_{t \in T} 2^{-\ell_t} = 1$.*

Proof. This claim directly follows from Theorem 5.2.1. in Cover and Thomas (2006) and its proof. As in that proof, one can convert an instantaneous code into a binary so that the lengths of paths to the terminal nodes correspond exactly to the codeword lengths. We have an equality here instead of inequality because in our inquiry tree every non-terminal node branches down to two further nodes. □

Claim A.2. *Let $K \geq 1$. If $\ell = (\ell_1, ..., \ell_{K+1}) \in \mathbb{N}^{K+1}$ satisfies $\sum_{k=1}^{K+1} 2^{-\ell_k} = 1$, then there exists $\ell' = (\ell'_1, ..., \ell'_K) \in \mathbb{N}^K$ such that $\ell'_k \leq \ell_k$ for all $k = 1, ..., K$, $\ell'_{k_0} < \ell_{k_0}$ for some $k_0 \in \{1, ..., K\}$, and $\sum_{k=1}^{K} 2^{-\ell'_k} = 1$.*

Proof. Without loss of generality assume that $\ell_1 \leq \cdots \leq \ell_{K+1}$. It follows that $\ell_K = \ell_{K+1}$; for otherwise the terminal node corresponding to ℓ_{K+1} must be the only successor of its predecessor. Let $\ell'_k = \ell_k$ for $k = 1, ..., K-1$ and let $\ell'_K = \ell_K - 1$. Thus,

$$\sum_{k=1}^{K} 2^{-\ell'_k} = \sum_{k=1}^{K-1} 2^{-\ell_k} + 2^{-\ell'_K} = \sum_{k=1}^{K-1} 2^{-\ell_k} + 2^{-\ell_K + 1} = \sum_{k=1}^{K+1} 2^{-\ell_k} = 1,$$

where the second last inequality follows from the fact that $\ell_K = \ell_{K+1}$. □

Claim A.3. *Let $I = \{I_k\}_{k=1}^K$ be a partition of X into K elements, let $D = \{d_1, ..., d_K\} \subset A$, and let $\ell = (\ell_1, ..., \ell_K) \in \mathbb{N}^K$ be a length profile such that*

$$\sum_{k=1}^{K} 2^{-\ell_k} = 1. \tag{16}$$

Then, there exists an inquiry $Q = \langle N, T, \sigma, \chi, d \rangle$ with a set $T = \{t_1, ..., t_K\}$ of terminal nodes such that

$$I_{t_k}(Q) = I_k \quad \text{and} \quad \ell_{t_k}(Q) = \ell_k \quad \text{for all } k = 1, ..., K. \tag{17}$$

Proof. By Theorem 5.2.1. in Cover and Thomas (2006) (with the argument as in the proof of Claim A.1 that translate instantaneous codes into binary trees), (16) implies that there exists a finite binary tree with a set of nodes N and a successor relation over N, with K terminal nodes labeled $t_1, ..., t_K$, such that, for each $k = 1, ..., K$, the length of the path from the root to each terminal node t_k is exactly ℓ_k.

We now construct an inquiry $Q = \langle N, T, \sigma, \chi, d\rangle$ that satisfies (17). Let N be as above, and let $T = \{t_1, ..., t_K\}$. For each nonterminal node $n \in N - T$, let us associate two edges leading out of n with $true$ and $false$, and define the map σ so that $\sigma(n, true) = n^{true}$ if $n \rightsquigarrow n^{true}$ along the edge labelled $true$ and $\sigma(n, false) = n^{false}$ if $n \rightsquigarrow n^{false}$ along the edge labelled $false$. Let decision rule d be given by the choice of d_k in terminal node t_k for each $k = 1, ..., K$.

It remains to construct a proposition mapping χ that yields the partition I in the terminal nodes. First, we associate each node in N with a set, $I_n(Q)$, as follows. For each $k = 1, ..., K$, let $I_{t_k}(Q) = I_k$. Then, by backward induction, for each nonterminal node $n \in N - T$, let $I_n(Q) = I_{\sigma(n,true)}(Q) \cup I_{\sigma(n,false)}(Q)$. This implies that $I_{n^o}(Q) = X$ at the root n^o, since $\{I_k\}_{k=1}^{K}$ is a partition.

Finally, define a proposition map χ as follows. For each nonterminal node $n \in N - T$, let $\chi(n) = I_{\sigma(n,true)}(Q)$. By induction from the root of the tree, it is straightforward to verify that χ satisfies (1), so, for each $n \in N$, $I_n(Q)$ is indeed the information set induced by Q at node n. □

We now prove Lemma 3.2. Let $Q = \langle N, T, \sigma, \chi, d\rangle$ be an optimal inquiry. Suppose, by contradiction, that $d_{t'} = d_{t''}$ for some $t', t'' \in T$ with $t' \neq t''$. Let $K = |T| - 1$, and let us label the terminal nodes consecutively, $T = \{t_1, ..., t_K, t_{K+1}\}$, such that $t_K = t'$ and $t_{K+1} = t''$.

Now we construct an alternative inquiry, $Q' = (N', T', \sigma', \chi', d')$, with $|T'| = K$ terminal nodes that leads to a strictly higher expected value to the DM. Let

$$I'_k = I_{t_k}(Q) \text{ for each } k = 1, ..., K - 1, \quad \text{and} \quad I'_K = I_{t_K}(Q) \cup I_{t_{K+1}}(Q), \tag{18}$$

and let

$$d'_k = d_{t_k} \text{ for each } t = 1, ..., K.$$

Now, by Claim A.1, we have $\sum_{k=1}^{K+1} 2^{-\ell_{t_k}(Q)} = 1$. By Claim A.2, there exists $\ell' \in \mathbb{N}^K$ such that

$$\ell_{t_k}(Q) \le \ell'_k \text{ for all } k = 1, ..., K, \ \ell_{t_k}(Q) < \ell'_k \text{ for some } k \in \{1, ..., K\}, \tag{19}$$

and $\sum_{k=1}^{K} 2^{-\ell'_k(Q)} = 1$. By Claim A.3 applied to $I' = \{I'_k\}_{k=1}^K$, $\ell' = (\ell'_1, ..., \ell'_K)$, and $d' = (d'_1, ..., d'_K)$, there exists an inquiry $Q' = \langle N', T', \sigma', \chi', d' \rangle$ with $T' = \{t_1, ..., t_K\}$ such that

$$I'_{t_k}(Q') = I'_k \ \text{ and } \ \ell'_{t_k}(Q') = \ell'_k \ \text{ for all } k = 1, ..., K. \tag{20}$$

Thus, we obtain

$$\begin{aligned} W(Q'; \lambda) &= \sum_{k=1}^{K} \int_{I_{t_k}(Q')} (u(d'_k, x) - \lambda \ell'_{t_k}(Q')) G(\mathrm{d}x) = \sum_{k=1}^{K} \int_{I'_k} (u(d'_k, x) - \lambda \ell'_k) G(\mathrm{d}x) \\ &> \sum_{k=1}^{K+1} \int_{I_{t_k}(Q)} (u(d_{t_k}, x) - \lambda \ell_{t_k}(Q)) G(\mathrm{d}x) = W(Q; \lambda), \end{aligned}$$

where the first and last equalities are by (2), the second equality is by (20), and the inequality is by (18), (19), and that, by Lemma 3.1, all the terminal nodes in T are reached with positive probability under Q. □

A.4. **Proof of Lemma 3.3.**

Necessity. Suppose that an outcome profile (D, ℓ, I) is implementable by an inquiry $Q = \langle T, N, \sigma, \chi, d \rangle$. Let $(D, \ell, I) = (T, \ell(Q), I(Q))$. By Lemma 3.2, $D \subset A$, and, by Claim A.1, (D, ℓ) satisfies (6).

Sufficiency. Immediate by Claim A.3. □

A.5. **Proof of Lemma 3.4.** Let (D, ℓ) be given. For any partition $I = \{I_d : d \in D\}$, let

$$W(I; D, \ell) = \sum_{d \in D} \int_{I_d} [u(d, x) - \lambda \ell_d] G(\mathrm{d}x).$$

Now, by (7), for any I and any $d \in D$, if $x \in I^*_d(D, \ell) \cap I_{d'}$ with $d \ne d'$ then

$$[u(d, x) - \lambda \ell_d] > [u(d', x) - \lambda \ell_{d'}].$$

Thus, since $\mathbb{P}(X - \cup_{d \in D} I_d^*) = 0$ by (A_3) and the fact that G has full support,

$$
\begin{aligned}
& W(I^*; D, \ell) - W(I; D, \ell) \\
= & \sum_{d, d' \in D} \int_{I_d^* \cap I_{d'}} \{[u(d, x) - \lambda \ell_d] - [u(d', x) - \lambda \ell_{d'}]\} \, G(\mathrm{d}x) \\
- & \sum_{d \in D} \int_{I_d \cap (X - \cup_{d \in D} I_d^*)} [u(d, x) - \lambda \ell_d] G(\mathrm{d}x) \\
= & \sum_{d \neq d' \in D} \int_{I_d^* \cap I_{d'}} \{[u(d, x) - \lambda \ell_d] - [u(d', x) - \lambda \ell_{d'}]\} \, G(\mathrm{d}x) \geq 0,
\end{aligned}
$$

and the inequality is strict if $\mathbb{P}(I_d^* \cap I_{d'}) > 0$ for some $d \neq d'$. This proves the result. □

A.6. **Proof of Theorem 3.2.** By Lemma 3.4, if (D, ℓ, I) is the outcome of an optimal inquiry, then $W(I; D, \ell) = W(I^*; D, \ell)$. To be optimal, it then must solve (8).

A.7. **Proof of Proposition 3.1.** (a) If (D, ℓ) solves (8), given the partition, the length profile must deliver the lowest average length and hence must be given by Huffman coding.

(b) Let $Z = (D, \ell, I)$ be the outcome of an optimal inquiry. First we show that if $\ell_d < \ell_{d'}$, then $\mathbb{P}(I_d) \geq \mathbb{P}(I_{d'})$. Suppose, by contradiction, that $\mathbb{P}(I_d) < \mathbb{P}(I_{d'})$. Now, let $\ell'_d = \ell_{d'}$ and $\ell'_{d'} = \ell_d$, and keep other outcomes unchanged. Note that the new outcome still satisfies (6) and hence can be induced by some inquiry. But now

$$
\begin{aligned}
[\mathbb{P}(I_d)\ell'_d + \mathbb{P}(I_{d'})\ell'_{d'}] - [\mathbb{P}(I_d)\ell_d + \mathbb{P}(I_{d'})\ell_{d'}] &= [\mathbb{P}(I_d)\ell_{d'} + \mathbb{P}(I_{d'})\ell_d] - [\mathbb{P}(I_d)\ell_d + \mathbb{P}(I_{d'})\ell_{d'}] \\
&= -[\mathbb{P}(I_{d'}) - \mathbb{P}(I_d)](\ell_{d'} - \ell_d) < 0.
\end{aligned}
$$

Thus, the new inquiry decreases the average length but keeps the utilities unchanged. This is a profitable deviation and a contradiction to the optimality of the original inquiry. □

A.8. **Proof of Proposition 4.1.** Let $\lambda_1 < \lambda_2$. Let Q_{λ_j} be an optimal inquiry for $j = 1, 2$. We show that $\bar{\ell}(Q_{\lambda_1}) \geq \bar{\ell}(Q_{\lambda_2})$, and that $D(Q_{\lambda_1})$ cannot be a strict subset of $D(Q_{\lambda_2})$

By the optimality of Q_{λ_j} given λ_j, for each $j = 1, 2$, we have

$$
\bar{u}(Q_{\lambda_1}) - \lambda_1 \bar{\ell}(Q_{\lambda_1}) \geq \bar{u}(Q_{\lambda_2}) - \lambda_1 \bar{\ell}(Q_{\lambda_2})
$$

$$
\bar{u}(Q_{\lambda_2}) - \lambda_2 \bar{\ell}(Q_{\lambda_2}) \geq \bar{u}(Q_{\lambda_1}) - \lambda_2 \bar{\ell}(Q_{\lambda_1}).
$$

Combining these inequalities yields

$$\lambda_1\left(\bar{\ell}(Q_{\lambda_1})-\bar{\ell}(Q_{\lambda_2})\right)\leq\bar{u}(Q_{\lambda_1})-\bar{u}(Q_{\lambda_2})\leq\lambda_2\left(\bar{\ell}(Q_{\lambda_1})-\bar{\ell}(Q_{\lambda_2})\right) \tag{21}$$

Thus, by (21) we have $\bar{\ell}(Q_{\lambda_1})\geq\bar{\ell}(Q_{\lambda_2})$ whenever $\lambda_1<\lambda_2$.

Clearly, for λ sufficiently large, the benefit from more information does not justify any inquiry and hence optimal Q is the degenerate one. Now we prove that for sufficiently small λ we have optimal $D=A$. Clearly, for any $D\subsetneq A$ and any ℓ for D satisfying (6), $W(D,\ell;0)<W(A,\ell^*;0)$, where ℓ^* is obtained from the Hufffman coding with the distribution $(p_k^B)_{k=1}^K$. Since W is continuous in λ, for sufficiently small λ the inequality still holds strictly. By (8) optimal inquiry has $D=A$. □

A.9. **Proof of Proposition 4.2.** Let (D,ℓ,I) be the outcome of an optimal inquiry Q. Suppose that $\delta(a',a'')<\lambda$ for some $a',a''\in A$. Suppose by contradiction that $a',a''\in D$. There are two cases.

Case 1. Suppose that $\ell_{a'}\neq\ell_{a''}$. W.l.o.g., let $\ell_{a'}<\ell_{a''}$. By Lemma 3.4 and assumption (A_3), $a'\in D$ implies that the set

$$I_{a'}=\{x\in X:u(a',x)>u(a,x)-\lambda(\ell_{a'}-\ell_a)\text{ for all }a\in D-\{a'\}\} \tag{22}$$

has nonempty interior. Therefore, because $a''\in D$, we must have

$$u(a',x)>u(a'',x)-\lambda(\ell_{a'}-\ell_{a''})\geq u(a'',x)+\lambda\ \text{ for each }x\in I^*_{a'},$$

where the first inequality is by (22), and the second inequality is because $\ell_{a'}<\ell_{a''}$ and both $\ell_{a'}$ and $\ell_{a''}$ are integers. This contradicts the assumption that $\delta(a',a'')<\lambda$.

Case 2. Suppose that $\ell_{a'}=\ell_{a''}$. Consider an inquiry $\hat{Q}$ with the outcome $(\hat{D},\hat{\ell},\hat{I})$ given by $\hat{D}=D-\{a''\}$, $\hat{\ell}_{a'}=\ell_{a'}-1$, $\hat{\ell}_a=\ell_a$ for all $a\in D-\{a'\}$, $\hat{I}_{a'}=I_{a'}\cup I_{a''}$, and $\hat{I}_a=I_a$ for all $a\in D-\{a'\}$. In words, $\hat{Q}$ is the same as Q except that $\hat{Q}$ merges actions a' and a'' and removes the question that distinguishes these actions. Because

$$\ell_{a'}=\ell_{a''}=\hat{\ell}_{a'}+1, \tag{23}$$

we obtain $2^{-\ell_{a'}}+2^{-\ell_{a''}}=2^{-\ell_{a'}}$. Since $\sum_{d\in D}2^{-\ell_d}=1$, we obtain that

$$\sum_{d\in\hat{D}}2^{-\hat{\ell}_d}=\left(\sum_{d\in\hat{D}-\{a'\}}2^{-\hat{\ell}_d}\right)+2^{-\hat{\ell}_{a'}}=\left(\sum_{d\in D-\{a',a''\}}2^{-\ell_d}\right)+2^{-\ell_{a'}}+2^{-\ell_{a''}}=1.$$

Thus, by Lemma 3.3, there exists an inquiry $\hat{Q}$ constructed as above. As Q and $\hat{Q}$ differ only for $x \in I_{a'} \cup I_{a''}$, we obtain

$$\begin{aligned}
W(\hat{Q};\lambda) - W(Q;\lambda) &= \int_{I_{a'}} \Big((u(a',x) - \lambda\hat{\ell}_{a'}) - (u(a',x) - \lambda\ell_{a'}) \Big) G(\mathrm{d}x) \\
&\quad + \int_{I_{a''}} \Big((u(a',x) - \lambda\hat{\ell}_{a'}) - (u(a'',x) - \lambda\ell_{a''}) \Big) G(\mathrm{d}x) \\
&= \int_{I_{a'}} \lambda G(\mathrm{d}x) + \int_{x \in I_{a''}} \Big(u(a',x) - u(a'',x) + \lambda \Big) G(\mathrm{d}x) \\
&> 0,
\end{aligned}$$

where the first equality is by (**??**), the second equality is by (23), and the inequality is because $\delta(a',a'') < \lambda$ and $I_{a'} \cup I_{a''}$ has nonempty interior. We thus obtain a contradiction to the optimality of Q. □

A.10. **Proof of Proposition 4.3.** Let (D,ℓ,I) be the outcome of an optimal inquiry such that $D \subseteq A' \subsetneq A$. It is immediate by Theorem 3.2 that the DM's maximal expected payoff when restricted to A' is attained by the same outcome. □

A.11. **Proof of Theorem 5.1.** Let (D,I,ℓ) be an outcome of an optimal inquiry. By (A_2) and (A_3), for each $d \in D$, the set

$$I_d^B = \{x \in X : \; u(d,x) > u(d',x) \text{ for all } d' \in D - \{d\}\}$$

has nonempty interior.

Let $\bar{K} = |D|$, and let $d_1, ..., d_{\bar{K}}$ be an order of actions in D that satisfies (9). By Proposition 3.1b, $\mathbb{P}(I_{d_k}) > \mathbb{P}(I_{d_{k+1}})$ implies $\ell_{d_k} \leq \ell_{d_{k+1}}$. Thus we have

$$\ell_{d_1} \leq \ell_{d_2} \leq ... \leq \ell_{d_{\bar{K}}}. \tag{24}$$

For each k and K such that $1 \leq k \leq K \leq \bar{K}$, we have

$$\begin{aligned}
I_{d_k}^B &= \big\{x \in X : u(d_k,x) > u(d_m,x) \; \forall m \in \{1,...,\bar{K}\}, \, m \neq k\big\} \\
&\subseteq \big\{x \in X : u(d_k,x) > u(d_m,x) \text{ for all } m \in \{K+1,...,\bar{K}\}\big\} \qquad (25) \\
&\subseteq \big\{x \in X : u(d_k,x) > u(d_m,x) - \lambda(\ell_{d_m} - \ell_{d_k}) \text{ for all } m \in \{K+1,...,\bar{K}\}\big\},
\end{aligned}$$

where the equality is by definition of $I_{d_k}^B = I_{d_k}^*(D,0)$, the first inclusion is because the constraint for $x \in X$ is weaker, and the second inclusion is because $\ell_{d_m} \geq \ell_{d_k}$ by (24), so the constraint for $x \in X$ is weaker. Next, observe that

$$\bigcup_{k=1}^{K} \{x \in X : u(d_k, x) > u(d_m, x) - \lambda(\ell_{d_m} - \ell_{d_k}) \text{ for all } m \in \{K+1, ..., \bar{K}\}\} = \bigcup_{k=1}^{K} I_{d_k},$$

because the left-hand side is the set of all x where at least one of the actions in $\{1, ..., K\}$ is strictly better than all actions in $\{K+1, ..., \bar{K}\}$, and the right-hand side is the same expression by definition of $I_{d_k} = I^*_{d_k}(D, \ell)$. Thus, taking the union over $k = 1, ..., K$ in (25), we obtain that (10) holds.

Suppose that the inquiry is not uniform, that is, there exist K and m with $K < m$ such that $\ell_{d_K} < \ell_{d_m}$. Then, the second inclusion in (25) is strict for all $k \leq K$, because $\lambda > 0$ and, by (24), $\ell_{d_m} - \ell_{d_k} \geq \ell_{d_m} - \ell_{d_K} > 0$. Moreover, by ($A_3$), the difference between the third and second line in (25) has nonempty interior. Thus, taking the union over $k = 1, ..., K$ in (25), we conclude that (11) holds for the specified K. □

References

Caplin, Andrew, Mark Dean, and John Leahy (2019): "Rational inattention, optimal consideration sets, and stochastic choice," *Review of Economic Studies*, 86 (3), 1061–1094. 5, 9

Cover, Thomas M and Joy A Thomas (2006): *Elements of Information Theory*, John Wiley & Sons, 2nd ed. 4

Jung, Junehyuk, Jeong Ho Kim, Filip Matějka, and Christopher A Sims (2019): "Discrete actions in information-constrained decision problems," *The Review of Economic Studies*, 86 (6), 2643–2667. 5, 9

Manzini, Paola and Marco Mariotti (2014): "Stochastic choice and consideration sets," *Econometrica*, 82 (3), 1153–1176. 5

Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Y Ozbay (2012): "Revealed attention," *American Economic Review*, 102 (5), 2183–2205. 5

Matějka, Filip and Alisdair McKay (2015): "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 105 (1), 272–298. 5, 9

Nickerson, Raymond S (1998): "Confirmation bias: A ubiquitous phenomenon in many guises," *Review of general psychology*, 2 (2), 175–220. 4

Nimark, Kristoffer P and Savitar Sundaresan (2019): "Inattention and belief polarization," *Journal of Economic Theory*, 180, 203–228. 5

SIMS, CHRISTOPHER A (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50 (3), 665–690. 5

——— (2010): "Rational inattention and monetary economics," in *Handbook of Monetary Economics*, ed. by Benjamin M. Friedman and Michael Woodford, Elsevier, vol. 3, 155–181. 2

WILSON, ANDREA (2014): "Bounded memory and biases in information processing," *Econometrica*, 82 (6), 2257–2294. 4, 5